

Chapter 6

Systematic Uncertainties

6.1 Introduction

The $P_\mu^\pi \xi$ uncertainties are summarised in Table 6.1. The entries will be described fully in the current chapter. The table has two classes of $P_\mu^\pi \xi$ *systematic* uncertainties: those related to the accuracy of the P_μ simulation, and a separate group from the decay positron and its reconstruction. The largest uncertainty is from the P_μ simulation: the muon beam and magnetic field map together determine the muon's polarisation at the time of decay, and our knowledge of these limited the accuracy of the final result. The decay positron reconstruction uncertainties are evaluated simultaneously for ρ , δ and $P_\mu^\pi \xi$, by exaggerating an effect in the simulation or the analysis software. The table has two *statistical* uncertainties; these could be reduced by accumulating more data and/or simulation under the same running conditions.

There are four corrections to $\Delta P_\mu^\pi \xi$ (the difference in $P_\mu^\pi \xi$ between the data and a hidden simulation value):

1. The simulation does not include depolarisation of the muons while leaving the muon production target, resulting in a correction of $+0.9 \times 10^{-4}$ to the nominal sets, and $+5.9(5.2) \times 10^{-4}$ for the sets at $\langle p \rangle = 28.75 \text{ MeV}/c$ ($28.85 \text{ MeV}/c$). The uncertainty from these corrections is included in Table 6.1.
2. The simulations were generated with an incorrect rate for the time dependent depolarisation. A correction is applied rather than re-generating the simulation, since this would require several months of computer processing. The correction is $+2.9 \times 10^{-4}$ for silver (2006) and $+2.4 \times 10^{-4}$ for aluminium (2007). The uncertainty of this correction is negligible.
3. The statistics in the simulation spectra exceed the data by a factor of 1.8 to 4.0, depending on the set. This introduces a very small bias in the central value of $\Delta P_\mu^\pi \xi$ during the spectrum fitting (see Section 3.4 for a description of the fitting

technique). As a result, a set-independent correction of -0.5×10^{-4} is applied. The uncertainty of this correction is $\pm 0.5 \times 10^{-4}$, which is too small to enter Table 6.1.

4. The excess statistics in the simulation also causes a bias in the energy calibration procedure. This is applied a set-dependent correction of between $+1.3 \times 10^{-4}$ and $+2.3 \times 10^{-4}$. The exact numbers will be given in Table 7.1. The uncertainty of this correction is $< 0.5 \times 10^{-4}$, which is again too small to enter Table 6.1.

Corrections 1 and 2 will be described in this chapter. Further detail on corrections 3 and 4 can be found elsewhere[79, 97].

Table 6.1: Uncertainties for $\Delta P_\mu^\pi \xi$, the difference in $P_\mu^\pi \xi$ between the data and a hidden simulation value. These are systematic unless marked (stat.).

Category	Section	$\Delta P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)		
		This analysis	MacDonald '08 [10, 83]	Jamieson '06 [21, 57]
Polarisation				
Magnetic field map and μ^+ beam	6.2.1 to 6.2.9	-4.0,+15.8	Not eval. ^a	34.0
Stopping material				
Systematic effects	6.2.10	3.2	Not eval.	12.0
λ (stat.)	6.2.10	2.4	Not eval.	Not eval.
μ^+ production target	6.2.11	0.3	Not eval.	2.1
Background muons	6.2.12	1.0	Not eval.	1.8
Chamber response				
DC space-time-relationship	6.3.1	0.9	6.0	Not eval.
DC geometric effects	6.3.2	1.3	0.7	2.2
US-DS efficiency	6.3.3	1.4	1.1	1.9
Crosstalk	6.3.4	0.5	Not eval.	Not eval.
Wire time offsets	6.3.5	0.8	0.4	8.9
Detector alignment				
z length scale	6.4	0.0	0.7	2.2
u/v width scale	6.4	0.2	0.2	Not eval.
Positron interactions				
δ -electron rate	6.5.1	0.1	1.4	2.9
Bremsstrahlung rate	6.5.1	0.5	0.03	
Outside material	6.5.2	0.4	0.6	0.2
Resolution	6.6	1.5	0.7	Not eval.
Momentum calibration				
Tracking B-field	6.7.1	0.3	1.1	0.9
Kinematic endpoint				
Model uncertainty	6.7.2	1.4	Not eval.	Not eval.
Propagation to bulk	6.7.2	0.5	0.01	1.6
Beam stability				
Muon beam intensity	6.8.1	0.3	0.2	1.8
External				
Radiative corrections	6.9.1	0.5	0.5	1.0
η correlation	6.9.2	1.1	1.1	Not eval.
Extraction of $\Delta P_\mu^\pi \xi$ (stat.) ^b	7.1	2.5	3.7	6
Total systematic		-6.3,+16.5	-	38
Total statistical		3.5	-	6

^a In the most recent TWIST analysis (MacDonald '08), the polarisation uncertainties were not re-evaluated since it was a measurement of only ρ and δ .

^b This includes the statistical uncertainty from the kinematic endpoint parameters (see Section 6.7.2).

6.2 Polarisation

6.2.1 Overview

The simulation transports the muon spin from the end of the M13 beam line to the metal stopping target. This relies on the accuracies of the muon beam measurement and the magnetic field map, both of which will contribute to the assessment of the $P_\mu^\pi \xi$ uncertainty. The term “fringe field” will be widely used; this refers to the magnetic field from the end of the M13 beam line up to the first drift chamber.

The systematic uncertainties from the muon beam and fringe field are summarised in Table 6.2, in the order they appear in this chapter. First, the uncertainties in the alignment of the muon beam and the magnetic field map will be described. These will be used simultaneously to establish an uncertainty of ${}_{-1.2}^{+6.4} \times 10^{-4}$. Second, the uncertainties in the transverse magnetic field components will be evaluated using three consistency tests, and these will set an uncertainty of ${}_{-0.0}^{+13.9} \times 10^{-4}$. Third, uncertainty in the width of the muon beam angular distributions will lead to three uncertainties that total $\pm 3.8 \times 10^{-4}$.

Table 6.2: Summary of muon beam and fringe field uncertainties, for sets with a nominal beam tune.

Description	$\Delta P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)
μ^+ beam and magnetic field alignment	-1.2,+6.4
Magnetic field transverse components	-0.0,+13.9
μ^+ beam angular distributions	
Simulation of multiple scattering	± 3.1
Noise from TEC electronics	± 1.7
Aging of TEC sense planes	± 1.5
Quadratic sum	-4.0,+15.8

6.2.2 Measures of polarisation

Throughout this chapter the muon beam’s polarisation is always considered to be an average of the z -components of the muons’ spins, *not* the average of the spin projected onto the momentum vector. Two measurements of polarisation will now be described³⁶.

³⁶A third measurement is possible using the time dependence of the forward-backward asymmetry. This was described in Section 3.6: $P_\mu(t) = P_\mu^*(0) \exp(-\lambda t)$ is fit to the asymmetry. The difference in

In the simulation, the absolute polarisation is available *before* any time dependent target depolarisation has taken place. This is the polarisation after passing through the solenoidal magnetic field, and is calculated by simply averaging the z -components of the muons' spins. From here on the simulation's absolute polarisation is written as $P_\mu(0)$, where the (0) is a reminder that it is effectively determined at $t = 0$. Since the simulation knows the spin of each muon exactly, $P_\mu(0)$ is known with high precision using a relatively small number of muons.

A fit between two decay positron spectra (data or simulation) measures $\Delta P_\mu^D \xi$, where P_μ^D is the decay polarisation of the muon *after* time dependent depolarisation has taken place. If the spectra have the same ξ and time dependence, then this quantity is directly comparable to changes in $P_\mu(0)$. Spectrum fits between the data with a nominal and steered beam will later be used to evaluate the leading systematic uncertainty for this measurement.

6.2.3 Initial beam position and angle alignment

A muon beam measurement was made with the TECs at the beginning and end of most data sets. These two measurements have small differences in position and angle, resulting in different values of $P_\mu(0)$. Since the GEANT simulation can only use one of the beam measurements, the differences must be assessed as a systematic uncertainty.

The changes in average position and angle between the two muon beam measurements are listed in Table 6.3, where differences of up to 0.18 cm in position and 3 mrad in angle are observed. Note that these changes are smaller than the observations from the previous TWIST $P_\mu^\pi \xi$ analysis. The cause of the changes could be muon beam instability, a variation in the TEC drift cell response, or a limitation in the reproducibility of the TECs. Each of these possible causes will now be discussed.

Muon beam instabilities could originate from an instability in the M13 beam line elements (*i.e.* quadrupoles, dipoles, slits, jaws, asymmetric currents for quadrupole steering), or a change in the proton beam upstream of the muon production target. We do not believe there were measurable M13 beam line instabilities since all elements were monitored with a slow control system, and runs where an element fluctuated were not analysed (see Section 5.4). The proton beam was ruled out as a significant source of instability. In detail, a special test was carried out with the proton beam displaced at the production

$P_\mu^*(0)$ between two data sets is then a measure of their polarisation difference, as long as λ is the same for each set. Polarisation differences from $P_\mu^*(0)$ are *not* used in this chapter.

Table 6.3: Muon beam differences for the beginning and end of set TEC measurements. The temperature difference between the measurements is denoted by ΔT .

Set	Target	Description	$\Delta \langle x \rangle$ (cm)	$\Delta \langle y \rangle$ (cm)	$\Delta \langle \theta_x \rangle$ (mrad)	$\Delta \langle \theta_y \rangle$ (mrad)	ΔT^a ($^{\circ}\text{C}$)
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.11	-0.05	0.2	-3.2	-0.3
70	Ag	B = 1.96 T	0.03	0.00	1.0	-0.4	-1.2
71	Ag	B = 2.04 T	0.09	-0.05	0.0	0.1	2.4
74	Ag	Nominal A ^b	-	-	-	-	-
75	Ag	Nominal B	0.04	-0.10	-0.5	1.5	3.2
76	Ag	Steered beam	-0.04	-0.06	-0.6	1.9	1.3
83	Al	Downstream beam package in place	0.12	-0.09	0.6	0.7	-0.3
84	Al	Nominal C	0.18	-0.15	0.2	1.4	-0.4
86	Al	Steered beam B	0.04	-0.01	1.0	-0.01	-0.4
87	Al	Nominal D	0.13	-0.11	-0.1	0.7	-1.3
91/92/93	Al	Lower momentum ^b	-	-	-	-	-

^a $\Delta T = T_{\text{end}} - T_{\text{start}}$, so that $\Delta T > 0$ indicates a temperature rise between measurements.

^b These sets only had one TEC measurement.

target by ± 0.1 cm vertically, which is about five times larger than the beam could have moved during normal operation³⁷. For this test, the largest observed TEC changes in the muon beam were $\Delta \langle y \rangle = \pm 0.07$ cm in position and $\Delta \langle \theta_y \rangle = \pm 1.0$ mrad in angle, which are negligible after scaling down by a factor of five.

We have good evidence that the muon beam was sufficiently stable throughout a set, and was therefore not the cause of the variations in Table 6.3. First, we acquired a whole week of data with the TECs in place throughout (set 72), and found that the average muon beam position and angle were stable to < 0.02 cm and < 1 mrad, respectively (see Fig. 5.2). Second, when the TECs were not in place the muon beam measurement from the wire chambers was used to monitor the beam's stability; a typical monitoring plot was shown in Fig. 5.3. The sensitivity of this measurement was calibrated by deliberately changing the currents in each quadrupole and dipole by $\pm 5\%$. This allowed us to conclude that the muon beam instabilities that were observed are all small and uncorrelated with any beginning/end of set changes in Table 6.3.

³⁷The proton beam was surrounded by four monitor plates (top, bottom, left, right). In order to steer the beam vertically by 0.1 cm and avoid destroying the top or bottom plate, the proton beam current was reduced from the nominal setting of $\approx 100 \mu\text{A}$ to $\approx 20 \mu\text{A}$. This suggests that during normal operation the proton beam could not have moved by more than $\approx 20/100 \times 0.1 \text{ cm} = 0.02 \text{ cm}$.

The space-time-relationship (STR) in the TEC drift cells depended on temperature. This was ruled out as the cause of the differences in Table 6.3. In detail, new STRs were prepared that corresponded to a $\pm 3^\circ\text{C}$ variation, and all the data from the TECs were then re-analysed. The average reconstructed positions changed by between 0.028 cm and 0.050 cm, depending on the proximity of the beam to the sense plane of the TEC³⁸. The predicted changes in angle were all < 0.05 mrad, except for set 76, which was still predicted to change by only 0.4 mrad. Clearly these changes in position and angle cannot explain the larger differences observed in Table 6.3.

Lastly, the insertion and removal of the TECs required the beam line elements to be switched off, and a breaking of the vacuum in the beam line, which then had to be pumped down again before data could be taken with the TECs. This process exerted significant forces on the beam line components and the box containing the TECs, and these forces are the prime candidate for the measured variation in initial position and angle. Therefore we conclude that the muon beam itself was stable and no systematic uncertainty is necessary for its fluctuations. Instead an uncertainty from the initial position and angle of the muon beam is needed, since we cannot be sure that the alignment of the TECs was reproducible.

Later we will evaluate the $P_\mu(0)$ sensitivity to position changes of ± 0.2 cm and angle changes of ± 3 mrad. These are the limits of the observations in Table 6.3, but they are not overly conservative for a number of reasons. First, there are not enough entries in Table 6.3 to establish whether the changes in position and angle follow a predictable distribution. Second, the TECs were aligned to the drift chambers with a systematic uncertainty of about 2 mrad. And third, the long term stability measurements of the TECs found movements of 0.1 cm, but only had an accuracy of about 0.1 cm (see Section 2.13 for more information on the alignments and stability measurements).

6.2.4 Magnetic field position and angle alignment

The fringe field map for the analysis was generated with the `Opera` software package[85]. The generated map was compared to B_z measurements that were taken with Hall probes (see Appendix D). This comparison could not produce a precise translational alignment in x and y . Instead the position of the muon beam inside the detector was used. On a set-by-set basis, a field translation was determined such that the data and simulation positions

³⁸The x -positions increased with temperature and the y -positions decreased; see Ref. [76] for further detail.

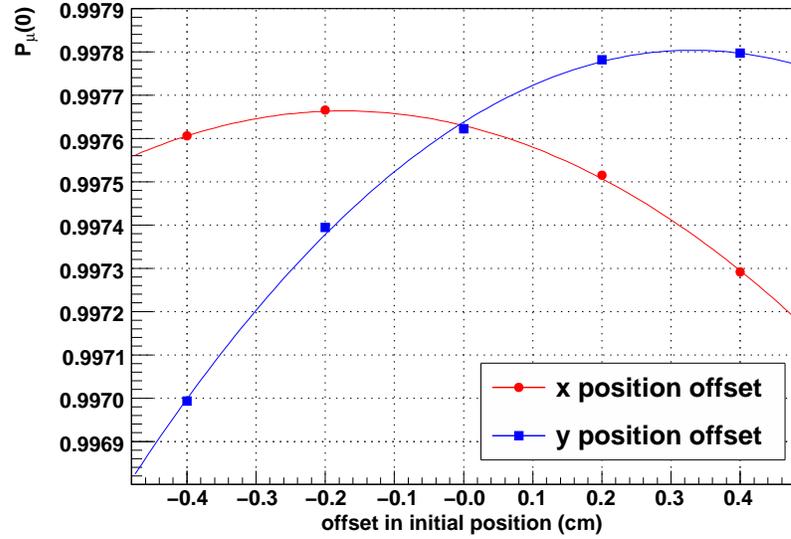
matched. On average this required a translation of the entire map by³⁹ $(\Delta x, \Delta y) = (0.21, 0.15)$ cm. The uncertainties in this translation from the muon beam and an error in the strength of the magnetic field components were evaluated, and found to be ± 0.04 cm. This is smaller than the uncorrelated uncertainty from the Hall probe alignment; these were aligned in the yoke's coordinate system system to < 0.1 cm in x and y and < 1 mrad in θ_x and θ_y [98]. The next section will use the uncorrelated uncertainties from the Hall probe alignment to determine the systematic uncertainty for $P_\mu(0)$.

6.2.5 Alignment uncertainties from beam and magnetic field

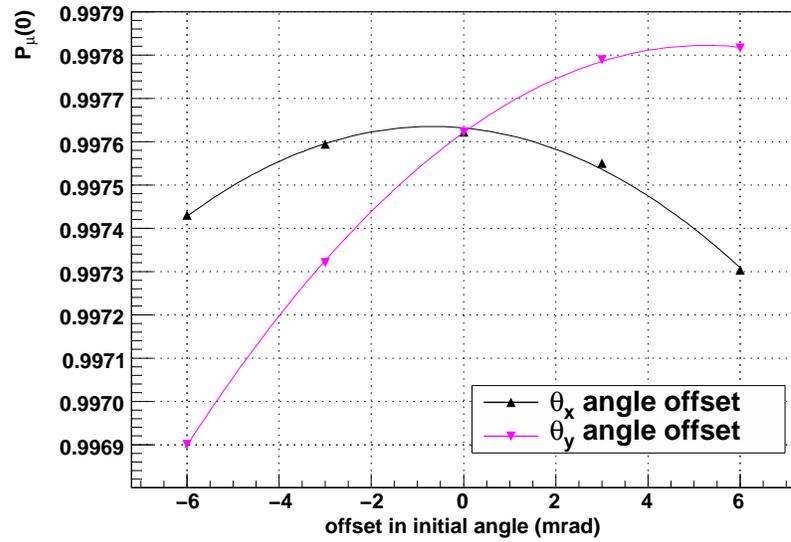
Section 6.2.3 established that the alignment of the muon beam is uncertain at the level of ± 0.2 cm in both the x and y position, and ± 3 mrad in both the θ_x and θ_y position. Similarly, Section 6.2.4 established the alignment of the magnetic field map to be uncertain at the level of ± 0.1 cm in x and y , and ± 1 mrad in θ_x and θ_y . All of these uncorrelated uncertainties must be combined in order to determine the sensitivity of $P_\mu(0)$ to possible misalignments.

