

University of Alberta

A PRECISION MEASUREMENT OF THE MUON DECAY PARAMETERS ρ AND δ

by

Robert Paul MacDonald



A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of **Doctor of Philosophy**.

Department of Physics

Edmonton, Alberta
Fall 2008



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-46372-7

Our file Notre référence

ISBN: 978-0-494-46372-7

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

"Sometimes, if you pay real close attention to the pebbles you find out about the ocean."

—*Lords and Ladies*, by Terry Pratchett

To Nate Rodning

Abstract

The *TWIST* collaboration at TRIUMF has performed new measurements of two of the Michel parameters of muon decay: ρ , which governs the shape of the overall momentum spectrum, and δ , which governs the momentum dependence of the parity-violating decay asymmetry. As with all of the measurements from *TWIST*, these were carried out using a blind analysis. We find $\rho = 0.75014 \pm 0.00017(\text{stat}) \pm 0.00044(\text{syst}) \pm 0.00011(\eta)$, where the last uncertainty represents the correlation between ρ and the decay parameter η , and $\delta = 0.75067 \pm 0.00030(\text{stat}) \pm 0.00067(\text{syst})$. These are consistent with the value of $3/4$ given for both parameters in the Standard Model of particle physics. The new measurements are a factor of two more precise than the measurements previously published by *TWIST*, and five times more precise than measurements before *TWIST*. Combined with other muon decay data, an improved indirect limit on the decay parameter $P_\mu \xi$ of 0.99726 ± 0.00132 is obtained, also consistent with its Standard Model value of $P_\mu \xi = 1$. These results set new model-independent constraints on the possible weak interactions of right-handed particles. Specific implications for left-right symmetric models are discussed.

Acknowledgements

I would like to thank my current supervisors, Art Olin and Roger Moore, and my previous supervisors at the University of Alberta, Manuella Vinciter and the late Nathan Rodning. Their guidance, support, and shared love of physics has been invaluable throughout this process.

I would like to thank the members of the *TWIST* collaboration and the staff at TRIUMF, without whom the experiment obviously could not have been done. I am particularly grateful to Glen Marshall, David Gill, Dick Mischke, and Carl Gagliardi for their guidance and teaching; to Konstantin Olchanski, Renée Poutissou, and Robert Openshaw for detailed discussions of aspects of the experiment and the concepts behind them; and to my fellow graduate students on the experiment, Jim Musser, Andrei Gaponenko, Blair Jamieson, Ryan Bayes, James Bueno, and Anthony Hillairet, whose collaboration helped us each to learn and achieve more than we could alone. I would also like to thank all those who provided feedback and editing on this thesis, including several people mentioned above and Alexander Grossheim. Many thanks, too, to the members of my supervisory committee.

I would like to thank my parents, Paul and Mary MacDonald, and the rest of my family, for their advice, support, and encouragement. Special thanks to my wife, Claire MacDonald, for her love, friendship, and understanding.

The funding from the National Science and Engineering Research Council, and the support of the University of Alberta, are gratefully acknowledged. The *TWIST* experiment was also funded by grants from the National Science and Engineering Research Council, and from the U.S. Department of Energy. The Westgrid computing facility was an important factor in the success of this work.

Table of Contents

1	Introduction	1
2	Muon Decay and the Weak Interaction	4
2.1	The Weak Interaction and the Standard Model	4
2.2	Kinematics	6
2.2.1	Kinematics of Pion Decay	6
2.2.2	Kinematics of Muon Decay	7
2.3	General Form of Muon Decay	7
2.4	The Parameterized Muon Decay Spectrum	9
2.4.1	Asymmetry at the Spectrum Endpoint	11
2.5	Extensions to the Standard Model	12
2.5.1	Right-Handed Muon Decay	12
2.5.2	Left-Right Symmetric Models	12
3	Review of Previous Measurements	14
3.1	ρ : Peoples 1966 (unpub.), Derenzo 1969	15
3.2	δ : Balke 1988	16
3.3	$P_\mu \xi \delta / \rho$: Jodidio 1986	17
3.4	η : Derenzo 1969, Burkard 1985, Danneberg 2005	17
3.5	Global Analysis: Gagliardi <i>et al.</i> 2005	18
4	<i>TWIST</i> Experimental Setup	20
4.1	TRIUMF Beamlines	20
4.1.1	Proton Beam and Production Target	20
4.1.2	Muon Beamline	22
4.1.3	Muon Beam Monitor	27
4.2	<i>TWIST</i> Detector	27
4.2.1	Magnet	27
4.2.2	Trigger and Other Scintillators	28
4.2.3	Tracking Chambers	29
4.2.4	Stopping Target	31
4.2.5	<i>TWIST</i> Coordinate Systems	32
4.2.6	<i>TWIST</i> Control Software	33
4.3	Calibration	33
4.3.1	Alignment	33
4.3.2	Wire Time Offsets (t_0)	34
5	<i>TWIST</i> Analysis	35
5.1	Michel Spectrum Fitting	37
5.2	Chamber Drift Time Calculations	41
5.3	Data Analysis	41
5.3.1	Description of Analysis	41

5.3.2	Improvements in Analysis Software Since Previous <i>TWIST</i> Measurements	48
5.3.3	Error Found Since Completion of Analysis	48
5.4	Positron Spectrum Summation	48
5.4.1	Spectrum Summation Method	49
5.4.2	Event Cuts and Track Selection	49
5.4.3	Improvements in Spectrum Summation Software Since Previous <i>TWIST</i> Measurements	52
5.5	Energy Calibration	52
5.5.1	Relative Energy Calibration at the Endpoint	54
5.5.2	Application of the Energy Calibration	55
5.5.3	Absolute Energy Calibration	55
5.6	Fiducial Region Considerations	58
6	<i>TWIST</i> Simulation	61
6.1	Michel Spectrum Generation and Blind Analysis	61
6.2	Monte Carlo Simulation	63
6.2.1	Simulation Technique	63
6.2.2	Muon Decay Spectrum	65
6.2.3	Optional Output Data Banks: "Truth Banks"	66
6.2.4	Calibration Files in Simulation	66
6.2.5	Improvements in Simulation Software over Previous Analysis	66
6.3	Validation of <i>TWIST</i> Simulation	67
6.3.1	Positron Interactions	67
6.3.2	Reconstruction Efficiency	75
6.3.3	Outside Material	78
6.3.4	TDC Spectra	84
7	<i>TWIST</i> Data	88
7.1	Description of Data Sets	89
7.1.1	Matching Simulation to Data	90
7.2	Data Quality	91
7.2.1	Online Monitoring and Regulation	91
7.2.2	Consistency of Measurements	92
8	Corrections and Systematic Uncertainties	94
8.1	Chamber Response	97
8.1.1	Chamber Drift Times (STRs): Time-Independent Component	97
8.1.2	Chamber Drift Times (STRs): Gas Density Variations	102
8.1.3	Chamber Drift Times (STRs): Foil Bulge Variations	104
8.1.4	Foil Bulge: Geometric Effects	105
8.1.5	Efficiency Asymmetry	107
8.1.6	Wire Time Offsets (t_0 's)	107
8.2	Energy Scale	114
8.2.1	Magnetic Field Map Shape	114
8.2.2	Energy Calibration	115
8.2.3	Momentum Dependence of Energy Calibration	117
8.3	Positron Interactions	118
8.3.1	Simulated Bremsstrahlung and Delta Ray Production Rates	119
8.3.2	Outside Materials	126
8.4	Reconstruction Resolution	127
8.4.1	Momentum and Angle Resolution	129
8.5	Spectrometer Geometry	139
8.5.1	Alignment	139

8.5.2	Spectrometer Length Scales	142
8.6	Beam Intensity	144
8.6.1	Muon and Positron Rates	144
8.7	External Uncertainties	149
8.7.1	Radiative Corrections	149
8.7.2	Correlations with η	151
8.8	Negligible Systematic Uncertainties	152
9	Results	153
9.1	Results Revealed	153
9.1.1	“White Box” Consistency Check	156
9.2	Results in Context	157
9.2.1	Global Analysis of Muon Decay	157
9.2.2	Implications for Physics Beyond the Standard Model	158
9.2.3	Future Work	159
9.3	Conclusion	159
	Bibliography	161
	A Window Type Definitions	164
	B Event Type Definitions	166
	Index	169

