

Chapter 6

Systematic Uncertainties

6.1 Introduction

There are 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 latter are evaluated simultaneously for ρ , δ and $P_\mu^\pi \xi$, by exaggerating an effect in the simulation or the analysis software; 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 a statistically meaningful MP change, 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 uncertainties in the MP differences are scaled down until the reduced χ^2 from the fitting procedure is equal to one.

6.2 Overview

The $P_\mu^\pi \xi$ uncertainties for the nominal sets are summarised in Table 6.1, which shows that the measurement is limited by the accuracy of the muon beam and the magnetic field map. The entries in this table will be described fully in the current chapter. The table indicates three statistical uncertainties; these could be reduced by accumulating more data and/or simulation under the same running conditions. The most recent TWIST analysis (MacDonald '08 in the table) did not re-evaluate the polarisation uncertainties, since it was a measurement of ρ and δ . A selection of the systematic uncertainties are set-dependent, and in these cases Table 6.1 contains the uncertainty for the nominal sets only.

Three corrections to $\Delta P_\mu^\pi \xi$ (the difference in $P_\mu^\pi \xi$ between the data and a hidden simulation value) will be described. First, the simulation does not include depolarisation of the muons while leaving the 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 $\langle p \rangle = 28.75 \text{ MeV}/c$ ($\langle p \rangle = 28.85 \text{ MeV}/c$). Second, the time dependent depolarisation in the simulation used a preliminary relaxation

rate, resulting in a correction of $+6.3 \times 10^{-4}$ for silver (2006) and $+0.9 \times 10^{-4}$ for aluminium (2007). Third, since we are not able to choose between two applications of energy calibration, a correction of $+0.7 \times 10^{-4}$ is made to place the central value of $P_{\mu}^{\pi} \xi$ half way between the two choices.

Table 6.1: Summary of $P_\mu^\pi \xi$ uncertainties. The statistical uncertainties are marked (stat.); otherwise the uncertainties are systematic. For this analysis, (0) indicates the uncertainty is no longer evaluated.

Category	Thesis section	New eval.?	$\Delta P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)		
			This analysis	MacDonald '08 [10, 18]	Jamieson '06 [21, 57]
Extraction of $\Delta P_\mu^\pi \xi$ (stat.)	7.1	✓	2.4*	3.7	6
Polarisation					
μ^+ beam/ fringe field	6.3.1	✓	$^{+15.0}_{-4.4}$	34.0	34.0
Production target	6.3.2	✓	0.3	2.1	2.1
Stopping material					
λ (stat.)	6.3.3	✓	3.0	Not eval.	Not eval.
$P_\mu(t)$ model	6.3.3	✗	(0)	12	12
Background muons	6.3.4	✓	1.0	2	1.8
Beam intensity	6.3.5	✓	0.8	0.2	1.8
Chamber response					
DC STR	6.4.1	✓	0.0	6.0	Not eval.
Wire time offsets	6.4.2	✓	0.3	0.4	8.9
US-DS efficiency	6.4.3	✓	1.3	1.1	1.9
Dead zone	6.4.4	✓	0.2	0	0.1
Foil bulge	6.4.5	✓	0.5	0.7	2.2
Cell asymmetry	6.4.6	✗	(0)	0	2.2
Density variations	6.4.7	✗	(0)	0.2	0.2
Detector alignment					
DC alignment	6.5	✗	0.02	0.02	2.2
z length scale	6.5	✗	0.3	0.7	2.2
u/v width scale	6.5	✗	0.2	0.2	Not eval.
B-field to axis	6.5.1	✗	0.3	Not eval.	0.3
Positron interactions					
δ -electron rate	6.6.1	✓	0.1	1.4	2.9
Bremsstrahlung rate	6.6.1	✓	0.7	0.03	
Outside material	6.6.2	✓	0.5	0.6	0.2
Multiple scattering	6.6.3	✗	(0)	0	0.8
Energy loss	6.6.4	✗	0.01	0.01	0.1
Resolution	6.7	✓	*	0.7	Not eval.
Momentum calibration					
Tracking B-field	6.8.1	✓	0.3	1.1	0.9
Kinematic endpoint					
Parameters (stat.)	6.8.2	✓	1.4	0.5	1.6
Propagation	6.8.2	✓	0.7	0.09	
External					
Radiative corrections	6.9.1	✓	0.01	0.5	1.0
η correlation	6.9.2	✓	1.1	1.1	Not eval.
Total systematic			$^{+15.2}_{-5.1}$	-	38
Total statistical			4.1	-	6

6.3 Polarisation

6.3.1 Muon beam and fringe field

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 systematic uncertainties from the muon beam and fringe field are summarised in Table 6.2, in the order they appear in this chapter. Note that the muon beam uncertainties are separated into two approximately orthogonal parts: contributions from the average position/angle of the beam, and those from the width of the beam’s angular distributions. Before describing the uncertainties, explanations of polarisation measures, the fringe field map and the beam tunes will be given.

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 average position/angle	
Initial position/angle	2.3
Magnetic field map	
Translational alignment	1.3
Rotational alignment	0.9
Transverse field components	+14.1 -0
μ^+ beam angular distributions	
Simulation of multiple scattering	3.6 -2.6
Noise from TEC electronics	1.7
Aging of TEC sense planes	+1.6 -1.4
Quadratic sum	+15.0 -4.4

Measures of polarisation

Here $P_\mu(0)$ is the average z -component of the muon’s spin in the simulation, at the metal target, *before* any time dependent depolarisation has taken place. This is used to evaluate changes in $\Delta P_\mu^\pi \xi$ (the difference in $P_\mu^\pi \xi$ between data and a hidden simulation value) from uncertainties in the magnetic field and muon beam measurements, and should not be taken

as the absolute polarisation of the muons. Since the simulation knows the spin of each muon exactly, $P_\mu(0)$ can be determined with high precision using a relatively small number of muons. The $P_\mu(0)$ values in this section use simulations with 0.1×10^6 muons, which results in a statistical uncertainty of 0.1×10^{-4} for nominal sets, and 0.2×10^{-4} for sets with a steered beam profile.

For the data, the difference between two decay positron spectra can measure $\Delta P_\mu^D \xi$, where P_μ^D is the average polarisation at the time of decay. The quantity $\Delta P_\mu^D \xi$ was validated to be directly comparable to the difference between two $P_\mu(0)$ values from the simulation.

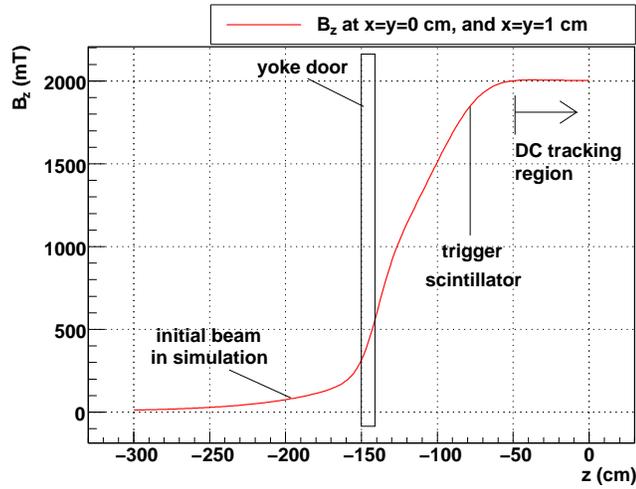
Magnetic field map

The term “fringe field” will refer to the magnetic field from the end of the M13 beam line up to the first drift chamber ($-200 \text{ cm} < z < -50 \text{ cm}$). The three components of the field map used for the analysis are shown in Fig. 6.1. These were generated using the OPERA software package[83], which will be described in more detail later. The longitudinal components (B_z) increase steadily until the drift chamber (DC) tracking region. The on-axis transverse components (B_x, B_y) are less than 1.5 mT, but become more significant away from the axis. These components are radially symmetric, and are maximised just inside the door of the yoke. The transverse components are closely related to the depolarisation, which is shown for the simulation of a nominal profile in Fig. 6.1(c). For example, the onset of rapid depolarisation coincides with the maximisation of the transverse components, and the field’s inflection at $z = -100 \text{ cm}$ is accompanied by an inflection in the depolarisation.

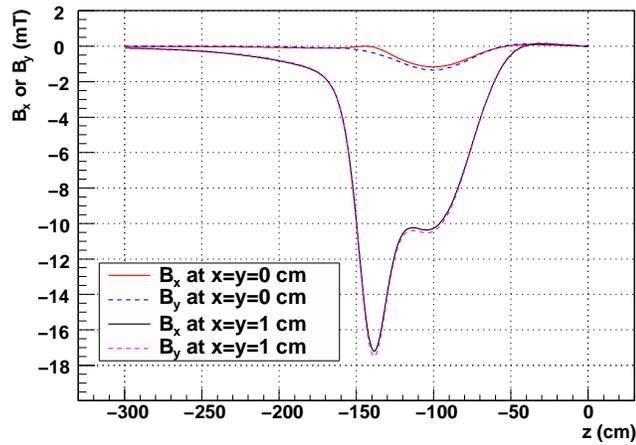
Clearly the quality of the fringe field downstream of the door is important since it controls the rate of depolarisation. In addition, the field upstream of the door must be known since it affects which part of the fringe field the beam is transported through.

Muon beam tunes

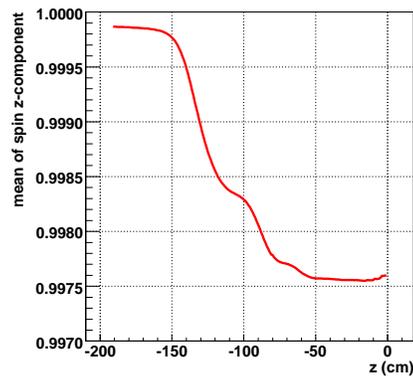
The nominal beam tune was described in Section 5.2; initially the position of the beam at the TECs was steered to be close to $x = y = 0$, with the angles θ_x and θ_y minimised, but ultimately the tune was chosen to place the muon beam spots within the detector along a straight line, corresponding to minimised transverse momentum. The “envelope” of a simulated nominal muon beam, defined as the mean plus or minus one root-mean-square, is shown in Figs. 6.2(a) and 6.2(b). The bulk 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 $z = -100 \text{ cm}$. After this focus the envelope develops oscillations in its mean position and size.



(a) B_z component.



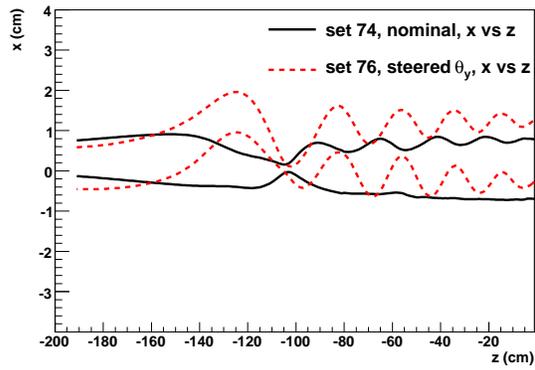
(b) B_x and B_y components.



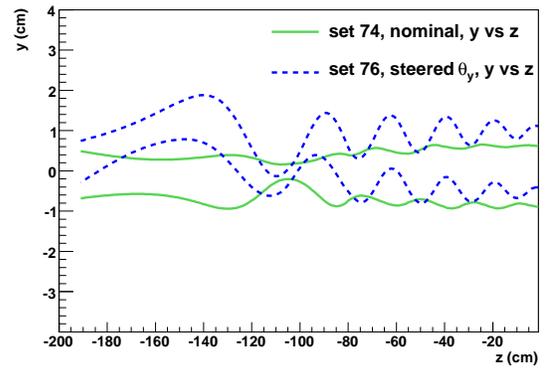
(c) Average of spin z -component, for a nominal beam profile.

Figure 6.1: Fringe field components from the OPERA finite element analysis. The B_x and B_y components indicate radial symmetry. The average spin is also shown since its behaviour is closely related to the B_x and B_y components.

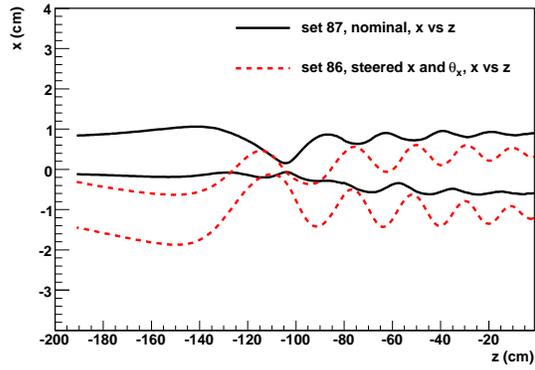
Three additional beam tunes are shown in Fig. 6.2. These will be used to assess the systematic uncertainties for the fringe field. They are set 76, where the muon beam was steered to have $\langle\theta_y\rangle \approx 28$ mrad at the TECs, set 86, where the beam was placed off axis and pointed away from the aaxis ($\langle x\rangle \approx -1.0$ 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 set 76 and set 86 sample the fringe field at a different location to the nominal beam, with a non-zero average angle. They undergo a dramatic 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 data's depolarisation will provide a stringent test of the fringe field model; such a validation was not available for the previous $P_\mu^\pi \xi$ analysis.



