Searching for a Muon Electric Dipole Moment with the Straw Trackers at the Fermilab Muon $g-2$ Experiment

Dominika Vasilkova

Submitted in fulfilment of the requirements of the degree of
Doctor of Philosophy (PhD)

Department of Physics and Astronomy
University College London

March 2023

Declaration

I, Dominika Vasilkova, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Dominika Vasilkova, March 2023

The Muon $g-2$ Experiment at Fermilab has been designed and operated by a large team of over 200 people. The introductory chapters of this thesis contain a general review of the physics behind the experiment and setup of the experiment. Besides the choices in presentation, this is not directly my work.

Similarly, the straw tracking detectors are the end result of years of effort from the $g-2$ tracker team, with my contributions during this PhD described in Chapter 4. In Chapter 5, the standalone simulation and associated studies were performed solely by me, with the GM2RINGSIM comparison coming from data I produced using the collaboration-built simulation software.

The radial field measurements and Electric Dipole Moment (EDM) analysis methods were developed in tandem with S.Grant's Run 1 analysis, with any overlap indicated via references. The GM2RINGSIM simulation is also used here, with the initial production of the simulation data split between several students including myself. Improvements including the study of limit setting methods, the Run 2/3 specific corrections, the base blinding, updated blinding and the implementation and optimisation of the acceptance maps were performed fully by me. Additionally, the studies for the EDM systematics for Run 2/3 are also solely my work, and are independent to Run 1.

Acknowlegdements

This thesis and the work in it would have never reached completion without the support of a whole crowd of people, supporting both academically and otherwise. Firstly, I would like to thank my supervisor Becky Chislett, for unwavering support and guidance, for always having Really Good Ideas whether that's for projects or how to get unstuck, for introducing me to the culinary wonders that are arancini and elotes, and for making a pretty darn effective attempt at squashing my impostor syndrome over the years!

The whole Muon $g-2$ collaboration deserves thanks for being such a great group of people to work with as well as an interesting experiment! Special thanks goes to Mark Lancaster, for providing both entertainment during our mutual covid lockdown month (who knew there was so much to learn about corrie and golf?!) as well as general good advice. I'd also like to thank Joe Price, James Mott, and Gavin Hesketh for their expertise and help untangling both tracker and EDM issues, and Simon Corrodi for helping with the radial field.

Within $g-2$ but also UCL, I would like to thank my friends and fellow (now ex!) PhD students, Gleb Lukicov and Sam Grant. Gleb, I thank for the infectious enthusiasm about anything $g-2$, for being my python buddy in a sea of C++ aficionados, and for showing me the ropes of data analysis, which has clearly grown to be a huge part of my PhD. Sam I would like to thank for the 'superhappyfuntime' we had, it's been a pleasure to work with you on the EDM and radial field, I've really appreciated being able to bounce ideas around with you in a way that's really helped me understand things better - and of course, all the much-needed jokes when everything we thought we understood falls apart!

During my time at Fermilab, I've had the good fortune to meet many lovely people that also deserve a mention here for making my time in the states both welcoming and fun, including Brendans Casey and Kiburg, Mackenzie Devilbiss, Adam Schreckenberger, Matteo Sorbrara, and the Jag Club: Lily Asquith and Saskia Charity. Particular shout-outs to the 'free food alert' crew – Meghna, Sudeshna and Lorenzo: may there always be an assortment of snacks left out unguarded in your near vicinity, Sophie for helping keep me sane during run-coordinating by providing lunch escapes and sushi Sundays, and Pawel for being a great flatmate while we were in Naperville – thanks for all the lifts, the soundtracks, and the steaks! A final FNAL-area shoutout to Esra, Mete and Jason for providing much-needed distractions from work via our weekly board games – chrrr chrr, tweet

tweet!

This section would be very incomplete if I didn't also thank those responsible for setting me on this path in the first place – Steve Maxfield and Barry King, who first snared my attention with a tale of a little particle with an intriguing anomaly. I would also like to thank my secondary school physics teacher, Mr Gibbs, for his endless patience with a student that would not stop asking questions. I wouldn't be here if it wasn't for you, so thank you.

Outside of work, it would be remiss to not thank the friends that have been providing some all-important non-physics fun and supporting me when things got hard. Particular thanks to Colin Beveridge and Elizabeth Williams for their constant encouragement (and for not being /too/ mean to physics even when it was being mean to me), and to Susannah Felstead for adventures in London and reminding me of the importance of dabbling in artsy things to keep myself sane. A special thank you to Vials and Co (Eshita, Nick, Juan, Rosie and Tim) and co-evil-plotter Tom, for letting me shake off the mantle of physicist once a week to roll dice and plan how to best mess with you. I would also like to thank my family for putting up with all my physics throughout the years as well as my new tendency to vanish abroad for months at a time!

Finally, I would like to thank my tiny feline overlord, Persia, without whom this thesis would probably have been written a lot faster. In the interests of preserving her contributions to cat science, I have reproduced here her contributions to various chapters:

On searching for BSM physics: Zde5tgh67ull,p pok00000000000k

On identifying crosstalk hits in the trackers: .,kmnjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjjbhjnnnnn099900000o0knh

And finally, on the new fit function needed to fit the EDM oscillation: jjjnssdfhp;'# lkjhuyhgq23WASE1RFTVGYHBU =======PPPPPPPPPPPPPP PPPPPPPPPPPPPPPPPPPPPL

Abstract

As well as measuring the muon magnetic dipole moment (MDM), the Muon $g-2$ experiment at Fermilab is aiming to make a new world leading measurement of the muon electric dipole moment (EDM), with a target of 10^{-21} e·cm, an improvement of two orders of magnitude compared to the previous best limit set at Brookhaven National Laboratory (BNL) [1]. This thesis presents the analysis of the Run 2/3 data, focussing on using the straw tracker detectors to measure an oscillation in the vertical angle of the muon precession plane to search for an EDM.

A standalone simulation is developed to study the experimental EDM sensitivity, enabling fast evaluation of the main systematics. The radial field is found to be the dominant error if not measured to better than 1 ppm precision. A new measurement technique is developed to measure this, with the results of the Run 5 radial field scan successfully achieving a 0.46 ppm uncertainty, better than the target precision. The EDM limit setting methods themselves are discussed and compared in order to select the optimal one for the analysis.

Since the trackers are a vital part of this measurement, characterisation and optimisation of their performance is also considered. Various studies are performed, including an investigation into the optimal operating voltage and an in-depth study of crosstalk. A new time-dependent tracker effect is discovered which directly impacts the EDM analyses, which is evaluated and a correction developed.

Run 2/3 data is analysed to perform a blinded search for a muon EDM, including corrections and systematic uncertainties. The results, if found to be zero when unblinded, will set a new world limit on the muon EDM of 5.9×10^{-20} e·cm. This improves on the previous value by a factor of 3.

Impact statement

The main work presented in this thesis is towards an analysis to measure the electric dipole moment (EDM) of the muon. This can be considered a measure of the intrinsic charge distribution within it, where a zero EDM would correspond to the muon being perfectly spherical.

The Standard Model (SM) of Particle Physics has shown great success in describing the Universe, but known issues with it remain. The field is therefore searching for any new physics which could help explain these anomalies on many frontiers. A non-zero muon EDM measurement would be a sign of new physics, as the SM predicts it is small, beyond the reach of current experiments. However, some theories predict a larger EDM. A large non-zero EDM also provides an additional source of CP violation, which is needed to explain the matter-antimatter asymmetry of the early Universe, one of the biggest unsolved problems in Physics.

This thesis presents an analysis of the Run 2/3 data from the Fermilab Muon $g-2$ experiment, with a tenfold improvement in statistics compared to the BNL analysis, leading to $3\times$ greater statistical sensitivity. As the analysis uses data from the tracking detectors to measure a vertical angle tilt, it is beneficial to optimise the tracking to obtain the best result. Work in this thesis is performed to achieve this, featuring studies informing the current running state and also improvements to our understanding of the crosstalk in the detectors, with a percentage and likely cause identified for the first time. Similarly, a technique is developed to measure the radial field, a key systematic uncertainty for the analysis, reducing it to subdominant levels. Improvement is also achieved in the calculation of the correction for tracker acceptance, reducing the uncertainty by a factor of four. Combined, these improvements allow the final data to have the maximum possible statistical sensitivity.

Commonly in Particle Physics, work like this can lead to real-world applications, with examples including proton therapy and better uncertainty of radiation detection and shielding. While the muon EDM search is currently very far from this, any new physics found could lead to new technology in the future. Additionally, the work in this thesis involves processing and analysing a large amount of data from the experiment. Developments from this could lead to improvements in data analysis for industry, where 'Big Data' is becoming increasingly important.

Contents

List of Figures

List of Tables

0.1 List of Acronyms

SM Standard Model (of Particle Physics)

BSM Beyond the Standard Model

QED Quantum Electrodynamics

MDM Magnetic Dipole Moment

BNL Brookhaven National Laboratory

ppm/b parts per million/billion

EDM Electric Dipole Moment

HVP Hadronic Vacuum Polarisation

QCD Quantum Chromodynamics

HLbL Hadronic light-by-light

CKM matrix Cabibbo-Kobayashi-Maskawa matrix

SUSY Supersymmetry

LHC Large Hadron Collider

MSSM Minimally Supersymmetric Standard Model

MRF Muon Rest Frame

ESQ Electrostatic Quadrupoles

IBMS Inflector Beam Monitoring System

DAQ Data Acquisition

CBO Coherent Betatron Oscillation

VW Vertical Waist

SCC Surface Correction Coils

NMR Nuclear Magnetic Resonance

CTAG Calorimeter Tags

SiPM Silicon Photomultiplier

DCA Distance of Closest Approach

ASDQ Amplifier Shaper Discrimination with Charge

TDC Time-to-Digital Converter

MIDAS Maximum Integrated Data Acquisition System

DQC Data Quality Cuts

MC Monte Carlo

FFT Fast Fourier Transform

FNAL Fermilab National Accelerator Laboratory

HV High Voltage

WC Width Cut

RMS Root Mean Square

RWC Relative Width Cut

PDG Particle Data Group

PDF Probability Density Function

FC Feldman-Cousins

VB Vertical Betatron

NDF Number of Degrees of Freedom

Chapter 1

Introduction

Despite the great success of the Standard Model (SM) in both describing and predicting the existence of a wide range of phenomena, it is known to be incomplete. For example, neutrinos are massless in the SM, and it doesn't include dark matter or gravity. These known gaps in the theory motivate experimental searches for physics beyond the Standard Model (BSM).

The search for new physics is divided into two frontiers. The energy frontier, which is focussed at colliders, aims to probe increasingly more energetic interactions to look for new physics. The precision frontier is dedicated to making precise measurements of the SM in order to compare these to theoretical predictions. Past measurements of fundamental magnetic moments of particles have a long history of being used for precision tests of the SM, with the magnetic moment of the electron contributing directly to the development of Quantum Electrodynamics (QED) [2]. It is therefore logical to also study the heavier leptons which are sensitive to a wider range of interactions.

The Muon $g-2$ experiment at Fermilab is the latest in a long line of experiments designed to precisely measure the muon's magnetic dipole moment (MDM) to high precision. Of particular interest is the 'anomalous' part of the this moment, which is the difference from the Dirac equation prediction of g = 2 [3]. This arises due to virtual particles interacting with the muon, and therefore includes the full panoply of the SM in the required theoretical calculations. Hence, the anomalous magnetic moment is a powerful probe of new physics. If a discrepancy is found between theory and experiment, this indicates that there is either something missing from the theory, or that the interactions within the SM are not fully understood.

The muon is a particularly good tool as the strength of the contributions of heavier particles to the magnetic moment scale with the lepton mass squared [4]. The election is 200 $\times$ lighter than the muon, so is less sensitive to potential BSM

physics, and the tau lepton, despite having high sensitivity due to its mass, has a very short lifetime and more complex decay channels, making the measurement significantly more challenging. This leaves the muon in the ideal central spot: heavy enough to have good sensitivity to new physics, but with a relativistic lifetime long enough that muons can be stored and studied.

The first measurement of the muon magnetic moment was performed at CERN in 1965. The CERN-I experiment used parity violation to generate a highly polarised muon beam and studied the angles of the decay products in order to analyse the spin [5]. As techniques improved, a storage ring was added to the experimental setup [6] in order to monitor the muons for a longer period of time. This is still used as the basic experimental design for muon $g-2$ experiments today. The CERN-I experiment motivated the inclusion of the first virtual particle loop corrections in the theoretical calculations. Several more iterations of the experiment led to additional corrections, motivating the need to measure this very precisely. The muon magnetic moment was also a very useful tool for studying the properties of the W and Z, which were difficult to study in colliders [7]. As a result, a new experiment at Brookhaven National Laboratory (BNL) was designed to improve the precision of the measurement, with a new single-piece ring magnet design which increased the uniformity of the magnetic field significantly.

The BNL experiment measured the anomalous magnetic moment to a precision of 0.54 ppm (parts-per-million) [8], with the central value found to be in disagreement with the SM with a significance of 3.7 σ. This sparked a new wave of efforts from both the theoretical and experimental communities to reduce uncertainties in order to determine if this anomaly was real. The most recent effort on the theory side comes from the Muon $g-2$ Theory Initiative, with the 2020 White Paper [9] defining a consensus single value for the anomaly. However, some recent results complicate the picture, which are discussed in detail in Section 1.1.1. On the experimental side, the Muon $g-2$ experiment at Fermilab was designed, built, and operated, with a target of 140ppb (parts-per-billion), a fourfold better precision compared to the BNL result [10]. If the central values do not move, this would lead to a > 7 σ difference, a conclusive sign of new physics.

Muon $g-2$ experiments are also capable of measuring the muon electric dipole moment (EDM). This arises due to the intrinsic charge distribution within the muon interacting with the electric fields generated by relativistic motion in a magnetic field. Unlike the MDM, a permanent EDM in a fundamental particle is predicted to be vanishingly small in the SM, of order 10^{-36} e·cm for the muon [11], which is well below the reach of current experiments. However, several BSM exten-

sions predict larger EDMs, making an EDM search a direct search for new physics. This is particularly interesting as a non-zero EDM would provide a new source of CP violation in the lepton sector and thus potentially help explain the matter-antimatter asymmetry in the early Universe. More details about the EDM and its implications are discussed in Section 1.1.2. The previous best measurements of the muon EDM were performed at BNL [1] and more recently the preliminary results from the Run 1 dataset of the Fermilab experiment [12]. The BNL EDM results limit the muon EDM to a value $< 1.9 \times 10^{-19}$ e·cm, and the Fermilab Run 1 results, if they measure a central value of zero, will limit the EDM at $< 2.0 \times 10^{-19}$ e·cm. This however only uses a small fraction of the total data the Fermilab experiment has collected, with a target final limit of order 10^{-21} e·cm [10].

This thesis focuses on work towards the EDM analysis at the Fermilab Muon $g-2$ experiment. The MDM and EDM searches are introduced alongside the theory behind the measurements, in particular lepton dipole moments, the SM predictions for the MDM and EDM, and a brief overview of BSM models which motivate both searches. Chapter 2 covers the experimental principles used to make the measurements themselves. Chapter 3 focusses on the experimental setup at Fermilab, discussing how the muon beam is generated, stored, and measured. Moving on from the introductory chapters, Chapter 4 describes the work of the author towards improving the performance of the tracking detectors, a key component of the EDM analysis. Chapter 5 discusses the implementation and results of a standalone EDM simulation written by the author to study the sensitivity of the Fermilab experiment to a muon EDM, including a consideration of which systematic uncertainties may limit the final results. The knowledge gained from this is used to help design a new measurement of the radial field, found to be the dominant systematic, detailed in Chapter 6. Finally, Chapter 7 pulls all these threads together into a full EDM analysis of the Run 2/3 data from the Fermilab Muon $g-2$ experiment. While the results of this analysis remain blinded, if a central value consistent with zero is measured, this analysis would improve the muon EDM limit by a factor of three. Finally, Chapter 8 summarises the results across the thesis and discusses the future outlook of muon EDM measurements.

1.1 Theory of muon dipole moments

1.1.1 Magnetic dipole moments

A charged particle in a magnetic field will interact with the field via its magnetic moment μ, which arises due to quantum mechanical particles having an intrinsic angular momentum, spin. This means that when placed in a magnetic field, they experience a torque, which aligns the spin direction of the particle with the field in order to minimise the potential energy [13]. This can be expressed with a Hamiltonian as:

$$H = -\vec{\mu} \cdot \vec{B} \tag{1.1}$$

, where $\vec{B}$ is the magnetic field vector. The magnetic moment $\vec{\mu}$ is defined as:

$$\vec{\mu} = g\left(\frac{Qe}{2m}\right)\vec{s}. \tag{1.2}$$

Here, Q is the charge, $\pm$ 1 for the muon, e the electron charge, m the particle mass, $\vec{s}$ the spin vector and g the dimensionless factor known as the gyromagnetic ratio or the landau g-factor. Classically, g is equal to one [14], however for quantum mechanical particles, the Dirac equation predicts g to be exactly equal to 2 for particles with spin 1/2 [3]. This 'leading order' prediction corresponds to the tree-level diagram interaction as shown in Figure 1.1 a). However, the g factor is modified due to the effects of virtual particles. These interact with the vertex in Figure 1.1 a) via loops, with an example involving a virtual photon being shown in Figure 1.1 b). The higher-order diagrams lead to the true value of g being slightly larger than 2, defining the anomalous magnetic moment a_μ:

$$a_\mu = \frac{g-2}{2}. \tag{1.3}$$

The anomalous part of the moment is predicted by theory and measured experimentally in order to test the SM [15]. The simplest correction to the tree level diagram is the addition of a single virtual photon, known as the Schwinger correction, which is illustrated in Figure 1.1 b). This correction leads to a contribution to a_μ equal to $\alpha/2\pi$, or 0.00116, where α is the fine-structure constant [16].

This correction was sufficient to make the theory and the experimentally measured value by Kusch and Foley consistent for the electron anomalous magnetic moment [2]. However, as theoretical knowledge and experimental measurements progressed, additional corrections were required, with the current state of the theory requiring the summation of thousands of higher-order diagrams.

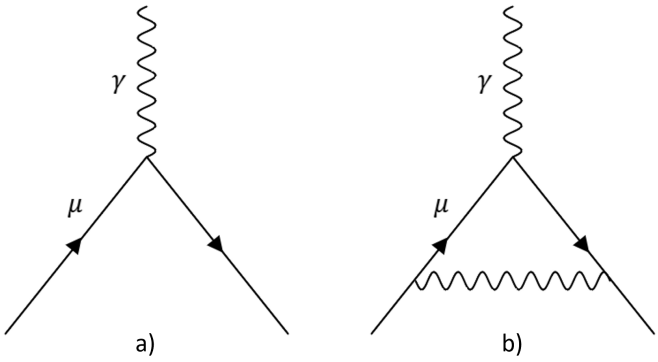

Figure 1.1: The tree level diagram of the muon-photon vertex, corresponding to g = 2 (left), and the largest contribution to the anomalous moment, which comes from a single photon loop (right).

The full theoretical calculation includes three distinct categories of corrections to g: QED, electroweak, and hadronic contributions:

$$a_\mu^{SM} = a_\mu^{QED} + a_\mu^{EW} + a_\mu^{HAD}. \tag{1.4}$$

Of these, the QED corrections are the largest and also the most well understood of the three groups, with an uncertainty of order 10^{-12} [9], well below the current experimental uncertainties. These interactions are loops containing only leptons and photons, leading to radiative corrections like the Schwinger correction. A higher order correction example is shown in Figure 1.2. These have been calculated up to tenth order [17] in the fine structure constant α, with order n diagrams being proportional to α^{n+1} and therefore suppressed by this factor.

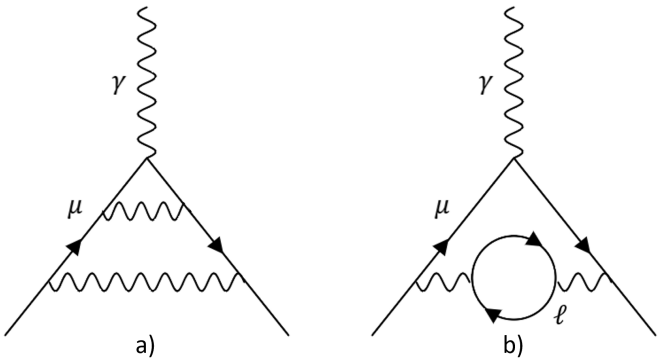

Figure 1.2: Two examples of higher-order QED contributions to the vertex interaction.

Electroweak corrections arise from interactions with the W, Z and Higgs bosons,

and are the smallest corrections to a_μ due to their large masses. The leading order contributions are shown in Figure 1.3, comprising a single Z boson or Higgs replacing the photon in the Schwinger correction diagram, and a muon converting into a W boson which is then recaptured with a muon neutrino. These corrections are also very well understood, only contributing an uncertainty of 10^{-11} [18].

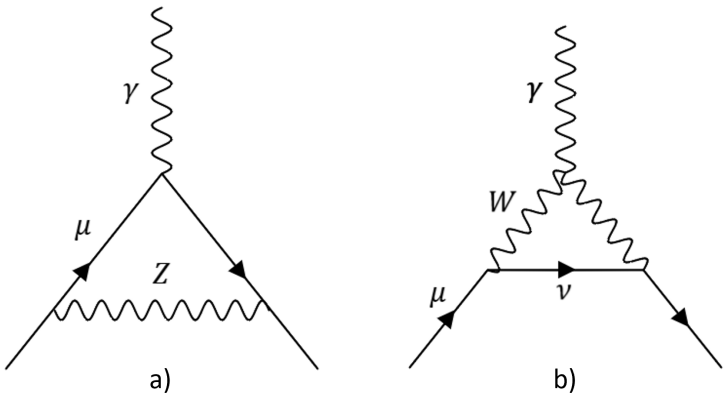

Figure 1.3: Diagrams showing the leading-order EW corrections, where a) can also feature a virtual Higgs.

The final set of corrections to a_μ are the hadronic corrections, which have a relatively small contribution to the actual value but dominate the uncertainty. There are two main types of diagrams possible, both illustrated in Figure 1.4. First is the hadronic vacuum polarisation (HVP), similar to the higher order QED loops (which can also be referred to as leptonic vacuum polarisation) featuring hadrons being generated by pair production rather than leptons. This is much harder to calculate than the leptonic version due to QCD meaning they cannot be calculated perturbatively. As a result, the evaluations rely on data-driven approaches using the experimentally measured cross-sections of $e^+e^- \rightarrow$ hadrons from experiments like BaBar and KLOE, with the current best calculation having an uncertainty of 4×10^{-10} [9].

More recently, lattice QCD methods have also been used to calculate these corrections. These have had significant improvements in precision in recent years. While the majority of lattice results agree with the data-driven methods, albeit with large uncertainties, in 2021 the BMW collaboration released a new lattice result which disagreed with all previous calculations and is more consistent with the experimentally measured value for a_μ [19]. This value is in tension with the data-driven methods at 2.7 σ and introduces a second discrepancy within the theory itself. This may suggest the true anomaly is in electron cross-sections

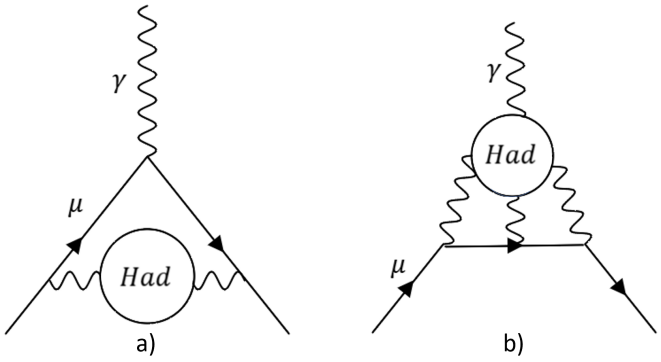

Figure 1.4: The two types of hadronic contributions to $g-2$, the hadronic vacuum polarisation and the hadronic light-by-light.

rather than the muon MDM. Other lattice groups are currently working to verify this result. Additionally, the CMD-3 experiment, which studies the e^+e^- cross-section to two pions, was recently found to be in tension with other data-driven results [20]. Work is ongoing to understand this and the impact on theoretical predictions.

The second type of hadronic diagram is the hadronic light-by-light (HLbL) corrections, where the vertex interaction involves an 'extra' three photons, as shown in Figure 1.4 b). This also has a large contribution to the uncertainty of a_μ, 19 $\times 10^{-11}$, but a smaller contribution to the value itself due to the three-loop structure. The HLbL correction is also calculated using a data-driven approach or lattice QCD.

The final combination of all these corrections as defined by the Theory Initiative combines all these corrections into a single theoretical value for a_μ [9]:

$$a_\mu^{\mathrm{WP2020}} = 116591810(43) \times 10^{-11}. \tag{1.5}$$

This is the most up-to-date consensus of the theory community, so is the value used for comparison with experimental results in this thesis.

The first experimental results from the Muon $g-2$ collaboration were released in 2021, based on the Run 1 dataset, which comprises 6 % of the total data so far:

$$a_\mu^{\mathrm{exp,FNAL}} = 116592061(41) \times 10^{-11}. \tag{1.6}$$

These results compared to the 2020 WP are shown in Figure 1.5, agreeing with the BNL result, with the discrepancy with the SM rising to 4.2 σ when combined [21]. While this is not yet above the 5 σ threshold, it is intriguing and suggests

that some new physics may be required, either in the magnetic moment or in the electronic cross-sections.

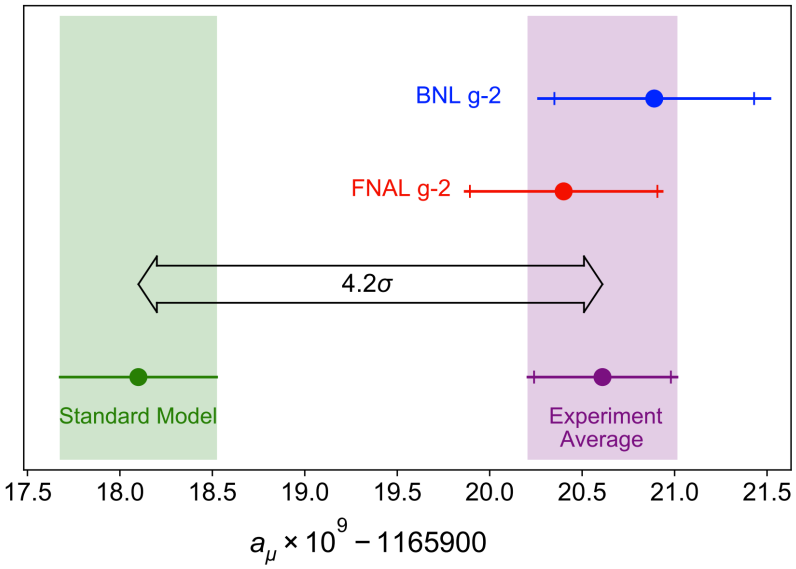

Figure 1.5: The results of the first 6 % of the data from the Fermilab Muon $g-2$ experiment, showing a good agreement with the BNL result.

1.1.2 Electric dipole moments

Similarly to how a particle with spin interacts with a magnetic field, it could also interact with an electric field. This would modify the Hamiltonian, giving it the full form:

$$H = -\vec{\mu} \cdot \vec{B} - \vec{d} \cdot \vec{E} \tag{1.7}$$

, where $\vec{d}$ is the electric dipole moment, defined as:

$$\vec{d} = \eta \left(\frac{Qe}{2mc} \right) \vec{s}. \tag{1.8}$$

Here, η is a dimensionless constant, analogous to g, which describes the strength of the coupling between the electric field and the particle spin. It can be expressed in terms of fundamental constants, including the magnitude of the EDM along the spin vector $d = \vec{d} \cdot \vec{s}$ as [15]:

$$\eta = \frac{4dmc}{Qe\hbar}. \tag{1.9}$$

The intrinsic properties of the electric and magnetic fields lead to the moments having differences in their fundamental transformation properties under charge

(C), parity (P) and time reversal (T). Both $\vec{\mu}$ and $\vec{d}$ are proportional to the spin vector, which is axial, so they both change sign under time reversal. The magnetic field is also an axial vector, so also changes sign, meaning that the combination $\vec{\mu}\cdot\vec{B}$ is overall invariant under time reversal. However, the electric field $\vec{E}$ is a polar vector and does not change sign under T, so the combination $\vec{d}\cdot\vec{E}$ is not invariant. Since CPT is assumed to be a global symmetry, as it is required for Lorentz invariance [22], this means that the EDM interaction must be CP-violating. This is particularly interesting as the only known CP-violating processes found so far have been in the quark sector ([23], [24]), arising due to the complex phase in the CKM matrix. CP violation is one of the mechanisms required for the matter-antimatter asymmetry of the early Universe [25], however the CP violation in the SM is currently insufficient to explain this. This means that any searches for new sources of CP violation are very well-motivated.

Particle EDMs have been studied in both fundamental and composite particles, with the earliest search being the neutron EDM to test parity symmetry in the 1950s [26]. Since then, the proton [27], electron [11] and muon [1] as well as updated neutron EDM measurements [28] have all been performed, with no permanent non-zero EDM discovered in any of these particles. In the SM, the muon EDM is predicted to be at most 10^{-36} e·cm, by assuming that the electron and muon are identical other than their masses and scaling the SM prediction of the electron EDM. This is very small due to multi-loop diagrams being required to get the first non-zero term, and is well below the reach of current experiments. This means that any observation of an EDM would be a sign of new physics. Various BSM theories, particularly those with a different mass scaling of interactions between generations of particles, predict a much larger EDM than the SM prediction, as discussed in Section 1.2.

1.2 Searching for BSM physics

Potential BSM contributions to a_μ are very varied, due to the loop structure involving the full SM and any other possible new particles. The sensitivity of the muon MDM to new physics is proportional to:

$$\delta a_{mu}^{NP} \propto \frac{m_\mu^2}{\Lambda^2} \tag{1.10}$$

, where m_μ is the muon mass and Λ the mass scale of new physics [15]. Using the BNL result, the upper bound of possible energies the muon can probe in this way is around 2 TeV. This is roughly comparable to the effective energies being

probed at collider experiments.

Many theories that attempt to explain this anomaly propose extensions to the SM, including the addition of supersymmetric (SUSY) particles, new bosons, for example dark photons, and other exotic particles like leptoquarks.

Supersymmetric theories extend the symmetries of the SM, introducing a correspondence between fermions and bosons, with each particle in the SM having an opposite superpartner. The possible theories of this type that could explain the anomaly are strongly limited by the lack of observation of SUSY in collider experiments. However, some parameter space still remains unexplored, including contributions from particles like the supersymmetric partners of the Higgs and W [29]. Additionally within this umbrella, theories such as the Two Higgs Doublet Model could also help explain the discrepancy [30]. Another popular theory is to introduce a new boson, for example a dark photon or a Z' boson [31], but many of these theories, including the dark photon, are disfavoured as the direction of the anomaly is not in the right direction to describe a_μ [32]. Finally, some theories of leptoquarks, which are particles that carry both lepton and baryon number, allowing them to couple with both quarks and leptons, could also explain the anomaly. For example [33] considers scalar leptoquarks which couple to both the muon and the hadronic part, modifying the photon vertex. While these theories can explain the anomaly, they show disagreements with LHC measurements, where direct searches for a particle decaying into quark-lepton pairs have placed tight constraints on possible leptoquark masses.

A large muon EDM could be generated by decoupling the EDM value from the mass scaling between generations of leptons. These BSM models also require a large CP violating phase. Otherwise, the electron EDM measurement, which is very precise, tightly constrains the muon EDM, as does the current a_μ anomaly. In general, theories that satisfy this either predict light new particles, or introduce a much stronger Higgs coupling than the SM has, a process called chiral enhancement [34].

The Two Higgs Doublet Model could explain a muon EDM, with the scaling of the election and muon EDM depending on $\tan(\beta)$, the ratio of the vacuum expectation values of the Higgs fields in the model. The extra particles introduced here allow the EDM to become non-zero with only two loops, generating an EDM of order 10^{-23} e·cm [35]. Other SUSY models that scale with $\tan(\beta)$ also work, like the Minimally Supersymmetric Standard Model (MSSM), which introduces three generations of supersymmetric lepton counterparts which decouple effects in electrons and muons, again removing the mass scaling [36]. Scalar leptoquarks

would also generate a larger EDM via virtual loops at the photon vertex, with the largest being of order 10^{-22} e·cm [34]. Finally, models which predict extra dimensions interact differently with different generations of leptons. This could generate an EDM as large as 10^{-20} e·cm, which is within the reach of the Fermilab experiment. However, these theories are less able to explain the magnetic moment anomaly, with very small predictions for a_μ, so are mostly disfavoured [37].

This brief overview of some possible new physics scenarios demonstrates that the muon EDM search is well motivated. A more comprehensive coverage is provided in [34]. Future experiments like muEDM at PSI [38] are planned to perform a measurement with an even higher sensitivity, narrowing the possible parameter space down even further.

Chapter 2

Experimental principles

This chapter focusses on the physical principles that underpin the Muon $g-2$ experiment, including details on both the a_μ measurement and the EDM search. A discussion of experimental choices made to maximise the sensitivity for both analyses is also presented.

2.1 a_μ measurement principles

The Fermilab Muon $g-2$ experiment uses a storage ring method, where a beam of polarised positive muons is injected into a ring with a vertical magnetic field. These muons decay via the weak force into positrons and neutrinos, and the positrons are measured to study the spin direction. The method and associated hardware is described in detail in Chapter 3, but the core of the measurement is that muons in a magnetic field have their spin precess around the field vector, with the frequency proportional to a_μ. Since spin direction is not a directly observable quantity, the experiment is instead designed to measure how the average direction of emitted positrons from the muon decay changes as a function of time. With a highly polarised beam of muons and only considering high-energy positrons, this is correlated with the spin direction due to parity violation.

2.1.1 Parity violating decays

Particles like pions and muons both decay via the weak force, which allows parity-violating interactions, first measured by Wu using the beta decay of Cobalt-60 in 1957 [39]. This means that rather than having an isotropic emission direction, the decays have a preferred direction. Since the weak force only couples to left-handed (LH) fermions and right-handed (RH) anti-fermions, this limits the possible kine-

matic outcomes. Considering the example of pion decay, the dominant decay of the positive pion is into positive muons and a muon neutrino [40]

$$\pi^+ \rightarrow \mu^+ + \nu_\mu. \tag{2.1}$$

In the massless limit, chirality (the handedness when interacting with the weak force) and helicity (the projection of the spin onto the momentum vector) are equal, so a LH fermion is one where the spin and the momentum vectors are antiparallel to each other. The neutrino mass is much smaller than the masses of the other particles in this interaction, allowing it to be considered massless, so the neutrino created will almost certainly be left-handed. Since pions are spin 0 particles, the muon must be emitted with a spin direction equal and opposite to the neutrino in order to conserve angular momentum. Similarly, the momentum vectors of the neutrino and muon must be equal and opposite in the rest frame of the pion to conserve linear momentum. The net result is that the muon is emitted with its spin and momentum vectors antiparallel in the pion rest frame. In the lab frame, this introduces a correlation between the spin direction and the momentum. Therefore, by generating relativistic pions, a highly polarised muon beam can be produced. The decay in the pion rest frame is shown in Figure 2.1, illustrating the constraints on the vectors.

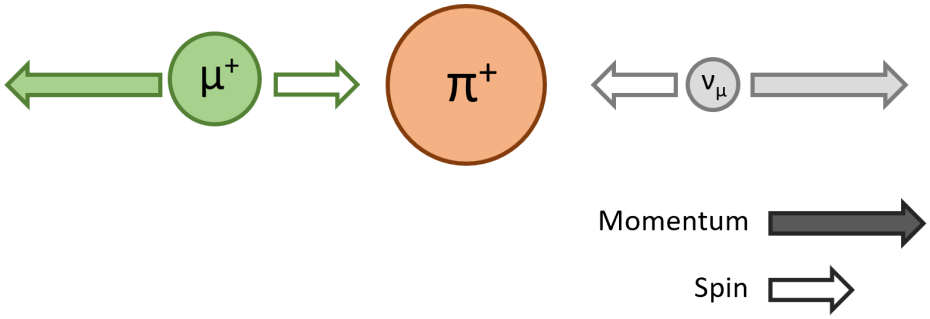

Figure 2.1: A diagram showing a pion decaying into a muon and a neutrino in the pion rest frame.

The highly polarised muon beam is then injected into the storage ring, where almost all of the muons decay into positrons. This is a weak decay, so is also parity-violating:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu_\mu}. \tag{2.2}$$

The decay is a three-body decay, with two neutrinos. The antineutrino must

be right-handed to interact with the weak force, which means it has spin and momentum parallel, while the neutrino must be left-handed as before. Considering the special case where the muon has the maximum possible energy, the positron must be emitted with spin in the same direction as the muon to conserve spin, with both neutrinos emitted in the opposite direction to the positron. The end result is that the highest energy positrons are emitted with their spin and momentum vectors parallel in the muon rest frame. This leads to a correlation between the energy of the positrons and their spin direction. The kinematics of this interaction are shown in Figure 2.2.

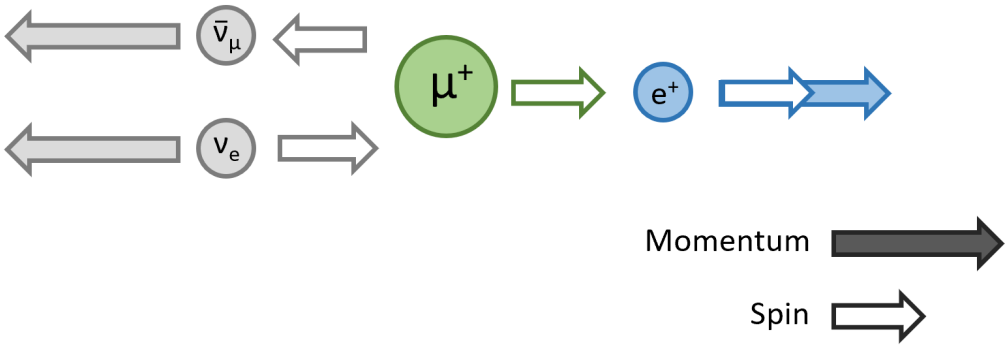

Figure 2.2: A diagram showing the muon decaying into a positron and two neutrinos in its rest frame.

This is what allows the experiment to measure the spin precession, achieved by applying an energy cut to only look at high energy positrons. The details of this cut are covered further in Section 2.1.3.

2.1.2 Larmor precession and cyclotron motion

A muon in a magnetic field experiences a torque, causing the spin direction to precess around the magnetic field direction, at a frequency ω_s:

$$\omega_s = g\frac{eB}{2m_\mu} \tag{2.3}$$

, where g is the gyromagnetic ratio, B the the magnetic field strength, and m_μ the muon mass.

As well as the spin precession, the muons will undergo motion due to the Lorentz force. Since the magnetic field provided by the storage ring is vertical, the force acts perpendicular to this, causing the muons to travel in a circle with

an angular frequency ω_c.

$$\omega_c = \frac{eB}{m_\mu}. \tag{2.4}$$

This motion causes the momentum vector to also precess around the field in the muon rest frame. Therefore, the frequency which can be measured is actually the difference between the two oscillations, the so-called anomaly precession frequency ω_a.

$$\omega_a = \omega_s - \omega_c = \left(\frac{g}{2} - 1\right)\frac{eB}{m_\mu} = a_\mu \frac{eB}{m_\mu}. \tag{2.5}$$

This means that, very conveniently, the measurable precession frequency ω_a is proportional to the anomalous part of the muon magnetic moment, allowing it to be measured directly.

Another consequence of Equation 2.5 is that if g was exactly equal to 2, as predicted by Dirac, the two frequencies would be equal, leading to the two vectors staying parallel while the muons travel around the ring. Both cases are illustrated in Figure 2.3.

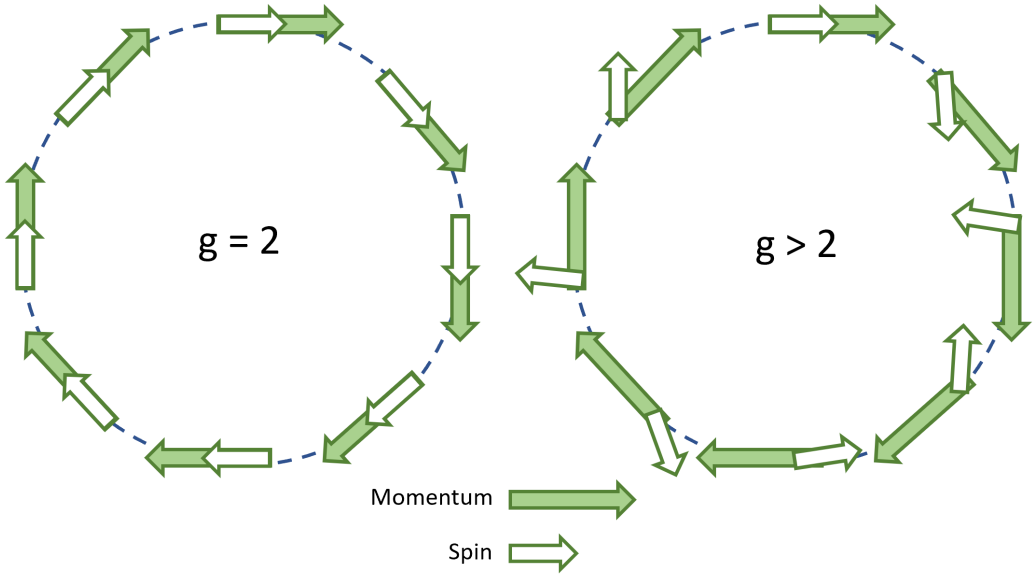

Figure 2.3: Diagram illustrating the anomalous precession in a storage ring. If g = 2, the momentum and spin would precess at the same rate, however for g larger than 2 they go out of phase. The difference between the two is the frequency ω_a. The effect over one orbit of the ring is exaggerated for the purpose of illustration.

2.1.3 Maximising the sensitivity to a_μ

If all positrons were considered with no limitations on the energy, no oscillation would be observed. When an energy cut is applied, the number of positrons seen above this energy threshold oscillates at the anomalous precession frequency ω_a. The choice of threshold is a trade-off between sensitivity to the oscillation and statistics, since setting the cut too high will remove a large fraction of the measured positrons, and setting it too low will lead to a lower amplitude.

Following the methodology of J. Miller [42], in the muon rest frame, assuming the positrons are relativistic, the differential probability of emission into a solid angle Ω is given by:

$$dp(\lambda, \phi) \propto N(\lambda)[1 + A(\lambda)\cos\theta]d\lambda d\Omega \tag{2.6}$$

, where $\lambda = E/E_{max}$ is the fractional energy compared to the maximum muon energy, θ the angle between the momentum vector and the spin vector, N(λ) the number distribution of the muons as a function of energy, and A(λ) an energy-dependent 'decay asymmetry' which quantifies the amplitude of the oscillation as a function of the energy fraction.

The number distribution N(λ) and the decay asymmetry A(λ) are given by the Michel parameters to 0th order:

$$N(\lambda) = 2\lambda^2(3 - 2\lambda), \tag{2.7}$$

$$A(\lambda) = \frac{2\lambda - 1}{3 - 2\lambda} \tag{2.8}$$

, both using the assumption of E >> m_ec. These functional forms are plotted in Figure 2.4 a). Transforming into the lab frame and introducing a time-dependent variation in the angle θ in Equation 2.6, with frequency ω_a and a phase ϕ, the number distribution is:

$$N(\lambda, t) = N_0 \exp(-t/\gamma\tau)[1 + A(\lambda)\cos\omega_a t + \phi(\lambda)] \tag{2.9}$$

, where τ is the boosted muon lifetime, γ the Lorentz factor, N_0 the initial number of muons and $A(\lambda)$ the boosted decay asymmetry function. This equation is key to the ω_a analysis, serving as the basic '5-parameter fit' used to extract the precession frequency, with the 5 parameters being ω_a, N_0, τ, A and ϕ. The lab frame boosted forms for $A(\lambda)$ and $N(\lambda)$ are given by equations 2.10 and 2.11, and

plotted in Figure 2.4 b).

$$N(\lambda) = \frac{1}{3}(\lambda - 1)(4\lambda^2 - 5\lambda - 5) \tag{2.10}$$

$$A(\lambda) = \frac{1 + \lambda - 8\lambda^2}{4\lambda^2 - 5\lambda - 5} \tag{2.11}$$

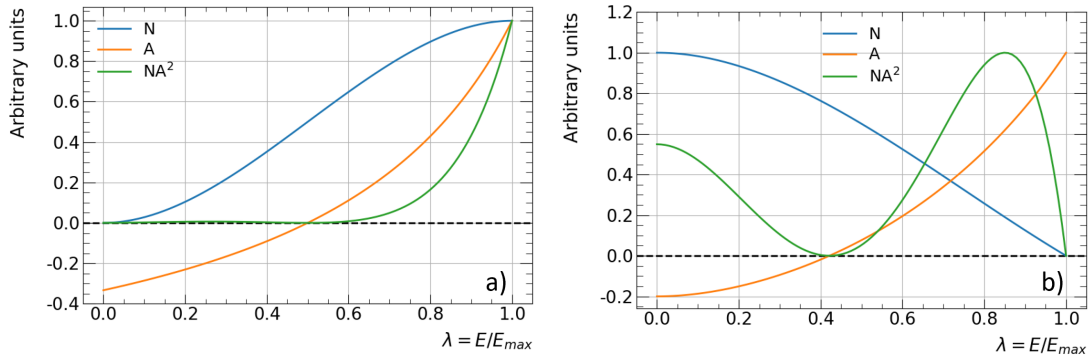

Figure 2.4: Plots of the number distribution and decay asymmetry functions as a function of the fractional energy, for the muon rest frame (a,left) and lab frame (b,right). The statistical figure of merit NA^2 is also plotted.

Since the aim is to extract the optimal energy cut for the analysis, one final step is needed, which is to integrate these expressions between some lower cut energy fraction, λ_{cut} and the maximum energy $\lambda = 1$. The final expressions obtained for the number density and the decay asymmetry function are plotted in Figure 2.5.

$$N(\lambda_{cut}) = \frac{1}{3}(\lambda_{cut} - 1)^2(-\lambda_{cut}^2 + \lambda_{cut} + 3) \tag{2.12}$$

$$A(\lambda_{cut}) = \frac{\lambda_{cut}(2\lambda_{cut} + 1)}{3 + \lambda_{cut} - \lambda_{cut}^2} \tag{2.13}$$

Finally, the statistical uncertainty on the precession frequency is considered, which can be obtained by fitting the above-threshold number oscillation with the five parameter fit from Equation 2.9:

$$\frac{\delta\omega_a}{\omega_a} = \frac{\sqrt{2}}{\omega_a \tau \sqrt{NA^2}}. \tag{2.14}$$

Assuming a fixed precession frequency due to a perfectly uniform field, this uncertainty is inversely proportional to the quantity NA^2, which is defined as the 'figure of merit' [43]. The ideal energy threshold is therefore at the point where NA^2 is maximised. Considering Figure 2.5, the optimal threshold is found to

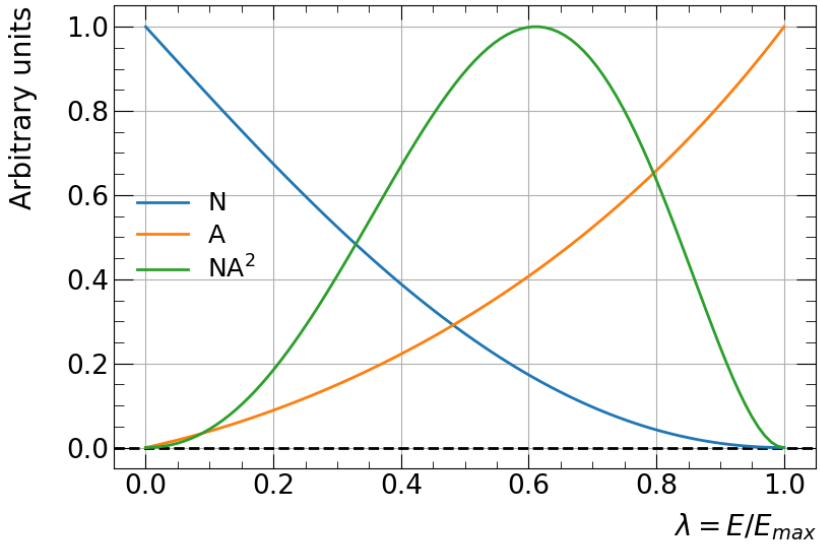

Figure 2.5: The number distribution, asymmetry function, and figure of merit from integrating between a defined cut point and the maximum energy. The optimal cut is a high-energy cut, as would be expected from the underlying physics.

be around 0.6 E_{max}, which corresponds to $\approx$ 1700 MeV [10]. These functional forms are ideal, with the real world values slightly impacted by detector effects like acceptance. Therefore, the true optimal cut value is shifted closer to 1900 MeV.

2.1.4 Extracting a_μ

Relativistic corrections

Since the muons are relativistic, they undergo an extra precession due to the effect of the rotating frame known as Thomas precession [44]. This leads to a small correction to the spin precession frequency ω_s:

$$\omega_s' = g\frac{eB}{2m_\mu} + (1-\gamma)\frac{eB}{\gamma m_\mu}. \tag{2.15}$$

The cyclotron frequency also gains a factor of γ in the denominator to account for relativistic effects:

$$\omega_c' = \frac{eB}{\gamma m_\mu}. \tag{2.16}$$

These two expressions are used to calculate the impact on ω_a:

$$\omega_a' = \omega_s' - \omega_c' = g\frac{eB}{2m_\mu} + (1-\gamma)\frac{eB}{\gamma m_\mu} - \frac{eB}{\gamma m_\mu} = \left(\frac{g}{2} - 1\right)\frac{eB}{m_\mu} = \omega_a. \tag{2.17}$$

Conveniently, the relativistic factors of γ fully cancel in the final expression, so the non-relativistic and relativistic anomalous frequencies are equal.

Electric field corrections

As part of the experimental setup, E fields are used to focus the muon beam, applied by electrostatic quadrupoles. The motivation and implementation of this is covered in more detail in Section 3.4. The introduction of this field leads to both ω_s and ω_c being altered, leading to a new expression for ω_a [10]:

$$\vec{\omega_a} = \frac{e}{m_\mu}\left[a_\mu \vec{B} - a_\mu\left(\frac{\gamma}{1+\gamma}\right)(\vec{\beta}\cdot\vec{B})\vec{\beta} + \left(a_\mu - \frac{1}{\gamma^2-1}\right)(\vec{\beta}\times\vec{E})\right]. \tag{2.18}$$

This expression can be simplified in two ways. Firstly, the second term in Equation 2.18 tends to zero under the assumption that the magnetic field and the muon momentum are perpendicular to each other. This is generally a very good approximation, however it is not perfect. The details of the correction required due to this are covered in Section 3.14.

Secondly, the third term in Equation 2.18 can also be set to zero by a careful choice of momentum. This 'magic momentum' is defined using the γ value, and sets the scale of the experiment along with the magnetic field, as the radius is proportional to it and the momentum.

$$\gamma_{magic} = \sqrt{1 + \frac{1}{a_\mu}} \tag{2.19}$$

This gives a value for gamma of 29.3, which corresponds to an operating muon momentum of 3.094 GeV and a ring radius R_0 of 7.112 m for a 1.45 T field. In practice there is a small momentum spread, which introduces another correction to ω_a. This is discussed in more detail in Section 3.14.

Final extraction of a_μ

While the extraction of a_μ using Equation 2.18 is possible, it requires knowledge of the muon mass, which is known to 50 ppb [45]. In general, ratios of fundamental constants are known to a higher precision than the values themselves, so it is desirable to write the expression for a_μ in terms of ratios rather than having it directly depend on the mass. Therefore, the expression for a_μ in Equation 2.5 is

rewritten in terms of more precisely known ratios:

$$a_\mu = \frac{m_\mu}{m_e}\frac{\mu_p}{\mu_e}\frac{g_e}{2}\frac{\omega_a}{\tilde{\omega_p}}. \tag{2.20}$$

Here, μ_p/μ_e is the proton-to-electron magnetic moment ratio, g_e the electron gyromagnetic factor and $\tilde{\omega_p}$ the Larmor precession of a free proton in the magnetic field of the experiment. The ratio of masses m_μ/m_e is known to 22 ppb, from muonium hyperfine splitting [45], and g of the electron is known to 0.17 ppt from single trapped electron measurements [46]. Finally, the ratio of the proton magnetic moment to the electron magnetic moment, which is measured using hydrogen maser measurements, is known to 0.3 ppb [47].

The introduction of $\tilde{\omega_p}$ allows for greater field accuracy as a frequency is measured, and is determined using an array of field measurement tools described in Section 3.6.1. All additional terms combined give a total uncertainty added in quadrature of 22 ppb.

2.2 Muon EDM measurement principles

The electric dipole moment also causes a torque on the muon spin vector, due to interactions with the electric field generated from relativistic motion in the magnetic field of the ring. This modifies the spin precession, adding an extra component and introducing a time-varying tilt of the precession plane $\pi/2$ out of phase with the $g-2$ spin precession. The size of this tilt is directly proportional to the size of the muon EDM, d_μ, as defined in Equation 1.9, and is illustrated in Figure 2.6.

This adds a new term to the equation for ω_a, which, operating under the two simplifying assumptions discussed in Section 2.1.4, becomes:

$$\vec{\omega_{a\eta}} = \vec{\omega_a} + \vec{\omega_\eta} = \frac{e a_\mu \vec{B}}{m_\mu} + \frac{e\eta}{2m_\mu}\left[\frac{\vec{E}}{c} + \vec{\beta}\times\vec{B}\right] \tag{2.21}$$

, where η is the dimensionless constant from the EDM definition in Equation 1.8. The tilt of the precession plane, δ, is proportional to the ratio between the two torques, which can be written in terms of η in this simplified case. Here, $\vec{\beta}$ and the $\vec{B}$ field are assumed perpendicular, and the quadrupole $\vec{E}$ field is 0 for a perfectly centred beam at the magic momentum:

$$\delta = \arctan\left(\frac{\omega_\eta}{\omega_a}\right) = \arctan\left(\frac{\eta\beta}{2a_\mu}\right), \tag{2.22}$$

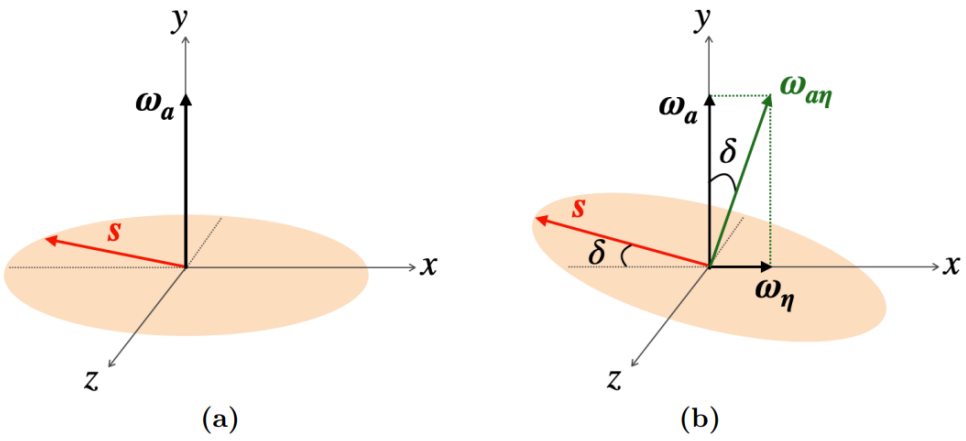

Figure 2.6: Illustration of the precession plane of the muon, first without an EDM in a), then with a large EDM in b) [48]. The ring magnetic field acts in the y direction, with the muon momentum along z, so the tilt is towards the centre of the ring.

This tilt is transformed into the lab frame, reducing it due to the boost. Defining y as the direction of the magnetic field and therefore the direction of ω_a, and x as the direction of the EDM tilt, the tilt angle is given by the ratio:

$$\delta = \arctan\left(\frac{\Delta x}{\Delta y}\right). \tag{2.23}$$

Applying a Lorentz boost, the y direction is perpendicular to the boost but Δx is parallel, so picks up a factor of $1/\gamma$:

$$\delta' = \arctan\left(\frac{\Delta x'}{\Delta y'}\right) = \arctan\left(\frac{\Delta x}{\gamma \Delta y}\right) = \arctan\left(\frac{\tan(\delta)}{\gamma}\right) \tag{2.24}$$

, where primed quantities are in the lab frame. Using the small angle approximation, which is valid as $\omega_a >> \omega_\eta$, this can be interpreted as an approximate $1/\gamma$ reduction in the tilt angle seen.

Using Equation 2.24 and substituting in for the original tilt angle from Equation 1.9, this gives a final conversion factor between the measured tilt in the lab frame and the value of the muon EDM d_μ, as:

$$d_\mu = \frac{e\gamma\hbar a_\mu}{2m_\mu c\beta}\tan\delta'. \tag{2.25}$$

In practice, the measured vertical angle of the measured positrons is further reduced by two effects. Firstly, both detector systems in the experiment are finite in size and do not measure or reconstruct every decay positron. The largest decay

angles are missed completely, leading to a net reduction in tilt by a factor that must be determined from simulation. Studies of this are presented in Sections 5.3.2 for a simple Monte-Carlo simulation and Section 7.2.8 for the Run 2/3 EDM analysis.

The second effect is known as the dilution, which is intrinsic to the decay. The measured positron vertical angles are momentum-dependent, as shown in Figure 2.9. As a result, not all positrons are emitted at the maximum possible angle, leading to a net reduction. This is discussed further in Section 2.2.3.

2.2.1 Searching for an increase in precession frequency

From Equation 2.21, an EDM not only tilts the precession plane, but also increases the overall precession frequency. Using the SM value, this is small and can be considered negligible. The current best limit on the muon EDM is 1.9×10^{-19} e·cm, which is also too small for this to be a viable search method.

This can be illustrated by calculating what EDM would be needed to explain the a_μ discrepancy between theory and experiment. The fractional change in the precession frequency is given by [41]:

$$\frac{\omega_{a\eta}}{\omega_a} = \sqrt{1 + \frac{\omega_\eta}{\omega_a}} = \sqrt{1+\delta}. \tag{2.26}$$

Expanding this to first order,

$$\frac{\omega_{a\eta}}{\omega_a} \approx 1 + \frac{\delta^2}{2} = 1 + \frac{\eta^2\beta^2}{8a_\mu^2}. \tag{2.27}$$

Setting a_μ to the Run 1 result, and taking the full difference to the 2020 WP as the discrepancy, this gives an upper bound on the EDM of:

$$|d_\mu| \leq 2.3 \times 10^{-19} e \cdot cm. \tag{2.28}$$

This value is larger than the EDM limit set by BNL [1], confirming that the a_μ discrepancy is unlikely to be from a non-zero EDM. This also demonstrates the low sensitivity of this method for an EDM search, motivating other methods which utilise direct measurements of the tilt of the precession plane. More detail about these methods and their relative strengths and weaknesses is covered in Section 3.15.

2.2.2 Momentum dependence of the maximum tilt angle

A further consequence of the Lorentz boost is the reduction in maximum measurable vertical angle as the energy/momentum of the positrons increases. This leads to lower momentum positrons having greater sensitivity to the EDM than might be expected, and requires an adjustment to the EDM decay asymmetry in the boosted frame discussed in the next section.

The maximum allowed vertical angle after a boost can be found using kinematic arguments. This angle can be defined in terms of the momentum components of the decay as:

$$\delta' = \arctan\left(\frac{P'_T}{P'_z}\right) \tag{2.29}$$

, where P'_T is the transverse momentum, and P'_z the momentum in the z direction, with a total momentum lab frame momentum P'. As the Lorentz boost is along the z direction, the transverse momentum before and after the boost is equal:

$$P'_T = P_T = \sqrt{E^2 - P_z^2}. \tag{2.30}$$

To maximise the angle, the transverse momentum must be maximised, so the energy in the rest frame must be E_{max}, equal to $m_\mu/2$. Similarly, P'_z can be expressed in terms of the rest frame P_z as:

$$P'_z = -\beta\gamma E + \gamma P_z \approx \gamma\frac{m_\mu}{2} + \gamma P_z \tag{2.31}$$

, under the assumption that the muons are highly relativistic, so $\beta \approx 1$. Rearranging this for P_z and substituting into Equation 2.30, this gives the functional form for P'_T in terms of the z momentum. Finally, since the majority of the momentum is in the z direction, $P'_z \approx P'$, so the maximum vertical angle as a function of momentum becomes:

$$\delta' \approx \arctan\left(\frac{\sqrt{m_\mu P'/\gamma - (P'/\gamma)^2}}{P'}\right). \tag{2.32}$$

This function is plotted in Figure 2.7 overlaid onto the distribution seen in data, with good agreement seen.

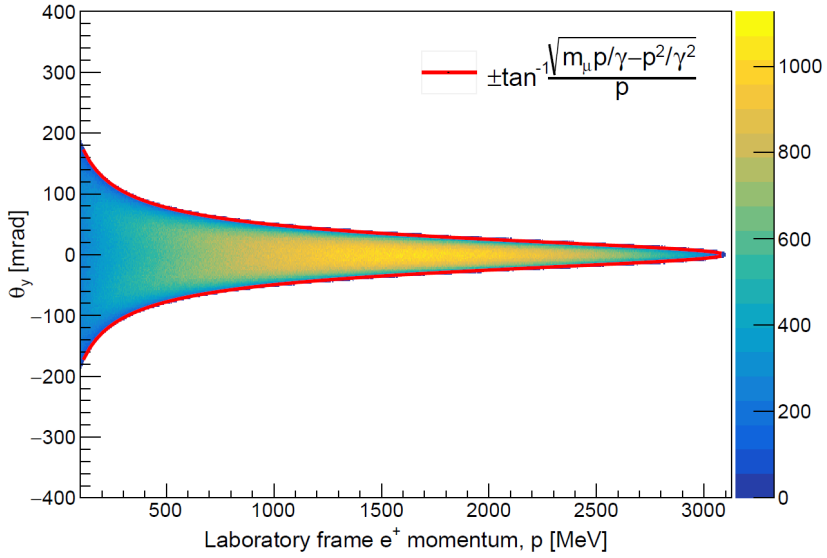

Figure 2.7: The variation of maximum vertical angle as a function of momentum, with colour corresponding to counts/bin and the expected functional form from this derivation plotted on top, reproduced from [12]. This limits the EDM sensitivity at high momentum values somewhat.

2.2.3 Maximising the EDM sensitivity

Similarly to the ω_a analysis, the tilt angle has a momentum-dependent asymmetry, so optimal cuts must be chosen in order to maximise the EDM sensitivity. The procedure by P. Debevec in [49] is followed. An EDM decay asymmetry function A(λ) is defined based on the up/down vertical asymmetry, with N(λ) remaining as for ω_a. Following the same procedure as before and integrating between an energy threshold λ and the maximum energy, the following expression is found for the asymmetry in the muon rest frame:

$$A_{EDM}(\lambda) = \frac{\sqrt{\lambda(1-\lambda)}(1+4\lambda)}{5+5\lambda-4\lambda^2} \tag{2.33}$$

, with the statistical figure of merit, still defined by NA^2, having the functional form:

$$NA^2_{EDM} = \frac{1}{3}\,\frac{\lambda(1-\lambda)^2(1+4\lambda)^2}{5+5\lambda-4\lambda^2}. \tag{2.34}$$

These functions, along with the number distribution, are plotted in Figure 2.8. This shows a different shape to the ω_a sensitivity plots, with the figure of merit being maximised for central energy values.

These distributions are then boosted into the lab frame, with the width of the vertical distribution as a function of momentum taken into account. The lab frame

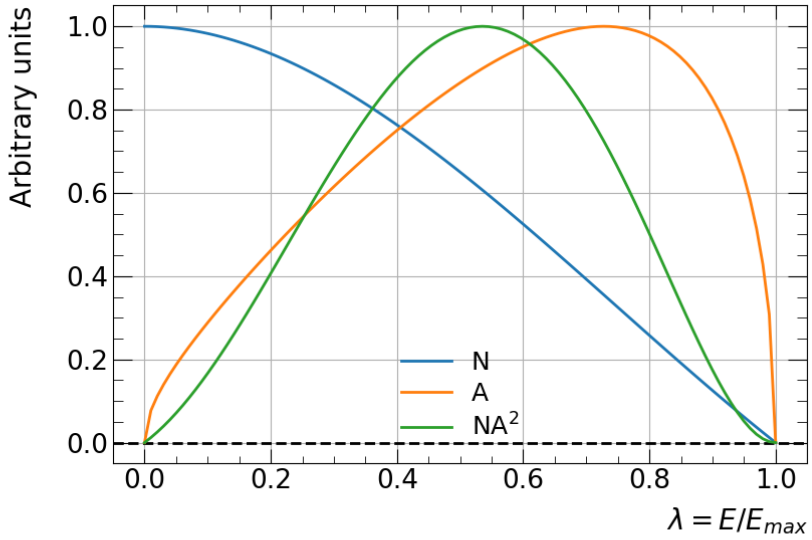

Figure 2.8: The number distribution and asymmetry for the EDM in the muon rest frame, showing that the optimal momentum cuts keep the central region.

decays at the point of maximum tilt are described by [50]:

$$< \theta_y(\lambda) >= \frac{(\lambda - 1)(2\lambda + 1)}{(4\lambda^2 - 5\lambda - 5)} \frac{\sin \delta}{\gamma} \tag{2.35}$$

, where $< \theta_y(\lambda) >$ is the average vertical angle as a function of energy fraction, and $\sin(\delta)/\gamma \approx \delta'$, the lab frame tilt, under the small angle approximation. This is plotted in Figure 2.9. In an ideal case, this function would exactly describe the maximum tilt angle as a function of the muon's energy. This energy dependence leads to a reduction in the measured tilt which must be corrected to extract the EDM, using the dilution function, defined as:

$$d_{EDM} = \frac{1}{\delta'} < \theta_y(\lambda) > . \tag{2.36}$$

This function combines the energy dependence of the decay with the reduction in angle due to boosting into the lab frame from the muon rest frame. In practice, the muon decay is slightly impacted by radiative corrections, which reduces the dilution function by a constant factor, as covered in Section 7.2.8.

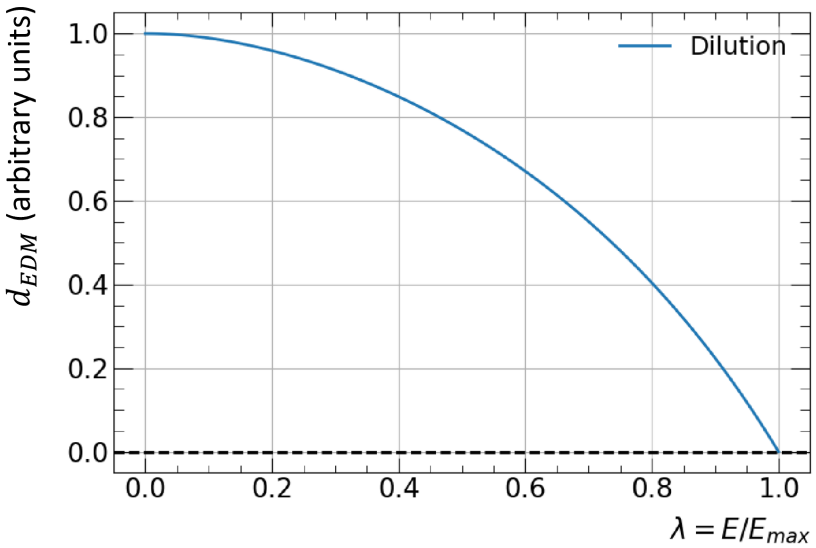

Figure 2.9: The functional form of the momentum-dependent dilution of the tilt angle, as a function of λ. A value of 1 indicates there is no reduction in average vertical angle from this effect.