The dependence of $P_\mu(0)$ on each alignment is non-linear. This is demonstrated for a nominal beam in Fig. 6.1. In addition, since the true beam alignment is optimised, the quantity $P_\mu(0)$ is also optimised, and introducing a combination of misalignments is therefore more likely to reduce $P_\mu(0)$ than increase it. As a result of these complications, the uncertainties are combined using a Monte Carlo approach. One hundred simulations are run with each misalignment selected from an independent Gaussian distribution with the standard deviation equal to the alignment uncertainty. For example, the initial beam position in x is drawn from a Gaussian with a mean of 0.0 cm, and standard deviation equal to 0.2 cm. The resulting $P_\mu(0)$ distribution from these simulations is asymmetric, and the width above and below the central value is used to establish a $P_\mu(0)$ systematic uncertainty of ${}_{-6.4}^{+1.2} \times 10^{-4}$ for the nominal sets. The uncertainty for $\Delta P_\mu^\pi \xi$, the difference between the data and a hidden simulation value, is then ${}_{-1.2}^{+6.4} \times 10^{-4}$.

³⁹In practice, an early analysis found a preliminary translation of $(\Delta x, \Delta y) = (0.20, 0.20)$ cm, and this was used to generate all the simulations. The correction between the preliminary translation of $(\Delta x, \Delta y) = (0.20, 0.20)$ cm and the final value of $(\Delta x, \Delta y) = (0.21, 0.15)$ cm was evaluated to have a negligible on $P_\mu(0)$ ($< 0.6 \times 10^{-4}$ for the nominal sets.)



(a) Initial position sensitivity.



(b) Initial angle sensitivity.

Figure 6.1: The polarisation of the muons after the fringe field, $P_{\mu}(0)$, has a quadratic dependence on changes in the initial position and angle of the beam. The results of simulating a nominal beam profile are shown.

6.2.6 Overview of the magnetic field map

The magnetic field map for the analysis was generated using the `Opera` software package[85]. The resulting components of \vec{B} are shown in Figs. 6.2 and 6.3. The longitudinal components (B_z) increase steadily up to the drift chamber (DC) tracking region. The transverse components (B_x , B_y) are less than 1.5 mT while on-axis ($x = y = 0$), but increase significantly off-axis; Fig. 6.3(a) demonstrates this by including the components for ($x = y = 1$ cm). The transverse components have approximate radial symmetry and are maximised just inside the door of the yoke. They are also strongly linear within ≈ 5 cm of the axis, which fully encloses the volume through which the muons pass (see Fig. 6.3(c)). The transverse components are closely related to the depolarisation, which is shown for the simulation of a nominal profile in Fig. 6.3(b). For example, the onset of rapid depolarisation coincides with the maximisation of the transverse components, and the field's inflection at $z = -100$ cm is accompanied by an inflection in the depolarisation. Clearly the quality of the fringe field downstream of the door is important since it controls the rate of depolarisation. In addition, the field upstream of the door must be known since it affects which part of the fringe field the beam is transported through.

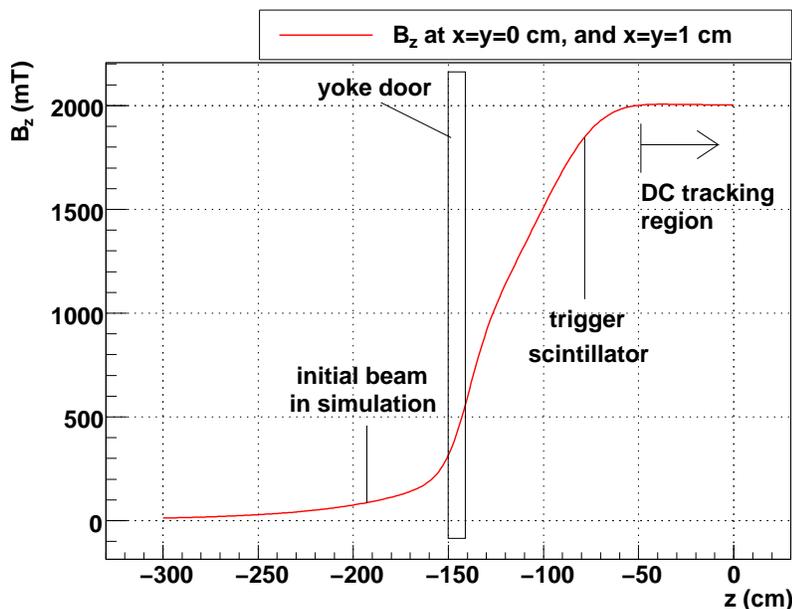
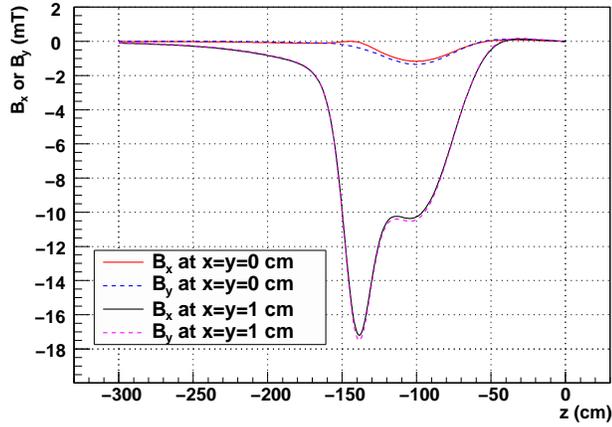
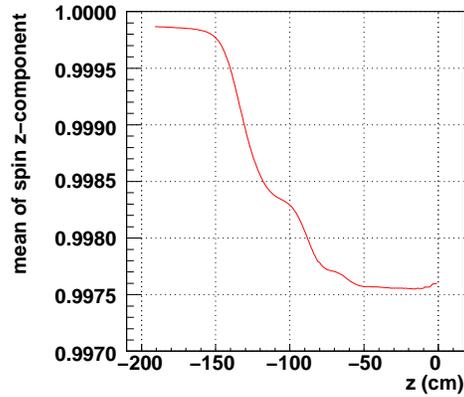


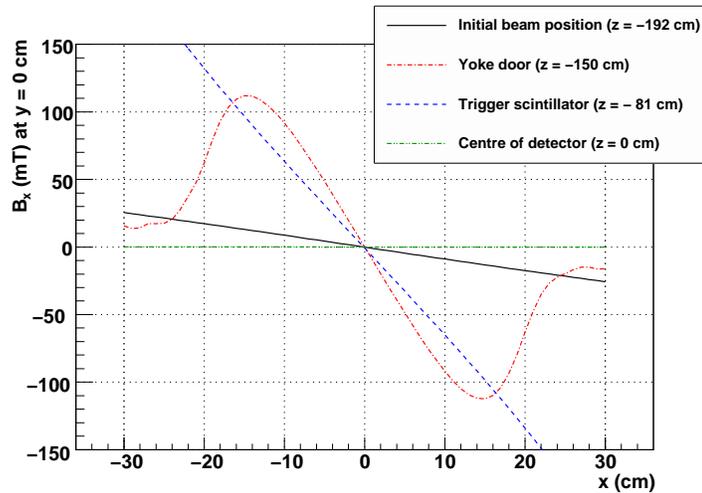
Figure 6.2: Longitudinal (B_z) components from the field map used for the analysis. These are produced from an `Opera` finite element analysis[85].



(a) B_x and B_y components.



(b) Average of spin z -component, for a nominal beam profile.



(c) B_x and B_y components for $y = 0$ cm.

Figure 6.3: Transverse fringe field components, B_x and B_y , from the field map used for the analysis. The average spin is also shown since its behaviour is closely related to the transverse components.

The B_z components from `Opera` are compared to Hall probe measurements in Fig. 6.4, where discrepancies of up to 6 mT are observed. The same comparison is not possible for the smaller B_x and B_y components since the Hall probes were single axis and did not measure them. We have identified three possible causes for the discrepancies. First, the `Opera` software did not include all the boundary conditions, such as the steel in the floor of the M13 area and the final M13 quadrupoles; an attempt was made to include these quadrupoles, but it did not find a significant effect for the B_z components. Second, there were several inputs to `Opera` that had to be tuned within their measured accuracy. For example, adjustments were made to the solenoid's coil positions in $x/y/z$, the radii of these coils and their current densities, the $B - H$ curve for the iron yoke, and the position of the yoke door in z [99, 100, 101, 102]. There was freedom of up to 0.2 cm in the positions of the solenoid coils[99], since their positions were not measured, and were instead initially placed according to a sketch from the magnet manufacturer. The previous $P_\mu^\pi \xi$ measurement found that variations in the `Opera` inputs affected $P_\mu(0)$ by 3×10^{-4} at most[57]. As a result of this low sensitivity the `Opera` inputs were not tuned further for the current measurement. Third, the `Opera` software used a finite element method to solve Maxwell's equations. This is expected to have limitations when modelling the 40 cm diameter circular hole in the yoke door at $z = -150$ cm, which is at a critical region for the *transverse* field components and hence the depolarisation. This is because there are several orders of length scale involved, and it is reasonable to expect that finite element analyses will suffer from accuracy problems under these conditions. In this case, the muons are within $\lesssim 5$ cm of the axis, there is a circular hole in the yoke of diameter 40 cm, and the whole map must be determined over a z -length of about 5 m.

If the discrepancy in the B_z components is eliminated, then the $P_\mu(0)$ estimates are changed by $< 1 \times 10^{-4}$ for all profiles; this test does not obey Maxwell's equations, but it does demonstrate that a discrepancy of a few mT in only the B_z components is unimportant at the 10^{-4} level. In contrast, a change of a few mT in the B_x and B_y components has a significant effect on $P_\mu(0)$. Since there are no Hall probe measurements of the B_x and B_y components, the uncertainty from them must be evaluated indirectly. We carry out this evaluation by comparing how well the simulation reproduces large differences in polarisation between sets with a nominal and a steered muon beam. The next sections will describe this technique.

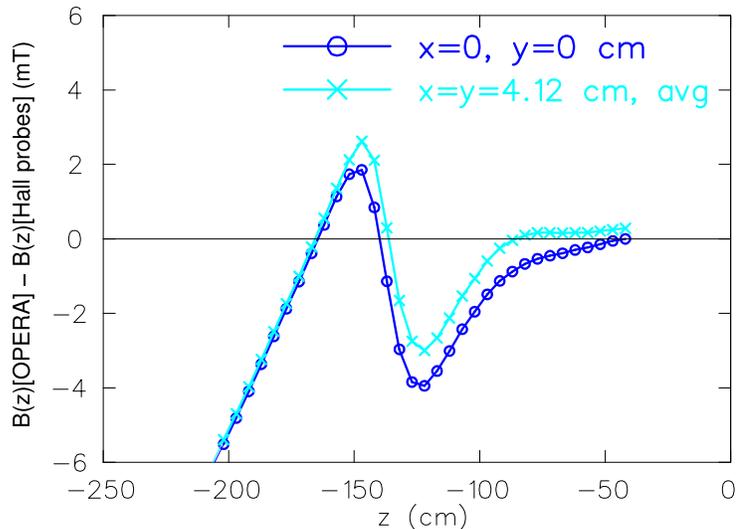


Figure 6.4: Difference in B_z between the magnetic field maps from `Opera` and the Hall probes. Two comparisons are shown: on-axis ($x = y = 0$) and an off-axis average of $x = \pm 4.12$ cm, $y = \pm 4.12$ cm. Note that muons start in the simulation at $z = -191.944$ cm, and stop in the target at $z = 0$ cm.

6.2.7 Muon beam tunes for evaluating the uncertainty

The nominal beam tune was described in Section 5.2; initially the position of the beam at the time expansion chambers (TECs) was steered to be close to $x = y = 0$, with the angles θ_x and θ_y minimised, but ultimately the tune was chosen to place the muon beam spots within the detector along a straight line, corresponding to minimised transverse momentum. The “envelope” of a simulated nominal muon beam, defined as the mean plus or minus one standard deviation, is shown in Figs. 6.5(a) and 6.5(b). Most of the beam remains within 1.0 cm of the solenoid’s axis, and is focussed by the field to a few millimetres in extent at $z = -100$ cm. After this focus the envelope develops oscillations in its mean position and size.

Three additional beam tunes are shown in Fig. 6.5. These will be used to assess the systematic uncertainties for the fringe field. They are set 76, where the muon beam was steered to have $\langle \theta_y \rangle \approx 28$ mrad at the TECs, set 86, where the beam was positioned off-axis and pointed away from the axis ($\langle x \rangle \approx -1.0$ cm, $\langle \theta_x \rangle \approx -10$ mrad), and set 72 where the TECs were in place throughout, which increased the muon beam’s emittance due to the additional multiple scattering. In the range where rapid depolarisation takes place, ($-150 < z < -100$) cm, the beams for sets 76 and 86 sample the fringe field at

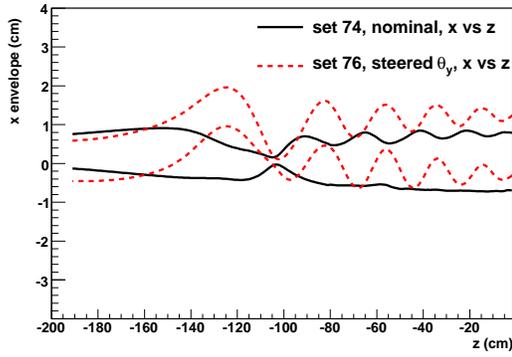
a different location to the nominal beam, with a non-zero average angle. They undergo a considerable focus, corresponding to an intersection of the field lines at a steep angle. As a result, these beams will be seen to undergo significantly more depolarisation, and the ability of the simulation to reproduce the depolarisation observed in the data will provide a stringent test of the fringe field model; such a validation was not available for the previous TWIST $P_\mu^\pi \xi$ analysis.

The three pairs of sets are listed in Table 6.4, and from here-on they are labelled as comparisons I, II and III. The first column in Table 6.4 shows the differences in polarisation from data. These are determined by fitting the nominal and steered decay spectra against each other, and the uncertainties are purely statistical. The next column in Table 6.4 shows the simulation’s prediction for the polarisation difference using the *Opera* magnetic field map. These are evaluated using a Monte Carlo technique that is similar to Section 6.2.5. For each Monte Carlo trial, the initial position and angle of the beam are selected independently for the nominal and steered set, since we believe that variations originate from the reproducibility of the TECs. In contrast, the magnetic field position and angle must be the same for a nominal and steered set, so the same random values are used for both sets in each Monte Carlo trial.

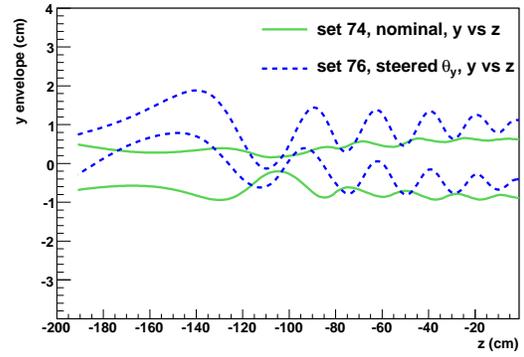
In terms of standard deviations, the simulation is seen to *underestimate* the depolarisation by 2.0σ and 0.6σ for comparisons I and II, and to *overestimate* the depolarisation by 1.0σ for comparison III. The result for comparison I would not usually be a cause for concern; however, there is good reason to believe that an additional systematic uncertainty from the B_x and B_y components is missing.

Table 6.4: Validation of the fringe field map by establishing how well the simulation reproduces large polarisation differences in the data. Ideally the value of simulation minus data should be consistent with zero.

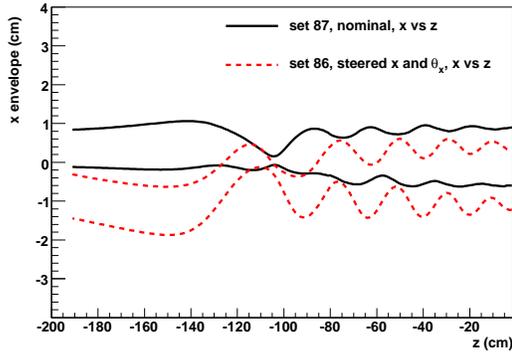
Comparison	Sets	Difference in polarisation ($\times 10^{-4}$)		
		Data	Simulation	Simulation minus data
I	74,76 (steered in θ_y)	105 ± 9	56_{-18}^{+23}	-48_{-20}^{+24}
II	87,86 (steered in x, θ_x)	62 ± 8	47_{-16}^{+23}	-15_{-18}^{+24}
III	74,72 (TECs-in)	18 ± 9	28_{-5}^{+12}	$+10_{-10}^{+15}$



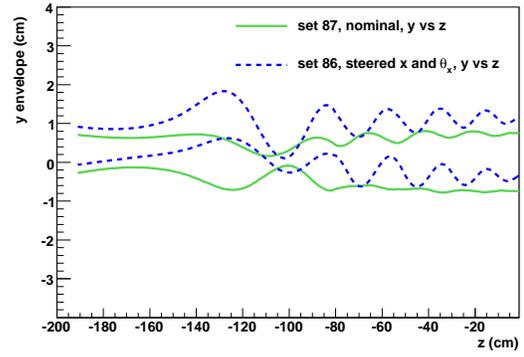
(a)



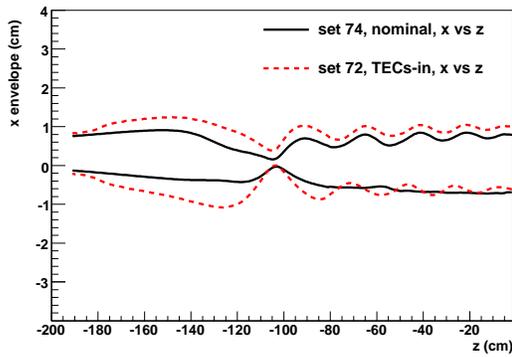
(b)



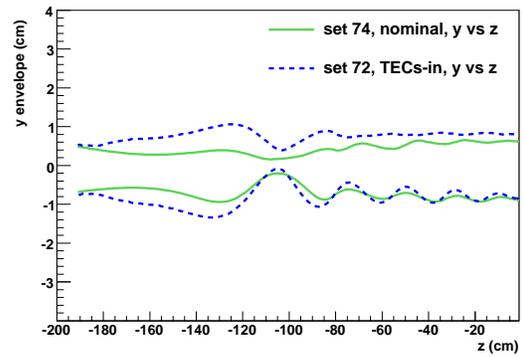
(c)



(d)



(e)



(f)

Figure 6.5: Muon beam *envelopes* from the simulation. These are defined as the mean position plus or minus one standard deviation.