List of Tables

2.1	The fermions of the Standard Model.	5
3.1	Standard Model values and direct experimental limits on the muon decay parameters.	14
3.2	Limits on the weak coupling constants.	15
3.3	Results of the 2005 global analysis of muon decay data.	19
5.1	Fiducial cuts used during spectrum fitting.	58
6.1	Peak positions and widths of energy loss and scattering distributions.	71
6.2	Hard interaction probabilities, in data and simulation.	73
6.3	Linear fits to peak energy loss vs. $1/\cos\theta$ for bins of momentum.	74
6.4	Integrated inefficiencies over the fiducial regions, for data and simulation.	78
6.5	Normalized backscatter counts for data and simulation.	84
7.1	Description of data sets used for Michel parameter measurements.	89
7.2	Standard M13 beamline element settings used for 2004 data.	90
7.3	Mean values and RMS variations of several experimental variables.	91
8.1	Summary of <i>TWIST</i> systematic uncertainties.	96
8.2	List of corrections to final Michel parameter measurement.	96
8.3	Summary of Chamber Response systematics.	97
8.4	Fits between spectra reconstructed using measured STRs to those using Garfield STRs, for data and MC. The differences are corrections to the measured decay parameters.	102
8.5	Results of comparative energy calibration between simulation with low-density STR file and standard.	103
8.6	Unscaled Michel parameter fit for the gas density variation systematic.	103
8.7	Systematic uncertainty due to variations in the drift chamber gas density.	103
8.8	Systematic uncertainty due to variations in the drift chamber foil bulge.	105
8.9	Unscaled Michel parameter fits for the geometric foil bulge systematic.	106
8.10	Michel parameter fits for the geometric foil bulge systematic, after applying the scale factor.	106
8.11	Unscaled Michel parameter fit for the reconstruction efficiency asymmetry.	107
8.12	Systematic uncertainty due to asymmetric reconstruction efficiency.	107
8.13	Effects of doctored t_0 files on the χ^2/dof results from the helix fitter, for measurement bias and time variation tests, for different scale factors.	113
8.14	Unscaled Michel parameter fits for t_0 systematics, based on measurement bias and on time variations	113
8.15	Systematic uncertainties due to uncertainties in measurement bias and time variations of the t_0 measurements.	113
8.16	Summary of Energy Scale systematics.	114
8.17	Unscaled Michel parameter fits for field map systematic.	115

8.18	Systematic uncertainty due to uncertainties in the shape of the magnetic field map.	115
8.19	Unscaled Michel parameter fits for the energy calibration systematic.	116
8.20	The error matrix for a standard energy calibration.	116
8.21	Systematic uncertainty due to uncertainties in the energy calibration.	117
8.22	Michel parameter fits for systematic uncertainty due to the momentum dependence of the energy calibration.	118
8.23	Summary of Positron Interaction systematics.	118
8.24	Results of the relative energy calibration between simulation with increased delta ray or bremsstrahlung production and standard surface muon simulation.	123
8.25	Unscaled Michel parameter fits for the delta ray production rate systematics.	123
8.26	Unscaled Michel parameter fits for the bremsstrahlung production rate systematics.	123
8.27	Systematic uncertainty due to errors in the simulation of the delta ray production rate.	126
8.28	Systematic uncertainty due to errors in the simulation of the bremsstrahlung production rate.	126
8.29	Unscaled Michel parameter fits for the effect of outside materials.	127
8.30	Systematic uncertainty due to errors in the simulation of outside materials.	127
8.31	Summary of Resolution systematics.	128
8.32	Constants for parameterization of data-MC resolution differences.	135
8.33	Unscaled Michel parameter fits for the angle resolution systematic.	136
8.34	Systematic uncertainties due to the data-MC angle resolution mismatch.	136
8.35	Unscaled Michel parameter fits for momentum smearing, before and after energy calibration.	136
8.36	Systematic corrections and uncertainties due to the data-MC momentum resolution mismatch.	137
8.37	Summary of Spectrometer Geometry systematics.	139
8.38	Effects of shear and corkscrew exaggerations on the helix fit χ^2/dof	141
8.39	Unscaled Michel parameter fits for translational and rotational alignment systematics.	142
8.40	Systematic uncertainties due to uncertainties in translational and rotational alignments.	142
8.41	Unscaled Michel parameter fits for length and width scale systematics.	143
8.42	Systematic uncertainties due to uncertainties in the length and width scales of the <i>TWIST</i> detector.	144
8.43	Summary of Beam Intensity systematics.	144
8.44	Multiple muon fractions in data and simulation, for each data set.	145
8.45	Beam positron fractions in data and simulation, for each data set.	146
8.46	Unscaled Michel parameter fits for the rate systematics, based on high muon and beam positron rates.	149
8.47	Systematic uncertainties due to errors in the simulation of the muon and beam positron rates.	149
8.48	Summary of External Uncertainties.	149
8.49	Unscaled Michel parameter fits for the radiative correction systematic.	150
8.50	Systematic uncertainties due to the radiative corrections of the decay spectrum.	151
8.51	Systematic uncertainties due to correlations between η and the other Michel parameters.	152
9.1	Measured differences in Michel parameters between data and simulation, for each data set.	153
9.2	Measured values of the Michel parameters, with black box values and corrections applied.	155
9.3	Results of a new global analysis of muon decay data.	157

9.4	New limits on the weak coupling constants from this analysis.	158
-----	---	-----

List of Figures

2.1	Michel Spectrum of muon decay.	9
2.2	Effect of the radiative corrections on the Michel spectrum shape.	10
4.1	Schematic view of the <i>TWIST</i> detector.	21
4.2	M13 beamline layout.	22
4.3	Muon flux vs. dipole magnetic field, for calibrating the M13 beamline momentum.	24
4.4	Time structure of the M13 beam (CPTOF) at 29.6 MeV/c.	26
4.5	Side view of the <i>TWIST</i> spectrometer, showing the arrangement of the tracking chambers.	29
4.6	Construction of the muon stopping target.	32
5.1	Flowchart of <i>TWIST</i> analysis chain.	36
5.2	Momentum dependences of the decay spectrum derivatives.	38
5.3	Two-dimensional distributions of the Michel spectrum and its derivatives.	39
5.4	Events rejected by each cut during spectrum summation, for data and simulation.	50
5.5	Normalized momentum edges used for energy calibration.	55
5.6	χ^2 of comparative energy calibration as MC is shifted.	56
5.7	Comparative energy calibration example.	57
5.8	Diagram of the fiducial region	59
6.1	Distribution of events accepted in upstream stops studies, for data and simulation.	68
6.2	Integrated energy loss distribution, near the peak energy loss region.	70
6.3	Integrated scattering distribution, near the peak scattering region.	70
6.4	Integrated energy loss distribution, for large energy loss events.	72
6.5	Integrated scattering distribution, for large scattering events.	72
6.6	Peak energy loss as a function of $1/\cos\theta$ for several bins of momentum.	74
6.7	Upstream reconstruction inefficiency $(1 - \epsilon)$ as a function of $(p, \cos\theta)$, for data and simulation.	76
6.8	Downstream reconstruction inefficiency $(1 - \epsilon)$ as a function of $(p, \cos\theta)$, for data and simulation.	77
6.9	Projected upstream inefficiencies as functions of p and of $\cos\theta$, for data and simulation.	79
6.10	Difference between data and MC inefficiencies, as a function of $(p, \cos\theta)$	80
6.11	Projected differences between data and MC inefficiencies, as functions of p and $\cos\theta$	81
6.12	Difference between downstream and upstream reconstruction inefficiencies, for data and simulation.	82
6.13	Projected differences between downstream and upstream reconstruction inefficiencies as functions of p and $\cos\theta$, for data and simulation.	83
6.14	PC Time-of-Flight histograms for data and simulation.	85