Chapter 3

The Muon $g-2$ experiment at Fermilab

The Fermilab Muon $g-2$ experiment is the newest in the series of storage ring experiments designed to measure the anomalous precession frequency of the muon. It uses a similar setup to the BNL experiment, with improvements from the beamlines at Fermilab and to the experimental components themselves. The target is a fourfold improvement in the uncertainty of a_μ compared to the previous measurement at BNL, reducing the uncertainty to 140 ppb [10]. This was predicted to require around 21$\times$ the data collected by BNL, which as of the writing of this thesis, has just been achieved. Another aim of the experiment is to measure the muon EDM to world-leading precision, which can be done parasitically using the same data as the ω_a analysis.

The base technique of the measurement is to inject a beam of highly polarised positive muons into a storage ring, which applies a uniform vertical magnetic field, causing the muon spin direction to precess. The muons decay into positrons, which curve inward due to the magnetic field before being detected by an array of detector systems around the inside of the ring. This chapter covers the technical details of the experimental setup, including the production, injection and storage of the muon beam, as well as details about the trackers, calorimeters and other detector systems needed to make the measurement. An outline of the methodology of the MDM and EDM analyses using these systems is also presented and discussed. A photograph of the experiment in the experimental hall is shown in Figure 3.1, with a schematic diagram showing the location of all components in Figure 3.2.

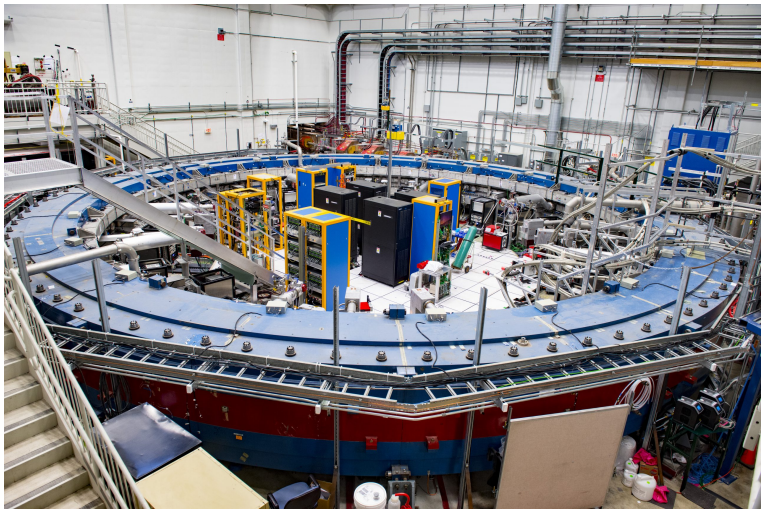

Figure 3.1: A photograph of the $g-2$ experiment, located at MC-1 at Fermilab's Batavia site.

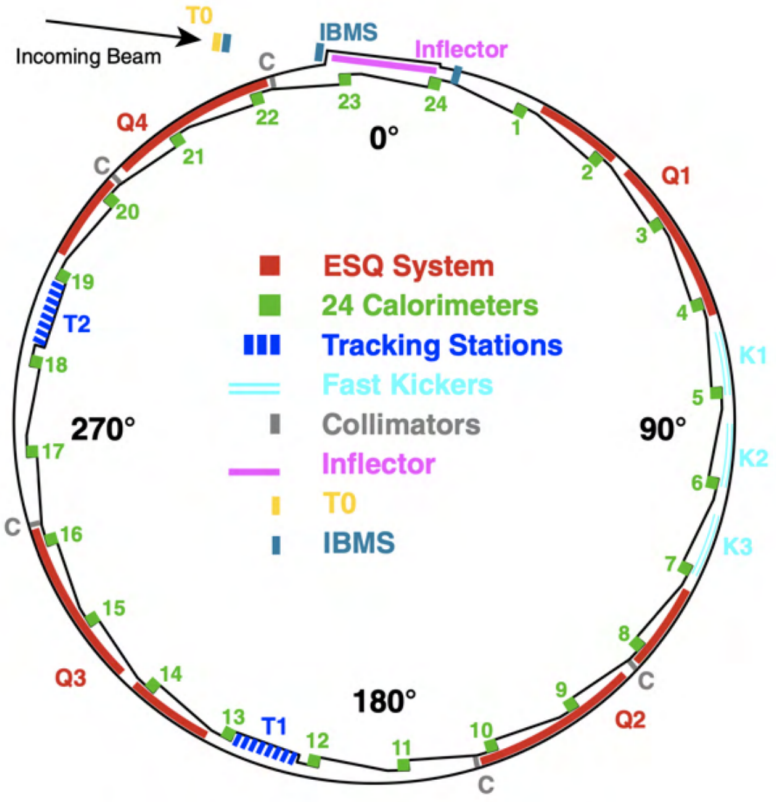

Figure 3.2: A diagram of the $g-2$ ring, showing the locations of the ring systems (kickers, quadrupoles (ESQ), inflector and collimators) and detectors (calorimeters, trackers, T0 and IBMS) discussed in this chapter. The 180° point lies under the stairs going down into the ring centre in the above photo. Reproduced from [52].

3.1 The muon beamline at Fermilab

The Fermilab accelerator complex uses protons as the starting point for the muon beamline, generated by stripping electrons from hydrogen atoms. A diagram of the complex is shown in Figure 3.3. The protons are accelerated by a linear accelerator (Linac) into a synchrotron called the Booster ring. This produces a bunched beam of 8 GeV protons, with each bunch containing around 4×10^{12} protons. This beam is passed to a second synchrotron known as the Recycler, where the bunches are further separated into smaller clusters each containing about 1/4 of the initial protons. This process is performed to reduce the eventual expected flux of muons to a manageable rate for the experiment, which prevents systems like the data acquisition (DAQ) and detectors from becoming overloaded. As part of this re-bunching process, the width in time of each bunch is set to 120ns, a value chosen to be less than the cyclotron period of the $g-2$ ring, which is 149.2 ns. This ensures clean injection into the ring without overlap. To ensure proper time separation between these bunches, they are delivered in groups of 8 to the target hall separated by a time gap of at least 10ns. The final beam delivery is 16 bunches per 1.4s [51].

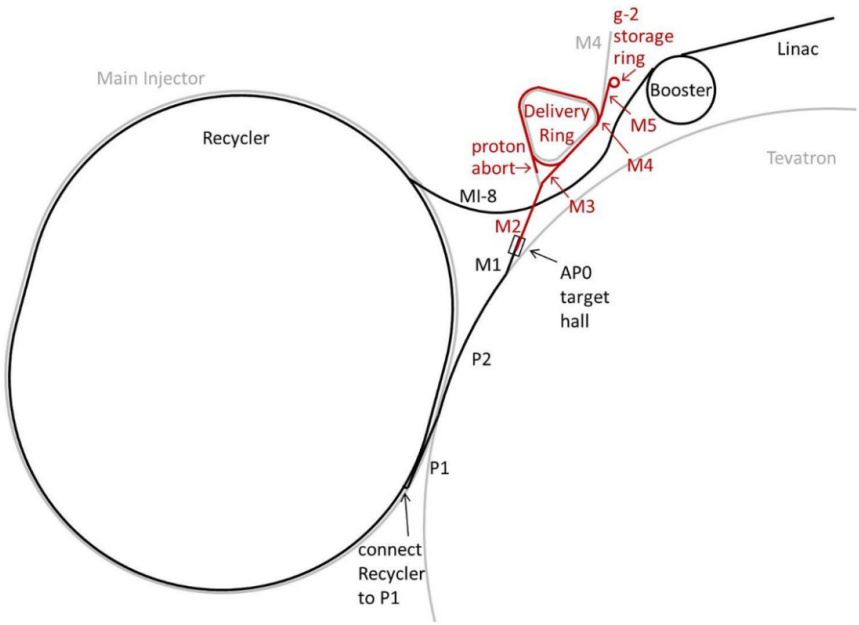

Figure 3.3: A diagram showing the Fermilab accelerator complex, with the path of protons outlined in black and the resulting secondary beams (pions and muons) outlined in red. Protons are accelerated in the Linac and Booster before being re-bunched in the Recycler and sent to the target hall. This generates a beam of pions which travel around the delivery ring until they decay into muons, which are injected into the experimental hall. Image from [10].

Inside the target hall, the bunches are incident on a pion production target, made of Inconel, a nickel-iron alloy. Here, positive pions are generated as a result of the proton-nucleus interactions:

$$p + p \rightarrow p + n + \pi^{+}. \tag{3.1}$$

Other interactions are possible, producing scattering products like negative pions and deuterons, so the positive pions must be separated from these. The scattering products are focussed using a lithium lens and sent through a pulsed bending magnet. This allows for the selection of positive pions of a certain momentum, chosen to be within 10% of 3.1 GeV, sending them into the M2 and M3 beamlines leading to the delivery ring. This momentum is chosen to be slightly above the final desired 'magic' momentum for the muons in the $g-2$ storage ring. As the pions travel down the beamlines and around the delivery ring, they decay into positive muons, and any residual protons or other particles are removed using a kicker using their time separation from the pions. As a final step, muons with the desired magic momentum are selected, with a momentum spread of 10% achieved. These are directed into the Muon $g-2$ experimental hall down the M5 beamline. The beam passes through four electrostatic quadrupoles to focus the beam, which is then injected into the $g-2$ storage ring.

Fermilab was chosen as the site for the new experiment for a number of reasons, including the design of this accelerator complex, which can deliver a very high intensity muon beam. Additionally, the design with long pion beamlines and delivery ring is key for reducing the initial pion contamination. This allows experimental systems to start collecting data much earlier in a fill than previous experiments [10].

3.2 Injection

The muon beam entering the ring would undergo a large deflection upon entering the storage ring magnetic field. Therefore, a superconducting niobium-titanium-copper-aluminium coiled magnet is used to almost completely counteract the ring field, allowing the beam to enter undeflected [53]. A diagram of the inflector, and the field it provides, is shown in Figure 3.4. This marks the start of a 'fill' of muons, a time used to synchronise other ring components, which is measured using a T0 detector. Additionally, the beam profile is monitored before and after injection using the IBMS (inflector beam monitoring system) detectors. Section 3.9 gives more detail about these auxiliary detectors.

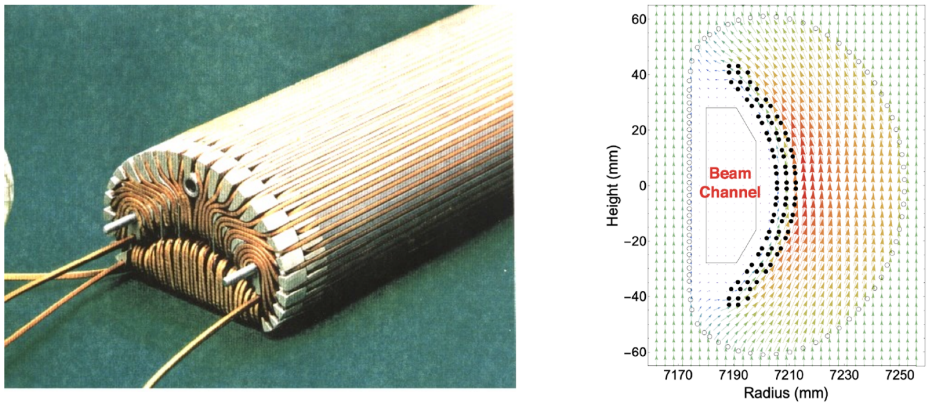

a) External view of the inflector, showing the coils and geometry

b) The magnetic field topology generated by the combination of the inflector and the ring magnet

Figure 3.4: A photo of the inflector coils and a plot of the magnetic field around the inflector, with contributions from both the ring magnet and the inflector itself. The beam channel region is field-free, which is required for the beam to enter without deflection. Photo and plot reproduced from [53].

3.3 The kicker system

The inflector injects on a wider orbit to prevent it blocking part of the storage region, so the muons must be moved towards the centre to avoid colliding with the inflector after one orbit. Three electromagnetic kickers located downstream from the inflector apply a very short lived but strong magnetic field to 'kick' the muons onto the optimal orbit R_0. These each consist of two long aluminium plates positioned inside the vacuum of the storage ring. Although aluminium is not as good for this purpose as other materials like iron, it is important for the kicker to not perturb the magnetic field, so non-ferric materials must be used. The duration of the kicker field is chosen to be 120ns, which covers the full width of the fill [54]. The shape of the kicker pulse is shown in Figure 3.5, demonstrating how the main peak of the pulse is contained within the required 120ns window.

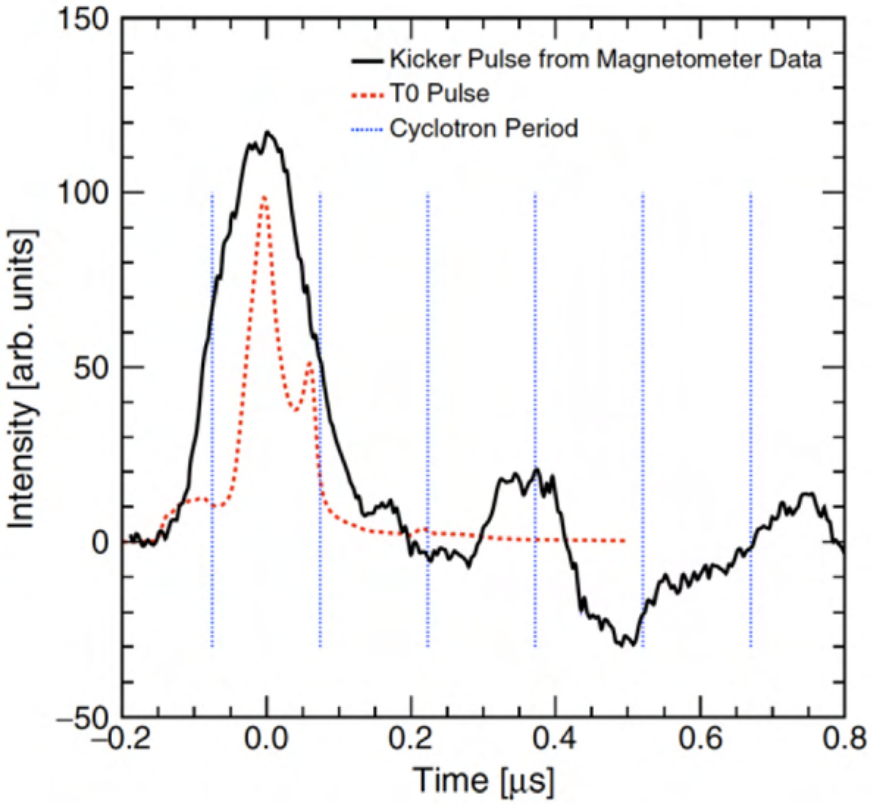

Figure 3.5: The intensity pulse shape for the kickers. The light blue lines on the plot show the cyclotron period, illustrating how the main peak intensity lies between the point of injection and the first revolution of the muon beam. The T0 pulse is the incoming beam time shape measured by the T0 detector as detailed in Section 3.9, which also shows good overlap. Plot from [54].

3.4 The electrostatic quadrupoles

The ring magnet does not provide vertical focussing, without which the muons would spread out and drift out of the ring. To counteract this, four electrostatic quadrupoles (ESQ) are arranged symmetrically around the ring, generating a field which provides vertical focussing of the beam. The strength of this is characterised as a function of vertical magnetic field strength B_0 by the field index n:

$$n = \frac{\kappa R_0}{v B_0} \tag{3.2}$$

, where κ is the electric field gradient in the vertical direction, $\delta E_y/\delta y$, R_0 the ideal storage radius and v the muon velocity [55].

The quadrupoles are also used to remove the tails of the momentum distribution of the muon beam at early times, since the muons are injected with a 10% spread around the optimal value. Muons not at the magic momentum impact the analysis

directly by making the $\beta.B$ term in Equation 2.18 slightly non-zero. To reduce the impact of this it is beneficial to remove these muons. Additionally, since the ring is designed to store muons at a specific curvature, any deviation from this makes it more likely that the muons will fall out of the storage region. These 'lost muons' also impact the ω_a analysis. The ESQ are operated asymmetrically to achieve this, moving the beam vertically and horizontally in the first 7 μs of the fill. This moves the outer edges of the beam into copper collimators designed to scatter the incident muons. This process is called 'scraping' with the resulting distribution having a momentum spread of 0.15 % [56].

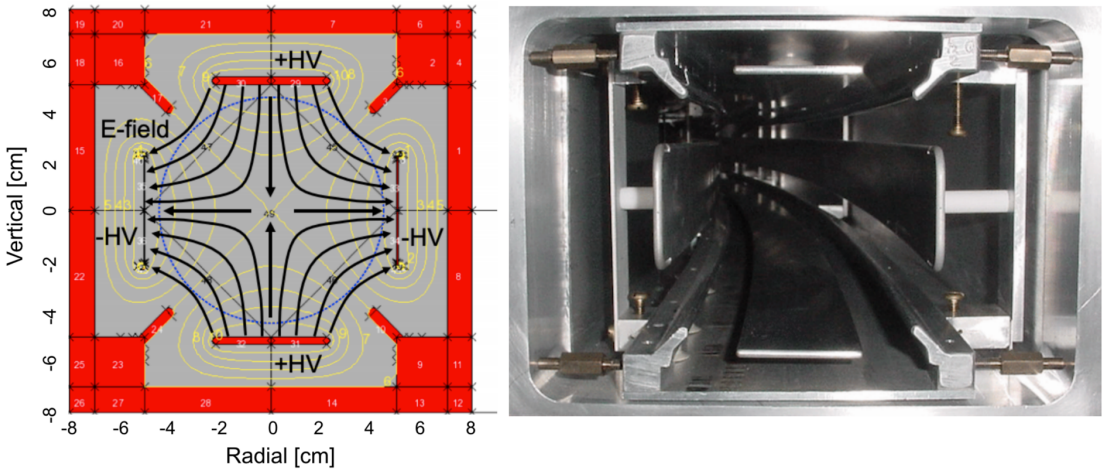

a) The electric field of the ESQ

b) A photo of the quadrupole plates

Figure 3.6: a): a diagram of the electric field lines generated by the quadrupoles, with the field itself drawn in black and the yellow lines denoting equipotentials. b) a photo of the plates installed in the ring that generate this electric field. Images from [10].

3.5 Beam dynamics

Due to the varying forces on the beam in the storage region, it is not static, instead evolving dynamically as a function of time. Since much of the motion is oscillatory, it is important to fully understand the beam motions and their impact on the measurements. The required corrections from these motions are discussed for the ω_a analysis in Section 3.14, and for the EDM analysis in Section 7.2.4. In this section, the relevant motions to these analyses are discussed. A more comprehensive discussion of beam motions can be found in [57].

3.5.1 Betatron motions

The restoring forces provided by the electrostatic quadrupoles cause the beam to undergo simple harmonic motion in both the vertical and radial directions. In the vertical direction, there is only the linear restoring force from the quadrupoles, whereas in the radial direction there are contributions from both the ESQ field and the vertical magnetic field. These motions can be expressed using the following equations of motion:

$$\frac{d^2x}{dt^2} = -\omega_c(1-n)x, \tag{3.3}$$

$$\frac{d^2y}{dt^2} = -\omega_c(n)y, \tag{3.4}$$

where ω_c is the cyclotron frequency introduced in Equation 2.4 and n is the field index defined in Equation 3.2. . These equations of motion can be solved by substituting oscillatory solutions of the form:

$$x(t) = A_x \cos(\omega_x t - \phi_x), \tag{3.5}$$

$$y(t) = A_y \cos(\omega_y t - \phi_y). \tag{3.6}$$

This shows that the resultant motion is an orthogonal set of oscillations, with frequencies of $\omega_c\sqrt{n}$ in the vertical direction and $\omega_c\sqrt{1-n}$ in the radial direction.

The radial betatron frequency is much larger than the precession frequency, so alone would not impact the ω_a analysis. However, the oscillations combine and introduce a coherent betatron oscillation (CBO) which is at a similar frequency to the muon precession frequency [56]. This arises due to a combination of two effects. Firstly, all the muons in the beam exist in a very similar phase space due to their common injection momentum. Secondly, as the detectors can only sample the beam once per rotation and $\omega_c\sqrt{1-n} \neq \omega_c$, each measurement sees the beam at a slightly different modulated position. The oscillation seen in the detectors therefore has a frequency:

$$f_x^{CBO} = \omega_c - \omega_x = \omega_c(1-\sqrt{1-n}). \tag{3.7}$$

If this frequency was identical to the precession frequency, it would be indistinguishable, making the precession frequency very difficult to disentangle and measure. Therefore, part of the experimental design requires choosing focussing strengths that give a distinct frequency for the CBO.

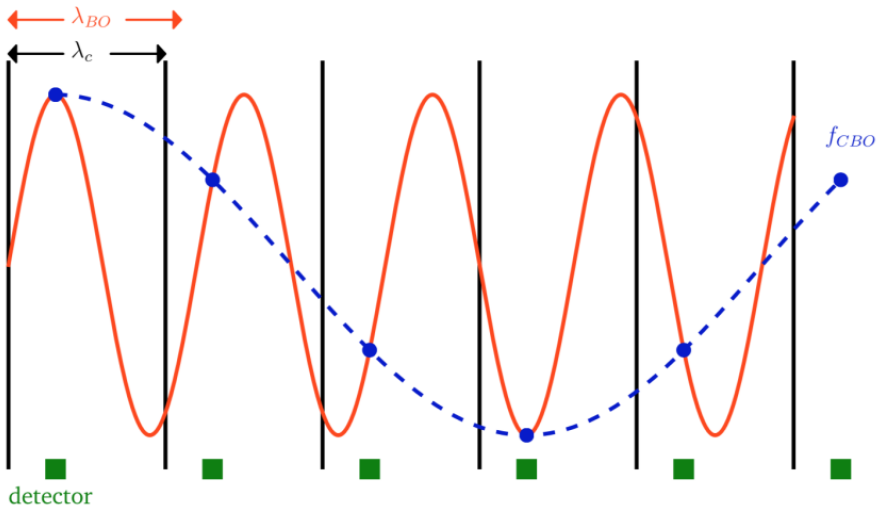

Figure 3.7: Diagram showing how the CBO arises from selective sampling of a higher-frequency oscillation. Each blue point corresponds to a time measurement in a detector separated by the cyclotron frequency, which leads to a slow aliased oscillation being seen rather than the fast one. Reproduced from [41].

The coherent oscillation in the vertical direction is much higher frequency than the radial direction oscillations, and also impacts the analyses, particularly the EDM.

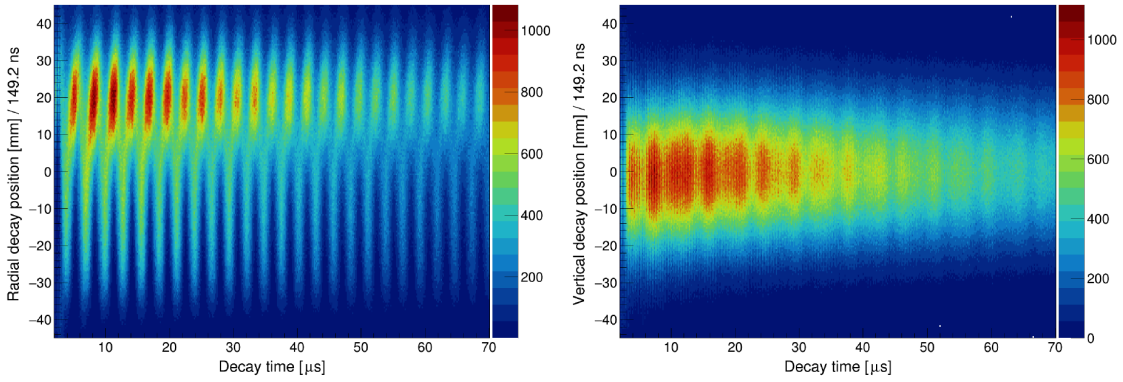

Figure 3.8: Plots showing the radial (left) and vertical (right) CBO oscillations, as measured using tracker data. The difference in frequency is quite clearly visible. Plots made by S.Grant [12].

The vertical width of the beam also oscillates, at a frequency of twice the vertical betatron frequency. This oscillation is seen at an aliased frequency relative to the cyclotron frequency, and in this form is called the vertical waist (VW):

$$f_{VW} = \omega_c - 2\omega_y. \tag{3.8}$$

Since the EDM analysis measures the average vertical angle as a function of time, the VW can also contribute either to the central values or to the uncertainties as a function of time, as well as impacting acceptance. A consideration of this, and other beam dynamics for the EDM analysis, is presented in Section 7.2.4.

3.5.2 Fast rotation

Due to the momentum spread of the beam, not all muons in the beam will take the same length of time to perform one orbit of the ring. As a result, muons at a smaller radius will start to overtake those at a larger radius over a fill. This leads to a rapid modulation at early times in the detectors, with the effect decaying away as the muons spread to fill the entire ring. This is called the fast rotation, and while it does rapidly decay, it does still impact the analyses. For the EDM analysis, this can be removed by combining data in cyclotron periods and applying a time randomisation, discussed in Section 7.2.4.

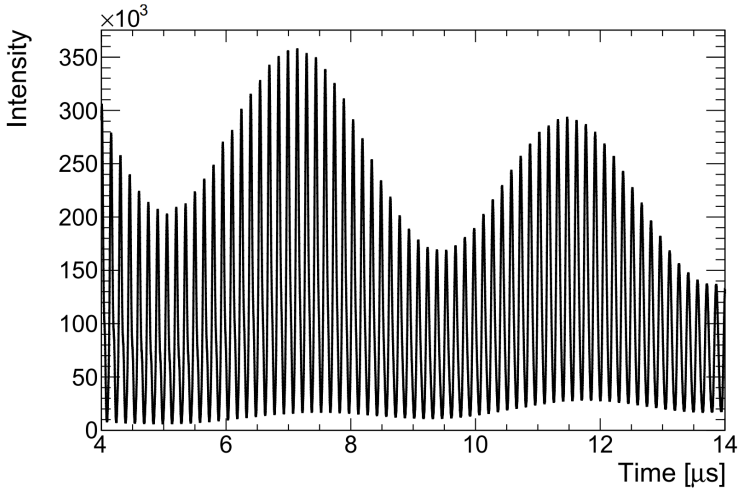

Figure 3.9: The number of positrons seen as a function of time between 4 and 14 μs in a fill, binned at 1 ns. A modulation is seen due to the fast rotation effect. Image from [57].

3.5.3 The closed orbit distortion

Small perturbations of the vertical magnetic field can lead to much more complex oscillatory shapes of the beam, with the equilibrium radius changing for different azimuthal positions. These changes in shape are called distortions to the closed orbit, and can be caused by changes in electric and magnetic fields as well as any local inhomogeneities. An illustration of this effect is shown in Figure 3.10.

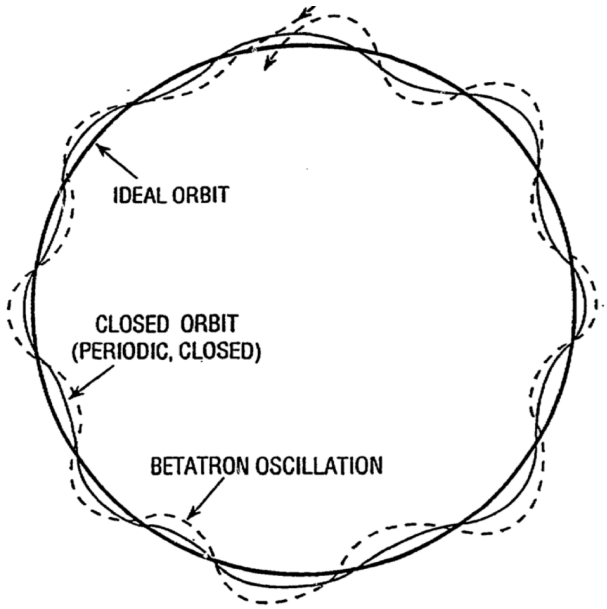

Figure 3.10: A diagram of the azimuthal position around the ring, with an exaggerated comparison of the betatron oscillation and the final orbit with the distortions. Reproduced from [58].

These changes play a large part in measuring the dominant systematic effect of the EDM analysis, the radial field. This is covered in much more detail in Chapter 6.

3.6 The storage ring and magnetic field

The storage ring is the central component of the entire experiment, and has been repurposed from BNL. It is designed to confine muons in a circular orbit at the magic momentum ($R_0 = 7.112$ m) with a 1.45 T dipole magnetic field. The main magnet is designed in a c-shaped yoke pattern, with the muon storage region at the centre. One side is open to allow component positioning close to the beam. Vacuum chambers are installed along the inside of the magnet to maintain a vacuum in the muon storage region, with ports along the inside edge. The magnet itself is powered by superconducting coils and a current of ≈ 5170 A, with two coils placed around the inside of the storage region and one placed on the outside [59]. A cross section of the magnet yoke, showing the location of these coils and other beam shimming and monitoring systems, is shown in Figure 3.11.

The uniformity of the magnetic field is a very important part of the precession frequency analysis, requiring a 1 ppm maximum variation in the storage region. The magnetic field would ideally be fully vertical with no significant radial or

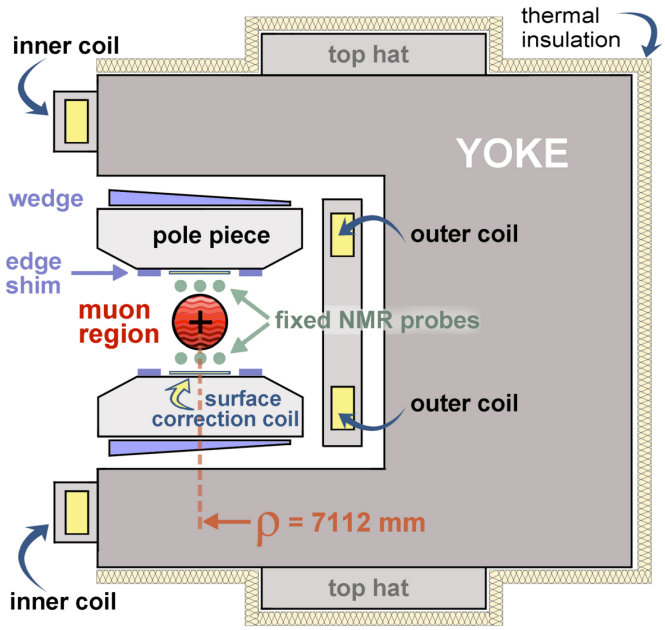

Figure 3.11: A schematic showing a cross-section of the $g-2$ storage ring, including all the field application and tuning subsystems as well as the shimming and monitoring systems. Reproduced from [60].

longitudinal field components in the storage region. As a purely vertical magnetic field is not physically possible within a finite space, the magnet is designed to apply a close approximation, with various components adjusted to minimise any higher order moments. These adjustments, known as shimming, can be done both actively, with systems monitoring and adjusting the field in real time, or passively, where a physical knob is turned to adjust the field and then left in that configuration for operation.

To decouple the magnetic field from any nonhomogenities in the steel of the ring magnet, separate pole pieces are inserted above and below the storage region. Above and below the pole pieces is a small air gap, which is what provides the decoupling and allows small pieces of aluminium called 'shims' to be inserted in this space to further angle the pole piece and fine tune the field. This allows control over the dipole and quadrupole moments of the field. Additional shimming is also provided by thin iron laminations on the pole pieces. Over 1000 of these pieces of aluminium were inserted during the shimming campaign performed in 2016-2017, which reduced the final field uniformity to better than 50 ppm, as shown in Figure 3.12.

After the passive shimming is complete, the field is controlled further by active methods using magnetic coils placed all the way around the ring, known as the

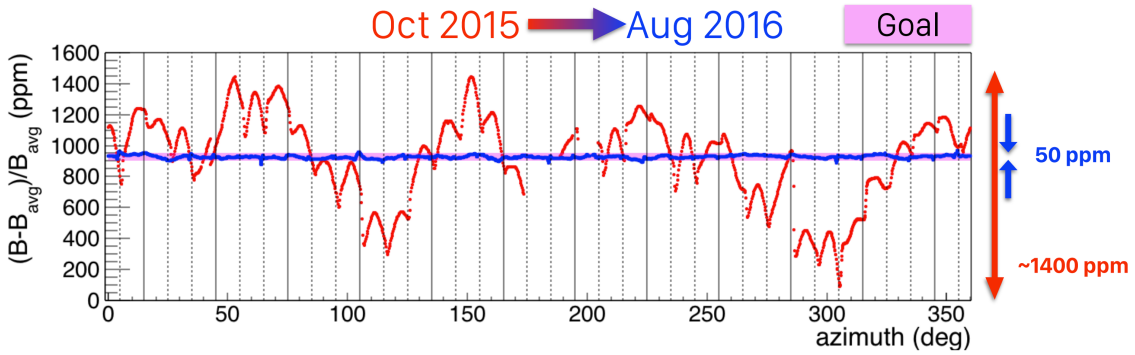

Figure 3.12: A comparison of the azimuthal field variations before and after the shimming campaign [61].

surface correction coils (SCC). The SCC are adjustable, allowing them to be used to correct local and global changes over time. The SCC can also be used to investigate the impact of non-zero radial and azimuthal field components, and are important for the method for measuring the radial field presented in Chapter 6.

3.6.1 Measurement of the magnetic field

A precise way to measure the magnetic field is required for the precession frequency analysis, with an overall target precision of 70 ppb. There are several systems designed to do this. The first is the trolley, a cart containing an array of NMR probes pulled around the inside of the muon storage region by non-magnetic fishing line. This measures the magnetic field in the region of the muon beam, generating magnetic field maps for each point as shown in Figure 3.13. Each NMR probe contains petroleum jelly, a proton-rich medium, which undergoes Larmor precession with the frequency proportional to the strength of the field. Since the 'free' proton frequency is required for the analysis, these probes are calibrated against a 'plunging probe' which contains ultra-high purity water. The plunging probe itself is also calibrated against several other water probes and a He-3 probe [61]. This allows the 'free' proton precession frequency to be extracted from the petroleum probes in the trolley to give the best determination of $\tilde{\omega_p}$.

Since the trolley moves in the storage region, it cannot be used to constantly monitor the field. Therefore, it is only used during dedicated 'trolley runs', which occur every few days during operation. In between these times, the field is monitored by a suite of fixed NMR probes located all around the ring, which are visible on the cross-section in Figure 3.11. These are also calibrated and allow interpolation between trolley runs to give maps of the field for all running times [61].

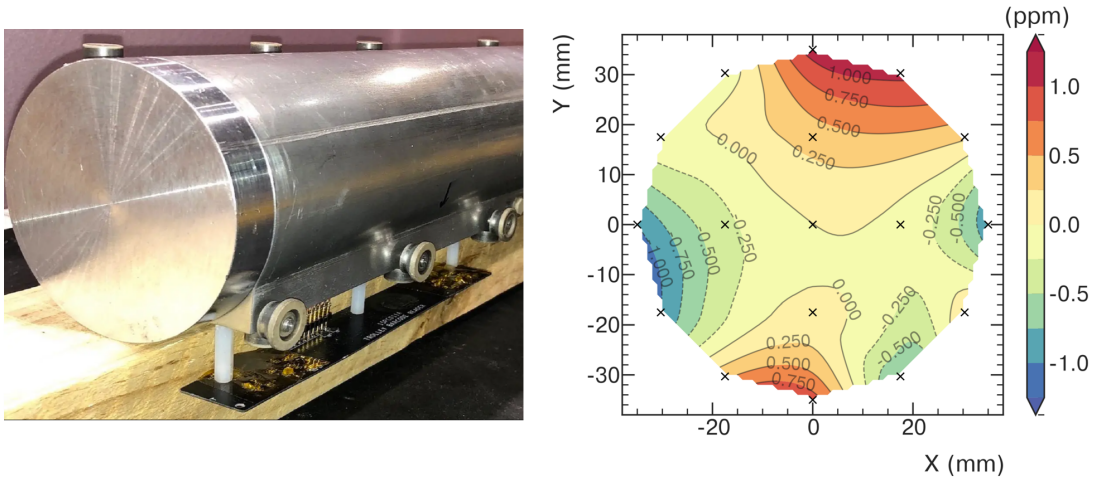

a) A photo of the front face of the trolley where the NMR probes are housed

b) An interpolated magnetic field map at a single azimuthal position around the ring

Figure 3.13: A photo of the trolley and a single field map generated by it. The 'x' markings on the right plot indicate the locations of the NMR probes. The field at other locations is interpolated from these probe measurements. Images reproduced from [60] and [61].

3.7 The calorimeters

The main detectors for the ω_a analysis are the calorimeters, of which there are 24 uniformly spaced around the inside of the ring. These are designed to register the hit time of positrons and measure their energy, counting the number of positrons seen above a defined energy threshold. This produces a time-varying signal, which gives the characteristic 'wiggle plot' from which the precession frequency can be extracted. The total number of hits seen in the calorimeters is used to quantify how much data the experiment is collecting, referred to as 'calorimeter tags' or CTAGs. Due to the segmented structure of the calorimeters, they can also be used to measure the hit position of the positrons. While this is not their primary purpose, their azimuthal coverage make them useful for studying any variations in the muon beam around the ring, for example the closed orbit distortion.

Each calorimeter is comprised of 54 PbF_2 crystals, arranged in a 9×6 configuration. These act as scintillators, releasing Čerenkov light when a positron passes through, with the intensity of the light generated proportional to the energy deposited [10]. This is detected by silicon photomultipliers (SiPM), chosen for readout due to their ability to be placed directly onto the crystals, their high photodetection efficiency, and their ability to function in a strong magnetic field. However, they are very sensitive to changes in bias voltage and temperature, so the crystal response to energy (also called the gain) must be carefully calibrated

[62]. This is achieved using a laser system, which sends pulses of light simultaneously to all 1296 crystals to track any gain changes. The expected energy deposit is compared to what the SiPMs measure, allowing any differences in the gain to be measured and logged. The laser system can also be used to measure the time resolution of the calorimeters by sending short pulses in quick succession [63]. The time resolution of the calorimeters was determined to be approximately 20 ps for 3 GeV positrons [64].

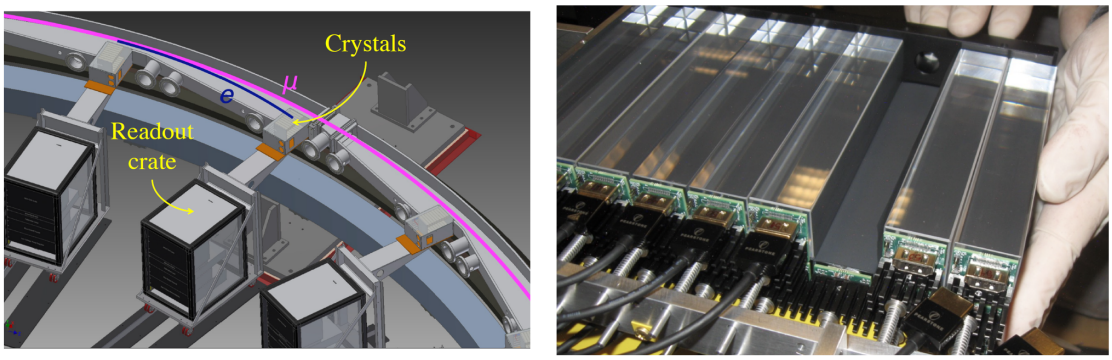

a) Schematic diagram showing the location of the calorimeter crystals relative to the ring

b) Individual calorimeter crystals, with attached SiPMs, photographed during construction.

Figure 3.14: The calorimeter crystals are placed directly in the path of the positrons, drawn here in blue, using specially designed extruded sections in the vacuum chambers, as shown in the diagram on the left. This avoids any scattering or losses due to travelling through material. On the right, a photo of how the PbF_2 crystals and their SiPMs are arranged. Images reproduced from [10] and [64].

PbF_2 was chosen as the crystal material due to its low magnetic susceptibility, short radiation length and low Moliere radius, which means that all scattering products are emitted within a small transverse direction. This ensures that the majority of the energy of the positrons will be deposited inside the detector and therefore recorded accurately. Each crystal is wrapped in a thin non-reflective and opaque foil to prevent photons travelling between the crystals. Although this reduces the overall intensity of light in the calorimeters, which slightly reduces the energy resolution, the ability to spatially separate clusters improves the time resolution substantially.

3.8 The straw trackers

Another key detector system for the experiment are the straw trackers, often just referred to as 'trackers' in this thesis and other works. Although the calorimeters

are ideal for counting positrons and measuring their energy, they give limited spatial resolution, and cannot measure the positron momentum directly. Therefore, it is beneficial to have another detector designed to reconstruct positron trajectories and extrapolate back to the point of decay.

By placing the trackers outside of the beam, the measurement is non-destructive, allowing for a time-varying measurement of the muon beam profile. This is very important for a full understanding of the beam dynamics and associated corrections and uncertainties. The beam profile measurements are also required to determine the magnetic field felt by the muon beam over time. Additionally, by reconstructing the trajectories of the positrons, the curvature can be assessed to measure the momentum of the positrons. This combined with the calorimeter data enables a 'matching' between the two, serving as an important cross-check and also allowing estimates of the efficiencies of the two detectors. This track-calorimeter matching gives better determination of detector effects like pileup in the calorimeters, where two positrons hit the same crystal at once, an understanding of which is very important for the ω_a analysis. Finally, precise position data at the vertex level is a fundamental part of the tracker EDM analysis, the main focus of this thesis. The trackers are designed to operate within the vacuum of the storage ring, since this greatly improves the acceptance of the system. As a result, they must have no magnetic footprint.

The trackers themselves are modular, with each tracker station being comprised of 8 modules with 4 layers of 32 straws. The layers in each module are tilted at a 7.5° angle, with the direction of tilt alternating between the first two layers (the 'U view') and the last two layers (the 'V view') to allow precise 3D hit position determination. The straws are held between two manifolds, as shown in Figure 3.15.

There are two tracker stations positioned around the ring, 'Station 12' at approximately 180 degrees and 'Station 18' at approximately 240 degrees from the injection point, as shown in Figure 3.2. The beam seen by these two stations is subtly different due to azimuthal effects like the closed orbit distortions, so it is often informative to compare data between the two. Each station is located within an extruded part of the vacuum chamber, upstream of a calorimeter, as shown in Figure 3.16.

Each straw is an individual drift time chamber made of a thin mylar outer shell with a 25 μm diameter wire held tensioned at the centre of the straw. The straws are filled with gas at a pressure of 1 atm [66]. The wire is held at a positive voltage of 1650 V, acting as a cathode while the neutral straw walls act as an anode. This

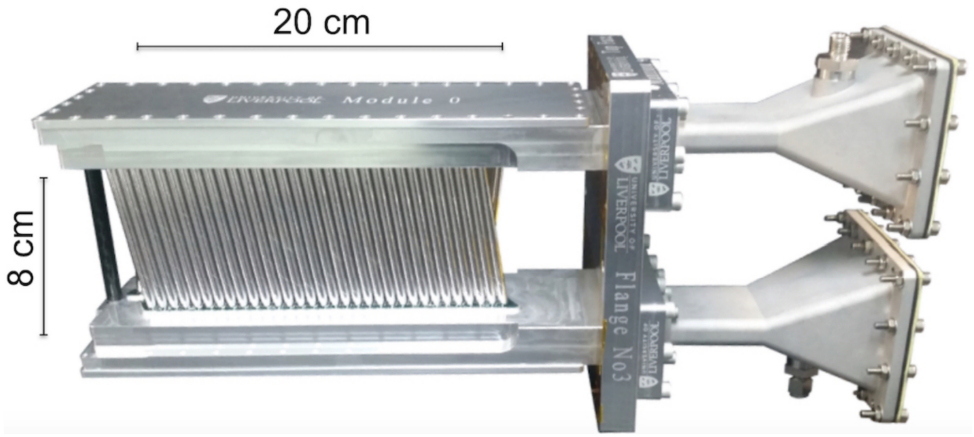

Figure 3.15: A photograph of a single tracker module, showing the tensioned structure and the crossing layers of straws. Reproduced from [66].

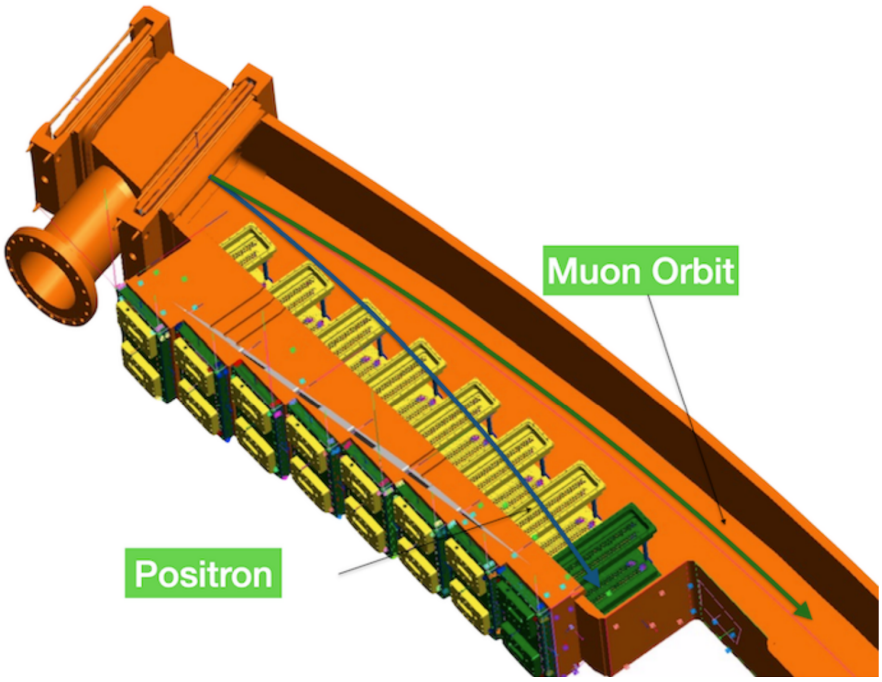

Figure 3.16: The location of tracker modules within a station. The blue line is an example positron path, travelling through several modules and into the calorimeter located behind the green module (in front of the green module in this view). Reproduced from [10].

fills the straws with a strong radial electric field. When a charged particle passes through the straw, it ionises the gas, leaving a trail of charged particles in its wake. The electrons from the ionisation drift towards the wire at the centre, causing an avalanche of secondary ionisations in the higher field region near the straw. These accumulate on the wire causing a current to flow [65].

The time taken for the ionisation products to reach the wire is roughly pro-

portional to the distance the particle passes away from the wire, also called the distance of closest approach (DCA). This defines a circular region where the original hit could have passed through the straw. Since the straws are operating in the ring's magnetic field, the trajectory of the positrons is not a straight line but rather a curved circular path, as shown in Figure 3.17.

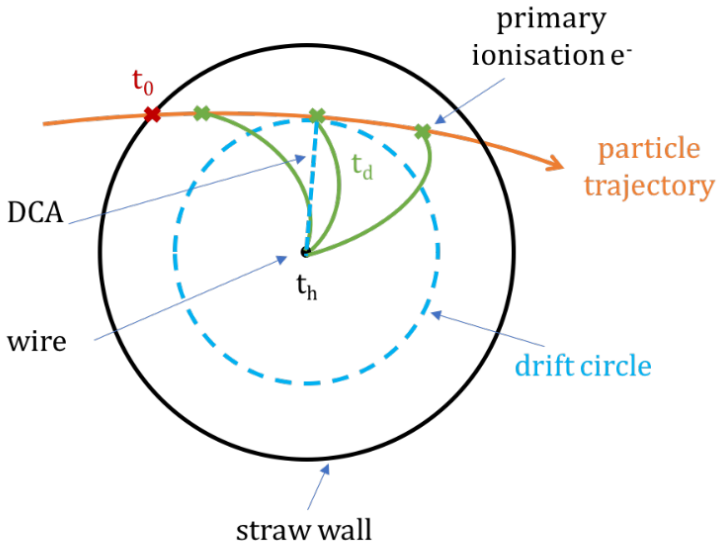

Figure 3.17: Diagram showing the response of a straw to a charged particle, with the ionisation products travelling in curved paths, and the definition of the distance of closest approach, or DCA. Image from [41].

The gas chosen for the tracker system is a 50-50 mix of argon and ethane. This mix includes a good ionisation medium (the argon) and a 'quencher' gas to absorb unwanted excess photons (the ethane). The photons are produced by primary ionisation similarly to the electrons, but are undesirable as they further ionise the gas as well as having a long path length. This can lead to a large shower of ionisation products which leaves the straw unable to register other signals for some time. This 'gas breakdown' point also limits the maximum operation voltage for the straws. This gas mix was tested against other commonly used mixes like argon and carbon dioxide, and was found to give the best performance, as well as minimising the leak rate of the gas through the straw walls into the vacuum [41].

Each straw hit is digitised by a series of electronics before being recorded. The first component is an ASDQ (Amplifier Shaper Discrimination with Charge) board, which triggers a 'hit' if the signal crosses a threshold. The ASDQ board also shapes the signal, combining the messy avalanche pulse into a smooth signal and recording the time above threshold [67]. These signals are passed to a TDC,

which manages 16 straws and time-stamps the hits, then to a logic board, which combines the signals from multiple TDCs with external clock timing information. These collections of hits are combined into 'events' which are then combined with data from other systems.

The 'digits' produced by this system are reconstructed into clusters of hits and fitted into tracks. First, the time stamped hits within an 80 ns window are grouped into a 'time island', since the maximum drift time for a straw is around 70 ns. This combines hits that are likely to have come from the same particle. They are then grouped spatially into clusters within the modules, and these are grouped into 'track candidates' across all modules containing the hits. The timing of these hits is used to estimate when the first hit in the tracker was detected, and therefore the associated drift times of the other hits in the track candidate. Using the relationship between drift time and radius, the DCA can be determined for each hit, defining a circle within each straw. These circles are finally fitted using a GEANE algorithm [68] that minimises the chi-squared of the fitted track [69].

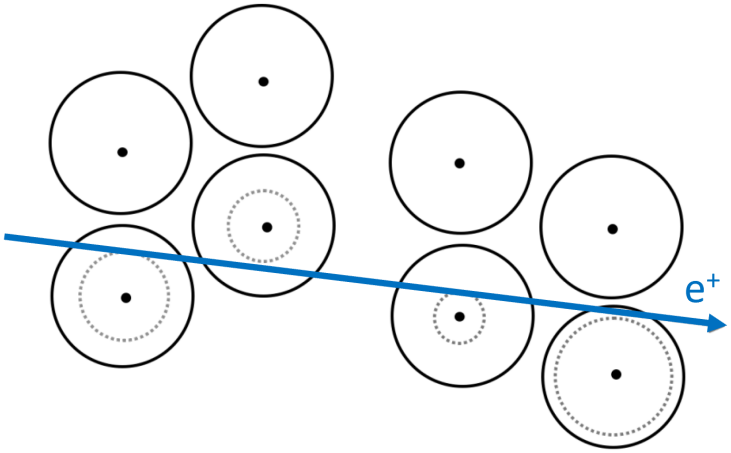

Figure 3.18: The circles defined by the DCA are used to fit the most likely positron path to reconstruct it. In 2D, the fit to these circles define vertical lines down the full length of the straw. The vertical hit position is defined by the overlap point between the two angled sets of straws. Adapted from [41].

Once the tracks are fitted, they are extrapolated both forwards towards the calorimeters, and backwards towards the decay vertex. This is achieved by using a Runge-Kutta algorithm [41], to account for the trajectory of the positrons in a magnetic field varying with position between 1 T and 1.45 T.

Before the tracks and vertices can be used for analysis, a series of quality cuts are applied to remove badly reconstructed tracks. A summary of the track quality cuts is given in Table 3.1. A track which passes all of these is termed a 'Quality

Track'. Similarly to the track-level cuts, an extra layer of cuts is applied to the vertices extrapolated from the tracks. These are shown in Table 3.2.

Track variable	Quality track requirements
Number of straws hit	> 12
Fit p-value	> 5 %
Drift time	$0\,\text{s} < t < 70$ s
Track x entrance point	$60\,\text{mm} < x < 150\,\text{mm}$
Track y entrance point	$-40\,\text{mm} < y < 40\,\text{mm}$
Track residuals	< 500 μm
Fraction of missed layers	< 30 %
|U-V| layers	≤ 4

Table 3.1: Table summarising the requirements for a track to pass the quality cuts. Cut definitions from [70].

Vertex variable	Quality vertex requirements
Volume hit	No volumes hit
Extrapolation distance	> 40 cm
Vertex uncertainty in R	$0.5\,\text{mm} < \sigma_Y < 5.0\,\text{mm}$
Vertex uncertainty in Y	$0.5\,\text{mm} < \sigma_R < 3.5\,\text{mm}$

Table 3.2: Table summarising the requirements for a tracker-obtained vertex to pass the quality cuts. A track is defined as having hit a volume if the extrapolation passes through material, for example the vacuum chamber wall, on its way to the trackers. The extrapolation distance is the distance between the vertex and the first hit in the tracker. Cut definitions from [70].

3.9 Auxiliary detectors

In addition to the calorimeters and trackers, there are three other detector systems used during operation of the experiment. The Inflector Beam Monitoring System (IBMS) is a subsystem consisting of three detectors placed near the injection point of the beam into the ring. Their purpose is to measure the incoming beam distribution before and after injection into the ring, allowing it to be monitored. This also allows tuning of the beamline to maximise the stored muons in the ring. These detectors are a grid of thin scintillating fibres read out by SiPMs, designed

to detect the beam without scattering too many muons [71]. This gives a 2D measurement of the beam profile in the x and y directions, with the y distribution shown in Figure 3.19.

It is also informative to understand the timing distribution of the muon beam as it enters the ring, achieved by the T0 detector, located right before the inflector. This consists of a scintillator connected to two photomultipliers. As the beam enters, the injection time is precisely measured by this detector, allowing for synchronisation of all other ring and detector systems. The output is a time distribution pulse, which measures the shape and width [72]. The integral of this pulse also gives the total injected muons into the ring, which can be cross-checked with CTAGs to ensure the injected beam is being stored efficiently.

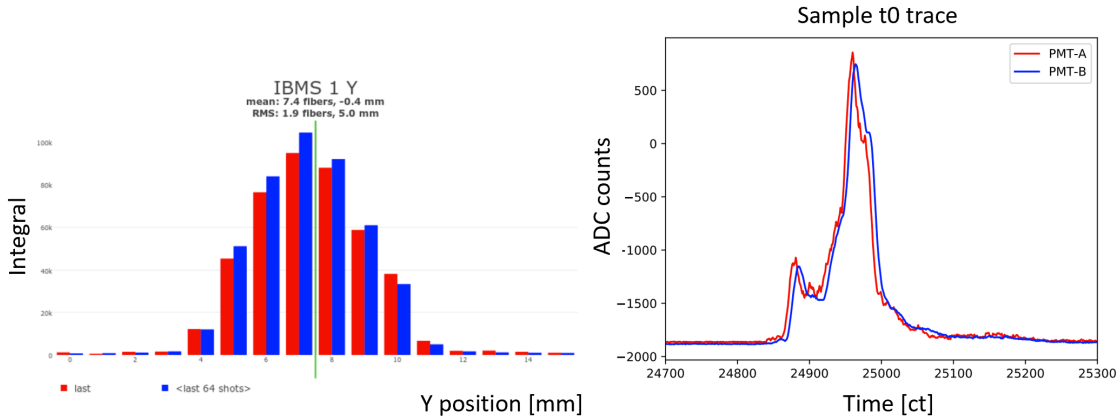

Figure 3.19: Two plots showing a beam profile measured by the IBMS upstream of the inflector (left) and an example pulse measured by the T0 detector (right) [72]. A and B are the two PMTs in the T0 detector.

Finally, a destructive measurement of the beam distribution is possible using the fiber harps. These serve a similar purpose to the tracker, allowing for a cross-check of the reconstructed beam as well as also extracting beam parameters such as the momentum and the CBO frequency. Each harp consists of two sets of scintillating fibres, with one oriented vertically and the other horizontally. These can be moved within the storage region to measure the beam at a range of radii [73]. Since this is a destructive measurement, the fiber harps are not used during normal operation but instead for dedicated systematic runs.

3.10 The clock and hardware blinding

All time measurements within the experiment are defined relative to a central clock. This utilises a rubidium source to give a short time period stability 'tick' which is linked to a GPS signal to provide longer term stability. This signal defines the time base for two synthesisers, one for the field systems at 60 MHz and another for the calorimeter and tracker timing at 40 MHz [74].

To avoid unconscious bias, the true clock tick timing of the 40 MHz clock is blinded, with the final tick frequency being set to a hidden value 40 - ϵ, with ϵ between 0 and 3 kHz. The clocks are located in a locked cabinet, with the blinding frequency shifts written down and put in two envelopes, kept by Fermilab employees who are not part of the collaboration. Once the ω_a analysis is complete, with all systematics calculated and cross-checks performed, the analysis is deemed 'frozen' and the decision can be made to unblind at this point. Each running period is given a separate blinding frequency to allow results to be sequentially unblinded as the analyses conclude.

The stability of the clock is monitored by using a third time synthesiser. This is used to generate another signal with the same time base as the original two at 30 MHz, blinded in the same way as the 40 MHz clock. This new signal is mixed with the original calorimeter signal in a way that hides the original value of the clock tick. This allows the stability of the clock to be monitored by the collaboration without revealing what the blinding frequency shift is.

3.11 Run periods

The Muon $g-2$ experiment at Fermilab has been actively collecting data since 2017, with a shutdown each summer for upgrades. This naturally splits the the experiment into several periods called 'Runs', with the 2017 run time being used for commissioning. At the time of writing, the experiment is nearing the end of its final run, Run 6. The integrated CTAGs for each run period are shown in Figure 3.20.

Due to maintenance and upgrades on the experiment itself between these run periods, there are often subtle differences in the running conditions. For example, the quadrupole voltage was different in Run 1 vs the other runs, and there was also an issue with some of the resistors in the quadrupoles, leading to a much slower charging time in Run 1 [21]. This introduced an early-to-late effect in many analyses, including the ω_a and EDM analyses. These resistors were fixed and the effect no longer exists from Run 2 onwards. As a result, all analysis methods, even

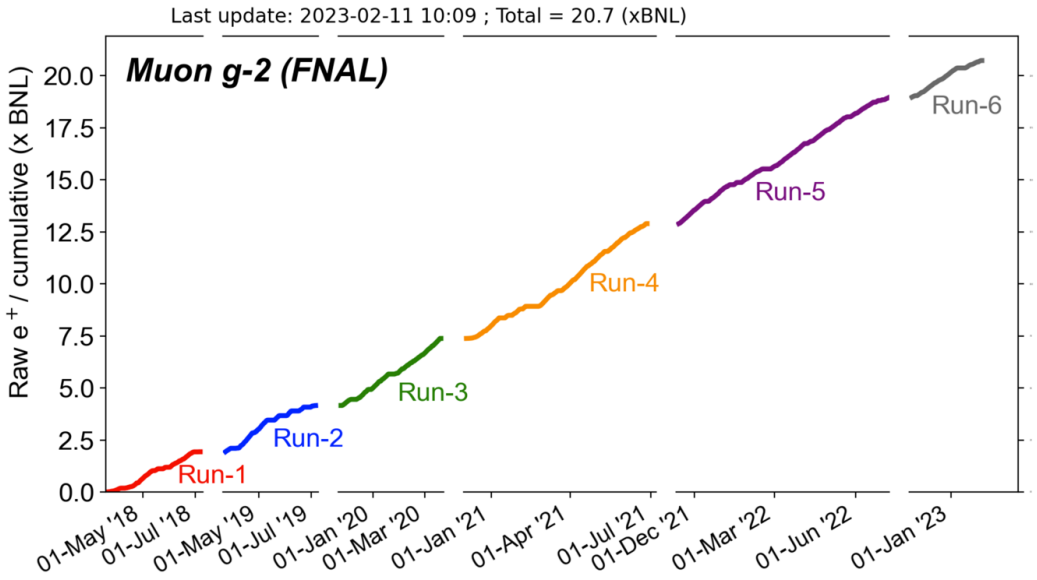

Figure 3.20: Total number of positrons collected in each of the 6 run periods, with Run 6 currently ongoing and therefore incomplete, in units of total number of datasets of the size of the Brookhaven experiment collected.

if they have worked previously, must be checked on new data to ensure no odd behaviour is seen.

The data used in this thesis predominantly comes from the Run 2/3 run periods, which are analysed together due to similar run conditions and being processed at similar times. This contains almost ten times as many tracks as Run 1. As a result, while Run 1 was purely statistically limited, systematic effects will start to have more impact in Run 2/3.

3.12 Simulation

Due to the limited coverage of detectors, not all effects needed for analyses can be fully extracted from data. Instead, a highly detailed simulation of the experiment is used. Several simulation tools exist, but the most relevant one for the analyses in this thesis is the GM2RINGSIM simulation, which is briefly described here.

GM2RINGSIM is a GEANT4-based [75] simulation which includes the full ring geometry, the positions of all ring systems and detectors as well as the injection beamline. The positions of detector systems are defined based on alignment surveys of the experiment. Electric and magnetic fields are included for all systems, with some being optional like radial magnetic field components or higher order multipoles in the ring field. Experimentally measured perturbations from components like the kicker are also included. The quadrupole fields are dynamic to

simulate the impact of scraping as well as allowing for studies in the 'damaged quadrupole resistor' mode for Run 1 [57].

The muon beam in the simulation can be run in two modes, the 'beam gun' and the 'gas gun'. The beam gun models the beam from the injection point and dynamically evolves it around the ring. This is useful for studying the impacts of injection as well as checking the expected levels of contamination from other particles. The gas gun skips this step and instead fills the whole ring with a uniform muon beam, which decays at the defined position. The distribution is defined to match the measured storage shapes as well as oscillations like the CBO. Filling the ring in this way improves the statistics of the simulation, speeding up the generation, and is useful for studying detector effects. This thesis makes use of the gas gun both to help tune an independent EDM simulation in Chapter 5, and to extract tracker acceptance corrections for the EDM analysis in Chapter 7.

While GM2RINGSIM is tuned to match data as closely as possible, there are some notable differences between data and MC that are occasionally relevant. For example, the simulation includes three trackers rather than two, and does not include all tracker effects such as the time width of hits. Part of Chapter 4 of this thesis focusses on an effort to include crosstalk, a known detector effect which was not present in this simulation, in order to study the impact it has on the tracking in a cleaner environment than in real data.

3.13 Data processing overview

The data processing at $g-2$ is split into two main sections. 'Online' data processing includes things that are performed in real time while the experiment is running, often with a focus on speed. 'Offline' data production is done at a later time to properly process the data and apply quality cuts to generate the final analysis datasets.

The main online systems are the components of the real-time data acquisition system (DAQ). This is a MIDAS [76] based system which has to handle a very high rate of data during each fill, around 18 GB/s, which comes from all the detector systems in the experiment, with the highest data flow coming from the calorimeters. The DAQ itself is made up of GPUs, which handle the high data rate and reduce it to 300 MB/s by removing calorimeter data without pulses, and frontends which handle the data flow to and from a particular system [77]. As well as the data collection frontends, there are data quality cut (DQC) frontends for both the calorimeters and the trackers, where a subset of the online data

is processed and reconstructed in real time. This allows for live data quality monitoring via real-time plots. The field systems have their own DAQ to collect data from the fixed probes and trolley when they are operational. A subset of this online data is run through the full data processing, generating a small set of ROOT files within a few hours of data collection. This is useful for monitoring systems and performing small systematic analyses to either better understand the state of the experiment or fine tune operations.

The online data is temporarily stored locally on an array of disks before being copied to tape storage using Fermilab's dCache system [78]. This provides permanent storage at the cost of having to 'stage' any data before it can be accessed for analysis. The offline production splits this data up into datasets, each bookended by two specific trolley runs, removing any data which is not production quality. The cleaned up data is run through a preproduction stage, involving the pre-staging of the relevant dataset, which is then run through the full calorimeter, tracker and field reconstruction. The output of this process is a series of ROOT ntuples. As a final step, a series of DQC cuts are applied, for example removing data where the field was not stable or removing any data where the muon beam storage was behaving abnormally. The data passing these cuts is collated into the official production datasets ready for analysis.

3.14 Measuring ω_a and $\tilde{\omega_p'}$

Although the main focus of this thesis is an EDM analysis, the measurement of the precession frequency includes a series of techniques which are useful for the EDM analysis. Additionally, some of the work in Chapter 4 contributes to the evaluation of systematics for this, so a very brief summary is presented here for the sake of completeness. A full discussion can be found in the Run 1 results paper from the experiment [52].

As described in Equation 2.20, in order to extract a_μ, the experiment measures the ratio $\omega_a/\tilde{\omega_p'}$. Accounting for the analysis corrections and combinations needed, this can be expressed in a 'master formula' as:

$$\frac{\omega_a}{\tilde{\omega_p'}} = \frac{f_{clock}\,\omega_a^{meas}(1 + C_e + C_p + C_{ml} + C_{pa})}{f_{calib}\langle\omega_p' \times M(x, y, \phi)\rangle(1 + B_k + B_q)}. \tag{3.9}$$

Starting with the numerator, f_{clock} is the hardware blinding frequency introduced by the shifting of the clock tick, described in Section 3.10. During the analysis itself, a software blinding is also applied, different for each analysis group

to allow an intermediate consistency check before the final unblinding.

ω_a^{meas} is the precession frequency measured by the calorimeters. This is extracted by plotting the number of positrons above the energy threshold, and fitting the decaying oscillation with a 22- or 24-parameter fit as shown in Figure 3.21, where the parameters correspond to various beam motions and time-varying effects. There are six different analysis groups, with slightly different reconstruction and analysis algorithms. Although the theoretical shape is simple, only requiring the 5-parameter fit from Equation 2.9, extra parameters are needed to account for things like the CBO and lost muons, which change the shape over time. The Fourier transform in the background of Figure 3.21 illustrates the difference between the two fits, with clear unfitted frequencies due to beam dynamics effects.

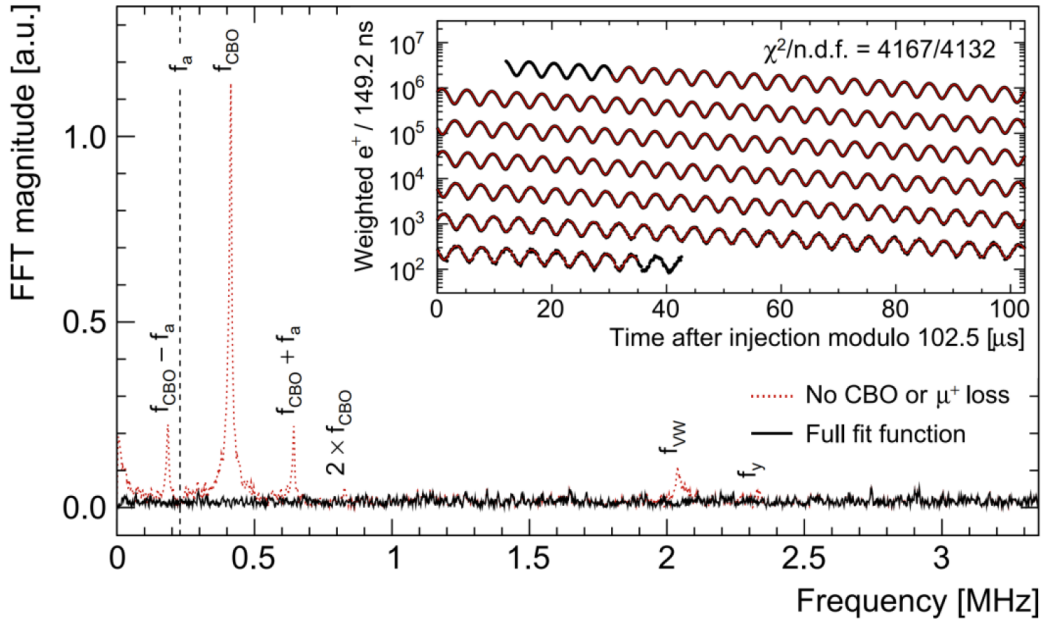

Figure 3.21: A nested plot showing the Run 1 wiggle plot 22-parameter fit, and the FFT of the fit with and without two of the most important 'extra' beam dynamics terms. A large difference is seen in the FFT with many peaks being removed due to more accurate fitting. Reproduced from [52].