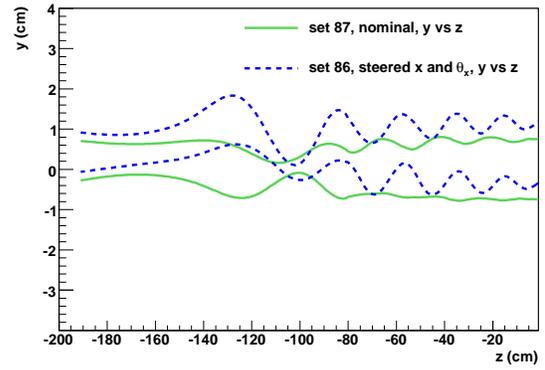
(a)



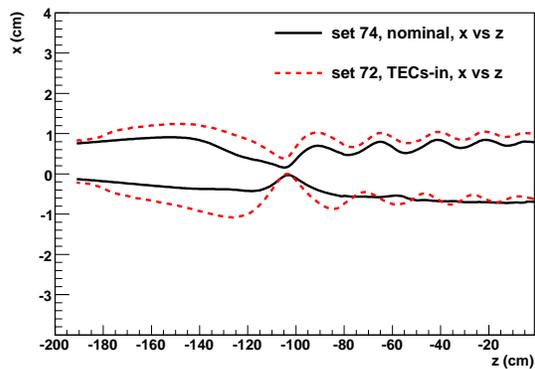
(b)



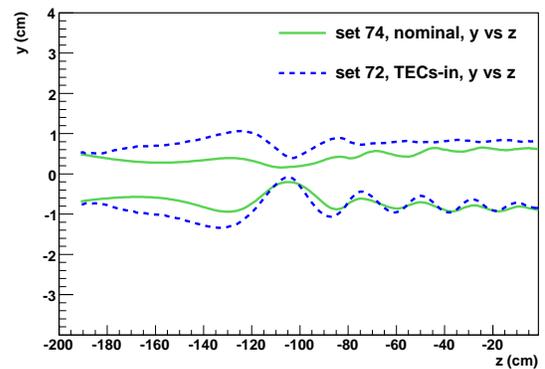
(c)



(d)



(e)



(f)

Figure 6.2: Simulation of muon beam envelopes, defined as the mean position plus or minus one root-mean-square.

Muon beam: uncertainty in initial position and angle

A muon beam measurement was made with the TECs at the beginning and end of most data sets. The 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, these 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. The origin of the changes could be muon beam instability, a limitation in the reproducibility of the TECs, or an instability in the TEC drift cell response. Each of these will now be discussed.

Table 6.3: Muon beam differences for the beginning and end of set TEC measurements.

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 (°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 $\Delta T = T_{\text{end}} - T_{\text{start}}$. $\Delta T > 0$ indicates a temperature rise between measurements.

^b These sets only had one TEC measurement.

Muon beam instabilities were caused by a change in the proton beam upstream of the production target, or an instability in the M13 beam line elements. A special test displaced the proton beam at the production target by ± 0.1 cm vertically, which is about five times larger than the beam could have moved during normal operation³⁷. The largest observed

³⁷The proton beam was surrounded by four counters (top, bottom, left, right). In order to steer the beam vertically by 0.1 cm and avoid destroying the top or bottom counter, 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 $\approx 20/100 \times 0.1 \text{ cm} = 0.02 \text{ cm}$.

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.

The settings of the M13 beam line elements (*e.g.* quadrupoles, dipoles, slits, jaws, asymmetric currents for quadrupole steering) were all monitored with a slow control system, and runs with an instability were eliminated (see Section 5.4). Sets 72 had the TECs in place throughout, and found that the average muon beam position and angle were stable to < 0.02 cm and < 1 mrad respectively over the period of a week (see Fig. 5.2). The muon beam measurement from the wire chambers was used to monitor the stability for the nominal sets, which did not have the TECs in place. The sensitivity of the internal beam to M13 instabilities was determined by adjusting the currents in each quadrupole and dipole by $\pm 5\%$. An examination of the internal beam found that the observed variations corresponded to negligible changes at the TECs, and did not correlate with the differences from Table 6.3.

The space-time-relationship in the TEC drift cells depended on temperature. A change of $\pm 3^\circ\text{C}$ altered the average reconstructed positions by between 0.028 cm and 0.050 cm, depending on the average position of the beam within the module. (The x -positions increased with temperature and the y -positions decreased; see Ref. [76] for further detail.) The changes in angle were all < 0.05 mrad, except for set 76, which still only changed by 0.4 mrad. The temperature differences in Table 6.3 are not correlated with the change in average beam parameters, ruling out temperature as the dominant cause of the beginning/end of set differences.

The insertion/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 the TECs could be used. 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 average initial position and angle.

Using the OPERA field map, the simulation determined the $P_\mu(0)$ sensitivity to position and angle changes of ± 0.2 cm and ± 3 mrad respectively. These are the limits of the observed changes 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). The dependence of $P_\mu(0)$ on changes in initial position and angle is well approximated by a quadratic form. This is demonstrated in Fig. 6.3 for a nominal beam profile. The systematic uncertainty is

assessed by conservatively taking half the range of $P_\mu(0)$ values, resulting in the numbers in Table 6.4. An uncertainty of 2.3×10^{-4} is assigned to all the nominal sets, and 12.8×10^{-4} to the sets with the steered beams.

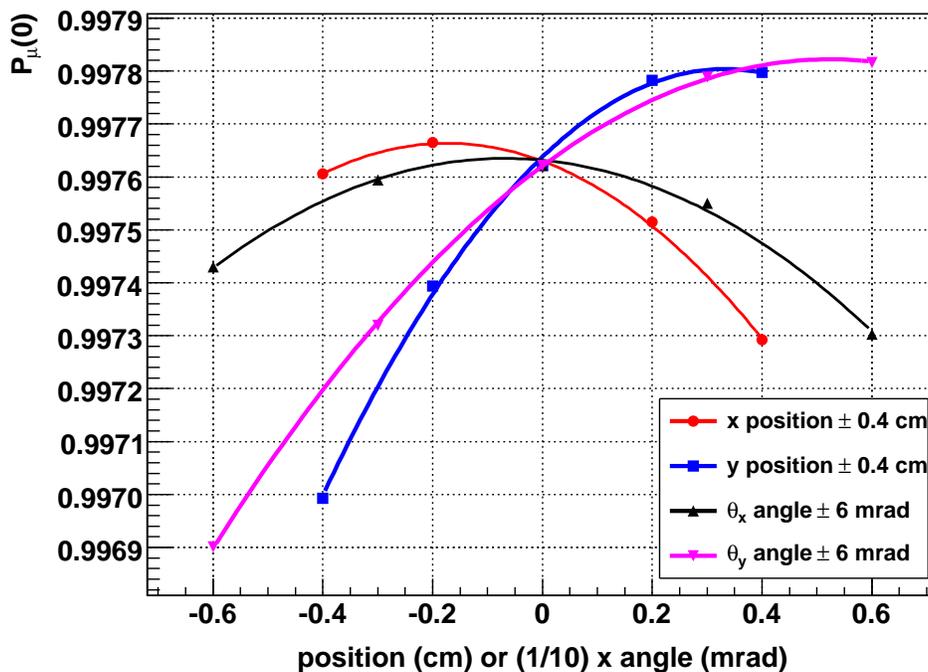


Figure 6.3: The polarisation of the muons after the fringe field depends quadratically on changes in the initial position and angle of the beam. The results of simulating a nominal beam profile are shown. The initial beam position has been displaced by ± 0.2 cm, ± 0.4 cm, and the initial beam angle by ± 3 mrad, ± 6 mrad.

Table 6.4: Uncertainty from the initial position and angle of the muon beam.

Set num.	Target	Description	$P_\mu(0)$ from OPERA	Uncertainty ($\times 10^{-4}$) from ± 0.2 cm in x, y , ± 3 mrad in θ_x, θ_y
72	Ag	TECs-in, nominal beam	0.99486	2.3
74	Ag	Nominal A	0.99766	2.3
76	Ag	Steered beam A	0.99172	12.8
86	Al	Steered beam B	0.99248	12.3
87	Al	Nominal D	0.99785	1.0

Magnetic field map uncertainties

The analysis used a fringe field map that was generated with the OPERA software package[83], which had accuracy limitations. Most importantly, the software used a finite element method to solve Maxwell’s equations, and the 40 cm diameter circular hole in the yoke door was expected to introduce difficulties in this modelling³⁸. (The hole was in a critical region for the transverse field components, and hence the depolarisation.) Also, the field map used for the analysis did not include the steel in the floor of the M13 area and the final M13 quadrupoles.

The B_z components from OPERA are compared to measurements using Hall probes in Fig. 6.4(a). Only the z components are shown since the Hall probes were single axis, and did not measure B_x and B_y . Extensive efforts were made to resolve the observed disagreement by adjusting the OPERA inputs, such as 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 door in z . The previous $P_\mu^\pi \xi$ measurement found that variations in these inputs affected $P_\mu(0)$ by 3×10^{-4} at most[57]; as a result of this low sensitivity, and the limitations of the finite element analysis, the tuning of OPERA was not pursued further.

The apparatus that supported the Hall probes had alignment limitations: the Hall probes were attached to an arm that was deflected by gravity, introducing a vertical misalignment of up to 0.1 cm, and the whole mapping device was aligned in the yoke’s coordinate system to about 0.2 cm in x and y . A comparison of the Hall probes and OPERA could not produce a precise translational alignment in x and y , but the position of the muon beam inside the detector did have sensitivity. On a set-by-set basis, a field translation was determined such that the data and simulation positions matched. On average, this required a translation of the entire map by $(\Delta x, \Delta y) = (0.18, 0.19)$ cm. Although this translation was determined precisely, we cannot be sure that it was accurate for a number of reasons. First, the position of the internal muon beam was also sensitive to the solenoid’s coil positions³⁹; their positions were not measured, and were instead initially placed according to a sketch from the manufacturer (Oxford Magnet Technologies Limited UK), and then tuned to match the B_z components in the tracking region, but not in the fringe field. Second, the field through the hole in the yoke should be constrained to have its symmetry axis pass through the centre

³⁸Specifically, finite element analyses are expected to encounter difficulties when there are scales involved that are several orders of magnitude apart. In this case, the important region for the depolarisation of the muons is within $\lesssim 4$ cm of the axis (see Fig. 6.2), 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.

³⁹A change in the solenoid coil positions by ≈ 1 cm caused the internal muon beam to move by about ≈ 0.3 cm. The coil positions were only good to about ≈ 0.2 cm[98].

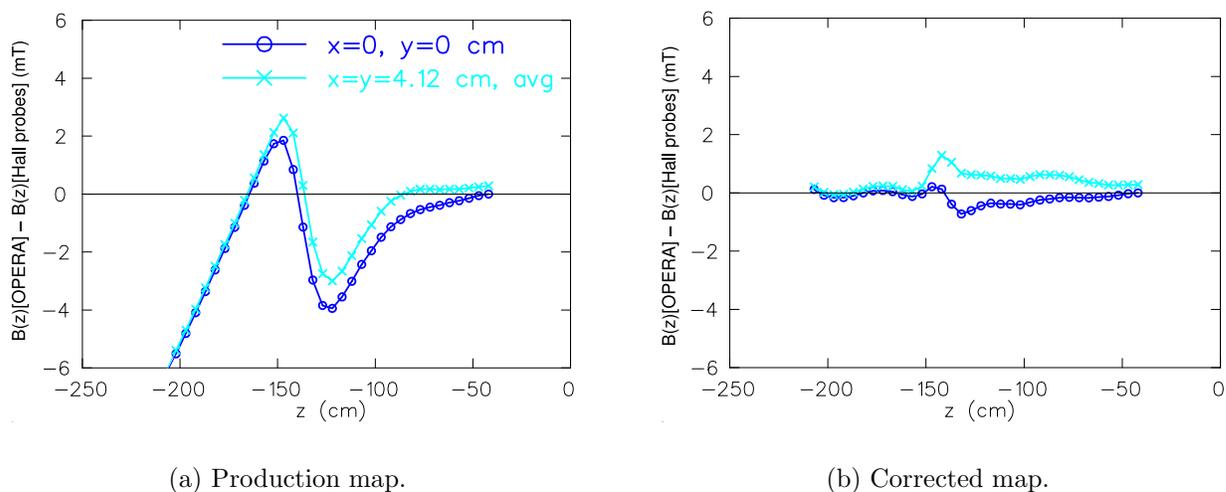


Figure 6.4: Difference in B_z between the magnetic field maps from OPERA and the Hall probes. Two comparisons are shown: the on-axis ($x = y = 0$) and an off-axis average of $x = \pm 4.12$ cm, $y = \pm 4.12$ cm. The corrected map has three current loops added. Muons start in the simulation at $z = -191.944$ cm.

of the hole. Third, the translation was determined from beam profiles that have already been shown to suffer from their own alignment uncertainties. Fourth, the translation may be compensating for the problems in the B_x and B_y components of the field map. For these reasons, the magnetic field translation is treated as an additional uncertainty.