6.15	PC Time-of-Flight histograms for data and simulation, with background removed.	86
6.16	Leading edge of TDC signal time distributions for data and simulation. . . .	87
7.1	Fitted values of $\Delta\rho$, $\Delta\delta$, and $\Delta P_\mu\xi$ for each data set.	93
8.1	STRs generated by Garfield software.	98
8.2	STRs measured from data, and from Monte Carlo.	99
8.3	Differences between data- or Monte Carlo-driven STRs and Garfield STRs. .	100
8.4	Hit occupancies within a drift cell, for data and MC.	100
8.5	Effects of analyzing data and MC using measured STRs, vs using Garfield STRs.	101
8.6	Number of degrees of freedom available to the helix fitter, with standard geometry and with a 500 μm foil bulge.	106
8.7	Average t_0 values for each wire in a DC plane, from MC.	109
8.8	Average t_0 values for each DC plane, as measured from MC.	110
8.9	Average differences in end-of-run and start-of-run t_0 values for each wire in a DC plane.	111
8.10	Average differences in end-of-run and start-of-run t_0 values for each DC plane.	112
8.11	Measured delta rate $R_\delta = N_{12}/N_{11}$ for different delta rate boost factors. The vertical axis is in units of 10^{-2}	120
8.12	Measured hard energy loss rate $R_B = N_h/N_{tot}$ for different bremsstrahlung rate boost factors.	122
8.13	Residuals of Michel parameter fits for increased delta ray production, against standard surface muon simulation.	124
8.14	Residuals of Michel parameter fits for increased bremsstrahlung production, against standard surface muon simulation.	125
8.15	Normalized backscatter counts for simulation with an aluminum plate added downstream.	128
8.16	Angle resolution vs $1/\sin\theta$, for four 5 MeV/c momentum bins, for data (red circles) and simulation (blue triangles). Vertical error bars are included but are very small.	130
8.17	Momentum resolution vs $1/\sin\theta$, for four 5 MeV/c momentum bins, for data (red circles) and simulation (blue triangles). Vertical error bars are included but are very small.	131
8.18	Data-MC quadrature differences in angle resolution, vs $1/\sin\theta$, for four momentum bins.	132
8.19	Data-MC quadrature differences in momentum resolution, vs $1/\sin\theta$, for four momentum bins.	132
8.20	Slopes of $\Delta\sigma_{Q\theta}$ vs $1/\sin\theta$, from Fig. 8.18, with constant fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Q\theta}$ values, and are underestimates.	133
8.21	Slopes of $\Delta\sigma_{Qp}$ vs $1/\sin\theta$, from Fig. 8.19, with constant fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Qp}$ values, and are underestimates.	133
8.22	Offsets of $\Delta\sigma_{Q\theta}$ vs $1/\sin\theta$, from Fig. 8.18, with fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Q\theta}$ values, and are underestimates. .	134
8.23	Slopes of $\Delta\sigma_{Qp}$ vs $1/\sin\theta$, from Fig. 8.19, with constant fit.	134
8.24	Momentum edge measured by comparative energy calibration, for momentum smearing.	137
8.25	Changes in Michel parameters vs endpoint resolution, for several momentum smearings.	138
8.26	Event type distributions for standard and high muon rate simulations. . . .	147
8.27	Event type distributions for standard and high beam positron rate simulations.	148
8.28	Reconstructed radiative correction spectrum.	150

9.1	Measured differences in ρ and δ between data and simulation, for each data set.	154
9.2	Measured differences in $P_\mu\xi$ between data and simulation, for each data set. .	155
A.1	Window type distributions.	165
B.1	Event type distributions.	167

List of Symbols

Names and Acronyms

- BL1A** A primary proton beamline at TRIUMF (Sec. 4.1)
- CERN** The European Organization for Nuclear Research, located near Geneva on the border between Switzerland and France.
- CP TOF** “CP Time-of-Flight”; the time between the arrival of a particle at the *TWIST* trigger and the next Capacitive Probe pulse from BL1A (Sec. 4.1.2)
- DC** Drift Chamber, a tracking chamber where the drift time can be used to precisely determine the location of a particle’s passage (Sec. 4.2.3)
- EPICS** Experimental Physics and Industrial Control System, for operating the M13 beamline at TRIUMF (Sec. 4.1.2)
- Geant** Simulation package produced by CERN (Sec. 6.2)
- M13** Secondary beamline at TRIUMF (Sec. 4.1.2)
- MC** Monte Carlo simulation, usually referring to *TWIST*’s Geant implementation (Sec. 6.2)
- Mofia** *TWIST* primary event reconstruction software (Sec. 5.3)
- MWPC** see PC
- NMR** Nuclear Magnetic Resonance; *TWIST* uses NMR probes to measure magnetic fields
- PACT** Pulse Amplitude Charge-to-Time converter, used to measure the amount of energy deposited by a passing particle in e.g. a scintillator
- PC** Multi-Wire Proportional Chamber, a fast tracking chamber whose signal size is proportional to the amount of energy deposited by a passing particle (Sec. 4.1.2)
- PSI** Paul Scherrer Institute, a laboratory in northern Switzerland
- ROOT** Data analysis package produced by CERN, used by *TWIST* for data analysis (Sec. 5.4)
- STR** Space-Time Relation, the table of electron drift times as a function of position within a drift cell (Sec. 5.2)
- T1ION** Pulser whose rate is proportional to BL1A proton beam current (Sec. 4.1)
- TDC** Time to Digital Converter, essentially an electronic clock
- TRIUMF** TRI-University Meson Facility; a cyclotron facility in Vancouver, B.C./ and the home of the *TWIST* experiment
- TWIST** TRIUMF Weak Interaction Symmetry Test—TRIUMF experiment E614

Symbols

The Michel Parameters

- ρ Governs the isotropic shape of the momentum distribution
- δ Governs the momentum dependence of the $\cos\theta$ distribution asymmetry
- ξ Governs the level of asymmetry of the $\cos\theta$ distribution (measured as $P_\mu\xi$)
- η Contributes to the isotropic shape of the momentum distribution; not measured directly by *TWIST*.

Coordinates

The coordinate systems are described in more detail in Sec. 4.2.5.

- (x, y, z) Standard laboratory coordinate system: z is approximately along the beam direction, y is up, and x completes the right-handed coordinate system.
- (u, v, z) The main *TWIST* coordinate system: z is along the beamline axis, and u and v are rotated 45° from x and y , corresponding to the orientation of the *TWIST* DCs and PCs.
- (θ, ϕ) *TWIST* polar coordinates: θ is the angle between the positive z axis and particle momentum vector, and ϕ is the azimuthal angle measured from either the positive x axis (ϕ_{xy}) or the positive u axis (ϕ_{uv}); usually $\phi = \phi_{uv}$.
- θ_s The angle between the decay positron's momentum vector and the muon's spin vector, used in theoretical calculations (Chap. 2).

Kinematic Variables

The kinematic variables used to describe positron tracks include:

- \vec{p} Total momentum (with $p \equiv |\vec{p}|$)
- p_z The longitudinal component of the momentum; $p_z = p \cos\theta$
- p_t The transverse component of the momentum; $p_t^2 = p^2 - p_z^2$
- E Particle total energy (in lab frame)
- m Particle rest mass
- P_μ Muon beam polarization
- x The positron's "reduced energy": $x = E/E_{max}$, where E is the energy of the decay positron/electron, and E_{max} is the maximum kinematically permitted energy (the decay cutoff)
- x_0 The kinematic minimum value of x (above): $x_0 = m_e/E_{max}$, which corresponds to a positron at rest.
- v Particle speed (in lab frame)
- β Particle speed with respect to c ($\beta = v/c$)

Energy Calibration Parameters

See Sec. 5.5.1 for details.

q_{US}, q_{DS} Slope of linear fit to relative endpoint position vs $1/\cos\theta$, for upstream or downstream calibration.

s_{US}, s_{DS} Intercept of linear fit to relative endpoint position vs $1/\cos\theta$, for upstream or downstream calibration.

Constants

α Fine structure constant, $\sim 1/137$

g_{ij}^κ Generic weak coupling constant (Sec. 2.3)

Chapter 1

Introduction

Wider expanses and greater depths are now exposed to the searching eye of knowledge, regions of which we had not even a presentiment. It has brought us much nearer to grasping the plan that underlies all physical happening.

- Hermann Weyl, Space, Time, Matter

Most laws of nature are preserved under a parity transformation—the reversal of all three spatial axes, the effect one would see by observing a system in a mirror. In 1956, Lee and Yang published a proposal [1] that, although parity was known to be conserved under electromagnetic and strong interactions, this symmetry might be violated under the weak interaction. (This proposal won them the Nobel Prize in Physics.) This was confirmed within a year by Wu *et al.* [2] in a measurement of the asymmetry of beta decay in polarized ^{60}Co . That experiment suggested that the violation of parity was very large.

Another sector proposed by Lee and Yang in which to search for parity violation was in the decays of pions and muons: if parity is not conserved, the muons produced by the decay of stopped pions would be polarized, and further the positrons from the muon decay would be emitted asymmetrically with respect to the direction of muon spin. This was confirmed by Garwin *et al.* [3], in a paper immediately following ref. [2] in the same journal.