The final terms in the numerator are the beam dynamics corrections, some of which have already been briefly touched on. C_e is the E-field correction, which arises due to the momentum spread of the beam. In Equation 2.18, the momentum of the muons has been chosen at the magic momentum to simplify the expression. However, not all muons are at the magic momentum, even after scraping, so a correction is required. Similarly, some vertical motion still remains, leading to the pitch correction, C_p .

The final two terms are both corrections applied to the phase. Due to this and

ω_a being very highly correlated, any bias to the phase of the positrons detected will directly impact ω_a. C_{ml} is the lost muons correction term, arising from muons which are lost during the fill, either due to scattering or having a non-ideal momentum. The loss of these can lead to a shift in the overall phase which must be accounted for to prevent a bias in ω_a.

Finally, C_{pa}, the phase acceptance, arises due to different parts of the beam cross section having slightly different phases, with the phase in the ω_a analysis being the average of these. As the beam is not stationary, this leads to different sub-regions of the beam being included in the analysis, which can introduce a time-varying effect to the phase and therefore ω_a.

All of these beam dynamics corrections are calculated and controlled using a combination of tracker data, MC simulations and dedicated systematic runs. As the trackers only measure the beam at two locations around the ring, the data must be evolved to cover the full ring.

To extract ω_p, the magnetic field as measured by the calibrated trolley and fixed probes must be combined with tracker-obtained beam distribution $M(x, y, \phi)$ to extract the magnetic field that the muons feel as they travel around the ring. In addition to this, any field perturbations must be corrected for, as these will change the precession frequency by a small amount. The two main corrections are transient fields from the kickers and the electrostatic quadrupoles. The kicker transient B_k arises due to eddy currents in the aluminium components nearby. These take slightly longer to die off than the kicker pulse itself, leading to a correction being needed. Similarly, the quadrupole magnets take some time to charge and discharge, which introduces a mechanical vibration into the system, accounted for in the quadrupole transient B_q. Finally, the ω_p must be converted into the form for a free proton, which is done via a collection of calibration factors f_{calib} as described in Section 3.6.1.

3.15 Measuring a muon EDM

As discussed in Section 2.2.1, looking for an increase in the muon precession frequency does not have sufficient sensitivity to search for the existence of a small EDM. Therefore, a different method must be used which searches for evidence of a vertical tilt in the precession plane directly. There are three main methods which were used to do this at BNL, two using calorimeter data, and one using tracker data. Common to all EDM methods, if the measured tilt is zero or consistent with zero, a limit can be set.

The first calorimeter method looks for a time varying oscillation by counting the number of positrons detected in the upper and lower halves as a function of time. The resulting data is plotted modulo the $g-2$ period to average out any slow beam dynamics effects, and fitted with a sinusoid to extract the amplitude. This method was very effective at BNL, but was purely systematically limited, with the dominant uncertainty being alignment of the calorimeters and the vertical reconstruction resolution [79]. This is improved in the Fermilab experiment by the introduction of segmented calorimeters, which allow for better vertical resolution of hits.

As the precession plane is tilted, the distance a positron travels before being detected changes, as shown in Figure 3.22. When they are detected, the difference in distance leads to a slightly different phase. Therefore, a second way to look for a precession plane tilt is look for a time-varying phase shift between the two sides. This method is also systematically limited, being affected by any tilt of the calorimeters as well as any misalignment between them [80].

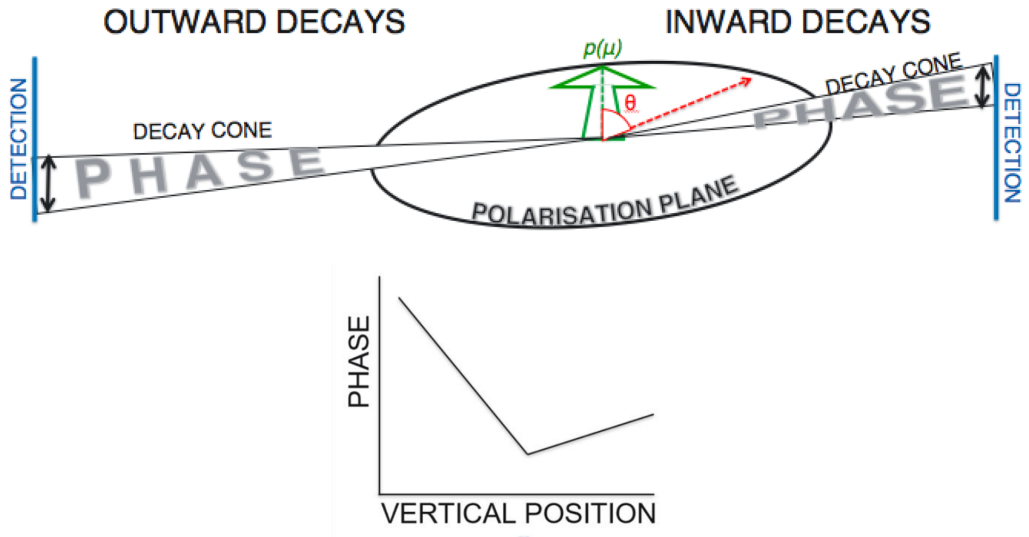

Figure 3.22: An illustration of the differing path lengths in the calorimeter phase method, and the resultant vertical position-phase plot that shows the tilt of the plane. Reproduced from [81].

The trackers can also be used to look for a vertical tilt 'directly' by measuring the trajectory of the positrons and extrapolating back to the decay vertex. This allows for a direct measurement of the vertical angle and therefore the plane tilt at the time of decay. This is plotted modulo the $g-2$ period similarly to the calorimeter analysis and fitted with a sinusoid to extract the amplitude, as shown in Figure 3.23.

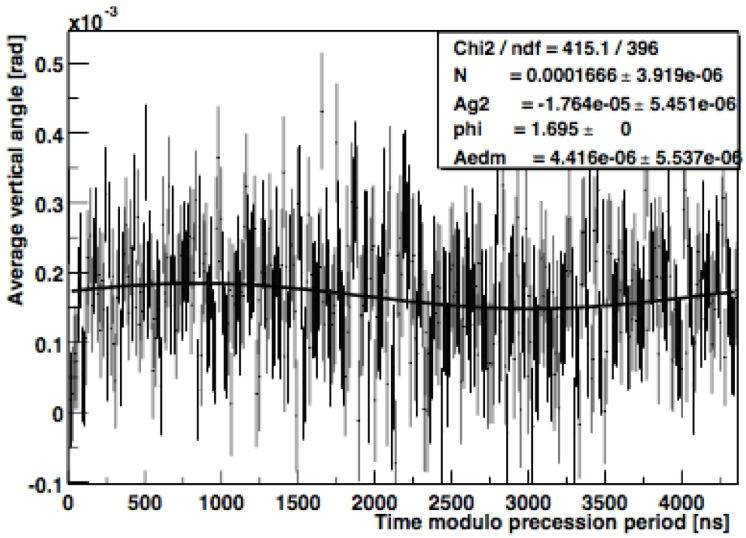

Figure 3.23: One of the BNL tracker EDM plot results, illustrating the a fit to a vertical angle oscillation. Sourced from [1].

At BNL, this analysis was statistically limited, so with the large increase in statistics collected at FNAL, this will be the analysis where the largest improvements can be made. This is therefore the main focus of this thesis. Achieving the full possible improvement requires an in-depth understanding of the impacts and relative sizes of the systematic uncertainties, something which was not required at BNL due to the statistical limitations [82]. Therefore, a new MC study is needed to understand the potential sensitivity and what the dominant uncertainties are likely to be. This is presented in Chapter 5.

The analysis itself is presented in full detail in Chapter 7, which follows on from S.Grant's analysis of Run 1 [12] to analyse Run 2 and Run 3, including several improvements. Finally, since this analysis uses tracker data, it is beneficial to fully understand the impact of detector effects on the reconstructed vertex data, presented in the next chapter.

Chapter 4

Characterisation and optimisation of the straw trackers

The straw trackers are a key component of the EDM analysis, providing positron tracks which are extrapolated back to the decay vertex. Therefore, optimising and fully understanding the behaviour of the trackers directly contributes to the EDM analysis, with any improvements in precision or efficiency directly improving the final sensitivity.

This chapter focusses on studies of the tracker performed by the author, including an in-depth study of crosstalk and the impacts of this on the quality of reconstruction, and a study varying the voltage across the straws to determine if they are still at the optimal setting and to look for signs of aging. It also covers the discovery of a new time-dependent effect in the tracker which must be accounted for in the EDM analysis.

The tracker data is also an important part of the precession frequency analysis, particularly for studying the beam dynamics. Therefore, the knowledge gained from the studies in this chapter is also used to calculate some of the systematic error contributions to the Run 2/3 ω_a analysis.

4.1 Tracker HV scans

Straw trackers perform best across a specific range of wire voltages, where the voltage is high enough to provide a sufficiently strong electric field to form good avalanches, but not so high that the ionisation products produced by the first hit continue to ionise the gas, known as refiring. This balance results in a 'plateau' region where the operation of the tracker is optimal [65]. An example of this behaviour is shown in Figure 4.1, using data collected with two different gases

during the design phase of the tracker. One of the reasons Argon-Ethane was chosen as the best gas for the tracker is due to its wide plateau.

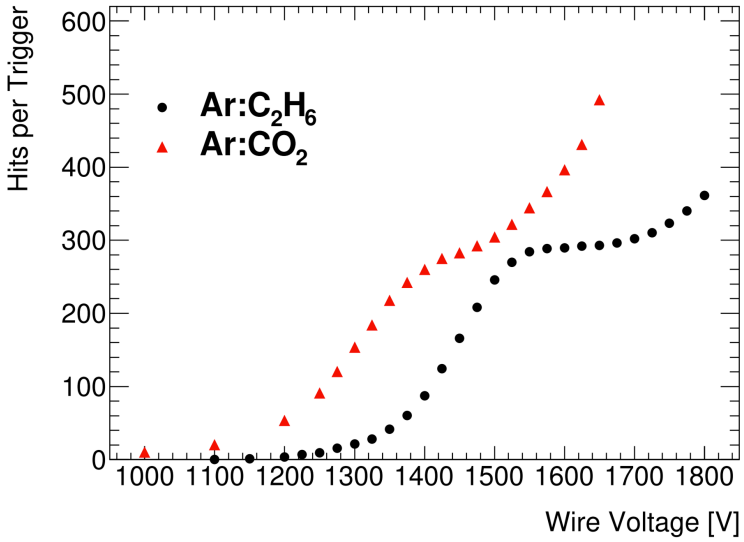

Figure 4.1: The hit rate of straws with two gases, Argon-CO_2 and Argon-Ethane. Argon-Ethane was chosen for its larger plateau region. Adapted from [66].

The existence of this plateau does, however, make the exact optimal running voltage for the tracker ambiguous. Other track properties, such as the efficiency of the tracker at detecting hits, the resolution of those hits, and the crosstalk must also be considered. These could change over time due to ageing of the straws, so monitoring is required.

The voltage of one module in the tracker was varied during a period of standard data collection to study these properties. This was done to check the tracker was still operating in its optimal regime, whether the HV could potentially be lowered, and to check for any signs of ageing. One tracker module is changed with the others left at 1650 V to allow any differences in the base running conditions to be monitored. The full plateau region is scanned across between 1650 V down to 1500 V, first scanning down in steps of 50 V, then the intermediate points back up to 1625 V, to allow an extra check of whether conditions changed during the scan.

The track detection efficiency can be determined by considering how many hits are missing in the path of fitted tracks. This is calculated as a percentage and plotted as a function of the DCA for each voltage point in Figure 4.2. Similar behaviour is seen for all voltages, with a large flat region and drops at low and high DCA. The high DCA drop is caused by a short path length through the straw, leading to insufficient initial ionisation to cause an avalanche, whereas the smaller

drop at low DCA is from multiple small ionisations which are spaced too far apart in time to combine to pass the detection threshold.

As the voltage is lowered, these drops become more and more pronounced, as lower voltages lead to smaller avalanches which are unable to compensate for the low ionisation effects. The peak efficiency at 1500 V is only around 90 % compared to 99.3 % at 1650 V. Therefore, the optimal running voltage based on efficiency alone remains 1650 V. Since this peak efficiency is identical to that seen during commissioning [66], this indicates that the trackers do not show any signs of ageing. A higher voltage is expected to age the trackers faster, but as there is no evidence of this seen, a reduction is not needed.

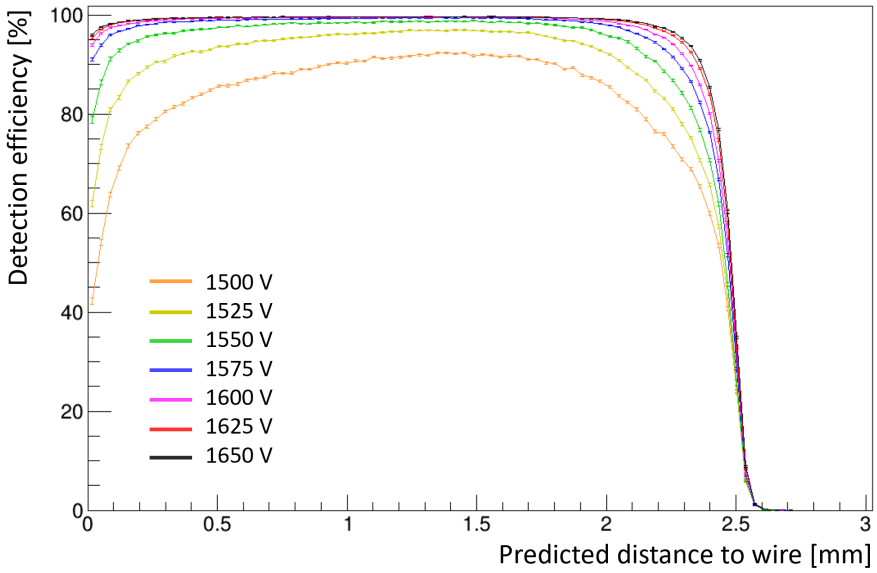

Figure 4.2: The tracking efficiency, plotted as a function of the distance of closest approach (DCA) for each HV setpoint. The efficiency increases as the voltage is increased.

The track resolution is calculated by removing a hit in the tracker and comparing the predicted position from the fitted track to the measured straw hit position. A histogram of this difference is fitted with a Gaussian to extract the width, which defines the resolution. This is plotted against the HV setting in Figure 4.3.

This shows an almost linear decrease in width as the voltage increases, corresponding to an improvement in tracker resolution. This is due to a higher voltage increasing the acceleration of electrons, meaning ionisation products neutralise faster. This allows the straw to register another hit faster, giving better time resolution and therefore position resolution. The xx25 V points also lie cleanly between the xx50 V points, indicating that there were no large changes in conditions during the scan, as otherwise there would be a relative shift between the two sets

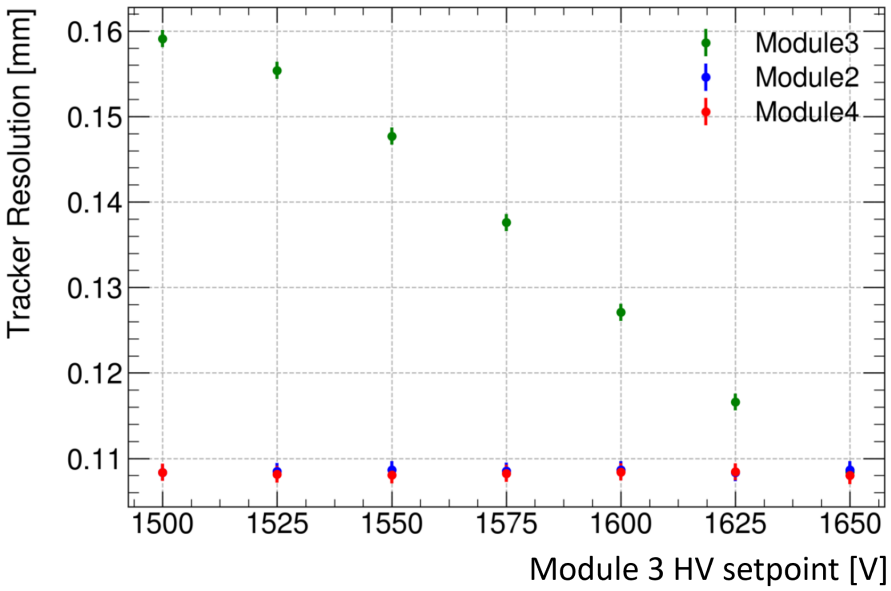

Figure 4.3: The hit resolution for different operating voltages. Module 3 had the voltage changed while modules 2 and 4 were left at 1650 V for the sake of comparison. A clear trend is seen of the resolution improving as the voltage is increased.

of points. The 1650 V point was also collected twice, at both ends of the scan, and the two resolutions extracted are consistent. As a bonus check, since the voltage of only one module was changed, the resolution can be plotted for other modules in the tracker, which would be expected to be constant, as observed in modules 2 and 4. Since dropping to 1625 V would make the resolution significantly worse, this further supports that it is optimal to continue operating at 1650 V.

4.2 Crosstalk

Another tracker property known to change with voltage is the level of crosstalk. Unlike the resolution and efficiency, the identification of crosstalk hits requires a new study, which is detailed in this section.

Crosstalk is an electrical coupling between components of a system, where E and B fields from normal operation induce fake signals. It is therefore important to quantify the rate of these signals and understand the impact they have on the tracking.

In the trackers, crosstalk manifests as fake straw hits in neighbouring straws. These extra signals introduce inaccuracies in track reconstruction, as there is an ambiguity between the true (primary) hit and the crosstalk hit. As a result, the tracking aims to remove as much crosstalk as possible prior to reconstruction,

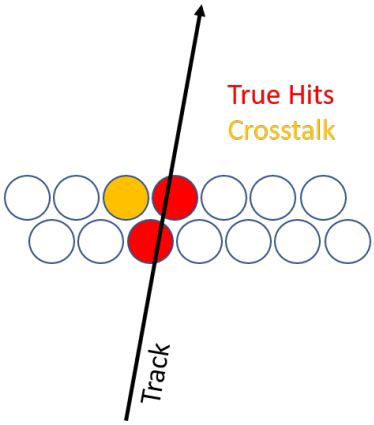

Figure 4.4: A diagram showing a crosstalk hit in relation to two real hits.

which requires an understanding of its characteristics.

One way to identify crosstalk is to consider the width of the straw response above the detection threshold. True hits occur when the gas in the straws becomes ionised, leading to an avalanche in the straw. This will spend a length of time above the detection threshold, called the hit width. A crosstalk hit will in general have a much smaller width than a real hit. This is because each real hit will cause multiple ionisations, whereas crosstalk will have at most one. The low width hits form their own separate peak in the width distribution, as shown in Figure 4.5.

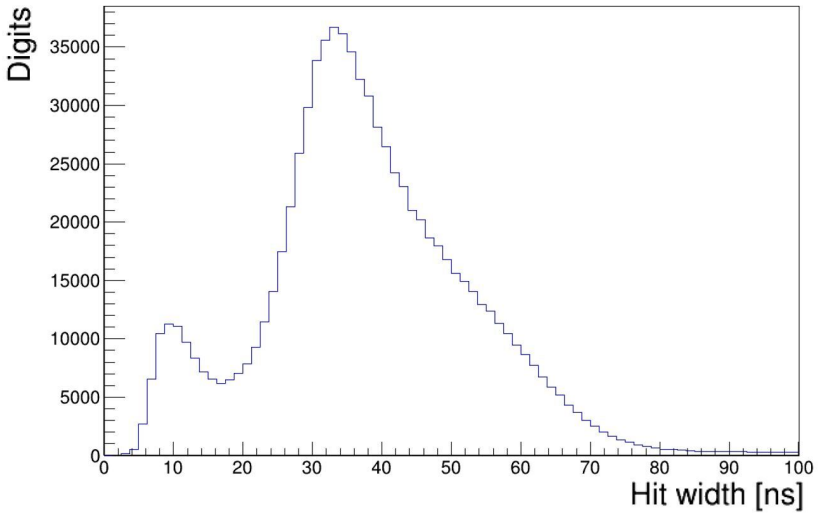

Figure 4.5: The full hit width distribution. The smaller peak at low widths is from crosstalk hits.

This allows the removal of most but not all of the crosstalk by applying a cut on the width at 14 ns. At this point, hits have been grouped into tracks based on their time differences, so each layer of the candidate can be checked for duplicate

hits. If any duplicates are found, for the sake of having clean data, all hits on this layer are removed from the data. This combined with the initial cut is very effective at removing crosstalk from data, although it comes at the cost of losing a fraction of quality tracks due to the cut on the number of hits. Removing a layer also increases the fit uncertainties, which leads to further tracks being removed and larger reconstruction errors. The worse vertex resolution from extrapolating these directly impacts physics analyses like the EDM.

In order to fully understand and minimise the effects of crosstalk, a series of dedicated studies was performed to devise an algorithm that would allow identification and removal of crosstalk hits. The impact of these extra hits is also considered and investigated.

4.2.1 Identifying crosstalk hits

To study the properties of crosstalk, the hits must first be distinguished from a number of known background processes that also cause multiple hits in a layer. These include multiple overlapping tracks within the same time island, secondary ionisations from real particles coming off the main track, and stray hits with no correlation to the track itself, for example from cosmic rays. The crosstalk hits occur in neighbouring straws with two slightly different signatures and properties. They can either occur in the same tracker layer or in an adjacent layer to the primary hit, referred to as 'same-layer crosstalk' and 'cross-layer crosstalk' henceforth. Once a crosstalk instance has been identified, the real hit in the layer, called the primary hit, must be identified and removed to leave only crosstalk hits in the final study sample.

Since a crosstalk hit is caused by another nearby hit, one potential identification method would be to compare the time both hits occurred, with crosstalk expected to be the later of the two hits. This would work well for same-layer crosstalk, however the existence of cross-layer crosstalk complicates this method. A crosstalk hit on an adjacent layer could register before the true hit in an adjacent layer if it has a shorter drift time than the real hit. This means that cross-layer crosstalk cannot be identified by time difference, and same layer crosstalk must be carefully selected to avoid cross-layer hits.

To remove crosstalk-like hits from overlapping tracks, time islands with only one track candidate are considered when collecting crosstalk. This removes most cases, but only works if all hits are successfully formed into track candidates. Therefore, this sample is is further cleaned by removing clusters of hits with no corresponding track candidate hit in that layer.

Since the mechanism behind crosstalk is a field, it is significantly more likely to occur in adjacent straws rather than over any larger distance. Therefore, the only hit pattern considered as potential crosstalk is two or three hits on the same layer. Additionally, only hits in the front layer of a view are considered, with either one hit or no hits in the layer behind it. This clean signature allows both same-layer and cross-layer crosstalk to be investigated.

For same-layer crosstalk, the hit in the back layer, if it exists, is required to arrive later than both hits in the front layer to ensure it is not cross-layer crosstalk. Therefore, in this case or the case where there is no hit in the back layer at all, the hit time can be used to identify and remove the primary hit. For cross-layer crosstalk, the requirement is flipped, with the hit in the back layer arriving before both hits in the front layer. Selection of the primary hit cannot be done using the hit time, as it is ambiguous whether the primary hit came first or second, so the hit with the largest width is removed from the sample.

These selection methods are used to isolate the crosstalk hits, allowing them to be counted. Although only the front layer in the pair of crossed layers is considered, crosstalk should be equally likely across layers, so the rate seen in the front layer is representative of both layers. The crosstalk percentage is defined as the ratio:

$$\text{Crosstalk percentage} = \frac{\text{Number of crosstalk hits}}{\text{Number of total hits}} \times 100 \quad \%. \tag{4.1}$$

The effectiveness of the cuts can be evaluated by plotting the width of all hits included in the crosstalk sample. Any residual primary hits will be visible as much higher-width hits compared to the crosstalk peak. A comparison of the initial width distribution and the final width distribution after crosstalk hits are selected, is shown in Figure 4.6.

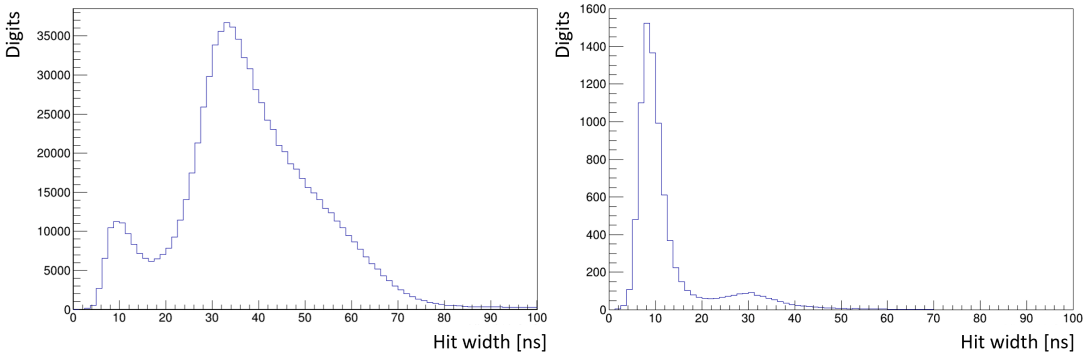

Figure 4.6: A comparison of the hit width distribution before (left) and after (right) crosstalk hits were selected for. Most high-width hits have been removed, leaving only low-width crosstalk hits.

This shows that the majority of high-width hits have been removed successfully, at a cost to the overall number of hits. The fractional loss will not bias the crosstalk percentage as the broad cuts remove both crosstalk and true hits equally.

Since crosstalk is an intrinsic property of the straws and electronics, it is expected to have an equal probability of occurring at any time, and also in any layer of the tracker. This provides another way to check hits selected by the cuts, since other crosstalk-like noise has time and/or layer variations. Tracks failing the candidate cuts due to not having enough hits are more likely to be found near the ends of the tracker. Similarly, overlapping tracks are more frequent at early times in the fill due to pileup. The layer dependence of the crosstalk percentage is shown in Figure 4.7, and is found to be stable across all layers within uncertainties. This shows the removal of track-like fake signals has been performed successfully.

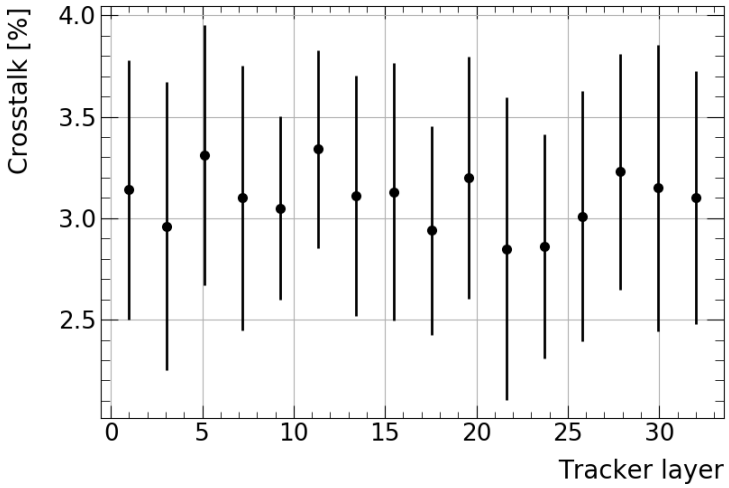

Figure 4.7: Variation of same-layer crosstalk across the layers of the tracker, sampled every two layers. This is flat, as would be expected.

In comparison, Figure 4.8 shows the variation of the total crosstalk percentage with the time in fill. Rather than being flat, this shows an exponential-type rise, which is fitted in order to characterise it and determine its origin. Further investigation reveals that the width distribution and tracker resolution also change with time in fill, with very similar lifetimes of about 50 μs, suggesting they originate from the same effect. This time dependence is discussed in more detail in Section 4.3.

Since this rise is caused by a non-crosstalk effect, evidenced by the lack of contamination seen across tracker layers, a time cut is applied at 100 μs to remove the rise. This cut is used for the calculation of the final crosstalk percentage, as this is the best representation of the true value in the trackers. The total crosstalk

percentage is calculated and plotted for same-layer and cross-layer crosstalk separately, and can be seen over the early-to-middle part of a fill in Figure 4.9.

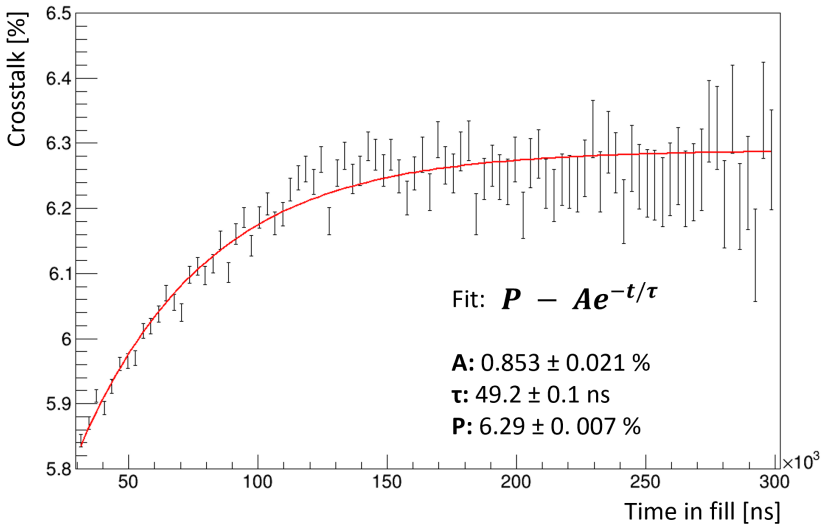

Figure 4.8: Variation of total crosstalk with time in fill. This is expected to be flat, but is not, showing a rise at early times, hinting at an extra time-dependent tracker effect.

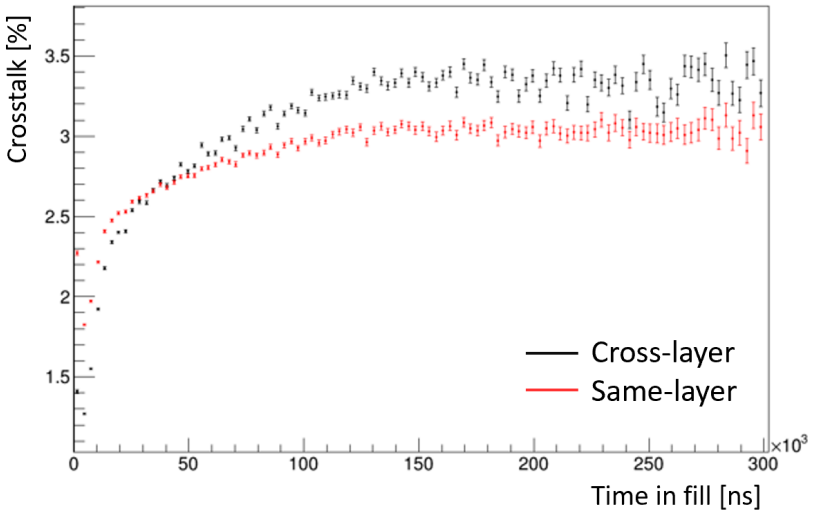

Figure 4.9: Final curves for same-layer and cross-layer crosstalk, using a full subrun of data. This shows that while they have slightly different behaviour as a function of time, in the true percentage there is slightly more cross-layer crosstalk than same-layer crosstalk.

The overall percentage of crosstalk is found to be 6.29 ± 0.03 %, with 3.00 ± 0.02 % same-layer and 3.29 ± 0.02 % cross-layer contributions. All uncertainties are purely statistical and could hypothetically be reduced by running the crosstalk

selection over additional tracks. However, for practical purposes, these percentages are only needed for comparison to each other, and are later used to simulate the impact of crosstalk. Since the uncertainties are small enough to show that the same-layer and cross-layer crosstalk are not equal, the precision achieved is sufficient.

The split between same- and cross-layer crosstalk hints at what is causing it. Crosstalk can occur either between the straws themselves, either in the wire or in the gas, or between components on the ASDQ boards [67]. Straw-to-straw crosstalk is equally likely to cause same- and cross-layer crosstalk since the straws are spaced equidistantly. However, crosstalk in the electronics of the tracker is more likely to be cross-layer crosstalk due to the design of the ASDQ boards, where the capacitors for cross-layer straws are closer together than for same-layer straws. This is explored further using dead straws in Section 4.2.3.

The crosstalk percentage is cross-checked with different selection methods. The simplest of these methods is a width-based method, where all hits below a certain threshold are counted as crosstalk. Setting the threshold at 20 ns, the percentage found is 6.43 $\pm$ 0.04 %. This value is slightly higher than the previous value, but also includes the tail of the real hit distribution, so is an overestimate. The ratio of the number of hits removed by the width cut to those removed by the more complex selection method can also be compared, with the results of this 'relative efficiency' shown in Figure 4.10. This shows the efficiency does start low at early times, but past 100 μs flattens out in the $\geq$ 90 % range. This shows that no large bias has been introduced as the final percentages are consistent.

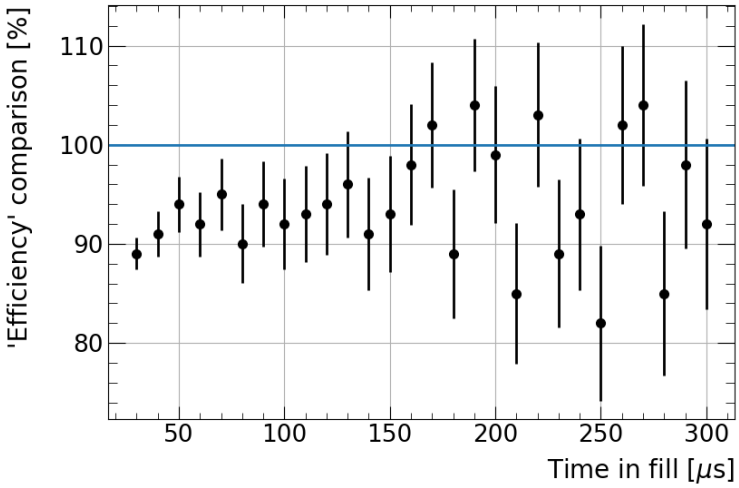

Figure 4.10: 'Efficiency' comparison between the simple width-based selection method and the more complex final selection method.

A template-based method is also used to cross-check the percentage. The hits in already identified good tracks are taken to generate an expected distribution for real hits. Subtracting this from the full width distribution therefore leaves only crosstalk or other false hits behind. Finally, most non-crosstalk noise is removed by requiring the hit to be adjacent to the track candidate hit. The crosstalk estimate given by this method is 6.18 ± 0.06 %, which confirms that the crosstalk is in the low 6-7 % range. This is also checked using MC generated data, which is discussed further in Section 4.2.6.

All these percentages are for data without any crosstalk removal. However, the crosstalk selection method can be applied to data that has been cleaned, to see what percentage of crosstalk remains. The results are shown in Figure 4.11, which indicates that the actual running level of crosstalk is around 0.51 ± 0.04%. This is likely an overestimate of the true value, as a significant number of hits are high-width contamination indistinguishable from crosstalk by their hit pattern. As a result, it is sensible to take this value as an upper bound rather than the true value.

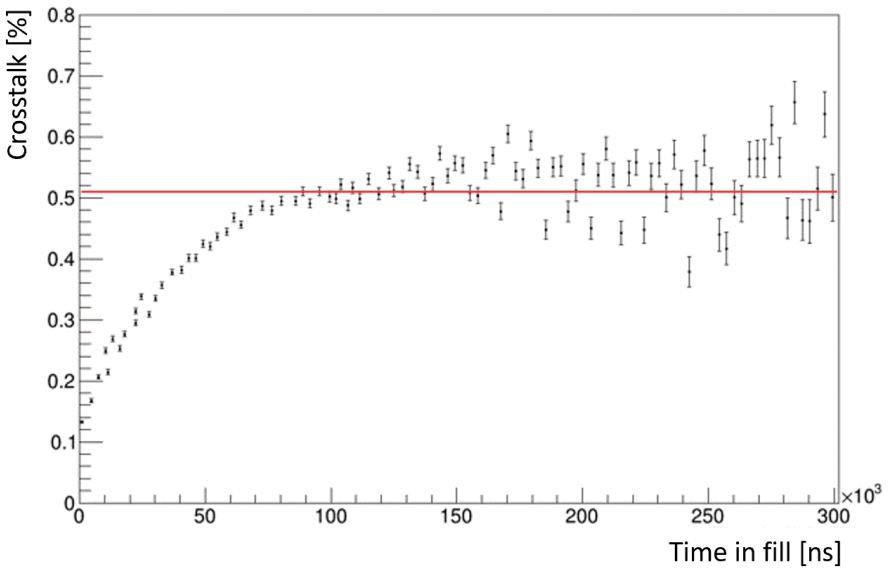

Figure 4.11: The total crosstalk percentage as a function of time in-fill, with the width cut applied. This is a better indication of the real level of crosstalk in final tracker data. The line is fitted from 100 ns onwards, but has been extended to the axis for readability.

4.2.2 Crosstalk HV scan results

The probability of crosstalk is expected to increase with tracker voltage, as a higher voltage leads to a larger avalanche in the gas and therefore a larger signal in the wire. However, the tracker efficiency and resolution are negatively impacted by a lower HV setting, so a trade-off point must be chosen. The data from the HV scan as detailed in Section 4.1 is used to find how large an impact the voltage has on the crosstalk percentage. The results can be seen in Figure 4.12.

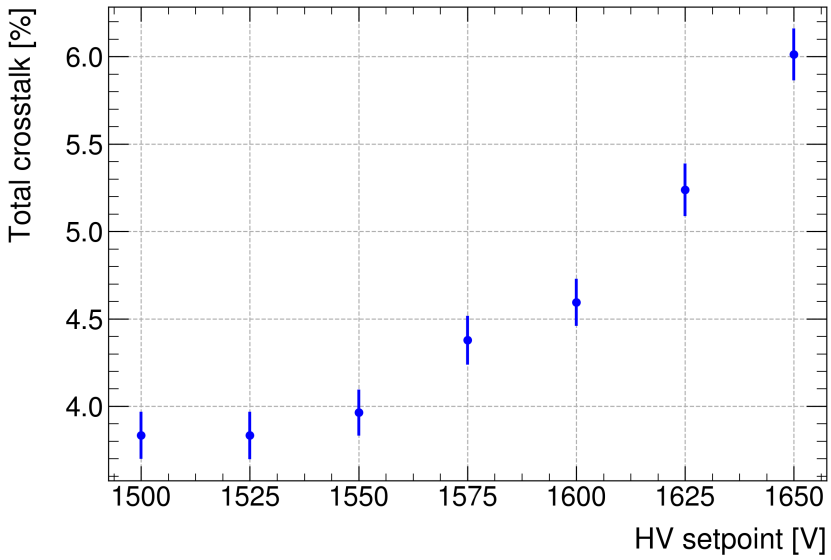

Figure 4.12: The variation of crosstalk probability with tracker HV setting.

The crosstalk increases with voltage in a non-linear fashion, with the shape matching an increasing straw response to higher voltages. Based on this plot, a small reduction in crosstalk could be achieved by lowering the HV to 1625 V, however the efficiency and resolution losses shown in Figures 4.2 and 4.3 are too large to motivate this change.

4.2.3 Dead straw crosstalk estimates

Some straws in the trackers are no longer functional, but still physically present in the modules. These straws failed leak tests and were not suitable to operate in the vacuum, so had the wire removed and ends plugged up with a plastic insert and epoxy, as shown in Figure 4.13.

These provide a unique opportunity to study purely electronic crosstalk in the trackers, since all components for straw-to-straw crosstalk have been removed. Since tracker modules are sometimes swapped, determining the exact positions

Figure 4.13: A photograph of a single dead straw, showing the plastic cap and epoxy covering. Provided by K. Thompson.

of the dead straws in any data is important, so hit rates were checked across every straw and compared to the documentation about which straws should be plugged. The positions of the dead straws found are shown in Figure 4.14. Two straws, marked in green in the diagram, were found to not be completely dead, with an order of magnitude higher hit rate than the other dead straws, so are not considered for this analysis.

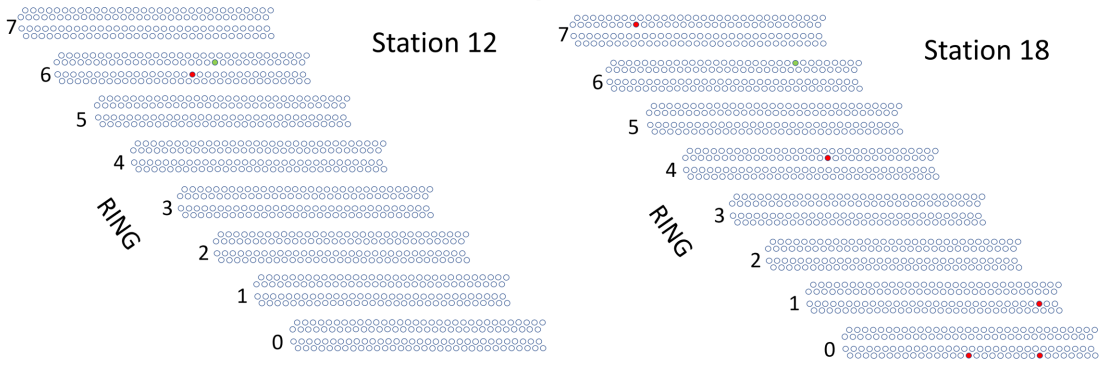

Figure 4.14: The positioning of dead straws in the trackers in Jan 2021, when these studies were performed. The red straws are fully dead, while the green straws still register a significantly higher hit rate.

The electronic crosstalk percentage is obtained by counting the number of hits in the fully dead straws and comparing it to the number of hits in adjacent straws. The same-layer/cross-layer split is determined by checking each time island to find which hit is closest in time to the dead straw hit. This is taken as the corresponding primary hit.

If there was any ambiguity in this determination, for example two real hits with very similar hit times, the whole island would have been dropped. However, in the data sample used all hit origins were clear and unambiguous. The final counts and resultant percentages are presented in Table 4.1.

Straw	Adjacent hits	Dead straw hits	Overall %	Same-layer %	Cross-layer %
1867	18593	429	2.31	0.42	1.89
2055	22472	481	2.14	0.39	1.75
2064	22627	501	2.21	0.40	1.81
2210	17403	396	2.28	0.41	1.86
2637	16418	344	2.10	0.38	1.71
3021	20244	437	2.16	0.39	1.77
		Averages:	2.20	0.40	1.80

Table 4.1: Counts for dead straws and resultant crosstalk percentages.

This indicates that electronic crosstalk accounts for 2.2 % of the overall crosstalk percentage, with 0.4 % coming from same-layer hits and 1.8 % coming from cross-layer hits. This suggests the main source of electronic crosstalk is the capacitors on the ASDQ board, as these are closer together for cross-layer straws than same-layer straws. It also follows that the majority of crosstalk in the trackers generally comes from straw-to-straw interactions, caused by something that the dead straws do not have, so either the straw wire itself, the pins that hold the wire in place, the board connections, or some combination of these.

4.2.4 Track-level crosstalk rates

While a hit-level estimate of crosstalk is useful, it is also informative to count how many tracks have at least one crosstalk hit, since this allows an estimation of how many tracks will be of lower quality due to reconstruction ambiguities. Tracks have between 5 and 32 hits, and the probability of having a crosstalk hit on each is 6.29 $\pm$ 0.03 %. Therefore, as well as simply counting the impacted tracks, a binomial estimate can be made based on the probability and the distribution of the number of hits, shown in Figure 4.15.

The time islands associated with each track are used to make sure all hits are included in the counting. Without the width cut, 62 % of tracks are found to have at least one crosstalk hit, with the average number of crosstalk hits per track being 0.98. Once the width cut is applied, the majority of the crosstalk hits are removed, with only 7.1 % of tracks still having a crosstalk-type hit. The mean number of crosstalk hits per track is reduced to 0.07, and no tracks have more

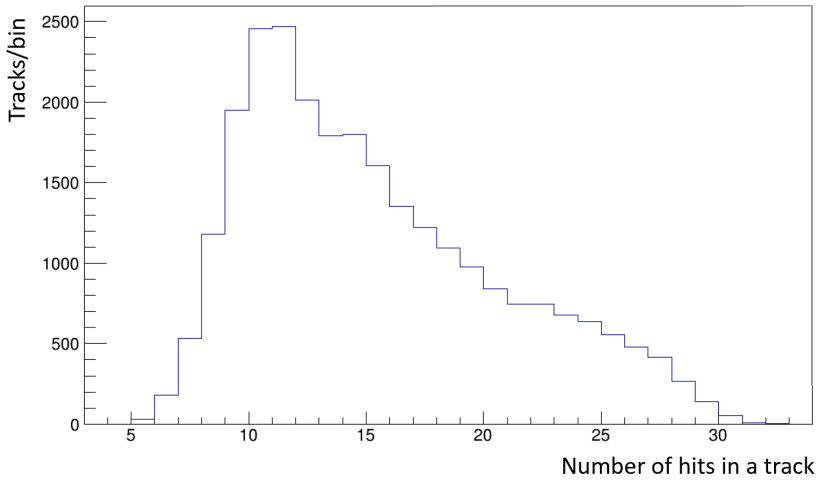

Figure 4.15: The distribution of number of hits per track. This can be used to estimate how many tracks will have at least one crosstalk hit.

than one crosstalk hit. Similarly to the width cut per-hit percentage, these percentages should be taken as an upper bound as some non-crosstalk contamination is expected. Using the probability of a single crosstalk hit, these percentages can also be estimated by assuming a binomial probability distribution. The results of this and a comparison with the directly counted numbers are shown in Table 4.2. The numbers are found to be in good agreement with each other.

	Counting	Counting (+WC)	nHits	nHits (+WC)
% tracks with xtalk	62	7.1	60	7.3
Average xtalk hits	0.98	0.078	0.94	0.081
Average xtalk hits if ≥ 1	1.3	1.0	1.3	1.0

Table 4.2: Comparison of track-level crosstalk percentages and averages from both counting raw tracks and estimating values using the nHits distribution, with and without the width cut (WC) active. The numbers found using both methods are similar.

4.2.5 Properties of primary hits

It is also interesting to consider the primary hits themselves, to study whether there is anything different that makes them more likely to cause crosstalk. Comparing the width distributions between hits identified in tandem with crosstalk, and hits in the tracker with only one hit per layer, it can be seen in Figure 4.16 that the primary hits have a slightly larger width, with a more prominent high-width tail than a non-crosstalk causing hit. However, there is a significant portion of hits at

lower width, so these are capable of causing crosstalk as well. This means the size of the hit is not the key mechanism in determining whether crosstalk occurs.

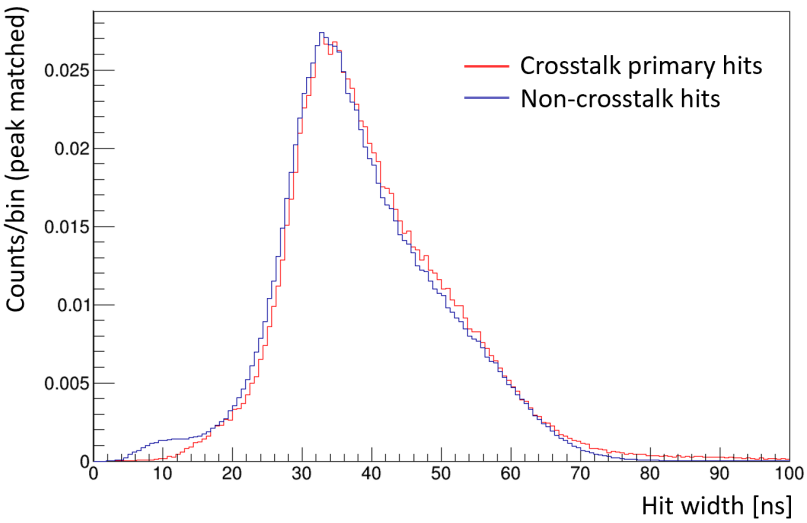

Figure 4.16: Comparison of the hit width distributions for real hits which cause a crosstalk hit, and those that do not. A small difference is seen at high widths, but in general both look similar.

The drift time can also be compared to study the structure of primary hits, as shown in Figure 4.17. There is a much larger visible difference in shape, with the primary hits having fewer occurrences in the very low and very high drift times.

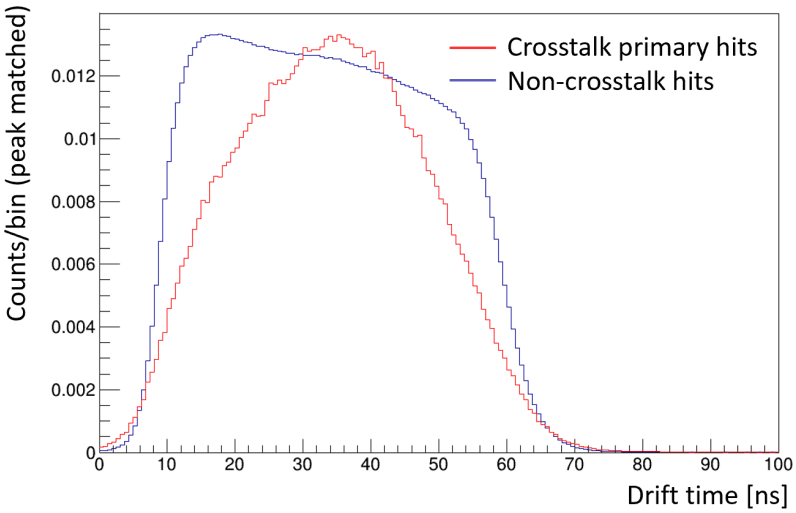

Figure 4.17: Comparison of the drift time, scaled for shape comparison. A much larger difference is seen, with crosstalk-causing hits preferring mid-range hit times.

The high end shape of this distribution is explained by the correlation of the drift time to the width, since if the width of a pulse is large, the first ionisation to

reach the wire must be earlier on. The low width end of the distribution is more interesting, as it indicates that primary hits do not occur close to the straw wire but instead at a larger radius. Combined with the hit width tail, this suggests that hits are more likely to cause crosstalk if they are single, large ionisations rather than a combination of multiple smaller charges.

4.2.6 Crosstalk in simulation

Although crosstalk can be studied well in data, the impact of it is difficult to fully disentangle from other effects, motivating a complimentary study of crosstalk using MC-generated data. The crosstalk is manually injected into either generated or real tracks, allowing a study of how varying the crosstalk percentage impacts tracking and a cross-check of the selection method in data. If the injected percentage is retrieved, this supports the reliability of previous results.

Both MC data and real data can be used for these studies. The pure MC approach uses the Gas Gun as the starting point, adding crosstalk to real hits before fitting tracks. One limitation of the MC is that it does not model the hit width, a key identifying property of crosstalk. This motivates a 'hybrid' method using real data as a base with crosstalk hits added. Both sets of hits are tracked up to the time island stage, then all time islands with more than one hit per layer are removed. This provides a completely clean set of hits to add crosstalk to.

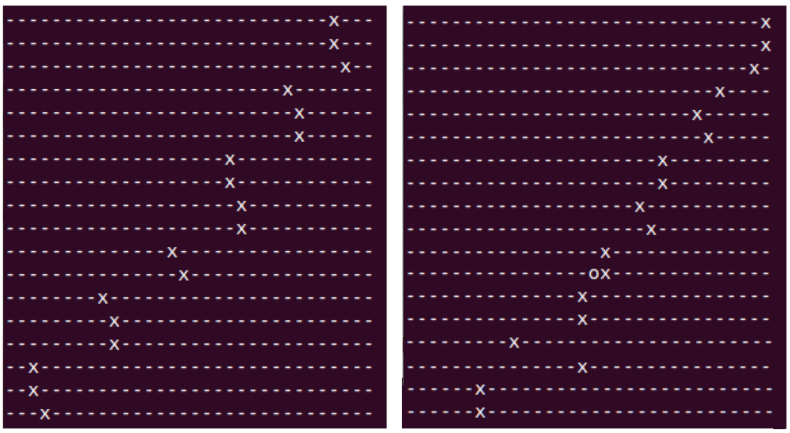

Figure 4.18: Two tracks displayed in a simple event display written by the author. The 'x' markers indicate real hits. The track on the right has had a crosstalk hit injected and identified, shown by the 'o' marker.

The crosstalk injection is performed by looping over all raw cleaned hits, adding new hits adjacent to these based on a user-defined percentage. The cross-layer and

same-layer split is also implemented, allowing the MC crosstalk to match data as closely as possible. The final step needed is to decide whether the new hit should be to the left or right of the primary hit. Since both are equally likely in data, this is set within the module as a 50-50 split. If a new hit would be defined outside the tracker due to the primary hit being on an edge, it is recalculated using a different primary hit in the same track to keep the overall percentage constant. The position of the injected hits are checked visually on an event display to check the new hits are in the expected locations. An example of this is shown in Figure 4.18.

The physical attributes of crosstalk hits must be defined as these are needed for track fitting. For simplicity, most of these are copied from the primary hit, with the hit width, hit time, drift time and DCA sampled from distributions seen in data. The particle ID flags are also updated to make it clear that these are not primary hits. The hit time for the simulated crosstalk hit is defined as the sum of the primary hit time and an added hit time difference, which is sampled randomly from the distribution seen in data. Since the total drift time must be less than 70 ns, any samples that would push the sum above this limit are resampled until they lie below it. The DCA is calculated from the drift time using a parametrisation from Run 1 data, as detailed in [66]. Finally, the hit width is defined by sampling from a Gaussian fitted to the low width peak in data. As shown in Figure 4.19, the resultant distributions do not match real data perfectly. As a check, since the width cut will be investigated, a cut is applied to make sure similar hit fractions are removed. A cut at 14 ns removes 6.23 % in data and 6.05 % in MC, which are similar enough, but this difference must be kept in mind when interpreting results.

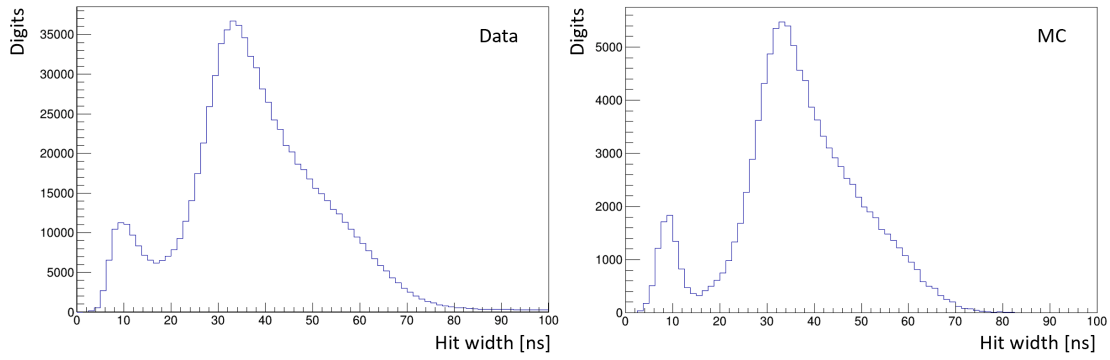

Figure 4.19: Comparison of hit distribution for real data (left) and hybrid method MC (right). With a width cut of 14 ns, the Gaussian approximation removes similar fractions of the overall distributions.

To test the crosstalk selection method used for data, two datasets were generated, one with the input percentage set to a large, clearly observable value of 50 %,

with a 50-50 same-layer/cross-layer split, and the other with percentages chosen to match real data. The selection cuts are then applied, with the results shown in Figure 4.20. A flat line is fitted to the points to extract the final percentage, with 46.3 ± 8.77 % being extracted from the dataset with 50 % input, and 6.10 ± 3.19 % from the dataset set to mimic data. Both of these values are consistent with the input, indicating that the percentage determination is accurate. The error bars in Figure 4.20 are purely statistical and are large due to the small size of the MC dataset, so the difference between same- and cross-layer crosstalk is not visible. Also interesting to note is that both of these plots are flat with time in fill, and do not show the characteristic rise at early times seen in the data. This further confirms that the selection method is not the cause of the time dependence, but rather something present in data but not in the simulation.

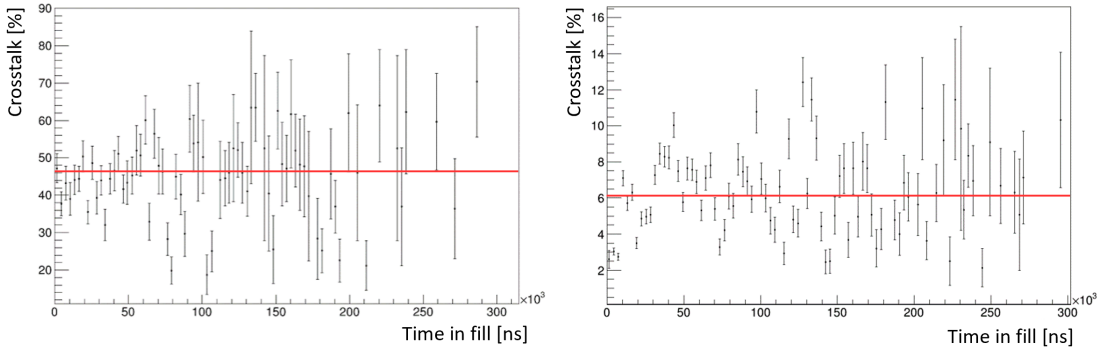

Figure 4.20: Outputs from the data selection method, for 50 % crosstalk and then data levels of crosstalk. The output percentages were 46.3 ± 8.77% and 6.10 ± 3.19%, both consistent with the injected input.

The impact of increasing the crosstalk percentage is investigated, considering 3,6,9 and 12 % to scan over a range near the measured value in data. The hybrid method is used here to ensure everything is as close to real data as possible, with the width cut removed to allow an evaluation of the full impact. Distributions are plotted with and without the crosstalk hits to compare. For 50 % crosstalk, the impacted track properties were found to be the p-value, degrees of freedom, tracker resolution and vertex resolution from the extrapolation to decay vertices. A large reduction in the final number of tracks was also seen, implying that crosstalk hits cause the track finding to fail more often. The track failure modes for different input percentages can be seen in Figure 4.21.

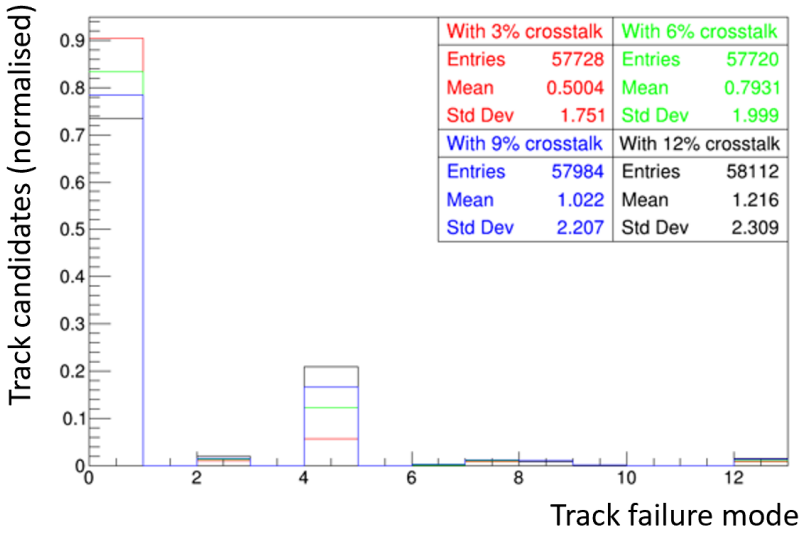

Figure 4.21: All track candidates and the track failure mode for different crosstalk input percentages. All failure modes are detailed in Table 4.3.

Failure mode	Cause of tracking failure
0	No failure, succesful track fit
1	Negative or NaN χ^2 value
2	Divergent χ^2 in fit
3	Unphysically recurling track
4	>1 hit in a layer
5	Start position defined before first hit
6	Geant4 error propagation failure
7	Track step length 0
8	Track step length overflow
9,10	No convergence in track fit
11	Non-positron initial hit in MC
12	Track outside vacuum chamber

Table 4.3: The definitions of different track failure modes used in plots in this thesis. There are two errors for lack of convergence due to two different track fitting modes available as options.

The total number of track candidates increases as the crosstalk percentage increases. Since the base dataset used here is the same, this increase is due to extra hits making the tracking break up single candidates into multiple candidates due to ambiguity about which hits belong to each track. In Figure 4.21, a failure

mode of 0 indicates a successful track has been made. This bin shows a steady decrease as the crosstalk percentage increases, with failed events primarily shifted to bin 4, which corresponds to track candidates with more than one hit per layer dropping the layer and therefore failing cuts. There is also a small increase in the number of track candidates where the fitting χ^2 does not converge. This is likely a consequence of the crosstalk hit perturbing the track, leading to a failed fit. This loss, which amounts to approximately 2.2 % of candidates per 1 % absolute crosstalk, could hypothetically be regained by identifying and removing the crosstalk hits rather than dropping the layer.

As well as the reduction in the overall number of tracks, the ones which do survive the cuts are on average lower quality than with no crosstalk. Figure 4.22 (left) shows that tracks with crosstalk have a worse p-value on average, so more are removed by the quality cuts.

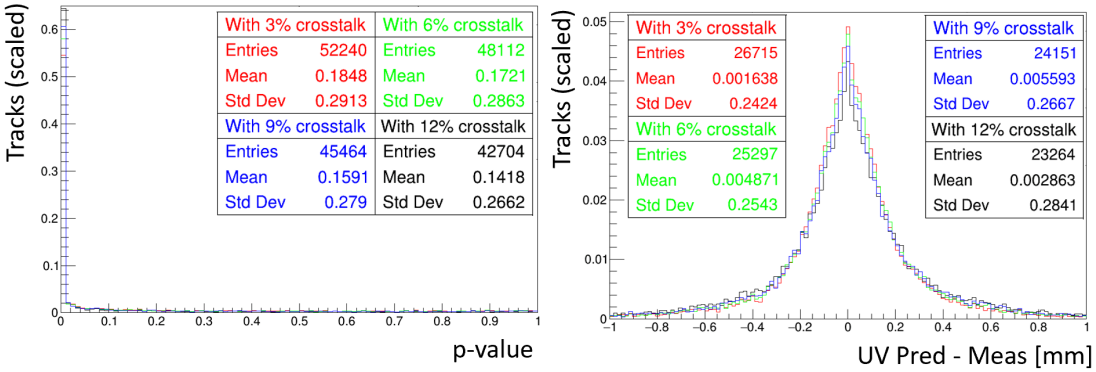

Figure 4.22: Variation of p-value (left) and tracker resolution (right, considering the standard deviation of the predicted-measured hit positions) with increasing levels of crosstalk.

The tracker resolution, which in this case is calculated as the RMS of the residual distribution shown in Figure 4.22 (right), also worsens at higher crosstalk percentage, likely due to the larger spread caused by the true hit ambiguity. This worsens at the rate of around 1.6 % per 1 % of crosstalk.

Finally, the impact of lower quality tracks can be seen when extrapolating to the decay vertices. As shown in Figure 4.23, both the radial and vertical distributions have a larger RMS with increasing crosstalk. This has a direct knock-on effect on the systematic uncertainties for the ω_a and EDM analyses.

The hybrid method is also tuned to match data to quantify the impact seen in real data. As before, crosstalk is added and is compared to the data without crosstalk. The resultant distributions can be seen in Figures 4.24, 4.25 and 4.26.

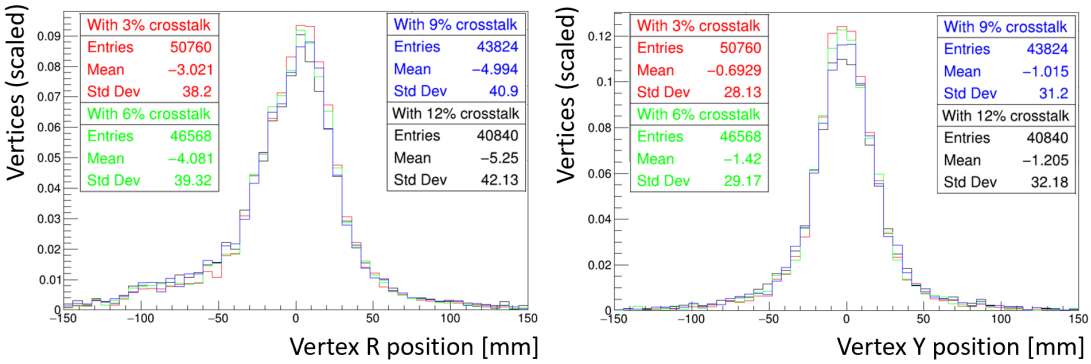

Figure 4.23: Variation of R and Y distributions as crosstalk is increased. More crosstalk leads to a larger uncertainty in both directions.

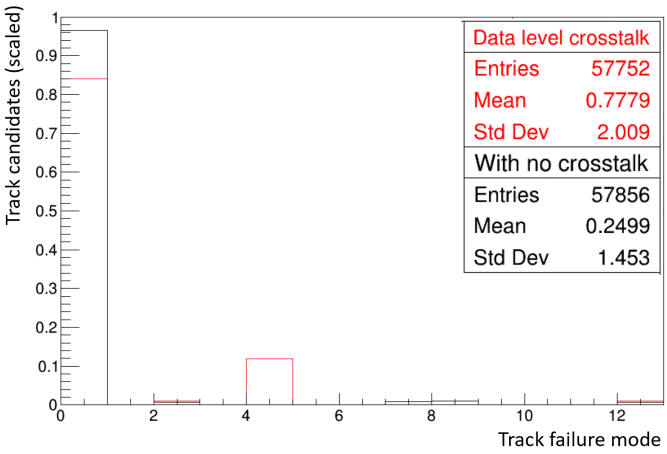

Figure 4.24: Track failure modes for data-level crosstalk.

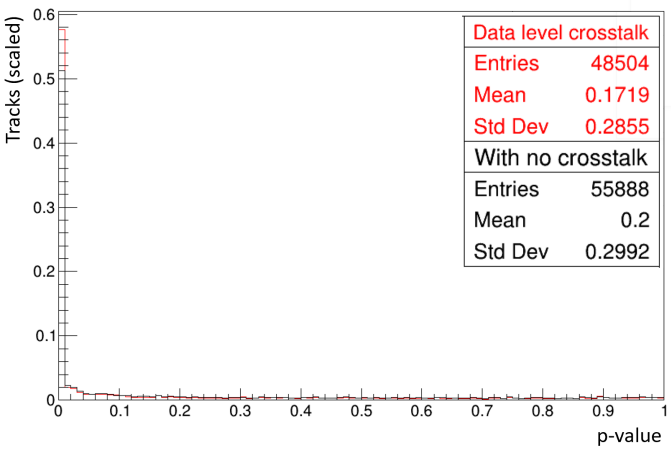

Figure 4.25: p-value variation for data-level crosstalk.

From the track failure mode plot, is it seen that 13.2 % of track candidates are lost to the candidate layer cut, consistent with the predicted numbers from the

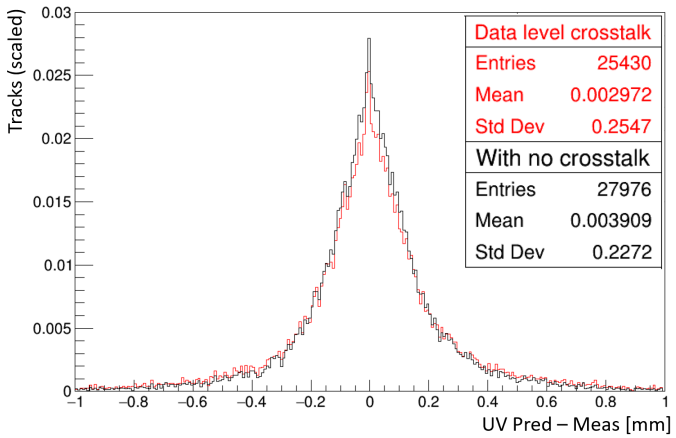

Figure 4.26: Track resolution loss for data-level crosstalk.

rough scan. The p-values are on average 14 % worse, which equates to a loss of 4.2 % of the quality tracks. The raw track resolution is 12 % worse, compared to the estimate of 11.3 % from the percentage change.