The central value for $P_\mu^\pi \xi$ uses the translated map, but an uncertainty is assigned as the difference in $P_\mu(0)$ between translation and no-translation. The values for these uncertainties appear in Table 6.5; the average $P_\mu(0)$ change for the nominal sets (1.3×10^{-4}) will be assigned as the systematic uncertainty. The increased sensitivity for the aluminium target sets is due to the lower quality of the muon beam⁴⁰.

There is also an uncertainty from the rotation of the magnetic field map. The drift chamber tracking region requires a well determined rotation in (θ_x, θ_y) of (0.3, 1.2) mrad, which is applied by rotating the entire field map. However, there is no guarantee that the fringe field should also be rotated by this amount. The change in $P_\mu(0)$ from making this rotation is shown in Table 6.6. Again, the average $P_\mu(0)$ change for the nominal sets (0.9×10^{-4}) is used as the systematic uncertainty.

⁴⁰For the data accumulated with the aluminium target, a vertical aperture was in place within the M13 beam line. As a result the slits and jaws at the frontend of M13 were opened wider, and the beam was not as well focussed at F3. In addition, muons with $y > 1.0$ cm at the TECs did not converge towards the solenoid's axis.

Table 6.5: Uncertainty from the translational alignment of the magnetic field map.

Set num.	Target	Description	$P_\mu(0)$ from OPERA simulation		Difference ($\times 10^{-4}$)
			(x, y) translation of (0.18, 0.19) cm No rotation	No translation No rotation	
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.99777	0.99770	-0.7
70	Ag	B = 1.96 T	0.99762	0.99755	-0.7
71	Ag	B = 2.04 T	0.99714	0.99710	-0.4
72	Ag	TECs-in, nominal beam	0.99492	0.99492	0.0
74	Ag	Nominal A	0.99762	0.99766	-0.4
75	Ag	Nominal B	0.99767	0.99755	-1.2
76	Ag	Steered beam A	0.99174	0.99068	-10.6
83	Al	Downstream beam package in place	0.99788	0.99773	-1.5
84	Al	Nominal C	0.99776	0.99755	-2.1
86	Al	Steered beam B	0.99244	0.99334	+9.0
87	Al	Nominal D	0.99787	0.99760	-2.7
91	Al	Lower momentum I	0.99691	0.99672	-1.9
92	Al	Lower momentum II	0.99669	0.99648	-2.1
93	Al	Lower momentum III	0.99675	0.99665	-1.0

Table 6.6: Uncertainty from the rotational alignment of the magnetic field map. Both maps are translated in (x, y) by $(0.18, 0.19)$ cm.

Set num.	Target	Description	$P_\mu(0)$ from OPERA simulation		Difference
			No rotation	With rotation ($\times 10^{-4}$)	
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.99777	0.99778	+0.1
70	Ag	B = 1.96 T	0.99762	0.99745	-1.7
71	Ag	B = 2.04 T	0.99714	0.99694	-2.0
72	Ag	TECs-in, nominal beam	0.99492	0.99480	-1.2
74	Ag	Nominal A	0.99762	0.99749	-1.3
75	Ag	Nominal B	0.99767	0.99767	0.0
76	Ag	Steered beam A	0.99174	0.99221	+4.7
83	Al	Downstream beam package in place	0.99788	0.99780	-0.8
84	Al	Nominal C	0.99776	0.99763	-1.3
86	Al	Steered beam B	0.99244	0.99311	+6.7
87	Al	Nominal D	0.99787	0.99782	-0.5
91	Al	Lower momentum I	0.99691	0.99686	-0.5
92	Al	Lower momentum II	0.99669	0.99660	-0.9
93	Al	Lower momentum III	0.99675	0.99673	-0.2

Now that the uncertainties from the initial beam and magnetic field alignments are evaluated, the fringe field validation can be revisited. The difference in polarisation between the nominal and steered beams from a decay parameter fit are shown in the first row of Table 6.7, where the uncertainties are statistical. The next row shows the OPERA predictions for these polarisation differences, where the uncertainties are the quadratic sum of the initial beam position/angle and field translation/angle uncertainties. The OPERA predictions are consistent for sets 87/86, and sets 74/72, but underestimate the polarisation difference for sets 74/76 by 43×10^{-4} , which is well outside of the systematic uncertainties already established. There are no additional uncertainties from the TECs that can explain this discrepancy, so we are left with a problem in the components of the OPERA field map.

Table 6.7: Validation of the fringe field map by establishing how well the simulation reproduces large polarisation differences in the data.

Difference in polarisation	Sets 74,76	Sets 87,86	Sets 74,72
Data	102 ± 8	57 ± 7	16 ± 8
Simulation with OPERA field	59 ± 18	54 ± 17	28 ± 4
Corrected map, with translation	78	133	61
Corrected map, no translation	117	112	63
OPERA field, match A	104	75	27
Corrected field map, match A	108	73	62

An attempt has been made to overcome the limitations of the OPERA finite element method, and resolve the discrepancy in Fig. 6.4(a). An additional field from three on-axis coils has been added to the OPERA map; the coils are located at $z = -265.1$ cm, -147.1 cm, -127.1 cm, with radii 55.1 cm, 25.1 cm, 25.1 cm and their central field strengths are +20 mT, -5.5 mT and +5.5 mT. This approach is motivated by observing that Fig. 6.4(a) resembles the field from a pair of gradient coils, and the final result is shown in Fig. 6.4(b). This approach obeys Maxwell's equations over the region that the muons passed through, but is unlikely to satisfy boundary conditions at the yoke. We feel it is not a coincidence that the coils' z -locations correspond to the outer and inner sides of the yoke door where the circular hole was located, and the last M13 quadrupole, which is not included in OPERA.

If only the B_z components from the OPERA map are replaced with those from the corrected map, then the $P_\mu(0)$ estimates are changed by $< 1 \times 10^{-4}$ for all profiles (this approach does not obey Maxwell's equation). However, the corrected field map significantly alters the transverse (B_x and B_y) components, and changing these has a dramatic effect on $P_\mu(0)$, with decreases in $P_\mu(0)$ of 10 to 20×10^{-4} for the nominal beam profiles, and up 93×10^{-4} for the

steered beam profiles; the $P_\mu(0)$ changes are listed for each set in Table 6.8.

The large changes in $P_\mu^\pi \xi$ can be understood by comparing the transverse components of the OPERA and corrected field maps. Along the symmetry axis ($x = y = 0$) there is no difference, but off-axis there are significant changes; this is demonstrated for $x = y = 1$ cm in Fig. 6.5. (The distance $x = y = 1$ cm is chosen since the beam envelopes in Fig. 6.2 already demonstrates that a large fraction of muons are within this region.) Unfortunately a direct validation of the transverse components is not possible since they were not measured for the real field. Instead we rely on guidance from the simulation's ability to match the large differences in $P_\mu(0)$ from the data. The results from the corrected field appear in Table 6.7; the increase in the transverse components using this map allows the simulation to match the difference in $P_\mu^\pi \xi$ between sets 74 and 76, but the agreement is then worse for the other two entries in the table. These results alone do not allow a strict limit on the B_x and B_y components.

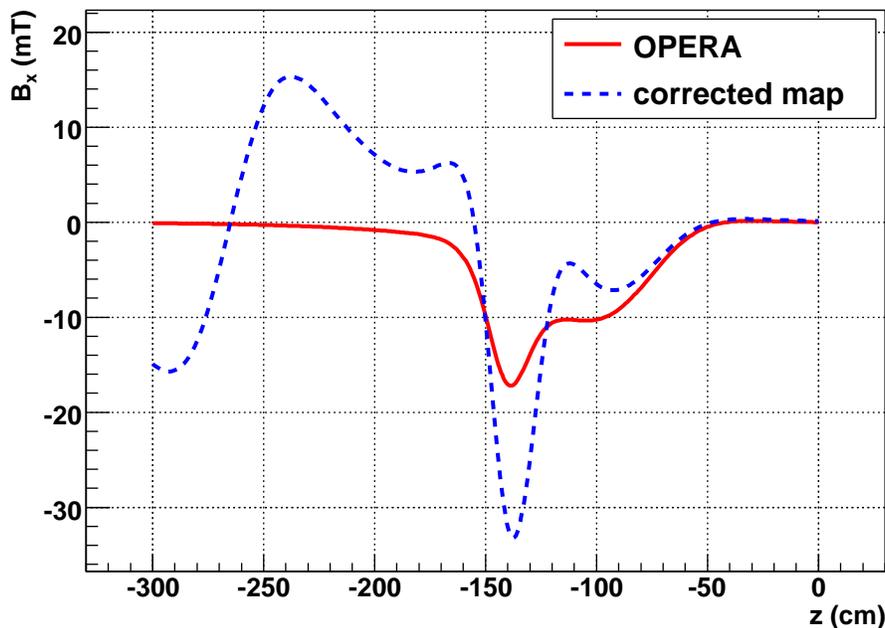


Figure 6.5: A comparison of the B_x components from OPERA and the corrected map, for $x = y = 1$ cm. The maps are identical for $x = y = 0$ cm.

For further guidance on the field components we used the internal muon beam measurements that were described in Section 3.7. The most relevant observation was a strong link between $P_\mu(0)$ and A , which describes the amplitude of muon beam oscillations (*i.e.* how

Table 6.8: Relative polarisation for the simulation, for the OPERA and corrected fringe field maps.

Set num.	Target	Description	$P_\mu(0)$ from simulation		Difference ($\times 10^{-4}$)
			OPERA map With translation	Corrected map With translation	
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.99777	0.99681	-9.6
70	Ag	B = 1.96 T	0.99762	0.99628	-13.4
71	Ag	B = 2.04 T	0.99714	0.99596	-11.8
72	Ag	TECs-in, nominal beam	0.99492	0.99021	-47.1
74	Ag	Nominal A	0.99762	0.99633	-12.9
75	Ag	Nominal B	0.99767	0.99659	-10.8
76	Ag	Steered beam A	0.99174	0.98857	-31.7
83	Al	Downstream beam package in place	0.99788	0.99586	-20.2
84	Al	Nominal C	0.99776	0.99595	-18.1
86	Al	Steered beam B	0.99244	0.98317	-92.7
87	Al	Nominal D	0.99787	0.99646	-14.1
91	Al	Lower momentum I	0.99691	0.99548	-14.3
92	Al	Lower momentum II	0.99669	0.99533	-13.6
93	Al	Lower momentum III	0.99675	0.99515	-16.0

much the mean position of the muon beam moves around). This relationship is shown in Fig. 6.6 for the pairs of sets from Table 6.7. For all combinations of magnetic field translation and initial beam position/angle, the points are well approximated by a quadratic function. When the field map is changed from OPERA to the corrected map, introducing larger B_x and B_y components, the quadratic curve moves down and the points spread out; in other words, A and $P_\mu(0)$ become more sensitive to the alignment of the magnetic field and initial beam. For sets 76/74 and 86/87, a match is possible in A somewhere between the OPERA map and the corrected map.

We note an additional interesting result: if the quadratic relationship between A and $P_\mu(0)$ is used to match the A values in data and simulation, then the difference in $P_\mu(0)$ between sets 76 and 74 can be matched in data and simulation; these results have been added to Table 6.7. Unfortunately this does not work for the other pairs of sets.

A systematic uncertainty from the fringe field must be assigned to the nominal sets, for which the beam was carefully steered onto our beam axis to avoid the larger uncertainties associated with sets 76 and 86, and had the TECs removed in order to minimise the beam size through the field. Therefore simply taking discrepancies from Table 6.7 as an uncertainty would be a gross overestimate. Instead we note from Table 6.7 and the above discussion that the true field is probably closer to OPERA (supported by the results from sets 87/86 and 74/72), but the B_x and B_y components could be as large as those from the corrected field map, since these are needed to accommodate the result for sets 76/74. The strong evidence that OPERA is closer to reality makes it wrong to place the $P_\mu^\pi \xi$ central value half way between the two field maps, so we instead use the OPERA map for the central value, and the differences between OPERA and the corrected map as an asymmetric systematic uncertainty. After averaging over the nominal sets, Table 6.8 gives the uncertainty from our knowledge of the fringe field map as ${}_0^{+14.1} \times 10^{-4}$. The OPERA magnetic field map is used to evaluate the remaining polarisation systematic uncertainties in the next section.