There are now known to be two varieties of weak interaction: *charged-current* and *neutral-current*. The charged-current weak interaction proceeds through an intermediate state involving a charged particle, called W^+ or W^- , while the intermediate state of the neutral-current weak interaction involves a neutral Z^0 particle. The muon decays through the charged-current interaction, in which the charge of the muon is carried through to the electron.

All evidence from muon decay, and indeed from all studies involving the charged-current weak interaction, are in favour of *maximal* parity violation—specifically, that only left-handed particles and right-handed anti-particles are affected by the charged-current weak interaction, leaving right-handed particles and left-handed anti-particles entirely unaffected. (The neutral current weak interaction is known to violate parity as well, but not maximally.)

This is the form in which the weak interaction has been incorporated into the Standard Model. Parity violation remains unique to the weak interaction. If there are “right-handed” contributions to the charged-current weak interaction, they must be very small at the energy scales which have been studied.

Muon decay is a particularly good system in which to study the charged-current weak interaction to high precision, since it is a purely leptonic process. As such, only the weak and electromagnetic forces participate directly, and these processes can be calculated unambiguously from the theory. The strong interaction only appears in higher order radiative corrections;¹ its contribution is negligible compared to current theoretical and experimental limits.

The muon was first discovered in cosmic rays in 1937, in a cloud chamber experiment by Street and Stevenson [4], and was found to decay into three light particles rather than two from measurements of the decay electron’s energy spectrum in the late 1940’s [5, 6].

The distribution of momentum and direction of the positrons from muon decay is precisely predicted by the Standard Model. More fundamentally, the Standard Model includes only a single specific form of the weak interaction: the interaction of only left-handed particles, mediated by a vector (spin 1) boson. However, the most general description consistent with Lorentz invariance and locality includes a wide variety of interaction forms, and the Standard Model excludes these only by construction, for a lack of experimental evidence supporting them. Careful study of a weak process such as muon decay can strengthen the Standard Model’s exclusion of these other interactions, or reveal “new physics” that must be included after all.

Early calculations of the shape of the muon decay spectrum were performed by Michel [7] in 1950, and by Kinoshita and Sirlin [8] in 1957 with parity violation included. The spectrum can be described by a number of shape parameters, often called the Michel parameters, including: ρ , which describes the angle-integrated momentum dependence; ξ , which describes the momentum-integrated asymmetry—the difference between the decay rate in the direction of muon spin and the rate in the opposite direction; and δ , which describes the way in which the spectrum asymmetry depends on momentum. There are additional shape parameters, as well. The values of these parameters depend on the actual form of the weak interaction. Prior to the *TWIST* experiment, the parameter ρ has not been successfully measured since the 1960’s, and δ and ξ since the 1980’s—see Chap. 3 for more on these and related muon decay experiments.

The *TWIST* experiment is designed to study the shape of the muon decay spectrum to high precision, particularly through the measurement of the three shape parameters

¹In addition to corrections to muon decay, the strong force is involved more directly in pion decay, since pions are composed of quarks. This influences the muon polarization, but the effect is negligibly small for *TWIST*.

mentioned above. Because of its large acceptance in energy and angle, the experiment is able to derive measurements for all three parameters from the same data.

The present work is a new direct measurement of the ρ and δ parameters. Chap. 2 presents the theory describing the decay of the muon, in terms of kinematics and of fundamental interactions, and discusses some proposed extensions to the Standard Model that could affect the shape of the decay spectrum. Chap. 3 describes previous experiments to measure the muon decay parameters, as well as a global analysis which combined all available muon decay data including measurements of ρ and δ from the *TWIST* experiment prior to this work. Chap. 4 explains the apparatus of the *TWIST* experiment, and Chap. 5 explains the analysis techniques. Chap. 6 discusses the simulations used by *TWIST* for comparison to data. Chap. 7 describes the various conditions under which the data were acquired, and presents some measures of data quality. Chap. 8 explores the systematic uncertainties of the experiment. Finally, Chap. 9 presents the new measurements of the ρ and δ parameters, and their implications to some of the proposed extensions to the Standard Model of the weak interaction.

The author's primary work included operating the experiment during the taking of data (Chap. 7) including periods as Run Coordinator, testing and validating the simulation (Sec. 6.3), and the determination of systematic errors and corrections (Chap. 8). The author also made significant contributions to the construction of the detector (Chap. 4), and to the development of the simulation (Chap. 6) and analysis (Chap. 5) software.

Chapter 2

Muon Decay and the Weak Interaction

It is not necessary to assume the sensory information corresponds to an objective reality only that the studying of it is interesting or useful.

—Byron K. Jennings, “On the Nature of Science”

2.1 The Weak Interaction and the Standard Model

The Standard Model of Particle Physics is a very successful description of the matter and forces. Many aspects of the Standard Model—types and strengths of interactions, values of particle masses, etc.—are included simply because they are observed to be so, and experiments are continually testing these aspects to search for “new physics” not included in the model. *TWIST* is one such experiment. A brief overview of the Standard Model is presented here, with an emphasis on the properties of the weak interaction. Note that many of the statements presented here are modified by the various proposed extensions of the Standard Model, such as those discussed in section 2.5. For more information, see e.g. Ref. [9].

Forces, or interactions, are mediated by particles called *bosons*: the strong force by gluons, the electromagnetic force by photons, and the weak force by the charged W^+ and W^- bosons and the neutral Z boson. Matter is composed of particles called *fermions*. There are two basic types of fermions: *quarks*, which can interact with the strong force; and *leptons*, which cannot. The fermions of the Standard Model are listed in Tab. 2.1. Each type of fermion can be organized into three generations; particle masses increase with each generation. Each generation of quarks includes an “up-type” quark, with charge $+2/3e$, and a “down-type” quark, with charge $-1/3e$, where e is the magnitude of the charge of the electron. Each generation of leptons includes a charged lepton (charge $-e$) and an uncharged neutrino. The variety of each quark (up, down, strange, etc.) is called a “flavour”; each

Quarks		
up (u)	charm (c)	top (t)
down (d)	strange (s)	bottom (b)
Leptons		
electron (e)	muon (μ)	tau (τ)
e neutrino (ν_e)	μ neutrino (ν_μ)	τ neutrino (ν_τ)

Table 2.1: The fermions of the Standard Model.

generation of leptons (electron, muon, tau) is called a flavour. An antiparticle is identical to its corresponding particle but with the sign of the charge and all other quantum numbers reversed. Charged antileptons are usually denoted by their sign, e.g. the e^+ , or positron, is the antiparticle of the e^- ; other antiparticles are denoted by a bar, e.g. $\bar{\nu}_e$.

The charged weak interaction is able to convert charged leptons into neutral leptons and vice versa, or up-type quarks into down-type quarks and vice versa; it is the only¹ interaction known to do this. Leptons are only converted within a single generation— $e^- \leftrightarrow \nu_e$, etc.—but for quarks this conversion can cross generations. The strengths of the charged weak interactions between different types of quarks are represented as elements of the Cabbibo-Kobayashi-Maskawa (CKM) matrix,

$$V_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (2.1)$$

In the Standard Model these values are free parameters, subject to the constraint that the CKM matrix must be unitary—meaning, among other things, that there are no more than three generations of quarks—and so they must be measured. The values of the CKM matrix are often represented using sines and cosines of “mixing angles.”

Each particle has an intrinsic *spin*. For the particles considered here, the fermions have spin 1/2, and the bosons have spin 1. The spin of the boson determines the “type” of interaction it mediates: a spin 0 boson would give a “scalar” interaction, a spin 1 boson gives a “vector” interaction, and a spin 2 boson would give a “tensor” interaction.

Each particle also has an intrinsic *chirality*, left-handed or right-handed. Chirality is essentially the particle’s “weak charge,” in that left-handed particles are affected by the weak interaction and right-handed particles are not; the opposite is true for anti-particles. Note that in the non-Standard-Model case that the weak interaction *does* couple to right-handed particles at some undetected level, it is possible that the right-handed quarks have their own CKM matrix.

Another way to say this is that the weak interaction violates both *charge conjugation* and *parity* symmetries, which are conserved by both the strong and electromagnetic interactions. Parity (P) symmetry states that physical laws remain the same under the reversal of all

¹Neutrinos will spontaneously oscillate between flavours, but this is not really an “interaction” as no bosons are involved. Moreover, neutrino oscillation is not included in the Standard Model.

three physical axes, the effect one would see by observing a system in a mirror; this would reverse a particle’s chirality, which changes its coupling to the weak interaction. Charge conjugation (C) symmetry states that the same physical laws apply to both matter and antimatter; the weak interaction violates this by interacting only with *left*-handed particles and *right*-handed antiparticles. However, this says that the weak interaction conserves “ CP ” symmetry: the interaction is the same if both C and P are reversed simultaneously. Rare examples of the weak interaction violating even CP symmetry exist, such as in the decay of a neutral kaon [10, 11], but CP violation has not been seen in muon decay, or in lepton interactions in general.