A full comparison of the quality cuts failed is shown in Figure 4.27. This shows that while the p-value cut has a large impact, the largest loss comes from the cut on the residuals, removing 24.5 % of tracks.

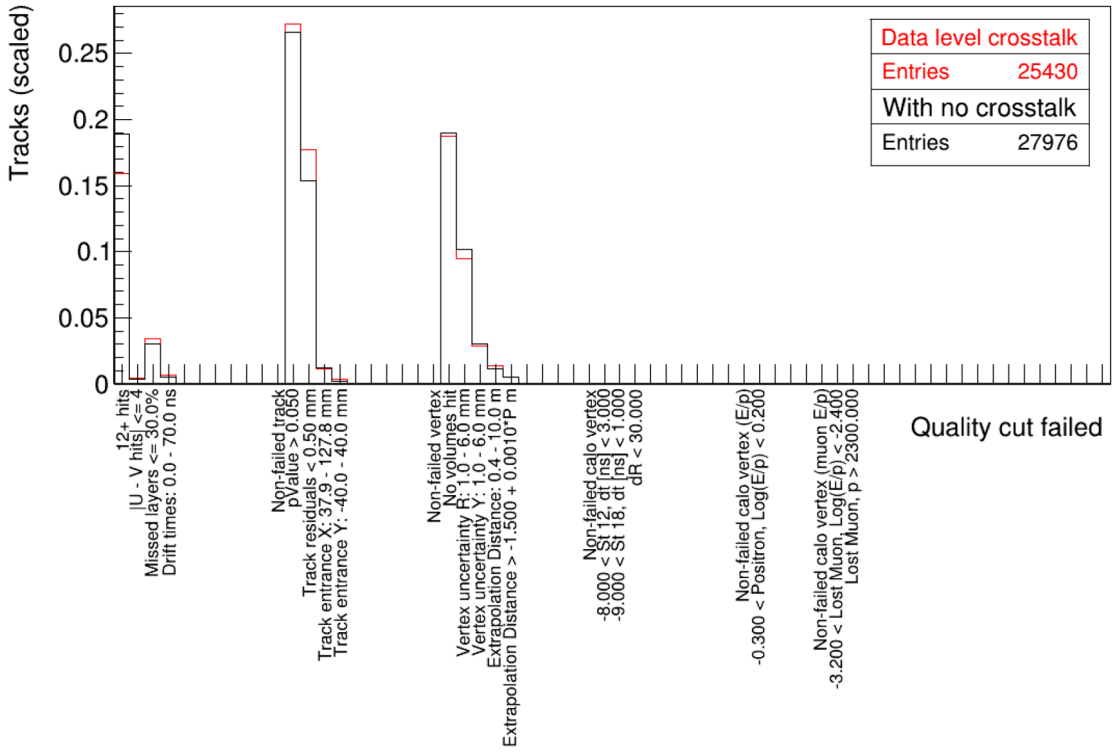

Figure 4.27: Tracks lost to quality cuts with and without data-level crosstalk.

Despite the loss in statistics, the quality cuts do a good job removing crosstalk-impacted tracks. This can be checked by plotting the above comparisons after the cuts have been applied, as shown in Figure 4.28.

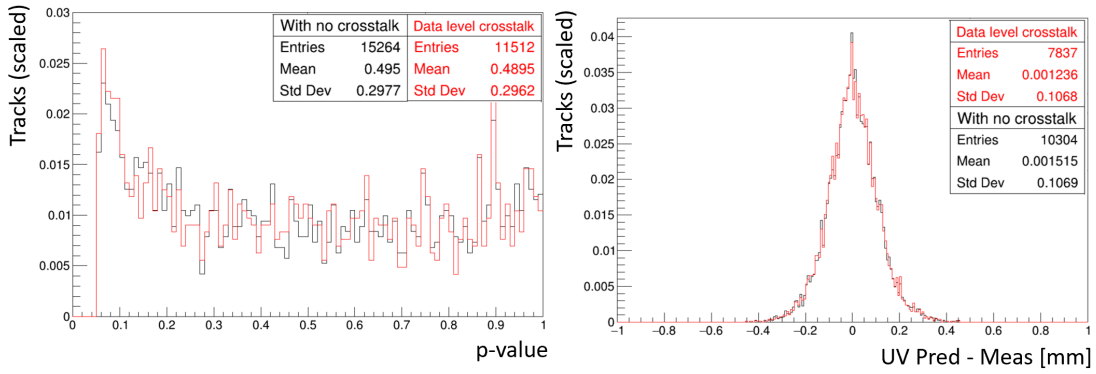

Figure 4.28: A comparison of tracks with and without crosstalk, after the quality cuts have been applied to both. The cuts have removed the majority of any bad tracks affected by crosstalk.

The MC can also be used to check how effective the width cut is at removing crosstalk, and whether this is sufficient to recover tracks. Figures 4.29, 4.30 and 4.31 show the results for all tracks with the width cut applied, but no quality cuts.

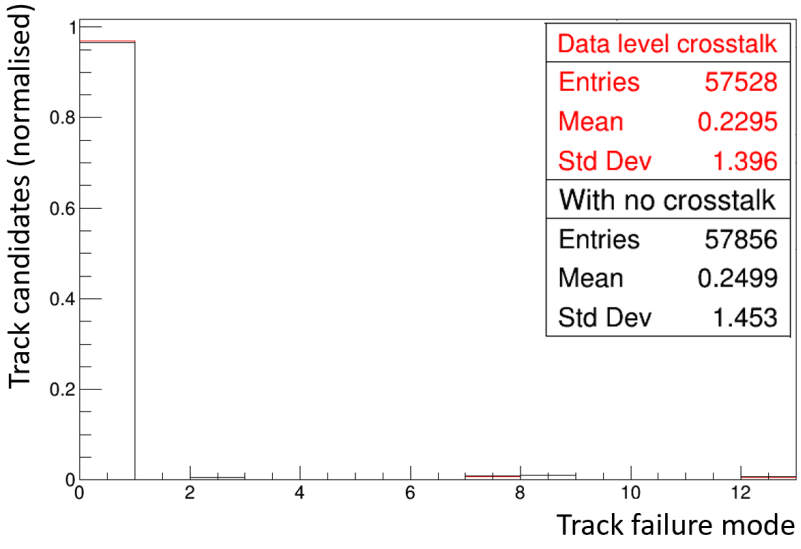

Figure 4.29: Track failure modes with and without crosstalk, after the width cut has been applied.

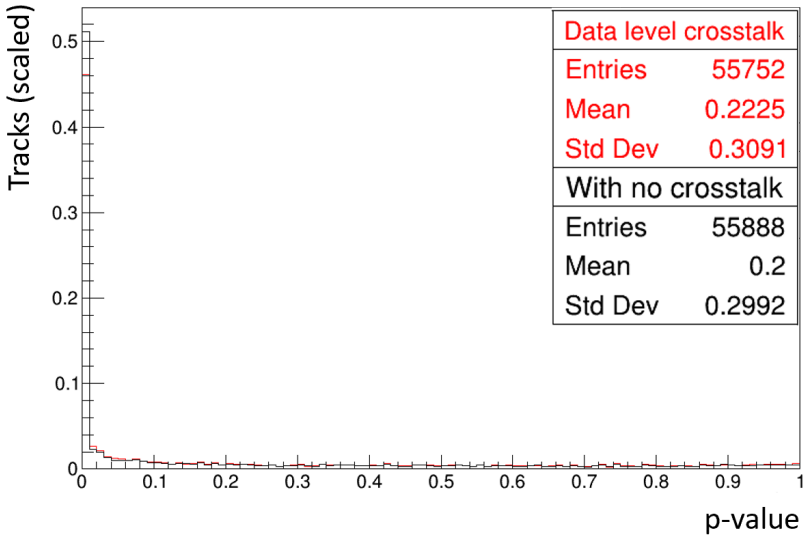

Figure 4.30: Track p-value with and without crosstalk, after the width cut has been applied.

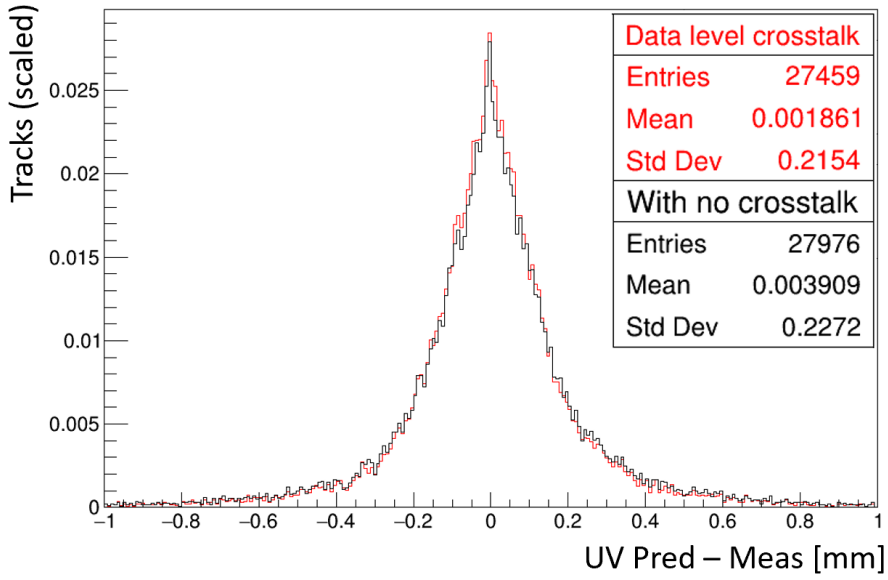

Figure 4.31: Track resolution with and without crosstalk, after the width cut has been applied.

The width cut successfully removes all effects from the crosstalk in simulation, with very little difference both in shape and number of tracks. When the quality cuts are applied as well, shown in Figure 4.32, the difference becomes even smaller, with almost the same tracks passing between the two datasets. From this, it can be concluded that at least for the MC generated crosstalk, the width cut does a very good job of mitigating the effects of the fake hits.

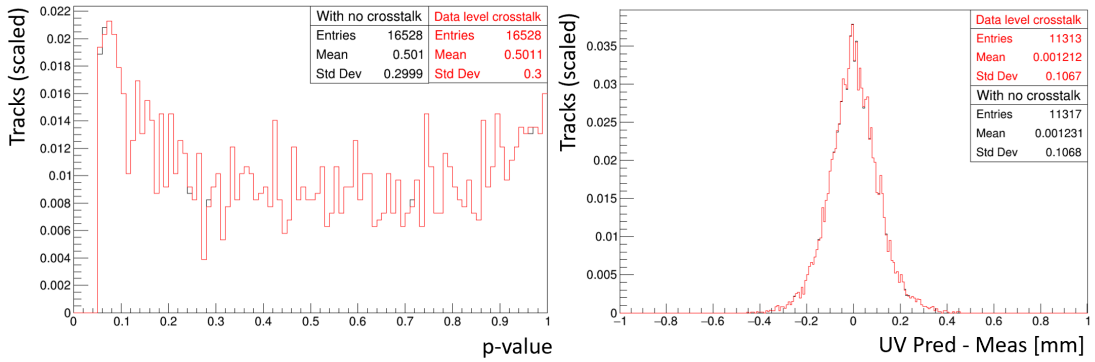

Figure 4.32: Track p-value and resolution, with both a width cut and the track quality cuts applied. The crosstalk tracks are completely removed, leaving almost the same tracks in both plots.

The MC however does not include all effects that are seen in real data, demonstrated by the lack of rise in Figure 4.20. The origins of this rise are discussed more in Section 4.3, but as the width distribution is changing rapidly at early times, it is possible that this is allowing more crosstalk though the width cut. This would mean it may not be as effective as the MC would suggest. Nonetheless, this is likely to be a small change, since the final level of crosstalk in data doesn't increase above the plateau percentage used for the MC studies.

4.2.7 Removing crosstalk from data

Despite being very effective at removing crosstalk, the width cut and candidate layer removal also remove a fraction of good hits and tracks. Therefore, other methods are investigated in an attempt to regain some of these. This would appear to be simple, but ultimately complications with crosstalk-like cases cause too many issues for more sophisticated methods to outperform the simple cut.

One alternative to a width cut is to try using time differences. This is expected to work well on same-layer crosstalk but less well with cross-layer. A similar comparison could also be done with the width, where instead of defining a fixed cut point, adjacent hits are compared, and the lowest width hit is removed. This should in theory deal well with both same- and cross-layer crosstalk, but will throw out real hits in the case of overlapping tracks. Finally, specific patterns of hits can be identified, for example triplet clusters of hits. If the two outside hits are crosstalk, this would mean the middle hit must be real. This, however, also has difficulties with additional hits on top of a track if not caused by crosstalk.

Figure 4.33 shows the width distribution after these various crosstalk removal

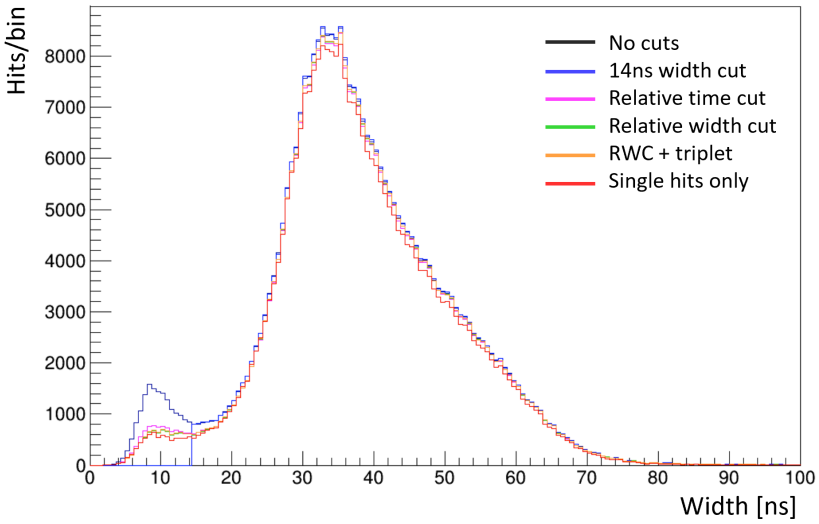

Figure 4.33: Width distribution after several crosstalk removal methods. All function very similarly, with the 'one hit per layer' representing the best possible removal without a width cut.

methods, compared with a hard width cut and throwing out any layers with more than one hit. This shows that most of these methods have very similar efficiency at removing crosstalk, other than the hard width cut, which is likely also removing good hits. The levels of crosstalk can be quantified by using a simulated crosstalk input and running it through the selection methods to count how many hits are removed. The results are shown in Table 4.4.

Method / [%]	Xtalk removed	Real hits removed	Remaining xtalk
14 ns width cut	89	0.7	0.69
20 ns width cut	100	2.2	0.00
Only isolated hits	85	4.6	0.94
Drop layer if > 1 hit	86	10.2	0.88
Relative width cut	77	1.4	1.44
RWC + triplet	75	1.6	2.01
Relative time cut	68	1.7	1.58

Table 4.4: Comparison of final systematics found using the two methods.

The alternatives are not as effective as the width cut, and often remove a greater number of real hits. Combining the relative width cut method and the triplet method does not help remove more crosstalk either, though it does lower the number of real hits incorrectly removed. The reason these methods are not as effective comes from the ambiguity in selecting crosstalk hits, as there are too many similar signatures. Therefore, the width cut at 14 ns remains the optimal removal method.

4.2.8 Conclusions

The impact of crosstalk was studied in the tracking detectors using both real and simulated data. A methodology for identification of crosstalk hits was developed and used to investigate the variation in number of hits with tracker HV, including the percentage at the current operational setup. An overall crosstalk percentage of 6.29 ± 0.03 % is found, with 3.00 ± 0.02 % coming from the same tracker layer and 3.29 ± 0.02 % from the adjacent layer. With cuts implemented this is reduced to $\leq 0.51\%$. A rise in crosstalk percentage is observed at early in-fill times, with a lifetime of approximately 50 μs.

Likely sources of crosstalk are discussed, using data from straws that can no longer detect real hits. This indicates that crosstalk occurs both in the straw components and on the ASDQ board, with the majority (4 %) coming from straw components. On the board, there is significantly higher cross-layer crosstalk compared to same-layer crosstalk, which suggests it is due to the capacitors on the board. For straw-to-straw, the remaining percentage must come from either the wire, the pins, or the connections to the board.

Finally, the results from data are used as simulation inputs to study how the tracking is impacted by crosstalk. Simulation studies reveal a total loss of 13.2 % raw tracks and 24.5 % quality tracks for the data-level crosstalk case compared to the no-crosstalk case. The tracks with crosstalk are lower quality, with a lower average p-value and larger residuals. These issues are magnified by extrapolation, with the vertex resolution similarly negatively impacted. The quality cuts do however successfully remove the tracks negatively impacted, leaving good quality data. The width cut is found to be very effective in MC, successfully mitigating the negative effects of the crosstalk by removing the hits before track candidates are made. However, these results do not include the time-dependent effect of the width, which may let more crosstalk through than this study would suggest. Nonetheless, the MC results are encouraging, as they indicate that other than the time dependence, the crosstalk is being removed near-optimally. Other methods were tested based on the knowledge learned from the studies, however none are as effective as the width cut.

4.3 The time-dependent resolution

During investigation of the crosstalk, a rise at early times was seen which did not match any known track effect. This section covers the investigations done into this effect, and potential theories of what is causing it.

Since the rise is seen in crosstalk, it is logical to check whether other reconstructed track parameters are impacted, such as the resolution and efficiency. The resolution is calculated as detailed in the HV scan section for slices of time between 0 and 500 ns, with the result of this shown in Figure 4.34. This is fitted with a falling exponential plus a constant to extract the lifetime.

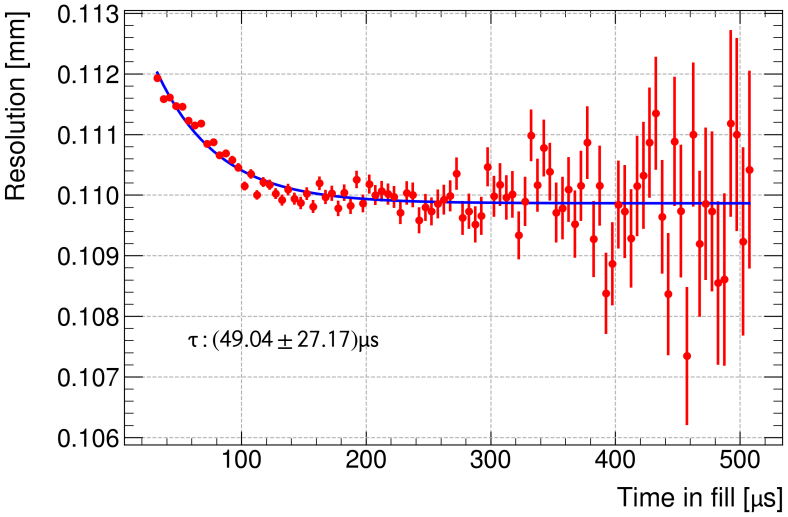

Figure 4.34: The resolution plotted as a function of time, and fitted with a falling exponential to extract the lifetime.

A clear trend is seen, with a worse resolution at early times which slowly improves until reaching a plateau at late times. This shows that not only the crosstalk is impacted, as if this was due to crosstalk, the resolution would be better at early times as the percentage is lower. The lifetime is consistent with the rise found in crosstalk, hinting at an underlying tracker effect. A drop in efficiency at early times is also seen, as shown in Figure 4.35 where several time slices have been plotted.

One potential explanation for all the observations would be a voltage drop at early times. Therefore, the HV scan data can be used to investigate whether the magnitude of this effect varies with the tracker voltage. Each subset of data collected during the scan is sliced in time and plotted to extract the resolution. The results for 1650 V vs 1600 V are shown in Figure 4.36.

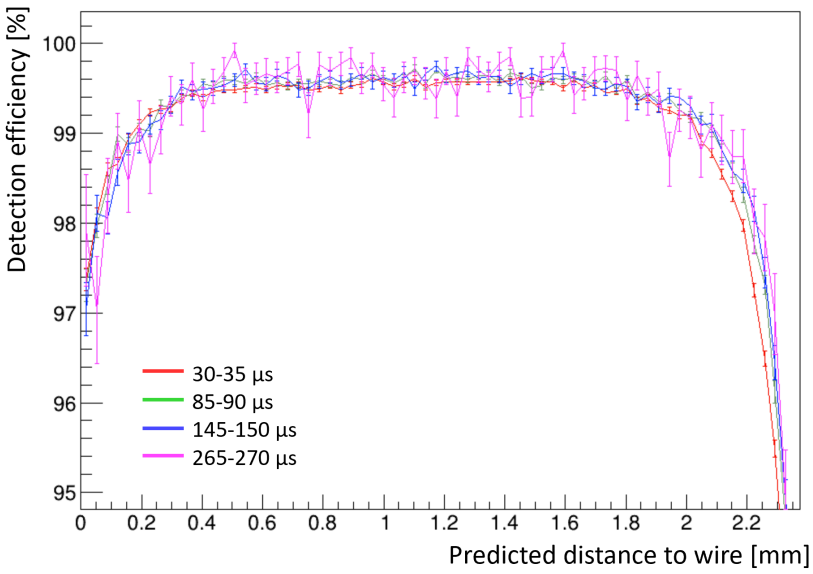

Figure 4.35: The efficiency as a function of DCA plotted for several different time slices. The drop in efficiency is shown by the red line being below the others at high DCA.

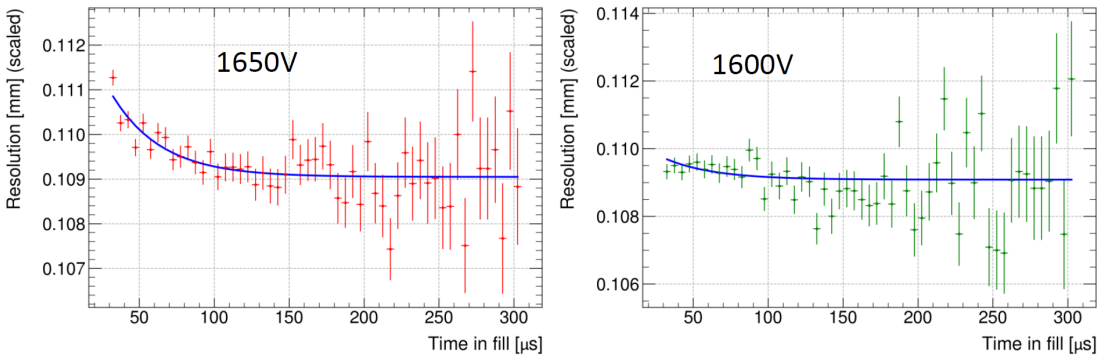

Figure 4.36: A comparison of the time-dependent resolution for 1650 V and 1600 V. The effect is much larger for 1650 V.

There is a clear difference in shape, with the higher voltage data having worse resolution at early times. Therefore, this effect must be caused by a lowering of the voltage seen by the electrons and ions in the straws. One potential explanation is a 'space-charge' effect, where a high rate of particles saturates the straws with positively charged ions, which shield the wire from the ionisation elections, temporarily reducing the effective voltage. If this is the cause, it would have two signatures which can be checked. Tracker Station 12 is closer to the injection point of the muon beam than Station 18, so it sees a higher rate of particles. If the space-charge effect is responsible, it should therefore be larger in Station 12 than

Station 18. A comparison of this is shown in Figure 4.37. Indeed, a difference is seen, with Station 12 having a significantly worse resolution at early times.

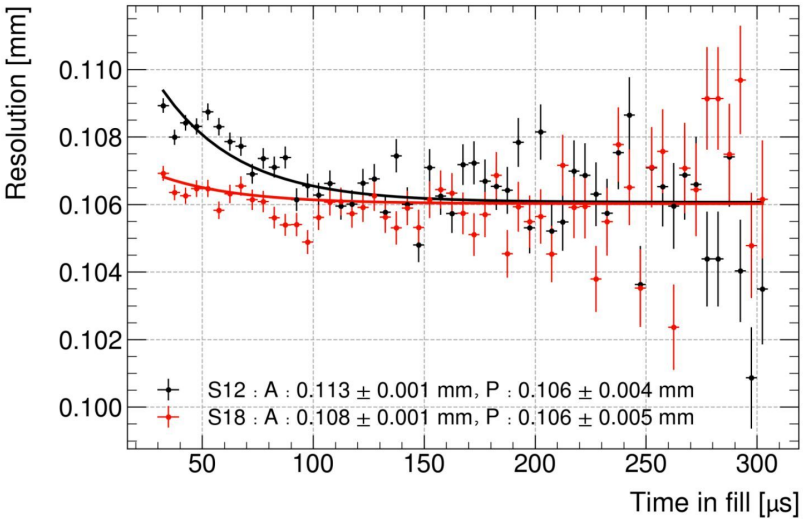

Figure 4.37: A comparison between tracker stations. Station 12 is found to have a much worse resolution than Station 18 at early times.

Additionally, special runs performed to investigate high and low momentum parts of the muon beam are used to check that intensity is the driving factor behind this effect, as the low and high momentum beams have worse storage so a lower intensity. A comparison of these two run conditions to the nominal running is shown in Figure 4.38. Again, the time dependent effect is much smaller in these special runs, supporting the space-charge hypothesis.

Finally, following a detailed mathematical parametrisation of this potential effect by J.Mott [83], the expected change in resolution can be compared to the change measured as a function of tracker voltage. This is shown in Figure 4.39, and matches the shape well. The slight offset is likely due to a free normalisation parameter that defines the base resolution in the space-charge model. This can be expected to not match data perfectly.

Since this is an early-to-late effect that will not be completely removed by time cuts around 30 μs, this must be accounted for in any analysis using tracker data. In most places, the impact is an extra part to the systematic uncertainty to account for slightly worse reconstruction at early times, but in some analyses, such as the EDM analysis, a correction is needed. The calculation of the systematic uncertainties and corrections are covered in Section 4.4.2.

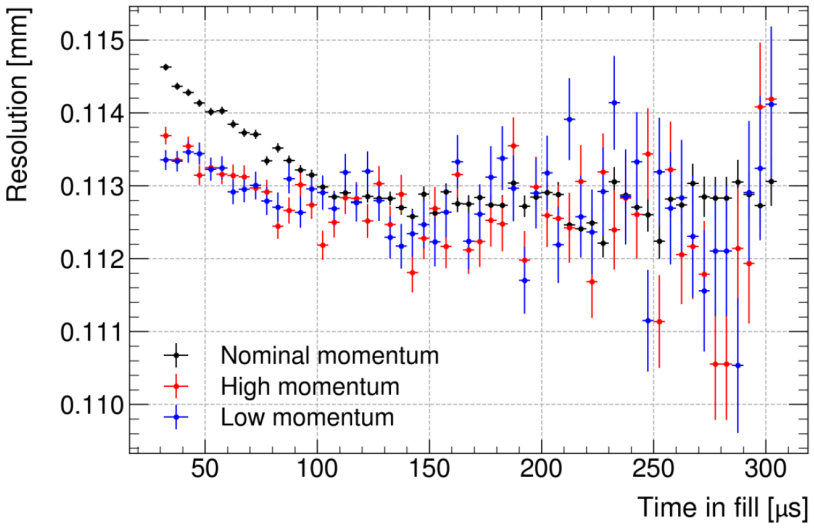

Figure 4.38: The resolution change for nominal running, and two special runs of high and low momentum muons which have 1/6 the usual beam intensity.

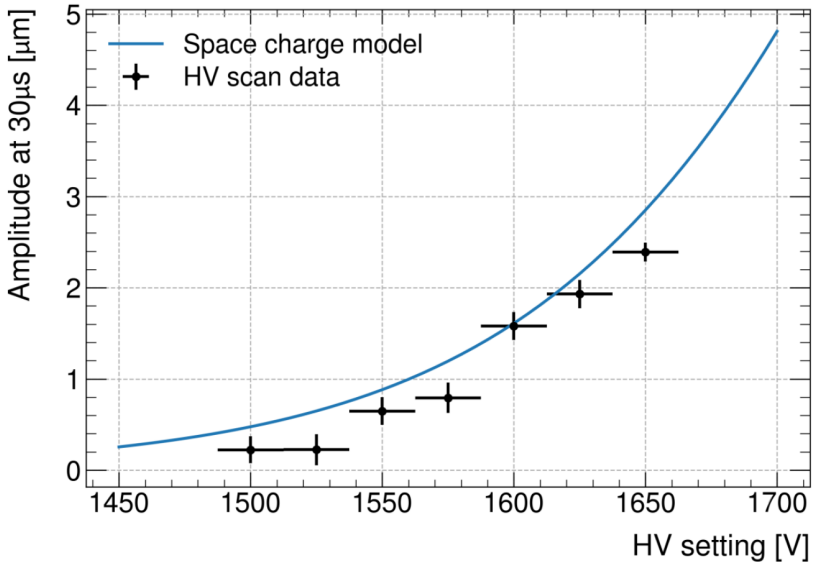

Figure 4.39: A comparison of the amplitude of the resolution rise (used as a proxy for the size of the rise) with the expected values from a space-charge model developed by J.Mott [83].

4.4 Tracker systematics

Due to the ability to monitor the beam, tracker data is a core part of the ω_a analysis. As a result, any uncertainties need to be carefully quantified. Here, the systematic uncertainties on the beam distributions due to crosstalk, measurement resolution, and time-dependent effects for the Run 2/3 time period are evaluated.

4.4.1 Crosstalk

The crosstalk systematic quantifies the impact of extra noise hits on the beam distributions. One possible procedure is similar to the method used in Section 4.2.6, where a sample of data from the Run 2/3 period is 'cleaned' by removing any tracks with more than one hit per layer. This will remove effects like overlapping tracks, giving a suitable baseline for purposes of comparison. This is run through the crosstalk MC generator, which assigns crosstalk hits randomly with the same probability as they appear in data (6.3 %). These two datasets are run through the full tracking and extrapolation, including all quality cuts, and the final beam distributions are compared to quantify the difference. The uncertainties on the differences are calculated as the sum in quadrature of the statistical uncertainties on the mean and RMS and the difference between data and MC. These distributions for both the cleaned hits and the hits with injected crosstalk are shown in Figure 4.40 for one subrun of data.

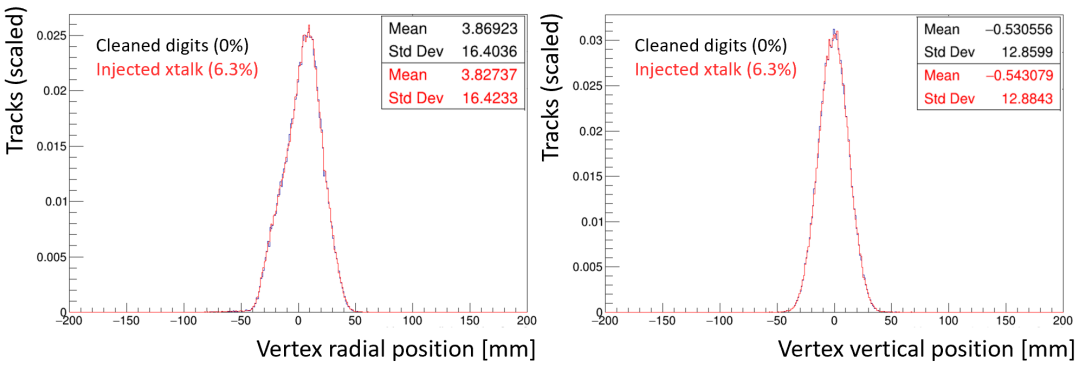

Figure 4.40: MC vertex distributions for the R and Y directions, with and without a data level of crosstalk injected. The difference allows the systematic uncertainty from the extra hits to be evaluated.

Crosstalk biases the mean in the negative direction both radially and vertically, and slightly increases the width of the beam in both directions. The radial bias is larger than the vertical bias as crosstalk hits cause fitting problems in the radial plane rather than the vertical plane. Nonetheless, as the quality cuts are effective at removing poorly fitted tracks, the overall impact is small compared to other systematics like the tracker alignment, which can shift the beam on the order of a mm.

The MC injection of crosstalk uses several approximations, in particular in the distributions of the hit width and time. To ensure these estimates do not introduce a bias, a data based method is designed to roughly cross-check the values. Since

there is no way to remove all crosstalk from data, the main cleaning methods discussed in Section 4.2.7 (width cut and relative width cut) are used, with the resulting data run through the tracking again to compare it to the raw data. The differences in beam shape are then plotted against the remaining percentage of crosstalk after each cleaning method, and a line is fitted to extrapolate to 0. The beam variables at this point are compared to those in the raw data to estimate the full impact.

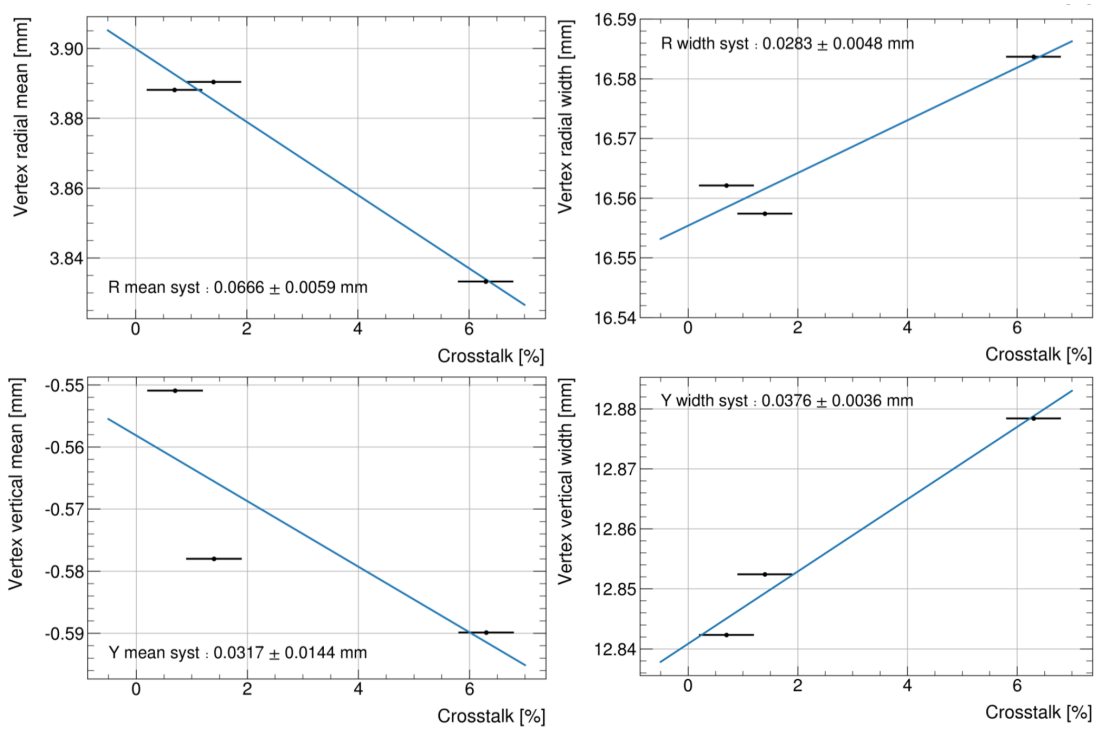

Figure 4.41: Fits to beam variables plotted against crosstalk percentage in data, determined by the known percentage left behind by removal methods.

As can be seen in Figure 4.41, this method works reasonably well, giving a set of points roughly consistent with a line. The variation likely comes from the fact that the cleaning methods are not perfect, and will remove hits that are not crosstalk, with the width cut being the harshest of these as it also removes real hits. The final set of systematics are in Table 4.5.

[mm]	Mean R	Width R	Mean Y	Width Y
MC method	-0.042 ± 0.008	0.020 ± 0.160	-0.012 ± 0.047	0.024 ± 0.007
Data method	-0.067 ± 0.005	0.028 ± 0.005	-0.032 ± 0.014	0.038 ± 0.004

Table 4.5: Comparison of final systematics found using the two methods.

The data method gives systematics that are higher than the MC method, however the same patterns are seen in both. Both are relatively very small compared

to other tracker systematic effects and general analysis uncertainties, so the more conservative guess of the two (the data method) is used for the final values.

4.4.2 Measurement resolution

An estimate is also needed of how well the raw hit resolution is known, and how this translates into the vertex resolution used in the analyses. To estimate the impact of an uncertainty in position, an artificial smearing in hit position is applied. This is performed by sampling from a Gaussian with a set smearing width, varied between 0.1 and 0.25 mm, and mean 0. Since the base data already has a tracker resolution, the overall resolution achieved is a sum in quadrature of the base tracker resolution and the width of the smearing. By applying different smearing widths, a range of resolutions can be plotted against the relevant beam distributions. One downside to this method is that it can only increase resolution, not decrease it. However, symmetry is a reasonable assumption as previous studies in MC found the change in vertex resolution to be approximately linear with hit resolution.

One effect which also complicates the evaluation of this systematic is that the resolution is known to be time dependent in the trackers. For clarity, the impact of the time dependence is separated out into a second systematic, and all resolution calculations are done in the plateau region where the resolution is constant. The change in resolution between early and late times can then be used to estimate the time dependence systematic. Based on Run 1 measurements, the uncertainty on the tracker resolution was estimated conservatively to be $\pm$ 30 μm [84]. This is checked for Run 2 and Run 3 by measuring the resolution for a randomly chosen sample of subruns to check the variation. The results of this can be seen in Figure 4.42.

This shows that there is some variation between subruns, particularly between Run 2 and Run 3, with the difference between runs at most around 5 μm, which is small compared to the 30 μm estimate. As a result, the conservative estimate is still valid, and is used for these two runs as well.

The hit position is smeared in fractions of the overall resolution, and the beam variables are plotted against it. The trends in beam vertical and radial mean are shown in Figure 4.43. Overall, these look linear, with the exception of one point at the largest smearing. As a result, this point is discounted as a random fluctuation and a line is fitted. This allows the beam bias to be quantified for any uncertainty on the resolution.

Figure 4.44 shows the vertical and radial beam width variation with increasing smearing. These follow a more quadratic trend, so are fitted with a polynomial of

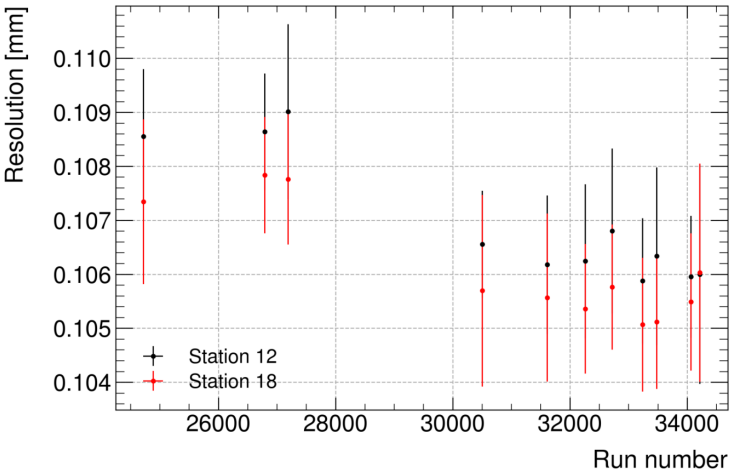

Figure 4.42: Variation of the resolution, with the first 3 points from Run 2 and all others from Run 3. All are consistent, even across runs.

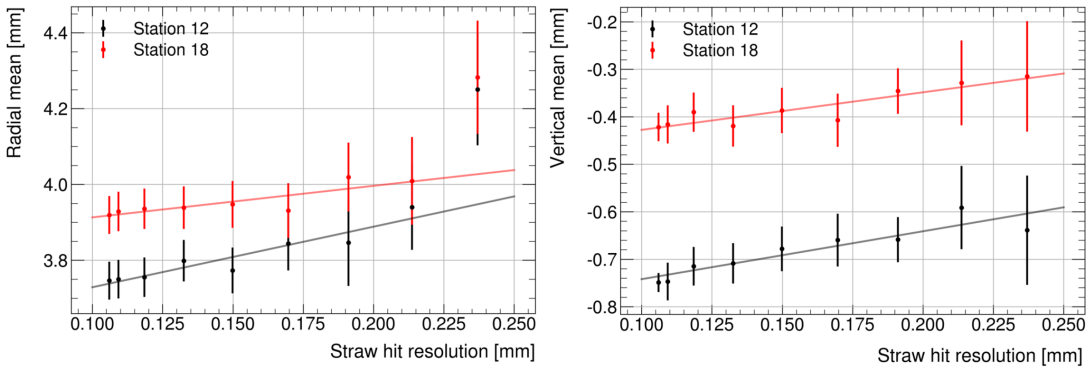

Figure 4.43: The variation of the radial and vertical means with increased smearing. The final point in the radial direction is not used for fitting.

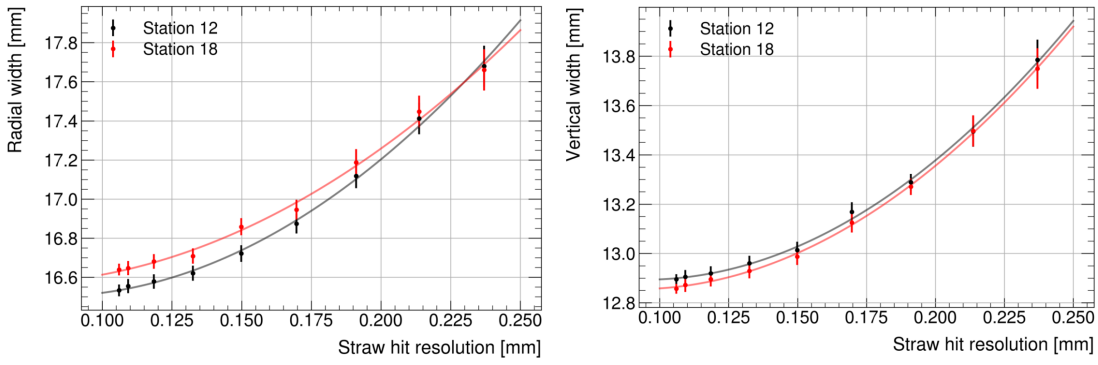

Figure 4.44: The variation of the radial and vertical widths with increased smearing.

order 2 in order to interpolate the required points. For the analyses, the impact of the beam width is a smearing in the vertex resolution, so this needs to be converted. At nominal tracker hit resolution, the vertex resolution is measured to be 2.8 mm vertically and 3.3 mm radially. This can be used to calculate the true width of the beam, which combined with different hit resolutions gives the corresponding vertex resolution. This conversion can be seen in Figure 4.45.

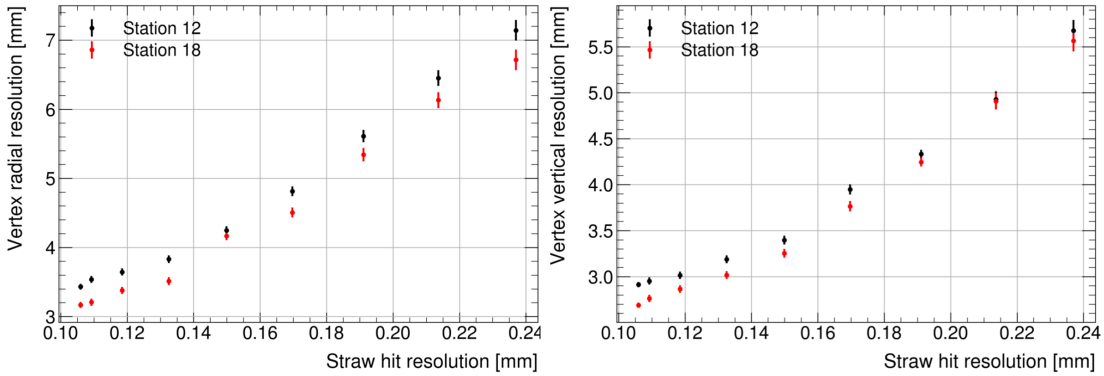

Figure 4.45: Change in the vertex resolution as a function of the straw hit resolution.

These plots are expected to be linear based on previous studies. However, there is a slight non-linearity in the Run 2/3 smeared data, likely due to tracker effects that the MC does not fully account for, for example a larger smearing leading to incorrect left/right assignment of hits. Since the range of resolutions generated for this study is quite large, the variation seen in this plot is high, with the true difference causing a small change near the origin. The points within the smaller 'physical' range are found to be sufficiently linear to fit a line through, as seen in Figure 4.46.

These plots are all used with a difference of 30 μm to give the final beam systematic values given in Tables 4.6 and 4.7.

Station / [mm]	Mean R	Width R	Mean Y	Width Y
Station 12	0.048 ± 0.012	0.119 ± 0.018	0.030 ± 0.010	0.087 ± 0.015
Station 18	0.025 ± 0.010	0.127 ± 0.016	0.024 ± 0.011	0.058 ± 0.012

Table 4.6: Final impact of a 30 μm track resolution change on the beam variables.

This assumption of linear behaviour in a short range of straw hit resolutions must be tested to make sure it doesn't bias the final systematic result. If the relation was perfectly linear, the measurement of the vertex resolution would not impact the systematic, since a difference is used, where both points would shift

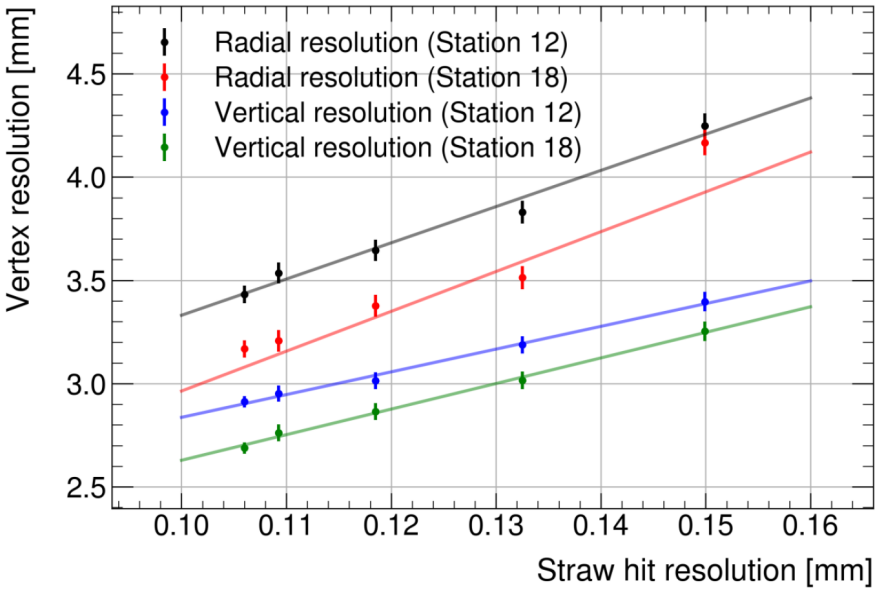

Figure 4.46: Fitting a line in the linear region of Figure 4.45 in order to extract the final systematic.

Station/[mm]	R	Y
Station 12	0.526 ± 0.019	0.330 ± 0.017
Station 18	0.611 ± 0.016	0.371 ± 0.015

Table 4.7: Final impact of a 30 μm track resolution change on the vertex resolution.

in one direction equally. However, if the relation is non-linear, this would not be the case. The conversion factor is varied to within just over a mm to identify any impact.

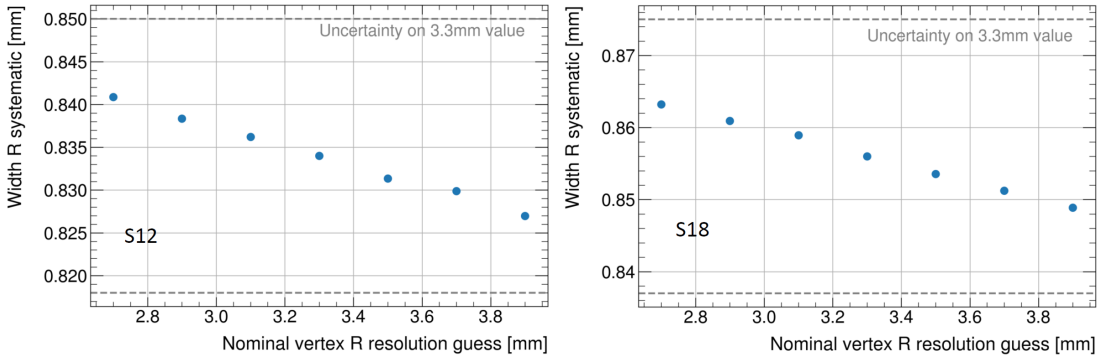

Figure 4.47: Variation of the conversion factor between tracker resolution and vertex resolution, and the resulting values for the R systematic. It is found to be not constant, but the change is negligible compared to the uncertainties.

Figure 4.47 shows that there is a change as the conversion factor is varied, likely from small non-linearity seen in the region of interest. However, this is very small, and within the quoted uncertainty on the systematics. Therefore the approximate linearity assumption is valid.

Finally, the change in resolution is used to estimate the time dependence systematic. A subrun of data is used to find the difference in the plateau resolution at t=0 and t=30 μs. The values were extracted by fitting a decaying exponential plus a constant to the data, as can be seen in Figure 4.37.

The curves in Figures 4.43 and 4.44 are then varied by the difference between t = 0 μs and the plateau to see how much the resolution variation changes the beam distributions. The choice of 0 μs as the start point is done to be conservative, as the difference will be larger. The results can be seen in Table 4.8.

Station/[mm]	Mean R	Width R	Mean Y	Width Y
Station 12	0.011 ± 0.010	0.112 ± 0.013	0.007 ± 0.009	0.077 ± 0.012
Station 18	0.002 ± 0.009	0.047 ± 0.012	0.002 ± 0.008	0.024 ± 0.012

Table 4.8: The impact of the early-time space-charge effect on the tracker resolution.

Station 12, as expected, has a larger effect than Station 18 due to a larger time dependence. Even using 0 μs as a start time, which gives differences much larger than any analysis will be impacted by, the change is small so will not dominate the systematics.

4.4.3 Summary

A subset of the tracker systematics relevant to the ω_a and EDM analyses have been evaluated, including the crosstalk, tracker measurement resolution, and the new time-dependence in the tracker resolution, both of which are found to be small compared to other systematics as seen in the ω_a Run 1 paper [52].

Two methods are compared for crosstalk, one comparing the beam distributions by injecting crosstalk directly and the other using the different crosstalk removal methods to give a purely data-driven value. Both methods give very small contributions to the overall tracker uncertainty, so the most conservative value is taken.

The impact of the resolution is quantified by artificially applying a Gaussian positional smearing to real data and retracking it. This is then translated to the impact on the vertex resolution, which is also found to be small compared to other systematics.

Finally, the impact of a time-dependent resolution on the beam distributions is investigated. Since this effect is more severe at earlier times, values were calculated for both 0 and 30 μs into a fill, with both found to be small but with a slightly larger value for Station 12 than Station 18. All three of these systematics have been included in the final tracker systematics which contribute directly to the Run 2/3 ω_a beam dynamics corrections detailed in Section 3.14. Some also contribute to the EDM analysis, but their impact is evaluated slightly differently in Chapter 7 as the beam distribution variation is not sufficient to quantify it.

Chapter 5

Development of a Monte Carlo simulation of the EDM

The target EDM limit from the Fermilab Muon $g-2$ experiment is an improvement of two orders of magnitude on the current world-leading measurement, to reach 10^{-21} e·cm. This goal is calculated by scaling the expected statistical uncertainty of the new experiment to the BNL dataset size of about 10 million tracks [1]. This gives a statistically-limited prediction of the sensitivity, but to fully understand the potential reach, consideration of the impact of systematic uncertainties is also needed. Additionally, knowing which of these will limit the final result is useful to guide experimental decisions, particularly for planning any studies which need to be performed before the end of experimental operations.

Ideally sources of uncertainty should be studied independently of each other, which motivates the use of Monte Carlo (MC) data. However, this requires the generation of a dataset similar in size to the final combined experimental dataset at Fermilab, which is not feasible with GM2RINGSIM, as it would take too long. GM2RINGSIM includes a lot of complexity that isn't required for the estimate, which also makes it more difficult to disentangle individual effects. Therefore, it is useful to develop a new MC with a focus on simplicity and speed for this purpose.

This new MC, implemented by the author, is designed to be simple while still giving accurate results, building on the muon decay physics to add complexity. Initially, the ring geometry is removed and replaced with a simpler fixed momentum direction coordinate system, with all muon beam motion and spread neglected in favour of a point-like 'perfect' beam at the magic momentum. Tracking is also assumed perfect, with a 1:1 correspondence between decays and detected positrons. This is particularly helpful for increasing the statistics of the simulation, as in both GM2RINGSIM and data most positrons miss the trackers. The base MC

therefore only needs to include muon decay physics, the $g-2$ precession, and the EDM oscillation. Additional complexity is added in a modular fashion on top of this to study the impact of the radial field and tracker acceptance.

5.1 Simulating muon decay

The muon decay is first simulated in the muon rest frame (MRF), using the assumption that $m_e << m_\mu$ and neglecting any kinematics not in the plane of the muon momentum. With this approximation, the 2D differential decay distribution for a muon decaying into a positron has the form [85]:

$$\frac{d^2\Gamma}{dx\, d\cos\theta} \sim x^2[(3-2x) + P_\mu \cos\theta(1-2x)] \qquad x = \frac{E_e}{E_e^{max}} \tag{5.1}$$

, where x is the energy fraction E_e/E_e^{max}, P_μ is the muon polarisation and θ is the angle between the muon spin direction and the momentum direction. The parity-violating correlation between energy and decay angle can be seen by this being a function of both variables.

Individual distributions are extracted by integrating over angle and energy:

$$\frac{d\Gamma}{dx} \sim (3x^2 - 2x^3), \qquad \frac{d\Gamma}{d\cos\theta} \sim 1 + \frac{1}{3} P_\mu \cos\theta. \tag{5.2}$$

The full derivations of these expressions can be found in Appendix A.1. A 'decay' class is defined with attributes for positron energy, momentum and decay angle. The energy is sampled from the integrated energy distribution in Equation 5.2, then the decay angle is sampled as a function of energy to accurately simulate the correlation. This is run for a large number of decays and compared to the expected functional forms, shown in Figure 5.1. Additionally, a 2D histogram of energy and decay angle is plotted in Figure 5.2 to check that the correlation has been correctly simulated.

Since positrons in data are not measured in the rest frame but rather the lab frame, each generated positron is boosted along the direction of momentum. For the sake of simplicity, it is assumed that all muons in the beam have the same momentum, set at the magic momentum of 3.094 GeV, so all positrons are boosted in the same way. This modifies the distributions, reducing the angular spread into a tight cone around the direction of momentum. This behaviour matches what is seen in experimental data. The post-boost distributions are plotted in Figure 5.3.

Finally, a time t is defined and sampled from a boosted exponential distribution

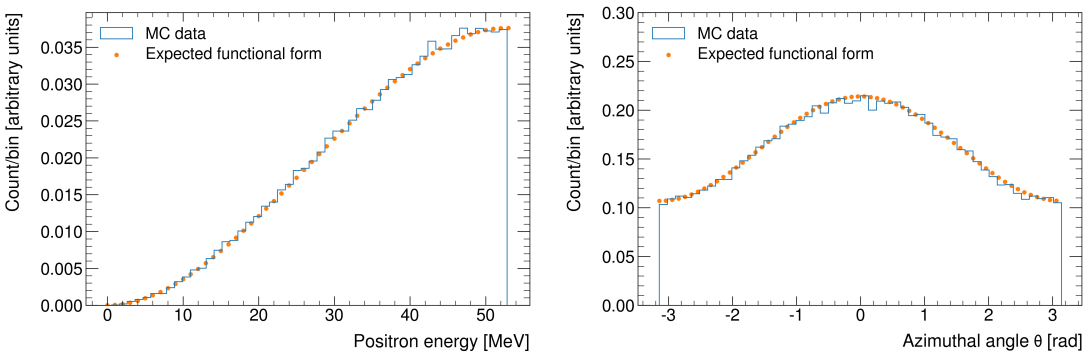

Figure 5.1: Energy and angular distributions of 1×10^5 MC-generated decays, overlaid with the expected functional distributions. Runtime is very fast, approximately 20s.

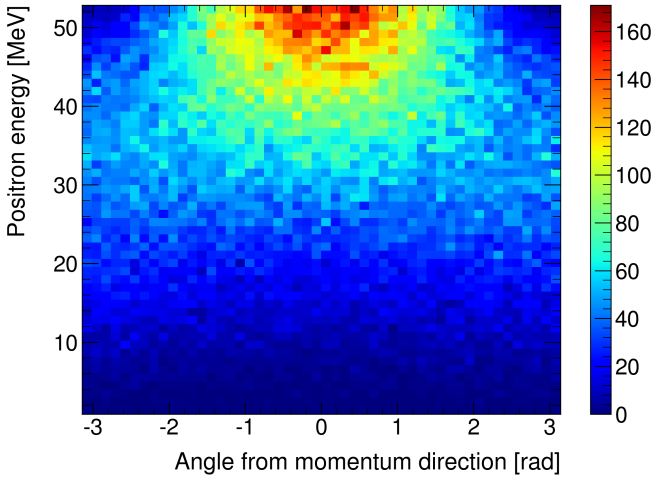

Figure 5.2: 2D histogram of MC data, confirming that higher energy positrons are indeed being emitted preferentially at low angle. This being correct is essential for the $g-2$ oscillation.

to include the muon decay lifetime. These combined aspects now fully describe a simple muon decay, with events generated very fast, taking approximately 20s to generate 10^5 events.

The oscillation of the spin coplanar with the momentum is added, corresponding to the $g-2$ precession. The frequency is input as the value measured at BNL since at the time Run 1 had not been unblinded. Then, an energy cut is applied, removing all positrons with energy < 2000 MeV, since for highly boosted positrons the energy and momentum are roughly equivalent. The remaining positrons are counted and binned in time. This generates a wiggle plot successfully, indicating that the base physics of the simulation is behaving as expected. The only uncer-

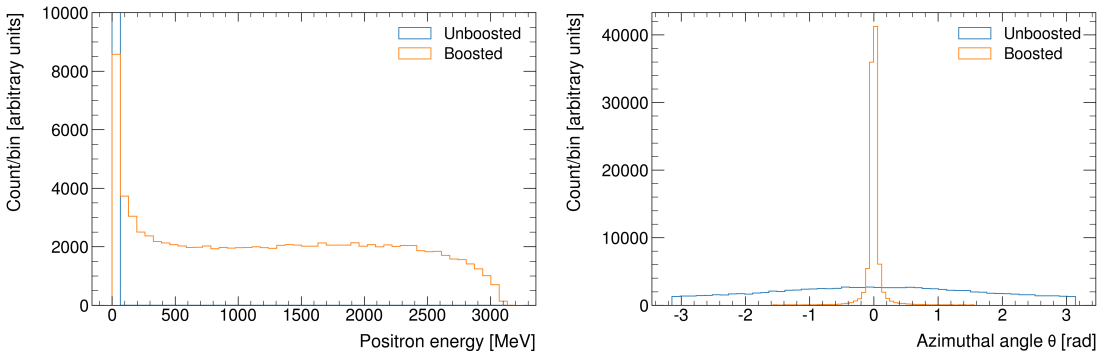

Figure 5.3: Comparisons of energy and angular distributions in the muon rest frame (unboosted) and the lab frame (boosted). The angular distribution becomes sharply peaked around 0, the momentum direction, whereas the energy increases.

tainties present at this point are statistical, so an error bar of $\sqrt{n}$ is calculated for each bin. The wiggle is finally fitted with the 5-parameter fit function from Equation 2.9 to check that the input frequency is extracted. An example is shown in Figure 5.4, where the precession frequency has been set to a large test value.

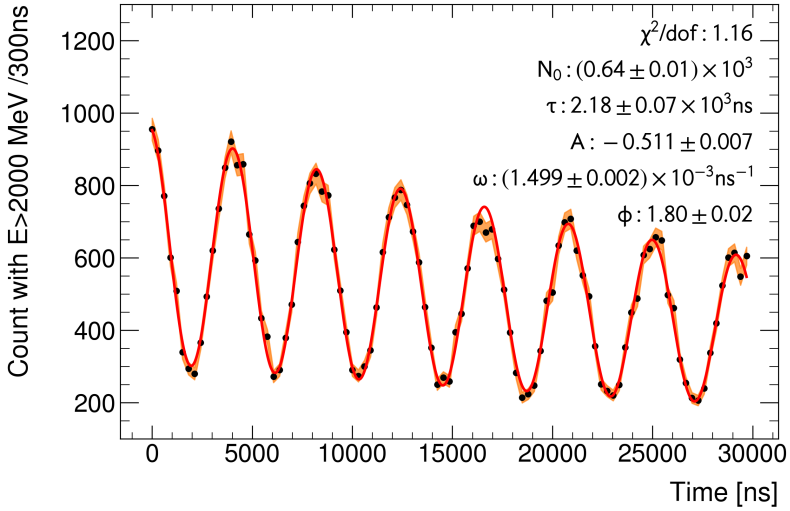

Figure 5.4: A fitted test wiggle plot from MC data. The errors, plotted as the coloured band, are largest where the function is changing rapidly, leading to some variation in fit χ^2. Most importantly, the fit ω_a matches the input of $1.5 \times 10^{-3}\text{ns}^{-1}$.

As a non-zero EDM causes an oscillation of the 2D precession plane, an oscillation in the average vertical angle of the decay positrons is the expected EDM signature. Therefore, a third spatial dimension is required if the EDM is to be studied. The differential decay in Equation 5.1 does not include the vertical angle ϕ, since the decay is symmetrical in this direction. However, extracting total decay

parameters like the muon lifetime requires integrating over solid angle, introducing a $\sin(\phi)d\phi$ term. ϕ is generally defined from 0 to π, but for the EDM it makes more sense to redefine this as -$\pi/2$ to $\pi/2$, shifting the phase by $\pi/2$ and making the expected final distribution a cosine. This is demonstrated in Figure 5.5. This spread is sampled to define the ϕ angle of each decay. Similarly to before, the vertical angle is also reduced by the boost.

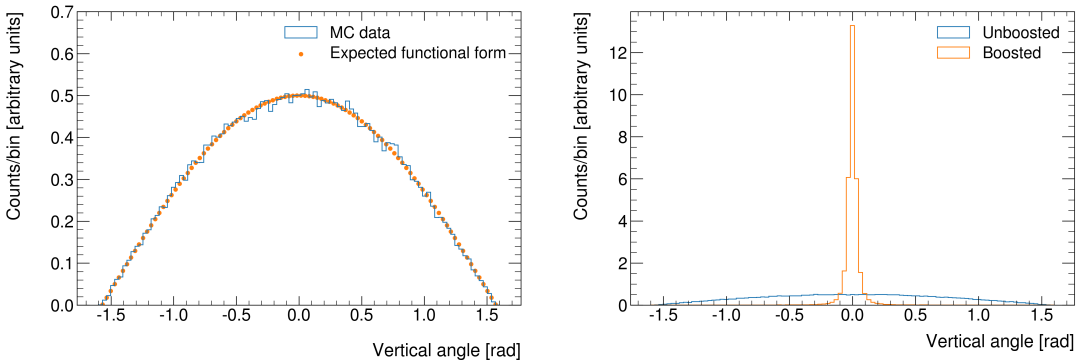

Figure 5.5: Distributions for the vertical angle before and after the boost. Similarly to the azimuthal angle, the distribution becomes sharply peaked after the boost.

Finally, the EDM signal needs to be injected. This is a vertical angle oscillation $\pi/2$ out of phase with the $g-2$ oscillation, with an amplitude that depends on d_μ. The tilt angle is defined in the lab frame using Equation 2.25.

As a test, a large EDM equivalent to 100× the BNL limit is injected. Again, errors here are purely statistical, so the vertical angular distribution is fitted with a Gaussian in each time bin to extract the mean and width. The average vertical angle is then plotted against time modulo the $g-2$ period. In real data, this is done mainly to destructively interfere any oscillations slower than the $g-2$ frequency, which are not present here, but it is still useful as it improves the statistics per bin. The data from Figure 5.4 is fitted to extract the frequency and phase with a function of the form:

$$A = A_{EDM}\sin(\omega_a t + \phi_{EDM}). \tag{5.3}$$

The phase is fixed based on the wiggle fit, allowing only the amplitude A_{EDM} to vary, as shown in Figure 5.6. This amplitude is compared to the input to confirm the EDM has been injected correctly.

The MC now has all the machinery in place to perform the required studies to determine the sensitivity in the case of no signal. The final result will be a limit set on the EDM, so the particular method of calculating the limit must be chosen for both these MC results and the experiment as a whole.

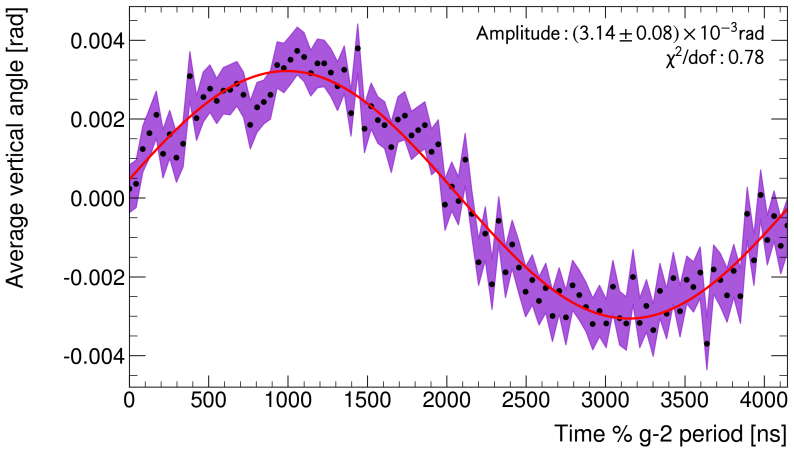

Figure 5.6: The average vertical angle binned modulo the $g-2$ period, with an error band, fitted to extract the amplitude. More variation is seen here than in the $g-2$ wiggle due to the average amplitude being similar in magnitude to the vertical angle spread.

5.2 Defining the EDM limit

The calculation of the EDM limit depends on two fit parameters, the final extracted precession plane tilt and its uncertainty. Assuming the distribution of d_μ values is Gaussian, these values are used as PDF parameters and a statistical method used to set the limit [82]. This assumes that no EDM is seen, as a non-zero EDM would count as a discovery at the measured level so limits are not appropriate.

The BNL EDM limit was set using a Gaussian limit-setting method [1], but modern experiments use more sophisticated methods. Methods in the field were investigated to determine which would be most suitable for the FNAL result. Of particular importance is how these methods behave near physical boundaries, whether they give sensible results for large fluctuations, and how easily the result can be compared to or combined with other results.

Searches for other EDMs were considered as a starting point. The limits set and methods used by previous experimental EDM searches are shown in Table 5.2. Theoretical predictions which have not yet been tested experimentally, for example the neutrino and tau lepton EDMs, are not included as the methods used are less relevant. The Fermilab experiment's Run 1 EDM preliminary result [12] is also not included, as the method used was directly informed by this work.

A variety of different methods have been used in the past, with the choice likely based on the age of the experiment and the Particle Data Group (PDG) recommendations at the time. The merits of each are considered separately.