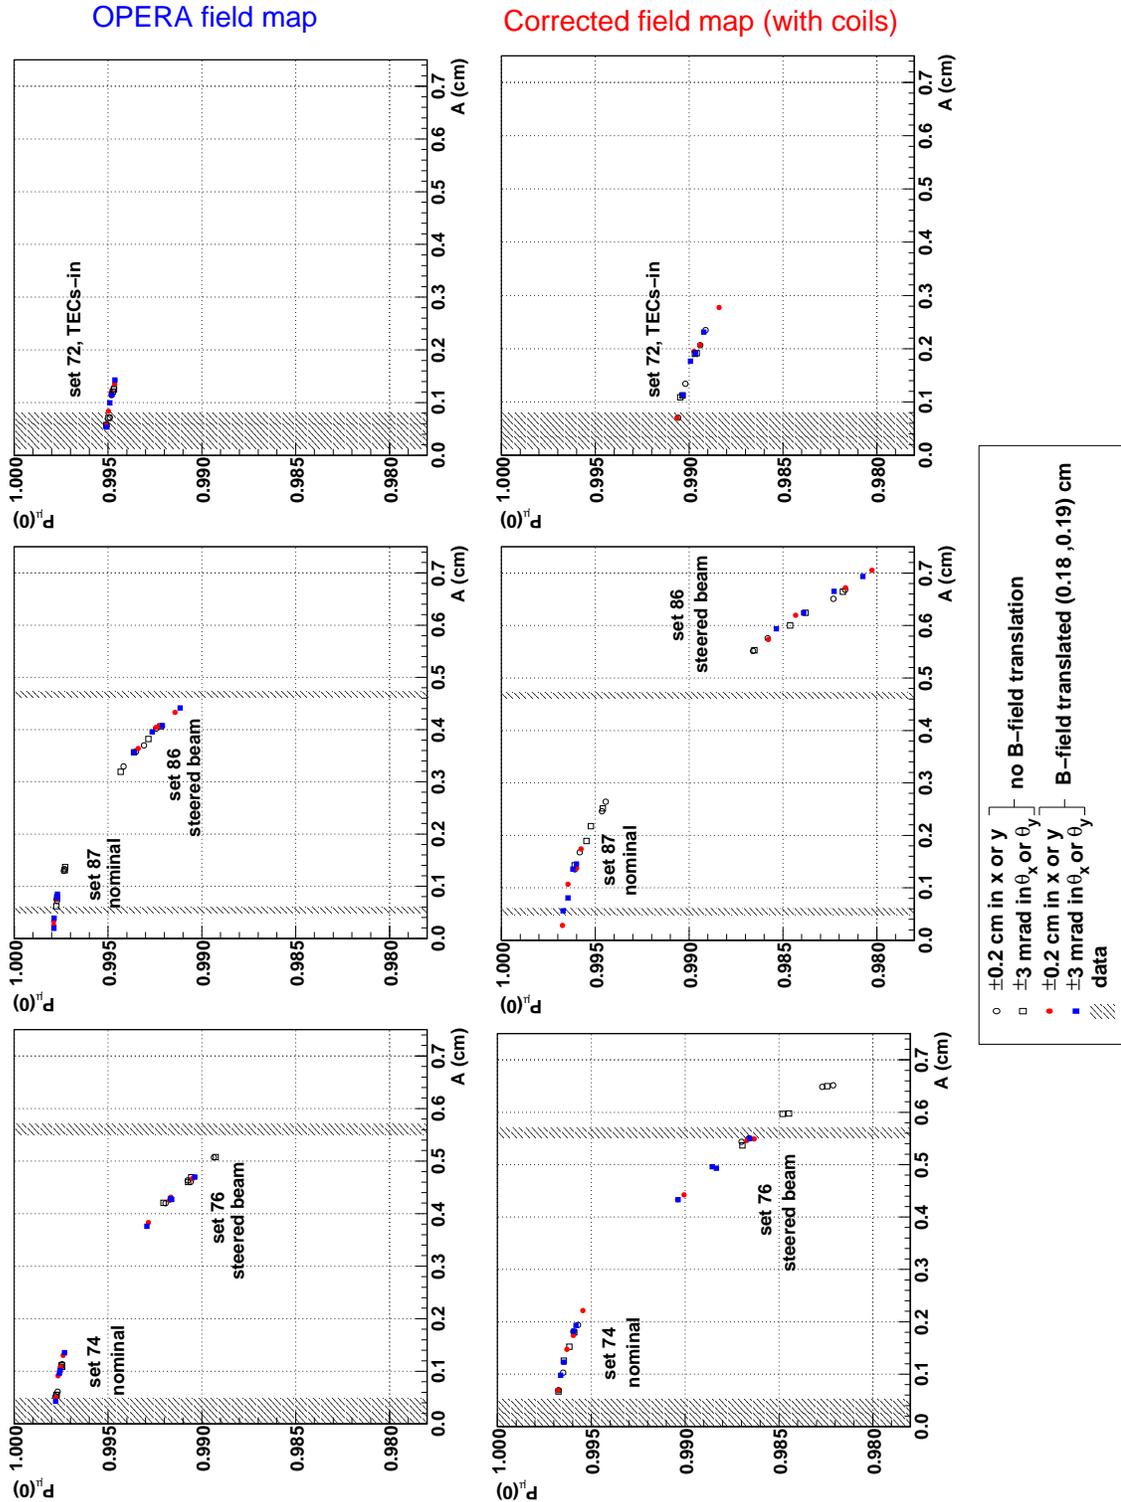


Figure 6.6: The apparent polarisation, $P_{\mu}(0)$ depends quadratically on the amplitude (A) of the internal muon beam oscillations. Each point in the figure corresponds to an alignment uncertainty. The top row shows the OPERA magnetic field map that was used for the production simulations. The bottom row shows the results from a special field map with larger B_x and B_y field components; see the text for more details.

Muon beam angular distribution width

The muons were multiple scattered as they passed through the TECs, resulting in a measured angle distribution that was 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 $c_x = 0.6391$ in the x -module, and $c_y = 0.4795$ in the y -module to account for this multiple scattering. The dependence of $P_\mu(0)$ on the choice of c_x is shown in Fig. 6.7 for a nominal and 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 $c_x = 0.6391$ 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 systematic uncertainties from these factors can be treated as orthogonal to those already evaluated.

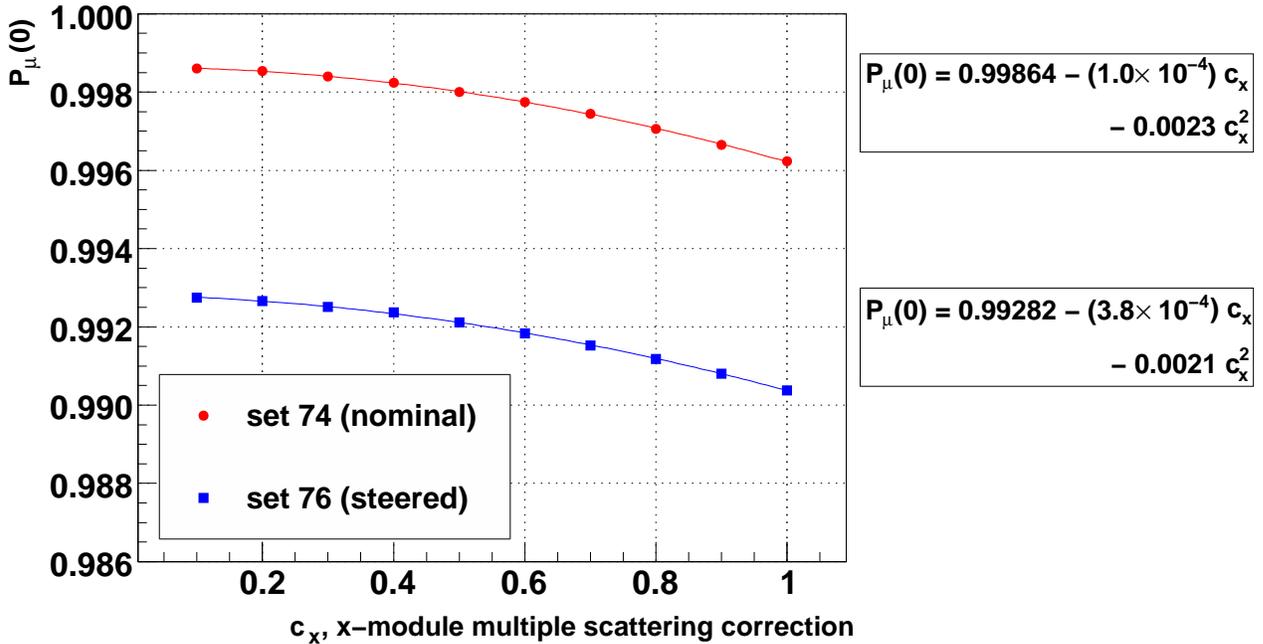


Figure 6.7: Sensitivity of $P_\mu(0)$ to c_x , the multiple scattering correction factor in the x -module. Note that the ratio $c_x/c_y = 63.91/47.95$ was maintained.

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 that

$$\theta_{\text{TECs}+25 \mu\text{m}} \approx \sqrt{\theta_{\text{TECs}+6 \mu\text{m}}^2 + \theta_{19 \mu\text{m}}^2}. \quad (6.1)$$

The results for $\theta_{19 \mu\text{m}}$ from data and simulation are shown in Table 6.9, 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 ${}_{+2.6}^{-3.6} \times 10^{-4}$. 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 6.9 that have not been evaluated.

Table 6.9: Width of reconstructed angle distributions for TEC Mylar windows of thickness $6 \mu\text{m}$ and $25 \mu\text{m}$. The bracketed number indicates the statistical uncertainty in the final digits.

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 (see Appendix G.3.4). Since an event-by-event investigation could not distinguish which variant was the most accurate, the difference between 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 1.7×10^{-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. For each set, $\langle n_x \rangle$ and $\langle n_y \rangle$ are shown in Fig. 6.8. 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 reduced by 7.9% in the x -module, and 3.7% in the y -module. If the larger of these is used, then a conservative systematic uncertainty for the nominal sets due to sense plane aging is ${}_{-1.6}^{+1.4} \times 10^{-4}$.

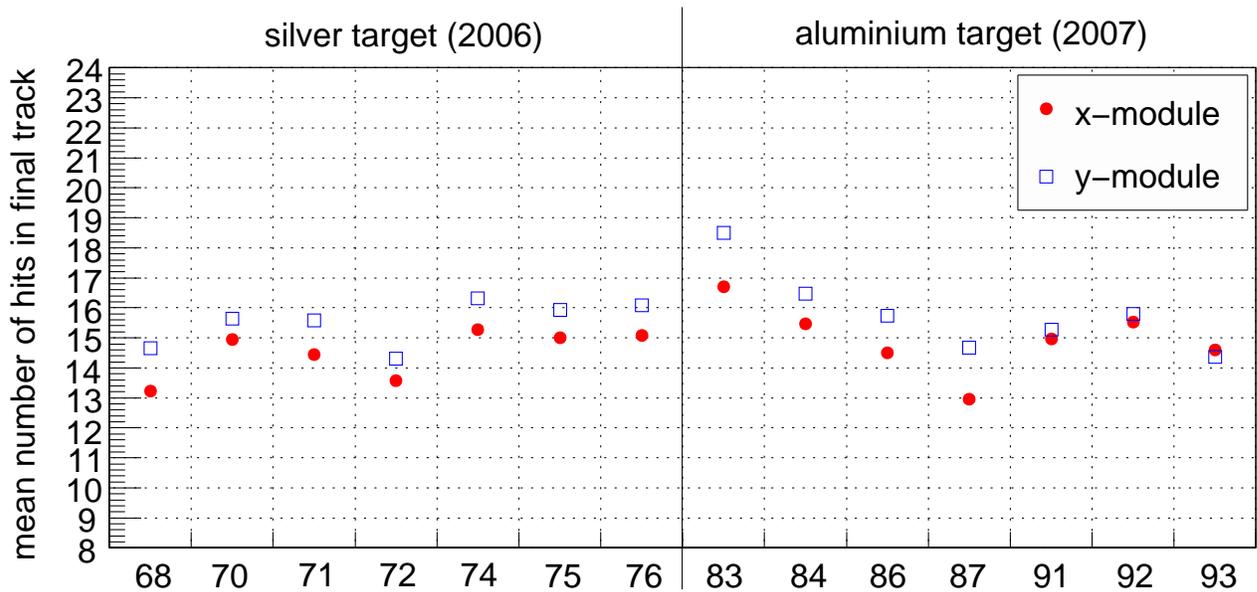


Figure 6.8: Number of hits in final TEC track, for each module. The c_x and c_y correction factors were tuned using set 75, which had $\langle n_x \rangle = 15.0$ and $\langle n_y \rangle = 15.9$.