2.2 Kinematics

2.2.1 Kinematics of Pion Decay

Pions are light “unflavoured” mesons.² The positive pion, whose decay produces muons for *TWIST*, is composed of the quarks $u\bar{d}$ and has a lifetime of 26.0 ns. Its dominant decay mode, with a branching ratio of more than 99.99%, is $\pi^+ \rightarrow \mu^+ \nu_\mu$. Since this is a two-body process, in the centre-of-mass reference frame the muon from pion decay is produced at a single momentum p_μ given by

$$p_\mu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} = 29.7920 \text{ MeV}/c \quad (2.2)$$

ignoring the neutrino mass. (Recent observations of the cosmic microwave background from the Wilkinson Microwave Anisotropy Probe (WMAP) experiment find a limit on the *sum* of the masses of all three neutrino types to be $< 2.11 \text{ eV}/c^2$ [12].) As a consequence of the weak interaction’s maximal parity violation in the Standard Model, neutrinos produced in weak processes (such as pion decay) are always of left-handed chirality, and hence (to the extent that the neutrino is massless) of left-handed helicity—the polarization is fixed to be opposite the momentum. Since pion decay is a two-body process, the neutrino and the muon have opposite momenta (in the pion rest frame). The pion is spinless, and the neutrino and muon are each spin 1/2, so conservation of angular momentum requires the muon to have spin opposite that of the neutrino, and hence the same helicity. The neutrino polarization is reduced by a factor of p/E ; since $p_\nu = p_\mu = 29.7920 \text{ MeV}/c$, the more conservative limit of $m_\nu < 0.19 \text{ MeV}/c^2$ gives $E_\nu < 29.7926 \text{ MeV}$, and $1 - p/E \lesssim 2 \times 10^{-5}$, which is negligible for the purposes of the *TWIST* experiment. Hence, muons created from the decay of pions at rest are essentially 100% negatively polarized.³

²A meson is a composite particle composed of two quarks. “Unflavoured” is a historical term denoting particles composed of only up and down quarks.

³If the Standard Model weak interaction is not assumed, however, then right-handed neutrinos can be produced in pion decay (through the same mechanisms which would affect the muon decay spectrum), which then have right-handed helicity and force the muon to have the same. This would reduce the polarization of the muon accordingly.

The radiative decay modes such as $\pi^+ \rightarrow \mu^+ \nu_\mu \gamma$ also contribute muons, but Pierre Depommier has shown in a series of *TWIST* technical notes [13] that their effect on the muon polarization is negligible.

No other pion decay modes have been observed to produce muons.

2.2.2 Kinematics of Muon Decay

The muon is a lepton, with a mass of $105.7 \text{ MeV}/c^2$ and a lifetime of $2.20 \text{ } \mu\text{s}$. The positive muon's dominant decay mode, with a branching ratio of nearly 98.6%, is $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. The positron from this decay can have a range of energies, since muon decay is a three-body process. The energy of the positron is maximum when both neutrinos are emitted in the direction opposite the positron direction. Using conservation of momentum and energy, it can be shown that the positron's energy in this case, E_{max} , is given by

$$E_{max} = \frac{m_e^2 + m_\mu^2}{2m_\mu} = 52.83 \text{ MeV} \quad (2.3)$$

again ignoring neutrino masses.

Other decay modes of the positive muon are the inner bremsstrahlung decay, $\mu^+ \rightarrow e^+ \gamma \nu_e \bar{\nu}_\mu$, with a branching ratio of $(1.4 \pm 0.4) \times 10^{-2}$, and $\mu^+ \rightarrow e^+ e^+ e^- \nu_e \bar{\nu}_\mu$, with a branching ratio of $(3.4 \pm 0.4) \times 10^{-5}$ [14]. If the secondary γ and e^\pm particles are not detected, these can be treated as radiative corrections.

It is useful to note that $p_e \approx E_e$ for the positron energies being considered. Even at the lower end of the energy spectrum, say $E_e = 15 \text{ MeV}$, we have $p_e = \sqrt{E_e^2 - m_e^2} = 14.99 \text{ MeV}/c$, and the approximation improves as energy increases.

2.3 General Form of Muon Decay

The most general local, Lorentz-invariant⁴ four-fermion description of muon decay is one which allows any type of weak coupling (scalar (S), vector (V), or tensor (T)) between any combination of left-handed or right-handed fermions—defined in this context as fermions whose wavefunctions are eigenfunctions of the $1 + \gamma^5$ and $1 - \gamma^5$ operators, respectively. The matrix element M describing this general interaction can be written as

$$M = \frac{4G_F}{\sqrt{2}} \sum_{\substack{i=L,R \\ j=L,R \\ \kappa=S,V,T}} g_{ij}^\kappa \langle \bar{\psi}_e | \Gamma^\kappa | \psi_{\nu_e} \rangle \langle \bar{\psi}_{\nu_\mu} | \Gamma_\kappa | \psi_{\mu_j} \rangle. \quad (2.4)$$

Γ^κ are the interaction matrices, given by combinations of the Dirac γ matrices:

$$\Gamma^S = 1, \quad \Gamma^V = \gamma^\mu, \quad \Gamma^T = \frac{1}{\sqrt{2}} \sigma^{\mu\nu} \equiv \frac{i}{2\sqrt{2}} (\gamma^\mu \gamma^\nu - \gamma^\nu \gamma^\mu). \quad (2.5)$$

⁴Lorentz invariance is the core symmetry of special relativity, stating that physical processes remain the same regardless of relative motion between the system and the observer.

The decay distribution in energy and angle is found from the matrix element by

$$\frac{d^2\Gamma}{dx d(\cos\theta_s)} \propto |M|^2. \quad (2.6)$$

where Γ is the decay probability, $x = E_e/E_{max}$ is the positron's reduced energy, and θ_s is the angle between the positron's momentum vector and the muon's spin vector. (The subscript s is to distinguish this angle from the angle θ between the positron momentum vector and the z axis, used for *TWIST*'s experimental measurements. See Sec. 4.2.5 for more on the *TWIST* coordinate system.) Note that, since the neutrinos are not observed, the above integrates over their degrees of freedom.

The coupling constants g_{ij}^κ are normalized [15] such that

$$\begin{aligned} & \frac{1}{4}(|g_{RR}^S|^2 + |g_{LR}^S|^2 + |g_{RL}^S|^2 + |g_{LL}^S|^2) \\ & + |g_{RR}^V|^2 + |g_{LR}^V|^2 + |g_{RL}^V|^2 + |g_{LL}^V|^2 \\ & + 3(|g_{LR}^T|^2 + |g_{RL}^T|^2) \\ & = 1 \end{aligned} \quad (2.7)$$

With this normalization, the probability that a j -handed muon will decay into an i -handed positron through a weak interaction of type κ is given by $n |g_{ij}^\kappa|^2$, where $n = 1/4, 1$, or 3 for $\kappa = S, V$, or T . We can then define Q_{ij} to represent the total probability for j -handed muons to decay into i -handed electrons as

$$Q_{RR} = \frac{1}{4} |g_{RR}^S|^2 + |g_{RR}^V|^2, \quad (2.8a)$$

$$Q_{LR} = \frac{1}{4} |g_{LR}^S|^2 + |g_{LR}^V|^2 + 3 |g_{LR}^T|^2, \quad (2.8b)$$

$$Q_{RL} = \frac{1}{4} |g_{RL}^S|^2 + |g_{RL}^V|^2 + 3 |g_{RL}^T|^2, \quad (2.8c)$$

$$Q_{LL} = \frac{1}{4} |g_{LL}^S|^2 + |g_{LL}^V|^2. \quad (2.8d)$$

“Pseudoscalar” and “axial vector” interactions are used as well in formalisms based in terms of parity eigenstates, instead of the weak flavour eigenstates used here; such interactions are covered in this formalism by combinations of interactions and handedness. The “ $V - A$ ” interaction of the Standard Model corresponds to the vector coupling between a left-handed muon and a left-handed positron.