Figures 6.3(a) and 6.3(c) show how the transverse field components depend on position. The muons are confined to within < 5 cm of the axis, and Fig. 6.3(c) indicates that the field has a linear dependence on x (and y) over this range. It is reasonable to expect that *Opera* has the correct linear shape, but may suffer from errors in the magnitude of the components. The sensitivity of comparison I to a scaling of B_x and B_y by a common factor, f , is shown in Fig. 6.6. The two sigma discrepancy between data and simulation can be reduced by increasing f , which increases the central value of the simulation, and enlarges the uncertainties in $P_\mu(0)$ from alignments.

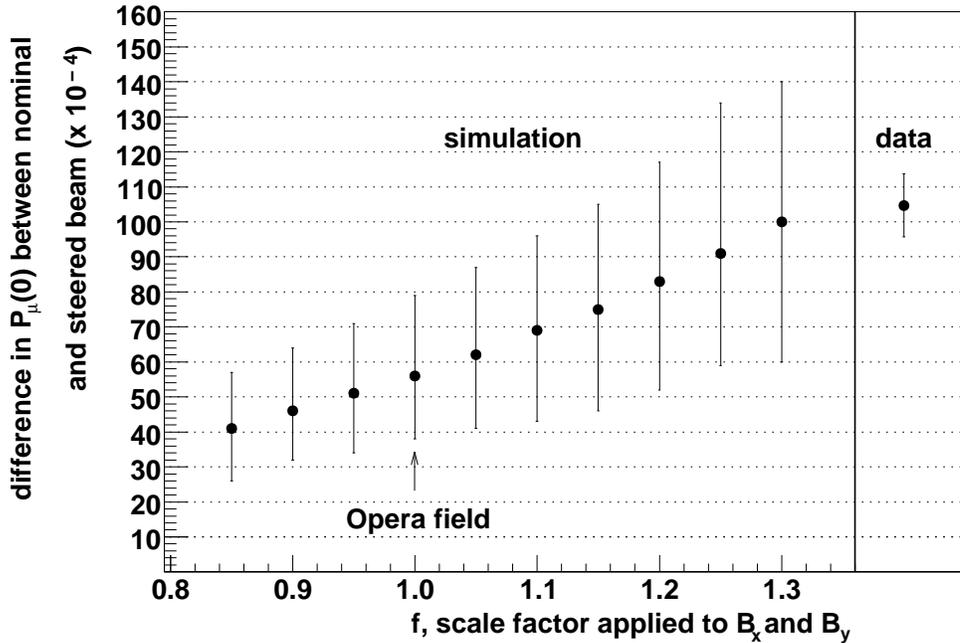


Figure 6.6: Comparison I is shown. With the original *Opera* field ($f = 1.0$) the data and simulation differ by about two standard deviations. The agreement improves as the B_x and B_y components are scaled by a factor, f .

The systematic uncertainty is evaluated by considering the effect of the scaling, f , on all three comparisons. The number of standard deviations by which data and simulation differ is shown in Fig. 6.7(a), as f is varied. A decrease in the transverse field components is limited by comparison I, since the simulation rapidly diverges from the data as f is decreased. Similarly, an increase in the transverse field components is limited by comparison III.

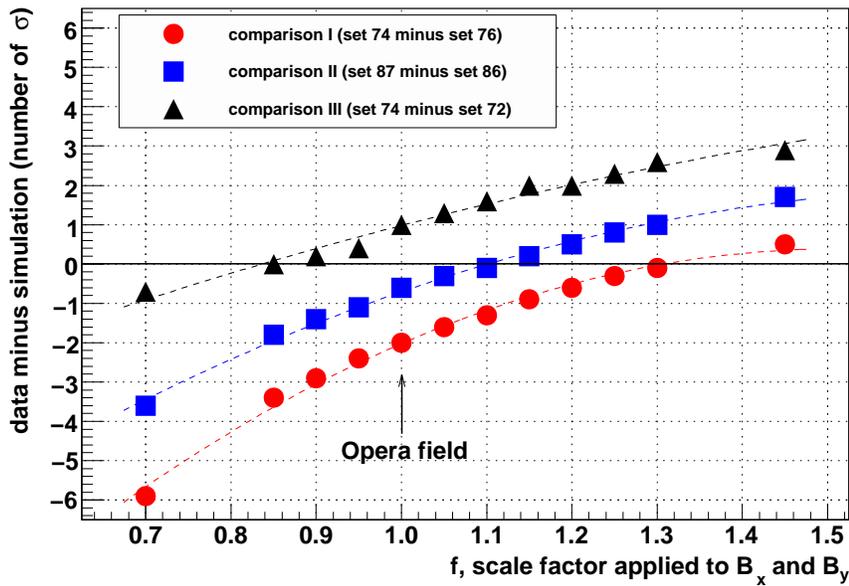
We have no reason to favour one of the comparisons over the others. Therefore we combine the differences between data and simulation for each comparison, and seek to

minimise this combination by scaling the transverse field components. Specifically the χ^2 statistic is calculated as

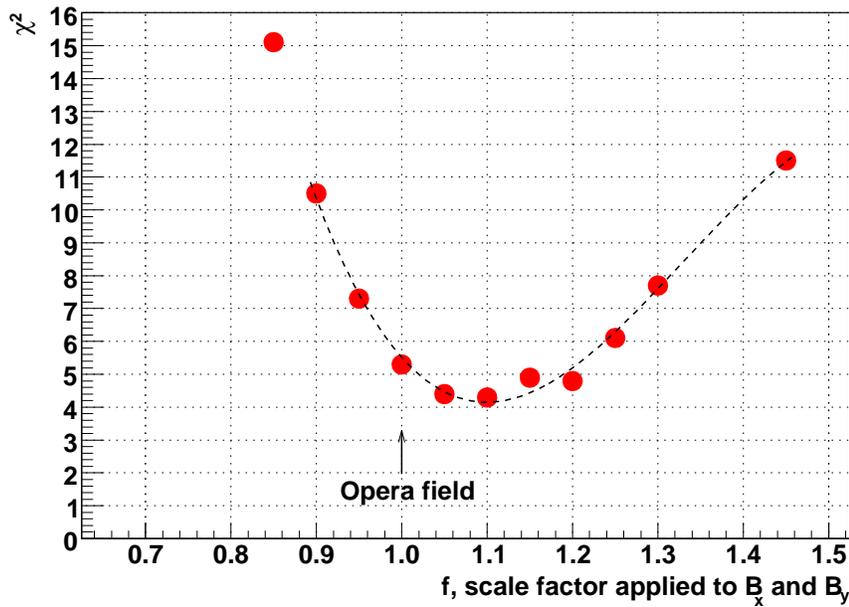
$$\chi^2 = \sum_{\text{comparisons}} \frac{(\bar{x}_{\text{data}} - \bar{x}_{\text{simulation}})^2}{\sigma_{\text{data}}^2 + \sigma_{\text{simulation}}^2} \quad (6.1)$$

where for convenience the definition $x = \Delta P_\mu(0)$ has been made, and $\Delta P_\mu(0)$ is the difference in polarisation between a nominal and a steered set. The resulting χ^2 statistic is shown in Fig. 6.7(b); a minimisation occurs when the transverse field components are scaled by $f = 1.10$. A correction of the result to this minimum is *not* made, since we cannot prove that the deviations between the **Opera** and true fields are fully covered by a simple scaling of the B_x and B_y components. Instead the f range corresponding to $(\chi_{\text{minimum}}^2 + 1)$ is taken to present a reasonable maximum of our error due to not measuring the transverse field components. This leads to a range in f of 1.01 to 1.20. All nominal sets are evaluated at these two f -values, and the average difference in $P_\mu(0)$ is 13.9×10^{-4} . This is assigned as a one sided uncertainty for $P_\mu(0)$ of ${}_{-13.9}^{+0.0} \times 10^{-4}$. The uncertainty for $\Delta P_\mu^\pi \xi$, the difference between the data and a hidden simulation value, is then ${}_{-0.0}^{+13.9} \times 10^{-4}$.

We note that early studies used a different field map to investigate the systematic uncertainty from the magnetic field. Three on-axis coils were added to the original **Opera** map, and their currents were tuned to almost completely eliminate the B_z mismatch in Fig. 6.4. The coils also increased the B_x and B_y components, and the effect on the nominal set polarisation was 14.1×10^{-4} , on average. This is very close to the result derived from the simple scaling approach.



(a) Data-simulation agreement for all three comparisons.



(b) Combining the three comparisons into a χ^2 .

Figure 6.7: The transverse magnetic field components (B_x and B_y) are multiplied by a factor, f . The difference in polarisation between a nominal and a steered set is sensitive to f . An overall better agreement is possible with $f = 1.1$, corresponding to a 10% increase in the strength of the transverse components.

6.2.8 Muon beam angular distribution width

The muons were multiple scattered as they passed through the TECs, resulting in a measured angle distribution that is larger than the distribution in the absence of the TECs. A GEANT3 simulation of the TECs finds that the root mean square of the angles should be reduced by a factor of $c_x = 0.64$ in the x -module, and $c_y = 0.48$ in the y -module to account for this multiple scattering. The correction in the y -module is larger since it is located downstream of the x -module. The dependence of $P_\mu(0)$ on the choice of c_x is shown in Fig. 6.8 for a nominal and a steered beam. To an acceptable approximation, $dP_\mu(0)/dc_x$ and $d^2P_\mu(0)/dc_x^2$ are independent of the beam steering. For example, if $c_x = 0.64$ then a variation of $\pm 10\%$ in c_x changes $P_\mu(0)$ by ${}_{+2.1}^{-1.8} \times 10^{-4}$ for the nominal case, and ${}_{+2.0}^{-1.9} \times 10^{-4}$ for the steered beam. As a result, the choice of c_x and c_y has no bearing on the comparison of polarisation differences between data and simulation, and the systematic uncertainties from these factors can be treated as orthogonal to those already evaluated.

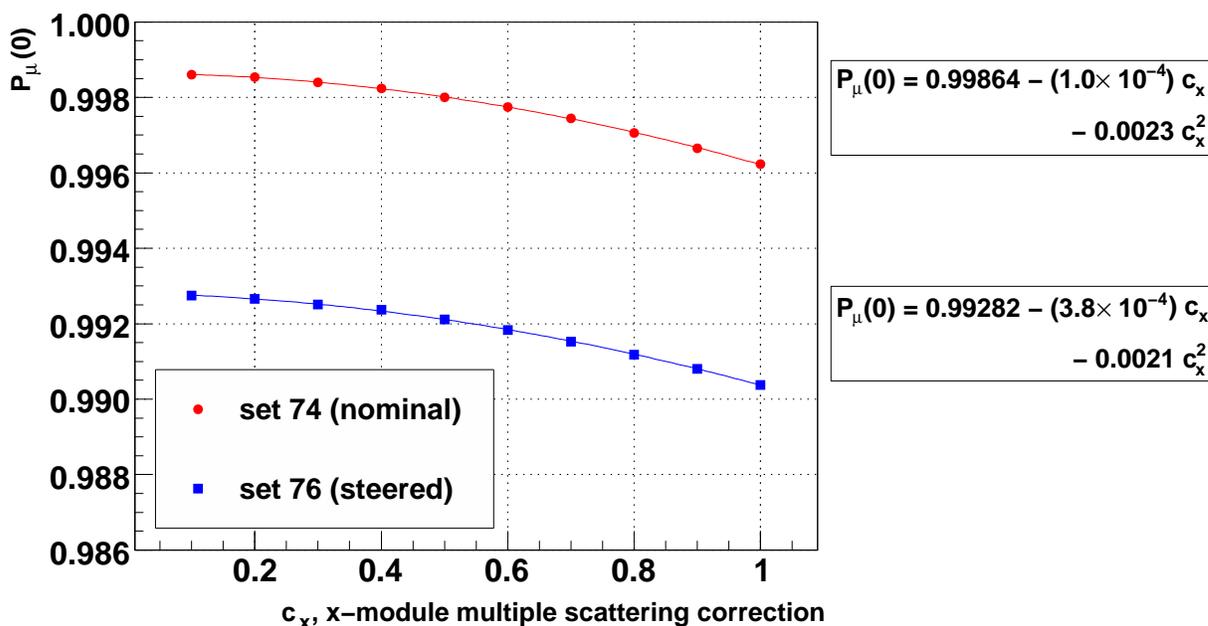


Figure 6.8: Sensitivity of the simulation's $P_\mu(0)$ to c_x , the multiple scattering correction factor in the x -module. The ratio $c_x/c_y = 0.64/0.48$ was maintained.

Note that although the simulation predicts the systematic uncertainties in this section are asymmetric, this asymmetry is at the 10^{-5} level, which is too small to be significant; the systematic uncertainties are therefore quoted as *symmetric*, with the magnitude set

to the average of the upper and lower bounds.

The parameters c_x and c_y relied on the accuracy of multiple scattering in **GEANT3**. The author is unaware of any validation studies for the multiple scattering of muons with $p \approx 30$ MeV/c in thin materials (the entire TEC apparatus was equivalent to just ≈ 7 mg/cm² of material). Our most direct test of the **GEANT3** accuracy used five runs with the upstream window on the TECs changed from the nominal $6 \mu\text{m}$ of Mylar to a thicker $25 \mu\text{m}$ (3.2 mg/cm²) window. The additional $19 \mu\text{m}$ of material increased the scattering distribution so that⁴⁰

$$\theta_{\text{TECs}+25 \mu\text{m}} \approx \sqrt{\theta_{\text{TECs}+6 \mu\text{m}}^2 + \theta_{19 \mu\text{m}}^2}. \quad (6.2)$$

The results for $\theta_{19 \mu\text{m}}$ from data and simulation are shown in Table 6.5, where the simulation overestimates the root mean square width of the scattering distributions by 18.3% in the x -module and 15.6% in the y -module. This implies that the c_x and c_y factors were reliable to 17.0% (the average of the overestimate in each module), which results in a systematic uncertainty of $\pm 3.1 \times 10^{-4}$. The observed discrepancy of 17.0% must not be taken as a formal validation of multiple scattering in **GEANT3**, since there are systematic uncertainties associated with the values in Table 6.5 that have not been evaluated.

Table 6.5: Width of reconstructed angle distributions for Mylar windows of thickness $6 \mu\text{m}$ and $25 \mu\text{m}$. These are placed on the upstream TEC. The bracketed number indicates the statistical uncertainty in the final digit.

Mylar window thickness (μm)	RMS of θ_x (mrad)		RMS of θ_y (mrad)	
	data	simulation	data	simulation
6	14.50 (5)	14.10 (7)	19.79 (7)	19.89 (6)
25	17.25 (8)	17.90 (9)	22.01 (6)	22.8 (1)
$\Rightarrow 19$	9.3 (2)	11.0 (2)	9.6 (2)	11.1 (2)

The TEC analysis code was reviewed for this measurement, and the accuracy of the reconstruction algorithm was found to be limited by noise from the electronics. This did not affect the mean position/angle, only the width of the angular distributions. An attempt to overcome this limitation resulted in two variants of the algorithm that are systematically different (see Appendix G.3.4). Since an event-by-event investigation could not distinguish which variant was the most accurate, the systematic difference between

⁴⁰The multiple scattering distribution is non-Gaussian, and adding the widths of layers in quadrature will systematically underestimate the total width[3]. However, in this case we are comparing data and simulation where the same error is made in both.

the two is taken as a systematic uncertainty. For all sets this difference in $P_\mu(0)$ was less than 1.7×10^{-4} , except for set 76 (steered) where $P_\mu(0)$ changed by 6.3×10^{-4} between the two variants of the algorithm. A conservative systematic uncertainty of $\pm 1.7 \times 10^{-4}$ is assigned for the nominal sets.

The width of the angular distributions depended on the mean number of hits in the final track ($\langle n_x \rangle$ in the x -module, $\langle n_y \rangle$ in the y -module), which decreased depending on the length of time that the sense planes were exposed to the beam. The same c_x and c_y correction factors were used for all muon beam measurements, despite differences in the age of the planes, and this resulted in a systematic uncertainty. The values of $\langle n_x \rangle$ and $\langle n_y \rangle$ for each set are shown in Fig. 6.9. The c_x and c_y factors were tuned using set 75, which had $\langle n_x \rangle = 15.0$ and $\langle n_y \rangle = 15.9$. For all the sets, the ranges of $\langle n_x \rangle$ and $\langle n_y \rangle$ were 13.0 to 16.7, and 14.4 to 18.5 respectively, which is almost symmetric about the values used for tuning. The set 75 data were reanalysed, with hits removed at random to reduce $\langle n_x \rangle$ to 13.0 and $\langle n_y \rangle$ to 16.7. The root-mean-square decreased by 7.9% in the x -module, and 3.7% in the y -module. If the larger of these is used, then a systematic uncertainty for the nominal sets due to sense plane aging is $\pm 1.5 \times 10^{-4}$.