6.3.2 Production target

The simulation generated muons with anti-parallel spin and momentum vectors, starting from the end of the M13 beam line. This neglected multiple scattering in the graphite production target and the beam line vacuum window, which changed 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 introduced an error of $< 10^{-8}$; see Appendix I).

Surface muons are produced with $p \approx 29.79 \text{ MeV}/c$, but the beam line was nominally tuned to accept muons with an average momentum of $\langle p \rangle = 29.6 \text{ MeV}/c$. Therefore the muons lost $0.19 \text{ MeV}/c$ of momentum on average, which is equivalent to $\approx 3.8 \text{ mg}/\text{cm}^2$ in graphite. (The $3 \mu\text{m}$ beam line vacuum window is safely neglected since it corresponds to just $0.3 \text{ mg}/\text{cm}^2$ of material.) The width of the resulting multiple scattering distribution⁴¹, θ_0 , was found to be 9.3 mrad using a **GEANT4** simulation. The uncertainty is conservatively estimated as $\pm 17\%$ based on the observed discrepancy for **GEANT3** from Section 6.3.1. 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 $-0.9_{-0.2}^{+0.3} \times 10^{-4}$, which is a correction to $\Delta P_\mu^\pi \xi$ of $+0.9_{-0.3}^{+0.2} \times 10^{-4}$.

The previous $P_\mu^\pi \xi$ analysis found a systematic uncertainty of 2×10^{-4} due to depolarisation in the production target; this was evaluated as a conservative upper limit, rather than making a correction[57].

The **GEANT4** simulations were repeated for the lower momentum sets. The widths were $\theta_0 = 24.2 \text{ mrad}$ for $\langle p \rangle = 28.75 \text{ MeV}/c$, and $\theta_0 = 22.9 \text{ mrad}$ for $\langle p \rangle = 28.85 \text{ MeV}/c$. These corresponded to $\Delta P_\mu^\pi \xi$ corrections of $5.9_{-2.2}^{+1.6} \times 10^{-4}$ and $5.2_{-1.9}^{+1.4} \times 10^{-4}$.

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

6.3.3 Stopping material

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. 2.14). Events were only accepted if the muon produced a signal in PC6, but not in PC7. Muons that stopped in the PC6 gas or wires were then removed by cutting on the muon pulse width in the chamber (see Section 3.3.3). This selected a clean sample of muons that stopped in the metal foil.

The weighted asymmetry was constructed according to the method described in Section 3.6, and each data set was fit with

$$P_\mu(t) = P_\mu(0) \exp(-\lambda t). \quad (6.2)$$

The results for the time range $(1.05 < t < 9.00) \mu\text{s}$ are shown in Table 6.10. A weighted average of these relaxation rates finds $\lambda_{\text{Ag}} = (0.840 \pm 0.072) \text{ ms}^{-1}$ and $\lambda_{\text{Al}} = (1.320 \pm 0.077) \text{ ms}^{-1}$.

Table 6.10: Relaxation rate λ for each data set. $P_\mu(t) = P_\mu(0) \exp(-\lambda t)$ has been fit over the nominal time range of $(1.05 < t < 9.00) \mu\text{s}$.

Set num.	Target	Description	λ (ms^{-1})	Fit quality χ^2/ndof	confidence
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.88 ± 0.21	$16.0/20 = 0.80$	0.720
70	Ag	B = 1.96 T	0.86 ± 0.17	$15.6/20 = 0.78$	0.742
71	Ag	B = 2.04 T	0.98 ± 0.18	$26.4/20 = 1.32$	0.153
72	Ag	TECs-in, nominal beam	0.93 ± 0.18	$25.9/20 = 1.29$	0.170
74	Ag	Nominal A	1.02 ± 0.22	$19.4/20 = 0.97$	0.497
75	Ag	Nominal B	0.86 ± 0.18	$12.9/20 = 0.65$	0.880
76	Ag	Steered beam A	0.32 ± 0.20	$12.5/20 = 0.63$	0.897
83	Al	Downstream beam package in place	1.41 ± 0.18	$32.7/20 = 1.63$	0.037
84	Al	Nominal C	1.26 ± 0.19	$24.4/20 = 1.22$	0.225
86	Al	Steered beam B	1.29 ± 0.16	$26.5/20 = 1.33$	0.149
87	Al	Nominal D	1.28 ± 0.18	$13.9/20 = 0.70$	0.833
91	Al	Lower momentum I	1.65 ± 0.33	$20.1/20 = 1.01$	0.449
92	Al	Lower momentum II	1.30 ± 0.29	$16.1/20 = 0.81$	0.708
93	Al	Lower momentum III	1.25 ± 0.22	$15.1/20 = 0.76$	0.770

The simulation used preliminary values of $\lambda_{\text{Ag}} = 0.732 \text{ ms}^{-1}$ and $\lambda_{\text{Al}} = 1.169 \text{ ms}^{-1}$. The weighted asymmetry analysis was applied to the simulation, and found $\lambda_{\text{Ag}} = (0.625 \pm 0.065) \text{ ms}^{-1}$ and $\lambda_{\text{Al}} = (1.104 \pm 0.076) \text{ ms}^{-1}$, using the nominal time range of $(1.05 < t <$

9.00) μs . These results are 1.6σ and 0.9σ below the true values in the simulation, which indicates a potential small and subtle bias in analysis. An independent investigation found that an unbiased muon lifetime measurement required a time fiducial of $(2.00 < t < 9.00) \mu\text{s}$. If the asymmetry analysis is applied to the simulation with a lower time cut of $2.00 \mu\text{s}$, then $\lambda_{\text{Ag}} = (0.614 \pm 0.087) \text{ms}^{-1}$ and $\lambda_{\text{Al}} = (1.19 \pm 0.10) \text{ms}^{-1}$, which are closer to the true values.

The data were reanalysed with the time range of $(2.00 < t < 9.00) \mu\text{s}$, yielding the experiment's most precise unbiased results for the relaxation rates,

$$\lambda_{\text{Ag}} = (0.94 \pm 0.10) \text{ms}^{-1}, \quad (6.3)$$

$$\lambda_{\text{Al}} = (1.20 \pm 0.10) \text{ms}^{-1}. \quad (6.4)$$

Note that these are consistent with the $\mu^+\text{SR}$ results from Appendix H.8:

$$\lambda_{\text{Ag}} = (0.9 \pm 0.2 \text{ (stat.)} \pm 0.2 \text{ (syst.)}) \text{ms}^{-1}, \quad (6.5)$$

$$\lambda_{\text{Al}} = (1.3 \pm 0.2 \text{ (stat.)} \pm 0.3 \text{ (syst.)}) \text{ms}^{-1}. \quad (6.6)$$

The simulation used an inaccurate λ value, and as a result $\Delta P_\mu^\pi \xi$ must be corrected. The effect on the spectrum of a change in λ can be calculated using

$$\frac{\int_{t_1}^{t_2} N(t) \cdot P_\mu(0) \exp(-\lambda_2 t) dt}{\int_{t_1}^{t_2} N(t) dt} - \frac{\int_{t_1}^{t_2} N(t) \cdot P_\mu(0) \exp(-\lambda_1 t) dt}{\int_{t_1}^{t_2} N(t) dt}, \quad (6.7)$$

where $N(t) = N(0) \exp(-t/\tau_\mu)$ and τ_μ is the muon lifetime, and λ_1 and λ_2 are the relaxation rates between which the correction is being made. 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) must be corrected by $+6.3 \times 10^{-4}$ for Ag ($\lambda_1 = 0.732 \text{ms}^{-1}$, $\lambda_2 = 0.94 \text{ms}^{-1}$) and $+0.9 \times 10^{-4}$ for Al ($\lambda_1 = 1.169 \text{ms}^{-1}$, $\lambda_2 = 1.20 \text{ms}^{-1}$). The statistical uncertainty in determining λ from the data causes a P_μ uncertainty of 3.0×10^{-4} for both targets, again using Eq. (6.7).

The simulation found that about 0.2% of muons entered PC7, but did not have enough energy to produce a signal. The depolarisation within the PC gas ($\text{CF}_4/\text{isobutane}$) and wires was about 3%. The systematic uncertainty due to these stops is therefore $\approx 0.2\% \times 3\% = 0.6 \times 10^{-4}$, which is negligible.

6.3.4 Background muon contamination

In the previous $P_\mu^\pi \xi$ analysis, the number of muons downstream of the stopping target did not agree in the data and simulation; this is demonstrated in Fig. 6.9(a). The stopping distributions were consistent if pion decays were simulated in the upstream “beam package” area. Improvements in the classification have removed most of the discrepancy, without having to include the additional pion decays; the modern agreement is demonstrated in Fig. 6.9(b).

The residual discrepancy in Fig. 6.9(b) introduced an uncertainty in the muon stopping distribution, which must be matched to prevent a bias in the muon polarisation, since high angle muons that undergo more depolarisation are preferentially stopped further upstream. Specifically, the simulation needed an extra 1.9 mg/cm² of material to match the stopping distribution in the data (see Section 2.11), and we could not be sure whether this was justified. Fortunately the effect on the polarisation was minimal: including an extra 1.9 mg/cm² in the simulation introduced a systematic uncertainty of just 1×10^{-4} for all the beam profiles except set 72 (TECs-in), which had an uncertainty of 4×10^{-4} .

6.3.5 Beam intensity

The most recent decay parameter analysis found that increasing the beam positron rate in the simulation by a factor of 10 only changed $\Delta P_\mu^\pi \xi$ by $(-5 \pm 7) \times 10^{-4}$ [18] (this was scaled down to obtain a systematic uncertainty; the point is that a very large exaggeration resulted in a sensitivity that was consistent with zero). This confirmed that beam positrons are efficiently removed by the analysis software, and therefore no systematic uncertainty due to the beam positron rate is assigned for this measurement.

The muon rate depended on the proton beam rate and the M13 beam line settings. We were careful not to adjust any M13 beam line setting while accumulating data. Variations in the proton beam rate had no measurable effect on the muon beam measurements. However, the rate did affect the reconstruction of decay positrons; if the rate was raised sufficiently high, then a large number of extra signals from the muons prevented the event classification from working effectively. The degree to which the rate in the simulation and data did not match will therefore introduce a systematic uncertainty. Increasing the muon rate from 2731 s⁻¹ to 27310 s⁻¹ changed $\Delta P_\mu^\pi \xi$ by $(23 \pm 8) \times 10^{-4}$ [18]. This was scaled down by repeating the procedure from Ref. [18] for the modern data and simulation. Specifically, the following steps were taken,

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

types:

$$R_\mu = \frac{(\text{more than one } \mu^+)}{(\text{more than one } \mu^+) + (\text{one } \mu^+, \text{ one decay } e^+)}. \quad (6.8)$$

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,

$$\frac{R_\mu(\text{sim.}) - R_\mu(\text{data})}{R_\mu(\text{data})}. \quad (6.9)$$

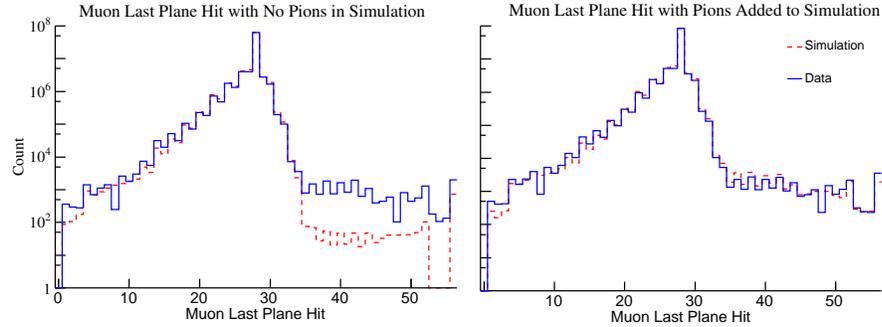
3. Multiply the relative ratio by the average trigger rate from the data, to estimate the error (in s^{-1}) made in the simulation.

4. Divide the simulation's error by the exaggeration ($27310 \text{ s}^{-1} - 2731 \text{ s}^{-1} = 24579 \text{ s}^{-1}$).