The coupling constants g_{ij}^κ must be measured. The Standard Model states that $g_{LL}^V = 1$ and all the other coupling constants are zero; this is done by fiat, because there is no evidence for any other type of weak interaction. (The exceptions are g_{LL}^T and g_{RR}^T , which are set to zero because their terms in the matrix element cancel out ($\langle \bar{\psi}_{e_R} | \sigma^{\mu\nu} | \psi_{\nu_e} \rangle \langle \bar{\psi}_{\nu_\mu} | \sigma_{\mu\nu} | \psi_{\mu_R} \rangle = 0$).)

See Tab. 3.2 on p. 15 for experimental limits on the values of the weak coupling constants.




Figure 2.1: Michel Spectrum of muon decay (Eq. (2.9)). Electron mass is included in the plot, but no radiative corrections. $\cos \theta_s = 1$ is in the direction of muon spin.

The strong interaction only appears in higher order radiative corrections, with a fractional contribution on the order of $0.07(\alpha/\pi)^2 \sim 4 \times 10^{-7}$ [16] (where α is the fine structure constant). These contributions are unimportant to the *TWIST* measurement.

2.4 The Parameterized Muon Decay Spectrum

The decay distribution described by Eqs. (2.4) and (2.6), as well as the decay kinematics, can be encapsulated using the parameterization of Michel [7] (this form follows that found in [17]) for the decay of a positive muon:

$$\begin{aligned} \frac{d^2\Gamma}{dx d(\cos \theta_s)} = & \frac{m_\mu}{4\pi^3} E_{max}^4 G_F^2 \sqrt{x^2 - x_0^2} \left(x(1-x) + \rho(4x^2 - 3x - x_0^2) + \eta x_0(1-x) \right. \\ & \left. + P_\mu \xi \frac{1}{3} \sqrt{x^2 - x_0^2} \left[1 - x + \frac{2}{3} \delta \left(4x - 3 + \left(\sqrt{1 - x_0^2} - 1 \right) \right) \right] \cos \theta_s \right) \end{aligned} \quad (2.9)$$

where $G_F \approx 1.17 \times 10^{-11} \text{ MeV}^{-2}$ is the Fermi coupling constant, θ_s is the angle between the muon spin and the positron momentum, $x = E_e/E_{max}$ is the positron's reduced energy, $x_0 = m_e/E_{max}$ is the minimum possible value of x , corresponding to a positron at rest, $P_\mu = |\vec{P}_\mu|$ is the degree of muon polarization, and ρ , δ , ξ , and η are the **Michel parameters**. (There are a number of additional Michel parameters which govern the polarization of the positrons, which *TWIST* does not measure.) This decay spectrum is plotted in Fig. 2.1.

Electromagnetic radiative corrections to the muon decay spectrum are significant, changing the spectrum by as much as 10% or more near $x = 1$ (see Fig. 2.2). Full $\mathcal{O}(\alpha)$ radiative




Figure 2.2: Effect of the radiative corrections on the Michel spectrum shape. Broken lines show the decay positron energy distributions at three values of $\cos\theta_s$, without radiative corrections; the solid lines show the distributions with radiative corrections included.

corrections have been calculated [18]. Additionally, $\mathcal{O}(\alpha^2 L^2)$, $\mathcal{O}(\alpha^2 L)$, and $\mathcal{O}(\alpha^2)$ radiative correction terms have been calculated [19–21], where $L = \log(m_\mu^2/m_e^2) \approx 10.66$ is a factor which significantly enhances some of the $\mathcal{O}(\alpha^2)$ terms; $\mathcal{O}(\alpha^2 L^2)$ and $\mathcal{O}(\alpha^2 L)$ are often referred to as the “leading logarithm” and “next to leading logarithm” terms. The $\mathcal{O}(\alpha^2)$ correction influences the spectrum at the level of $\sim 0.5 \times 10^{-4}$. The work of *TWIST* has been a significant motivator in the calculation of these radiative corrections. As demonstrated in Sec. 8.7.1, the precision of the available radiative corrections is sufficient not only for the current measurement but for the ultimate *TWIST* goals.

Note that radiative corrections must assume an underlying weak interaction model. All of the above calculations were performed assuming Standard Model values for the weak coupling constants, i.e. only left-handed particles interact, through a vector coupling. (In the implementation of this, however, *TWIST* multiplies the anisotropic radiative corrections by ξ , so that the whole spectrum depends on $P_\mu \xi$.) If deviations from the Standard Model values of the weak couplings are discovered, the radiative corrections would need to be re-examined.

The Michel parameters are given in terms of the general weak coupling constants by [22]:

$$\rho = \frac{3}{4} - \frac{3}{4} \left[|g_{LR}^V|^2 + |g_{RL}^V|^2 + 2|g_{LR}^T|^2 + 2|g_{RL}^T|^2 + \Re(g_{RL}^S g_{RL}^{T*} + g_{LR}^S g_{LR}^{T*}) \right], \quad (2.10a)$$

$$\xi\delta = \frac{3}{4} - \frac{3}{4} \left[|g_{LR}^V|^2 + |g_{RL}^V|^2 + 4|g_{LR}^T|^2 + 2|g_{RL}^T|^2 + 2|g_{RR}^V|^2 + \frac{1}{2}|g_{RR}^S|^2 + \frac{1}{2}|g_{LR}^S|^2 + \Re(g_{RL}^S g_{RL}^{T*} - g_{LR}^S g_{LR}^{T*}) \right], \quad (2.10b)$$

$$\xi = 1 - \left[\frac{1}{2}|g_{RR}^S|^2 + \frac{1}{2}|g_{LR}^S|^2 + 2|g_{RR}^V|^2 + 4|g_{RL}^V|^2 - 2|g_{LR}^V|^2 - 2|g_{LR}^T|^2 + 8|g_{RL}^T|^2 + \Re(g_{RL}^S g_{RL}^{T*} - g_{LR}^S g_{LR}^{T*}) \right], \quad (2.10c)$$

$$\eta = \frac{1}{2} \Re \left[g_{LL}^V g_{RR}^{S*} + g_{RL}^V (g_{LR}^{S*} + 6g_{LR}^{T*}) + g_{LR}^V (g_{RL}^{S*} + 6g_{RL}^{T*}) + g_{RR}^V g_{LL}^{S*} \right]. \quad (2.10d)$$

As mentioned, the Standard Model sets all weak coupling constants to zero except g_{LL}^V , so the Michel parameters take on the values $\rho = 3/4$, $\eta = 0$, $\delta = 3/4$, and $\xi = 1$. Measurements of the Michel parameters can be used to derive limits on the weak coupling constants, using Eqs. (2.10) in combination with other decay measurements such as inverse muon decay ($e^- \nu_\mu \rightarrow \mu^- \nu_e$) or the decay positron polarization; see Sec. 3.5 for a recent example of such an analysis.

2.4.1 Asymmetry at the Spectrum Endpoint

It is useful to note that the combination $P_\mu \xi \delta / \rho$ represents the asymmetry of the decay spectrum at the endpoint ($x = 1$). This can be seen from the definition of asymmetry, using the formula for the decay spectrum, assuming $x_0 = 0$ for this illustration. Let $N(\cos \theta)$ represent the shape of the spectrum; substituting $x = 1$ and $x_0 = 0$ into Eq. (2.9), we have

$$N(\cos \theta) \propto \rho + P_\mu \xi \delta \cos \theta. \quad (2.11)$$

The asymmetry A at the endpoint can be determined by

$$A = \frac{N(1) - N(-1)}{N(1) + N(-1)} \quad (2.12)$$

$$\begin{aligned} &= \frac{(\rho + P_\mu \xi \delta) - (\rho - P_\mu \xi \delta)}{(\rho + P_\mu \xi \delta) + (\rho - P_\mu \xi \delta)} \\ &= \frac{P_\mu \xi \delta}{\rho} \end{aligned} \quad (2.13)$$

Note that, by definition, $N(\cos\theta) \geq 0$. At $\cos\theta = -1$, then, this leads to the constraint that $\rho - P_\mu \xi \delta \geq 0$ or $P_\mu \xi \delta / \rho \leq 1$.

2.5 Extensions to the Standard Model

There are a number of extensions to the Standard Model which predict non-standard values of the Michel parameters, which can be constrained (or detected!) by the *TWIST* measurements. A selection of these extensions is described here.