Particle	EDM Limit	Method	Reference
e	$\lvert d_e \rvert < 1.1 \times 10^{-29}$	Feldman-Cousins	[11]
μ	$\lvert d_\mu \rvert < 1.9 \times 10^{-19}$	One-sided Gaussian	[1]
n	$\lvert d_n \rvert < 1.8 \times 10^{-26}$	Two-sided Gaussian	[28]
p	$\lvert d_p \rvert < 8.0 \times 10^{-25}$	Two-sided Gaussian	[27]

Table 5.1: EDM limits set on other particles and the methods used to set those limits.

Gaussian methods

There are several variations of classic Gaussian methods which can be used, the simplest of which is a two-sided confidence interval. This takes the Gaussian PDF and integrates outwards from the mean until the desired probability content is included. The limits are defined as the boundaries of this region. This has the advantage of being simple to implement and interpret, and always has correct coverage, which means it always contains the full desired probability content. However, it does not consider physical boundaries, like a negative EDM limit for a positive central value, and also may exclude the boundary, which in this case is an EDM measurement of zero.

One variation is to define a one-sided Gaussian interval, integrating out in only one direction from the physical boundary until the desired confidence level is reached. The limits are the boundary and the upper end point of the integral. This helps avoid unphysical values, but can lead to undercoverage and non-physical results from large fluctuations beyond the boundary. For absolute limits, a 'folded' Gaussian PDF may be used, with results beyond the physical boundary reflected into the physical region. This technique was employed by the electron EDM analysis [11].

The existence of one and two sided intervals leads to ambiguity about which method to use, with some analyses using both and swapping at some predetermined sigma. This is suboptimal as it leads to undercoverage for some values and overcoverage for others, a problem known as 'flip-flopping'.

Feldman-Cousins

The Feldman-Cousins (FC) method was designed to fix the issues of flip-flopping in a mathematically rigorous way by removing the choice of where to switch. It defines a new variable R:

$$R = \frac{P(x|\mu)}{P(x|\mu_{best})} \tag{5.4}$$

, where x is the measured value, μ is the true value, and μ_{best} is the best fit

value given the data and the allowed values for μ. Parameter regions are added to the confidence band in order of decreasing R [86]. This effectively renormalises the PDF, making it large for likely values in the physical region, and small in unphysical or distant regions. The integration for the limit then spreads from the most likely value outwards until a boundary is hit. The other limit continues to change until the region has correct coverage within the physical region. This allows the method to deal well with both boundaries and large fluctuations.

Comparison of methods

The methods mentioned above were implemented with their confidence bands plotted and compared as shown in Figure 5.7.

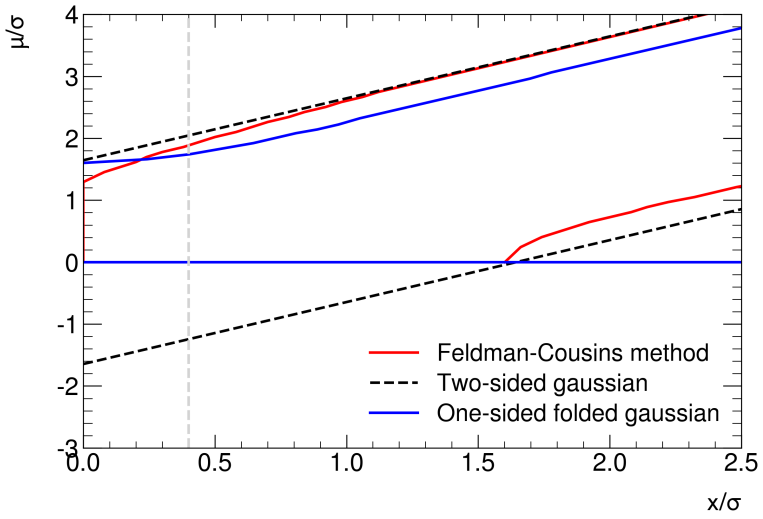

Figure 5.7: Confidence bands for the best versions of the 3 methods normalised by σ. Limits are calculated for measured value x as the points where the vertical line intersect the bands.

All methods gave broadly similar upper bounds for the limit, but varied in the lower bound. One-sided methods were decided to be less useful for the muon EDM, due to the influence of large fluctuations which set a high limit. This would not bound the available parameter space strongly.

Between the two-sided methods, Feldman-Cousins deals best with the physical boundary, especially if combined with the folded Gaussian method. This is a sensible procedure to follow as only the absolute value of the EDM limit is needed. Additionally, it has the correct coverage near a boundary, making FC the best method for the muon EDM. All limits quoted in the subsequent sections of this report are calculated using the Feldman-Cousins method. The MC results are

quoted at a 90 % confidence level, whereas the main data analysis uses 95 % to more closely mirror the electron EDM results.

5.3 Consideration of systematic uncertainties

5.3.1 The impact of a radial field

A radial field (B_r) component causes a tilt in the muon precession plane which mimics the signal of a non-zero EDM. This makes it likely to be a dominant systematic for the analysis unless it is measured and controlled to a high degree of precision. This, combined with the potential for improvement of the measurement, motivates it as the first effect to be studied. B_r can be treated as a background and subtracted off, but the uncertainty δB_r is most important, as any EDM causing a tilt smaller than this field uncertainty cannot be measured as it would lie within the radial field uncertainty. Therefore, this defines a 'lower bound' for the limits the experiment could set. It is useful to study where this lower bound lies for different values of δB_r, to gauge where it starts to become significant relative to the statistical uncertainty. From measurements made during commissioning, the uncertainty on the radial field if no new measurements are made is estimated to be approximately 50 ppm [88]. New measurements, either directly with field probes or indirectly using beam dynamics, could improve this to a sub-10 ppm level, potentially as good as 1 ppm, so these values are used to compare the final achievable limits.

Implementation and testing

Considering vertical and radial B field components, the tilt angle can be calculated using their ratio. For the field configuration of the $g-2$ experiment, the majority of the field is vertical, with the radial component being at most 100 ppm. This means the small angle approximation can be used to simplify the expression.

$$\delta = \arctan\left(\frac{B_r}{B_y}\right) \sim \frac{B_r}{B_y} \sim \frac{B_r}{B_{tot}}. \tag{5.5}$$

The radial direction must now be defined in the fixed coordinate system of the MC, which does not include the cyclotron motion of the muon momentum vector. As the true momentum direction rotates with time, and the radial field acts perpendicular to this, a rotating 'radial' direction is required. Both direction and tilt were implemented into the simulation and run for different input values

to check the output was as expected.

Two methods can be used to study the effect on the EDM limits. First, the field can be injected as a variable with a Gaussian distribution, then run through the full EDM analysis, including the known radial field uncertainty in the total error. This will give numerical limit values based on our knowledge of the radial field. The second option uses the maximum tilt caused by δB_r to determine the smallest possible limit measurable. This allows an injected radial field to act as a proxy for the uncertainty, and is slightly faster computationally. As a result, this second method was preferred, with a comparison performed to confirm the results from both methods match.

It is possible to make rough predictions for the lower bound of the limit using the radial field tilt angle and assuming a perfect measurement of the EDM at zero amplitude with negligible error. The statistical uncertainty can then be included by setting the uncertainty of this ideal measurement equal to the statistical uncertainty for a particular dataset size, with central value from the radial field tilt. This gives a lower bound as the measured amplitude is likely to skew above 0 due to the folded Gaussian used in limit setting.

To generate data, the full simulation for 10 million tracks was run 100 times for radial field values a factor of 10 apart between 0.1 ppm and 10,000 ppm. These numbers were chosen to span the full possible range of uncertainties while also including a very low value and very high value. Each dataset was run independently through the EDM analysis and the limits calculated using the Feldman-Cousins method. To produce a single result for each radial field value, the mean and standard deviation were calculated and plotted along with the prediction curves.

To facilitate comparisons for different numbers of tracks, datasets were combined up to a size of 1 billion tracks, chosen based on the number of tracks seen in Run 1 compared to the number of positrons seen in the calorimeters. Running with different numbers of tracks allows comparison with the purely statistical estimate, and gives numerical values for the minimum radial field uncertainty needed to achieve target limits. It also allows comparisons to be made between radial field and statistical uncertainty to determine which is more important.

Radial field results

The prediction curves for 10 million, 100 million, 1 billion and 10 billion tracks are shown in Figure 5.8. These include the statistical uncertainty, which scales as the square root of the number of tracks, and the radial field uncertainty, which is varied. These are combined in quadrature and used to calculate the limit.

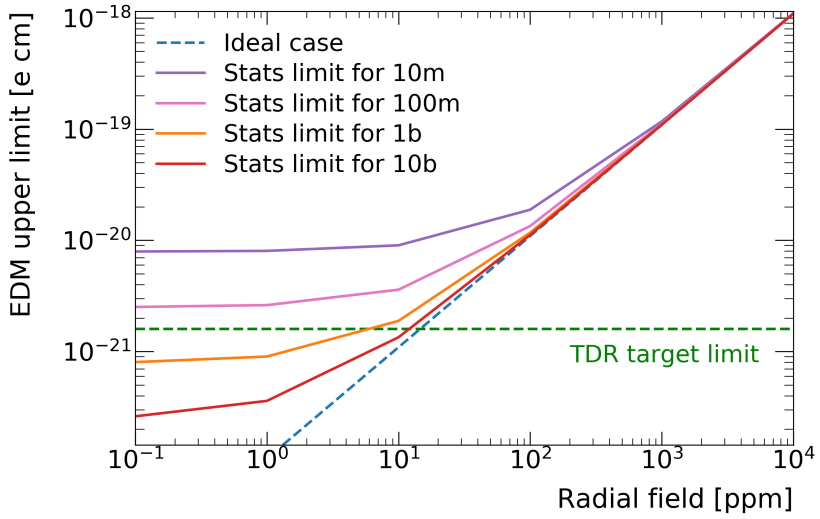

Figure 5.8: Prediction curves for the best possible limits as a function of the radial field uncertainty. These show the combined effect of the radial field uncertainty and statistical limits.

For large radial field uncertainty, the statistical error has a negligible effect, but as the radial field is reduced, the statistical error starts to dominate. Their relative sizes between 1 ppm and 10 ppm suggest that both will impact the final EDM limit.

The results of running for 10 million tracks are shown in Figure 5.9. The simulated data follows the predicted curves with slight upward bias from the folded Gaussian. This shows that the TDR limit is not obtainable with 10 million tracks, with a likely limit of 1.0×10^{-20} e·cm for $\delta B_r = 10$ ppm. This plot, despite having similar statistics to BNL, is below the BNL limit at 10 ppm, while it should match as the BNL result was statistically limited. However, the simulation does not include the tracker acceptance, which reduces the visible amplitude of the EDM oscillation. This would bring the line up by a factor of about 10, closer to the expected limit. Combining into a dataset of 1 billion tracks gives the results in Figure 5.10.

The increase in statistics now allows the target limit to be achieved, provided the radial field is known to 1 ppm for 1 billion tracks, or 6 ppm for 10 billion tracks. From this it can be concluded that the radial field is likely to be a large limiting factor, as the best possible estimate without extra measurement would have an uncertainty of around 10 ppm. With this level of precision and 1 billion tracks, the final EDM limit set in this simplified case would be $|d_\mu| < 2.9 \times 10^{-21} \mathrm{e} \cdot \mathrm{cm}$. Other systematics like the acceptance uncertainty are expected to further increase

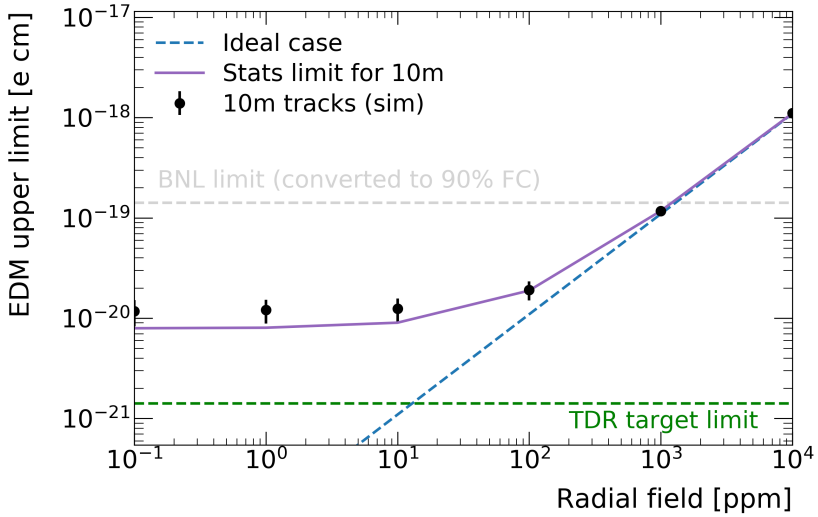

Figure 5.9: The limits set for the ideal case with only statistical errors and radial field uncertainty. 10 million tracks is similar in size to the combined BNL datasets.

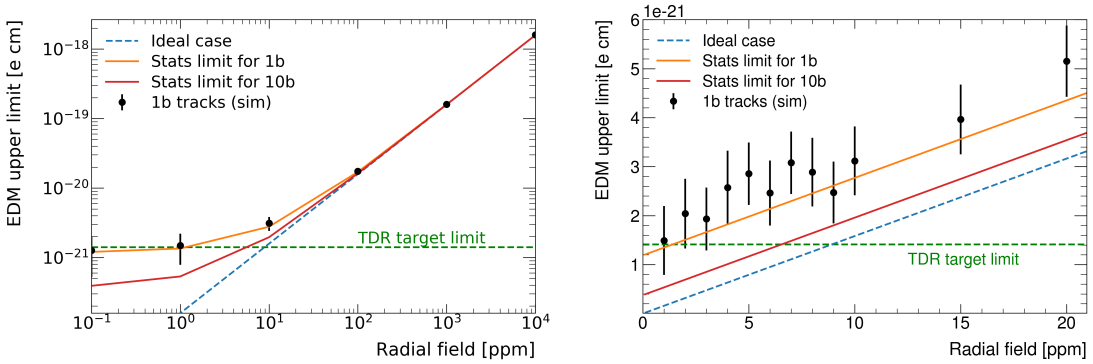

Figure 5.10: The limits set for combined datasets of 1 billion and 10 billion tracks. Additional datasets were run in the range 1 ppm - 20 ppm to closer investigate the region near the target limit. The shape is approximately linear.

this value.

To compare the relative impact of the statistics and the radial field, the limit is calculated for different sized datasets between 1 billion and 10 billion tracks and plotted in Figure 5.11. It can be concluded that the radial field has a larger impact than the statistical uncertainty, motivating a dedicated measurement of the radial field to get this systematic under control. This is detailed in Chapter 6.

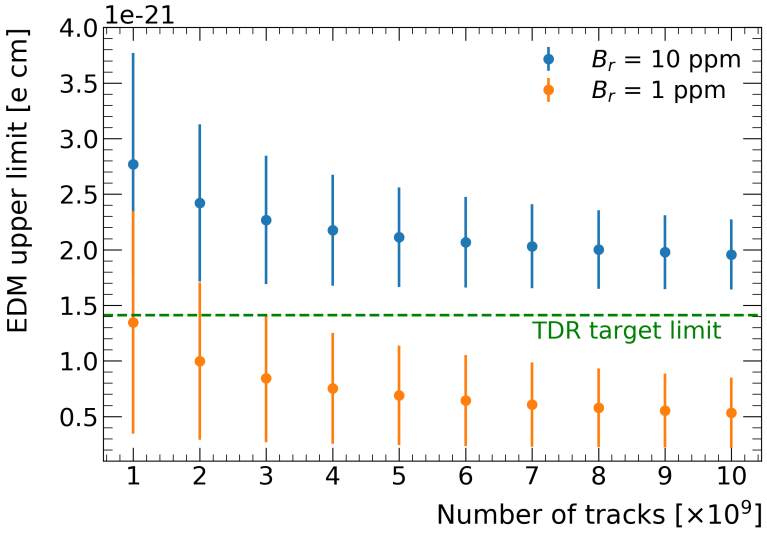

Figure 5.11: Limits set for different numbers of tracks between 1b and 10b, for the two extremes of δB_r. This shows that 10% of the data at 1 ppm gives better results than 100% at 10 ppm, so the radial field uncertainty dominates.

5.3.2 The impact of tracker acceptance

Only around 10-20 % of decay positrons hit the trackers. Whether a positron is detected depends on a large number of factors, including the angles of the decay, the position of the decay vertex, whether the positron travels though material and its momentum, which determines the curvature of the track and therefore how easy it is to reconstruct.

In the case of vertical acceptance, the physical size of the detector removes positrons with a large vertical displacement once they reach the trackers, which corresponds to a large decay angle at the vertex. This reduces the EDM amplitude, since removing the largest angles will lower the average. Similarly, low momentum positrons have highly curved trajectories so are less likely to meet the minimum 5-hit threshold in order to be tracked. High momentum positrons travel further before hitting the trackers so are more likely to scatter off other components before reaching the trackers. Overall, these effects make it harder to see an EDM, making the statistical error a more significant fraction of the amplitude and raising the statistical 'lower bound' for EDM limits.

Vertical acceptance

The simplest way to apply a vertical acceptance in the simultion is by removing all positrons above a certain vertical angle. Assuming a straw height of 10 cm, this angular cutoff would be $\phi > 0.05$ rad for a path length of 1 m and $\phi > 0.01$ rad

for 5 m. These two extremes were investigated to roughly bound the impact. The ratio of amplitudes before and after the cut is applied is compared to quantify the size of the reduction effect. A large EDM amplitude is injected for the purposes of the study.

The detector acceptance also depends on the position of the decaying muon in the beam, which requires simulating the beam distributions in the MC and calculating the final vertical displacement at the tracker. Based on data, the vertical distribution is modelled as Gaussian, centered at 0 mm, with width 12.9 mm. The radial distribution is more complex as it is momentum dependent, but as it does not contribute to the vertical acceptance, it is modelled as a skewed Gaussian using the stats.skewnorm function from SciPy [87], which defines the functional form in terms of a standard Gaussian PDF and CDF:

$$\text{skewnorm.pdf}(x, \mu, \sigma, a) = 2 \times \text{norm.pdf}(x) \times \text{norm.cdf}(ax) \tag{5.6}$$

, with mean μ, standard deviation σ and level of skew a defined based on a fit to the tracker-measured beam distrubution. A comparison can be seen in Figure 5.12.

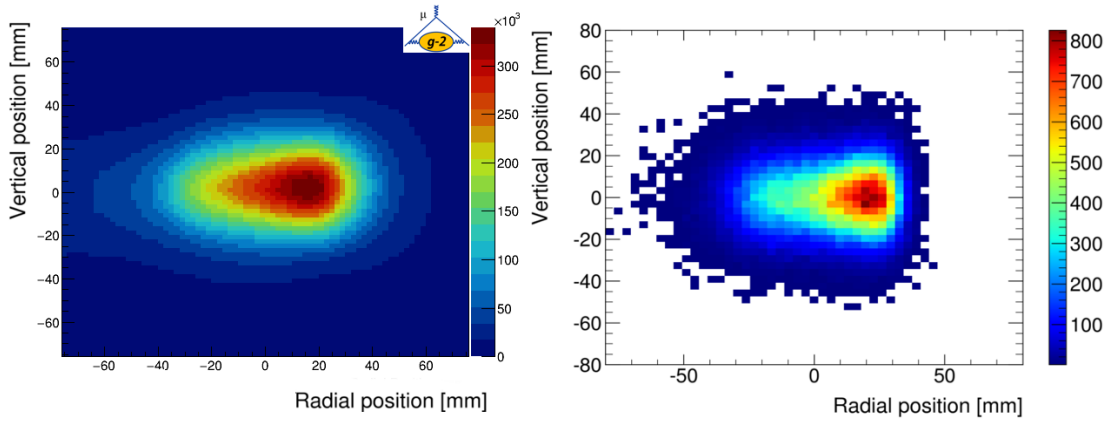

Figure 5.12: Comparison of the beam distribution in Run 1 data (Adapted from a plot by the $g-2$ tracker group) (left), to the MC beam distribution (right). The rough shape is the same, with the key parameters defined using data.

The final displacement of the positron at the tracker determines whether it is accepted, which requires the distance travelled between the vertex and the tracker to be calculated. This length is momentum dependent, since both decay position and radius of curvature determine the final position the positron has when it hits the tracker. In data, the distance peaks at 2.2 m away from the front of the tracker. To initially simulate this effect, the distance travelled by the positron is fixed as 2.2 m and combined with the displacement from beam offset and decay angles.

A cut is then applied on the tracker position, requiring less than 50 mm vertical displacement to count as having hit the tracker. This gives a decent initial estimate for the effect, but more detailed ring geometry is needed to properly simulate the impact of the distance travelled.

A separate toy MC was created to better quantify this, with basic ring geometry and a single tracker. Muons are placed uniformly around the ring, the momentum is defined, and the circular path of positrons decaying at each position calculated. This path is compared to the tracker coordinates to determine whether the track hits the tracker and is detected. The distance travelled between decay and tracker is calculated and stored along with the positron momentum. This allows the relation between the two to be studied in a similar idealised environment to the rest of the simulation. The ring MC being separate allows it to work both as a standalone simulation or as an add-on module to the main MC, though running as an add-on results in similar timing issues to GM2RINGSIM, so is not an ideal operation mode. The tracker is first modelled as the 1D front plane of the first tracker in the ring, and then as a 2D box, as shown in Figure 5.13. This allows hits to be registered for positrons entering through the sides of the tracker. The results are propagated through the analysis and the impact compared with the momentum-distance travelled relation results seen in tracker data.

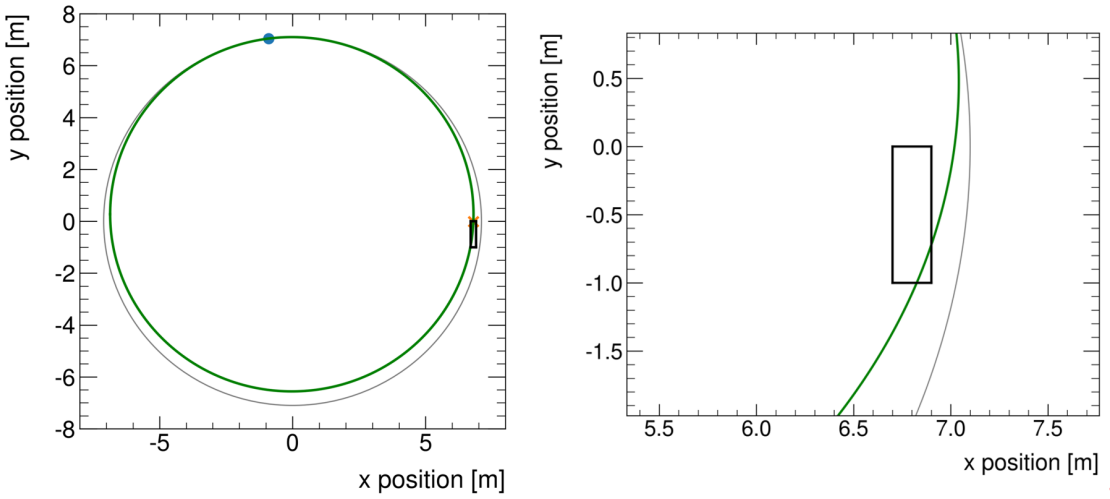

Figure 5.13: Plots showing the simple ring geometry implemented in the MC for the box tracker method. The plot on the left shows a high-momentum positron hitting the front face of the tracker and therefore being detected, while the right plot shows a zoomed in view of the tracker region being hit from the side.

The difference in distance travelled between the simple simulation and GM2RINGSIM is shown in Figure 5.14, with GM2RINGSIM's evaluation including beam width

and motion whereas the simple MC does not. Reasonably good agreement is seen between the two in the central momentum region important for the EDM analysis. A difference is seen in the low and high momentum behaviour, which arises due to the simple MC assuming every hit can be detected and tracked, which is not the case with real data. High momentum tracks in the simple MC also come from a large distance upstream around the ring, whereas in reality the azimuthal acceptance of the tracker is not this wide. Finally, the simple ring tracker geometry is a box roughly the size of the full tracker station. The real tracker is made of 8 modules that curve to follow the ring, so the acceptance will be different due to this as well.

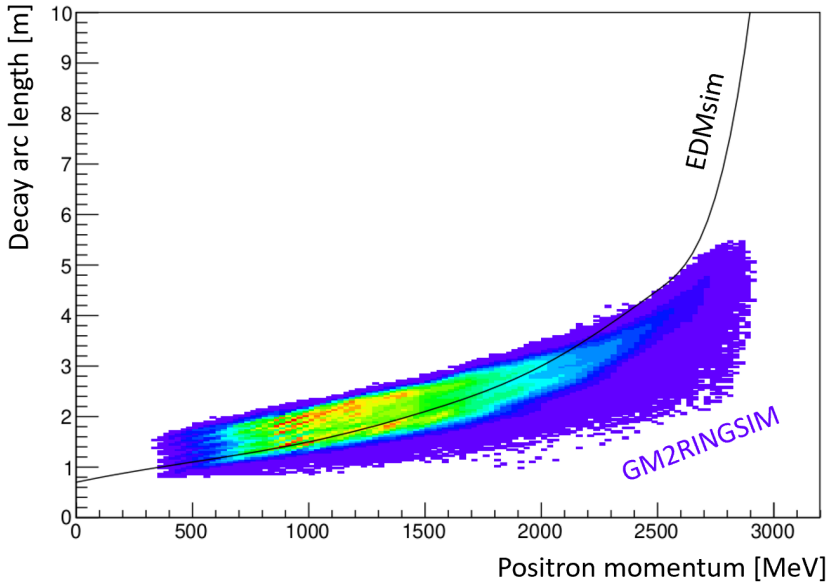

Figure 5.14: A comparison of the arc length as a function of momentum seen in the simple EDMsim and GM2RINGSIM. The EDMsim line is extracted by fitting a polynomial to the momentum distribution for easier comparison.

Momentum acceptance

As many acceptance effects are momentum-dependent, a simpler way to implement their effect is to cut on momentum directly. This removes the complication with correlations and ring-related geometric setups while giving a better approximation to the true tracker acceptance.

As some momentum dependence is already present from the vertical acceptance cuts, a comparison between this MC and GM2RINGSIM is made using the number of positrons removed at different vertical positions, which defines an efficiency and prevents double-applying any effects that have already been accounted for. The

Method	ϕ cut	Beam + ϕ cut	Plane tracker	Box tracker	Data
Ratio	7.1 ± 2.3	12.0 ± 6.1	25.0 ± 7.3	17.6 ± 6.8	22.0 ± 7.1

Table 5.2: Amplitude variations seen using the different acceptance cut methods. The plane tracker refers to the case where hits are only kept if they pass through a 1D tracker plane representing the front of the first tracker, and the box case adds a depth, allowing hits to be counted coming in from the side as well as the front.

missing acceptance shape is applied using a probability-based cut as a function of momentum to make the simple MC match with GM2RINGSIM. Additionally, a hard cut of P > 300 MeV is applied to remove all low momentum positrons, as these are not able to be tracked in data due to a low number of tracker hits. Combined, these two effects lead to good agreement between the GM2RINGSIM and simple MC momentum distributions.

Acceptance results

First, the vertical acceptance is considered. As the acceptance is known to reduce the EDM amplitude, the impact of each set of cuts is considered by finding the ratio between the inital amplitude with no acceptance cut, to the fitted amplitude with an acceptance cut. An example of the difference is shown in Figure 5.15.

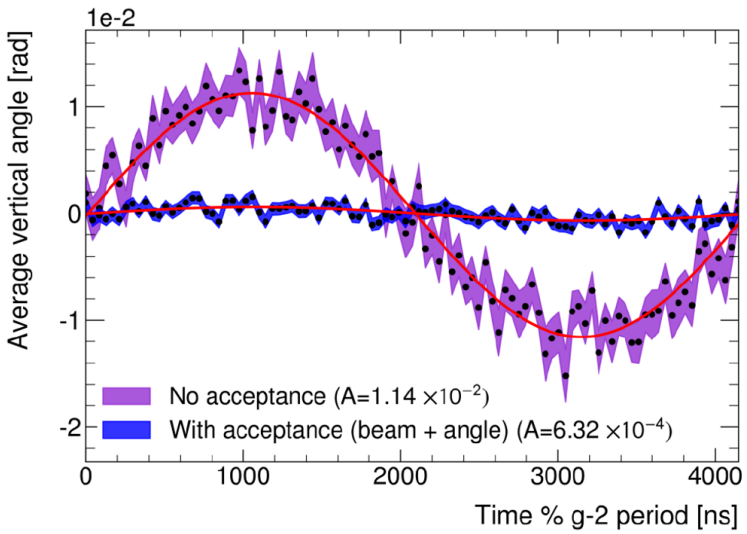

Figure 5.15: Comparison of the visible EDM amplitude with and without an acceptance cut applied.

A comparison of the ratio for the different methods can be seen in Table 5.2.

Compared to the simple cut, adding in the beam spread decreases the acceptance, as more positrons from the beam edges are removed. The plane tracker removes more positrons than the box case, as a larger fraction of positrons are missed,

however the two ratio values are similar within errors. The plane tracker only measures longer decay distances compared to data, as shorter distance positrons predominantly come through the side of the tracker, whereas the box tracker overestimates as the very high momentum positrons are unlikely to reach the tracker without scattering. Therefore, the expected ratio for data lies between the plane and box tracker Monte Carlo results. Comparing the fraction of positrons detected as a function of vertical angle, both the box tracker and data-informed methods match GM2RINGSIM's results well, as shown in Figure 5.16.

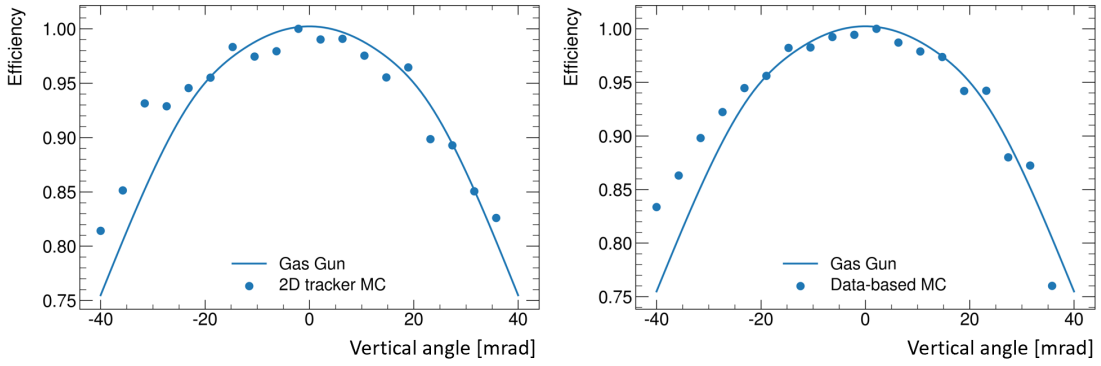

Figure 5.16: Comparison of the efficiency, defined as the ratio of positrons which pass the cut vs the total number of positrons, normalised to peak at 1. The data-based MC uses the arc length relation from fitting data, whereas the 2D tracker MC is pure MC. Both the 2D tracker and data-based arc length relations give good agreement.

The ratio between amplitudes also depends on the size of the oscillation injected, as higher angles result in larger fractions being removed. This effect is only seen in this MC as the injected EDM signal is unphysically large. For a more accurate ratio estimate, no EDM signal is injected. The ratio for different input amplitudes was studied and found to be roughly linear, as shown in Figure 5.17. For an input tilt of similar size to the radial field uncertainty, the ratio remains roughly constant.

The total numerical factor for the vertical acceptance ratio for EDM limits is therefore a factor of 3.6 $\pm$ 2.7. The limits in the previous section should be multiplied by this factor to give the new limit. Using the 10 ppm uncertainty, this gives a new limit of order 1.0 $\times 10^{-20}/:e \cdot cm$ for 1 billion tracks.

Considering the momentum acceptance, the efficiency difference between the vertical acceptance cut and GM2RINGSIM's results are shown in Figure 5.18.

There is a small region for mid-momentum values where they match, however there are large differences at the extremes of the momentum scale. Applying the

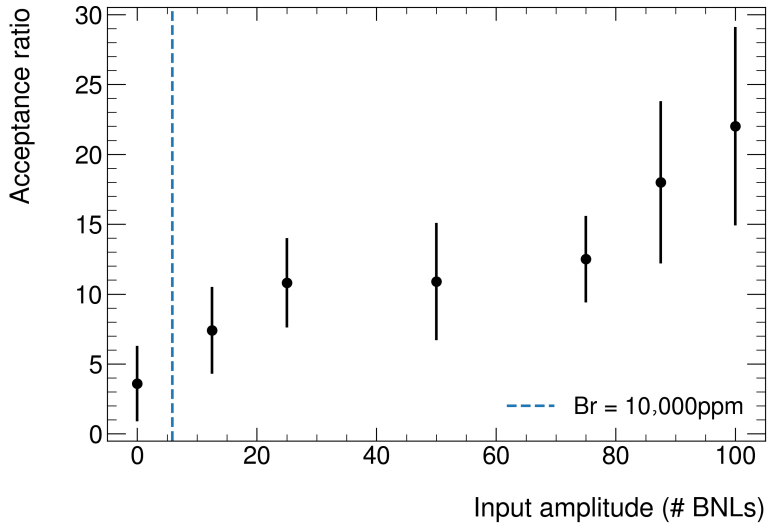

Figure 5.17: Plot illustrating how the acceptance reduction ratio changes with EDM amplitude in the MC. For the radial field studies, the largest input is 10,000 ppm, corresponding to $\sim$ 6 x the BNL dataset, marked with a dashed line. The variation between 0 and this amplitude is small, so the ratio can be assumed roughly constant across the scanned range and in data.

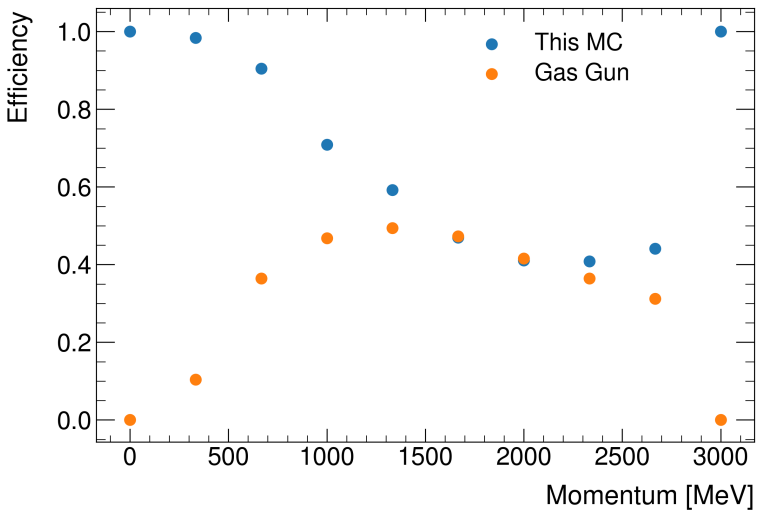

Figure 5.18: A comparison of the cut efficiency as a function of momentum between this MC and GM2RINGSIM's Gas Gun simulation.

probability cut to make the MC match data and comparing distributions with GM2RINGSIM now gives similar results by construction. The ratio was found to decrease to 3.5 ± 2.9.

While the combination of these various effects gives a good approximation to the tracker acceptance, it is not possible to make it perfect without introducing

all the complexity already in GM2RINGSIM. This is the greatest limitation of this standalone MC simulation method, and means for the data EDM analysis, GM2RINGSIM must be used to calculate the acceptance correction. Nonetheless, the simple MC has shown for the first time the expected impact of both the radial field and acceptance, and that they are both very important for the final results.

5.4 Conclusions, limitations, and further work

In conclusion, a standalone, simplified and fast MC was developed to study the impact of statistical and systematic uncertainties and estimate the experimental sensitivity to a muon EDM. The results indicate that the target limit should be achievable statistically, however the radial field and tracker acceptance both increase the uncertainty and therefore the limit. The tracker acceptance is a core part of the experimental setup, and cannot be improved, however this does help inform any future experimental designs for the measurement of an EDM in a storage ring, showing that the sensitivity of a future experiment could be improved with a larger vertical acceptance.

However, the uncertainty on the radial field measurement is large enough to have a direct impact on the uncertainty of the EDM, which can be improved by better measurement of the radial field. A good target uncertainty is found to be 1 ppm to ensure any final result from the full experimental dataset would not be limited by the radial field. These results motivated a dedicated measurement, which is covered in Chapter 5.

While the results of the MC allow for many important studies to be performed, improving our understanding, the simplicity does have limitations. Firstly, the simulation does not include any beam dynamics, which impact the extraction of the tilt angle and the overall tracker acceptance. Another effect that has since been found to be very important in data is the vertical beam width, which changes with momentum, becoming much narrower at higher momentum. This limits the spread of vertical angles, making it smaller and therefore harder to extract a tilt caused by an EDM. This could be added to the MC, and is expected to shift the limits upwards, making them slightly worse than the results here would suggest. Nonetheless, the radial field remains the dominant factor impacting the EDM limit the experiment can set, so improving this measurement is a priority.

Chapter 6

The Radial Field

While the $g-2$ ring magnet provides a mostly vertical field, there are also small components in the radial and longitudinal directions. Any non-zero radial field is a key systematic for the EDM analysis, as shown in Chapter 5 due to introducing a fake 'EDM-like' signal. This was measured before the commissioning run [88] with large variations around the ring, leading to an uncertainty of around 50 ppm. This would limit the final results of the EDM analysis, motivating a new set of measurements using a novel 'indirect' beam method to push this uncertainty down as far as possible.

A non-zero radial field also contributes to the MDM analysis a small amount, as it modifies the size of the overall field and therefore the precession frequency. This means a more accurate measurement benefits both physics goals of the experiment.

This chapter describes a series of measurements taken and analysed in Spring 2021 to meet the 1 ppm target. These are then used to estimate the radial field in each dataset of Run 2/3, in order to quantify the corrections in the EDM analysis.

6.1 Measurement principles

A non-zero radial field changes the average vertical position of the muon beam due to the Lorentz force. This change in beam position $\langle y \rangle$ is resisted by the electric fields generated by the quadrupoles, which apply a restoring force. The combined $\langle y \rangle$ from these two effects and any non-uniformities result in the closed orbit distortion, which depends on the quadrupole voltage and the size of the radial field. It can be described as a function of angle around the ring, θ, by Equation 6.1, where R_0 is the ideal non-perturbed storage radius, B_0 the vertical component of the field, the B_{rcN} and B_{rsN} components are orthogonal components of the radial magnetic field, and n is the ESQ field index, defined in terms of electric field

gradient in the y-direction κ and muon velocity v [89].

$$y(\theta) \approx \sum_{N=0}^{\infty} \frac{R_0}{B_0} \frac{B_{rcN}\cos(n\theta) + B_{rsN}\sin(N\theta)}{N^2 - n^2}, \qquad n = \frac{\kappa R_0}{vB_0}. \tag{6.1}$$

This complex functional form means the vertical position of the beam varies all the way around the ring, as shown in Figure 6.1.

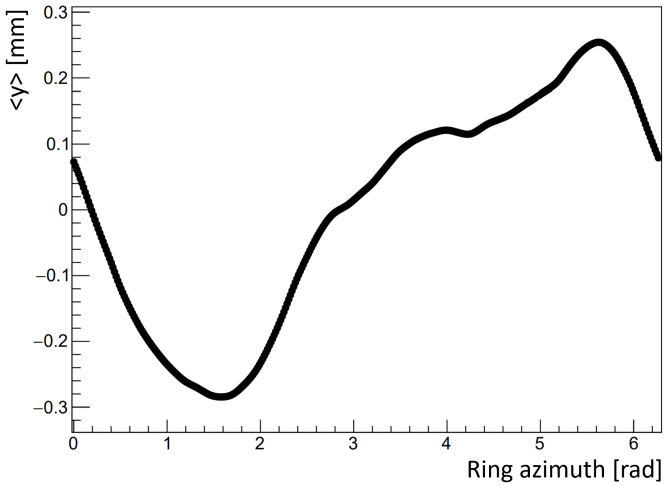

Figure 6.1: Simulation of the closed orbit distortion around the ring, generated by injecting a muon beam and applying the various magnetic and electric fields. Plot from [12], by D.Tarazona.

However, by measuring and averaging $\langle y \rangle$ all the way around the ring, the higher order terms in N cancel [58], leaving only the N = 0 term, with the details of this covered in Appendix A.2.

$$\langle y \rangle \propto \frac{R_0}{B_0} \frac{\langle B_r \rangle}{n} \propto \frac{\langle B_r \rangle}{\kappa} \tag{6.2}$$

The field gradient κ is proportional to the electric field, so this can be written in terms of the applied voltage of the electrostatic quadrupoles.

$$\langle y \rangle \propto \frac{\langle Br \rangle}{V} \tag{6.3}$$

Since $\langle y \rangle$ is now proportional to a simple ratio between the total radial field and ESQ voltage V, varying these and measuring the impact on $\langle y \rangle$ allows for an estimation of the background radial field.

6.1.1 Methodology of data collection

The vertical position of the beam is measured using calorimeter data, as although the trackers have better vertical resolution, a key requirement for this analysis is

the need to average around the ring azimuthally. The total radial field is a sum of two parts: the background radial field, which is what this study aims to measure, and the applied radial field, which is introduced by currents running through the surface coils (SCC). These are used to zero the radial field for normal operations by approximately centering the beam in the storage region, but can also be used to apply extra fields on top of this correction. If the total radial field seen by the muons is zero, then the field applied by the coils must be equal and opposite to the background radial field.

$$\langle B_r^{total} \rangle = \langle B_r^{app} \rangle + \langle B_r^{bkg} \rangle = 0, \qquad \langle B_r^{app} \rangle = -\langle B_r^{bkg} \rangle. \tag{6.4}$$

This point is found by scanning across different voltages and SCC settings. For a total radial field of zero, changing the quadrupole voltage will not change the vertical position of the beam. By plotting 1/V against $\langle y \rangle$ for a range of SCC settings, the gradients of these points are plotted against the applied field and interpolated with a line to find the applied field with zero gradient. The background radial field is then equal and opposite to this, given by the x-intercept of the interpolating line.

Cuts are applied to ensure that the $\langle y \rangle$ is measured only for stable muons. A time cut at 30 μs removes any instability due to quadrupole charging at early times, with late times after 300 μs discarded due to large statistical variation in this region. Additionally, a mid-range energy cut is applied to remove the impacts of pileup at high energies, and lost muons at lower energies. The final selected parameter range is identical to that used in the Run 4 scan [12], and is shown in Figure 6.2.

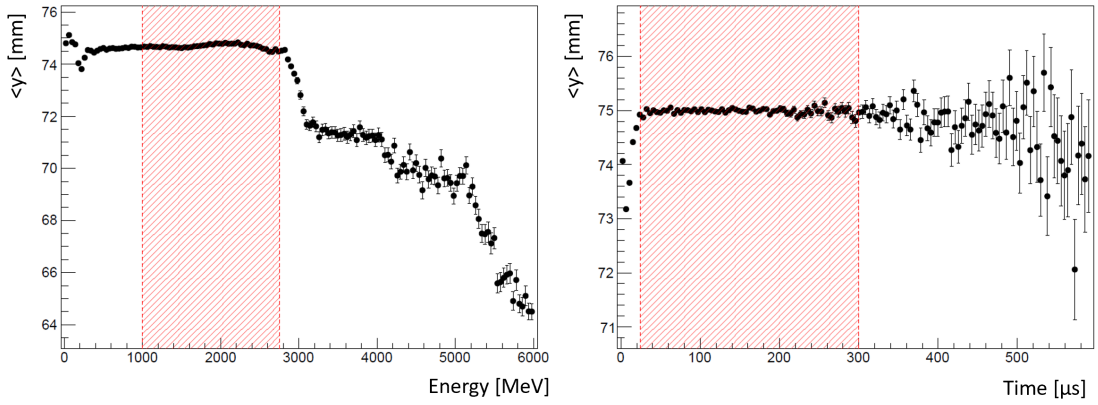

Figure 6.2: Plots showing variation of the mean vertical beam position, adapted from [12]. The red banded region is used as the 'stable' beam region for the scan.

6.1.2 Extrapolation to other running periods

The radial field scan takes a snapshot measurement of the field at one part of the run. However, it is known that the field may drift over time, particularly between runs or with different hall conditions, so the results of the scan must be extrapolated to other run periods. This can be achieved by comparing a reference mean vertical position, where the radial field is known, to the mean vertical position at a different time. This requires the calculation of the constant of proportionality k, which converts a position difference $\Delta\langle y\rangle$:

$$\Delta\langle y\rangle = \langle y_{meas}\rangle - \langle y_{ref}\rangle \tag{6.5}$$

to a radial field difference ΔB_r:

$$\Delta B_r = k\Delta\langle y\rangle. \tag{6.6}$$

The factor k is estimated empirically by reversing the radial field measurement procedure, fitting B_r against $\langle y\rangle$ for a range of voltages. The gradients of this are plotted against 1/V. The conversion factor k for a specific quadrupole voltage can then be calculated using a fitted line to extract the gradient m and the y-intercept c:

$$k = 1\Big/\frac{m}{V} + c \tag{6.7}$$

, with the total uncertainty on k being a function of m, the fitted line gradient, c, the fitted line intercept, their uncertainties and their covariance σ_{mc} [90]:

$$\delta k = \sqrt{\left(\frac{1}{Vk^2}\right)^2 \delta m^2 + \frac{1}{k^4}\delta c^2 + \frac{2}{Vk^4}\sigma_{mc}}. \tag{6.8}$$

The derivation of this can be found in Appendix A.3. Extrapolating provides a good estimation of the field, but has some limitations which are accounted for using systematic uncertainties. These are evaluated following the methodology in the Run 4 scan [12]. Firstly, changes in alignment between calorimeters will impact $\langle y\rangle$ and are likely to change between runs, so must be accounted for in the final uncertainty. The impact is estimated by comparing the vertical positions for two adjacent calorimeters to their defined reference positions, and taking the width of the resulting distribution as the uncertainty. Any estimates in run periods other than Runs 4, 5 and 6 can be expected to have a large contribution from this, as the only option is to estimate them by extrapolation.

Any drift in the vertical position will also impact the extrapolation, since the position is averaged over the dataset. The uncertainty from this is estimated conservatively as the total width of the difference in vertical beam distributions. Since each run is split into datasets, it is therefore beneficial to estimate the field in each one separately. The total uncertainty is taken as a sum in quadrature of these contributions, as they are independent:

$$\delta\Delta\langle y\rangle. = \sqrt{\delta\langle y_{\text{stat}}\rangle^2 + \delta\langle y_{\text{drift}}\rangle^2 + \delta\langle y_{\text{align}}\rangle^2}. \tag{6.9}$$

The uncertainty on the difference in radial field combines this with the uncertainty on the constant of proportionality k:

$$\delta\Delta B_r = \Delta B_r \sqrt{\left(\frac{\delta\Delta\langle y\rangle}{\Delta\langle y\rangle}\right)^2 + \left(\frac{\delta k}{k}\right)^2}. \tag{6.10}$$

The final total uncertainty on the radial field measurement is a combination of the uncertainty on the difference, and the uncertainty on the radial field at the reference position:

$$\delta B_r^{tot} = \sqrt{(\delta\Delta B_r)^2 + (\delta B_r^{ref})^2}. \tag{6.11}$$

6.2 Preliminary studies

A series of studies were performed to check the feasibility of this method and determine the optimal data collection conditions needed for the scan. These were considered in detail by S.Grant [12] but are included here to illustrate the reasoning behind experimental choices in the Run 5 scan.

A toy model is generated with 24 measurement setpoints, at voltages of 14 kV,16 kV, 18 kV and 20 kV and the applied radial field varying between ± 50 ppm in steps of 20 ppm. This is chosen to cover the full storage region without impacting storage too heavily. A background radial field of 8 ppm is injected, with the intention of studying how well this value can be recovered. The statistical uncertainties on the calorimeter cluster positions are defined based on Run 1 data, which contains 2.87 million CTAGs. The $\langle y\rangle$ position is randomly sampled from a Gaussian distribution, with width equal to this uncertainty and mean calculated using Equation 5.3. This is performed 1000 times to extract the spread in the measured radial field. The standard deviation of these results is taken as the uncertainty. The result is shown in Figure 6.3, where the injected background radial field is successfully extracted.

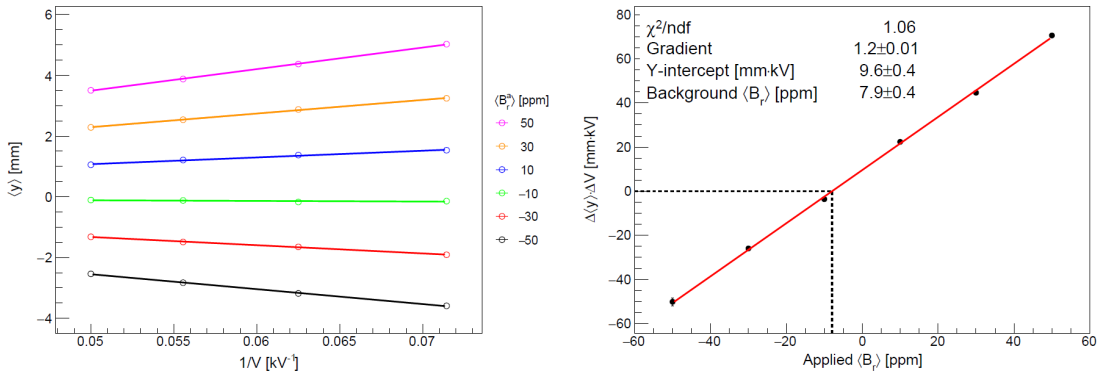

(a) Toy ESQ/SCC scans. **(b)** Toy $\langle B_r^b \rangle$ fit.

Figure 6.3: Results from the toy model of the radial field measurement, from [12]. The background radial field of 8 ppm is sucessfully extracted from the data, with an uncertainty better than the target of 1 ppm.

As the uncertainties are dominantly statistical, the chosen value in this toy model can be varied to find the minimum data collection time needed per point to reach the 1 ppm target precision. The is shown in Figure 6.4.

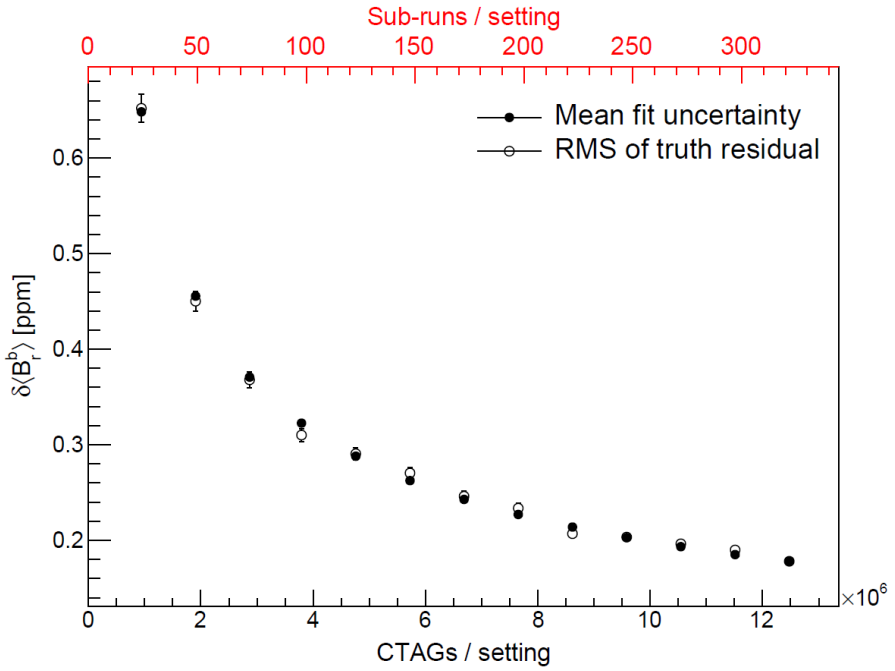

Figure 6.4: The radial field uncertainty achievable with increasing statistics per point, from [12]. It is found that a reasonably small set of data is needed per setpoint to reach the target radial field uncertainty.

From this, the target is found to be 9.5 $\times 10^5$ CTAGs per point. At the rate of data collection seen in Run 3, this was predicted to take around 4 hours of data

collection time or 10 minutes per point. In both Run 4 and Run 5, 12 hours of beam time was allocated to the scan, so more data was collected to ensure the uncertainty was comfortably below the target, and to minimise the extrapolation uncertainties. The intensity of the muon beam was also slightly lower than the Run 3 data used in the toy model, so 20 minutes per point was chosen as the minimum data collection length, corresponding to 3.8 $\times 10^6$ CTAGs per point.

A preliminary scan was taken for Run 4 with only two data points to check the calculations and method. The quadrupoles were found to be unstable at 20 kV, frequently sparking, so the upper voltage point was lowered to 19.5 kV.

6.3 Results of the Run 5 scan

The Run 5 scan ran over four quadrupole voltage settings (14 kV, 16 kV, 18 kV and 19.5 kV) and nine different applied radial fields (-90 ppm, -80 ppm, ±50 ppm, ±30 ppm, ±10 ppm, 0 ppm), with scraping turned off. The data was collected in two scan periods about a week apart. Three of these settings, the -90, -80 and 0 ppm points, were not originally intended to be part of the scan, due to the large beam displacements a very large radial field causes, which impacts storage of the beam. Despite their collection being in error, they provide useful insights, so are included in some plots but not used in the final radial field determination. During the data collection period, the CTAGs were monitored to check that the field was being changed, and that sufficient positrons were being collected per data point. This preliminary look can also be used to predict the sign and rough magnitude of the radial field, as the highest storage will occur when the total radial field is near zero. Therefore, the bin with the highest CTAGs will be equal and opposite to the true field. The CTAG map for this scan is shown in Figure 6.5. From this, it can be estimated that the background radial field is positive, and approximately 10 ppm (with a 10 ppm uncertainty).

The calorimeter data for each data point is then taken, the time and energy cuts are applied, and the average vertical angle is found combining data from all 24 calorimeters. This is plotted against the reciprocal of the ESQ voltage in Figure 6.6.

The quality of each fit is evaluated by considering the size of the fit residuals, which would show any anomalies between points. All are found to be small and similar in size, as shown in Figure 6.7. This indicates there is nothing unexpected or unexplained between the points.

The -90 and -80 ppm points are excluded from this final fit due to the extremely

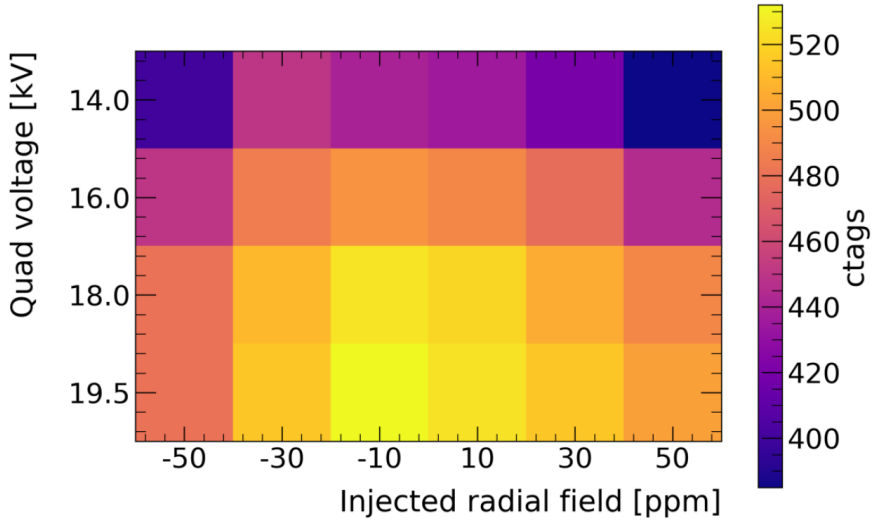

Figure 6.5: The CTAGs for each 'good' collected point in the scan. This does not include the -90 ppm, -80 ppm or the 0 ppm points as they were not part of the main scan.

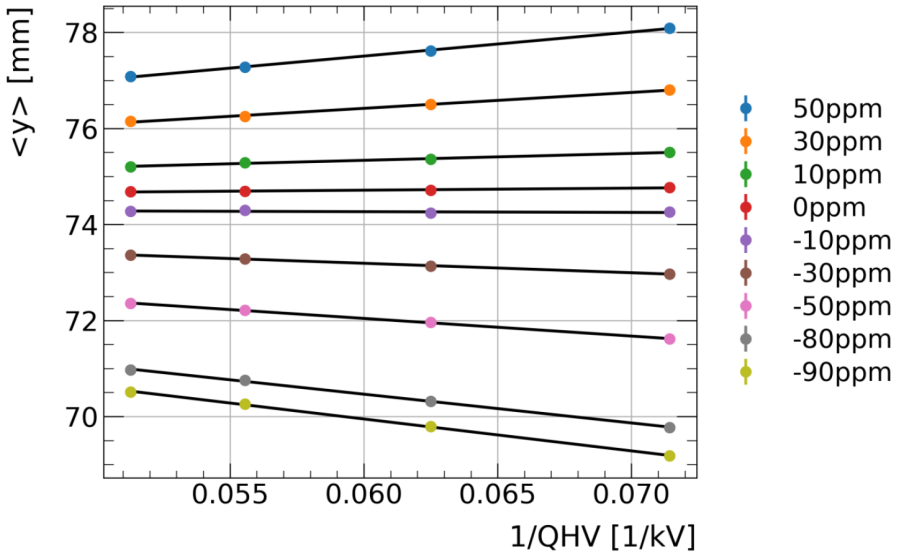

Figure 6.6: The mean vertical position around the ring, plotted against the reciprocal of the quad voltage. Each point is fitted with a straight line to extract the gradient, in order to find the radial field setting where the gradient would be zero.

large beam displacement they are expected to cause, leading to the loss of a disproportionately large part of the lower half of the beam. This would be expected to skew the average beam position. However, the 0 ppm point is expected to lie near the centre, so is still included.

Finally, the uncertainties must be evaluated. Since the applied voltage error is known to be very small, the statistical uncertainty on the mean vertical position

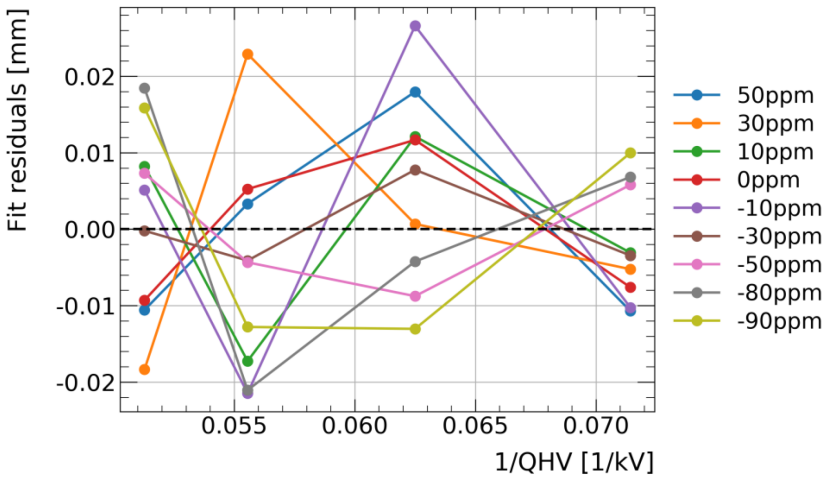

Figure 6.7: The residuals of the fits in Figure 6.6, showing that they are all small and similar in magnitude between the scan points.

will dominate the vertical position uncertainty. The applied radial field has two contributions. Firstly, the field applied by the surface coils can be expected to have some small non-uniformity across the ring. The impact of this is estimated by considering the ideal field map at the largest applied radial field (-90 ppm). The beam is then moving at most 5 mm away from the central value during the scan, which is taken as 10mm to be conservative. The field map for -90 ppm is shown in Figure 6.8, with the 10 mm range marked in the centre.

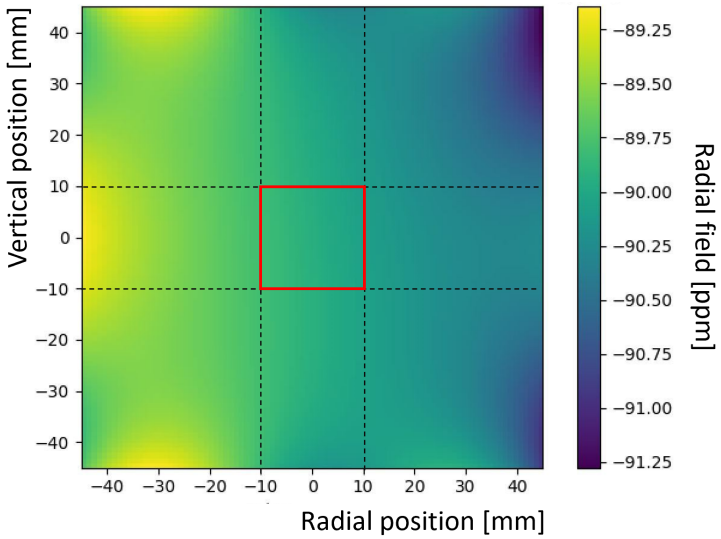

Figure 6.8: The applied radial field showing inhomogeneities from the surface coils. The centre square region is a conservatively-defined region which the scan beam moves in, so the uncertainty is taken as the maximum difference seen in this region.

The extremes of this region are found to differ by at most 0.1 ppm, which is taken as the final uncertainty. Additionally, the ring magnetic field may have an extra vertical non-uniformity. This is estimated using the variation seen in the commissioning field scans [88] as 2.5 % per ppm away from the zero crossing point. The final fit must be iterated until the uncertainties and parameters settle. The final fit result is shown in Figure 6.9.

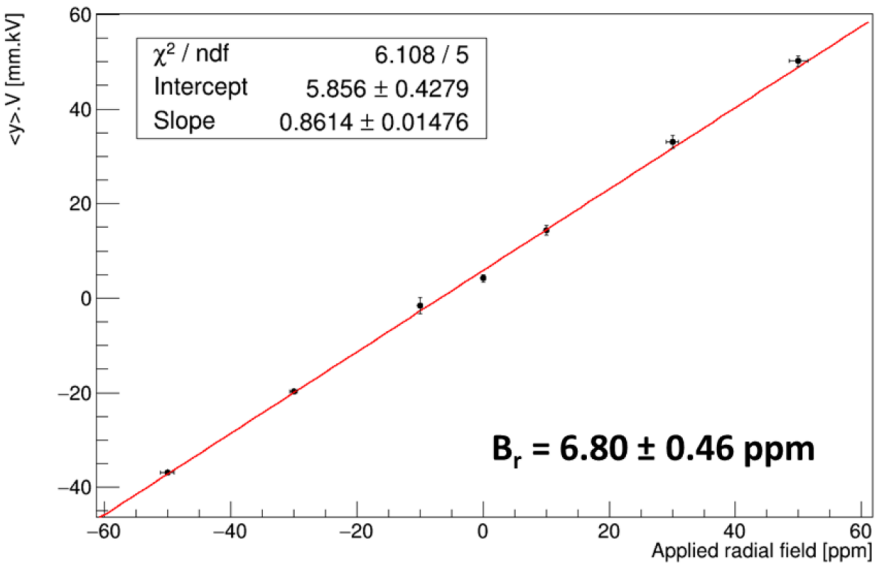

Figure 6.9: The final straight line plot to the Run 5 scan data. The final radial field is extracted as equal and opposite to the x-intercept of the fitted line.

From this, the final background radial field is extracted as 6.80 ± 0.46 ppm, well within the target precision of 1 ppm uncertainty, and consistent with the initial guess from the CTAG map of around 10 ppm. Since the central value is small, there is no need to zero this field, but it is informative to extrapolate this field across the rest of the run and back to estimate the radial field in Run 2/3.

6.4 Extracting the radial field for Run 2/3

The radial field result from the scan is extrapolated to other runs. First, a comparison setpoint must be defined, where the radial field is known. As the scan data was collected across two periods of time with about a week of normal running in between, data from this week is used as the reference, with a simple 0th order polynomial fit being performed to extract the mean vertical position across all 24 calorimeters. This is shown in Figure 6.10. This reveals a new difficulty which arises for Run 5, which was not seen in Run 4. There is a slight drift upwards

in the mean vertical position, even on the short timescale of a week. This is accounted for by the uncertainties, but will lead to a larger error on extrapolated values as a result. Despite this, Run 5 remains a better comparison setpoint due to including all 24 calorimeters, whereas in Run 4 calorimeters 5 and 9 had to be removed from the averages due to faulty SiPMs, biasing the averaging around the ring. The cause of this drift is investigated as part of the extrapolation process.

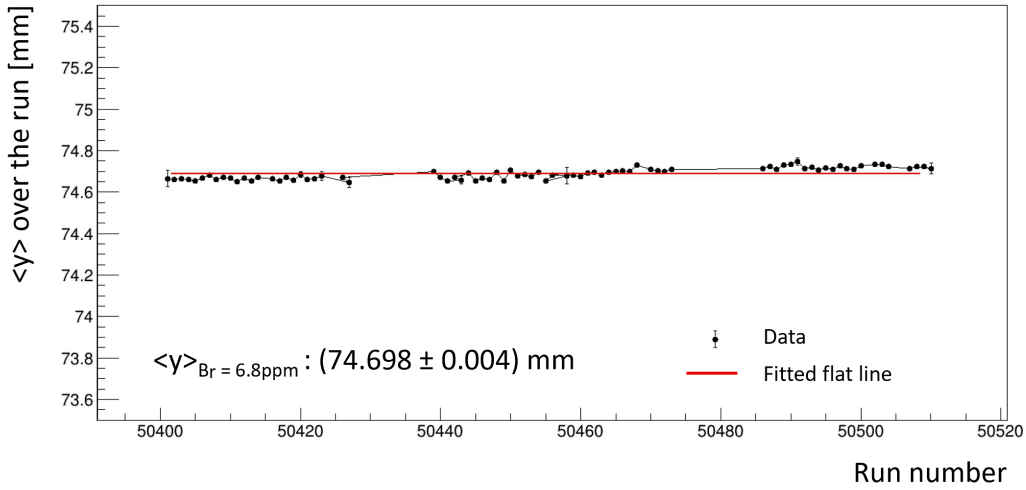

Figure 6.10: A fit to $\langle y \rangle$ over the week of normal conditions running data between the two radial field scans. This reference position, despite showing some drift, will give a good comparison point for the average of the two scans.

The next step is to extract the value of the constant k. This is expected to be very similar to Run 4, but changes in the run conditions could lead to a small deviation. The fits of mean position against applied vertical field, and the final straight line fit to extract the slope and intercept, can be seen in Figure 6.11.

The results are similar to those from the Run 4 scan [12], and consistent within errors, but the uncertainties in Run 5 are slightly larger due to the drift. These can then be used to extract the final value of k as 20.1 ± 0.4 ppm/mm. As a final check, this is compared to the differences in Figure 6.6, with the values found to differ with radial field as expected.

With the constant of proportionality calculated, the radial field can now be estimated for any dataset where the mean vertical position is known. Firstly, the whole run period of Run 5 is considered on a subrun-by-subrun basis, to study whether the upward drift seen is a trend across the full run. The mean vertical position is found and converted into a radial field difference, then the absolute radial field is calculated and plotted vs run number. The result of this is shown in Figure 6.12.