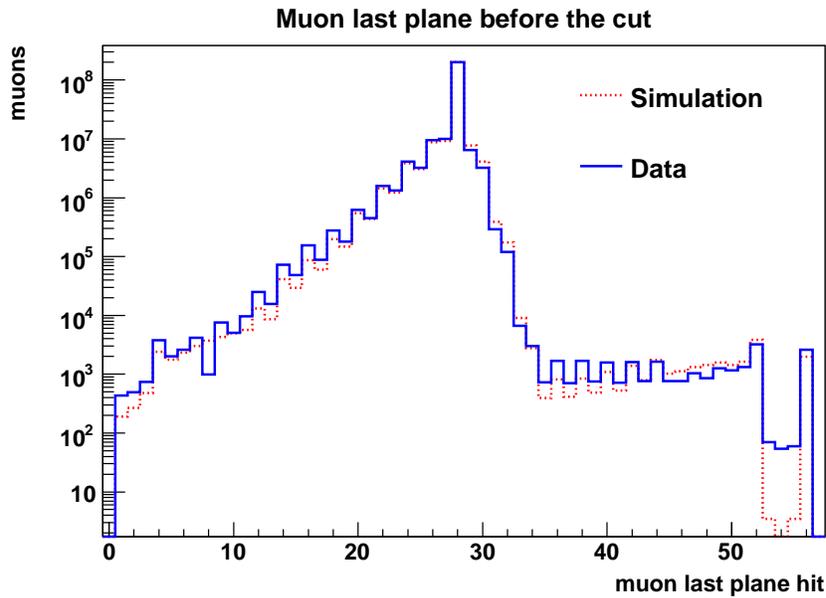
The calculations from each of these steps appear in Table 6.11. An upper limit on the effect of data-simulation rate discrepancies uses the smallest scale factor (28.7, from set 92), which leads to a very conservative uncertainty of $[(23 \pm 8)/28.7] \times 10^{-4} = (0.8 \pm 0.3) \times 10^{-4}$.

Table 6.11: 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. 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



(a) Comparison of muon stopping distribution from the previous analysis (originally Fig. 6.9 from Ref. [57]). The mismatch between data and simulation is resolved by adding muons from pion decays in the M13 beam line.



(b) The same figure re-made for the current analysis. In this figure there are no additional pion decays added.

Figure 6.9: Background muon contamination in the two $P_{\mu}^{\pi} \xi$ analyses.

6.4 Chamber response

6.4.1 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 the helix fitting bias 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 subtle 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. 6.10. 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 $v = 1.8$ cm), where there were few 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. 6.10. Fit this distribution with a fifth order polynomial function.
3. Exaggerate the polynomial function until the χ^2/ndof becomes a factor of two worse. This corresponded to the single hit resolution being degraded by $\approx (1 - \sqrt{2}) \approx 40\%$, and required an exaggeration factor of ten in the polynomial function.
4. Re-analyse the simulation with the scaled STRs, and compare to the nominal simulation.
5. Reduce the $P_\mu^\pi \xi$ change by a scale factor of ten.

This determined the uncertainty as $(0.0 \pm 0.4) \times 10^{-4}$.

Note that the drift cell resolution function and pattern recognition procedure do not have separate systematic uncertainties since they are included in the STR uncertainty.

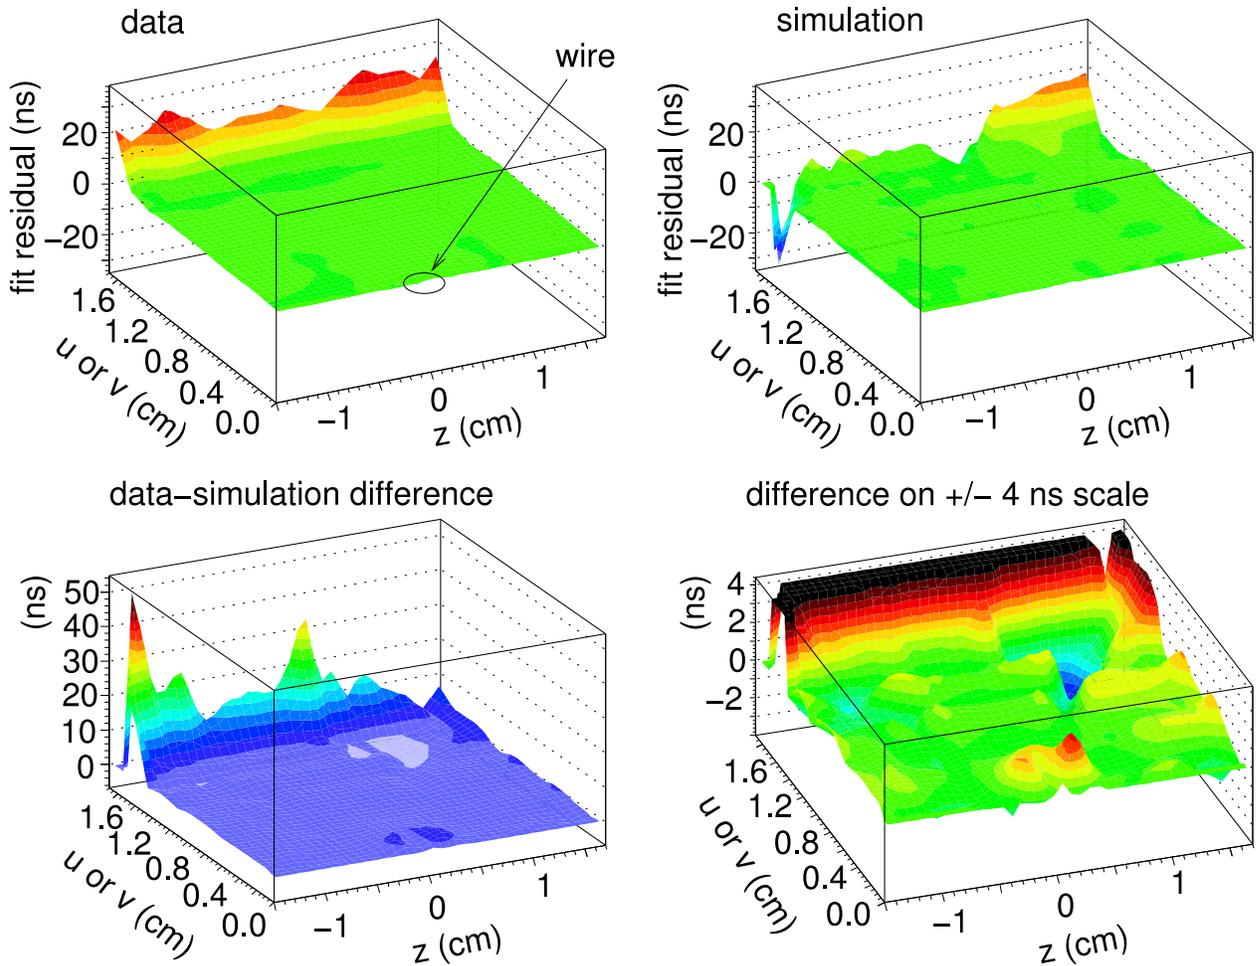


Figure 6.10: Time residuals of the helix fits, after refining the space-time-relationships. Only half a cell is shown, which in practice is reflected about $uv = 0$.

6.4.2 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.

The relative timing of the upstream and downstream halves was determined to 0.050 ns. An exaggerated upstream-downstream shift of 10 ns changes $\Delta P_\mu^\pi \xi$ by $(13 \pm 31) \times 10^{-4}$. After scaling down by $10 \text{ ns}/0.050 \text{ ns} = 200$, this results in a systematic uncertainty of $(0.07 \pm 0.16) \times 10^{-4}$.

For this analysis the wire time offset calibration was also applied to the simulation. This included calibration bias to first order and degraded the resolution of the simulation, but not by a measurable amount. The width of the simulation's time distributions differ from the data, and this has been estimated to introduce a negligible systematic uncertainty of 0.3×10^{-4} [99].

6.4.3 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 were stopped close to the trigger scintillator, and the decay positrons were 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 the other.

The difference in TRE between data and simulation is shown in Fig. 6.11, where a cut has been placed at $(23 < p < 29) \text{ MeV}/c$ to remove beam positrons. There is no evidence of a dependence on p , but there *is* statistical evidence of a linear dependence on $|\cos \theta|$. Specifically, the upstream half of the detector has a data-simulation difference of

$$[(5.1 \pm 1.7) + (-6.4 \pm 2.3)|\cos \theta|] \times 10^{-4}, \quad (6.10)$$

and the downstream half has

$$[(3.5 \pm 2.1) + (-5.3 \pm 2.8)|\cos\theta|] \times 10^{-4}. \quad (6.11)$$

The systematic uncertainty was measured by taking a nominal spectrum and multiplying the number of upstream counts by

$$1 - 10 \times (5.1 - 6.4|\cos\theta|) \times 10^{-4}, \quad (6.12)$$

and the number of downstream counts by

$$1 - 10 \times (3.5 - 5.3|\cos\theta|) \times 10^{-4}, \quad (6.13)$$

where the factor of ten is an exaggeration. This changed $P_\mu^\pi \xi$ by $(12.7 \pm 0.5) \times 10^{-4}$, which is a systematic uncertainty of 1.3×10^{-4} after scaling down by ten.

6.4.4 Dead zone

When a decay positron hit the same wire as the muon, close to where the muon left ionisation, there was reduced gain for the positron due to lingering muon ionisation (the wire had a “dead zone”).

The simulation included the dead zone by removing hits within 0.06 cm of the muon hit, during a recovery time of 3.0 μ s. These parameters were tuned according to the fraction of hits removed by the dead zone in data. In the nominal 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 used an enhanced dead zone: instead of only deadening the wire within 0.06 cm of the muon hit, the entire wire was made insensitive to decay positrons for the 3.0 μ s recovery time. The enhanced simulation 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 $(1/160) \times 25 \times 10^{-4} = 0.2 \times 10^{-4}$.

6.4.5 Chamber foil bulge

The chamber foils were always flat in the simulation, but in the real detector they could bulge when the gas system did not respond sufficiently quickly to ambient pressure changes (see Section 5.4.5). This introduced two systematic uncertainties. First, the electric field

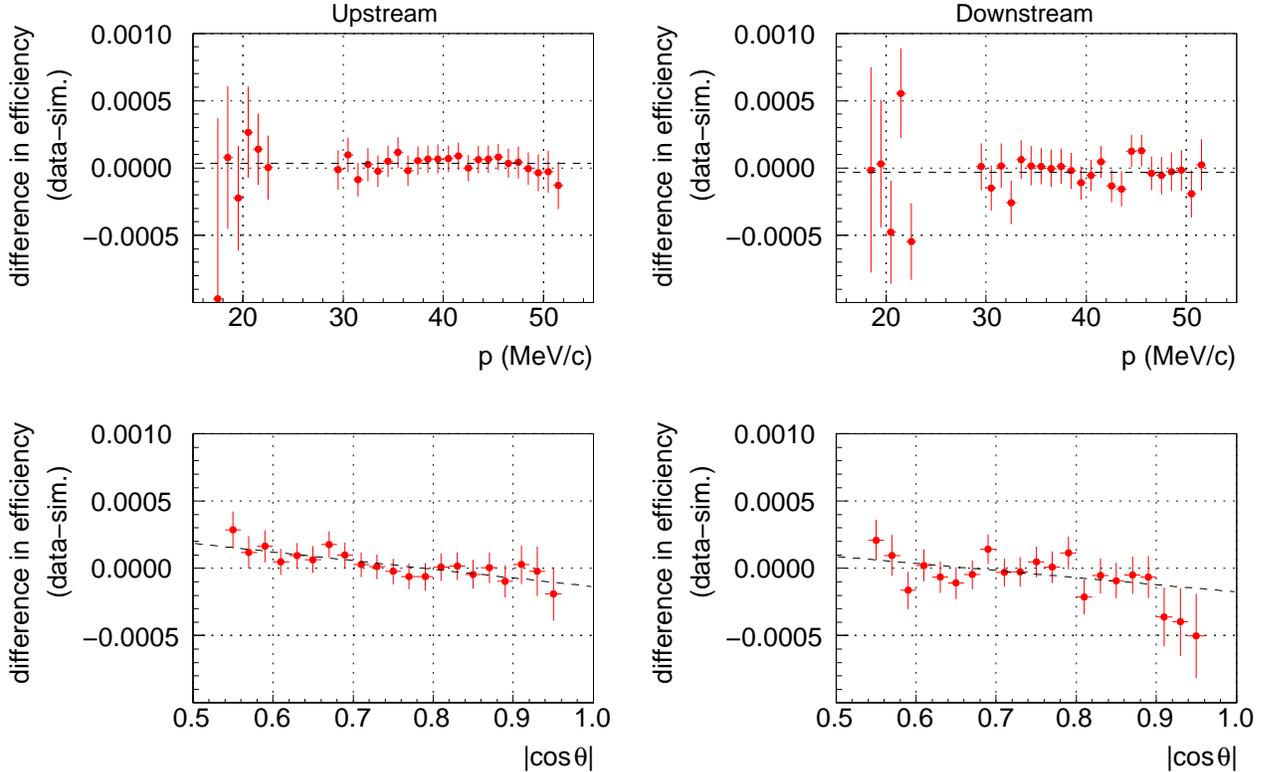


Figure 6.11: Track reconstruction inefficiency for the kinematic fiducial. The upstream inefficiency is determined by how often a positron is reconstructed in the downstream half of the detector, but not the upstream half. The definition is reversed for the downstream half. A cut has been placed at $(23 < p < 29)$ MeV/c in order to remove beam positrons.

between the wire and foil changed, which affected the space-time-relationships in the drift cells. Second, an inwards bulge made the chambers thinner, reducing the number of hits that were available for the helix fitting algorithm.