2.5.1 Right-Handed Muon Decay

Some proposed extensions result in the weak decay of the right-handed muon, forbidden under the Standard Model. To search for this, we can use Eqs. (2.8) to define

$$\begin{aligned} Q_R^\mu &= Q_{RR} + Q_{LR} \\ &= \frac{1}{4} |g_{LR}^S|^2 + \frac{1}{4} |g_{RR}^S|^2 + |g_{LR}^V|^2 + |g_{RR}^V|^2 + 3 |g_{LR}^T|^2 \end{aligned} \quad (2.14)$$

which represents the total probability of a right-handed muon decaying into any type of electron, a process forbidden in the Standard Model weak interaction. Q_R^μ can be written in terms of the Michel parameters [15]:

$$\begin{aligned} Q_R^\mu &= Q_{RR} + Q_{LR} = \frac{1}{2} \left(1 + \frac{1}{3} \xi - \frac{16}{9} \xi \delta \right) \\ &= \frac{1}{2} \left(1 + \frac{1}{3} \left(\frac{\xi \delta}{\rho} \right) \frac{\rho}{\delta} - \frac{16}{9} \left(\frac{\xi \delta}{\rho} \right) \rho \right). \end{aligned} \quad (2.15)$$

The 90% confidence limit on this quantity prior to the *TWIST* experiment was $Q_R^\mu < 0.014$, from measurements of ρ , δ , and the combination $P_\mu \xi \delta / \rho$ (which gave the relevant coupling limits in Tab. 3.2); the 2005 muon decay global analysis [23] gives $Q_R^\mu < 0.007$, with the improvement due mainly to the *TWIST* measurements of ρ [24] and δ [25] published in 2005.

2.5.2 Left-Right Symmetric Models

The Standard Model charged weak interaction includes only so-called $V - A$ currents, which is to say that only left-handed particles interact weakly through the charged current. A reasonable extension is to assume that right-handed particles *do* have a charged current interaction, through a $V + A$ current which is suppressed but nonzero. In this model, the Standard Model $SU(2) \times U(1)$ electroweak gauge group is extended to $SU(2)_L \times SU(2)_R \times U(1)$ [26]. The charged bosons in this group can mix:

$$W_L = W_1 \cos \zeta + W_2 \sin \zeta \quad (2.16)$$

$$W_R = e^{i\omega} (-W_1 \sin \zeta + W_2 \cos \zeta) \quad (2.17)$$

where W_1 and W_2 are the mass eigenstates (which have masses m_1 and m_2), W_L and W_R are the weak interaction eigenstates, ζ is the mixing angle, and ω is a CP -violating phase. In general, the right-handed and left-handed interactions are governed by separate coupling constants, g_R and g_L , corresponding to g_{RR}^V and g_{LL}^V in Eq. (2.4).

The presence of the $V + A$ interaction in the charged current would modify the decay spectrum. Following the notation of Herczeg [26], defining

$$t \equiv \frac{g_R^2 m_1^2}{g_L^2 m_2^2}, \quad (2.18)$$

$$t_\theta \equiv \frac{g_R^2 m_1^2 \cos \theta_1^R}{g_L^2 m_2^2 \cos \theta_1^L}, \quad (2.19)$$

where θ_1^R and θ_1^L are mixing angles from the left- and right-handed CKM matrices, and

$$\zeta_g \equiv \frac{g_R}{g_L} \zeta, \quad (2.20)$$

we can write the changes to the muon decay parameters as

$$\rho \simeq \frac{3}{4}(1 - 2\zeta_g^2), \quad (2.21)$$

$$\xi \simeq 1 - 2(t^2 + \zeta_g^2), \quad (2.22)$$

$$\frac{\xi\delta}{\rho} \simeq 1 - 2t^2, \quad (2.23)$$

$$\text{and } P_\mu \simeq 1 - 2t_\theta^2 - 2\zeta_g^2 - 4t_\theta\zeta_g \cos(\alpha + \omega) \quad (2.24)$$

where α is a CP -violating phase from the right-handed CKM matrix. To a good approximation, δ and η on their own are unaffected by the presence of a $V + A$ current.

Measurements of the decay parameters hence give limits on the generalized mass ratio t and mixing angle ζ_g . In particular, the limit on ρ gives a mass-independent limit on the mixing angle. Noting that, from Eq. (2.21), $\rho > 0.75$ is forbidden in left-right symmetric models, the measurements of ρ in Tab. 3.1 can be converted into 90% confidence limit lower bounds: the pre-*TWIST* value becomes $\rho > 0.7467$, and the published *TWIST* measurement becomes $\rho > 0.7487$. Using Eqs. (2.21), these give limits on the mixing angle ζ_g of $|\zeta_g| < 0.047$ for the pre-*TWIST* value, and $|\zeta_g| < 0.030$ for the published *TWIST* value.

Chapter 3

Review of Previous Measurements

Banks and banks of humming machinery. I've never seen so many knobs!
—Propellerheads, “Bang On”

This chapter presents the most recent measurements of selected Michel parameters, aside from the *TWIST* experiment’s results. Only “model-independent” results—that is, results which do not assume specific values of some decay parameters or g_{ij}^{κ} weak coupling constants in order to obtain values for others—are included. Furthermore, this review includes only results relevant to the present work; in addition to the ρ and δ measurements, η and the combination $P_{\mu}\xi\delta/\rho$ are important for constraining correlations in a global analysis.

Each experiment will be described briefly, its measurements stated, and its main systematic uncertainties listed. Note that all experiments gave values consistent with the Standard Model predictions. Also included is a brief section on the 2005 global muon decay analysis by Gagliardi *et al.* [23], which included *TWIST* measurements of ρ and δ .

Table 3.1 lists the Standard Model values of the Michel parameters, the experimental limits on the Michel parameters available before the *TWIST* experiment, and the results produced by *TWIST* prior to this work. Table 3.2 shows the Standard Model values of the weak coupling constants from Eq. (2.4), as well as their best limits from before *TWIST* and the limits derived using the global analysis discussed in Sec. 3.5. Note that the limits for

	Standard Model	Pre- <i>TWIST</i> [22]	2005 <i>TWIST</i> results [24, 25, 27]
ρ	0.75	0.7518 ± 0.0026	$0.7508 \pm 0.0003(\text{stat.}) \pm 0.0010(\text{syst.})$
δ	0.75	0.7486 ± 0.0038	$0.7496 \pm 0.0007(\text{stat.}) \pm 0.0011(\text{syst.})$
$P_{\mu}\xi$	1	1.0027 ± 0.0085	$1.0003 \pm 0.0006(\text{stat.}) \pm 0.0038(\text{syst.})$
η	0	-0.007 ± 0.013	—

Table 3.1: Standard Model values and direct experimental limits on the muon decay parameters. *TWIST* is not able to make a sensitive measurement of the parameter η . The present work is a new measurement of the parameters ρ and δ . Uncertainties listed for pre-*TWIST* values combine statistical and systematic uncertainties.

	Standard Model	Previous limits	2005 Global analysis [23]
$ g_{RR}^S $	0	< 0.066	< 0.067
$ g_{RR}^V $	0	< 0.033	< 0.034
$ g_{LR}^S $	0	< 0.125	< 0.088
$ g_{LR}^V $	0	< 0.060	< 0.036
$ g_{LR}^T $	0	< 0.036	< 0.025
$ g_{RL}^S $	0	< 0.424	< 0.417
$ g_{RL}^V $	0	< 0.110	< 0.104
$ g_{RL}^T $	0	< 0.122	< 0.104
$ g_{LL}^S $	0	< 0.550	< 0.550
$ g_{LL}^V $	1	> 0.960	> 0.960

Table 3.2: 90% confidence limits on the weak coupling constants, from a 2005 global analysis of all available muon decay data [23], including the 2005 *TWIST* measurements. (Previous limits are also taken from their listing in [23].) The major improvements on the coupling limits from this global analysis are for $|g_{LR}^S|$, $|g_{LR}^V|$, and $|g_{LR}^T|$, and these improvements arise from the inclusion of the *TWIST* ρ and δ measurements [24, 25]. $|g_{RR}^T|$ and $|g_{LL}^T|$ are identically zero and are not listed.

$|g_{RR}^S|$ and $|g_{RR}^V|$ are slightly less restrictive than the previous limits; an error was discovered in the value of $P_\mu \xi \delta / \rho$ [28, 29] used in the previous calculation, and the newer analysis uses the corrected, less restrictive value.