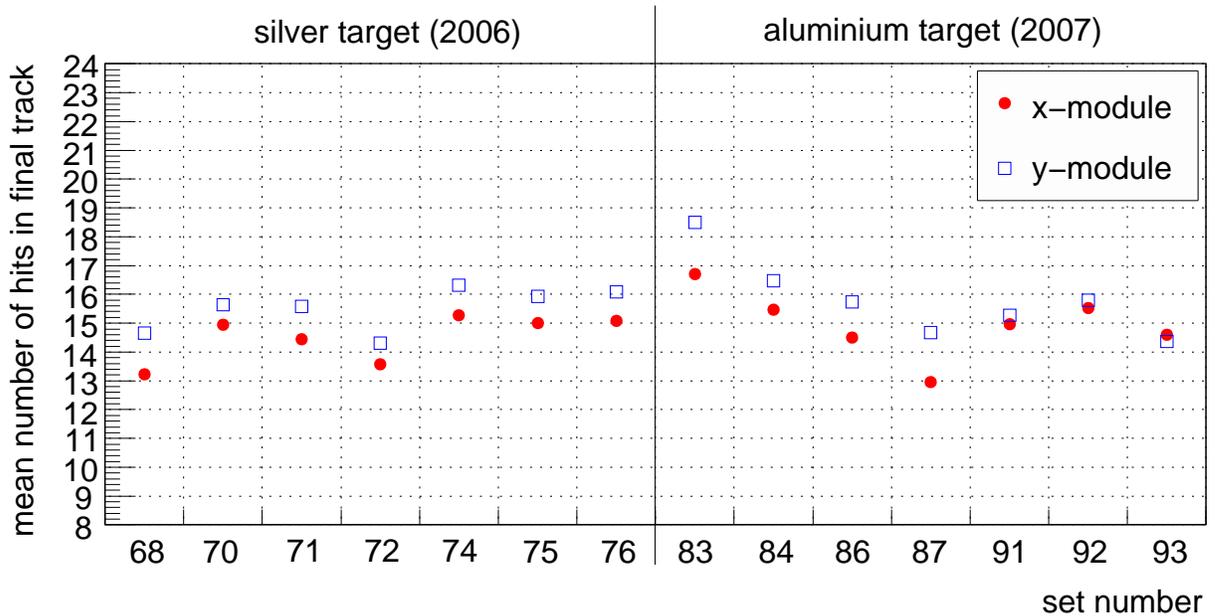


Figure 6.9: Mean number of hits in final TEC track, for each module. The c_x and c_y correction factors were tuned using set 75, which had $\langle n_x \rangle = 15.0$ and $\langle n_y \rangle = 15.9$.

6.2.9 Summary of muon beam and fringe field uncertainties

First the discrepancies between beginning and end of set muon beam measurements were considered (Section 6.2.3). These are blamed on a problem with the reproducibility of the TECs' alignment, rather than a muon beam instability or a variation in the TEC drift cell response. The resulting alignment uncertainties are ± 0.2 cm for position and ± 3 mrad for angle. Second, the alignment of the magnetic field in position and angle was considered (Section 6.2.4), leading to uncertainties of ± 0.1 cm for position and ± 1 mrad for angle. These are combined in Section 6.2.5 using a Monte Carlo approach, and produce an uncertainty of ${}_{-1.2}^{+6.4} \times 10^{-4}$ for $\Delta P_{\mu}^{\pi} \xi$.

The accuracy of the depolarisation through the magnetic field relies on the simulation's ability to reproduce large polarisation changes in three tests. The polarisation was lowered by either steering the muon beam away from the solenoid's symmetry axis (comparison I and II), or by inserting the TECs so that the muon beam's emittance was increased (comparison III); see Section 6.2.7 for a full description. A Monte Carlo approach was used to see how well the simulation reproduced the polarisation differences in the data. Using the original `Opera` field map, the largest discrepancy between data and simulation was 2.0σ , which would not usually be a cause for concern. However, the `Opera` map is expected to differ from reality due to the quality of the initial conditions given, and the accuracy of the finite element method that is used to solve Maxwell's equations. Also the transverse field components (B_x and B_y) from `Opera` could not be validated directly since they were not measured. Therefore although we believe that the original `Opera` map is close to the truth, we must assign an uncertainty that covers a potential error in the field map. This was evaluated by scaling the transverse magnetic field components, and observing where the agreement in the comparisons is minimised. The resulting uncertainty for $\Delta P_{\mu}^{\pi} \xi$ is ${}_{-0.0}^{+14.1} \times 10^{-4}$.

Finally there are orthogonal systematic uncertainties from the degree to which the angle distributions of the muons could be reconstructed. There are contributions from the accuracy of multiple scattering in `GEANT3` ($\pm 3.1 \times 10^{-4}$), from the algorithm used to reconstruct the muon trajectories ($\pm 1.7 \times 10^{-4}$), and from the aging of the sense planes in the TECs ($\pm 1.5 \times 10^{-4}$).

6.2.10 Stopping material

Overview

About 80% of the muons stopped in a metal target, which also served as the shared cathode foil for the proportional chambers PC6 and PC7 (see Fig. 2.13). Events were accepted only if the muon produced a signal in PC6, but not in PC7. Muons that stopped in the PC6 gas (a mixture of CF_4 and isobutane) or the wires were removed by cutting on the muon pulse width in PC5 and PC6 (see Section 3.3.3). This selected a sample of muons that mostly stopped in the metal foil. The relaxation was determined using a weighted asymmetry measurement (see Section 3.6). Each data set was fit with

$$P_\mu(t) = P_\mu^*(0) \exp(-\lambda t). \quad (6.3)$$

The relaxation rate from the data was used in the simulation.

There is one correction, one statistical uncertainty and four systematic uncertainties related to the stopping material. A correction to $\Delta P_\mu^\pi \xi$ is necessary since the simulation used an incorrect λ value that was derived from an early analysis; re-generating the simulation would have required several months, so a correction must be made instead. A statistical uncertainty results from the precision with which λ can be determined from the data ($\pm 2.4 \times 10^{-4}$). There is a systematic uncertainty from the degree to which stops in the PC6 chamber gas are eliminated ($\pm 0.3 \times 10^{-4}$), and another uncertainty from muons that pass through the target and stop in PC7, but do not have enough energy to produce a signal ($\pm 0.9 \times 10^{-4}$). A small fraction of muons are scattered from the target back into PC6, resulting in a systematic uncertainty of $\pm 0.2 \times 10^{-4}$. Lastly, there is an indication of a small bias in the asymmetry analysis, and this introduces a systematic uncertainty of ($\pm 3.0 \times 10^{-4}$). After adding in quadrature the total systematic uncertainty is $\pm 3.2 \times 10^{-4}$. Note that for this analysis there is no uncertainty due to the choice of model for $P_\mu(t)$ in the metals; see Section 1.6.3 for the theoretical arguments that demonstrate a strong preference for a single exponential form.

Correction and statistical uncertainty

The results of the asymmetry analysis applied to the data are shown in Table 6.6, where the fit is made over the nominal time range of $(1.05 < t < 9.00) \mu\text{s}$. A weighted average

of these relaxation rates finds

$$\lambda_{\text{Ag}}^{\text{data}} = (0.813 \pm 0.079) \text{ ms}^{-1}, \quad (6.4)$$

$$\lambda_{\text{Al}}^{\text{data}} = (1.279 \pm 0.086) \text{ ms}^{-1}. \quad (6.5)$$

Note that these are consistent with the μ^+ SR results from Appendix H.8:

$$\lambda_{\text{Ag}}^{\mu\text{SR}} = (0.9 \pm 0.2 \text{ (stat.)} \pm 0.2 \text{ (syst.)}) \text{ ms}^{-1}, \quad (6.6)$$

$$\lambda_{\text{Al}}^{\mu\text{SR}} = (1.3 \pm 0.2 \text{ (stat.)} \pm 0.3 \text{ (syst.)}) \text{ ms}^{-1}. \quad (6.7)$$

The simulation used incorrect values of $\lambda_{\text{Ag}} = 0.725 \text{ ms}^{-1}$ and $\lambda_{\text{Al}} = 1.169 \text{ ms}^{-1}$, and as a result the quantity $\Delta P_{\mu}^{\pi} \xi$ must be corrected. The effect on the spectrum of a change in λ is calculated using

$$\frac{\int_{t_1}^{t_2} N(t) \cdot P_{\mu}^*(0) \exp(-\lambda_2 t) dt}{\int_{t_1}^{t_2} N(t) dt} - \frac{\int_{t_1}^{t_2} N(t) \cdot P_{\mu}^*(0) \exp(-\lambda_1 t) dt}{\int_{t_1}^{t_2} N(t) dt}, \quad (6.8)$$

where $N(t) = N(0) \exp(-t/\tau_{\mu})$, τ_{μ} is the muon lifetime, and λ_1 and λ_2 are the relaxation rates between which the correction is being made. The time range is the same as the normal analysis ($1.05 \mu\text{s} < t < 9.00 \mu\text{s}$). The common $P_{\mu}^*(0)$ factor is close to 1.0, and its choice has a negligible impact on the correction. The quantity $\Delta P_{\mu}^{\pi} \xi$ (the difference between the data and a hidden simulation value) is then corrected by $+2.7 \times 10^{-4}$ for Ag ($\lambda_1 = 0.725 \text{ ms}^{-1}$, $\lambda_2 = 0.813 \text{ ms}^{-1}$) and $+3.3 \times 10^{-4}$ for Al ($\lambda_1 = 1.169 \text{ ms}^{-1}$, $\lambda_2 = 1.279 \text{ ms}^{-1}$). The statistical uncertainty in determining λ from the data causes a P_{μ} uncertainty of $\pm 2.4 \times 10^{-4}$ for both targets, again using Eq. (6.8).

Table 6.6: Relaxation rate λ for each data set. The model $P_\mu(t) = P_\mu^*(0) \exp(-\lambda t)$ has been fit over the nominal time range of $(1.05 < t < 9.00) \mu\text{s}$.

Set num.	Target	Description	λ (ms^{-1})	χ^2/ndof	Fit quality confidence
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.71 ± 0.24	1.03	0.43
70	Ag	B = 1.96 T	0.89 ± 0.19	0.62	0.90
71	Ag	B = 2.04 T	0.93 ± 0.20	1.58	0.05
72	Ag	TECs-in, nominal beam	0.85 ± 0.19	1.03	0.42
74	Ag	Nominal A	1.02 ± 0.24	0.63	0.90
75	Ag	Nominal B	0.85 ± 0.19	0.72	0.81
76	Ag	Steered beam A	0.38 ± 0.22	0.52	0.96
	Ag	Weighted average:	0.813 ± 0.079		
83	Al	Downstream beam package in place	1.46 ± 0.19	1.85	0.01
84	Al	Nominal C	1.10 ± 0.21	1.08	0.36
86	Al	Steered beam B	1.21 ± 0.18	1.08	0.36
87	Al	Nominal D	1.22 ± 0.20	0.85	0.65
91	Al	Lower momentum I	1.63 ± 0.40	1.06	0.39
92	Al	Lower momentum II	1.35 ± 0.35	0.98	0.49
93	Al	Lower momentum III	1.32 ± 0.27	0.75	0.78
	Al	Weighted average:	1.279 ± 0.086		

Muon stops in gas

The distribution of the pulse widths in the PCs immediately before the target is shown in Fig. 6.10. The majority of the distribution is made up of stops in the metal target, which primarily contribute to zones 1 and 2, and there is a “gas band” that runs approximately parallel to the cut-B line and contributes mostly to zones 2 and 3. The analysis selects only zone 1; the cut-B line removes almost all of the gas contamination, but a small fraction still makes it to zone 1.

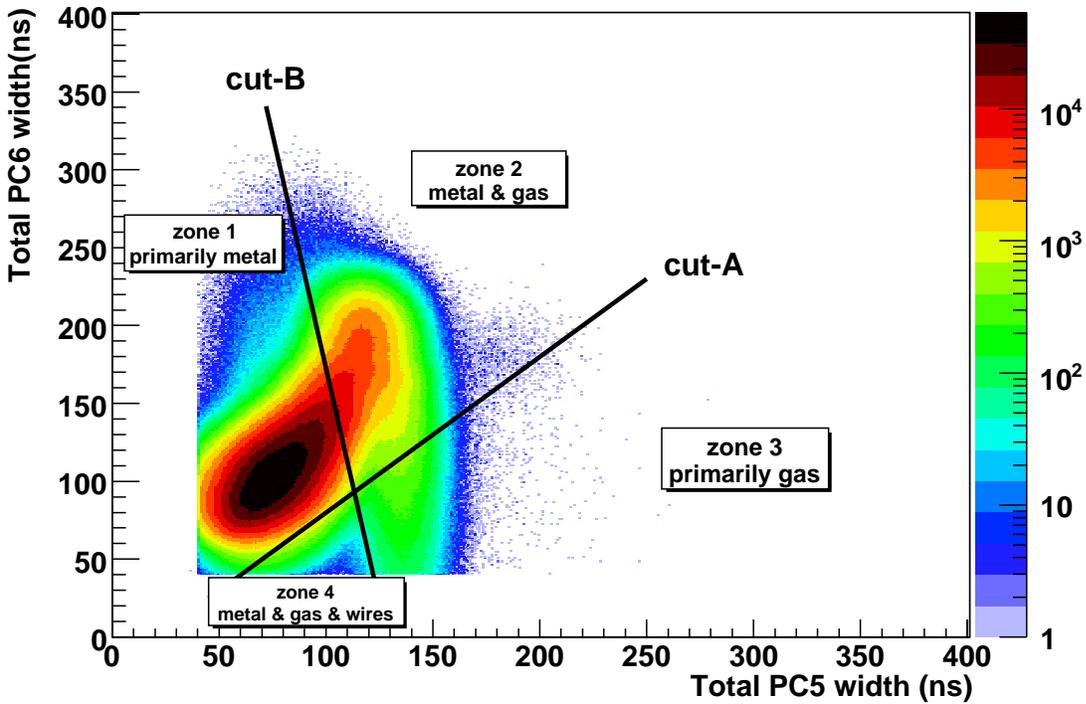


Figure 6.10: Muon pulse widths in PC5 and PC6 (the proportional chambers immediately before the metal stopping target). For data, the stops in gas appear as a band that is approximately parallel to cut-B, primarily contributing to zones 2 and 3.

A systematic uncertainty was estimated using only the data. This required three numbers that are described in Appendix J: the depolarisation in gas (8.0%), the fraction of the total PC6+target stops in gas (f_g), which is between 4% and 7% depending on the set, and the fraction of the gas distribution in zone 2 that leaks into zone 1 (tuned to be < 0.5% by adjusting cut-B). If the number of muons stopping in the gas in zone i is

denoted N_i^g , then

$$f_g N_{\text{total}} = N_1^g + N_2^g + N_3^g + N_4^g, \quad (6.9)$$

$$N_1^g < 0.5\% \times (N_1^g + N_2^g), \quad (6.10)$$

which gives

$$\frac{N_1^g}{N_{\text{total}}} < 0.5\% \times \left[f_g - \frac{N_3^g + N_4^g}{N_{\text{total}}} \right]. \quad (6.11)$$

If all the gas stops are conservatively assumed to be in zones 1 and 2 so that $N_3^g = N_4^g = 0$, and f_g is set to its maximum value of 7%, then an upper limit on the systematic uncertainty from muons stopping in PC6 gas is $8.0\% \times 0.5\% \times 7\% = 0.3 \times 10^{-4}$.

A special simulation found that no more than 0.11% of muons entered PC7 (the wire chamber *after* the metal stopping target), but did not have enough energy to produce a signal. These stop in the PC gas, and experience a depolarisation of 8.0% (see above). An upper limit for the systematic uncertainty due to these stops is therefore $0.11\% \times 8.0\% = 0.9 \times 10^{-4}$. The same simulation determined that 0.02% of muons enter the target volume, but are scattered back into PC6, and pass all the analysis cuts. These contribute a negligible systematic uncertainty of $0.02\% \times 8.0\% = 0.2 \times 10^{-4}$.

Analysis bias

The asymmetry analysis is applied to the simulation to confirm that the input values of λ can be recovered ($\lambda_{\text{Ag}}^{\text{input}} = 0.725 \text{ ms}^{-1}$ and $\lambda_{\text{Al}}^{\text{input}} = 1.169 \text{ ms}^{-1}$). The results of this study are shown in Fig. 6.11, where the simulation that accompanies set 74 (nominal) is seen to have a λ value that is 3.4 standard deviations below the input value. As a result, the *weighted average* of λ for the silver simulations is 2.4 standard deviations from the input; the weighted average for the aluminium simulations is consistent with the input λ .

An exhaustive investigation attempted to find the cause of the discrepancies for the silver simulations. The results were found to be robust to the binning, weighting strategy and time range used in the asymmetry analysis. The confidence levels from the asymmetry analysis were spread evenly between 0 and 1, indicating that the λ uncertainties are properly determined. A separate analysis that made fits to the upstream and downstream time distributions (*i.e.* an asymmetry was not constructed) confirmed the λ results, but with inferior statistical precision. The simulation for set 74 only differed from the other nominal sets in its input beam profile and muon/positron rates; the beam profile produced a reasonable value for $P_\mu(0)$, and the rate differences from the other sets were insignificant.

We note that an additional statistically independent simulation of set 74 (not used here in the weighted averages) produced $\lambda = (0.59 \pm 0.17) \text{ ms}^{-1}$, which is consistent with both the input value ($\lambda = 0.725 \text{ ms}^{-1}$) and the anomalous result from the earlier set74 simulation ($\lambda = (0.32 \pm 0.12) \text{ ms}^{-1}$); clearly significant additional statistics are needed to make a firm conclusion. Since the investigation could not find a reason to reject the anomalous λ result, we are forced to admit the possibility that a systematic λ uncertainty exists due to an analysis bias, or the λ statistical uncertainties from the asymmetry analysis are underestimates.