A previous analysis found these uncertainties were very small, so they are not re-evaluated for the current measurement. Specifically, a bulge of $50 \mu\text{m}$ was found to change $\Delta P_\mu^\pi \xi$ by $(0.2 \pm 0.2) \times 10^{-4}$ for the electric field effect, and $(0.7 \pm 0.7) \times 10^{-4}$ for the reduced hits effect[18]. Since the current analysis rejected runs if they corresponded to a bulge of more than $35 \mu\text{m}$, the total systematic uncertainty cannot be larger than

$$\frac{35}{50} \sqrt{0.2^2 + 0.7^2} \times 10^{-4} = 0.5 \times 10^{-4}. \quad (6.14)$$

6.4.6 Cell asymmetry

The previous $P_\mu^\pi \xi$ analysis used a simulation that centred the wires in z between 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 2.2×10^{-4} . The modern simulation includes the offset, eliminating this uncertainty.

6.4.7 Density variations

In the simulation the gas density inside the drift chambers is constant, but for the data it tracks the atmospheric pressure and exterior temperature, which results in drift cell space-time-relationships (STRs) that vary over the 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[18], 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 0.2×10^{-4} . Therefore the change in STRs due to gas density is a negligible effect, and is not evaluated here.

6.5 Detector alignment

An analysis in 2008 determined the $P_\mu^\pi \xi$ uncertainty from translational and rotational drift chamber alignment as 0.7×10^{-4} [18]. This is adopted here as the $P_\mu^\pi \xi$ uncertainty since the alignment procedures are just as accurate for the current analysis. A summary of this uncertainty will now be given.

Two systematic misalignments of the drift chambers were considered: 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 0.009 ± 10^{-4} and 0.020×10^{-4} for the shear and corkscrew respectively. Random misalignments were not considered, since these smear the tracking residuals and degrade the resolution, which is handled as an orthogonal systematic uncertainty in Section 6.7.

The length (z) and width (u, v) scales are used to determine the momentum components of the reconstructed helices. The z length scale is known to $25 \mu\text{m}$ out of 100cm , which is a fractional uncertainty of 2.5×10^{-5} . The wires were positioned to better than $5 \mu\text{m}$; a

wire plane was 32 cm in width, corresponding to a fractional uncertainty of 1.6×10^{-5} . The corresponding systematic uncertainties were 0.3×10^{-4} for the length scale, and 0.2×10^{-4} for the width scale.

6.5.1 Magnetic field to axis

This section addresses systematic uncertainties from the alignment of the magnetic field for the decay positron reconstruction, which are distinct from the uncertainties associated with determining the polarisation from the muons at the time of decay. The magnetic field is rotationally aligned to the drift chambers to better than 0.03 mrad, within the drift chamber tracking region. This level of uncertainty was previously found to change $P_\mu^\pi \xi$ by just 0.3×10^{-4} [57]; this negligible value is adopted for the current analysis.

A translation uncertainty in the magnetic field of 0.2 cm in x and y was described earlier. A special analysis used a field displaced by 2 cm in both x and y , and the change in $P_\mu^\pi \xi$ was $(1.1 \pm 7.4) \times 10^{-4}$. After scaling down by a factor of ten, this contributes a systematic uncertainty of just 0.1×10^{-4} .

6.6 Positron interactions

6.6.1 δ -electron and Bremsstrahlung rates

Uncertainties from the simulation’s continuous energy loss model are part of the energy calibration uncertainties in Section 6.8.2. The current section describes uncertainties from discrete processes; 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 are radiated using deceleration, which can subsequently undergo pair production). The simulation must accurately reproduce these processes since the extra tracks interfere with the reconstruction. Approximations in the **GEANT3** physics and uncertainties in material thicknesses will 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 was broken in two, with an additional electron track originating from the point where the track was broken. (This measurement is background free, confirmed using a special simulation with the δ -electrons disabled.) The momentum distributions of the reconstructed electrons are compared in data and simulation on the left of Fig. 6.12, where the reconstruction inefficiency 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 data has more δ -electrons than the simulation).

A special simulation increased the δ -electron probability 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. 6.12, where the ratio of counts is 2.80 ± 0.04 . The systematic uncertainty is then

$$\frac{1.007 - 1.0}{2.80 - 1.0} \times 25 \times 10^{-4} = 0.1 \times 10^{-4}. \quad (6.15)$$

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. 6.13. A special simulation increased the Bremsstrahlung probability by a factor of three, and 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. 6.13. 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 Brem rate simulations is 2.82 ± 0.02 . The systematic uncertainty is therefore

$$\frac{1.024 - 1.0}{2.82 - 1.0} \times 55 \times 10^{-4} = 0.7 \times 10^{-4}. \quad (6.16)$$

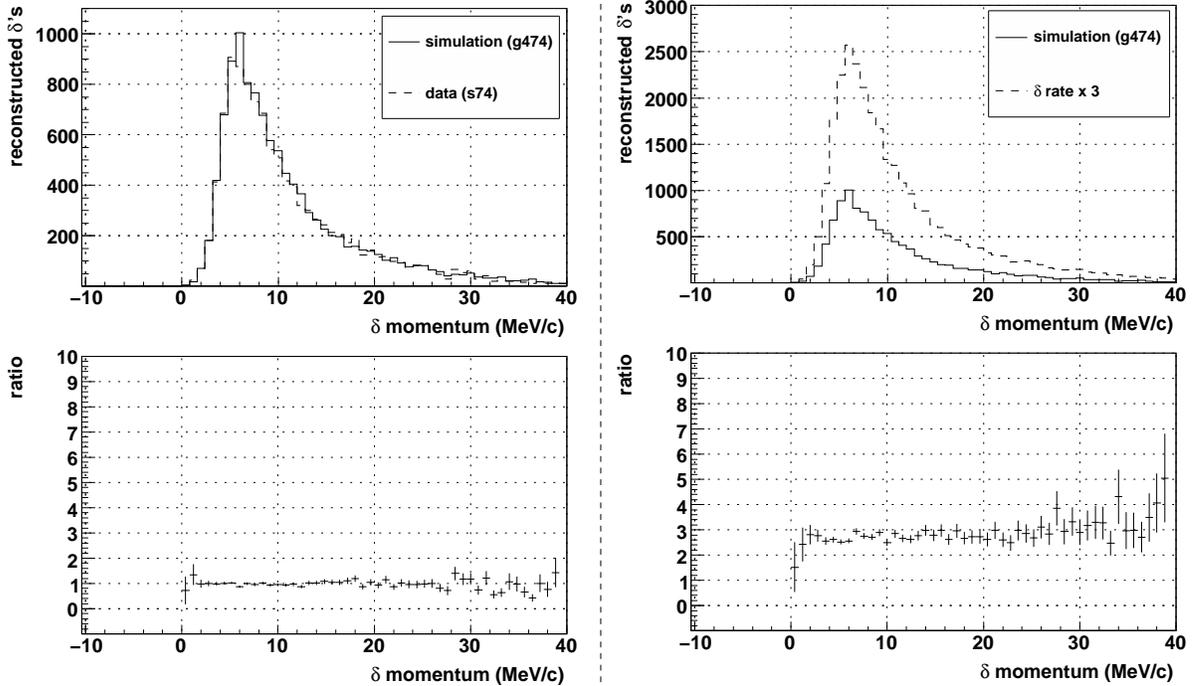


Figure 6.12: The left hand side compares the distribution of reconstructed δ -electrons in data and simulation. The right hand side shows the simulation where the δ production rate was increased by a factor of three.

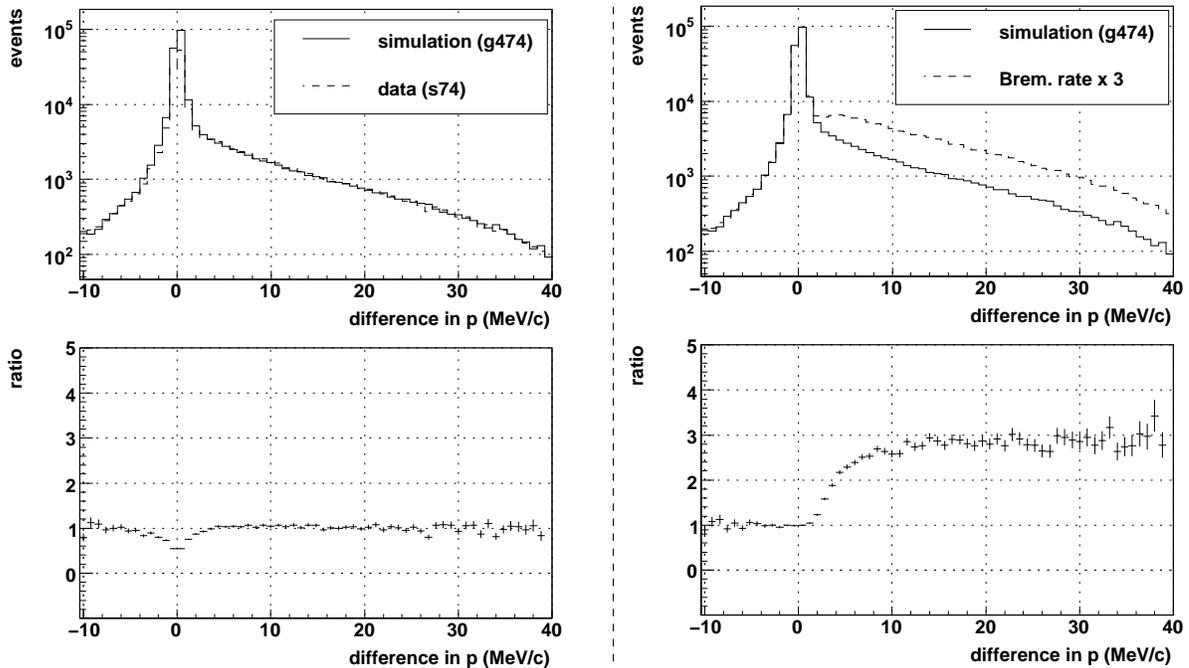


Figure 6.13: The left hand side compares the energy difference for the two halves of a broken track, for data and simulation. The right hand side shows the same distribution for a simulation where the Bremsstrahlung rate was increased by a factor of three.

6.6.2 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 for a description). 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. The backscatters cause extra hits that interfere with the reconstruction of decay positrons. The simulation includes the beam pipe and most of the upstream beam package (excluding, for example, the light guides) but does not include the steel of the yoke since the additional showering significantly increases the computation time. The simulation's lack of yoke, and imperfections in placement and/or thickness of materials is 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 the time distribution. The upper and middle distributions in Fig. 6.14 show this time distribution for the windows where the decay positron is upstream and downstream respectively. The simulation shows evidence of a surplus in upstream backscatters, and a deficit in downstream backscatters 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 backscatters. The difference in $P_\mu^\pi \xi$ between these simulations is $(3.2 \pm 4.0) \times 10^{-4}$, and the two time distributions are shown in the bottom of Fig. 6.14. 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 together quadratically, this results in a final systematic uncertainty of 0.5×10^{-4} .