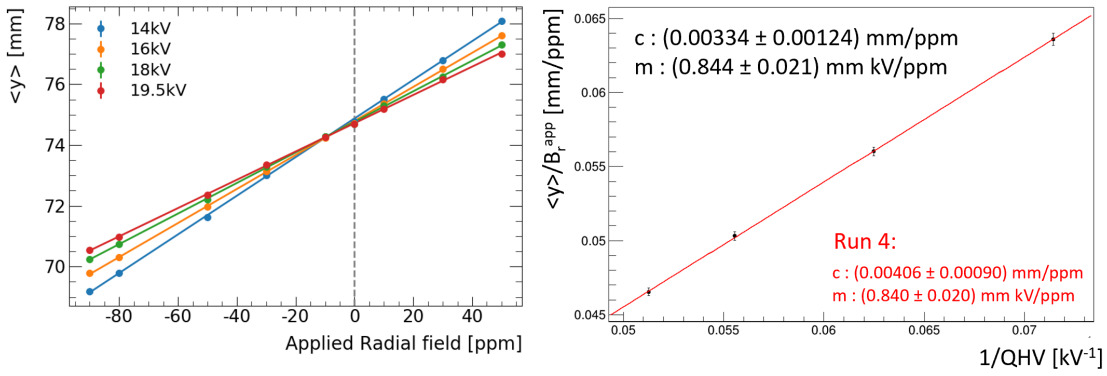

Figure 6.11: Using the radial field scan data to extract the components of the constant of proportionality k. The two parameters are compared to Run 4, and found to be consistent but not identical.

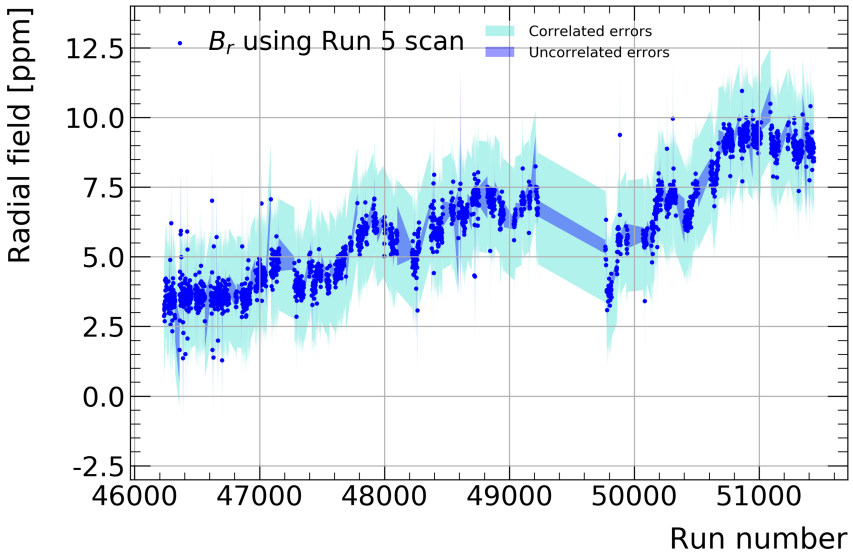

Figure 6.12: The radial field across the majority of Run 5. While small variations are seen locally, globally there is a trend of increasing radial field.

The most notable part of this is that the radial field is drifting upwards across Run 5. The uncertainties are separated into correlated and uncorrelated uncertainties, since only the statistical error is responsible for fluctuations of the points. The correlated uncertainties are dominated by the calorimeter alignment.

In the same way, the radial field is estimated for each dataset of Run 2 and Run 3, with an average taken across the whole dataset rather than by subrun.

Interestingly, Run 2 and Run 3 also drift up over time. Understanding the origin of this is important, particularly for the EDM analysis. Anything causing this must have variability on short timescales of hours to days, as well as a slow

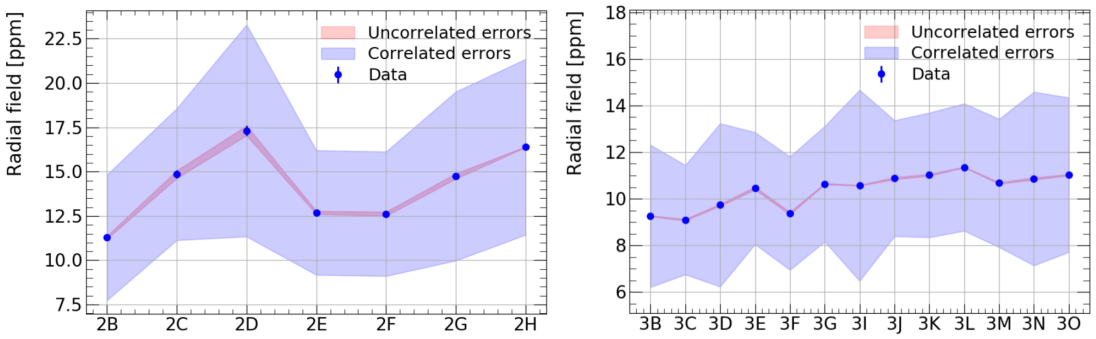

Figure 6.13: The radial field across the entirety of Run 2 and Run 3. Similar rise behaviour is seen in both run periods, with Run 2 having greater variability than Run 3.

change in one direction over a longer timescale. For this reason, the temperature is a possible cause, since this would match the behaviour seen with runs starting in the winter and ending in the summer. This is an easy correlation to check, as the hall and magnet temperatures are measured by sensors in the hall, and it is known that the temperature stability in the hall was improved prior to Run 3 with better cooling and a blanket on the magnet. Therefore, Run 2 would be expected to be more variable with a larger overall drift than Run 3. Additionally, very little drift is seen across the Run 4 radial field [12], so the temperature can be checked to see if it varied less during this time.

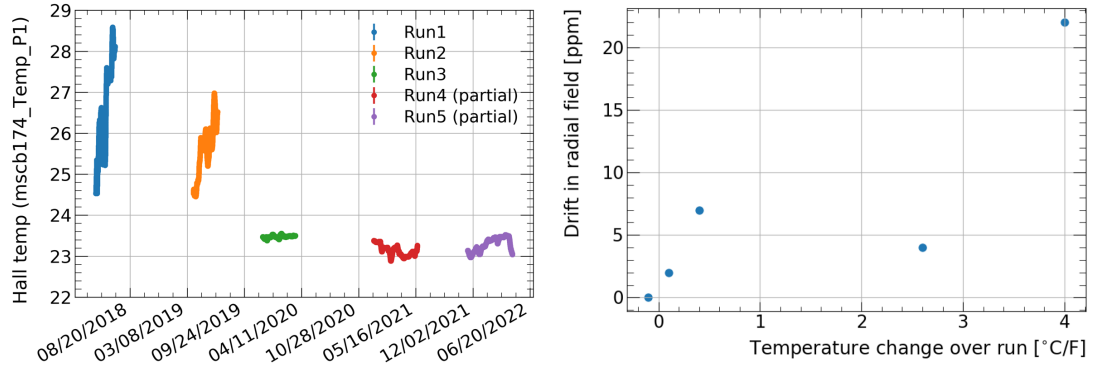

Figure 6.14: The variation in hall temperature across runs (left) and the change in radial field plotted vs the temperature change (right). While there are not many data points, a weak trend is seen, suggesting that the temperature is likely a major cause but not the only contributor.

All the runs that see drift in radial field also see a drift upwards in temperature, with varying severity. The shapes seen in Run 2 and 3 match the shorter term variations and show the improvement from better hall cooling. This suggests that the temperature is likely a key factor, but possibly not the only one.

6.4.1 Converting radial field to tilt angle

In order to use the results of this chapter in the EDM analysis, the radial field values must be converted into the tilt of the precession plane. This can be done simply using the ratio of the field components, as the radial field is defined relative to the main vertical field in ppm:

$$\delta B_r = \frac{B_r}{B_y}. \tag{6.12}$$

This tilt is then converted into an effective d_μ value using the conversion in Equation 2.25. Following this procedure for the radial field values across Run 2/3, the final tilt angle contributions and uncertainties are given in Tables 6.1 and 6.2. These will be subtracted from the final central value for the EDM, with the uncertainties directly contributing to the EDM uncertainty and therefore the limits that can be set.

Subrun	Radial field [ppm]	Equivalent $d_\mu(\times 10^{-20})$ [e · cm]
2B	11.3 ± 3.62	4.18 ± 1.34
2C	14.9 ± 3.93	5.50 ± 1.46
2D	17.3 ± 6.26	6.41 ± 2.32
2E	12.7 ± 3.60	4.70 ± 1.33
2F	12.6 ± 3.63	4.97 ± 1.34
2G	14.7 ± 4.92	5.46 ± 1.82
2H	16.4 ± 4.98	6.10 ± 1.84

Table 6.1: Equivalent d_μ tilts caused by the radial field in Run 2

Due to no direct scans being performed in Run 2 and Run 3, the necessary extrapolation means these uncertainties are larger than the 1 ppm target, and larger than the values measured directly in Run 4 and Run 5. However, the 1 ppm target was calculated considering the total experimental dataset, which is much larger than each dataset here, and indeed much larger than Run 2/3. Therefore, the statistical error dominates still, and the radial field has been estimated to sufficient precision to not limit the Run 2/3 EDM result. More discussion of the relative size of uncertainties can be found in Chapter 7.

Subrun	Radial field [ppm]	Equivalent $d_\mu(\times 10^{-20})$ [e · cm]
3B	9.26 ± 3.07	3.43 ± 1.14
3C	9.09 ± 2.37	3.37 ± 0.879
3D	9.73 ± 3.55	3.60 ± 1.31
3E	10.4 ± 2.46	3.89 ± 0.912
3F	9.38 ± 2.49	3.47 ± 0.921
3G	10.6 ± 2.50	3.94 ± 0.927
3I	10.6 ± 4.11	3.91 ± 1.52
3J	10.9 ± 2.53	4.03 ± 0.937
3K	11.0 ± 2.71	4.08 ± 1.00
3L	11.4 ± 2.75	4.20 ± 1.02
3M	10.7 ± 2.78	3.95 ± 1.03
3N	10.9 ± 3.77	4.02 ± 1.40
3O	11.0 ± 3.34	4.08 ± 1.43

Table 6.2: Equivalent d_μ tilts caused by the radial field in Run 3

6.5 Conclusions and future improvements

The radial field is an important systematic for the EDM analysis as a non-zero field looks like an EDM signal. This has been successfully measured during the Run 5 run period, to within the 1 ppm target precision, meaning the EDM analysis will not be radial field limited for any of the runs.

The central value of the field in Run 5 is found to be 6.80 ± 0.46 ppm. This is extrapolated out across Run 5 to study the variation of the field over time, and also to Run 2/3 for the main analysis in this thesis. The radial field across all these periods is found to be drifting, likely due to hall temperature changes. While this drift is accounted for in the extrapolation uncertainties, this could be improved in the future by looking at the data in smaller time steps. This would allow for better correction of the radial field, however the size of the dataset also contributes to the uncertainty. Therefore, for Run 2/3, which is statistically dominated, there is no need to switch to smaller time steps for the analysis, but this could still be useful for the full dataset.

The dominant uncertainty on the extrapolated field values is the calorimeter alignment. This could be removed or reduced by using known positional information to correct the vertical position rather than just assuming it is maximally incorrect. Additionally, the non-uniformity in the radial field for different vertical positions could be directly measured all the way around the ring rather than us-

ing a conservative estimate. A preliminary form of this measurement has recently been performed, showing the estimate in the vertical direction remains valid.

One final improvement would be a combination of setpoints between the Run 4, 5 and recently completed Run 6 scan. This would roughly triple the statistics, improving the fits in the vertical position and therefore reducing the final uncertainty even further. However, this is subdominant compared to the calorimeter alignment and drift uncertainties. Nonetheless, the measurements achieved so far have more than enough precision to not be limited by the radial field for the foreseeable future.

Chapter 7

The Run 2/3 EDM analysis

7.1 Introduction

The Run 2/3 tracker analysis follows on from the Brookhaven tracker analysis discussed in Section 3.15 [1] and S. Grant's Run 1 EDM analysis [12], and is expected to achieve a significant improvement in limit compared to these. There are three key differences compared to the Run 1 analysis: firstly, there is a large increase in statistics, with almost ten times as many tracks, so any systematic effects become more important. Secondly, Run 1 had damaged resistors in the quadrupole focussing magnets, which caused a slow early time rise in the average vertical angle. Finally, further investigations into the expected functional form due to results from a high-stats MC led to the discovery that a fit function modification is needed to accurately fit the data. This is discussed in detail in Section 7.2.7. Systematic uncertainties, particularly the tracker acceptance correction uncertainties, are also reduced due to the implementation of a new method for calculating the corrections. The final result is statistically limited as a result of these improvements.

Run 2 data was collected between Autumn 2018 and Summer 2019, and Run 3 between Autumn 2019 and Spring 2020. The number of tracks in each dataset after basic EDM analysis cuts are applied (detailed in Section 7.2.2) are shown in Figure 7.1.

There is significant variation in the size of datasets, with some being comparable to the Run 1 datasets while others are much larger. Since this data was collected over a long time period, during which conditions may have changed, each dataset is analysed separately, including all fits and corrections. The final EDM result is extracted by fitting across all datasets per run.

This chapter covers the methodology of the EDM analysis and a preliminary set of blinded results for the Run 2/3 data. Run 3B is chosen as a 'typical' dataset

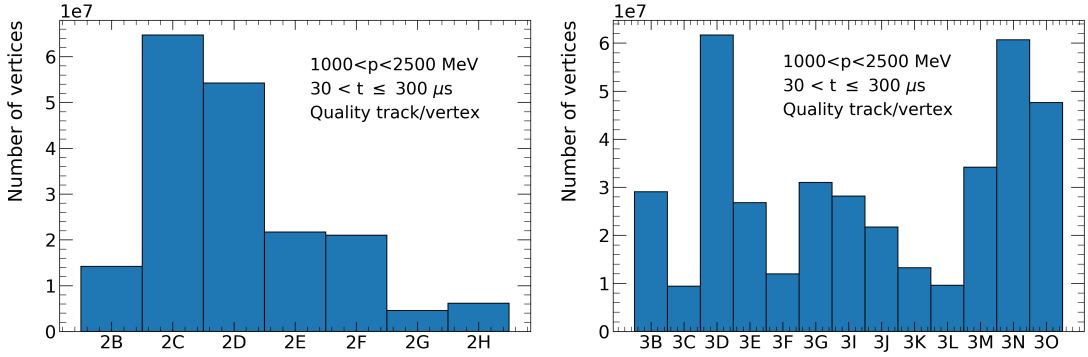

Figure 7.1: The number of vertices in each Run 2/3 dataset passing the basic EDM quality cuts. Large variation is seen in their respective sizes. The momentum and time cuts are discussed in the next section, and quality track/vertex cuts are described in Section 3.8.

to illustrate the steps of the analysis, as it is an average sized dataset and lies roughly in the middle of the two data collection periods.

7.2 EDM analysis methodology

7.2.1 Blinding data

When working with data, particularly when looking for a small signal among backgrounds, it is important to ensure there are no biases introduced by the analysis method. This is achieved by a process called 'blinding', where the true signal is hidden until the analysis is complete and frozen. In the case of the EDM analysis, there are two blinded quantities to consider: the muon precession frequency, and the EDM amplitude. The precession frequency is blinded as described in Section 3.10, as at the time of writing the Run 2/3 ω_a results were still blinded. However, this doesn't blind the EDM analysis so an additional blinding must be applied.

This is achieved by injecting a fake EDM signal on top of the data, produced using the ω_a blinding library. A Gaussian distribution is defined based on a range of large values for the EDM amplitude, in units of the limit set by BNL. An example distribution can be seen in Figure 7.2. The random number generated from the blinding string is used to sample from this distribution to extract the amplitude.

A large amplitude is chosen to ensure any signal comparable to the BNL limit is masked, and to make plots that are blinded easily identifiable. Therefore, the central value must be chosen to make it unlikely ($< 0.1\%$) that the injected amplitude is less than the BNL limit. The Gaussian is plotted for a range of central

values to find the minimum central value, which is found to be 3.75 times the BNL limit.

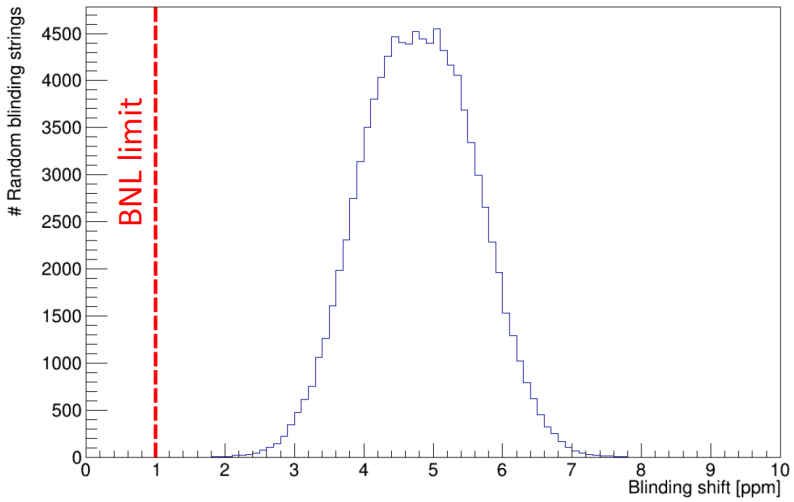

Figure 7.2: An example distribution for the EDM blinding. The injected EDM amplitude is sampled from this based on a random number generated from a blinding string.

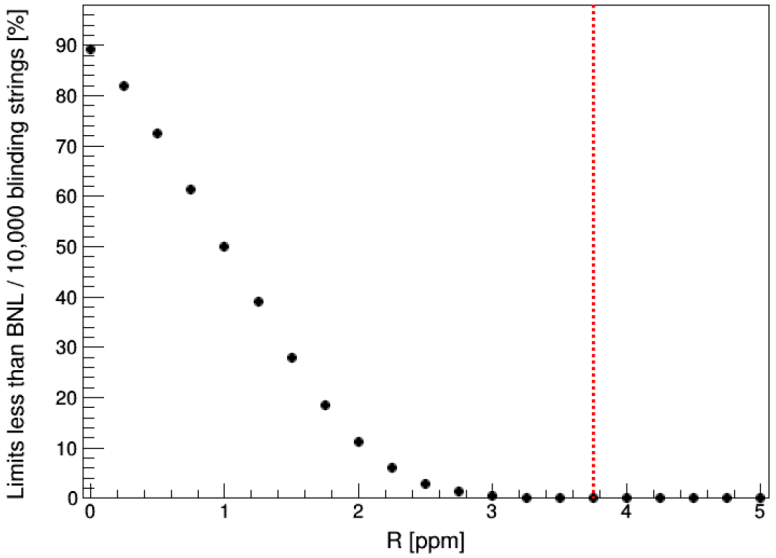

Figure 7.3: Sampling 10,000 blinding strings from a Gaussian distribution of width 1 x the BNL EDM limit and a moving central value, R, counting how many strings give values equal to or less than the BNL limit. The threshold at which 0.1 % of the strings fall below this boundary is marked with the red line.

It is important to check that the blinding doesn't impact or bias the fits in any way. This can be done by applying it to MC generated data with an injected EDM and comparing the fits with and without the blinding. The results of this can be

seen in Figure 7.4, showing that all other parameters, as well as the fit quality, remain the same within uncertainties both with and without the blinding.

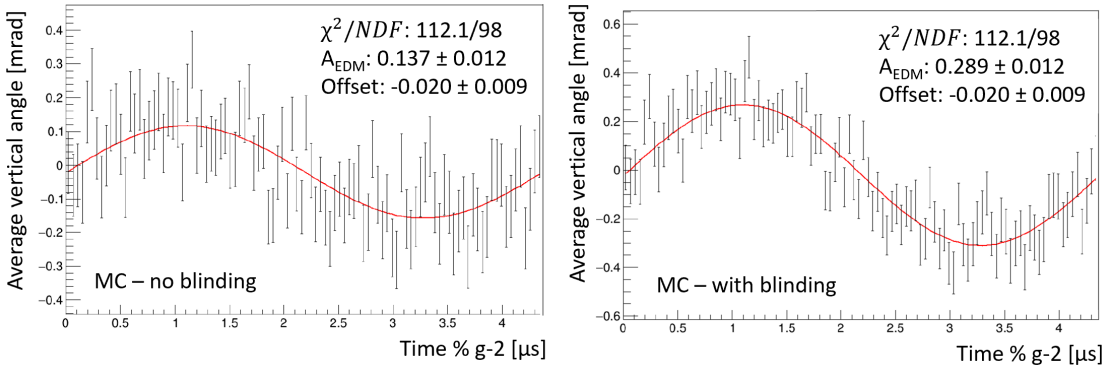

Figure 7.4: The impact of injecting a blinding EDM in MC data, before and after the injection. The EDM amplitude increases but all the other fit parameters and uncertainties remain the same.

This, however, only shows the impact of one particular blinding value, so a wider range of injected signals is scanned over. The χ^2/NDF becomes more variable as the blinding shift is increased, likely due to small effects in the data becoming artificially more magnified, however the variation is small compared to the expected statistical variation. The other important quantity to check is the uncertainty on the EDM amplitude, which should remain unchanged as the injected signal increases. These two scans are shown in Figure 7.5.

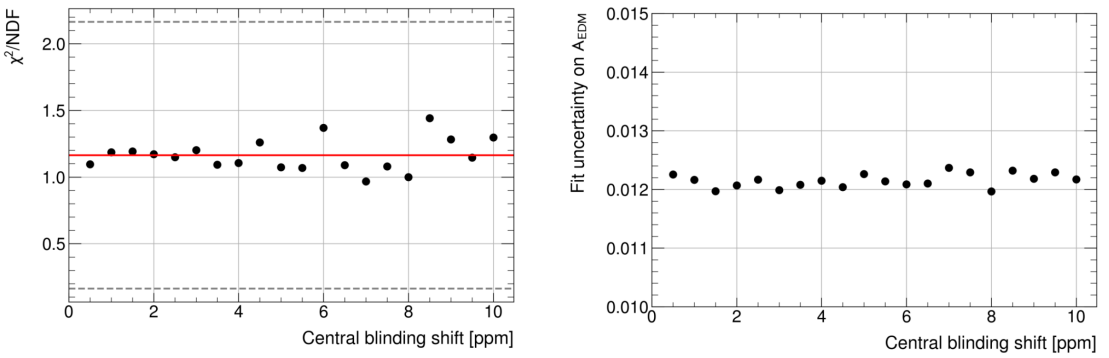

Figure 7.5: A scan of different blinding central values, showing the lack of impact on the EDM quality of fit and the fit amplitude uncertainty. The dashed grey lines in the left plot show the expected statistical variation of the fit quality, with all variations seen being much smaller than this range.

Given that the size of the blinding does not impact either quantity, the central value of the sampling distribution is chosen to be larger than the minimum value for safety. From Figure 7.3 any central value in the range 4-5 ppm would be

good, with 4.81 being selected. This blinding is used for all data in the Run 2/3 EDM analysis, with Run 2 and Run 3 separately blinded to allow for individual unblinding once the analysis is complete.

To inject an accurate signal, the shape in data is required. As a result, the blinding changes partway through the analysis due to a better understanding of the fit function. These changes are detailed in Section 7.2.7.

7.2.2 Extracting the average vertical angle oscillation

The number oscillation seen in the tracker data is plotted for positrons with a momentum larger than 1900 MeV. This wiggle plot is then fitted with a 5-parameter fit of the form:

$$N_0 \exp(-t/\tau) \left[1 + A\cos(Rt + \phi)\right] \tag{7.1}$$

,in order to extract the $g-2$ phase, ϕ, where N_0 is the initial amplitude, τ the boosted muon lifetime, A the asymmetry and R the blinded $g-2$ frequency, which is kept both software and hardware blinded. This is performed between 30.6 μs and 300.6 μs, which avoids beam instability at early times and larger statistical variations at late times due to muon losses. Results for each dataset are shown in Table 7.1, and plotted for Run 3B in Figure 7.6. Each value is cross-checked with the phases found in the main ω_a analysis to make sure the values are consistent.

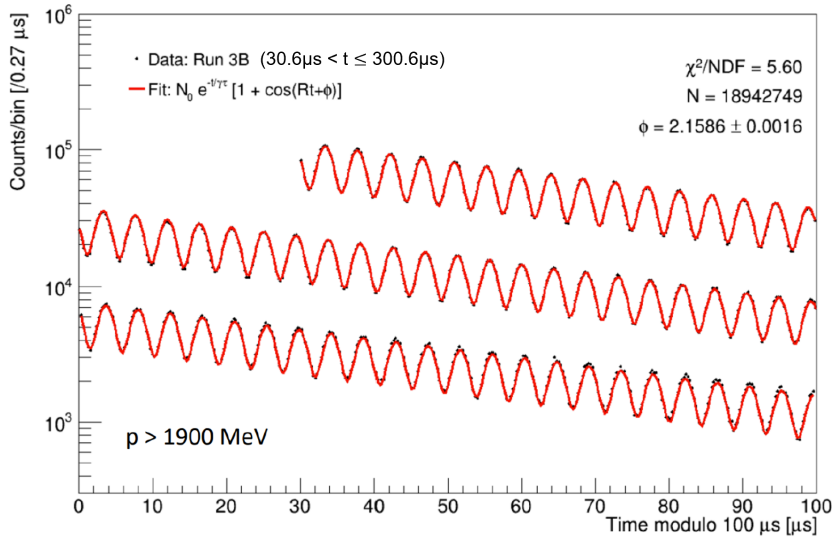

Figure 7.6: 5-parameter fit for Run 3B, extracting the phase parameter for the EDM fit. This is plotted modulo 100 μs and on a log scale for easier viewing.

Run	Phase [rad]	Run	Phase [rad]
2B	2.164 ± 0.002	3B	2.159 ± 0.002
2C	2.162 ± 0.001	3C	2.179 ± 0.003
2D	2.160 ± 0.002	3D	2.163 ± 0.001
2E	2.155 ± 0.002	3E	2.145 ± 0.002
2F	2.165 ± 0.002	3F	2.149 ± 0.003
2G	2.154 ± 0.004	3G	2.138 ± 0.002
2H	2.156± 0.003	3I	2.154 ± 0.002
		3J	2.151 ± 0.002
		3K	2.149 ± 0.003
		3L	2.150 ± 0.003
		3M	2.157 ± 0.002
		3N	2.165 ± 0.001
		3O	2.170 ± 0.001

Table 7.1: Fitted phases from the 5-parameter fits for each run within Run 2/3. All have small uncertainties and are roughly consistent with the ω_a analysis phases.

Next, the average vertical angle is plotted modulo the $g-2$ period, which reduces the statistical uncertainty in each bin and averages out any slow oscillations not at the $g-2$ frequency. A momentum cut of 1000-2500 MeV is applied to keep positrons with the greatest EDM sensitivity, as shown in Figure 2.8. The resulting data is then fitted for two sinusoidal components, fixing the phase.

$$A_{g-2}\ \cos\left(\omega_a t + \phi_{g-2}\right) + A_{EDM}^{blind}\ \sin\left(\omega_a t + \phi_{g-2}\right) + c \tag{7.2}$$

The first term is an oscillation in phase with the $g-2$ oscillation, which is the expected signature of a non-zero longitudinal magnetic field, arising from a tilt in the direction of travel of the muons. The exact amplitude of this fit does not matter for the EDM analysis, and has previously been measured to be consistent with zero [48], but prevents any sine component from being incorrectly absorbed into any oscillation which may be present. The second term is the blinded EDM oscillation, $\pi/2$ out of phase with the $g-2$ precession. Finally, the third term is a constant value which accounts for any offset in the central value from zero. This arises due to a combination of tracker alignment and tracker acceptance, and must be included to prevent amplitude suppression from fitting around zero. The fit for Run 3B is shown in Figure 7.7.

The uncertainty on the phase is propagated through to the impact on the

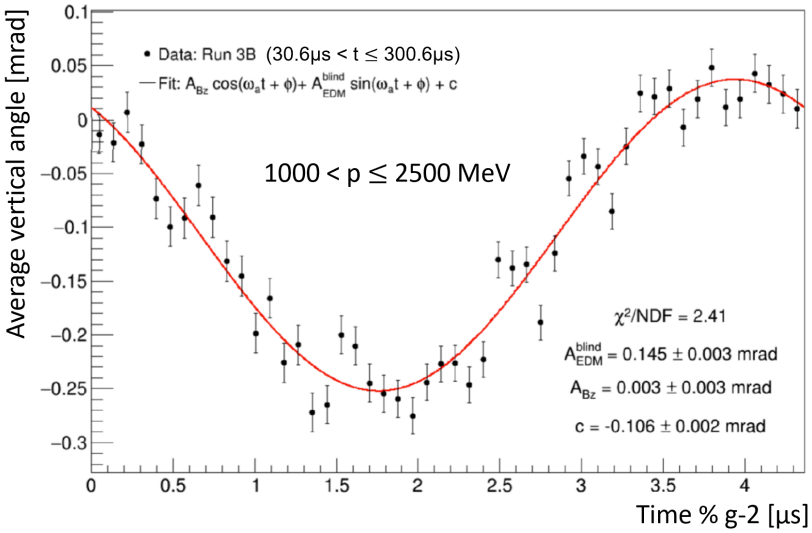

Figure 7.7: The blinded average vertical angle plotted modulo the $g-2$ period and fitted with the function in Equation 7.2 for Run 3B. The EDM amplitude parameter A_{EDM}^{blind} is the main focus of this analysis.

amplitude in Section 7.2.9 and found to be very small. This indicates that a more complex fit is not required.

7.2.3 Checking the analysis cuts

As part of the analysis, it is important to check that the cuts are optimal for the stability of the fit parameters. The offset parameter c is expected to have a momentum dependence and possibly also a time dependence. While the correlation between the offset and the two amplitudes is found to be low, the EDM amplitude is also momentum-dependent, so it is important to understand the behaviour of the offset to ensure it does not affect the fitted amplitude.

Figure 7.8 shows the variation of c with time and momentum. There is a sharp rise at early times due to scraping, which then levels off. This could also potentially be a consequence of high occupancy in the tracker at early times making reconstruction more difficult. This is stable by around 30 μs, so a 30.6 μs start time is appropriate.

For momentum, the offset has a central plateau region, outside of which it drops or increases sharply. This is likely a side effect of the tracker acceptance, as the tracking struggles with low momentum tracks due to them being highly curved. Therefore, the momentum cut at 1000 MeV is optimally placed, and cannot be lowered to increase the statistics.

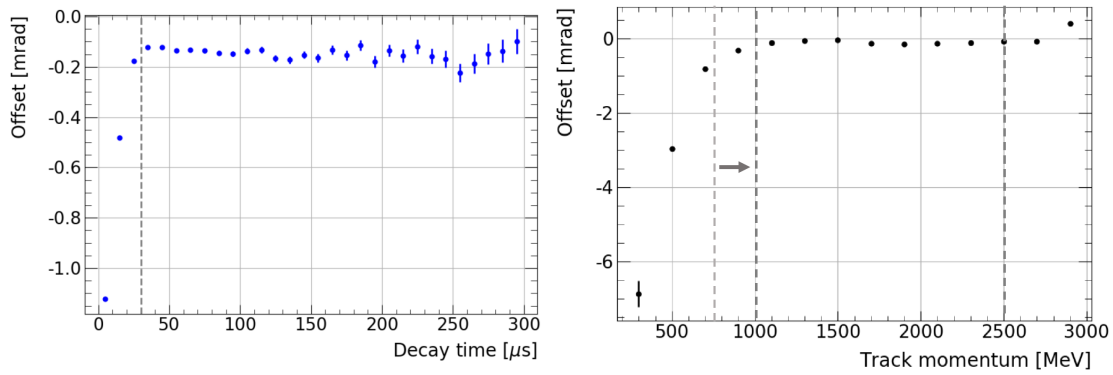

Figure 7.8: The variation of the offset with time cut and momentum cut. It is expected to vary with both due to being linked to tracker acceptance. These plots show that the cuts chosen give a stable value for the offset, and that the momentum cut cannot be reduced.

7.2.4 Beam dynamics corrections

Vertical betatron oscillations

Performing a sinusoidal fit on the Run 2/3 datasets produces poor results, with the chi-squared values in general found to be much greater than one, as shown in Figure 7.9.

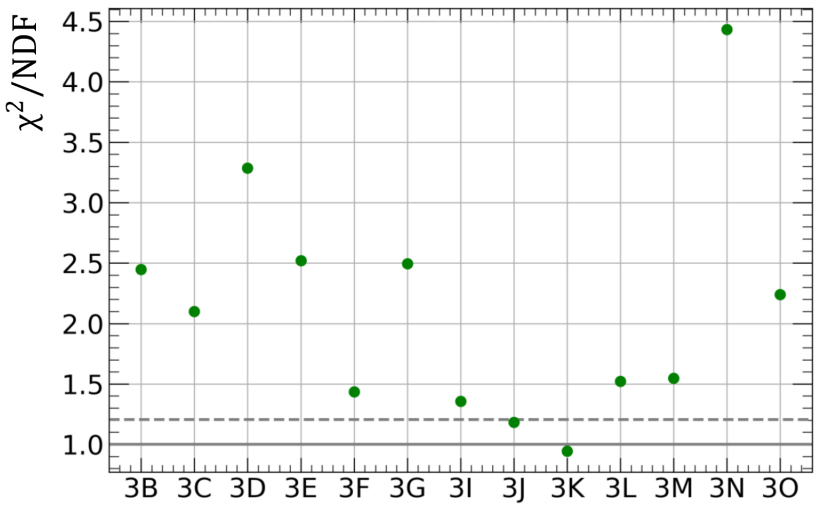

Figure 7.9: The chi-squared value divided by the number of degrees of freedom for all the datasets in Run 3 after following the above fit procedure. The fit quality for most is found to be suboptimal, suggesting there is something unfitted, for example a beam dynamics effect, impacting the results.

This suggests the fit in Equation 7.2 is missing some part of the functional shape in the data. This could be due to beam motions, so a fast Fourier transform

(FFT) is performed on the residuals of the EDM fit to find any unfitted oscillatory behaviour. The results of this for Run 3B can be seen in Figure 7.10.

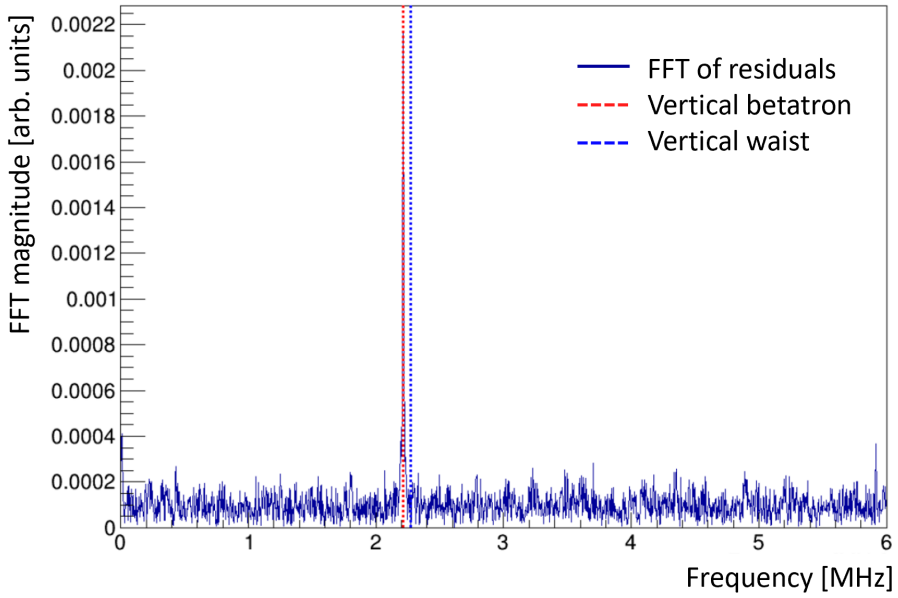

Figure 7.10: FFT of the fit residuals of Figure 7.7, showing a large peak at 2.21 MHz. This corresponds to the vertical betatron oscillation.

The FFT reveals a large peak at 2.21 MHz, which corresponds to the vertical betatron frequency (VB). The vertical waist frequency is also marked on the FFT for comparison, however no peak is seen. There are two possible approaches for dealing with this extra frequency, either a randomisation to remove it from the data entirely, or a simultaneous fit with the EDM. Previously, only the randomisation method was used, so the fitting method is investigated in case the increase in statistics leads to improvements.

The randomisation algorithm is based on the one used for the ω_a analyses, and involves shifting the the time of decay slightly by sampling from a uniform distribution in the range $\pm T_{VB}/2$, where T_{VB} is the vertical betatron period. This flattens any oscillation at the vertical betatron frequency. A FFT is performed again to ensure that the oscillation has been removed successfully, and to check for any other frequencies that may now be visible.

As Figure 7.11 shows, the peak has been successfully removed, and no other peaks appear. The χ^2/NDF values for the Run 2 and Run 3 datasets after this randomisation are shown in Figure 7.12, showing sizeable improvement compared to before the randomisation.

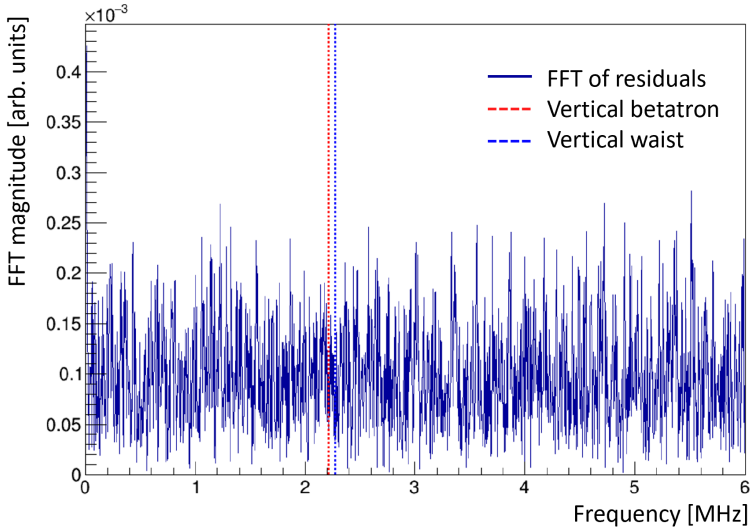

Figure 7.11: FFT of the fit with randomisation applied. The large peak is gone, and no other peaks are visible other than a small peak at 0, indicating the likely presence of a slow effect.

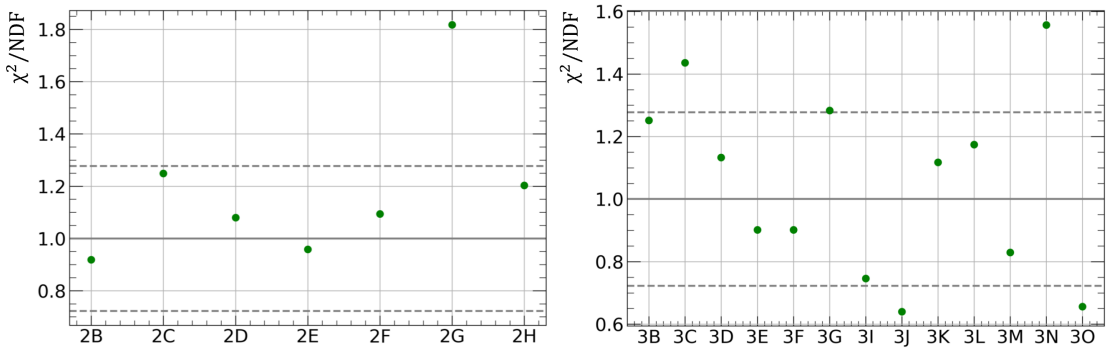

Figure 7.12: The fit quality for both Run 2 and Run 3 after randomisation, with the grey dashed lines indicating expected statistical variation. A large improvement is seen.

It is still beneficial to check whether fitting would give a better overall quality of fit. The expected form for the VB includes an oscillation and a decay, so an additional term of the form:

$$A_{VB} \exp(-t/\tau_{VB}) \cos(\omega_{VB} t + \phi_{VB}) + m \tag{7.3}$$

is tried, with A_{VB} being the initial amplitude of this oscillation, τ_{VB} the lifetime, ω_{VB} the frequency and ϕ_{VB} the phase. The m term is added as an offset term, though when fitted simultaneously with the EDM the c offset term will absorb this. Initially the EDM fit residuals are fitted as proof that this function works. Parameters are allowed to float with the lifetime constrained as positive.

The result of the fit is shown in Figure 7.13.

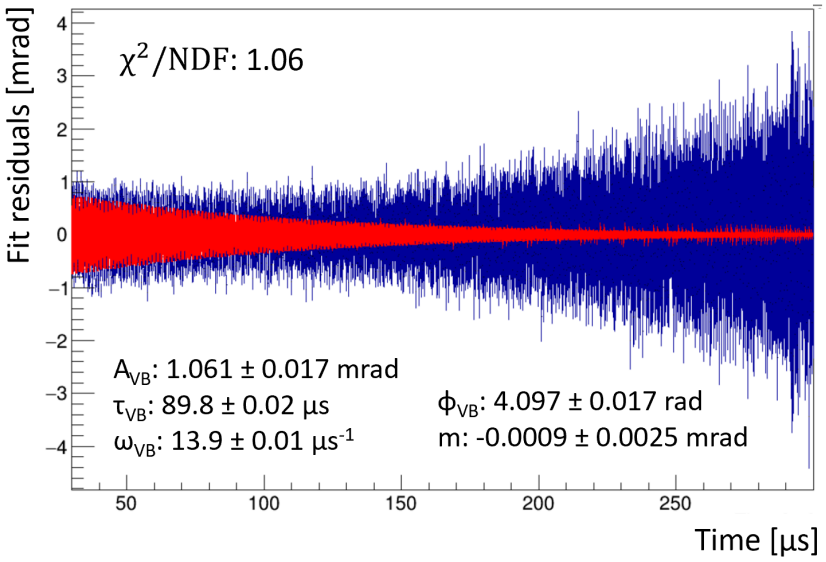

Figure 7.13: A fit (red) to the residuals (blue) of the EDM fit, extracting the VB oscillation parameters successfully. This fit shows that the majority of the spread at early times is due to the VB, which decays away at late times.

The oscillation is fitted successfully, with a good quality of fit. The correlation matrix (shown in Appendix C) shows all fit parameters are highly correlated. A toy model is used to investigate the expected shape of the VB in the modulo plot, shown in Figure 7.14. The total sum is overlaid in black.

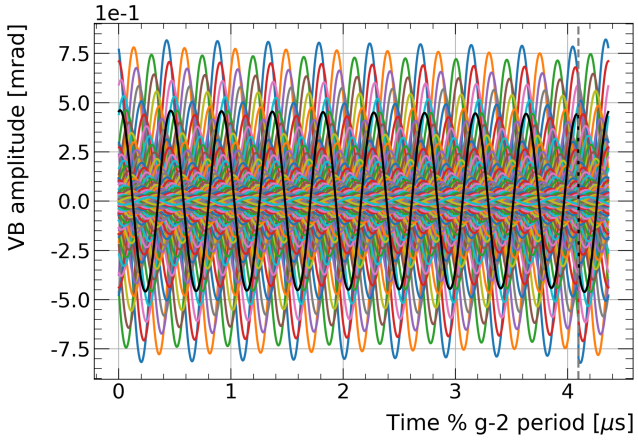

Figure 7.14: Multiple periods of the VB oscillation overlaid, modulo the $g-2$ period, with the sum of the coloured lines shown in black.

The summed shape of the modulated VB is still a full, fast oscillation. In this toy model, a discontinuity is seen due to starting at exactly 30 μs rather than an integer number of $g-2$ periods. The amplitude and phase of the final modulo VB

oscillation are both shifted relative to the initial oscillation, so this is investigated by varying the start time, which shifts the modulated components relative to each other. The results are shown in Figure 7.15.

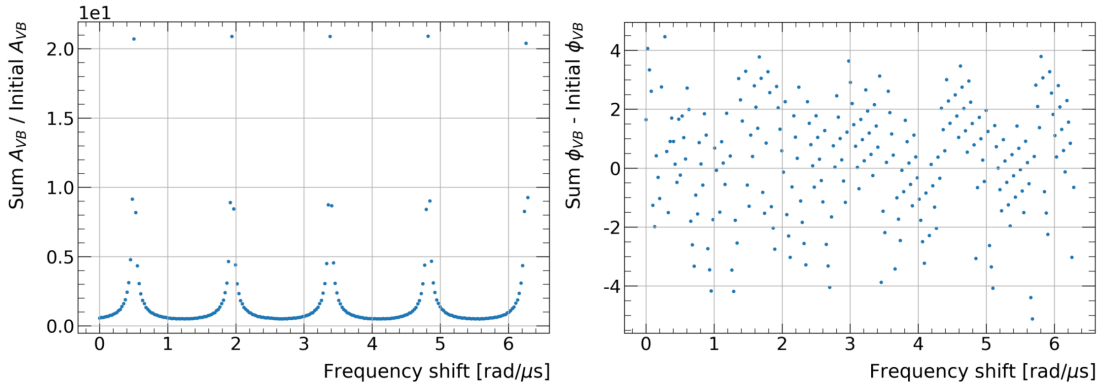

Figure 7.15: The variation of the final (black curve) amplitude and phase as the frequency of the underlying VB oscillation is varied.

The amplitude heavily depends on the start time, as this dictates the relative phases of the modulated components, and how they combine. On average, it is around half the maximum initial amplitude, however if the start time causes the segments to align, there is a large increase in amplitude as they constructively interfere. By starting at an integer number of $g-2$ periods, the amplitude is expected to be in the half maximum region. The phase shifts linearly between $\pm$ the original phase of the non-modulated oscillation. This information is used to perform the simultaneous fit in Run 3B, shown in Figure 7.16.

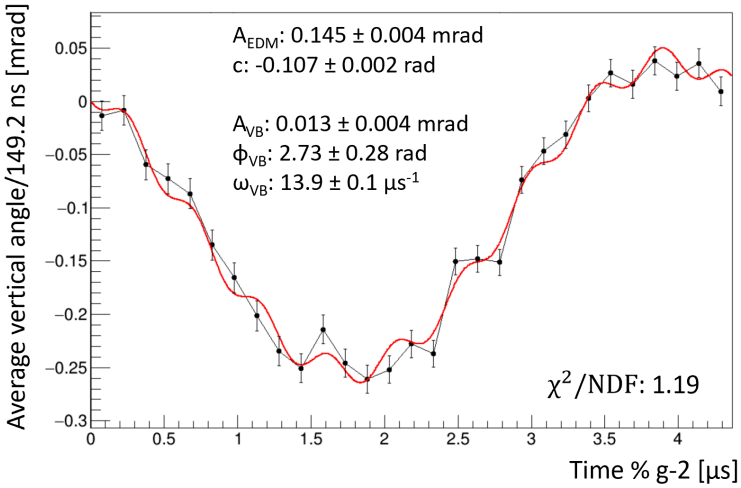

Figure 7.16: A simultaneous fit to the EDM and VB oscillations in Run 3B, using the sum of Equations 7.2 and 7.3 as the fit function. The fit quality is good, with expected fit parameters extracted successfully.

This fit agrees with the toy model from Figure 7.14, with a frequency consistent with the non-modulo fit and the amplitude reduced. The EDM parameters are also consistent with the simpler fit, suggesting the new term has absorbed the extra oscillation. This is confirmed by performing a FFT, shown in Figure 7.17.

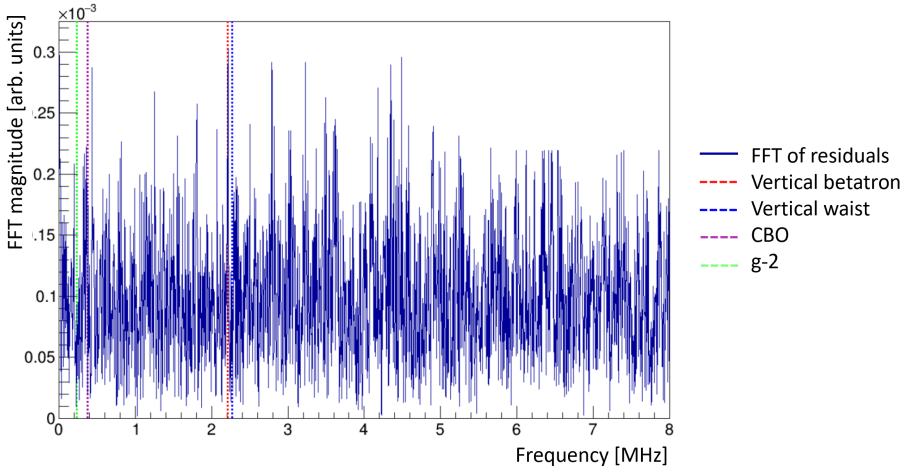

Figure 7.17: The FFT of the residual after fitting the EDM and VB simultaneously. No large peaks are seen, with several common frequencies marked with dashed lines.

To select the optimal method, the impact on the EDM fit parameters and the fit quality are compared, shown in Table 7.2. For both methods, the only changes seen are within the uncertainties of the parameters, so therefore negligible.

Randomistion	Parameters	Simultaneous fit	Parameters
χ^2/NDF	0.75	χ^2/NDF	1.19
NDF	4	NDF	8
A_{EDM}	0.146 ± 0.003 mrad	A_{EDM}	0.145 ± 0.003 mrad
A_{Bz}	0.0028 ± 0.0035 mrad	A_{Bz}	0.0028 ± 0.0035 mrad
c	-0.106 ± 0.002 mrad	c	-0.107 ± 0.002 mrad

Table 7.2: Comparison of fit parameters between the two methods of dealing with the VB, for Run 3B. They both give very similar values, with any differences consistent within uncertainties.

Comparing the quality of fit, small differences are seen, shown in Figure 7.18. Since the highly-correlated parameters make the fit method difficult to apply, with no significant improvement, randomisation is chosen as the superior method for Run 2/3.

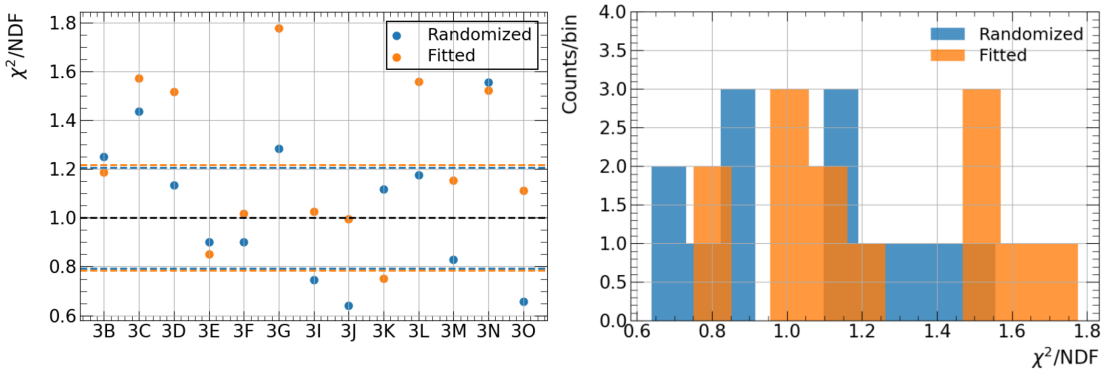

Figure 7.18: A comparison of fit quality across Run 3 between the two methods. Both perform well.

Fast rotation

The fast rotation frequency is higher than ω_a, so may be present in the modulated data, even though the choice of start time is expected to remove the worst of the impact. To remove any residual oscillation, the decay time is randomised again, shifting it by $\pm T_c/2$ and binning the vertical angle plot at the cyclotron period. As shown in Figure 3.9 for Run 3B, only a small improvement is seen, likely due to the muon beam being well-distributed around the ring at the time of the fit.

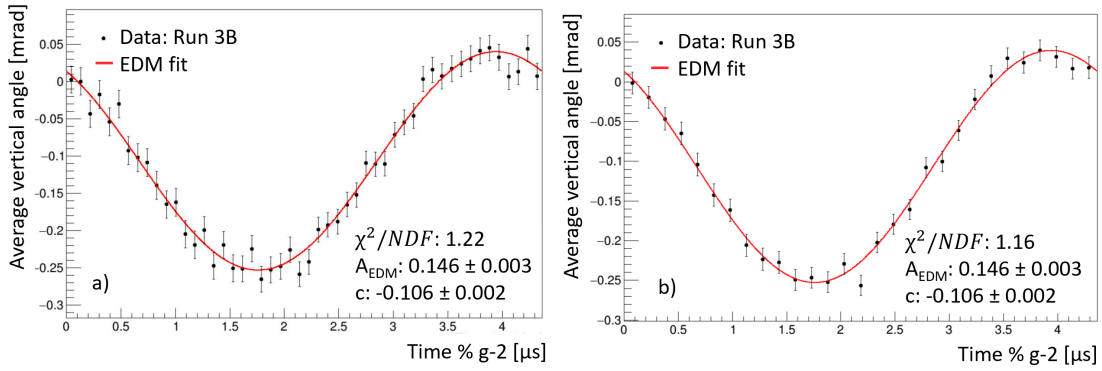

Figure 7.19: The impact of randomising out the fast rotation and binning the EDM modulo plot at the cyclotron frequency, with plot a) showing the fit before the process, and b) showing it after. A small improvement in the fit quality is seen, with parameters unchanged.

7.2.5 Early-time effects

As the muon beam evolves over time, there is potential for differences between early and late times in the fill. The majority of these effects can be removed with the choice of fit start time, but it is also important to check for any slow effects

which persist beyond this. In the Run 1 analysis, early time effects were negligible compared to the bad resistor impact, so Run 2/3 is the first opportunity to look for other required corrections.

First, the average vertical angle can be plotted as a function of time without the time modulation, for both Station 12 and Station 18 separately to check for any differences. This is necessary as the tracker stations measure different parts of the beam, shown in Figure 7.20.

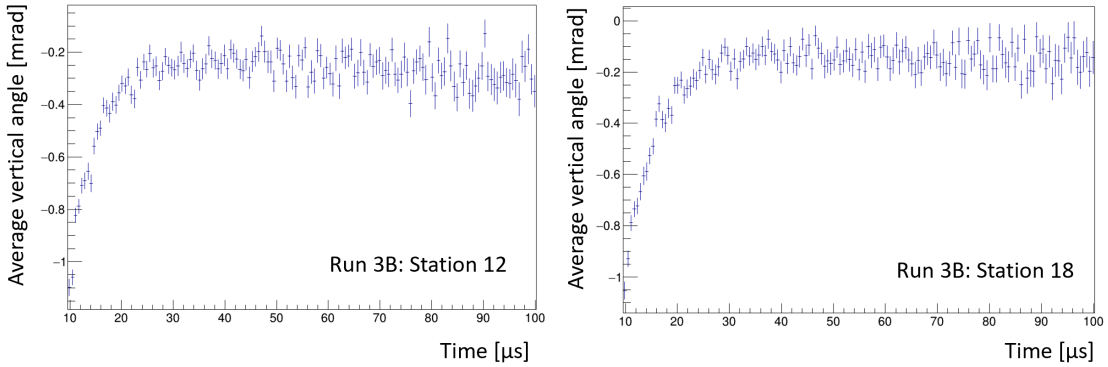

Figure 7.20: Non-modulo plots of the average vertical angle as a function of time, between 10 and 100 μs. A slow rise behaviour is seen.

There is a rise in the average vertical angle between 0 and about 80 μs, which extends beyond the 30 μs point. This was also hinted at by the peak at zero seen in the FFTs. This rise would be difficult to see modulo the $g-2$ period, as the rate of change is so small, but nonetheless will be present. To quantify the full extent, a simple straight line fit is performed on the plots in Figure 7.20, using a start time varying between 25 μs and 200 μs, with a fixed end time of 600 μs. The gradient is plotted as a function of time. If the effect decays away at late times, the gradient will reach zero, whereas if the effect persists it will remain. All datasets of Run 2 and Run 3 are combined in Figure 7.21, with the two stations plotted separately.

In both runs, the effect dies off by 100μs. The gradient is plotted per dataset to investigate any changes over a longer timescale. As Figure 7.22 shows, all the datasets are consistent with each other, indicating that this is a constant effect. This rules out some potential causes such as the radial field.

The gradient is also plotted as a function of the track momentum, as any momentum dependence would help identify the cause. However, Figure 7.23 shows that there is no momentum dependence other than some variation at very high and very low momenta. These are likely artefacts of the momentum acceptance of the tracker.

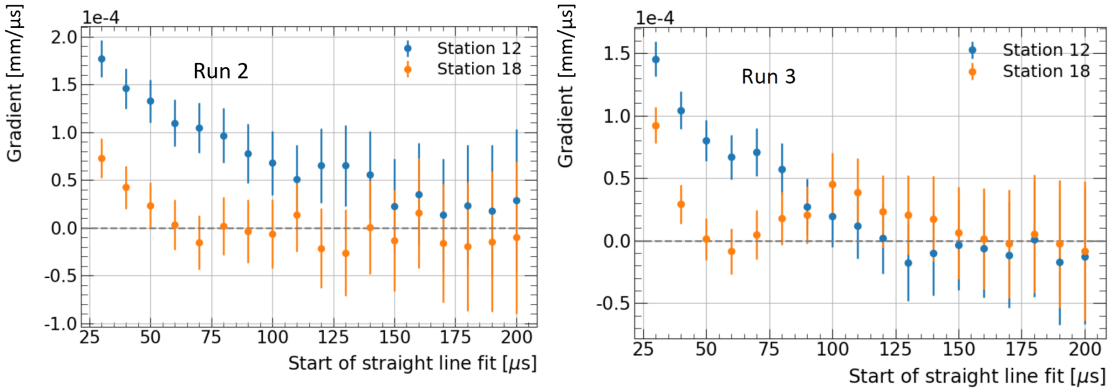

Figure 7.21: The gradient of a line fitted between the start time on the x-axis, and 600 μs. The gradient is found to tend towards zero, indicating that this is an early time effect and not a constant drift.

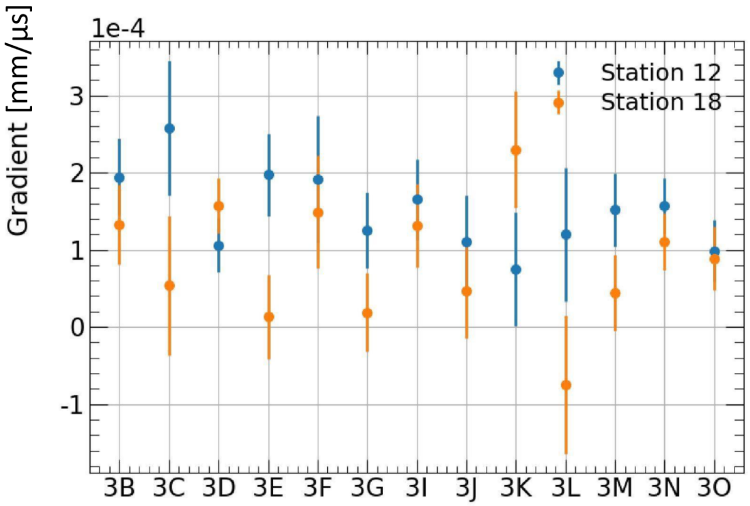

Figure 7.22: The run dependence of the gradient fit, fitted between 30 and 600 μs. All values are consistent, with no trend seen, which rules out causes like the drifting radial field.

Calorimeter data can be looked at to determine whether this is a beam motion or tracker-related effect. Run 3B is used for comparison, with the vertical betatron in the calorimeter data randomised out, and all 24 calorimeters combined and averaged to remove the the closed orbit distortion. Although the calorimeters do not have the same vertical resolution as the trackers, leading to larger uncertainties, all points are consistent with zero, with no rise seen. This is also compared across all calorimeters individually, in case another effect was cancelled out by the averaging, and to compare the calorimeters behind the trackers. Both the time variation and the azimuthal variation are shown in Figure 7.24.

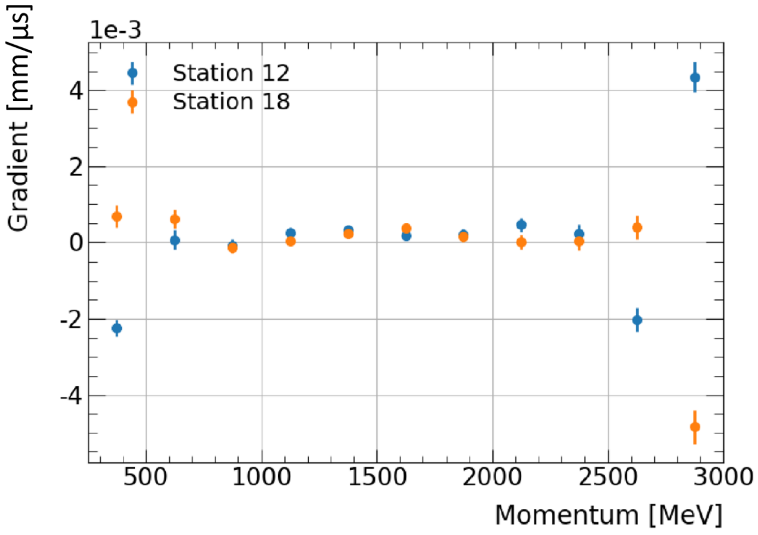

Figure 7.23: The momentum dependence of the gradient fit. The central region where fits are performed is flat, with some variation at very high and very low momenta. This behaviour does not help find the cause of the rise, since the tracker acceptance is likely impacting this at these extremes.

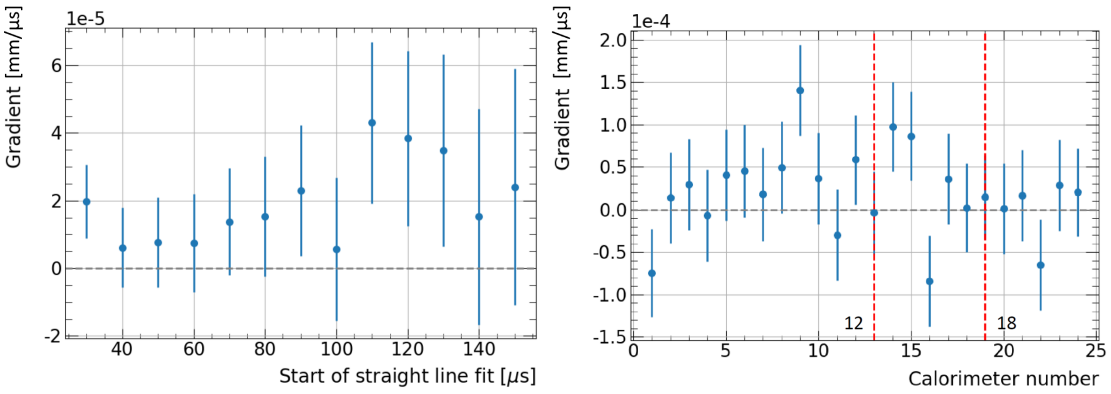

Figure 7.24: The fitted gradient for calorimeter data, as a function of time and calorimeter number, with both tracker station locations marked with dashed red lines. No evidence of the rise is seen, suggesting that the cause must be a tracker effect.

Although there is variation across calorimeters due to the closed orbit distortion, this does not change over time, so the early time rise must be from the tracker itself. A similar rise was seen in Section 4.3 with the space-charge effect, which makes the changing resolution a likely candidate. Additionally, there is significantly more rise in Station 12 than there is in Station 18, another property common to both this and the space-charge effect. It can therefore be removed by fitting for the rise up to 100 μs, then applied to data between the start time and 100 μs. This is performed for each station separately. The fit and resultant correction for Station 12 in Run 3B is shown in Figure 7.25.

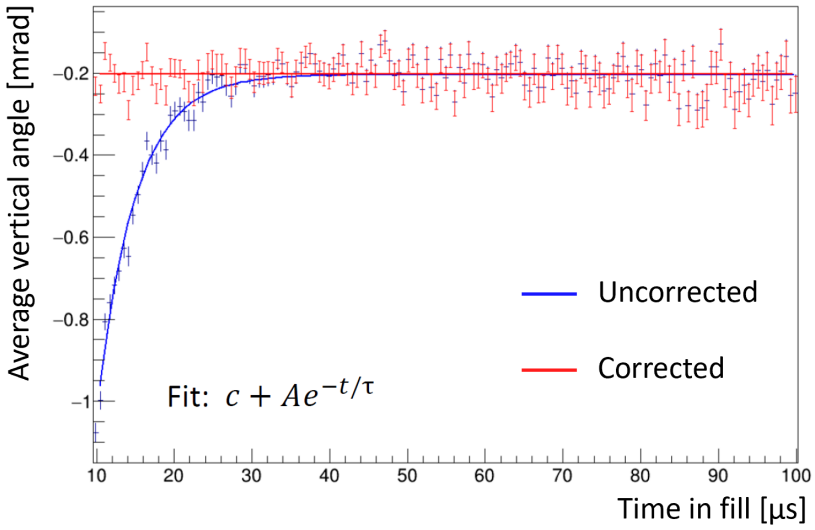

Figure 7.25: The drift in 3B Station 12 tracker data, fitted to extract the functional form, and then corrected to flatten out the vertical angle.

Finally, the impact of this correction is checked within the EDM plot itself. As shown in Figure 7.26, the EDM amplitude remains the same within errors with a small change seen in the offset, likely due to the shift upwards of the average angle.

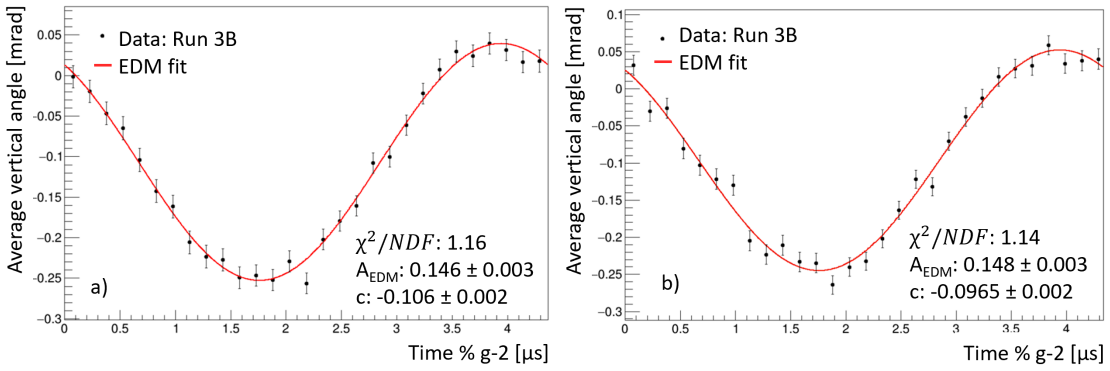

Figure 7.26: Comparison between the EDM fit with and without the early time correction. A small improvement in fit quality is seen.

Fit start time

One positive effect of the correction is that the fit start time is not limited to the region after the rise, allowing an earlier start to gain back more vertices. This must be done with caution to ensure other uncertainties from beam effects do not become significant, so the start time is reduced in integer units of the $g-2$

frequency to find the optimal starting point. As Figure 7.27 shows, pushing the start time back to 6 $g-2$ periods improves the uncertainty on the EDM by 3.4 % compared to the 30.6 μs start time, which is 7 periods. This is equivalent to the expected increase in statistical uncertainty from the extra tracks, showing that beam effects have not become dominant. However, pushing back to 5 $g-2$ periods only yields an increase compared to 6 of around 1 %, and at 4 $g-2$ periods the uncertainty becomes worse. Therefore, the optimal start time is at 6 $g-2$ periods, approximately 26.2 μs.

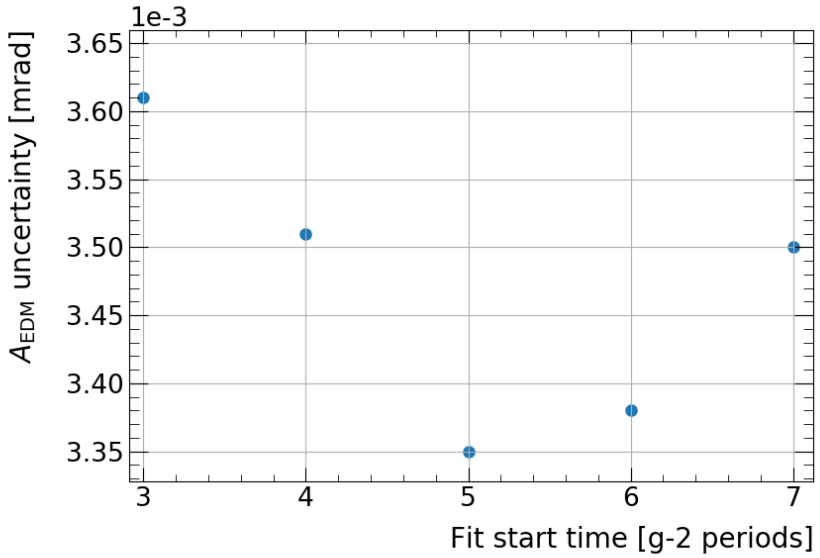

Figure 7.27: Variation of EDM uncertainty with the start time, with the early time correction applied. Pushing back to one $g-2$ period earlier gives an increase consistent with a purely statistical increase, but going any further early time beam effects start to dominate, eventually increasing the uncertainty.