We are not aware of a credible mechanism that would create a different bias for the silver and aluminium targets. Therefore the average discrepancy between the input and measured λ from the simulation ($0.5 \times (2.4 + 0.0)\sigma = 1.2\sigma$) is taken as an additional uncertainty that must be applied to the *data*. Using Eqs. (6.4), (6.5) and (6.8), this results in a systematic uncertainty of $\pm 2.9 \times 10^{-4}$ for silver and $\pm 3.1 \times 10^{-4}$ for aluminium. The average of these ($\pm 3.0 \times 10^{-4}$) is applied to all sets as the systematic uncertainty due to a potential analysis bias.

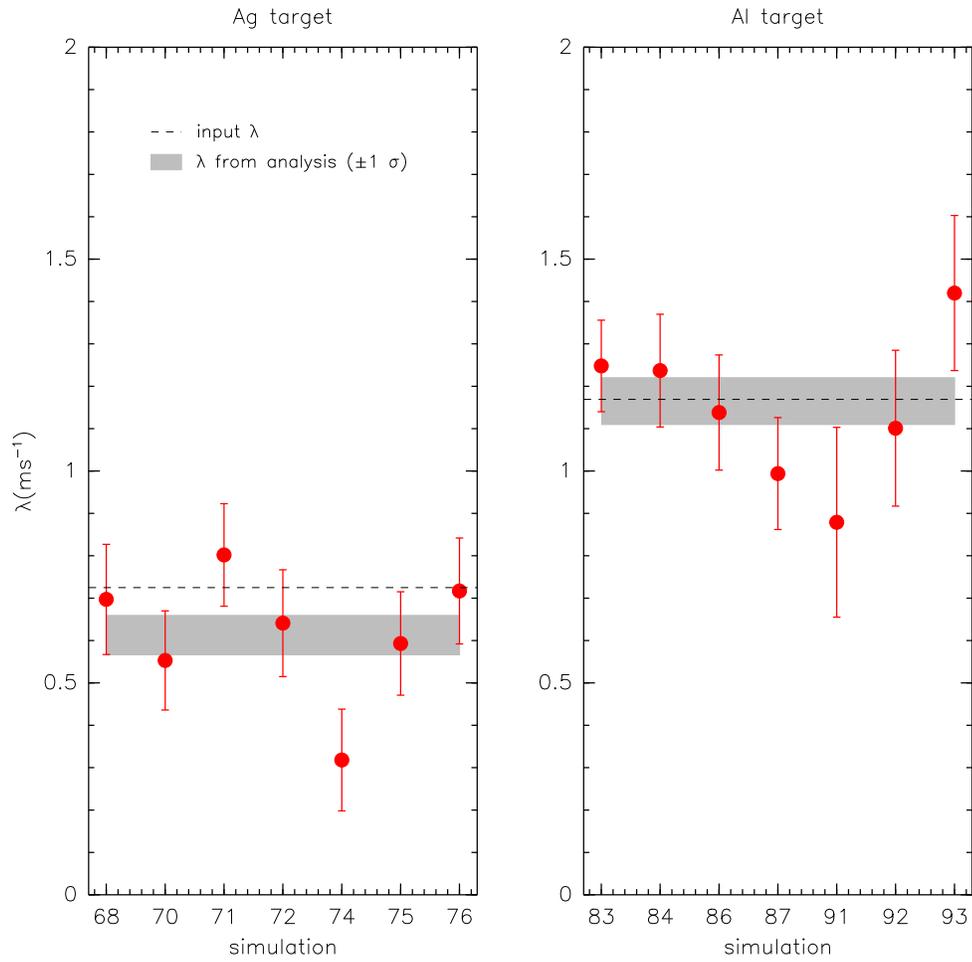


Figure 6.11: The result of the asymmetry analysis applied to the simulations[103]. The relaxation rate (λ) for the set 74 simulation is 3.4σ away from the input value.

6.2.11 Muon production target

The simulation generates muons with anti-parallel spin and momentum vectors, starting from the end of the M13 beam line. This neglects multiple scattering in the muon production target and the beam line vacuum window, which changes the momentum vector but not the spin. This is treated here as a systematic correction with an associated uncertainty. (The difference in precession frequencies of the momentum and spin through the M13 beam line is neglected since it introduces an error of $< 10^{-8}$; see Appendix I).

Surface muons are produced with $p \approx 29.79$ MeV/c, but the beam line was nominally tuned to accept muons with an average momentum of $\langle p \rangle = 29.6$ MeV/c. Therefore the muons lost 0.19 MeV/c of momentum on average, which is equivalent to ≈ 3.8 mg/cm² in graphite. (The 3 μ m beam line vacuum window can be safely neglected since it corresponds to just 0.3 mg/cm² of material.) The width of the resulting multiple scattering distribution⁴¹, θ_0 , is found to be 9.3 mrad using a **GEANT4** simulation[103]. The uncertainty is conservatively estimated as $\pm 17\%$ based on the observed discrepancy for **GEANT3** from Section 6.2.8. The degree to which the spin is depolarised with respect to the momentum is then estimated by $\cos(\theta_{\text{space}}^{\text{rms}})$, where $\theta_{\text{space}}^{\text{rms}} = \sqrt{2} \theta_0$. This results in a correction to the simulation's P_μ of $(-0.9 \pm 0.3) \times 10^{-4}$, which is a correction to $\Delta P_\mu^\pi \xi$ of $(+0.9 \pm 0.3) \times 10^{-4}$.

The **GEANT4** simulations are repeated for the lower momentum sets. The widths are $\theta_0 = 24.2$ mrad for $\langle p \rangle = 28.75$ MeV/c, and $\theta_0 = 22.9$ mrad for $\langle p \rangle = 28.85$ MeV/c. These correspond to $\Delta P_\mu^\pi \xi$ corrections of $(5.9 \pm 1.9) \times 10^{-4}$ and $(5.2 \pm 1.6) \times 10^{-4}$, respectively. The uncertainty for this correction will be applied set-dependently in Chapter 7, although in practice this detail makes no difference to the final answer.

⁴¹ θ_0 is the standard deviation of a Gaussian fit to the central 98% of the plane-projected multiple scattering distribution.

6.2.12 Background muon contamination

The true muon stopping distribution is not available for the data, and we can only measure the last plane that registered a muon hit. In the previous $P_\mu^\pi \xi$ analysis this distribution did not agree well in data and simulation unless an additional source of pion decays was included in the upstream beam package area (see Fig. 6.12(a)). An upper limit on the effect of these “background muons” was calculated as 1.8×10^{-4} , assuming they have opposite polarisation to the surface muons. For the current analysis the agreement is better *without* the additional pion source; this is demonstrated in Fig. 6.12(b), where improvements in the classification have increased the number of downstream events in the simulation relative to the data. If anything, the simulation now has *more* downstream muons for planes 35 onwards⁴², and even this excess is at the level of $< 10^{-5}$. Therefore the background muon contamination is now reduced to a negligible level, and this part of the systematic uncertainty is zero.

However, the data-simulation discrepancies in planes 35 to 53 (and in planes 10 to 20) do introduce a systematic uncertainty; the simulation’s stopping distribution must match the data, since high angle muons that undergo more depolarisation are preferentially stopped further upstream. Figure 6.12(b) indicates that the mean stopping position and/or its shape are not reproduced by the simulation. The most accurate measurement of the mean muon stopping distribution comes from the energy calibration procedure (see Section 3.5). The reconstructed spectrum endpoint depends on the thickness of target (and detector) material traversed, which results in a dependence on $1/\cos\theta$. The gradient of this relationship, α_{up} or α_{down} for upstream and downstream respectively, is therefore a measure of the material traversed by the upstream and downstream decay positrons. The stopping distribution can then be measured using

$$\alpha_{\text{diff.}} = \alpha_{\text{up}} - \alpha_{\text{down}}. \quad (6.12)$$

The energy calibration procedure finds the difference in reconstructed spectrum endpoint between data and simulation, which has a gradient equal to

$$\alpha_{\text{up}}^{\text{data}} - \alpha_{\text{up}}^{\text{sim.}}, \quad (6.13)$$

⁴²Planes 53 to 56 in Fig. 6.12(b) suggest there are more muons in the data than the simulation. However, these are beam positrons that are mistakenly identified as muons in the data due to an imperfect separation of pulse width in the proportional chambers. This does not occur in the simulation, where the separation between muons and beam positrons is perfect.

in the upstream half of the detector, and

$$\alpha_{\text{down}}^{\text{data}} - \alpha_{\text{down}}^{\text{sim.}} \quad (6.14)$$

in the downstream half. The difference between the gradients in the upstream and downstream half (*i.e.* Eq. (6.13) minus Eq. (6.14)) is then

$$(\alpha_{\text{up}}^{\text{data}} - \alpha_{\text{down}}^{\text{data}}) - (\alpha_{\text{up}}^{\text{sim.}} - \alpha_{\text{down}}^{\text{sim.}}) = \alpha_{\text{diff.}}^{\text{data}} - \alpha_{\text{diff.}}^{\text{sim.}}, \quad (6.15)$$

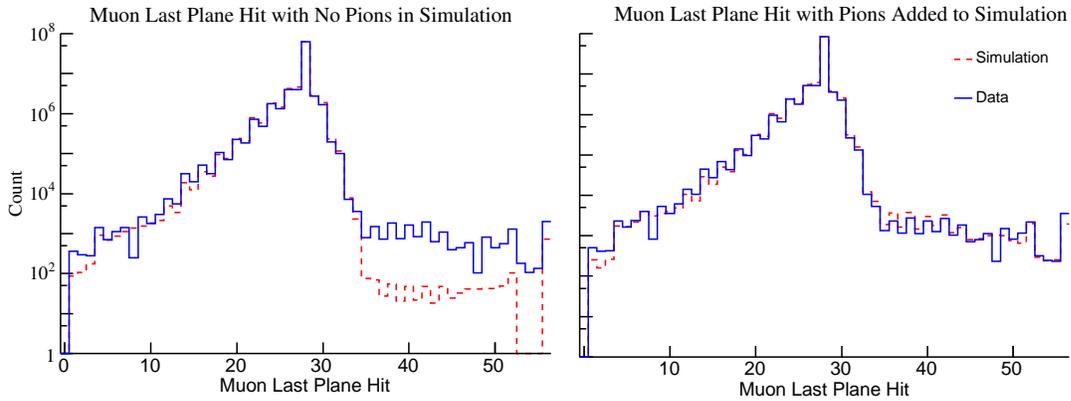
which measures how well the muon stopping distributions agree in data and simulation.

The stopping distribution measure from Eq. (6.15) is shown in Fig. 6.13. For five of the simulations, the mean of the muon last plane distribution was matched by adding an extra 1.9 mg/cm² of material to the simulation (see Section 2.11), and the resulting disagreement is (5.5 ± 1.4) keV/c on average. For the other simulations no additional material was added, and this improved the data-simulation disagreement to (0.2 ± 1.1) keV/c (We note that although $\alpha_{\text{diff.}}^{\text{data}} - \alpha_{\text{diff.}}^{\text{sim.}}$ is an accurate measurement of the stopping distribution, it is clearly not precise. For this reason it could not be used to tune the simulation so that it matched the data.)

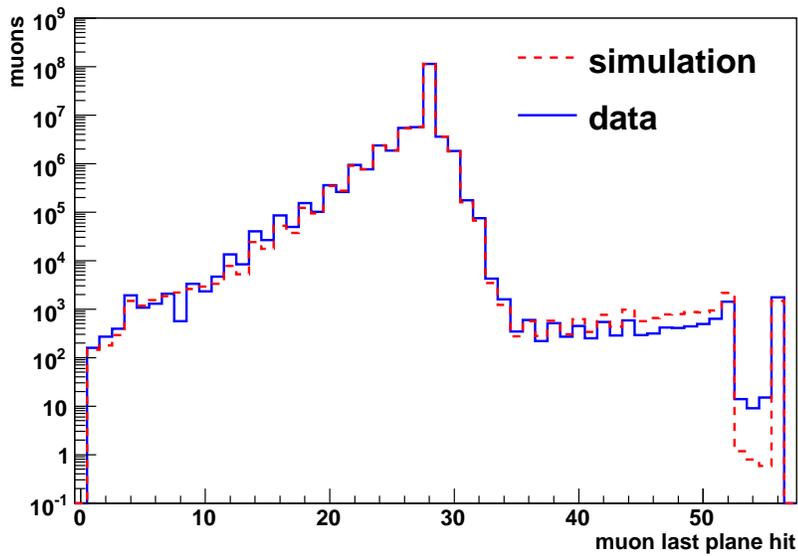
The simulation can determine the relationship between $P_{\mu}(0)$ and $(\alpha_{\text{diff.}}^{\text{data}} - \alpha_{\text{diff.}}^{\text{sim.}})$. This depends on the beam profile, and is within the range $0.04 \times 10^{-4}/(\text{keV}/c)$ to $0.25 \times 10^{-4}/(\text{keV}/c)$, with an average of $0.13 \times 10^{-4}/(\text{keV}/c)$. The systematic uncertainty from matching the stopping distribution is therefore $\pm 5.5 \times 0.13 \times 10^{-4} = \pm 0.7 \times 10^{-4}$ for the sets with the extra 1.9 mg/cm² of material. This result is an approximation since it averages over the silver and aluminium targets, where differences are expected in the data-simulation agreement. Since this is a small number and only applies to 5 of the 14 simulation, it is treated as a conservative systematic uncertainty rather than a correction. An additional piece of the same magnitude is added in quadrature to account for a small difference in stopping distribution shape, resulting in a total systematic uncertainty of $\pm 1.0 \times 10^{-4}$.

The remaining systematic uncertainties are related to the decay positron reconstruction. They are mostly evaluated by exaggerating a setting in the simulation or analysis. The exaggerated spectrum is then fit to the original spectrum in order to determine the change in the muon decay parameters (MPs). The exaggeration factors are made as large as possible to obtain statistically meaningful MP changes, while maintaining a linear relationship with the MPs. The changes in MPs are then scaled down according to how large

the effect could actually be, resulting in the systematic uncertainty. When the original and exaggerated spectra are highly correlated (i.e. they contain a large number of events with identical energy and angle), the spectrum fit will have a reduced χ^2 that is much smaller than one, and the uncertainty in the $P_\mu^\pi \xi$ change will therefore be too large. In this case the uncertainties in the MPs are scaled down so that the reduced χ^2 is equal to one, corresponding to multiplication by $\sqrt{\chi^2/\text{ndof}}$.



(a) Distribution of last plane that measures the muon, for data and simulation, from the previous analysis (originally Fig. 6.9 from Ref. [57]). The mismatch is resolved by including an additional source of pion decays at the location of the upstream beam package.



(b) The same figure re-made for the current analysis, where *no* additional pion decays are added.

Figure 6.12: Background muon contamination in the two $P_{\mu}^{\pi} \xi$ analyses.

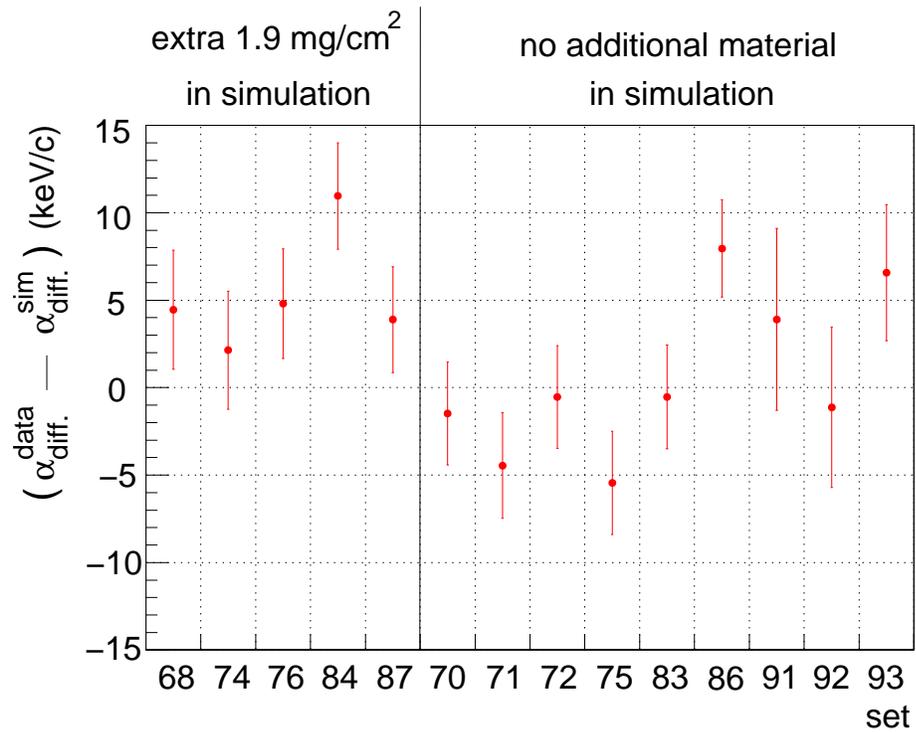


Figure 6.13: The muon stopping distribution is most accurately measured in data and simulation using the endpoint calibration; see the main text for a description of the ordinate. The data and simulation distributions are in agreement when $\alpha_{\text{diff.}}^{\text{data}} - \alpha_{\text{diff.}}^{\text{sim.}} = 0$.

6.3 Chamber response

6.3.1 Drift chamber space-time relationship

The space-time-relationships (STRs) in the drift cells were optimised by minimising the residual between the hit times from the drift cell, and the times that best fit the helix trajectories (see Section 3.2.7). In the simulation, where it was sufficient to use a single STR cell for all wires and planes, this procedure effectively absorbs a small bias from the helix fitting into the STRs. In data, where a separate STR cell was obtained for each plane, the procedure corrects for plane-to-plane construction and response differences, in addition to any small bias from the helix fitting algorithm.