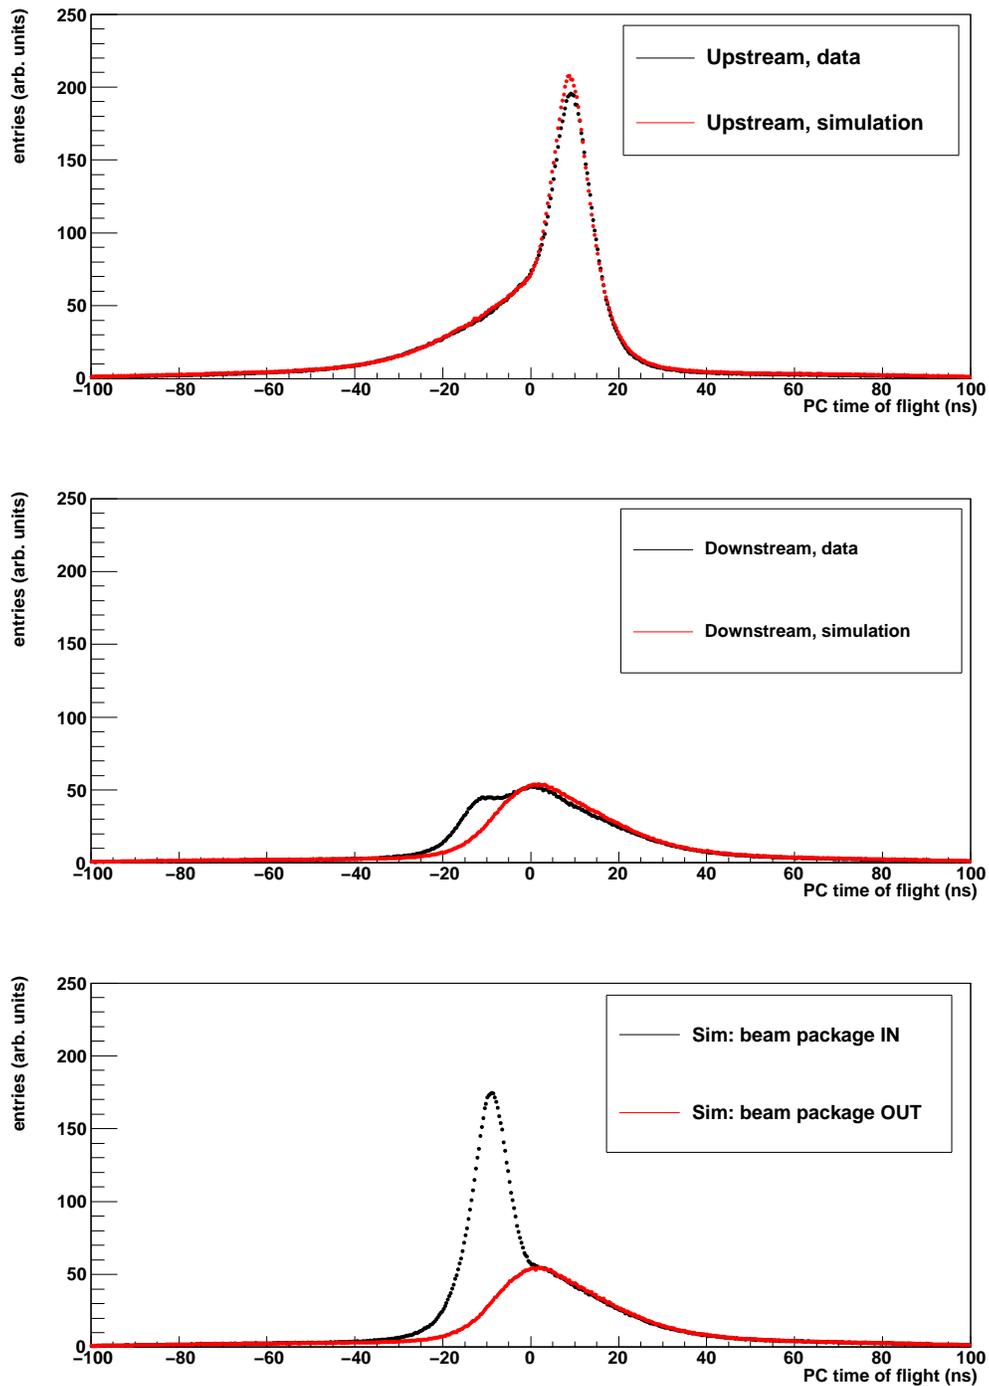


Figure 6.14: The distributions are the difference between the average upstream and downstream PC times. The upstream and downstream backscatters are compared for data and simulation in the upper and middle distributions. The effect of including a downstream beam package in the simulation is shown at the bottom.

6.6.3 Multiple scattering

The helix reconstruction software in the previous $P_{\mu}^{\pi} \xi$ analysis did not permit multiple scattering through the positron trajectory. As a consequence, a systematic uncertainty was assigned. The modern analysis allows for a change in track angle at each pair of planes in the “sparse stack”, and for one change of angle in the “dense stack” (see Section 2.8 for definitions of these terms). As a result of this improvement, there is no longer a systematic uncertainty due to multiple scattering of positrons.

6.6.4 Energy loss

The helix reconstruction software used the same energy loss during tracking as Ref. [18], which determined a negligible systematic uncertainty of 0.01×10^{-4} due to the model. This uncertainty was not re-evaluated.

6.7 Resolution

The momentum (p) and angle ($\cos \theta$) reconstruction resolutions differ in data and simulation. This difference 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 difference at this endpoint in resolution will introduce a systematic uncertainty.

The resolution is measured using the special “upstream 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 difference in momentum between these two reconstructions is shown for data and simulation in Fig. ??, for ?STATEPRANGE? and ?STATECOSTHETARANGE?. The distribution is Landau energy loss convolved with a Gaussian resolution function. If the distribution is fit with a Gaussian close to the peak, the difference in the σ width between data and simulation is a measure of the resolution. (This difference is also sensitive to discrepancies in target thickness and the simulation’s positron interactions physics, but here we conservatively blame the full difference on resolution).

A previous analysis found that the difference in resolution had a non-trivial dependence on momentum and angle. For this measurement the differences are smaller and approximately independent of momentum and angle; the evidence is shown in Fig. ?. These improvements are attributed to the use of better drift cell space-time-relationships, as described in Section 3.2.7.

TODO: explain that quadratic difference of MC minus data is averaged to get scale factor.
 TODO: make Figs. ?? and ??.

6.8 Momentum calibration

6.8.1 Magnetic field shape

The simulation was self-consistent since it used the same OPERA field map for generating and reconstructing the positrons. The data reconstructed the positrons with the OPERA field map, but this had known discrepancies with the mapper measurements within the tracking region, which introduced a systematic uncertainty.

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

$$\delta B_z = C_2 z^2 + C_3 z^3 + C_r r, \quad (6.17)$$

where the optimum C_2 , C_3 and C_r values are recorded in Table 6.12. 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[100].

Table 6.12: Coefficients of Eq. (6.17). 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	-2 ± 2	$+11.4 \pm 0.3$	-19.7 ± 0.5
C_3 ($\times 10^{-10}$ T/cm ³)	-4	-8 ± 4	$+2.0 \pm 0.5$	-2.2 ± 0.7
C_r ($\times 10^{-6}$ T/cm)	-12.5	-8 ± 2	-1.1 ± 0.6	-2.8 ± 0.9

^a The same parameters were used in Refs. [18, 57, 82, 85].

A new field was produced with the coefficients in Eq. (6.17) 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

$$\delta B_r = - \left(C_2 r z + \frac{3}{2} C_3 z^2 r \right). \quad (6.18)$$

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 systematic uncertainty for the nominal sets is 0.3×10^{-4} . Previous evaluations of this uncertainty used a smaller scale factor of ten, and did not apply Eq. (6.18), but still found an effect below 1×10^{-4} [18, 57].

In previous analyses the fields for analysing data at 1.96 T and 2.04 T were created by simply scaling the 2.0 T OPERA map by $\pm 2\%$. From the parameters in Fig. 6.12, this would have resulted in scale factors of 6 and 10, and systematic uncertainties of 1.5×10^{-4} and 2.6×10^{-4} for 1.96 T and 2.04 T respectively. However, for this measurement the 1.96 T and 2.04 T data were analysed with *corrected* maps, leading to negligible systematic uncertainties.

6.8.2 Use of 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, and this must be corrected by shifting or scaling the entire data spectrum. (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 two uncertainties from the energy calibration: a statistical part since only a limited region of the spectrum is used to establish the required correction, and a systematic part since either a shift or scale (or combination) must be used to propagate the correction to the rest of the spectrum.

The difference in kinematic endpoints between data and simulation is shown for a nominal set (fiducial only) in Fig. 6.15. In previous analyses the upstream ($1/\cos\theta < 1$) and downstream ($1/\cos\theta > 1$) points have been 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 parameters. In the current analysis the stopping distribution in data and simulation is better matched, introducing the possibility of a fit with a single intercept and slope (two parameters). Using all available data/simulation spectra there is no preference from the reduced- χ^2 for a two or four parameter fit, and the discussion of which model to use is ongoing since this choice is significant for the ρ and δ parameters.

There are two extreme choices for how to propagate the difference in endpoint to the rest

of the spectrum. In the first case the entire spectrum is shifted in momentum according to

$$p_{\text{corrected}} = p_{\text{reconstructed}} - \left(b - \frac{a}{|\cos \theta|} \right). \quad (6.19)$$

In the second case the spectrum is scaled depending on momentum so that

$$p_{\text{corrected}} = \frac{p_{\text{reconstructed}}}{1 + b/W_{e\mu}} + \frac{a}{|\cos \theta|}, \quad (6.20)$$

where $W_{e\mu}$ is the maximum kinematic positron momentum (52.83 MeV/c). The statistical uncertainties for combinations of shift/scale and two/four parameter fits are explored in Table 6.13. Since there is no firm decision on the number of parameters, or a preference for shift over scale, the statistical uncertainty is conservatively taken as 1.4×10^{-4} , which is the largest entry from Table 6.13. For the systematic part, the central value of $P_{\mu}^{\pi} \xi$ is placed half way between the shift and scale extremes, and a systematic uncertainty of $0.5 \times 1.4 \times 10^{-4} = 0.7 \times 10^{-4}$ is assigned to cover both possibilities.

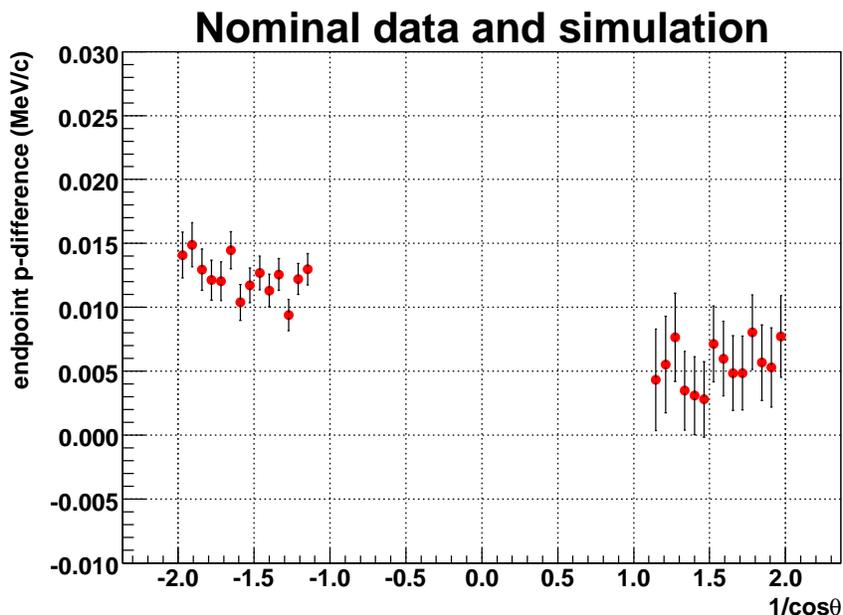


Figure 6.15: Difference between reconstructed momentum in data and simulation at the endpoint of the muon decay spectrum.

Table 6.13: Statistical and systematic uncertainties for the energy calibration, under the possible fitting scenarios that are described in the text.

Number of fit params.	$P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)		
	Shift model, Eq. (6.19)	Scale model, Eq. (6.20)	Systematic (difference between shift and scale)
4	0.6	1.4	0.2
2	0.5	0.3	1.4

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; see Section 1.4.5 for more details. Note that the contribution from the $O(\alpha^2 L^2)$ term is $\sim 10^2$ stronger than the $O(\alpha^3 L^3)$ term, making it the highest order term used by TWIST. The next highest term *not* in use is $O(\alpha^2 L^0)$, which was calculated in 2007 by Ref. [16]. The same reference shows the ratio of the $O(\alpha^2 L^1)$ to $O(\alpha^2 L^0)$ terms over the TWIST kinematic fiducial, and demonstrates that the $O(\alpha^2 L^0)$ term is at most 0.2 of the of $O(\alpha^2 L^1)$ term.

A pure $O(\alpha^2 L^1)$ spectrum was added to a nominal spectrum, and $P_\mu^\pi \xi$ changed by $(2.9 \pm 0.1) \times 10^{-4}$. The nominal spectrum had 4.79×10^8 events, and the combined spectrum had 2.28×10^{10} events, corresponding to an increase by a factor of 47.6. Therefore the effect on $P_\mu^\pi \xi$ of only adding 0.2 of the pure $O(\alpha^2 L^1)$ spectrum is

$$(0.2/47.6) \times 2.9 \times 10^{-4} = 0.01 \times 10^{-4}. \quad (6.21)$$

This is 50 times smaller than the uncertainty reported in Ref. [18], which has a mistake in the number of events in each spectrum[101].

6.9.2 η correlation

The muon decay spectrum does not allow a precise measurement of the parameter η . Therefore η was fixed to its world average value, $(-36 \pm 69) \times 10^{-4}$, 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 1.1×10^{-4} .