It is also important to check the stability of the EDM amplitude with start time, to check for bias. The expected statistical variation σ_{diff} of a parameter between two datasets, with dataset 1 a subset of dataset 2, is given by [8]:

$$\sigma_{diff} \approx \sqrt{\sigma_2^2 - \sigma_1^2}. \tag{7.4}$$

Any drift outside of this could suggest badly fitted parameters or an uncorrected time-dependent effect. In the context of the earlier fit start time, this allows another check that the early time beam effects are not negatively impacting the fits. The results for Run 3B are shown in Figure 7.28, and are checked for all the Run 2/3 datasets. In all cases, the drift seen is consistent with the 1 sigma band.

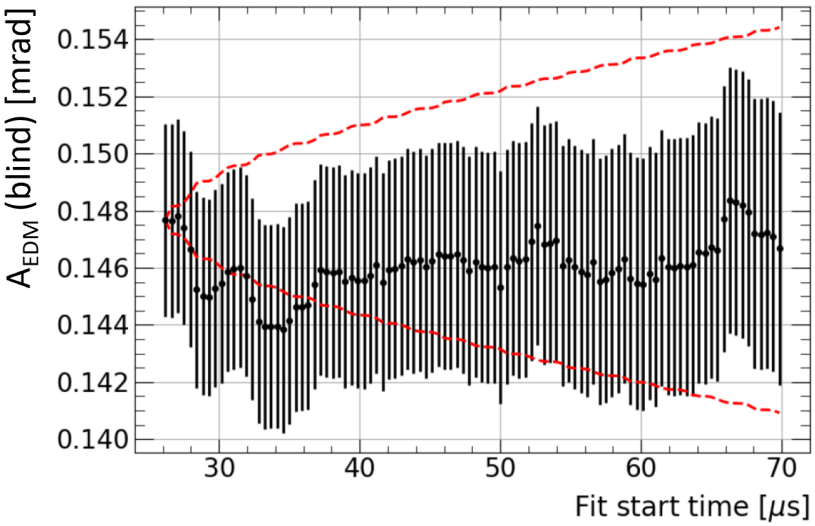

Figure 7.28: Start time scan for Run 3B, with the red band indicating the expected statistical variation relative to the first point. The parameter remains consistent with the band and does not drift away from it, showing that there are no uncorrected time dependences.

7.2.6 Momentum binned analysis

One large improvement in the Run 1 analysis is the introduction of binning the EDM fits in momentum, then fitting the amplitudes in each bin to extract the final value. Since many corrections are momentum dependent, this improves the accuracy as the true shape can be fitted better, leading to an increase in overall sensitivity. This is therefore also included for the Run 2/3 analysis, with the optimal binning for these new datasets needing determination. For the sake of easy combination, it is useful to have all datasets binned identically to each other and the MC. Run 3B is therefore an ideal test dataset as it is average in size for Run 2/3.

As the EDM amplitude is momentum-dependent, the corrections must be applied to fit for the final extracted value. This is covered in detail in Section 5.3.2, but as a starting point, only the dilution is corrected and the resulting points are fitted to study how the final statistical uncertainty varies. This momentum-binned fit is shown in Figure 7.29.

The final EDM amplitude is extracted by fitting a 0th order polynomial to the bins, with the result compared to the previous combined fit. There is a significant improvement in sensitivity, with an 11 % reduction in the error on A_{EDM}. The number of momentum bins is varied to determine the optimal binning, with Figure 7.30 showing how the EDM uncertainty varies as the number of bins is increased.

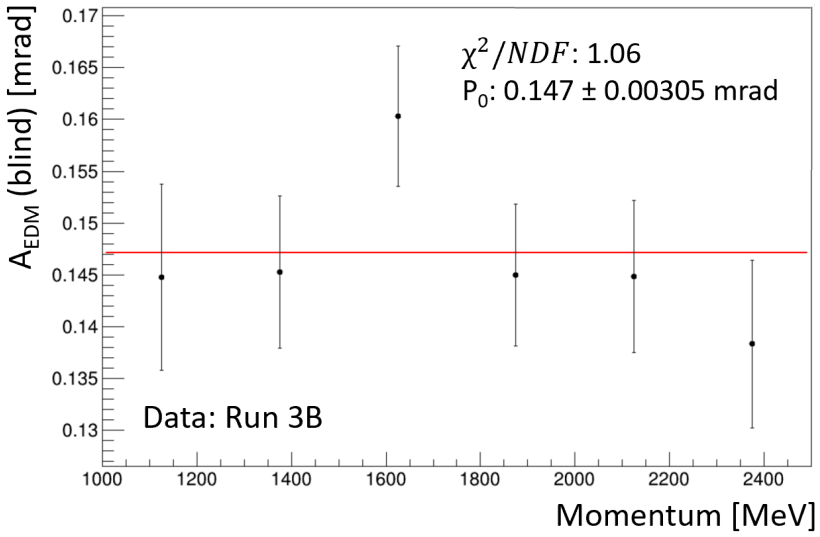

Figure 7.29: A momentum-binned extraction of the EDM amplitude for Run 3B, corrected roughly using the expected ideal momentum dependence. An 11 % increase in sensitivity is seen compared to the unbinned case, with a much better uncertainty due to applying a more accurate correction.

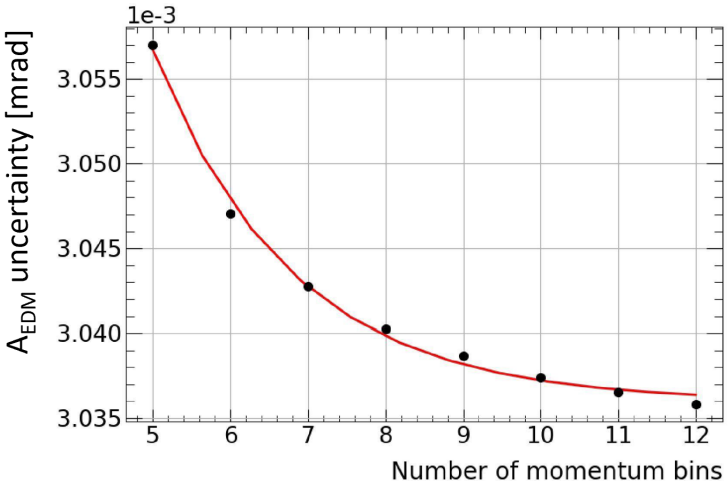

Figure 7.30: The increase in sensitivity seen as the number of momentum bins is increased, from being able to apply a more accurate correction. An exponential fit is applied in an attempt to characterise the shape, and found to be a reasonable approximation.

This shows that as the number of bins increases, so does the sensitivity, following a roughly exponential decrease in error. A fit is performed to allow estimation beyond the range scanned over. The trend is cross-checked by extrapolating back to one bin, where the uncertainty should be roughly consistent with the one-bin EDM fit. They are found to be within 0.0002 mrad of each other, an order of mag-

nitude smaller than the uncertainties themselves, demonstrating good agreement. This plot alone suggests that as many bins as possible will give the largest improvement in sensitivity, however the statistical uncertainty should also be considered. This scales as shown in Figure 7.31, with slight variations based on the dataset size shown in Appendix D. This indicates 6 or 7 bins is the optimal number. Since the MC needs to also be binned in the same way and is costly to generate, it makes sense to choose the lower of the two binnings. Run 2/3 momentum binned fits therefore all have 6 bins, of width 250 MeV across the range 1000-2500 MeV.

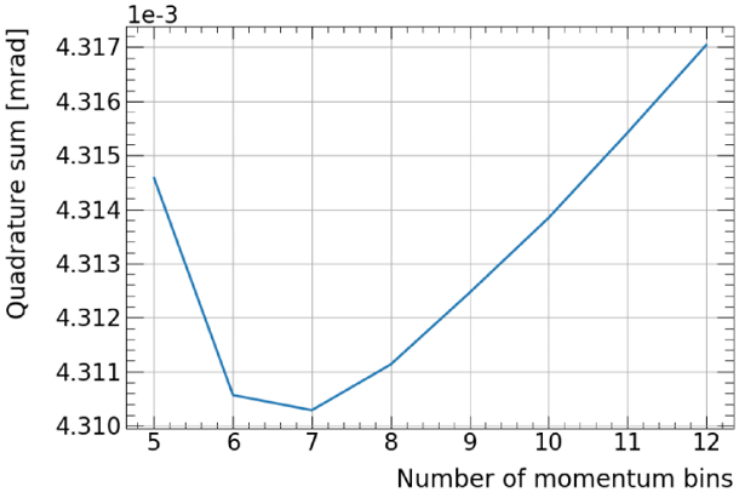

Figure 7.31: The combination of the uncertainty in Figure 7.30 and an uncertainty which scales as the square root of the statistics in the bin, for example how the drift correction would. An optimal binning is found to be around 6-8 bins.

7.2.7 The new fit function

The introduction of the momentum-binned analysis allows the quality of fit in the individual bins to be checked. As part of this process, it was found that the sinusoidal fit function does not correctly fit the oscillation in the high momentum bins. A full derivation of the expected shape by J.Price [50] revealed that the true shape of the EDM fit is not just a sinusoid, but also needs to include the number oscillation in the denominator to properly normalise it, as shown in Equation 7.5.

$$\frac{A_{g-2}\ \cos\left(\omega_a t + \phi_{g-2}\right) + A_{EDM}^{blind}\ \sin\left(\omega_a t + \phi_{g-2}\right) + c}{N(t)} \tag{7.5}$$

The momentum variation of N(t) is well understood from the ω_a analysis, as it is just the ω_a precession fitted with the 5-parameter fit from Equation 2.9, with amplitude momentum dependence as seen in Section 2.1.3. This means that for

low momentum bins, the fit is very similar to the old function due to the low asymmetry, with the high momentum bins seeing a larger change. The improvement in the highest momentum bin (2250-2500 MeV) can be seen in Figure 7.32. The shape matches significantly better, and the quality of fit is overall improved. This is used for the analysis moving forward.

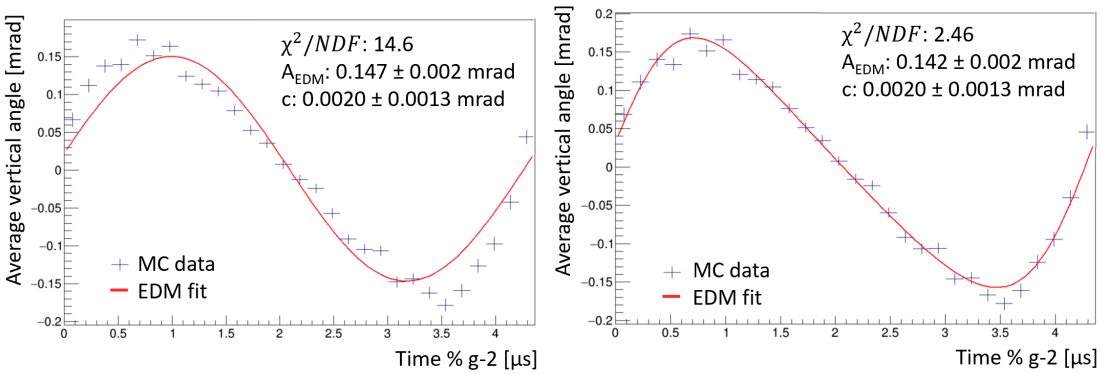

Figure 7.32: The 2250-2500 MeV momentum bin, fitted with the old fit function (left) and the new fit function (right). A clear improvement is seen with the new fit function.

Updating the blinding

The blinding requires the injection of an accurate fake signal on top of the data, so must be updated with this new fit function. Additionally, it needs to include the momentum dependence of the EDM amplitude for the momentum binned analysis. To make the blinding momentum-dependent, the expected distribution of the EDM amplitude in the lab frame from the dilution is taken from Figure 2.9. This is integrated in each bin to find the expected average, then multiplied by the blinding amplitude. The result is a momentum-dependent shape which hides the true EDM amplitude across all bins. The fake signal is then updated to use the new functional form rather than just a sinusoid. This requires the plotting and fitting of the number distribution for each momentum bin, since it defines the denominator of the function.

Both of these blinding updates are tested using MC data. Figure 7.33 shows a comparison between the unblinded and blinded MC in the 2250-2500 MeV momentum bin, chosen as this is where the shape differences are most visible. As can be seen, the EDM amplitude changes but no other fit properties or parameters are impacted.

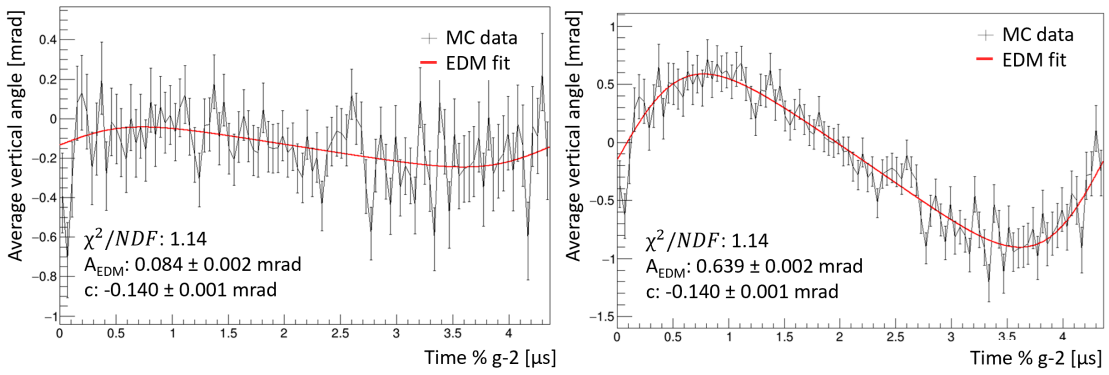

Figure 7.33: The 2250-2500 MeV momentum bin, with and without the blinding, which has been updated to use the new functional form and have momentum dependence. As before, only the EDM amplitude parameter is impacted.

7.2.8 Acceptance and dilution

Once the EDM amplitude is extracted, two key corrections must be applied to obtain the true tilt of the precession plane. The angle measured in the lab frame is reduced by both the momentum-dependent dilution effect in Figure 2.9, and the tracker acceptance, which also reduces the angle. MC data is used to evaluate both of these, with the MC scaled to better match data as described in Section 7.2.8. This sample is generated using the GM2RINGSIM Gas Gun with 8 billion events, which results in around 40 million quality tracks. An EDM signal 30× the BNL limit is injected, corresponding to a true tilt angle of 1.69 mrad.

The impact of various effects can be clearly seen in MC by comparing three different subsamples of data. The 'All Decays' data is a truth MC sample that includes all the positions generated, regardless of final trajectory. This represents what would be measured experimentally with perfect tracking. The subsample of truth data that does hit the tracker and could be reconstructed is called the 'Track Truth' sample. This removes the impact of resolution and inaccuracies in reconstruction. Finally, there is the 'Reco' sample, which is tracked in the same way as the data. This is expected to most closely match real data.

Dilution correction

The dilution function in Equation 2.35 describes the ideal case, however in practice a small correction must be included to account for radiative corrections to the muon decay. This introduces a small reduction factor to the dilution, which can be determined from simulation. The All Decays MC sample is fitted for the EDM amplitude in fine momentum bins and plotted against momentum. These points

are fitted with the theoretical ideal function from Equation 2.35 with a scale factor δ_0 to account for the reduction. The fit is shown in Figure 7.34. This is found to need a slightly larger correction than Run 1, due to the difference in fit function.

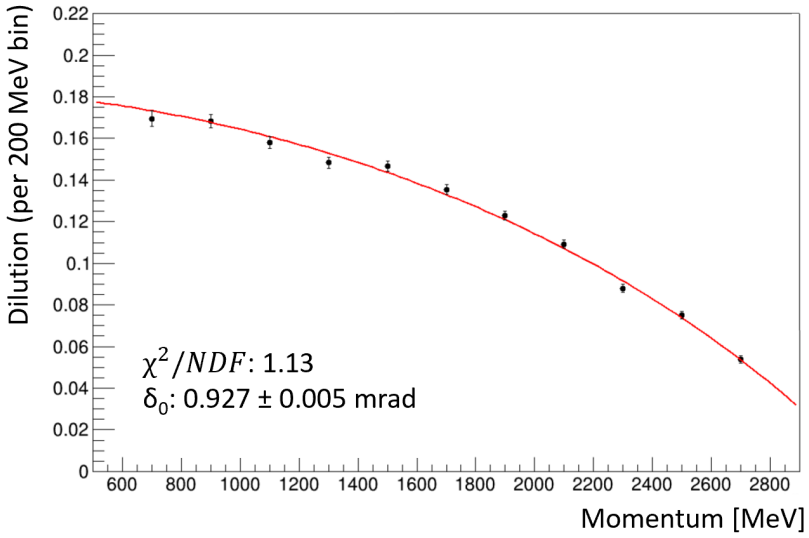

Figure 7.34: A fit of Equation 2.35 to the MC momentum dependence in the All Decays sample. A small reduction is seen from the ideal case due to radiative corrections to the muon decay, which is accounted for by including a floating constant in the fit.

Basic acceptance correction

Only a subset of positions can be reconstructed by the tracker, so a very specific area of phase space is seen when looking at tracker data. This leads to a net reduction in the measured tilt angle, which must be corrected to extract the true tilt in data. Since it is not possible to collect data without acceptance effects, this must be calculated using MC. The differences between MC and data also need to be accounted for to give an accurate correction factor. This is covered in Section 7.2.8.

A basic correction is found by comparing the EDM amplitudes between the All Decays and Reco MC samples, with the Reco having a factor of 10 fewer tracks. The ratio between the two is defined as the correction for each bin. As the uncertainty in this is purely statistical, the low events in the Reco sample is the limiting factor in the precision. The differences between the All Decays, Track Truth and Reco samples are shown in Figure 7.35.

The combined acceptance and dilution corrections should give an extracted tilt consistent with the tilt injected when the MC was generated. This is shown in

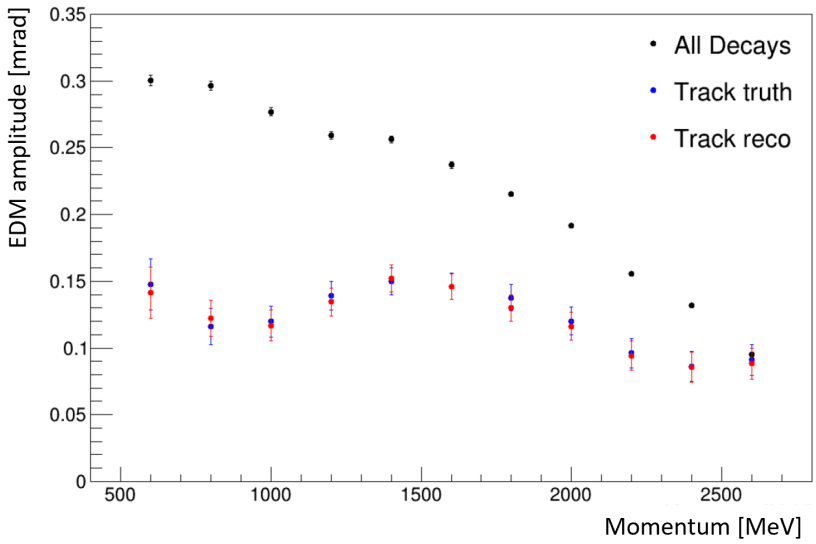

Figure 7.35: A comparison of the fitted EDM amplitudes across the three MC samples. Truth and Reco are consistent, which is good to see for the accuracy of reconstruction, with the All Decays much higher due to the tracker acceptance.

Figure 7.36, where the uncertainties are a combination in quadrature of the statistical uncertainties from the fits themselves and the uncertainty in the correction ratio. The input tilt of 1.69 mrad is successfully extracted.

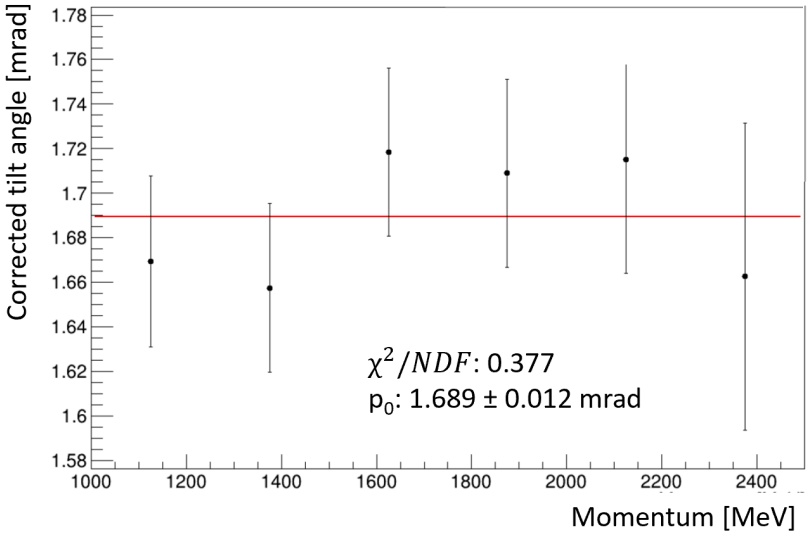

Figure 7.36: The corrected tilt angle, using the dilution correction and the simple acceptance correction. The injected tilt of 1.69 mrad is successfully extracted.

Acceptance map method

The comparison of the All decays and Reco fit results from MC gives a usable result for the acceptance correction, and is the method used for Run 1. However, due to the large uncertainty, this makes the acceptance the dominant uncertainty given the improvements made in measuring the radial field. Therefore, it is important to try and reduce this.

The total acceptance uncertainty is a combination of the statistical uncertainties of the datasets used to find the ratio, with the All decays uncertainty being about five times smaller than the Truth/Reco. As a result, if it would be possible to use just the All Decays sample to evaluate the correction, the uncertainty would be significantly improved. This motivates the implementation of a new scaling for the All Decays sample, using maps of the acceptance. A 2D histogram of beam vertical position against vertical angle is made for both the track Truth and the All Decays samples, an example of which is shown in Figure 7.37. While the All Decays sample uniformly covers the full beam region, the track Truth sample only includes decays from certain parts. To make an acceptance map, the ratio of these two histograms is found, and normalised based on the highest bin. This map is then applied as a probability of acceptance to the All Decays data. The expected reduction in uncertainty comes from the normalisation, which allows the map to apply the acceptance shape, with approximately 25 $\times$ the statistics.

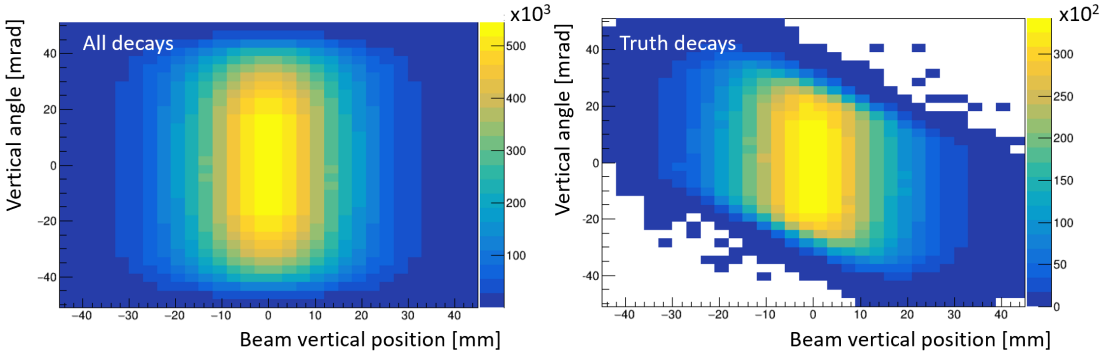

Figure 7.37: A comparison of the distribution of beam vertical position and vertical decay angle between the All Decays sample and the track Truth sample, showing the effect of tracker acceptance.

The beam radial and longitudinal position are also considered, however neither show a strong correlation with the vertical angle. Since the tracker acceptance is momentum-dependent, generating a single map is not sufficient to accurately apply the effect. A map is generated for each of the six bins used in the EDM analysis. The resultant maps can be seen in Figure 7.38.

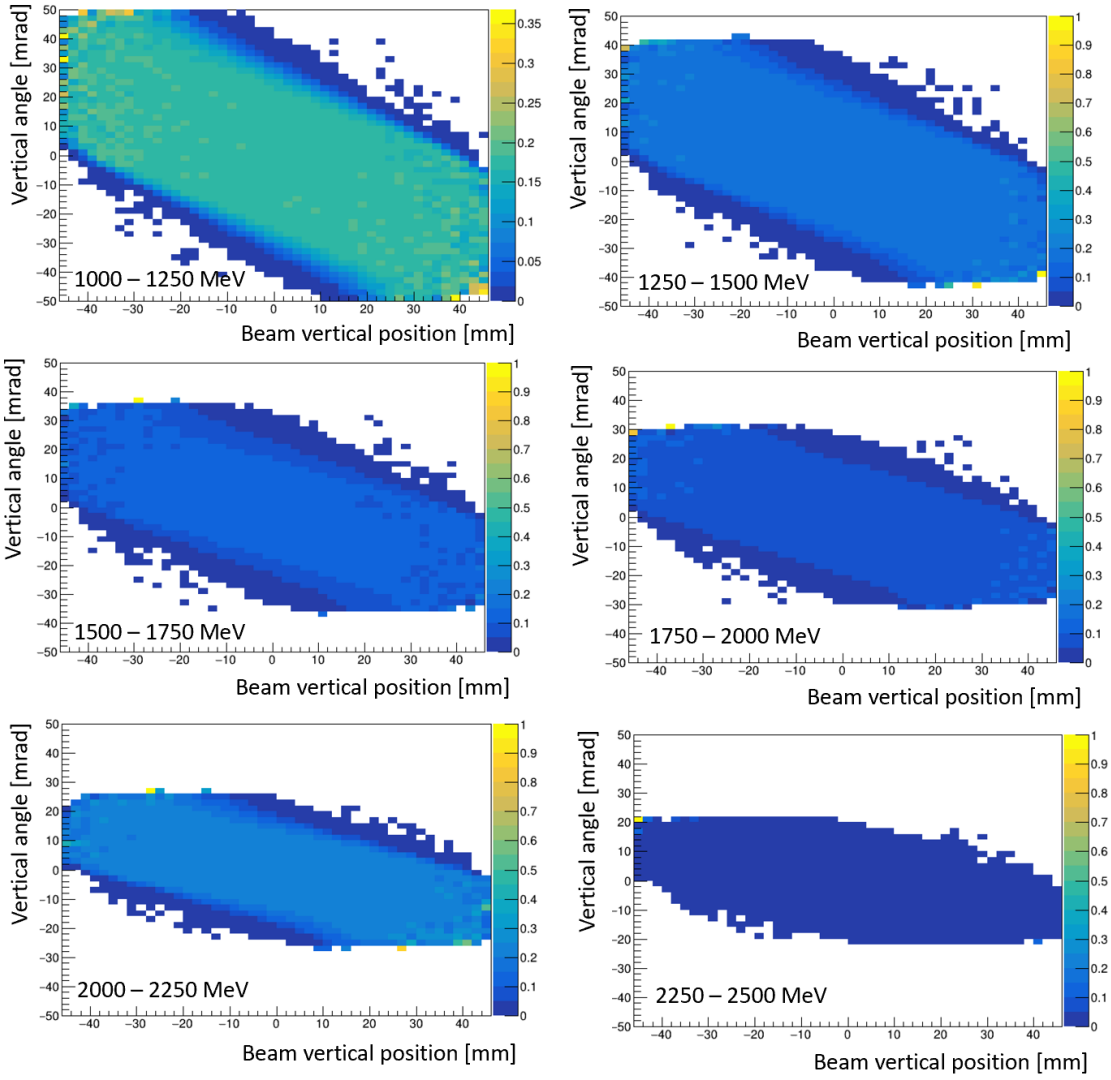

Figure 7.38: The acceptance maps generated by finding the ratio between the All Decays and Truth sample of vertices.

The track momentum is used to select which of the maps from Figure 7.38 to use, and then the probability of retention is determined from the 2D bin the track belongs to. This is compared to a uniformly-generated random number in the range 0 to 1, and the track is discarded if the number is larger than the probability. The data after this process is fitted and plotted to determine how successful the scaling has been. Figure 7.39 shows that even with this simple binning, the impact of tracker acceptance is well approximated. The uncertainties are evaluated and compared to the basic method from Section 7.2.8, with results shown in Table 7.3.

Despite the success in matching, the uncertainties are roughly equal to or larger than the original uncertainties. The reason is that despite the statistical increase, there is an approximation being made where each track is placed in a

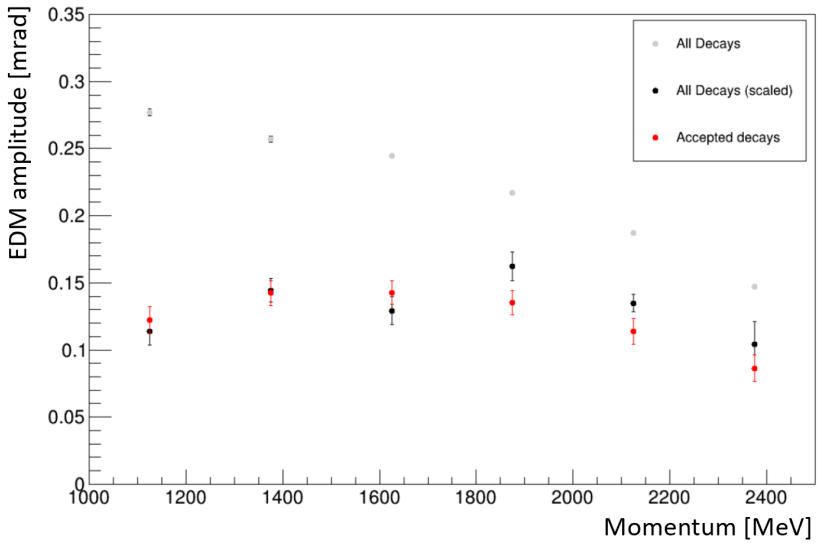

Figure 7.39: A comparison between the original All Decays amplitudes, the scaled All Decays amplitudes, and the Truth amplitudes the scaling is aiming to mimic. Good agreement is seen even with this simple binning.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
Old unc. [mrad]	0.037	0.037	0.037	0.041	0.050	0.069
Maps unc. [mrad]	0.036	0.035	0.044	0.049	0.035	0.115

Table 7.3: Comparison between the uncertainites from the simple ratio method and the acceptance map method.

fairly wide bin, introducing an extra uncertainty based on how well the momentum is approximated by the bin average. This is overriding any improvement from the increased statistics. In order to improve further, the determination of the probability must be more accurate. Since binning more finely is difficult with the low stats of the MC, interpolation is a more feasible option. To inform the interpolation, a rough uncertainty is estimated by manually shifting all results over by a bin in both vertical angle and vertical position, and comparing the difference seen. The vertical angle is found to have a two orders of magnitude larger impact than the vertical position. A linear interpolation of the maps is implemented in 2D using the Delauney triangle method [91]. This generates a set of triangles in 3D space linking the points, which collectively define a plane which is sampled from. Each point is guaranteed to not be within the circumcircle of any other point's triangles, which maximises the angles, keeping the interpolated surface smooth. One of the interpolated maps achieved in this way is shown in Figure 7.40. This works well for the majority of the map, but struggles at the edges, where the ratios are less stable.

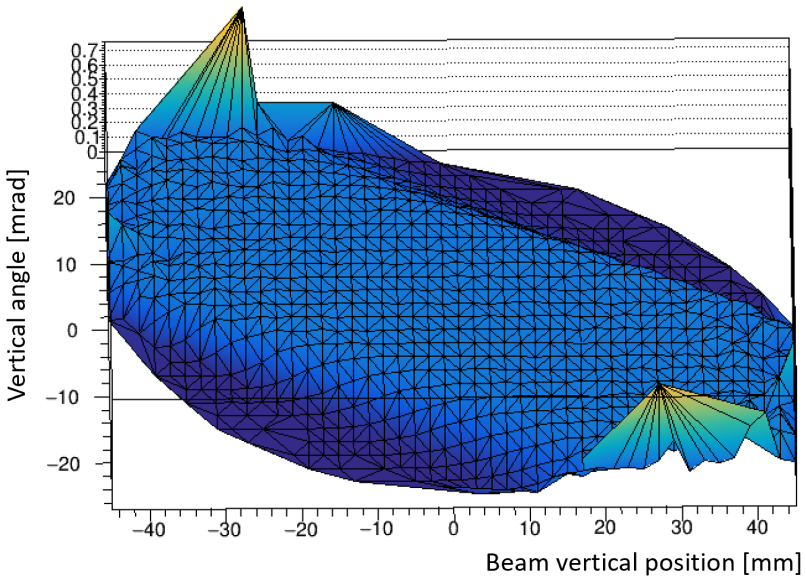

Figure 7.40: The Delauney triangulation of the 2000-2250 MeV acceptance map. A good smooth surface is achieved for most of the map, with some bins around the edges being outliers leading to spikes. Since it is very unlikely for a decay to end in these regions, the spikes do not impact the results, as shown in Figure 7.41.

While the probability of tracks lying in the edge regions is low, it is important to check the accuracy of the ratios. The results of applying this interpolated map to the All Decays data is shown in Figure 7.41, with uncertainties in Table 7.4.

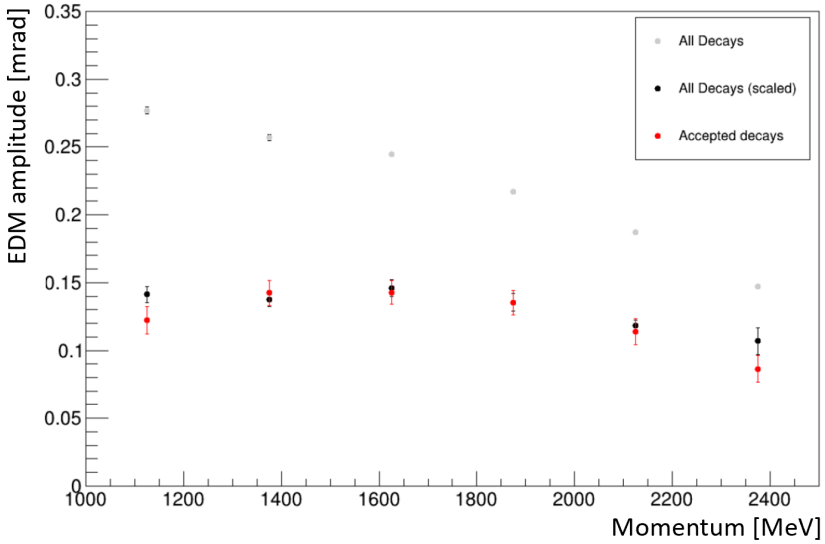

Figure 7.41: Comparison of the scaled and true EDM amplitudes, showing a better agreement with the interpolation added.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
Old unc. [mrad]	0.037	0.037	0.037	0.041	0.050	0.069
Maps unc. [mrad]	0.022	0.022	0.026	0.029	0.021	0.068

Table 7.4: Comparison between the uncertainties from the simple ratio method and the acceptance map method. An improvement is now seen in the size of the uncertainties.

A clear improvement is seen compared to the previous method, though this is not as large as the statistical increase would suggest, as there are still interpolation errors involved. To minimise these, the optimal binning is investigated. This needs to be optimised considering both the uncertainty and accuracy of the corrections.

Considering accuracy first, this is quantified by comparing the average difference between the scaled All Decays amplitude and the Truth amplitude to the Reco uncertainty. If the average difference is above the Reco uncertainty, the results of that binning are considered inconsistent within reconstruction errors, and therefore not accurate. Bins were varied between 5 and 140 bins in both vertical angle and vertical beam position, as shown in Figure 7.42. The accuracy is worst at bin extremes, with the low binned maps having insufficient resolution to accurately apply the tracker acceptance, and the high binned maps having too low statistics per bin. Based on this, a feasible range of binning between 40 and 100 bins is identified.

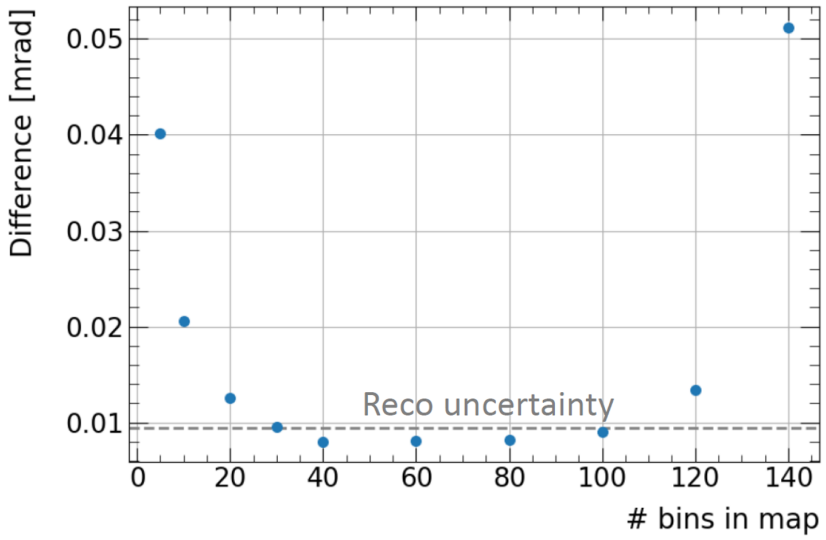

Figure 7.42: The difference between the scaled All Decays and Truth EDM amplitude, for different binnings of the acceptance maps. The maps considered sufficiently accurate are those with a difference below the reconstruction uncertainty.

The uncertainties must also be compared, with the averages shown in Figure 7.43. In general, the smaller the number of bins, the better the uncertainty, indicating that the dominant effect is the statistics per bin. Therefore, the optimal binning is the lowest binning in the feasible range defined by Figure 7.42, so 40 bins are used.

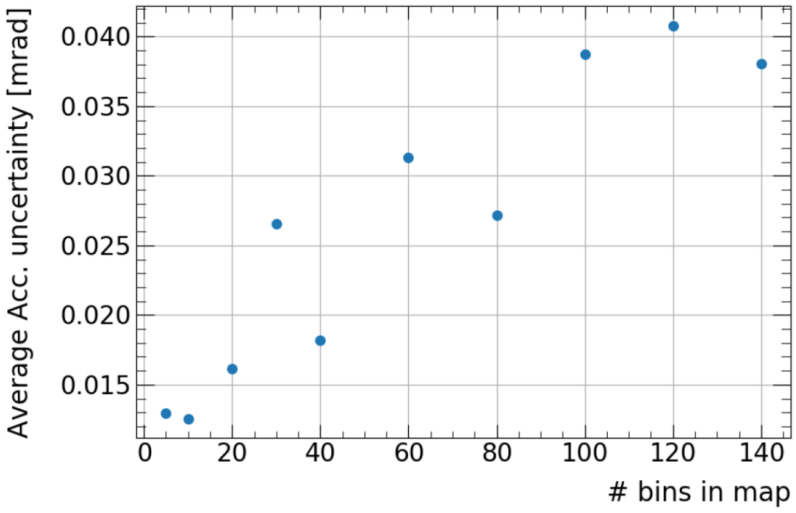

Figure 7.43: The average bin uncertainty for the acceptance correction, plotted as a function of the number of bins in the acceptance map used for scaling. This is dominated by the occupancy of the bins, with lower binned maps having smaller uncertainties.

The final uncertainty improvement compared to the simple acceptance correction is a factor of 3-4 across the bins, as shown in Table 7.5. This could hypothetically be improved by using a more sophisticated interpolation method, but given the relative scale of the new acceptance uncertainties to the radial field uncertainties, which are irreducible, no gain in sensitivity would result. Additionally, for Run 2/3, the statistical errors still dominate above all the systematics. The final comparison between the scaled All Decays and Truth for the 40-binned maps is shown in Figure 7.44.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
Old unc. [mrad]	0.037	0.037	0.037	0.041	0.050	0.069
Maps unc. [mrad]	0.012	0.019	0.014	0.024	0.016	0.022

Table 7.5: Comparison between the uncertainties from the simple ratio method and the acceptance map method with 40 bins. An improvement of factor 3-4 is seen across the bins.

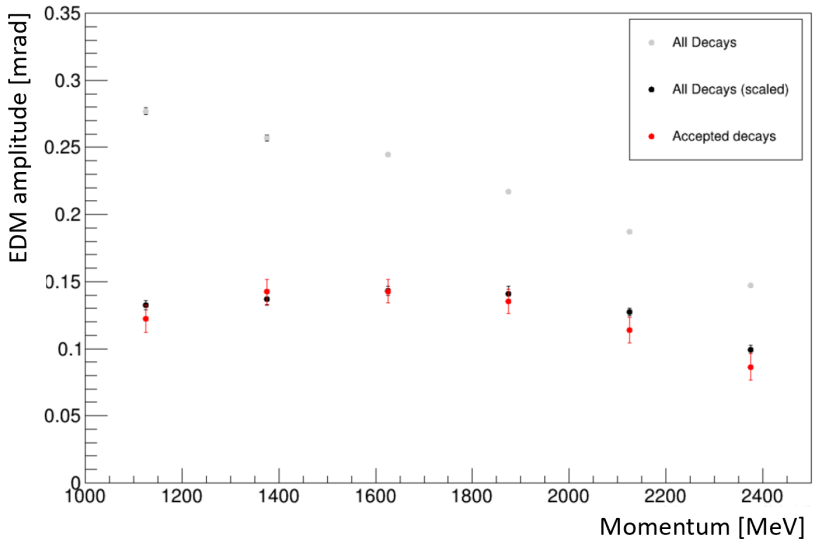

Figure 7.44: Final comparison between the scaled All Decays and Truth, for the optimally binned maps with 40 bins in both directions.

MC/data differences

Despite being a very good approximation, the MC is not perfect, which needs to be accounted for in the corrections. The largest impact comes from inconsistencies in the vertical angle distribution. The mean and RMS for each momentum bin in Run 3B and the MC are plotted in Figure 7.45, for each station individually to account for differences. In general, the width in data of the vertical angle distribution is narrower than the MC width, and the mean in MC is lower.

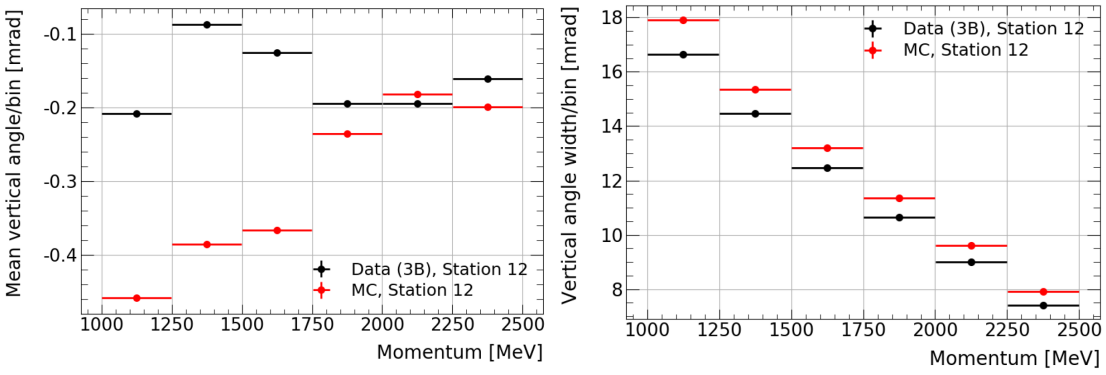

Figure 7.45: Vertical angle distribution mean (left) and width (right) differences between MC (red) and data (black) for Station 12 in Run 3B. The widths are narrower in data, with the mean value closer to zero.

Since the acceptance maps will be impacted by these differences, a weighting on the mean and width in the MC is applied to make it match data more accu-

rately. The acceptance maps are then regenerated to extract the new corrections. Run 3B is used as an example, however the width and mean of the data varies quite significantly between datasets, as shown in Figure 7.46. This process must therefore be repeated for each dataset to extract accurate corrections.

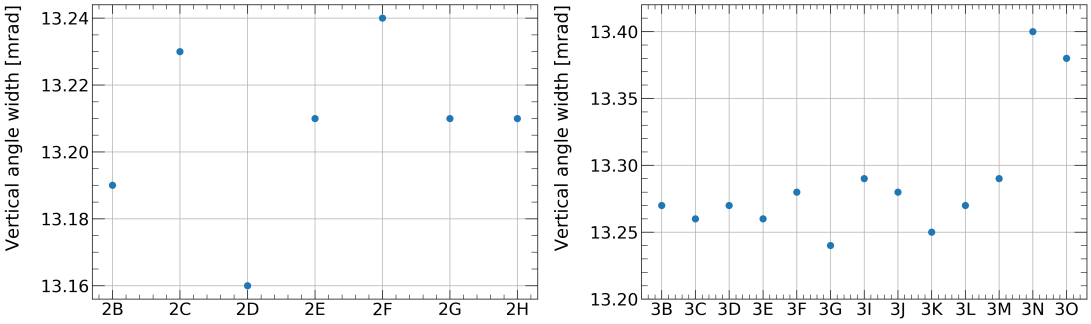

Figure 7.46: Width of the vertical angle distribution for all the Run 2/3 datasets. The large change for 3N and 3O comes from an adjustment in kicker settings, which led to a change in the beam shape.

As Figure 7.47 shows, after weighting the differences are now significantly smaller, with a perfect match of the RMS and mean being significantly closer. The slight difference in the means suggests there is an underlying extra difference in shape that is not so easily corrected. The impact of this on the EDM analysis is evaluated as a systematic uncertainty.

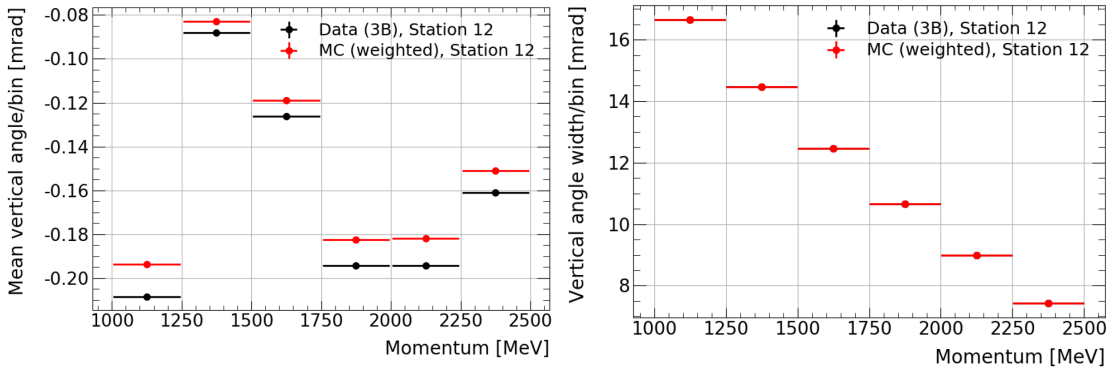

Figure 7.47: Figure 7.45 plotted again after a weighting based on data is applied. The widths now match exactly and the means have shifted to be closer.

Using the weighted MC, new acceptance maps are calculated, as shown in Figure 7.48 for 1000-1250 MeV. A large impact is seen here due to the weighted map being both more central in the storage region and narrower in width than the unweighted map, which improves the overall acceptance. This leads to slightly smaller uncertainties on the acceptance correction as a result.

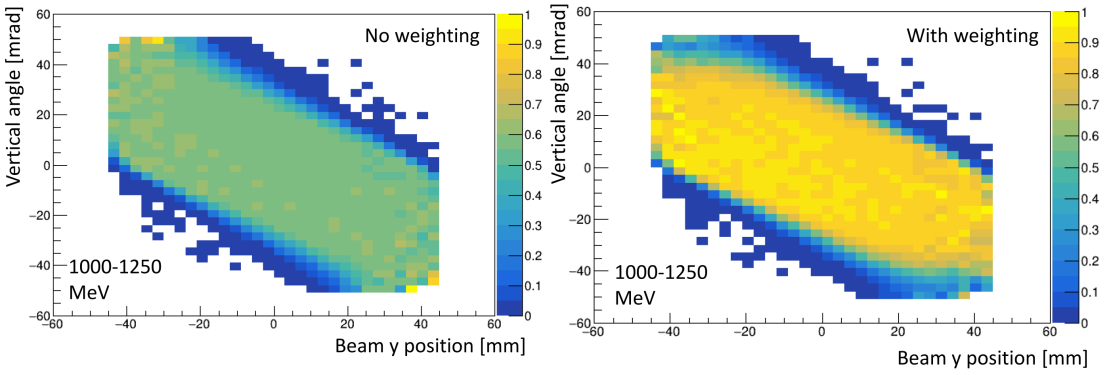

Figure 7.48: Acceptance maps for the 1000-1250 MeV momentum bin, with the left map being the unweighted MC and the right the MC after weighting. The weighted map is significantly less flat, with better overall acceptance.

Table 7.6 shows the impact of applying these new maps to the All Decays sample, and using it to recalculate the correction with the weighting for Run 3B. As expected, the uncertainties are improved due to the increase in statistics. The correction factors themselves also change substantially, as shown in Table 7.7. This shows that the differences between MC and data do need to be accounted for.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
Old unc. [mrad]	0.037	0.037	0.037	0.041	0.050	0.069
Maps unc. [mrad]	0.010	0.010	0.010	0.011	0.012	0.018

Table 7.6: Comparison between uncertainties on the amplitude in the ratio method and the maps method, with the All Decays sample scaled to match 3B data. A small improvement is seen with the data scaling applied due to more centred beam.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
Unweighted.	0.478	0.532	0.584	0.650	0.680	0.676
Weighted	0.407	0.499	0.509	0.569	0.599	0.634

Table 7.7: Acceptance correction factors, as calculated with unweighted and weighted maps to match Run 3B. A non-trivial difference is seen, motivating the inclusion of this weighting.

7.2.9 Systematic uncertainties

As well as the acceptance correction and radial field, there are other systematic uncertainties which need to be evaluated for the EDM analysis. This section covers how they are estimated and discusses their relative impacts, using Run 3B

to illustrate. A summary of the systematic uncertainties for all runs is located in Section 7.3.

Phase uncertainty

The first systematic considered is the impact of the ω_a phase uncertainty. This is evaluated by manually shifting the phase up and down by the uncertainty on ϕ_{g-2} in the vertical angle fits, taking the difference between the two as the uncertainty. The results are shown in Table 7.8, showing that this has a very small impact on the the EDM fit. This demonstrates that the 5-parameter wiggle fit in Figure 7.6 is sufficient for the extraction of this parameter.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.010 μrad	0.003 μrad	0.028 μrad	0.0009 μrad	0.031 μrad	0.014 μrad

Table 7.8: The impact of shifting the phase within its uncertainties on the fitted EDM amplitude, in μrad. This is a very small effect.

Dilution

The dilution is applied as a correction to find the final tilt angle, so any uncertainty on the scaling factor must be propagated through to the tilt. This is evaluated by generating 1000 random values for the dilution scaling factor, sampled from a Gaussian distribution with the mean set to the value from the fit in Figure 7.34 and the width set to the uncertainty. Each randomly sampled factor is used to correct the momentum binned values, with the final extracted tilt plotted for each momentum bin. The systematic is the error on the mean of the resultant distributions. An example plot for the 1000-1250 MeV bin is shown in Figure 7.49, with results for each momentum bin in Run 3B in Table 7.9.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0002 mrad	0.0003 mrad	0.0003 mrad	0.0003 mrad	0.0004 mrad	0.0004 mrad

Table 7.9: Dilution correction uncertainties for each of the momentum bins used in the analysis, in mrad. These are also very small compared to other uncertainties, so the impact on the fit is negligible.

In general, these are very small compared to the statistical and radial field uncertainties of Run 3B, so this does not have a large impact on the final result. As these errors are negligible, and as the dilution is calculated purely using MC, these are taken to be the same for all datasets in Run 2/3.

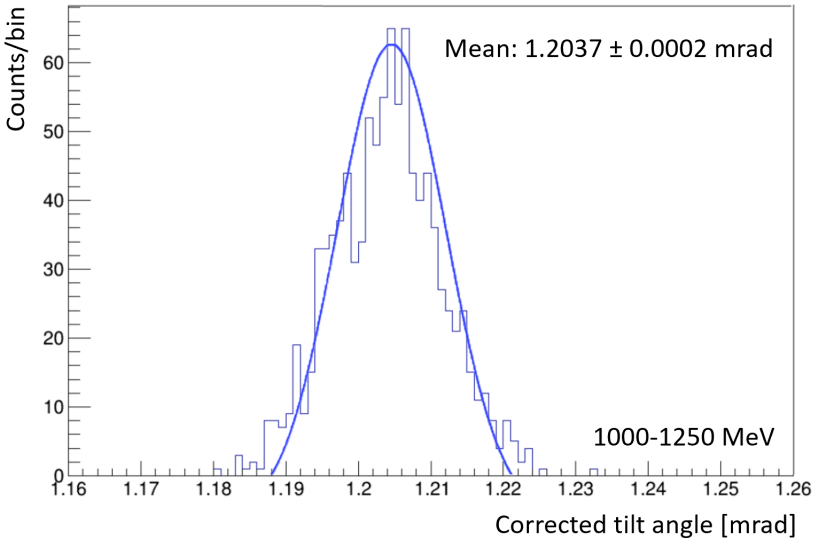

Figure 7.49: Example fit in a the 1000-1250 MeV momentum bin to extract the systematic uncertainty on the diltuion correction in that bin. The uncertainty is the error on the mean.

MC/Data differences

Given the large impact of weighting the MC, any remaining differences must be accounted for as a systematic uncertainty. After weighting, the widths match fully between MC and data, whereas the means of the vertical angle distribution do not. The differences seen in the means are shown in Figure 7.50. This is the same as Figure 7.47, without the widths plotted to make the differences more visible.

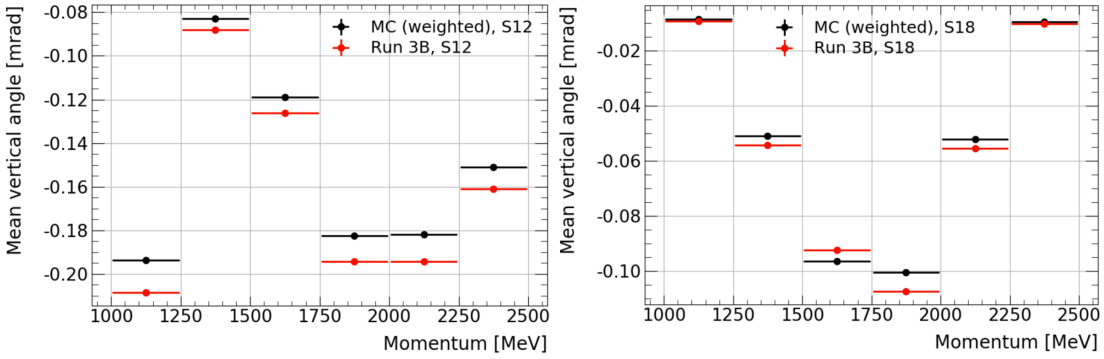

Figure 7.50: Figure 7.47 for both stations, showing the differences in the means after applying the Run 3B weighting. This is used as the metric to quantify the uncertainties introduced by the MC/data differences.

The impact is quantified using the mean difference. For each station, the mean of the distribution after scaling is shifted up and down by the difference between data and MC. New acceptance maps are calculated using the shifted data and

propagated through to the acceptance factors. The full difference between the up- and down-shift is taken as the uncertainty. The impact for Run 3B is shown in Figure 7.51, with the final uncertainties on the tilt in Table 7.10.

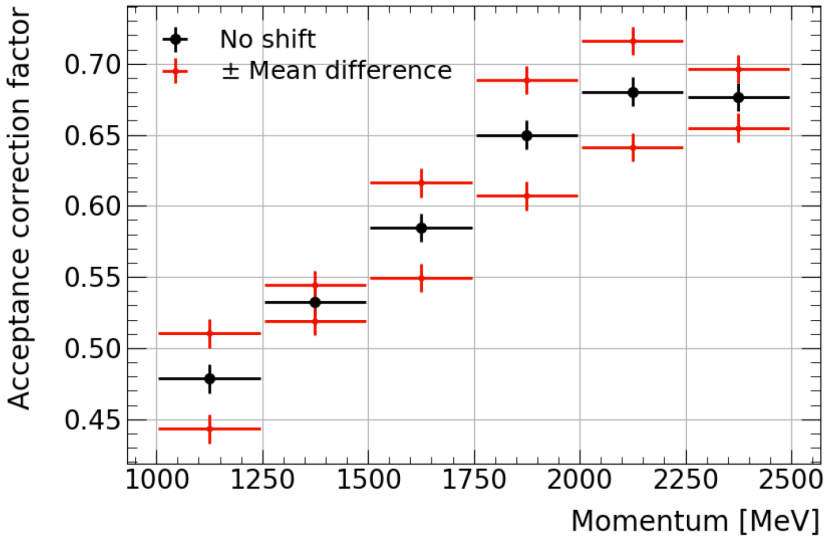

Figure 7.51: The impact on the acceptance factor of shifting the means of the vertical angle up and down by the MC/data difference in Run 3B. This is quite a large effect, and the shifts are not consistent with the unshifted value, so this must be included as a separate uncertainty.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0084 mrad	0.0032 mrad	0.0098 mrad	0.0003 mrad	0.0112 mrad	0.0003 mrad

Table 7.10: MC/data uncertainties in mrad, found by shifting the mean up/down and taking the full change in the acceptance factors as the uncertainty.

This leads to a large uncertainty, roughly comparable to the acceptance. Since the acceptance factors found with this method are not consistent within their individual statistical uncertainties, this necessitates the inclusion of this on top of the regular acceptance uncertainty.

While Run 2/3 is currently fully statistically dominated, it is likely that when analysing the full dataset, this uncertainty will need to be improved to reach the highest possible sensitivity. This could be achieved by better scaling the MC, or by looking at what other distributions vary and fixing the underlying differences in MC. In particular, the difference in the means suggests that the beam position in MC must be slightly different to data, which would be expected to impact the tracker acceptance and momentum distribution. Further work would be needed to fully identify the impacts and causes of this.

Early-time rise correction

The early-time rise correction is based on a fit as discussed in Section 7.2.5, used to calculate how much correction is needed. This includes three parameters: the amplitude of the rise, the lifetime of the rise, and the offset, all of which have associated uncertainties. These are propagated through to the EDM amplitude by shifting each parameter up and down by the uncertainty. This gives a very conservative estimate, as the amplitude and lifetime are correlated. This is performed for each station separately, shifting both either up or down to maximise the effect. The impact for all six momentum bins in Run 3B is shown in Table 7.11.

[MeV]	1000-1250	1250-1500	1500-1750	1750-2000	2000-2250	2250-2500
A [μrad]	0.036	0.052	0.022	0.042	0.022	0.042
τ [μrad]	0.11	0.12	0.096	0.11	0.10	0.12
c [μrad]	0.028	0.026	0.028	0.030	0.030	0.034

Table 7.11: The uncertainties from the early time rise correction, in μrad, for each fit parameter used for the correction.

The uncertainty on the lifetime is the dominant effect, but all three parameters are combined to give the final uncertainty in each momentum bin, shown in Table 7.12.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.115 μrad	0.135 μrad	0.102 μrad	0.123 μrad	0.107 μrad	0.130 μrad

Table 7.12: The uncertainties from the early time rise correction, in μrad. They are extremely small, orders of magnitude smaller than uncertainties like the radial field.

Since these corrections are calculated for each dataset, with slightly different results, this systematic is evaluated for each separately. All datasets are found to have a very small change in the EDM amplitude, which makes sense as the correction itself has a very small impact on the EDM fit.

Tracker vertex resolution

Although the reconstruction of tracks and vertices is shown to be very good, with similar EDM amplitudes seen for both sets of the MC, there is still an intrinsic uncertainty associated with this that needs to be accounted for. This is quantified by plotting the vertex vertical angle distributions for both Truth and Reco, subtracting the two from each other to give a 'difference' distribution. The error

on the mean of this distribution is found for each momentum bin and taken as the error on the vertex itself. The distribution for one momentum bin is shown in Figure 7.52, and the overall uncertainties for Run 3B are given in Table 7.13. Since this is very small and not expected to vary across runs, this is used for all datasets in Run 2/3.

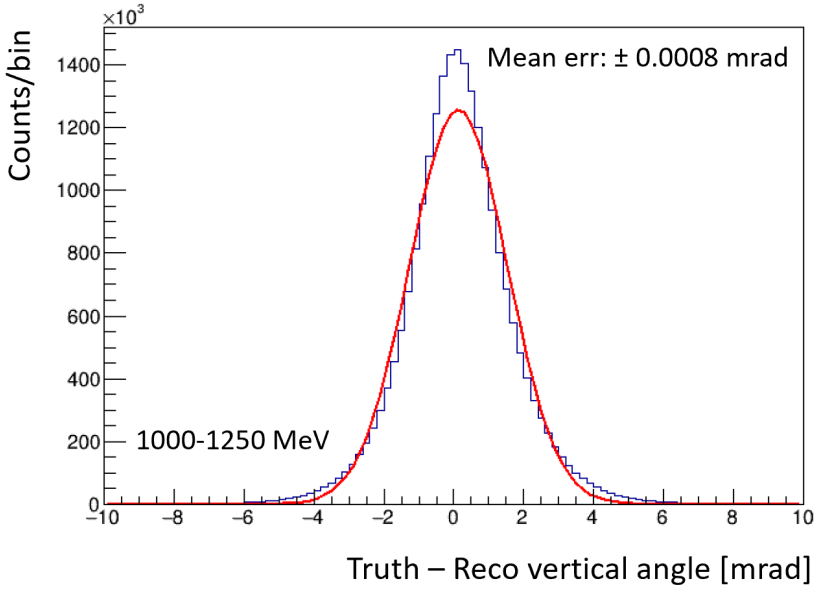

Figure 7.52: The Truth-Reco vertical angle values for quality vertices. This is fitted and the uncertainty on the mean is taken as the resolution systematic.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0008 mrad	0.0007 mrad	0.0006 mrad	0.0006 mrad	0.0007 mrad	0.0007 mrad

Table 7.13: The uncertainties from tracker vertex resolution, in mrad. They are also small.

Tracker alignment

Up to this point, the EDM analysis has assumed that the trackers are perfectly aligned with the centre of the beam in the vertical direction, whereas in reality the trackers in a module can be shifted both relative to the beam and relative to each other. The overall misalignment is referred to as the global alignment, and is determined using laser surveys, whereas the relative alignment of the trackers, known as the internal alignment, is calculated using fitted tracks and comparing the differences between true and expected hit positions [48]. For the EDM analysis, the greatest impact is expected to come from any vertical misalignment, or an angular tilt of the tracker in the YZ direction.

The high-stats MC used to determine the acceptance is run for four additional datasets, two with vertical displacements of the trackers by ± 1mm, and two with all trackers tilted by ± 0.1 degrees. These values are chosen as approximate values of the largest possible shift, allowing the shifted MC to be reused for other run periods. These are run through the methodology of the acceptance correction calculations, including recalculation of the maps, shown in Figure 7.53 for the vertical displacement shift, and Figure 7.54 for the angular tilt.

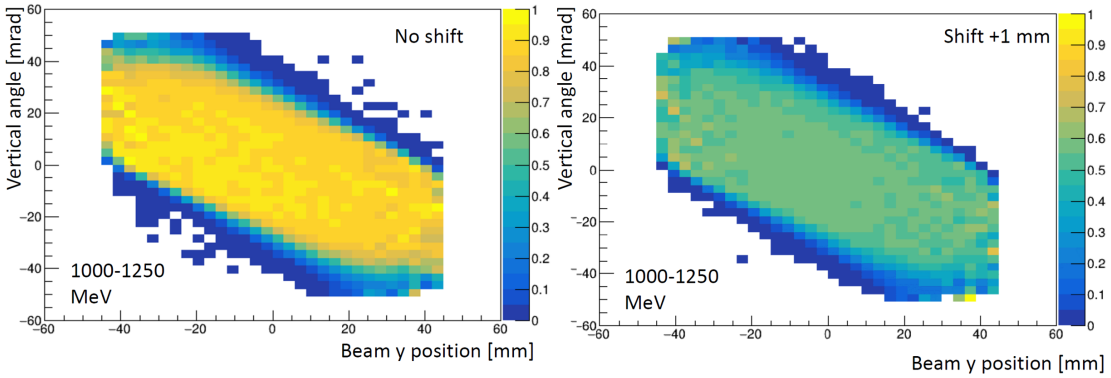

Figure 7.53: The acceptance map for the 1000-1250 MeV bin with nominal alignment compared to all the trackers shifted uniformly upwards by 1mm. The resultant map is flatter due to worse overall acceptance.

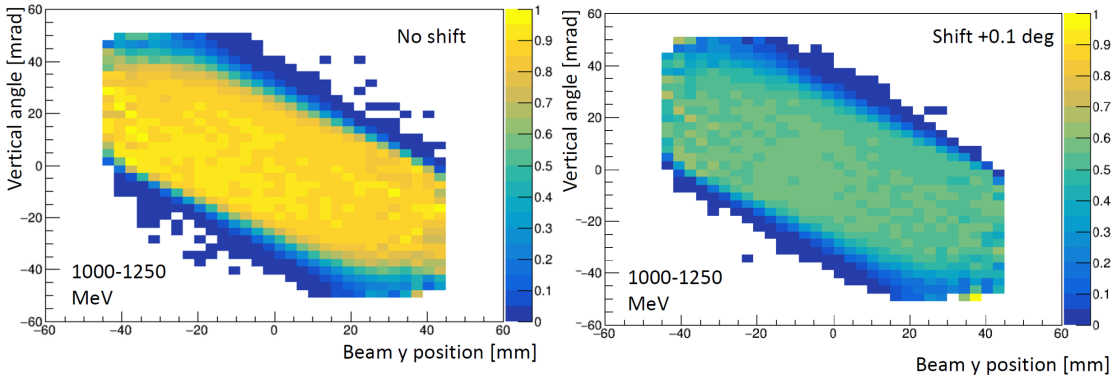

Figure 7.54: The acceptance map for the 1000-1250 MeV bin with nominal alignment compared to all the trackers tilted by 0.1 degrees in the YZ plane. Similarly to the 0.1 mm shift, acceptance is negatively impacted.

Both an angular tilt and a vertical displacement lead to flatter maps, as the measured beam is shifted away from the centre, giving an overall worse acceptance. To evaluate the impact on the EDM analysis, these maps are applied to the All Decays sample and the difference in final tilt angle is compared. The uncertainty in tracker position is estimated to be 0.6 mm in the vertical position and 0.02

degrees in the YZ tilt angle [92], so the impact is scaled down to match these smaller displacements and tilts.

The results of the 1 mm vertical shift are shown in Figure 7.55. Since the effect of shifting up/down is not symmetrical, as both worsen the acceptance, the uncertainty is taken as the largest difference between the nominal case and either the up or the down shift. The impact on the tilt angle, scaled to 0.6 mm, is shown in Table 7.14 for Run 3B.