The refinement procedure was carried out iteratively, with the STR forced to remain smooth at each step. After convergence, there were residuals in the drift cell corresponding to regions where manipulating the STRs could not bring the drift time closer to the fitted trajectory. The amount by which these residuals differ in data and simulation is the basis of the chamber response systematic uncertainty.

The difference between the data and simulation residuals at the final iteration is shown in Fig. 6.14. The data results are averaged over all planes. Only half a cell is shown, and in practice this is reflected about the line $uv = 0$. The largest discrepancies between data and simulation are at the edge of the cell (u or $v = 1.8$ mm), where there were low statistics and the single hit resolution was degraded. In the rest of the cell the variations are at the impressive level of < 4 ns.

The following approach exaggerated the differences between data and simulation without breaking the smoothness of the STRs:

1. Generate a separate refined STR cell for each plane in the simulation.
2. For each plane, find the difference in residuals at the final iteration, just like Fig. 6.14. Fit this distribution with a fifth order polynomial function.
3. Exaggerate the polynomial function until the χ^2/ndof of positron helix fits becomes a factor of two worse. This corresponds to the single hit resolution being degraded by $\approx (1 - \sqrt{2}) \approx 40\%$, and requires the polynomial function to be exaggerated by a factor of ten.
4. Re-analyse the simulation with the scaled STRs, and compare to the nominal simulation.

- Reduce the change in $P_\mu^\pi \xi$ by the scale factor of ten.

The change in $P_\mu^\pi \xi$ from this procedure is 2.0×10^{-4} . However, the exaggeration of errors in the STRs also worsens the reconstruction resolution, for which there is already an orthogonal systematic uncertainty. In order to avoid double-counting with the resolution systematic uncertainty, a subtraction of 1.1×10^{-4} is necessary, leading to a final systematic uncertainty from the STRs of $\pm 0.9 \times 10^{-4}$.

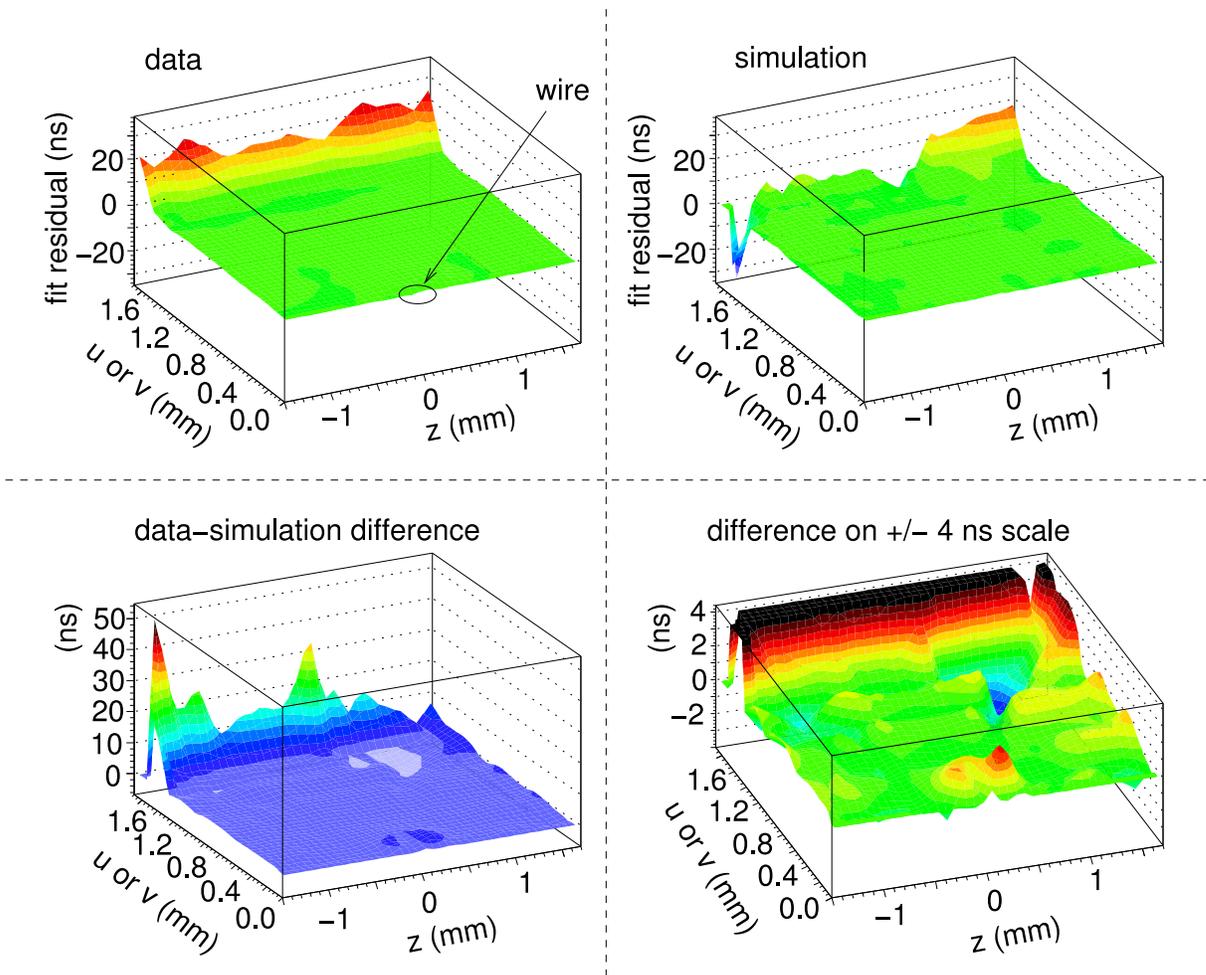


Figure 6.14: An example of the time residuals from the helix fit, after refining the space-time-relationships. Only half a cell is shown, which in practice is reflected about $uv = 0$.

6.3.2 Drift chamber geometric effects

In the simulation the drift chamber foils are always at the same, well known positions. In the real chambers there are three uncertainties in the shape and position of the foils: i) the absolute foil positions are only known to better than $100\ \mu\text{m}$, ii) the foils have an absolute bulge of up to $60\ \mu\text{m}$, and iii) the bulge in the foils varies by up to $35\ \mu\text{m}$ during data acquisition. An inward movement of the foils makes the chambers thinner, reducing the number of hits that are available for the helix fitting algorithm, and this introduces a systematic uncertainty.

An upper limit on the systematic uncertainty uses a simulation with the chamber foils moved inwards by $500\ \mu\text{m}$. This changes $\Delta P_\mu^\pi \xi$ by $(6.7 \pm 5.4) \times 10^{-4}$. Using the largest foil position uncertainty of $100\ \mu\text{m}$, this corresponds to a systematic uncertainty of $(100/500) \times 6.7 \times 10^{-4} = 1.3 \times 10^{-4}$.

6.3.3 Upstream-downstream efficiency

The simulation must reproduce the difference in track reconstruction efficiency (TRE) between the upstream and downstream halves of the detector. This is measured in data and simulation using a special ‘‘upstream stops’’ analysis, where muons are stopped close to the trigger scintillator, and the decay positrons are reconstructed independently in each half of the detector. The TRE is then calculated based on how often a positron is reconstructed in one half of the detector, but not in the other.

The difference in TRE between data and simulation is shown in Fig. 6.15, where a cut has been placed at $(23 < p < 29)\ \text{MeV}/c$ to remove the area of phase space that is contaminated with a beam positron background. There is evidence that the TRE depends on $|\cos\theta|$, but not on p . The systematic uncertainty is evaluated by multiplying the number of upstream counts by i) a constant to measure the p -contribution, and ii) a linear function to measure the $|\cos\theta|$ -contribution. After averaging over the aluminium and silver targets, the p -contribution is 1.3×10^{-4} , and the $|\cos\theta|$ -contribution is 0.2×10^{-4} . These are added in quadrature to give a total efficiency systematic uncertainty of $\pm 1.4 \times 10^{-4}$.

6.3.4 Crosstalk

The origin and removal of electronic crosstalk is described in Section 3.2.2. A limit was set on a systematic uncertainty from crosstalk by re-analysing a data set with the crosstalk

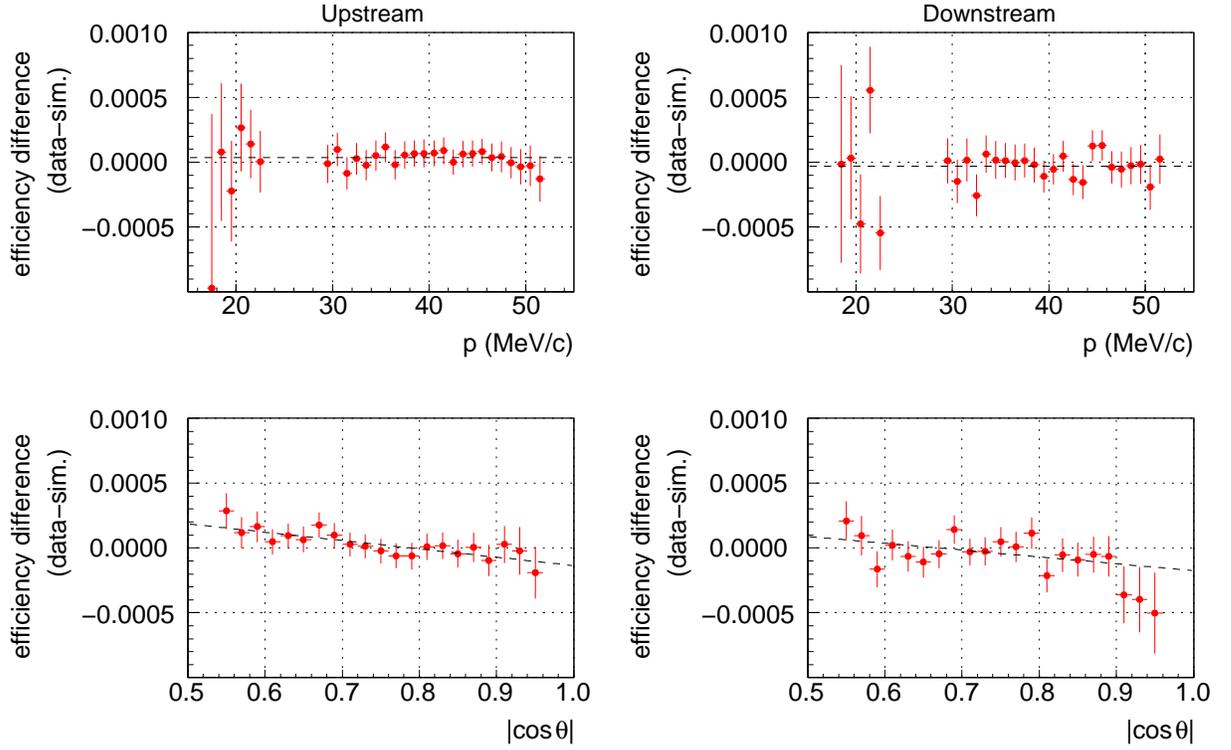


Figure 6.15: Track reconstruction efficiency for the kinematic fiducial. The results for the aluminium target are shown. The upstream efficiency is determined by how often a positron is reconstructed in the downstream half of the detector, but not in the upstream half. The definition is reversed for the downstream efficiency. A cut has been placed at ($23 < p < 29$) MeV/c in order to remove beam positrons. The dashed lines are simple linear fits that are used to characterise the dependence on angle.

removal routine disabled. The change in $\Delta P_\mu^\pi \xi$ was 0.5×10^{-4} , and this is used here as the systematic uncertainty.

6.3.5 Wire time offsets

In the previous $P_\mu^\pi \xi$ analysis the wire time offsets were only calibrated at the beginning and end of the run period. The calibration data were acquired with the magnetic field off, using 120 MeV/c pions and a special downstream trigger. There were significant differences between the two calibrations, and the $P_\mu^\pi \xi$ result changed by 9×10^{-4} depending on which calibration was used[57].

For this measurement a reliable downstream trigger was in place throughout data acquisition. The wire time offsets in each half of the detector were calibrated on a set-by-

set basis using the decay positrons (see Section 3.2.1). Beam positrons that passed through the entire detector were then used to determine the relative timing of the upstream and downstream halves of the detector.

There are two systematic uncertainties from the procedure: a bias from the calibration technique, and the degree to which the relative upstream/downstream timing could be established. The systematic uncertainty due to a bias was evaluated as $\pm 0.04 \times 10^{-4}$, and will therefore not be discussed here (see Ref. [79] for further details). The relative timing of the upstream and downstream halves was determined to better than 0.10 ns. An exaggerated upstream-downstream shift of 10 ns changes $\Delta P_{\mu}^{\pi} \xi$ by $(82 \pm 7) \times 10^{-4}$. After scaling down by $10 \text{ ns}/0.10 \text{ ns} = 100$, this results in a systematic uncertainty of $\pm 0.8 \times 10^{-4}$.

6.3.6 Negligible chamber response uncertainties

An additional three chamber response systematic uncertainties were considered, but were not large enough to require recognition in Table 6.1.

When a decay positron crossed the same drift cell as the muon, there was reduced gain for the positron due to lingering muon ionisation (the chamber had a “dead zone”). The simulation included this dead zone by removing hits within 0.06 cm of the muon hit along the wire, during a recovery time of $3.0 \mu\text{s}$. These parameters were tuned according to the fraction of hits removed by the dead zone in data. In the simulation this corresponded to 0.03% of hits, tuned to an accuracy of about 20% (*i.e.* $\pm 0.006\%$ of the total number of hits) A special simulation was run with an enhanced dead zone: instead of only removing positron hits within 0.06 cm of the muon hit, the entire wire was made insensitive to decay positrons for the $3.0 \mu\text{s}$ recovery time. Inclusion of the enhanced dead zone changes $\Delta P_{\mu}^{\pi} \xi$ by $(25 \pm 2) \times 10^{-4}$, and removes 0.99% of hits. Compared to the nominal case, this special simulation exaggerates the removed hits by $0.99\% - 0.03\% = 0.96\%$, resulting in a scale factor of $0.96\%/0.006\% = 160$. The change in $\Delta P_{\mu}^{\pi} \xi$ is then $(1/160) \times 25 \times 10^{-4}$, which is a systematic uncertainty of just $\pm 0.2 \times 10^{-4}$.

For the previous $P_{\mu}^{\pi} \xi$ analysis, the wires in the simulation were placed at exactly the centre of the cathode foils. In reality, they were determined to be offset from the centre by $150 \mu\text{m}$, and this introduced a systematic uncertainty of $\pm 2.2 \times 10^{-4}$. The current analysis includes the offset in the simulation, eliminating this uncertainty altogether.

In the simulation the gas density inside the drift chambers is constant, but for the data it tracks the atmospheric pressure and the exterior temperature, which results in drift cell

space-time-relationships (STRs) that vary over a data set. The previous analysis used the same STRs for all the data, and a systematic uncertainty of $(0.2 \pm 0.2) \times 10^{-4}$ was established based on the RMS variation in density[83], and $(1.7 \pm 1.0) \times 10^{-4}$ based on the largest variation in density[57]. For the current analysis, each run was analysed with STRs that were corrected for the atmospheric pressure and temperature, meaning the systematic uncertainty is much smaller than $\pm 0.2 \times 10^{-4}$. Therefore the change in STRs due to gas density is a negligible effect, and is no longer evaluated.

6.4 Detector alignment

Two systematic misalignments of the drift chambers have been previously considered[83]: a “shear”, where each detector plane is offset in u or v from the previous one by a constant amount, and a “corkscrew”, where each detector plane is rotated about the z -axis from the previous one by a constant angle. Stringent limitations from the detector design meant these systematic uncertainties had negligible values of $\pm 0.009 \times 10^{-4}$ and $\pm 0.020 \times 10^{-4}$ for the shear and corkscrew, respectively. These were not re-evaluated for the current measurement, and are simply discarded. Random misalignments are not considered, since these smear the tracking residuals and degrade the resolution, which is handled as an orthogonal systematic uncertainty (see Section 6.6).

The length (z) and width (u, v) scales are used to determine the momentum components of the reconstructed helices. The systematic uncertainties from these scales are re-evaluated for the current measurement. The z length scale is known to $50 \mu\text{m}$ out of 100 cm , which is a fractional uncertainty of 5.0×10^{-5} . A special analysis made a fractional change of 1×10^{-3} to the z -component of the momentum, which is a factor of 20 larger than the true uncertainty. The change in $P_\mu^\pi \xi$ is $(0.5 \pm 5.2) \times 10^{-4}$, which leads to a systematic uncertainty of $(1/20) \times 0.5 \times 10^{-4} = \pm 0.03 \times 10^{-4}$.

For the width scale, the wires were positioned in a plane to better than $5 \mu\text{m}$; a wire plane is 32 cm in width, corresponding to a fractional uncertainty of 2×10^{-5} . The sensitivity of $P_\mu^\pi \xi$ was determined in a similar way to the length scale: a special analysis made a fractional change of 1×10^{-3} in both the u - and v -components of the momentum, finding that $P_\mu^\pi \xi$ changes by $(7.8 \pm 5.2) \times 10^{-4}$. This time the systematic uncertainty is scaled down to $(2 \times 10^{-5}) / (1 \times 10^{-3}) \times 7.8 \times 10^{-4} = \pm 0.2 \times 10^{-4}$.