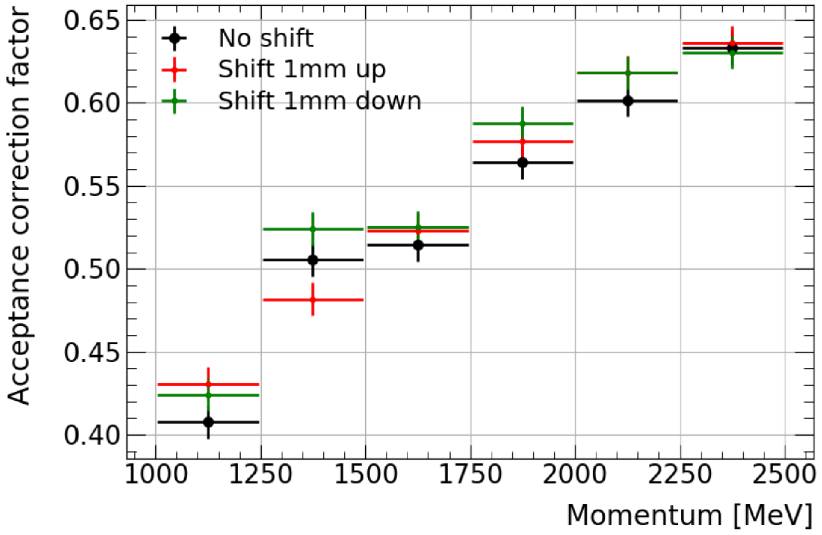

Figure 7.55: The impact of the 1 mm shift on the acceptance factors. This is no longer symmetric around the unshifted value, so the uncertainty is taken as the largest difference between the central and either up/down shifted value.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.017 mrad	0.017 mrad	0.0090 mrad	0.019 mrad	0.011 mrad	0.0060 mrad

Table 7.14: The uncertainties corresponding to a 0.6mm shift in global alignment, in mrad.

Similarly, the impact of a 0.02 degree tilt uncertainty is shown in Figure 7.56 and Table 7.15. This has a smaller contribution than the global position uncertainty as the fractional uncertainty is much smaller.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0003 mrad	0.0004 mrad	0.0006 mrad	0.0006 mrad	0.0002 mrad	0.0001 mrad

Table 7.15: The uncertainties corresponding to a 0.02 degree tilt in the YZ plane, in mrad.

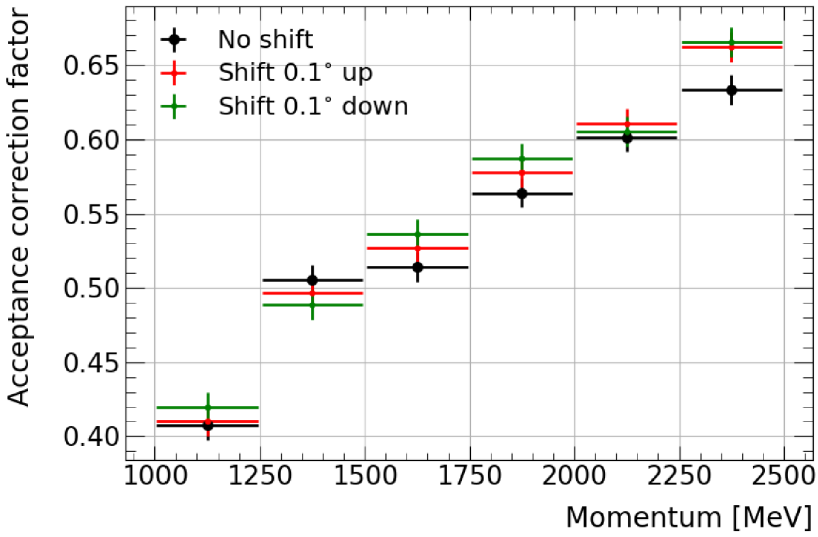

Figure 7.56: The impact of the 1 degree shift on the acceptance factors. This shows a smaller change than the displacement.

The internal alignment is known to a much higher precision than the global alignment, but must also be considered. This also contributes to the ω_a tracker systematics, so previous studies [93] of the impact on the vertex position and RMS are used. This is achieved by modifying the beam vertical position distributions in the high-stats MC by the largest change seen in the vertical mean and width to be conservative. The mean is shifted up and down by 1.711 μm, and the width is increased and decreased by 3.106 μm. The results are shown in Figure 7.57 and Tables 7.16 and 7.17.

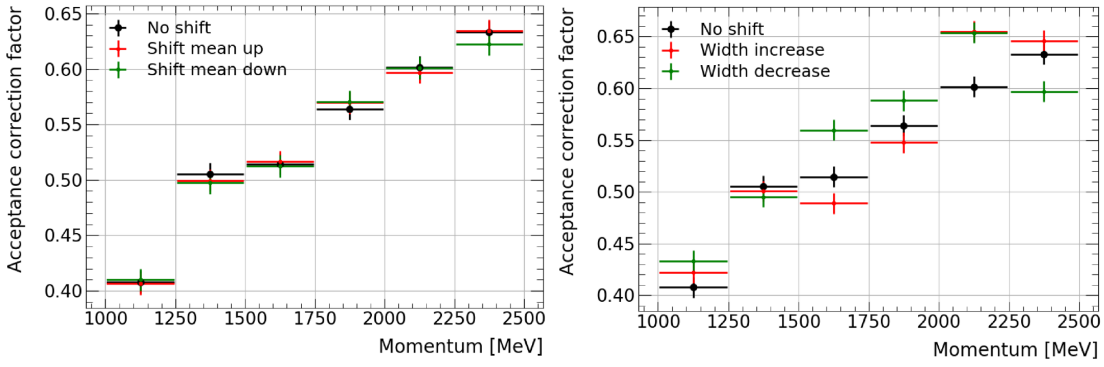

Figure 7.57: Shifts seen in the acceptance factors when shifting the beam to mimic the effect of internal alignment. The mean shift is seen to have little impact, but the width is large enough to be important for the EDM analysis.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0002 mrad	0.0010 mrad	0.0003 mrad	0.0008 mrad	0.0005 mrad	0.0001 mrad

Table 7.16: Systematics for the largest possible internal displacement between the trackers, in mrad. These are found to be negligibly small.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0018 mrad	0.0007 mrad	0.0038 mrad	0.0019 mrad	0.0035 mrad	0.0002 mrad

Table 7.17: Systematics obtained by shifting the width of the beam to mimic the maximum effect of internal misalignment, in mrad. A larger effect is seen here than with the mean, which is not surprising as the width is known to be important to the analysis.

The uncertainties are generally small, with the largest impact coming from varying the width. All four parts of the alignment uncertainty are combined in quadrature to give a total alignment uncertainty for each momentum bin. The results of this are summarised in Table 7.18. This uncertainty remains subdominant compared to the radial field.

1000-1250 MeV	1250-1500 MeV	1500-1750 MeV	1750-2000 MeV	2000-2250 MeV	2250-2500 MeV
0.0025 mrad	0.0021 mrad	0.0040 mrad	0.0029 mrad	0.0037 mrad	0.0006 mrad

Table 7.18: Final tracker alignment systematics for Run 3B for each momentum bin. These are large enough to not be negligible, but are not the largest systematic even despite the conservative evaluation.

The internal alignment does not consider relative tilts between the trackers. However, the largest impact would be seen with all 8 trackers tilted in the same way, as any other configuration would lead to some cancellation. Therefore, the largest impact is equivalent to the global tilt alignment, which has already been accounted for.

7.3 The Run 2/3 EDM analysis results

The procedures described in this chapter are applied to extract a vertical tilt angle for each dataset in Run 2 and Run 3. All uncertainties are added in quadrature, with everything except the radial field uncertainty calculated and applied in the six momentum bins of the analysis. A 0th order polynomial fit is applied to extract a central value and error. The radial field uncertainty is then added in quadrature to give the final result for each run.

A full list of preliminary uncertainties for the Run 2/3 EDM analysis is presented in Table 7.20 for Run 2 and 7.21 for Run 3, with Table 7.19 showing those

common to all datasets. The individual datasets are all dominated by statistical uncertainties, with the radial field, acceptance, and MC/data differences being the next most important. While the acceptance and MC/data differences could hypothetically be improved, the radial field is unlikely to see a large improvement, so acts as a floor for the final possible sensitivity.

Uncertainty [mrad]	Dilution	Tracker resolution	Phase
All runs	0.0000212	0.000140	0.0000274

Table 7.19: Final combined uncertainties for the dilution and tracker resolution, which do not change over runs, and the phase uncertainty, which is so small it can be assumed to be constant as well.

Unc. [mrad]	Statistical	Radial field	Acceptance	Drift	MC weighting	Alignment	Total
2B	0.0578	0.00363	0.00281	0.0000600	0.00245	0.00128	0.0581
2C	0.0265	0.00393	0.00273	0.0000504	0.00330	0.00208	0.0272
2D	0.0285	0.00626	0.00274	0.0000510	0.00409	0.00586	0.0295
2E	0.0459	0.00360	0.00208	0.0000641	0.00247	0.00178	0.0461
2F	0.0467	0.00363	0.00273	0.0000572	0.00281	0.00339	0.0471
2G	0.0992	0.00492	0.00276	0.0000947	0.00363	0.00508	0.0994
2H	0.0855	0.00498	0.00274	0.0000912	0.00271	0.00321	0.0858

Table 7.20: Final uncertainties for Run 2 datasets, in mrad. The total uncertainty combines the uncertainties listed here with those in Table 7.19.

Unc. [mrad]	Statistical	Radial field	Acceptance	Drift	MC weighting	Alignment	Total
3B	0.0437	0.00307	0.00257	0.0000211	0.00286	0.00120	0.0438
3C	0.0705	0.00237	0.00273	0.0000423	0.000940	0.00445	0.0708
3D	0.0275	0.00355	0.00270	0.0000414	0.00219	0.00180	0.0279
3E	0.0416	0.00246	0.00273	0.0000609	0.00246	0.00156	0.0419
3F	0.0624	0.00249	0.00275	0.0000932	0.00239	0.00689	0.0626
3G	0.0387	0.00250	0.00273	0.0000870	0.00268	0.00384	0.0392
3I	0.0402	0.00411	0.00276	0.0000634	0.00272	0.00257	0.0407
3J	0.0477	0.00253	0.00268	0.0000715	0.00277	0.00227	0.0480
3K	0.0589	0.00271	0.00275	0.0000732	0.00282	0.00191	0.0591
3L	0.0688	0.00275	0.00274	0.0000731	0.00271	0.00205	0.0690
3M	0.0382	0.00278	0.00274	0.0000652	0.00286	0.00244	0.0386
3N	0.0273	0.00380	0.00274	0.0000213	0.00259	0.00491	0.0278
3O	0.0310	0.00334	0.00276	0.0000344	0.00268	0.00221	0.0313

Table 7.21: Final uncertainties for Run 3 datasets, in mrad.

The 0th order fit for Run 3B is shown in Figure 7.58, with the fits for all other datasets shown in Appendix E. This shows good agreement with the expected final

flat line shape, indicating that the momentum dependence of the EDM amplitude has been well-corrected for in both the dilution and acceptance. A potential slight downward drift is seen across the momentum bins, which leads to a χ^2/NDF value larger that one, but all bins are consistent with each other.

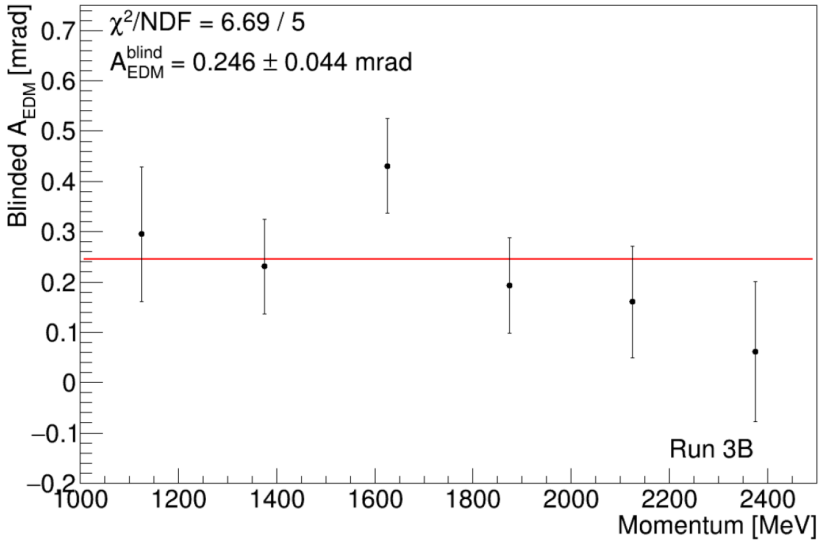

Figure 7.58: The final 0th order fit to the EDM amplitudes for Run 3B, with all corrections and uncertainties applied. A good agreement is seen with the fit line.

The rest of Run 2/3 in general is fitted well by a flat line, with a few datasets having χ^2/NDF values larger than one. A few also show the same downward drift behaviour as Run 3. Datasets with large variation are mostly the lower-stats datasets like Run 2G, which would suggest the impact is from statistical fluctuations. Each bin of these bins are checked to make sure the fit quality is good.

Finally, the extracted values for the tilt angle for each dataset are combined into a final plot, for Run 2 and Run 3 separately. The central value of each point is adjusted by subtracting the measured radial field from Tables 6.1 and 6.2, and finally fitted with a 0th order fit to extract a single blinded tilt angle. The points and fits are shown in Figure 7.59 for Run 2, and Figure 7.60 for Run 3.

In general, the results from all datasets are consistent with a flat fitted line, with variations within uncertainties. Since these results are blinded, the final EDM values cannot be extracted, but the uncertainties can be used to calculate the limits that would be set if the central value was zero.

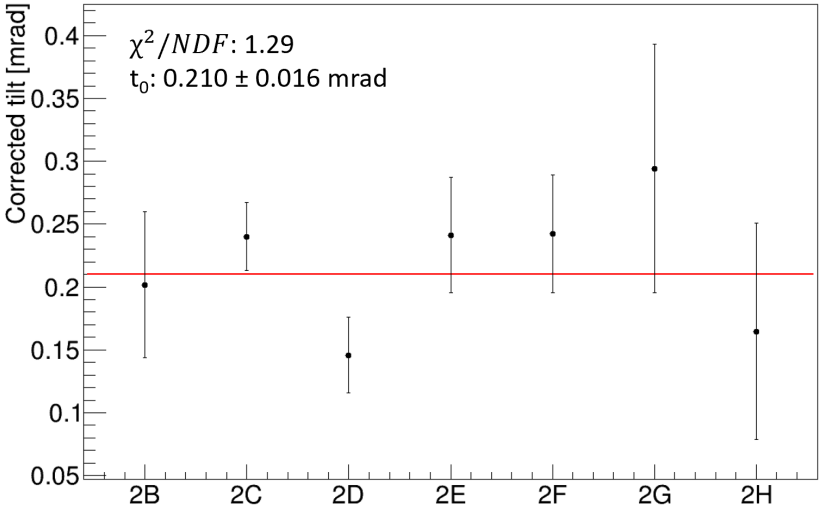

Figure 7.59: The final tilt extraction fit for the full blinded Run 2.

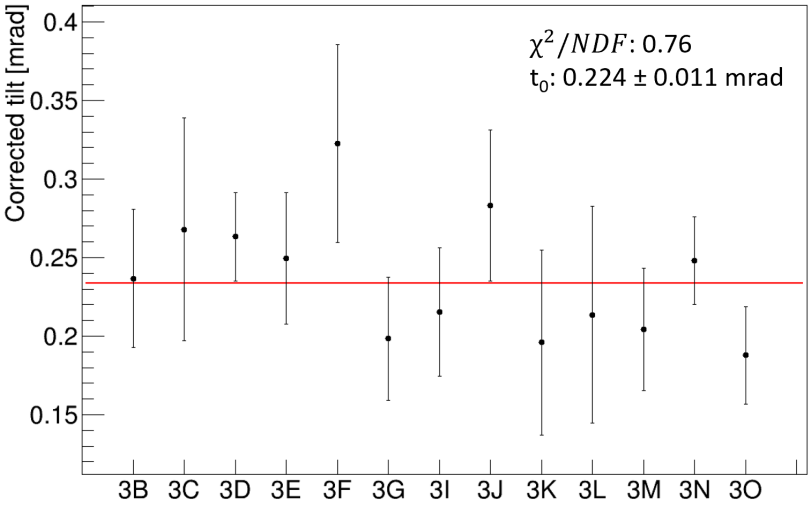

Figure 7.60: The final tilt extraction fit for the full blinded Run 3.

The final tilt angle uncertainties are converted into e·cm using Equation 2.25:

$$\text{Run 2}: \quad d_\mu^{unc} = \pm\, 5.064 \times 10^{-20}\, e \cdot cm, \tag{7.6}$$

$$\text{Run 3}: \quad d_\mu^{unc} = \pm\, 3.561 \times 10^{-20}\, e \cdot cm. \tag{7.7}$$

The limits that this would give at the 95 % confidence level, using the Feldman-Cousins method as decided in Section 5.2, gives the following preliminary limit

results for the Run 2/3 EDM analysis:

$$\text{Run 2}: \quad |d_\mu| < 8.41 \times 10^{-20}\, e \cdot cm, \tag{7.8}$$

$$\text{Run 3}: \quad |d_\mu| < 5.92 \times 10^{-20}\, e \cdot cm. \tag{7.9}$$

These represent a factor of three improvement compared to the world best limit on the muon EDM, set at BNL [1]. Compared to Run 1, the results follow the expected $\sqrt{n}$ statistical scaling, where n is the number of vertices. Therefore, this result can be considered to be statistically limited, which is a testament to the success of the radial field measurement campaign and the reduction of the acceptance uncertainty.

7.4 Conclusions/Further work

A blinded analysis searching for a muon EDM is presented, using the Run 2/3 data from the Fermilab muon $g-2$ experiment. The final results, if the central value is found to be consistent with zero when unblinded, will set a new world limit on the muon EDM, improving on the Run 1 analysis by a factor of two and the BNL result by a factor of 3. Many of the steps towards unblinding have been taken in this thesis, with a few cross-checks remaining before the blinding amplitude can be subtracted in each bin to reveal the measured central value and final limit. The necessary corrections for beam dynamics, early time effects, acceptance and dilution are tested and evaluated for all datasets within the run periods, with the early time effect being particularly interesting as it allows the analysis to start one $g-2$ period earlier. Since the result is statistically limited, this directly translates to an improvement on the final limits. The acceptance calculation and uncertainties are also improved with the introduction of a new acceptance map method, which allows a higher-statistics MC sample to be used. These maps have been fine-tuned to give the smallest possible uncertainty, improving on previous work by a factor of 4. This, combined with the dedicated radial field measurement, allows the final result to be statistically limited, which is a great success for the Run 2/3 analysis.

The results presented are preliminary, with some work still remaining until the analysis can be unblinded, which is discussed briefly here. Firstly, studies have shown that some of the vertex quality cuts could be loosened for the EDM analysis, which is demonstrated by the vertex resolution having a very small contribution

to the overall uncertainty. The loosening of a cut such as the requirement for the extrapolation to not pass through material could potentially drastically increase the number of vertices by a factor of 2-3. However, this must be done with caution to prevent this from introducing other dominant effects. If this is successful, the three leading systematic uncertainties, the acceptance, radial field and data/MC differences will become the uncertainties limiting the sensitivity. While the radial field is difficult to improve further for Run 2/3, the acceptance uncertainties can be improved by using a higher-statistics MC, or binning the acceptance maps more finely in momentum. This involves a more in-depth understanding of the data/MC differences. Preliminary studies have shown that the momentum distributions between data and MC do not match perfectly, particularly in the low momentum bins, which would impact the corrections and their uncertainties. Binning the acceptance maps more finely would help reduce the impact of this difference, or alternatively the differences themselves could be better understood and corrected in the MC rather than just scaling the mean and width. Finally, before unblinding it would be beneficial to perform a 'sideband' study, plotting the average vertical angle modulo a different frequency. The expected precession at a frequency not equal to the EDM is zero, so this provides a good check that no biases are introduced by the analysis steps and correction factors, and can be unblinded prior to the main EDM unblinding.

Chapter 8

Summary and conclusions

One of the physics goals of the Fermilab Muon $g-2$ experiment is to search for a muon electric dipole moment (EDM) by looking for a vertical tilt in the precession plane of the muon spin. This can be measured using both the calorimeters and the trackers. The tracker method is expected to be the most sensitive to an EDM due to being statistically limited, whereas the calorimeter methods are systematically limited. A blinded tracker EDM analysis is presented including studies to optimise and characterise the performance of the trackers, the development of a simulation to study the sensitivity, a measurement of the leading systematics, and analysis of the Run 2/3 data.

Three main tracker studies were performed; a HV scan, a study of crosstalk, and investigation into the time dependent resolution. It was found that both the resolution and the efficiency decreased substantially with any reduction in the wire voltage, with no evidence of ageing seen. As a result, the best voltage for the trackers was determined to be still 1650 V.

Crosstalk was studied in data, with an identification algorithm defined for crosstalk hits in order to study their properties. The selection method was tested using signals in dead straws, and used to estimate the prevalence of crosstalk in data as well as the likely causes. The overall crosstalk percentage was found to be around 6 %, with the majority coming from interactions between straw components. Crosstalk hits are simulated within the GM2RINGSIM experimental simulation to study the impact of crosstalk on tracking. The largest differences are seen in the number of tracks passing quality cuts, with those that pass having significantly worse vertex reconstruction. However, a width cut removes the worst of the impact, leaving the overall final crosstalk percentage as less than 0.5 %.

During the crosstalk studies, a time dependence was found, with similar lifetime behaviour also seen in the tracker resolution. One possible mechanism to explain

this is a space-charge effect, where the rate of hits is too high for the straws at early times in a fill. This effect is visible in other analyses that use tracker data, including the ω_a and EDM measurements, and must be corrected for.

A stand-alone MC is written to investigate the EDM limits the $g-2$ experiment should reach with the increase in statistical sensitivity, and to study the impact of the largest systematics, including the radial field and tracker acceptance. The results show that a large radial field dominates the total uncertainty, with the final EDM results being radial field limited unless it is measured to better than 1 ppm uncertainty. Given that at the time this was only known to at best 10 ppm, a dedicated radial field measurement was implemented in order to control this systematic. The impact of the acceptance was also seen as a worsening of the sensitivity, proportional to the reduction in amplitude seen.

Although the simple MC does not include sufficient complexity to fully model the acceptance, requiring the use of GM2RINGSIM for the full analysis, it remains a useful tool, both by the author and other collaborators. In future it may be useful to update the MC to use the full functional form of the EDM, as this would likely fix some of the inaccuracies with the momentum distributions/acceptance.

A dedicated radial field measurement is performed, successfully measuring the radial field to better than the needed 1 ppm target precision. The measurement is extrapolated to other run periods in order to use it for the Run 2/3 EDM analysis. This achieves sufficient precision for the Run 2/3 run periods to also not be limited by the radial field, making the next most dominant uncertainty the acceptance correction.

Finally, the Run 2/3 EDM analysis itself is performed. The EDM is blinded by injecting a large fake signal, and the data is plotted and fitted to extract the tilt angle. Corrections such as beam dynamics are studied, and the early time behaviour due to the space charge effect is corrected. This allows an earlier fit time, improving the statistical sensitivity by 3 %. The fit function itself is investigated and updated, and the analysis is binned in momentum bins, increasing the sensitivity by a further 11 %. The acceptance and dilution corrections are calculated using GM2RINGSIM, with a new acceptance map method used. This reduces the acceptance uncertainty by a factor of four, leaving the Run 2/3 EDM analysis fully statistically-limited. The systematics for the Run 2/3 analysis are evaluated for all 20 datasets across the two runs, including studies on tracker alignment, data/MC differences, uncertainties arising from resolution, the rise correction and the uncertainty on the EDM phase.

The final blinded results, assuming the central value is consistent with zero,

are a limit of $|d_\mu| < 8.41 \times 10^{-20}\, e \cdot cm$ for Run 2, and $|d_\mu| < 5.92 \times 10^{-20}\, e \cdot cm$ for Run 3, both using the Feldman-Cousins method at the 95% significance level. Eventually, these will be unblinded, and if consistent can be combined. For now, this is an improvement on the world leading EDM limit by a factor of 3.

Looking to the future, the Run 1 and Run 2/3 results are still only a fraction of the total data collected from the Fermilab experiment, with Runs 4,5 and 6 all larger than Run 3. This means the statistical uncertainty of the combination will continue to fall, with an expected final limit around $|d_\mu| < 3 \times 10^{-20}\, e \cdot cm$ assuming systematics like the acceptance uncertainty and the data/MC differences are adequately controlled [81]. The difference here from the initial TDR prediction arises due to the width of the beam limiting the overall sensitivity, an effect which was only recently quantified in the Run 1 analysis. Even further ahead, other experiments are also currently being designed to also measure the muon EDM, cross-validating the results from Fermilab and pushing down into the parameter space even further.

An experiment at J-PARC [94] is planning to measure both the magnetic and electric dipole moments, with a similar sensitivity to the Fermilab experiment but with an independent measurement method. Additionally, a dedicated muon EDM measurement is planned at PSI [38], using a 'frozen spin' technique where the magnetic precession is removed in order to achieve greater sensitivity to an EDM. The expected sensitivity is around 10^{-23} e·cm, which starts to push down far enough to verify or exclude some the BSM theories discussed in the introduction of this thesis. This means that through both Fermilab's $g-2$ experimental data and these future experiments, the muon EDM will continue to be a powerful search tool for new physics over the next few decades.

Bibliography

[1] Muon $g-2$ Collaboration (BNL), G. W. Bennett et al., *Improved limit on the muon electric dipole moment*, Phys. Rev. D **80** (2009) 052008.

[2] P. Kusch and H. M. Foley, *The Magnetic Moment of the Electron*, Phys. Rev. **74** (1948).

[3] P. A. M. Dirac, *The quantum theory of the electron*, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 117 no. **778**, (1928).

[4] A. Czarnecki and W. J. Marciano, *Muon anomalous magnetic moment: A harbinger for "new physics"*, Phys. Rev. D **64**, (2001) 013014.

[5] G. Charpak, P. J. M. Farley, E. L. Garwin, T. Muller, J. C. Sens, and A. Zichichi, *The anomalous magnetic moment of the muon*, Il Nuovo Cimento (1955-1965) **37 no. 4**, (1965).

[6] F. Combley, F. Farley, and E. Picasso, *The CERN muon ($g-2$) experiments*, Physics Reports **68 no. 2**, (1981).

[7] F. Farley and Y. Semertzidis, *The 47 years of muon $g-2$*, Progress in Particle and Nuclear Physics **52 no. 1**, (2004).

[8] Muon $g-2$ Collaboration (BNL), G. W. Bennett et al., *Final report of the E821 muon anomalous magnetic moment measurement at BNL*, Phys. Rev. D **73** (2006) 072003.

[9] T. Aoyama et. al (The Muon $g-2$ theory initiative), *The anomalous magnetic moment of the muon in the Standard Model*, Physics Reports **887**, (2020).

[10] Muon $g-2$ Collaboration, J. Grange et al., *Muon ($g-2$) Technical Design Report*, arXiv:1501.06858 [physics.ins-det].

[11] ACME Collaboration, *Improved limit on the electric dipole moment of the electron*, Nature **562(7727)**, (2018).

[12] S.Grant, *Towards a search for the Electric Dipole Moment of the muon at the Fermilab Muon $g-2$ experiment*, PhD Thesis, University College London. (2022).

[13] J. P. Miller, E. de Rafael, and B. L. Roberts, *Muon $g-2$: experiment and theory*, Reports on Progress in Physics **70 no. 5**, (2007).

[14] J. D. Jackson, *Classical Electrodynamics*, Wiley 3rd Edition. (1999).

[15] B.L. Roberts and W.J. Marciano, *Lepton Dipole Moments*, World Scientific, (2009).

[16] J. Schwinger, *A theory of the fundamental interactions*, Annals of Physics **2 no. 5**, (1957).

[17] T. Aoyama, M. Hayakawa, T. Kinoshita and M. Nio, *Tenth-order electron anomalous magnetic moment: Contribution of diagrams without closed lepton loops*, Phys. Rev. D **91**, (2015) 033006.

[18] C. Gnendiger, D. Stockinger and H. Stockinger-Kim, *The electroweak contributions to $(g-2))\mu$ after the Higgs-boson mass measurement*, Phys. Rev. D **88**, (2013) 053005.

[19] BMW Collaboration, *Leading hadronic contribution to the muon magnetic moment from lattice QCD*, Nature **593**, (2021).

[20] F. Ignatov et. al (The CMD-3 collaboration), *Measurement of the $e^+e^- \to \pi^+\pi^-$ cross section from threshold to 1.2 GeV with the CMD-3 detector*, arXiv:2302:08834 [hep-ex].

[21] B. Abi et. al (The Muon g-2 Collaboration), *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm*, Phys. Rev. Lett. **126**, (2021) 141801.

[22] G. Luders, *On the Equivalence of Invariance under Time Reversal and under Particle-Antiparticle Conjugation for Relativistic Field Theories.* Kongelige Danske Videnskabernes Selskab, Matematisk-Fysiske Meddelelser, **28 (5)**, (1954).

[23] 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).

[24] J. Libby, *Direct CP violation in hadronic B decays*, arXiv:1412.4269 [hep-ex].

[25] A. D. Sakharov, *Violation of CP invariance, C asymmetry, and baryon asymmetry of the Universe*, Pisma Zh. Eksp. Tcor. Fiz. **5**, (1967).

[26] J. Smith, E. Purcell and N. Ramsey, *Experimental limit to the electric dipole moment of the neutron*, Phys. Rev. **108**, (1957).

[27] V. Flambaum, I. Samsonov, *Limits on CP-violating hadronic interactions and proton EDM from paramagnetic molecules*, Journal of High Energy Physics **77**, (2020).

[28] C. Abel et. al, *Measurement of the Permanent Electric Dipole Moment of the Neutron*, Phys. Rev. Lett. **124**, (2020).

[29] L. Everett et. al, *Implications of muon g-2 for supersymmetry and for discovering superpartners directly*, Phys. Rev. Lett. **86**, (2001), 3484.

[30] S. Li, X. Li and Y. Yang, *Muon $g-2$ in a U(1)-symmetric Two-Higgs-Doublet Model*, Phys. Rev. D. **99**, (2019) 035010.

[31] A. Carcamo Hernandez, S. King, H. Lee and S. Rowley, *Is it possible to explain the muon and electron $g-2$ in a Z' model?*, Phys. Rev. D. **101**, (2020) 11.

[32] G. Mohlabeng, *Revisiting the dark photon explanation of the muon anomalous magnetic moment*, Phys. Rev. D. **99**, (2019) 115001.

[33] I. Dorsner et. al. *Muon $g-2$ and scalar leptoquark mixing*, Journal of High Energy Physics **89**, (2020).

[34] A. Crivellin, M. Hoferichter and P. Schmidt-Wellenburg, *Combined explanations of $(g-2)_\mu$, e and implications for a large muon EDM*, Phys. Rev. D. **98**, (2018) 113002.

[35] E. Chun, J. Kim and T. Mondal, *Electron EDM and Muon anomalous magnetic moment in Two-Higgs-Doublet Models*, JHEP. **12**, (2019) 068.

[36] X. Dong, S. Zhao, H. Zhang and T. Feng, *The two-loop corrections to lepton MDMs and EDMs in the EBL MSSM*, J. Phys. G. **47**, (2020) 045002.

[37] M. Beneke, P. Dey and J. Rohrwild, *The muon anomalous magnetic moment in the Randall-Sundrum model*, Journal of High Energy Physics **10**, (2013).

[38] A. Adelman et. al., *Search for a muon EDM using the frozen-spin technique*, arXiv:2102.08838 [hep-ex].

[39] C. Wu, E. Ambler, R. Hayward, D. Hoppes, and R. Hudson, *Experimental Test of Parity Conservation in Beta Decay*, Phys. Rev. **105** (1957).

[40] M. Tanabashi (Particle Data Group), *Review of Particle Physics*, Phys. Rev. D. **98**, (2018).

[41] S. Charity, *Beam profile measurements using the straw tracking detectors at the Fermilab Muon $g-2$ experiment, and a study of their sensitivity to a muon electric dipole moment*, PhD Thesis, Liverpool University. (2018).

[42] J. Miller, *Expressions for A(E) and N(E) in the Lab Frame*, BNL g-2 note #69.

[43] G. Bennett et. al. *Statistical equations and methods applied to the precision muon $g-2$ experiment at BNL*, Nucl. Instrum. Meth. A. **579**, (2007).

[44] L. Thomas, *The motion of the spinning electron*, Nature **117(2945)**, (1926).

[45] E. Tiesinga et. al., *CODATA recommended values of the fundamental physics constants: 2018*, Rev. Mod. Phys. **93**, (2021), 025010.

[46] D. Hanneke, S. Fogwell, and G. Gabrielse, *New Measurement of the Electron Magnetic Moment and the Fine Structure Constant*, Phys. Rev. Lett. **100**, (2008) 120801.

[47] W. Liu et al., *High Precision Measurements of the Ground State Hyperfine Structure Interval of Muonium and of the Muon Magnetic Moment*, Phys. Rev. Lett. **82**, (1999).

[48] G. Lukicov, *Alignment for the straw tracking detectors for the Fermilab Muon $g-2$ experiment and systematic studies for a muon electric dipole moment measurement*, PhD Thesis, University College London, (2020).

[49] P. Debevec, E989 Note 271: *Analytical expressions for $g-2$ and EDM asymmetry and Figure-Of-Merit*, DocDB: 25851, (2021).

[50] J. Price, E989 Note 285: *Vertical angle for EDM*, DocDB: 26584, (2022).

[51] D. Stratakis et. al., *Commissioning and first results of the Fermilab Muon Campus*, Phys. Rev. Accel. Beams. **22**, (2019), 011001.

[52] B. Abi et. al (The Muon g-2 Collaboration), *Measurement of the anomalous precession frequency of the muon in the Fermilab Muon $g-2$ experiment*, Phys. Rev. D. **103**, (2021) 072002.

[53] N. Froemming et. al., *Commissioning the superconducting magnetic inflector system for the Muon $g-2$ experiment*, Fermilab CONF-18-462-AD-PPD-1690576.

[54] A. Schreckenberger et. al., *The fast non-ferric kicker system for the Muon $g-2$ experiment at Fermilab*, Nucl. Inst. Mech. A. **1011**, (2021), 165597.

[55] Y. Semertzidis et. al., *The brookhaven muon $g-2$ storage ring high voltage quadrupoles*, Nucl. Inst. Mech. A. **503**, (2003).

[56] O. Kim et. al., *Reduction of coherent betatron oscillations in a muon $g-2$ storage ring experiment using RF fields*, New J. Phys. **22(6)**, (2020).

[57] T. Albahri et. al., *Beam dynamics corrections to the Run-1 measurement of the muon anomalous magnetic moment at Fermilab*, Phys. Rev. Accel. Beams. **24**, (2021), 044002.

[58] J. Grange, E989 Note 57: *Shift in ω_p due to couplings between longitudinal and transverse magnetic field inhomogeneities*, DocDB: 2548, (2015).

[59] G. Danby et. al., *The Brookhaven muon storage ring magnet*, Nucl. Instrum. Meth. A. **457**, (2001).

[60] D. Flay, *The muon $g-2$ experiment at Fermilab*, Proc. The 22nd International Spin Symposium (2016).

[61] T. Albahri et. al., *Magnetic field measurement and analysis for the Muon $g-2$ Experiment at Fermilab*, Phys. Rev. A. **103**, (2021), 042208.

[62] J. Kaspar et. al., *Design and Performance of SiPM-based readout of PbF_2 crystals for high-rate, precision timing applications*, JINST **12**, (2017).

[63] A. Anastasi et. al., *The laser-based gain monitoring system of the calorimeters in the Muon $g-2$ experiment at Fermilab*, JINST **14**, (2019).

[64] K. Shaw et. al., *Performance of the Muon $g-2$ calorimeter and readout systems measured with test beam data*, Nucl. Inst. Mech. A. **945**, (2019).

[65] W. Blum, W. Riegler, and L. Rolandi, *Particle detection with drift chambers*, Springer-Verlag Berlin Heidelberg, (2008).

[66] B. King. et. al., *The straw tracking detector for the Fermilab muon $g-2$ experiment*, JINST **17**, (2022).

[67] B. Bevenee et. al., *An amplifier shaper discriminator with baseline restoration for the ATLAS transition radiation tracker*, IEEE Trans. Nucl. Sci. **43**, (1996).

[68] V. Innocente, M. Maire, and E. Nagy. *GEANE: Average tracking and error propagation package*, Workshop on Detector and Event Simulation in High-energy Physics, (1991).

[69] N. B. Kinnaird. *Muon spin precession frequency extraction and decay positron track fitting in Run 1 of the Fermilab Muon $g-2$ experiment*, PhD thesis, Boston University, (2020).

[70] J. Mott., *Track quality cut service*, DocDB: 16444, (2019).

[71] P. Kammel and B. MacCoy, *IBSM manual*, DocDB: 9221, (2020).

[72] H. Binney, *T0 beam profiles for Run 3*, DocDB: 25458, (2021).

[73] F. Gray, *Auxiliary detectors*, DocDB: 1507, (2013).

[74] K. Labe, *Clock system documentation*, DocDB: 10995, (2019).

[75] S. Agostinelli et. al., *GEANT4, a simulation toolkit*, Nucl. Inst. Mech. A. **506**, (2003).

[76] S. Ritt and P. Amaudruz, *New components of the MIDAS data acquisition system*, IEEE conference on real-time computer applications for nuclear and particle physics, (1999).

[77] W. Gohn. et. al., *Data Acquisition for the new muon $g-2$ experiment at Fermilab*, J. Phys. Conf. **664**, (2015).

[78] R. A. Illingworth, *A data handling system for modern and future Fermilab experiments*, Journal of Physics: Conference Series **513**, (2014).

[79] S. Giron, *Measuring the electric-dipole moment of the muon at BNL E821*, PhD thesis, University of Minnesota, (2004).

[80] R. McNabb, *An improved limit on the Electric Dipole Moment of the muon*, PhD thesis, University of Minnesota, (2003).

[81] J. Price, *EDM: Update and projected sensitivity*, DocDB: 27468, (2022).

[82] M. Sossong, *A search for an Electric Dipole Moment of the positive muon*, PhD thesis, University of Illinois, (2005).

[83] J. Mott, *Tracker space charge model*, DocDB: 25786, (2021).

[84] J. Mott, *Tracker resolution*, DocDB: 21228, (2020).

[85] P.A. Zyla et al. (Particle Data Group), *2020 Review of Particle Physics*, Prog. Theor. Exp. Phys. (2020).

[86] G. J. Feldman and R. D. Cousins. *A Unified Approach to the Classical Statistical Analysis of Small Signals*. arXiv:physics/9711021 [physics.data-an], (1999).

[87] *https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.skewnorm.html*

[88] R. Osofsky, *Magnetic Field Determination for Run 1 of the Fermilab Muon $g-2$ experiment*, PhD Thesis, University of Washington, (2019).

[89] B. Morse, *Setting <Br> to zero*, DocDB: 23522, (2020).

[90] J. R. Taylor. *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements*, University Science Books, (1996).

[91] B. Delaunay, *Sur la sphère vide*, Bulletin de l'Académie des Sciences de l'URSS, Classe des Sciences Mathématiques et Naturelles. **6**, (1934).

[92] A. Keshavarzi, *Run 2/3 Tracker systematics: Global Alignment*, DocDB: 26602, (2022).

[93] G. Sweetmore, *Tracker internal alignment systematics*, DocDB: 26795, (2022).

[94] E34 collaboration, *Muon $g-2$/EDM experiment at J-PARC*, PoS KMI2017 **006**, (2017).

Appendix A

Full calculations and derivations

This section contains more detailed calculations and derivations of equations used in the thesis, specifically the decay distributions in Chapter 5, the uncertainty on k in Chapter 6 and the full form of the start time band in 7.

A.1 Integration of muon decay distributions

The differential decay distribution of the muon in 2D is given by:

$$\frac{d^2\Gamma}{dx d\cos\theta} \sim x^2[(3-2x) + P_\mu \cos\theta(1-2x)], \qquad x = \frac{E_e}{E_e^{max}} \tag{A.1}$$

, where x is the energy fraction E_e/E_e^{max}, P_u is the polarisation of the positron and θ is the angle between the muon spin direction and the momentum direction.

This must be integrated over energy and angle in order to give distributions in terms of just one variable in order to simulate it, which is non-trivial in the choice of limits so is presented here.

Integrating over angle to get the energy distribution, $\cos(\theta)$ varies between -1 and 1:

$$\frac{d\Gamma}{dx} = \left[3x^2\cos(\theta) - 2x^3\cos(\theta) + \frac{1}{2}P_\mu x^2(\cos(\theta))^2 - P_\mu x^3(\cos(\theta))^2\right]_{-1}^{1} \tag{A.2}$$

, which simplifies to:

$$\frac{d\Gamma}{dx} = 2(3x^2 - 2x^3) \propto (3x^2 - 2x^3). \tag{A.3}$$

Any multiplicative constants do not matter for the simulation, as the distributions are normalised.

Integrating over x, the energy varies between 0 and E_{max}, which changes x between 0 and 1:

$$\frac{d\Gamma}{d(\cos(\theta)} = \left[x^3 - \frac{1}{2}x^2 + \frac{1}{3}P_\mu x^3\cos(\theta) - \frac{1}{2}P_\mu x^4\cos(\theta)\right]_0^1 \tag{A.4}$$

, which simplifies to:

$$\frac{d\Gamma}{d\cos(\theta)} = \frac{1}{2}\left(1 - \frac{1}{3}P_\mu\cos(\theta)\right) \propto 1 - \frac{1}{3}P_\mu\cos(\theta). \tag{A.5}$$

A.2 Simplification of the closed orbit distortion for the radial field measurement

Starting with the form for the closed orbit distortion:

$$y(\theta) \approx \sum_{N=0}^{\infty}\frac{R_0}{B_0}\frac{B_{rcN}\cos(n\theta) + B_{rsN}\sin(N\theta)}{N^2 - n^2}, \qquad n = \frac{\kappa R_0}{vB_0} \tag{A.6}$$

This can be expanded, writing out the N terms explicitly and combining the sine and cosine together by introducing a phase shift ϕ_N:

$$y(\theta) \approx \frac{R_0}{B_0}\left[-\frac{B_{r,N=0}}{n} + \frac{B_{r,N=1}\cos(\theta+\phi_1)}{1-n} + \frac{B_{r,N=2}\cos(2\theta+\phi_2)}{4-n}...\right]. \tag{A.7}$$

Averaging around the ring, all the cosine terms average to zero, leaving only the first term:

$$\langle y\rangle = -\frac{R_0}{B_0}\frac{B_{r,N=0}}{n}. \tag{A.8}$$

The $B_{r,N=0}$ term is proportional to the average field around the ring, as it is the 0th order component of the total field, so this can be rewritten as:

$$\langle y\rangle \propto \frac{R_0}{B_0}\frac{\langle B_r\rangle}{n}. \tag{A.9}$$

which can then be related to κ and the ESQ field as detailed in the main text.

A.3 Propagation of uncertainties for k

Many results in the radial field chapter come from from performing a straight line fit, using both the gradient and the y-intercept to calculate the final values. These

are correlated, so the total uncertainty includes the covariance [90]:

$$\delta z^2 = \sqrt{\left(\frac{\delta z}{\delta m}\right)^2 \delta m^2 + \left(\frac{\delta z}{\delta c}\right)^2 \delta c^2 + 2\frac{\delta z}{\delta m}\frac{\delta c}{\delta c}\sigma_{mc}} \tag{A.10}$$

, where z is the value depending on both m and c.

Applying this to the conversion factor k:

$$k = 1 \Big/ \frac{m}{V} + c, \tag{A.11}$$

$$\begin{aligned}\delta k^2 = &\left(-\frac{1}{V}\frac{1}{\left(\frac{m}{V}+c^2\right)}\right)^2 \delta m^2 + \left(-\frac{1}{\left(\frac{m}{V}+c^2\right)}\right)^2 \delta c^2 \\ &+2\left(-\frac{1}{V}\frac{1}{\left(\frac{m}{V}+c^2\right)}\right)\left(-\frac{1}{\left(\frac{m}{V}+c^2\right)}\right)\sigma_{mc}.\end{aligned} \tag{A.12}$$

Which simplifies to:

$$\delta k = \sqrt{\left(\frac{1}{Vk^2}\right)^2 \delta m^2 + \frac{1}{k^4}\delta c^2 + \frac{2}{Vk^4}\sigma_{mc}}. \tag{A.13}$$

.

A.4 Start time scan simplification

The full form of the expected variation between two datasets, where one is a subset of the other, is given by:

$$\sigma_{diff} = \sqrt{\sigma_2^2 - \sigma_1^2\left(2\frac{A_1}{A_2}\cos(\phi_1 - \phi_2\right) - 1)} \tag{A.14}$$

, where A_1 and A_2 are fitted amplitudes and the ϕs are the 'analysing powers' of the two datasets. Where the two datasets are nested and contain mostly the same data, these are approximately equal, so this simplifies down into the form quoted:

$$\sigma_{diff} \approx \sqrt{\sigma_2^2 - \sigma_1^2}. \tag{A.15}$$

Appendix B

The Run 4 radial field

Although the author was not involved in the data collection and analysis of the Run 4 radial field scan data, it is referred to in the main text of this thesis several times, namely as Run 4 doesn't show a large drift in radial field across the run. The Run 5 setpoint is therefore used to extrapolate back to Run 4 to illustrate this.

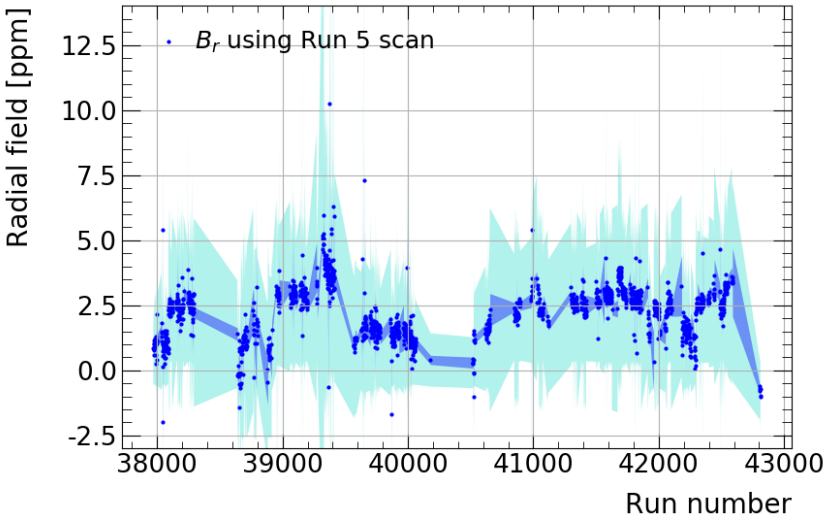

Figure B.1: The radial field across Run 4, calculated using the Run 5 setpoint.

The large variations seen for some runs comes from the fact that data from two faulty SiPMs have not been removed from the data. Nonetheless, this illustrates the flatness of Run 4 compared to Run 5.

Appendix C

Vertical betatron correlation matrix

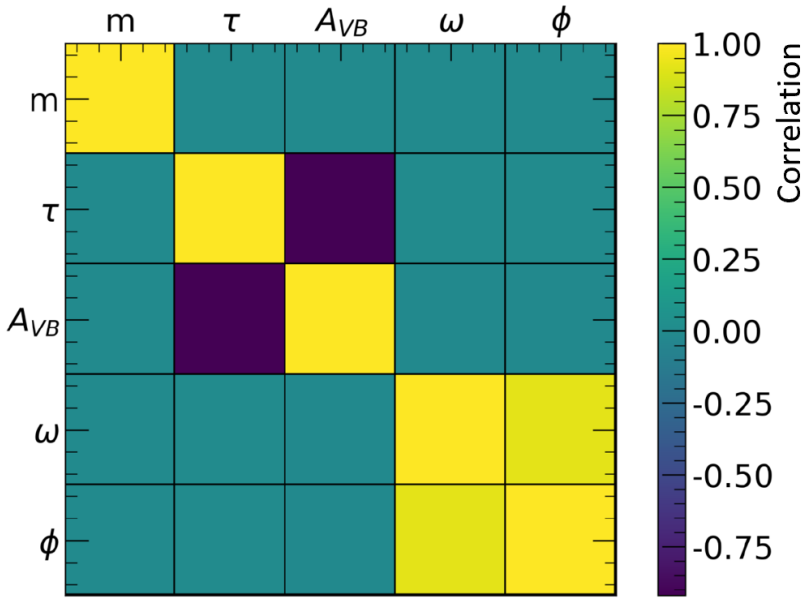

Figure C.1: Correlation matrix for the VB fit in Run 3B. Here, 1 corresponds to highly correlated, and -1 corresponds to highly anticorrelated.

As the fit parameters other than the offset m have high correlation with other parameters, this makes the fit very unstable and prone to finding local maxima/minima. Therefore, the randomisation method is much simpler to implement.

Appendix D

Optimal momentum binning for the EDM analysis

A slightly more in-depth coverage of the optimal momentum bin considers the different sizes of dataset across the Run 2/3 running period. Although using Run 3B as an average is good, the balancing of the statistical uncertainties per bin vs the improvement in sensitivity depends strongly on the initial statistics of the dataset. Here, an equivalent set of plots to Figure 7.30 is shown for a small dataset (Run 3J) and a large dataset (Run 3D).

In practice, the optimal binning varies, so this means it may be possible to gain some extra sensitivity in the full experimental result by binning each dataset differently in momentum. However, this would require a higher stats MC in order to calculate the corrections in each bin.

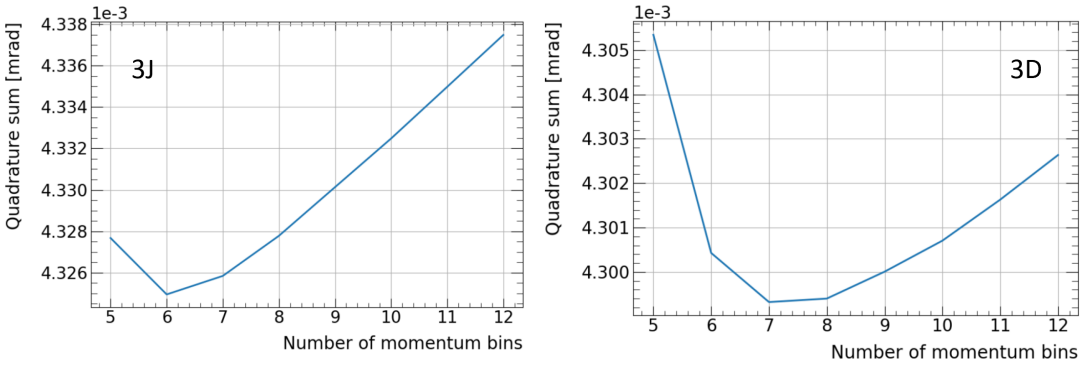

Figure D.1: A version of the quadrature plot for a small dataset, 3J and a large dataset, 3D. Both show a skew, with 6 being a good compromise between them.

Appendix E

Full Run 2/3 tilt fits

This appendix shows all 20 of the momentum-binned EDM tilt plots, with all corrections and uncertainties applied. In general, the fit quality is good, with no one clear shape in the data itself which would suggest some missed correction effect. There are some variations, particularly in the lower stats datasets, but these all look consistent with the size of the uncertainties of the bin in question.

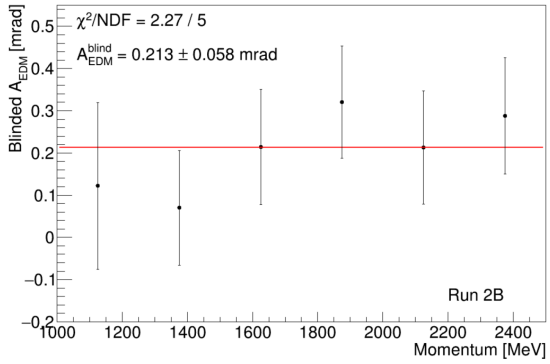

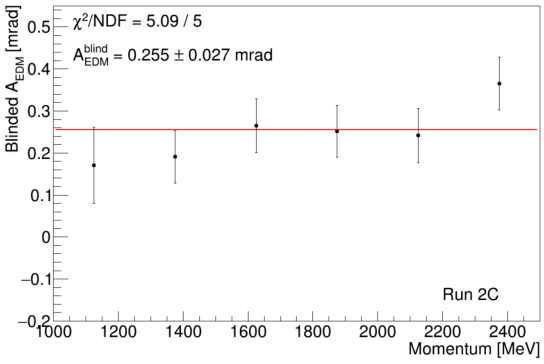

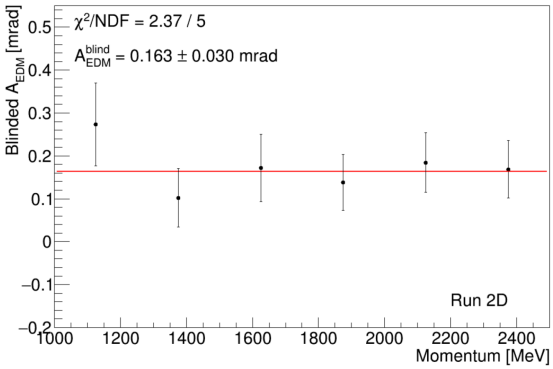

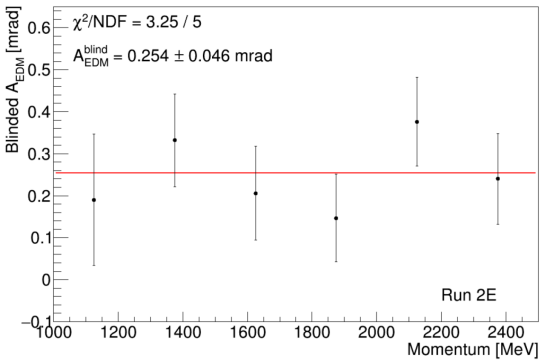

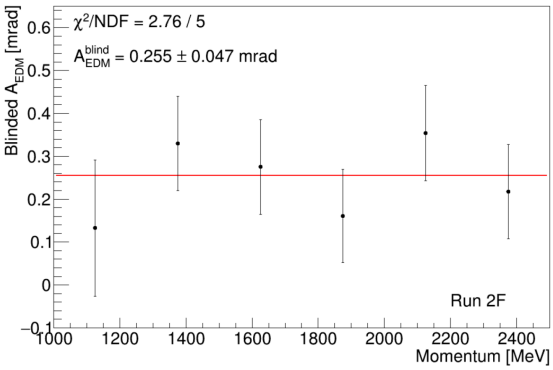
χ^2/NDF = 2.76 / 5
A_{EDM}^{blind} = 0.255 ± 0.047 mrad
Run 2F
Blinded A_{EDM} [mrad]
Momentum [MeV]

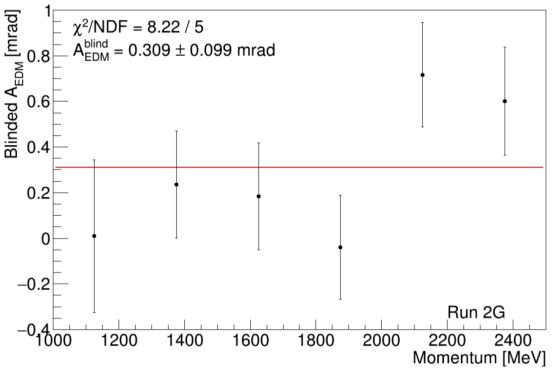
χ^2/NDF = 8.22 / 5
A_{EDM}^{blind} = 0.309 ± 0.099 mrad
Run 2G
Blinded A_{EDM} [mrad]
Momentum [MeV]

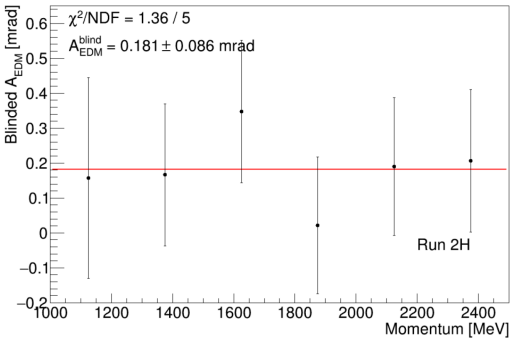
χ^2/NDF = 1.36 / 5
A_{EDM}^{blind} = 0.181 ± 0.086 mrad
Run 2H
Blinded A_{EDM} [mrad]
Momentum [MeV]

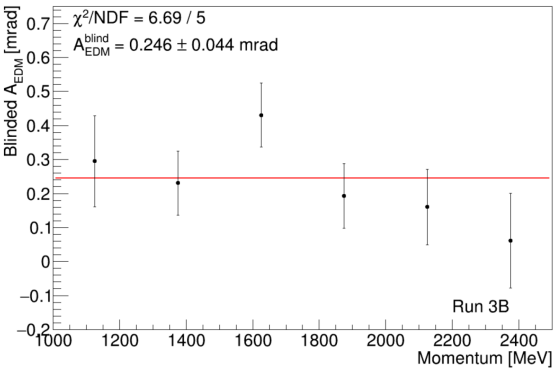
χ^2/NDF = 6.69 / 5
A_{EDM}^{blind} = 0.246 ± 0.044 mrad
Run 3B
Blinded A_{EDM} [mrad]
Momentum [MeV]

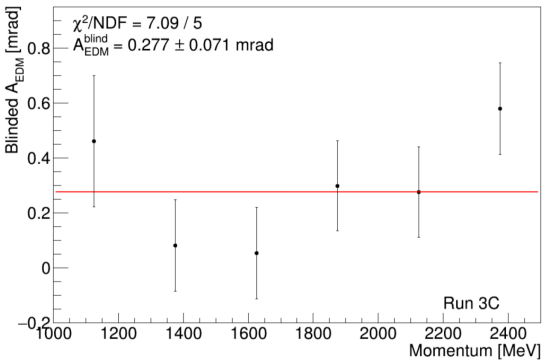
χ^2/NDF = 7.09 / 5
A_{EDM}^{blind} = 0.277 ± 0.071 mrad
Run 3C
Blinded A_{EDM} [mrad]
Momentum [MeV]

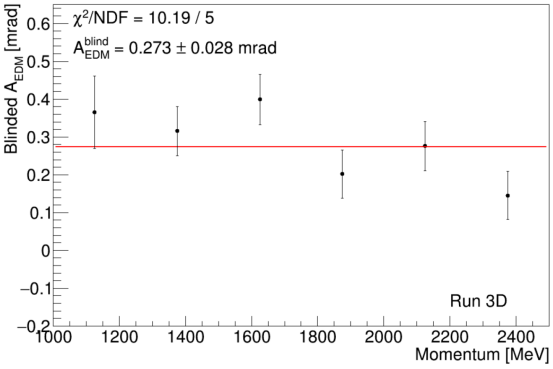
χ^2/NDF = 10.19 / 5
A_{EDM}^{blind} = 0.273 ± 0.028 mrad
Run 3D
Blinded A_{EDM} [mrad]
Momentum [MeV]

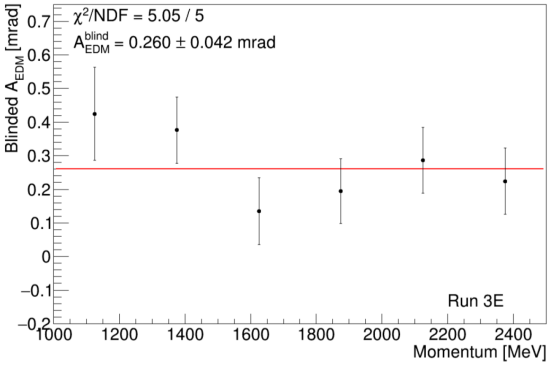
χ^2/NDF = 5.05 / 5
A_{EDM}^{blind} = 0.260 ± 0.042 mrad
Run 3E
Blinded A_{EDM} [mrad]
Momentum [MeV]

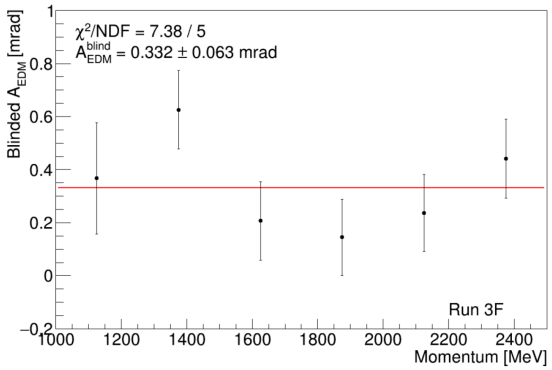
χ^2/NDF = 7.38 / 5
A_{EDM}^{blind} = 0.332 ± 0.063 mrad
Run 3F
Blinded A_{EDM} [mrad]
Momentum [MeV]

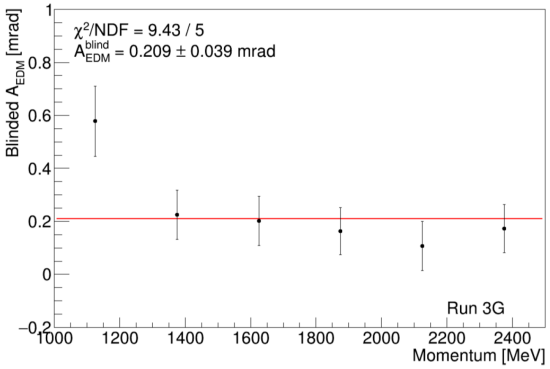
χ^2/NDF = 9.43 / 5
A_{EDM}^{blind} = 0.209 ± 0.039 mrad
Run 3G
Blinded A_{EDM} [mrad]
Momentum [MeV]

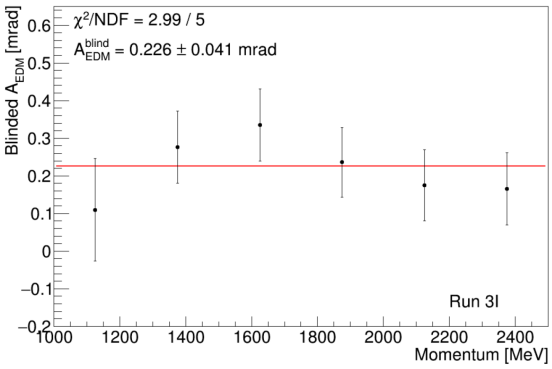
χ^2/NDF = 2.99 / 5
A_{EDM}^{blind} = 0.226 ± 0.041 mrad
Run 3I
Blinded A_{EDM} [mrad]
Momentum [MeV]

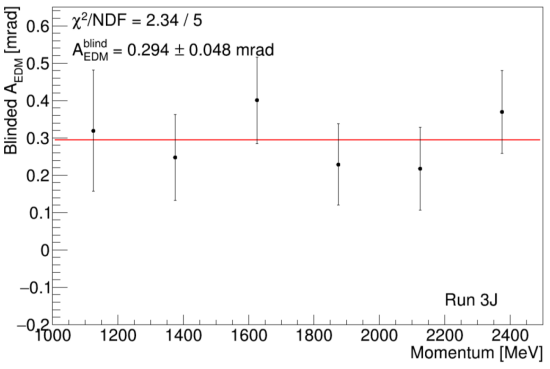
χ^2/NDF = 2.34 / 5
A_{EDM}^{blind} = 0.294 ± 0.048 mrad
Run 3J
Blinded A_{EDM} [mrad]
Momentum [MeV]

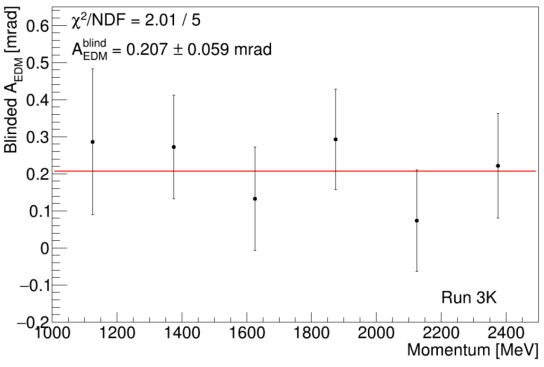
χ^2/NDF = 2.01 / 5
A_{EDM}^{blind} = 0.207 ± 0.059 mrad
Run 3K
Blinded A_{EDM} [mrad]
Momentum [MeV]

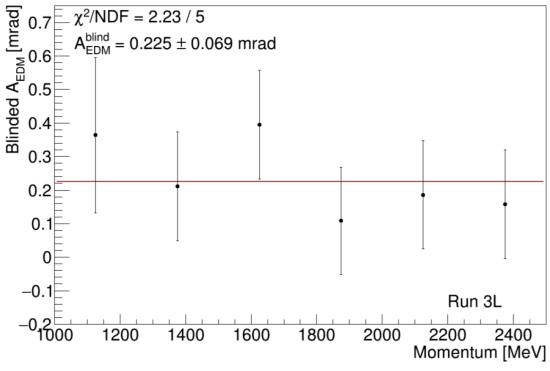
χ^2/NDF = 2.23 / 5
A_{EDM}^{blind} = 0.225 ± 0.069 mrad
Run 3L
Blinded A_{EDM} [mrad]
Momentum [MeV]

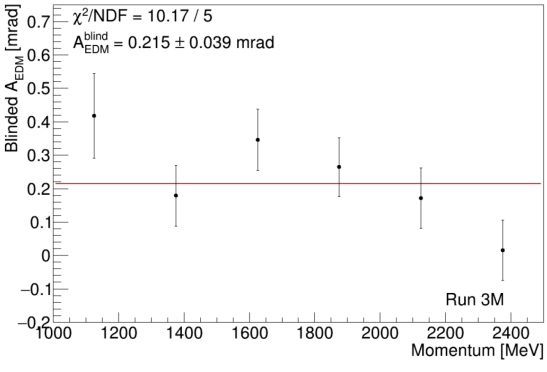
χ^2/NDF = 10.17 / 5
A_{EDM}^{blind} = 0.215 ± 0.039 mrad
Run 3M
Blinded A_{EDM} [mrad]
Momentum [MeV]

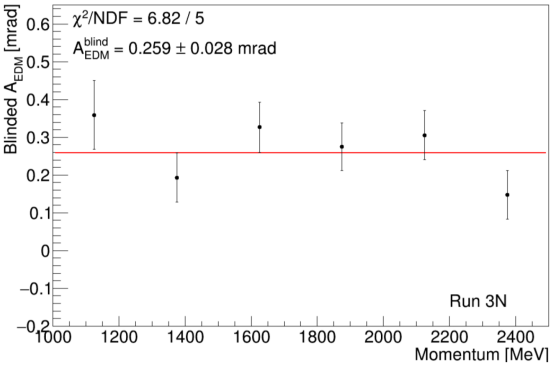
χ^2/NDF = 6.82 / 5
A_{EDM}^{blind} = 0.259 ± 0.028 mrad
Run 3N
Blinded A_{EDM} [mrad]
Momentum [MeV]

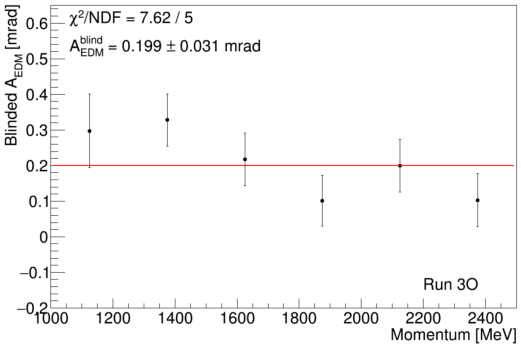
χ^2/NDF = 7.62 / 5
A_{EDM}^{blind} = 0.199 ± 0.031 mrad
Run 3O
Blinded A_{EDM} [mrad]
Momentum [MeV]
0.6
0.5
0.4
0.3
0.2
0.1
0
−0.1
−0.2
1000
1200
1400
1600
1800
2000
2200
2400