6.5 Positron interactions

6.5.1 Bremsstrahlung and δ -electron rates

Uncertainties from the simulation’s continuous energy loss model are part of the energy calibration uncertainties in Section 6.7.2. In this section the uncertainties from discrete processes are described; the most important of these are δ -electron production (where an electron is knocked out of an atomic orbital) and Bremsstrahlung (“braking radiation”, where one or more photons, radiated during deceleration, can subsequently undergo pair production). The simulation must accurately reproduce these processes since the extra tracks interfere with the reconstruction. Specifically there are three effects: first, the number of chamber hits will differ in data and simulation; second, the δ -electrons and Bremsstrahlung represent invisible energy loss contributions that cause the positron to be reconstructed with a different energy and angle than the original decay kinematics; third, the processes can “break” the track by introducing a large angle change, and this prevents reconstruction. Approximations in the GEANT3 physics and uncertainties in material thicknesses can cause the simulation’s rates to differ from data.

The δ -electron rate is compared in data and simulation by selecting events where the decay positron trajectory is broken in two, with an additional electron track originating from the point where the track is broken. The momentum of the electron track is well correlated with the momentum difference between the two broken track halves, indicating that genuine δ -electrons are being measured (see Fig. 6.16). A special simulation with the production of δ -electrons disabled confirmed that the background of the measurement is small.

The momentum distributions of the reconstructed electrons are compared for nominal data and simulation on the left of Fig. 6.17, where the reconstruction efficiency is seen to decrease below 6 MeV/c. Using the range ($6 < p_\delta < 16$) MeV/c and all the available data sets, the ratio of δ -electrons in data and simulation is 1.007 ± 0.009 (*i.e.* the central value indicates a deficit of δ -electrons in the simulation, but this is not statistically supported).

A special simulation is used with the δ -electron probability increased by a factor of three, changing $P_\mu^\pi \xi$ by $(25 \pm 7) \times 10^{-4}$. For this simulation, the reconstructed δ -electrons are compared for the nominal and special simulation on the right of Fig. 6.17, where the ratio of counts is 2.80 ± 0.04 ; the ratio is not exactly 3.0 due to second order effects. The

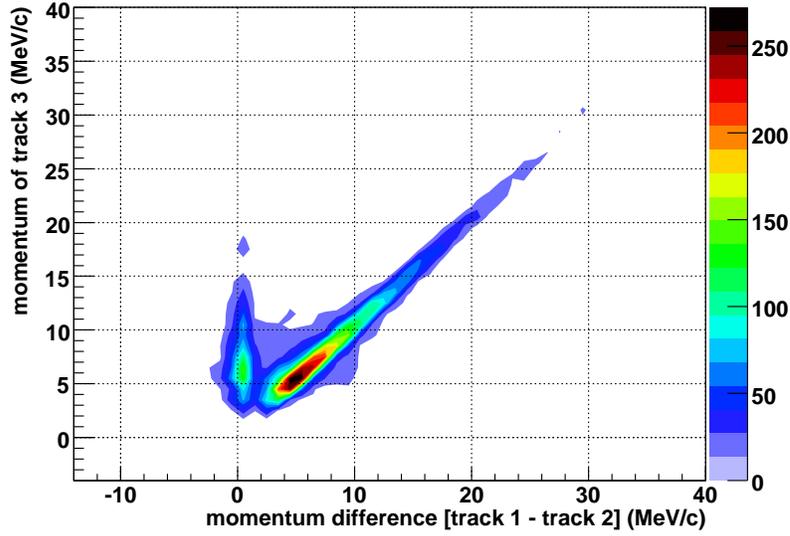


Figure 6.16: Validation of the δ -electron measurement method. Events are selected with a decay positron track that is broken in half, becoming tracks 1 and 2, with an additional negative particle originating from the break point (track 3). The observed correlation indicates that δ -electrons are being measured.

systematic uncertainty is then

$$\frac{1.007 - 1.0}{3.0 - 1.0} \times 25 \times 10^{-4} = \pm 0.1 \times 10^{-4}. \quad (6.16)$$

Note that using 2.80 instead of 3.0 in the denominator does not change the systematic uncertainty by a significant amount.

The Bremsstrahlung rate is compared in data and simulation using a similar approach to the δ -electrons: events are selected with a “broken” decay positron trajectory, and the momentum difference between the two halves of the track is shown on the left side of Fig. 6.18. Another special simulation is used with the Bremsstrahlung probability increased by a factor of three, and this changed $P_{\mu}^{\pi} \xi$ by $(55 \pm 7) \times 10^{-4}$; the effect on the broken track momentum difference is shown on the right side of Fig. 6.18. Using the momentum range of $(15 < p < 35)$ MeV/c, the ratio of Bremsstrahlung events in data and simulation averaged over all sets is 1.024 ± 0.004 . The ratio between the nominal and increased Bremsstrahlung rate simulations is 2.82 ± 0.02 . The systematic uncertainty is then

$$\frac{1.024 - 1.0}{3.0 - 1.0} \times 55 \times 10^{-4} = \pm 0.7 \times 10^{-4}. \quad (6.17)$$

However, exaggerating the Bremsstrahlung also degrades the reconstruction resolution. To avoid double counting the resolution uncertainty, the contribution from the resolution must be subtracted, which results in a final Bremsstrahlung systematic uncertainty of $\pm 0.5 \times 10^{-4}$.

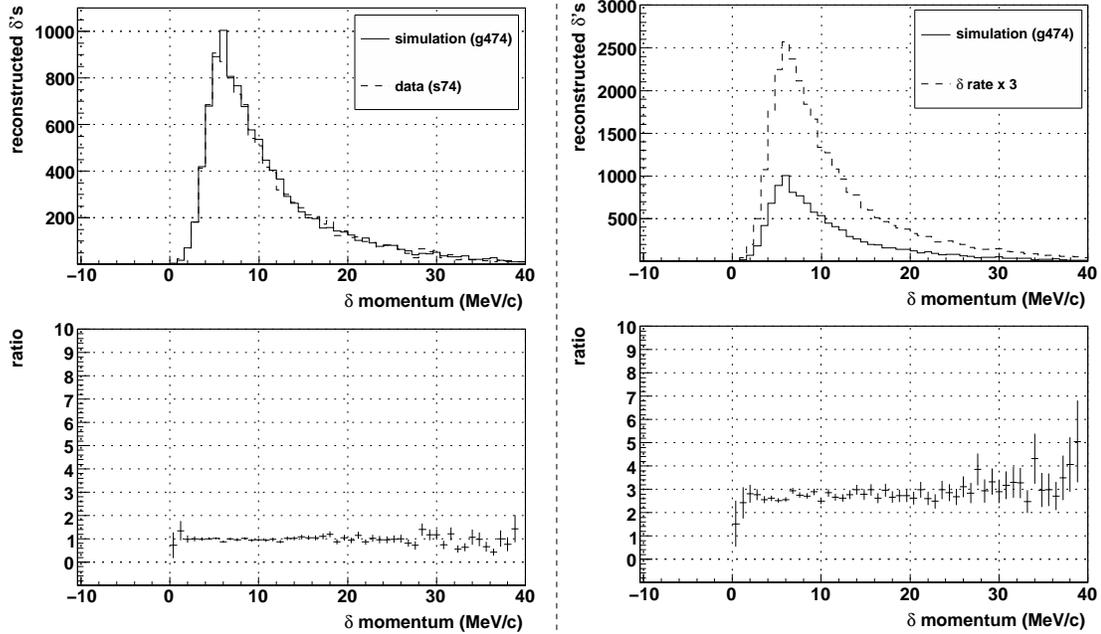


Figure 6.17: The left hand side shows the distribution of reconstructed δ -electrons for data and simulation. The right hand side shows the simulation where the δ production rate was increased by a factor of three. The range ($6 < p_\delta < 16$) MeV/c is used to evaluate the systematic uncertainty.

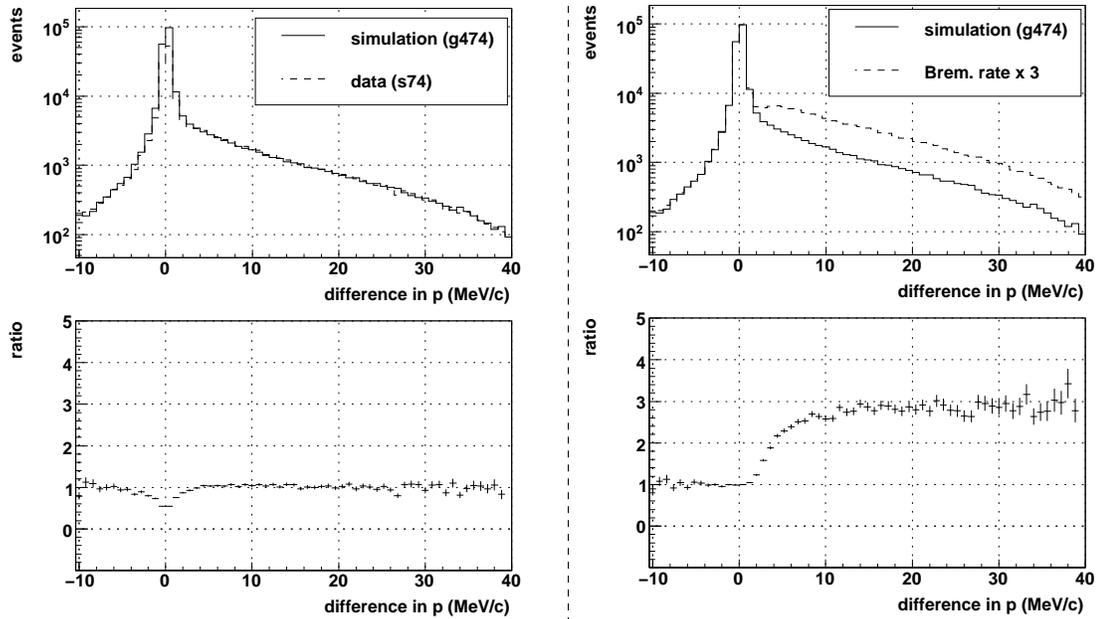


Figure 6.18: The left hand side shows the energy difference for the two halves of a broken track, for data and simulation. The right hand side shows the same distribution for a simulation where the Bremsstrahlung rate was increased by a factor of three. The range ($15 < p < 35$) MeV/c is used to evaluate the systematic uncertainty.

6.5.2 Outside material

In the *upstream* half of the detector the positrons could be backscattered from the beam pipe and the upstream beam package (see Section 2.9). The door of the steel yoke was also a source of backscatters, but most of the yoke was shielded by the upstream beam package. In the *downstream* half of the detector, during nominal operation there was no corresponding downstream beam package, so positrons could only be backscattered from the steel yoke.

These backscatters cause extra hits that interfere with the decay positron reconstruction. They are well reproduced in the simulation since the beam pipe and most of the upstream beam package (excluding, for example, the light guides) are included. However, the steel of the yoke is not included since the additional showering was expected to significantly increase the computation time. This lack of steel yoke and imperfections in placement and/or thickness of materials are expected to introduce a systematic uncertainty.

The degree to which backscatters match in data and simulation can be compared by selecting the time window containing the decay positron, and then finding the difference in average times between the PCs at the far upstream and downstream ends of the detector. A backscatter will cause extra hits in either the upstream or downstream PCs, resulting in an additional peak in this time distribution. The upper and middle plots in Fig. 6.19 show the time distribution for the windows where the decay positron is upstream and downstream respectively. The simulation shows evidence of a surplus in backscattered upstream decay positrons, and a deficit in backscattered downstream decay positrons that is consistent with the steel yoke being disabled.

The effect on $P_\mu^\pi \xi$ can be estimated using two simulations, with and without the *downstream* beam package in place. This exaggerates the number of downstream decay positrons that are backscattered. The difference in $P_\mu^\pi \xi$ between these simulations is $(2.4 \pm 4.0) \times 10^{-4}$, and the two time distributions are shown in the bottom panel of Fig. 6.19. The change in $P_\mu^\pi \xi$ is then scaled down according to the ratio of differences in counts; specifically, the difference between data and simulation is divided by the difference between the simulations with and without the downstream beam package. After averaging over all sets, this results in scale factors of 7 and 14 for the upstream and downstream backscatters respectively. When added in quadrature the systematic uncertainty is then $\sqrt{(2.4/7)^2 + (2.4/14)^2}$, which is $\pm 0.4 \times 10^{-4}$.

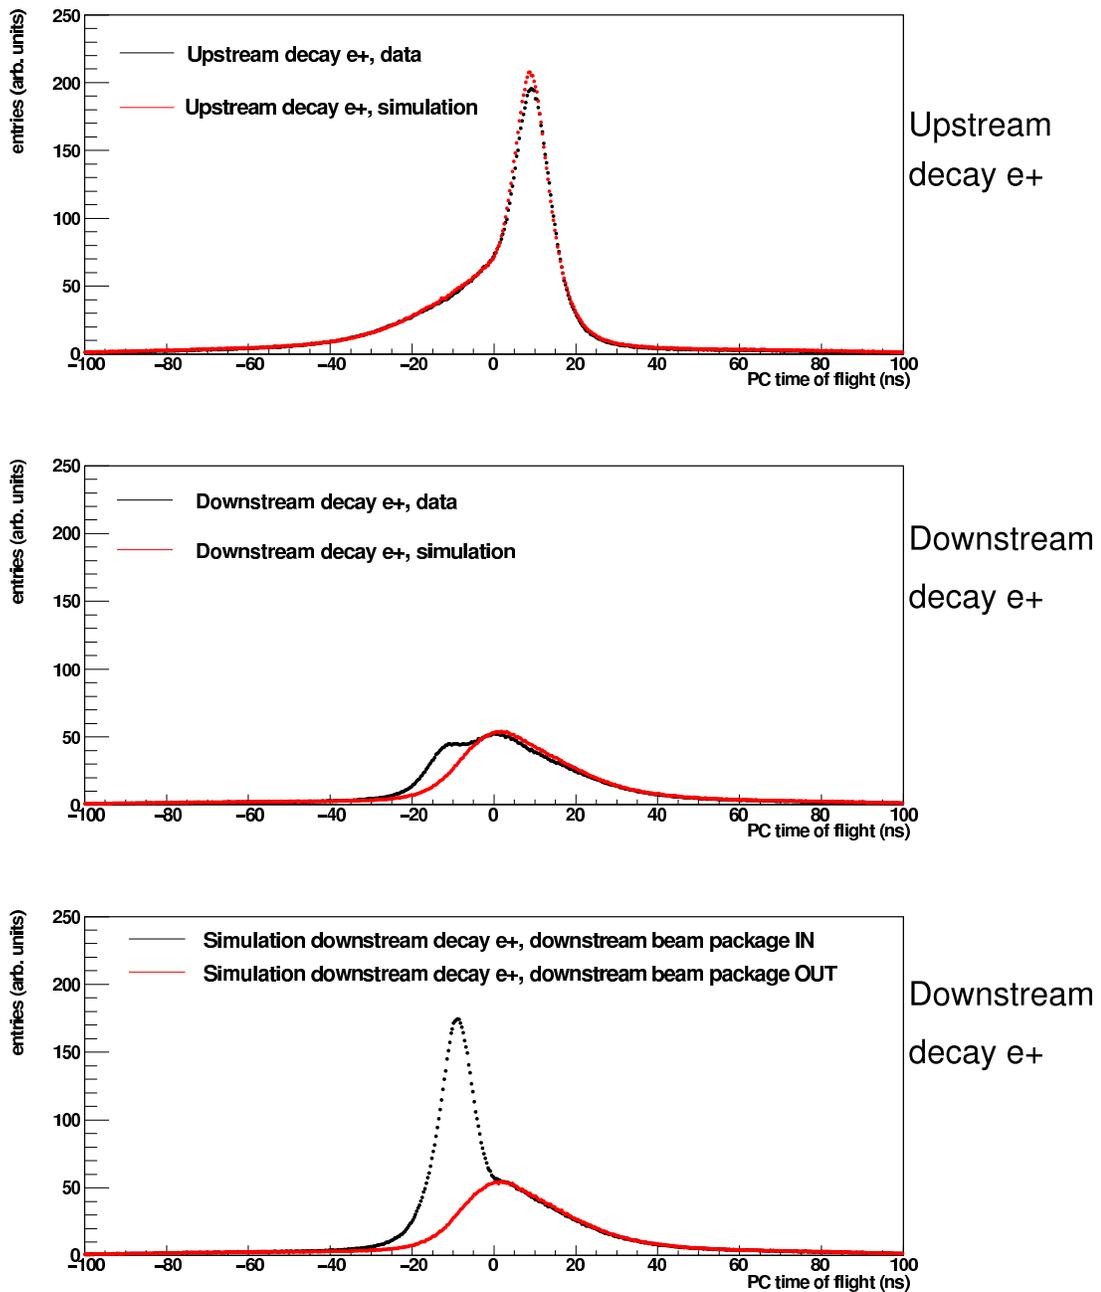


Figure 6.19: The distributions are the difference between the average upstream and downstream PC times. The backscatters from upstream and downstream decay positrons are compared for data and simulation in the upper and middle distributions. The effect of including a downstream beam package in the simulation is shown at the bottom.

6.6 Resolution

The momentum (p) and angle ($\cos \theta$) reconstruction resolutions might differ in data and simulation. This discrepancy is relatively unimportant over the kinematic fiducial since the decay spectrum is smooth. However, the analysis uses the sharp kinematic endpoint to energy calibrate the spectrum, and a significant difference in resolution at this endpoint will introduce a systematic uncertainty.

The resolution can be compared in data and simulation indirectly using special “up-stream stops” data, where muons are stopped at the entrance of the detector, and the decay positron is reconstructed independently in each half of the spectrometer. The reconstructed p and θ will differ in each half due to energy loss and multiple scattering through the target module. The distribution of the momentum/angle differences is dominated by a Gaussian resolution function, so that the difference in σ widths between data and simulation is a measure of how well the simulation reproduces the detector’s resolution. Even though this difference is also sensitive to discrepancies in target thickness and the simulation’s positron interactions physics, we conservatively assign the full difference as a resolution problem.

The Gaussian widths are shown for a limited p range in Fig. 6.20. The absolute width is seen to depend on $1/\sin \theta$, but the *difference* between data and simulation is well approximated by a constant for both the momentum and angle dependence. A previous analysis found this difference in width had a non-trivial dependence on p and $\cos \theta$; the current analysis is improved due to the use of superior drift cell space-time-relationships (see Section 3.2.7).

For each point in $(p, \cos \theta)$, the difference in p and θ width between data and simulation is constructed according to

$$\Delta\sigma = \begin{cases} \sqrt{\sigma_{\text{data}}^2 - \sigma_{\text{sim}}^2} & , \sigma_{\text{data}} > \sigma_{\text{sim}} \\ -\sqrt{\sigma_{\text{sim}}^2 - \sigma_{\text{data}}^2} & , \sigma_{\text{sim}} > \sigma_{\text{data}} \end{cases} \quad (6.18)$$

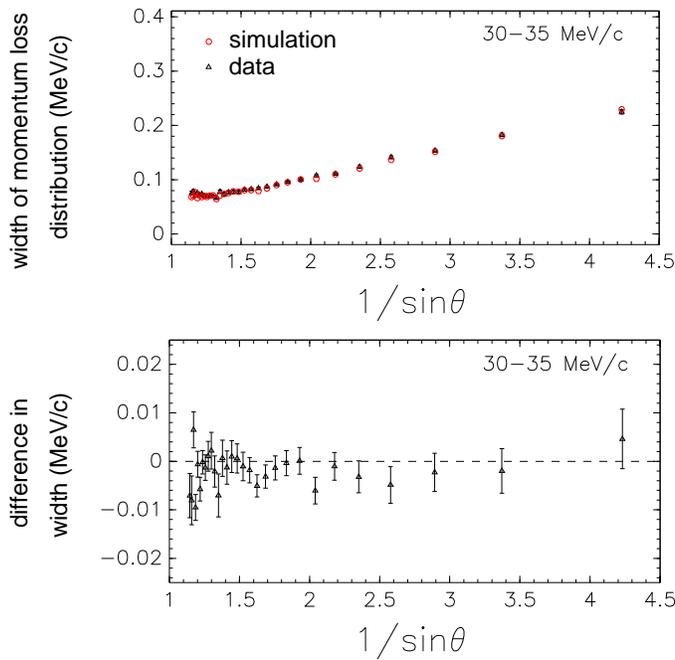
The weighted average of this quantity over all $(p, \cos \theta)$ is shown in Table 6.7. The systematic uncertainty is measured by exaggerating the largest differences from the table by a factor of five; specifically an additional smearing of 58 keV/c in momentum and 6 mrad in angle is added to the data. The momentum and angle smearing are carried out separately. The momentum smearing changes $P_{\mu}^{\pi} \xi$ by $(7.6 \pm 3.4) \times 10^{-4}$, resulting in a systematic uncertainty of $(7.6/5) \times 10^{-4} = \pm 1.5 \times 10^{-4}$. The angle smearing changes $P_{\mu}^{\pi} \xi$

by 0.2×10^{-4} , resulting in a negligible systematic uncertainty after scaling down by five.

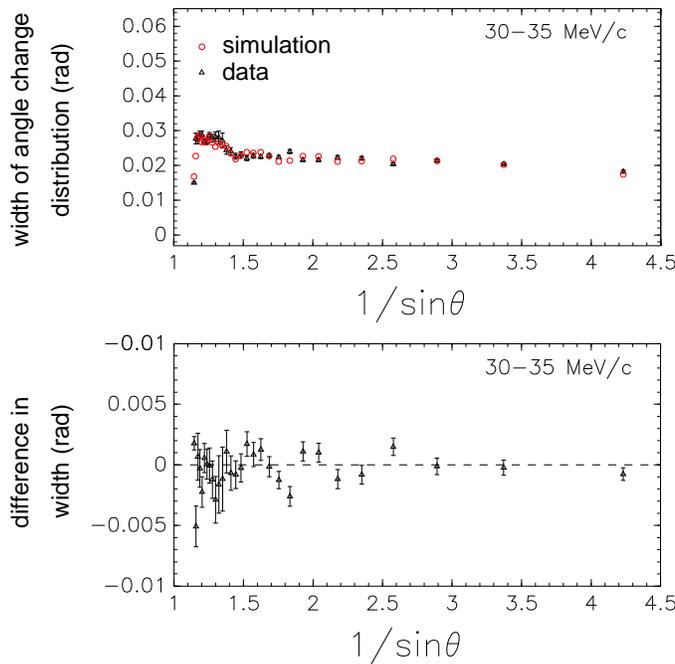
The width of the decay spectrum endpoint provides an additional independent measurement of the resolution at $\sin \theta = 1$ (extrapolated). The data-simulation difference in this width is (0.0 ± 0.3) keV/c for silver and (0.1 ± 0.3) keV/c for aluminium. Since these are smaller than the differences in Table 6.7, the systematic uncertainty of $\pm 1.5 \times 10^{-4}$ is already sufficient.

Table 6.7: Difference in resolution between data and simulation, as defined in the text.

Target	Difference in σ , defined by Eq. (6.18)	
	Momentum (keV/c)	Angle (mrad)
Aluminium	-6.3	-0.07
Silver	-11.5	1.1



(a) Momentum resolution.



(b) Angle resolution.

Figure 6.20: Indirect comparison of the resolution in data and simulation, derived from a special “upstream stops” analysis: muons stop at the entrance of the detector, and are reconstructed independently in each half of the spectrometer; the width of the energy loss and angle change distributions is sensitive to the resolution.

6.7 Momentum calibration

6.7.1 Magnetic field shape

The simulation was self-consistent since it used the same `Opera` field map for generating and reconstructing the positrons. In contrast, the data were analysed with the `Opera` field map, but this had known discrepancies with the mapper measurements within the tracking region, which can introduce a systematic uncertainty.

The difference in field shape between `Opera` and the mapper measurements is well approximated by

$$\delta B_z = C_2 z^2 + C_3 z^3 + C_r r, \quad (6.19)$$

where the optimum C_2 , C_3 and C_r values are recorded in Table 6.8. The three previous TWIST analyses have used the same coefficients for the nominal 2.0 T field. The coefficients were re-evaluated for the current measurement by weighting the mapper measurements more carefully[104].

Table 6.8: Coefficients of Eq. (6.19). These relate the `Opera` magnetic field to the results from the mapper measurements.

Parameter	Nominal 2.0 T field		1.96 T	2.04 T
	Previously ^a	This analysis		
C_2 ($\times 10^{-8}$ T/cm ²)	-6	-1.7 ± 0.4	$+11.4 \pm 0.3$	-19.7 ± 0.5
C_3 ($\times 10^{-10}$ T/cm ³)	-4	-7.8 ± 0.9	$+2.0 \pm 0.5$	-2.2 ± 0.7
C_r ($\times 10^{-6}$ T/cm)	-12.5	-8.3 ± 0.3	-1.1 ± 0.6	-2.8 ± 0.9

^a The same parameters were used in Refs. [57, 83, 84, 86].

A new 2.0 T field was produced with the coefficients in Eq. (6.19) exaggerated by a factor of 20. Maxwell's $\nabla \cdot \delta \vec{B} = 0$ equation was satisfied by modifying the radial field components according to

$$\delta B_r = - \left(C_2 r z + \frac{3}{2} C_3 z^2 r \right). \quad (6.20)$$

A data set was then re-analysed with the exaggerated field, and the change in $P_\mu^\pi \xi$ is $(5.1 \pm 7.8) \times 10^{-4}$. After scaling down by a factor of 20, the change in $P_\mu^\pi \xi$ is 0.3×10^{-4} .

This is too small to justify re-analysing all the data with a corrected map, and is instead taken as a systematic uncertainty of $\pm 0.3 \times 10^{-4}$. Previous evaluations of this uncertainty used a smaller scale factor of ten, and did not apply Eq. (6.20), but still found an effect below 1×10^{-4} [57, 83].

(The data sets with a central magnetic field strength of 1.96 T and 2.04 T are not included for $P_\mu^\pi \xi$. Therefore the systematic due to magnetic field shape for these sets is not discussed here. See Refs. [79, 97] for further information.)

6.7.2 Use of the kinematic endpoint

The motivation for an energy calibration and its implementation were described in Section 3.5. In summary, the reconstructed momenta of the data and simulation disagree at the kinematic endpoints by about 10 keV/c. This difference is shown in Fig. 6.21 for the kinematic fiducial. The upstream ($1/\cos\theta < 1$) and downstream ($1/\cos\theta > 1$) points are fit separately with a straight line, yielding two slopes ($a_{\text{up}}, a_{\text{down}}$) and two intercepts ($b_{\text{up}}, b_{\text{down}}$), for a total of four “endpoint parameters”. The calibration must be applied to the entire data spectrum by shifting or scaling each reconstructed momentum. (The same central value of $P_\mu^\pi \xi$ is obtained if both spectra are corrected to the true kinematic endpoint of $W_{e\mu} = 52.83$ MeV, instead of correcting the data relative to the simulation.) Note that the energy calibration procedure is applied to every systematic uncertainty test, which improves the robustness of the $P_\mu^\pi \xi$ measurement.

There are three uncertainties from the energy calibration: i) a statistical part since only a limited region of the spectrum is used to establish the required correction, ii) a systematic part since either a shift or scale must be used to propagate the correction to the rest of the spectrum, and iii) a systematic part due to the model imposed on the endpoint behaviour with angle. The statistical part is propagated separately for each set, and is applied later in Section 7.1. For the systematic part due to a shift or scale, the two extremes considered were a *shift* according to

$$p_{\text{corrected}} = p_{\text{reconstructed}} - \left(b - \frac{a}{|\cos\theta|} \right), \quad (6.21)$$

and a *scale* according to

$$p_{\text{corrected}} = \frac{p_{\text{reconstructed}}}{1 + \frac{1}{W_{e\mu}} \left(b - \frac{a}{|\cos\theta|} \right)}, \quad (6.22)$$

where $W_{e\mu}$ is the maximum kinematic positron momentum (52.83 MeV/c). The central value of $P_{\mu}^{\pi} \xi$ is placed half way between the shift and scale extremes, and a systematic uncertainty of half the difference is assigned to cover both possibilities ($\pm 0.5 \times 10^{-4}$).

An additional systematic uncertainty is assigned due to the linear model imposed on the endpoint behaviour with $\cos \theta$ (see Fig. 6.21). The reduced χ^2 for the linear fits average to 1.27. A decrease in the χ^2 to 1.00 corresponds to larger uncertainties for the endpoint parameters, and this ultimately results in a systematic uncertainty of $\pm 1.4 \times 10^{-4}$ for $P_{\mu}^{\pi} \xi$.

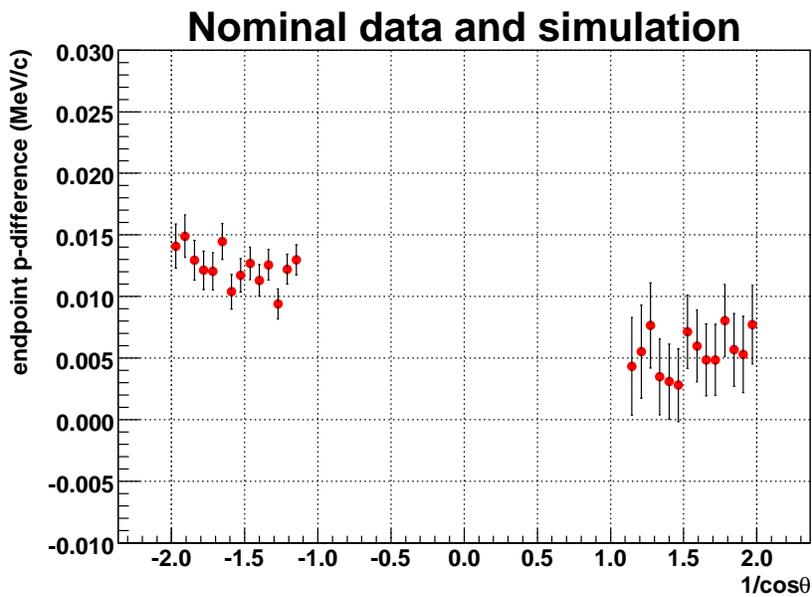


Figure 6.21: Difference between reconstructed momentum in data and simulation at the endpoint of the muon decay spectrum.

6.8 Beam stability

6.8.1 Muon beam intensity

The simulation's muon and beam positron rates were tuned to match the data; a disagreement can change the number of hits in the wire chambers, which affects the performance of the classification. The simulation's muon rate was tuned to approximately match the trigger rate in data. Similarly, the simulation's beam positron rate was tuned to match the probability of a beam positron entering an event.

Increasing the muon rate in the simulation by a factor of 10 changes $\Delta P_\mu^\pi \xi$ by $(10 \pm 7) \times 10^{-4}$. This must be scaled down by repeating the procedure from the previous analysis[83]:

1. For each data set and its accompanying simulation, calculate the following ratio of event types:

$$R_\mu = \frac{(\text{more than one } \mu^+)}{(\text{more than one } \mu^+) + (\text{one } \mu^+, \text{ one decay } e^+)}. \quad (6.23)$$

This is a measure of the probability of an event having more than one muon.

2. Calculate the relative ratio of R_μ for data and simulation,

$$\frac{R_\mu(\text{sim.}) - R_\mu(\text{data})}{R_\mu(\text{data})}. \quad (6.24)$$

3. Multiply the relative ratio by the average trigger rate from the data, to estimate the error (in s^{-1}) made in the simulation.
4. Divide the simulation's error by the exaggeration (24579 s^{-1}).

The calculations from each of these steps appear in Table 6.9. An upper limit on the effect of data-simulation rate discrepancies uses the smallest scale factor (28.7, from set 92), resulting in a systematic uncertainty of $(1/28.7) \times 10 \times 10^{-4} = 0.3 \times 10^{-4}$.

The systematic uncertainty due to a beam positron rate mismatch in data and simulation is just $\pm 0.2 \times 10^{-4}$, and is therefore not discussed further.

Table 6.9: Scale factors for the systematic uncertainty due to beam intensity. R_μ is a measure of the probability of an event with more than one muon. Sets 68-76 were accumulated in 2006 using a silver target. Sets 83-93 were accumulated in 2007 using an aluminium target, with a muon rate that was intentionally higher. Simulation is abbreviated as Sim.

Set	R_μ Data	R_μ Sim.	$\frac{R_\mu(\text{sim.}) - R_\mu(\text{data})}{R_\mu(\text{data})}$	Avg. data trigger (s^{-1})	Sim. error (s^{-1})	Scale factor
68	0.00553	0.00495	-0.10615	2066.0	-219.3	112.1
70	0.00630	0.00541	-0.14129	2324.8	-328.5	74.8
71	0.00660	0.00619	-0.06305	2582.9	-162.8	150.9
72	0.00718	0.00552	-0.23151	2674.2	-619.1	39.7
74	0.00672	0.00617	-0.08089	2592.9	-209.7	117.2
75	0.00739	0.00641	-0.13216	2686.6	-355.1	69.2
76	0.00903	0.00640	-0.29114	2740.9	-798.0	30.8
83	0.01220	0.01029	-0.15728	4221.8	-664.0	37.0
84	0.01291	0.01103	-0.14616	4452.3	-650.7	37.8
86	0.01353	0.01205	-0.10914	4972.3	-542.7	45.3
87	0.01188	0.00977	-0.17768	4024.0	-715.0	34.4
91	0.01218	0.01004	-0.17531	4202.5	-736.7	33.4
92	0.01251	0.01002	-0.19897	4305.3	-856.6	28.7
93	0.01077	0.01003	-0.06851	3809.5	-261.0	94.2

6.9 External

6.9.1 Radiative corrections

The simulation uses the following radiative corrections: full first order, $O(\alpha^2 L^2)$ and $O(\alpha^2 L^1)$ from the second order, and $O(\alpha^3 L^3)$ from the third order. The term $O(\alpha^2 L^0)$ was not used, despite becoming available in 2007 [16]; its effect on the decay parameters will now be shown as negligible.

Over the TWIST kinematic fiducial, the $O(\alpha^2 L^0)$ term has a similar shape to the $O(\alpha^2 L^1)$ term, and the ratio between the terms never exceeds 0.2; this is demonstrated in the paper where the $O(\alpha^2 L^0)$ term is calculated[16]. The $P_\mu^\pi \xi$ sensitivity to excluding the $O(\alpha^2 L^0)$ term was estimated by adding a pure $O(\alpha^2 L^1)$ spectrum to the nominal spectrum. This changes $P_\mu^\pi \xi$ by $(2.9 \pm 0.1) \times 10^{-4}$. The ratio of counts in the nominal and combined spectrum is 1.11. Therefore the effect on $P_\mu^\pi \xi$ of only adding 0.2 of the pure $O(\alpha^2 L^1)$ spectrum is

$$(0.2/1.11) \times 2.9 \times 10^{-4} = 0.5 \times 10^{-4}, \quad (6.25)$$

and $\pm 0.5 \times 10^{-4}$ is used as the systematic uncertainty.

6.9.2 Correlation with η

The muon decay spectrum does not allow a precise measurement of the parameter η . Therefore η was fixed to its world average value⁴³ of $(-36 \pm 69) \times 10^{-4}$ [7], and the correlation between η and $P_\mu^\pi \xi$ is assessed here as a systematic uncertainty. The correlation was found to be $d\xi/d\eta = 0.01528$, so that $\Delta\eta = \pm 69 \times 10^{-4}$ corresponds to a $P_\mu^\pi \xi$ systematic uncertainty of $\pm 1.1 \times 10^{-4}$.

⁴³The evaluation of η in Ref. [7] uses a global analysis that includes the TWIST experiment's result for ρ and δ from the 2002 datasets. The evaluation does *not* use the TWIST experiment's more recent ρ , δ and $P_\mu^\pi \xi$ measurements from the 2004 datasets.