3.1 ρ : Peoples 1966 (unpub.), Derenzo 1969

The 1966 experiment of J. Peoples *et al.* was a measurement of the momentum spectrum of positrons from unpolarized muon decay. (The final results are available in the Ph.D. thesis of Peoples [30] but were never published; a preliminary result was published in 1965 [31], pending final evaluation of systematic uncertainties.) Positive pions were brought to rest in an active stopping target, where they decayed to muons; 80% of the muons also stopped in the target. The apparatus, including the stopping target, was in a magnetic field. This caused the muons to precess before decaying, so that the time-integrated polarization of the muons was nearly zero. The magnetic field also caused the positrons from muon decay to travel a helical path, which was tracked by a series of four spark chambers for the measurement of positron momentum (with energy loss accounted for theoretically in the detector materials, and measured in the muon target). Events in which the positron scattering was too great, as determined by the track residual at the last spark chamber, were discarded. Due to the limited momentum acceptance of the apparatus, data were taken at several magnetic field values to cover overlapping ranges in momentum; in total the momentum spectrum was studied from 10 MeV/c to 53 MeV/c. The final measurement, including statistical and systematic uncertainties, was $\rho = 0.750 \pm 0.003$.

The main systematic uncertainties in this experiment included: errors in the momentum scale, due in part to magnetic field non-uniformities; track reconstruction biases; backgrounds and unidentified hard scatters; positron annihilation in flight; and uncertainties in chamber efficiencies. Another limitation in translating the momentum spectrum measurement into a measurement of the ρ parameter was the uncertainty on η at the time.

Derenzo [32] measured the η parameter in 1969 (see Sec. 3.4 for details). The correlations between ρ and η were much different in this experiment compared to the Peoples measurement, and Derenzo was able to combine data with the Peoples result to perform a two-parameter fit and obtain $\rho = 0.753 \pm 0.003$ and $\eta = -0.12 \pm 0.21$. (To be precise, all muon decay data available at the time were used in the fit, but the Peoples measurement dominated the input values for ρ .) These results were to provide the definitive experimental value of ρ for the next 35 years, and of η for 16 years.

3.2 δ : Balke 1988

In 1988 B. Balke *et al.* conducted a modified μ SR (“Muon Spin Rotation”) experiment [33] at TRIUMF to measure the decay positron asymmetry as a function of momentum. In a standard μ SR experiment,¹ muons are stopped in a strong transverse magnetic field. The muon spins precess about the magnetic field lines. Scintillators detect the decay positrons as a function of time since the muon’s arrival, and because of the precession and the asymmetric decay distribution the number of decays seen by the scintillator oscillates. The amplitude of these oscillations is proportional to the magnetic field, the muon polarization, and the asymmetry of the muon decay. In the Balke experiment, instead of detecting the decay positrons using scintillators, they were tracked as they passed through a solenoidal magnetic field to determine their decay direction, followed by a transverse field to measure the positron momentum. Drift chambers were used for tracking. Surface muons, from pions decaying at rest at the surface of the production target, were used for the experiment, providing nearly 100% polarization. Balke *et al.* found $\delta = 0.7486 \pm 0.0038$, including statistical and systematic uncertainties (both of which were comparable).

The leading source of error in this experiment was the need for multiple settings of the magnetic fields in order to cover the full momentum range, due to limited momentum acceptance of the spectrometer. The parameter δ governs the momentum dependence of the decay asymmetry, so measuring the momentum spectrum in parts could introduce significant distortions in the combined distribution. Momentum calibration of the spectrometer, and of the beamline itself, were also significant uncertainties, as were radiative corrections. In addition, the analysis had to assume a value of ρ ; the Derenzo result discussed above was

¹For more details, see the TRIUMF Centre for Molecular and Materials Science at <http://musr.org/>. The μ SR technique is used to study the properties of materials at the atomic and subatomic level.

used, with its associated uncertainty.

3.3 $P_\mu \xi \delta / \rho$: Jodidio 1986

Before the measurement of δ by Balke *et al.* described in Sec. 3.2 was performed, A. Jodidio *et al.* used the same apparatus and beamline to measure the decay rate at the momentum endpoint of the decay spectrum [28, 29], to measure the combined quantity $P_\mu \xi \delta / \rho$. The muon decay rate becomes proportional to $[1 - (P_\mu \xi \delta \cos \theta) / \rho]$ near $x = 1$ and $\cos \theta = 1$ (opposite the muon spin direction); the range $0.97 < x < 1$ and $0.975 < \cos \theta < 1$ was used. (It was found that the systematic uncertainty due to uncertainties in the momentum acceptance were quadratic in $1 - x$, so the selected momentum range was kept small to reduce this, but large enough to keep down the statistical uncertainties.)

Data were taken in two different modes. In the “spin-held” mode, used for the physics data, the muons were held in a strong longitudinal field which prevented depolarization until their decay. In the “spin-precessed” mode, the standard spin-precessing field was used, with appropriate decay time cuts, to produce effectively depolarized data for momentum and rate calibration. The $\cos \theta$ range was divided into five bins, and in each bin the momentum edge was parameterized simultaneously for both spin modes; the fitted relative rate was extrapolated to $\cos \theta = 1$ to yield $P_\mu \xi \delta / \rho$. The experiment found $P_\mu \xi \delta / \rho > 0.9968$ at 90% confidence.

Momentum calibration and acceptance consistency between the two modes were major sources of systematic uncertainty in this result. The degree of muon depolarization prior to decay was a significant uncertainty, as well. Other important sources of error included: chamber alignment, track reconstruction errors, and corrections applied to account for Coulomb scattering in the tracking.

3.4 η : Derenzo 1969, Burkard 1985, Danneberg 2005

In 1969 S.E. Derenzo *et al.* [32] studied the low energy region of the decay momentum spectrum, with positron momenta below 6.8 MeV/c, to measure the η parameter. This experiment used a hydrogen bubble chamber to stop π^+ and μ^+ and measure the decay positron momenta. As mentioned in Sec. 3.1, a two-parameter global fit was performed using the Peoples *et al.* measurement [30] and other available muon decay data to find $\rho = 0.753 \pm 0.003$ and $\eta = -0.12 \pm 0.21$. Main sources of error included normalizations, resolution uncertainty, momentum scale distortions, and efficiencies.

The 1985 experiment by H. Burkard *et al.* [34] at the SIN (Schweizerisches Institut für Nuklearforschung, the Swiss Nuclear Physics Institute; now the Paul Scherrer Institute, or PSI) facility in Switzerland measured the transverse polarization of positrons from muon

decay. Polarized muons were stopped in a magnetic field which caused the muon spins to precess. Decay positrons emitted along the direction of the magnetic field passed through scintillators and drift chambers and into a magnetized foil (whose electrons were polarized); those which annihilated in flight in the foil produced two γ rays which were detected by an array of crystal scintillators. The intensity distribution of the γ rays was at a maximum half way between the transverse e^+ polarization, which precessed with the muon polarization, and the constant e^- polarization in the foil. Analysis of the time structure of the event rates for pairs of crystals gave the transverse polarization of the decay positrons. This was used in a global analysis of all available muon decay data to obtain the value $\eta = -0.007 \pm 0.013$, as well as determining values for several other decay parameters not measured by *TWIST*. This η result was the leading model-independent measurement of η until 2005, and so was used as the input to the previous *TWIST* measurement of ρ [24].

In 2005 N. Danneberg *et al.* repeated the experiment [35] at the PSI facility in Switzerland, measuring the transverse polarization of positrons from muon decay as per Burkard. This measurement found $\eta = 0.071 \pm 0.037$, limited almost entirely by statistics. (A new analysis with additional data is in progress as of the time of this writing.) Energy calibration was one of the major challenges of this experiment; accounting for energy loss and background events were other significant sources of error.

3.5 Global Analysis: Gagliardi *et al.* 2005

A 2005 global analysis by Gagliardi, Tribble, and Williams [23] determined limits on the weak coupling constants g_{ij}^k . All available muon decay data were used in the analysis, including the 2005 ρ and δ measurements from *TWIST* [24, 25]. The resulting coupling constant limits are listed in Tab. 3.2. The analysis sampled the probability distributions for measured parameters (ρ , δ , etc.) and parameter combinations (such as $P_\mu \xi \delta / \rho$) using a Monte Carlo method to map out the resulting probability distributions for a set of 9 independent variables Q_{RR} , Q_{LR} , Q_{RL} , B_{LR} , B_{RL} , α/A , β/A , α'/A , and β'/A , each a bilinear combination of the weak coupling constants g_{ij}^k . The last four variables are related to the polarization of the decay positron. The Michel parameters and other muon decay parameters can be written in terms of these variables. A random set of values for these variables is selected, subject to certain constraints, and the decay parameters these give are compared to the input measurements to obtain the probability that is set of values is the correct set. The process is iterated to obtain the joint probability function for these variables, as well as for the muon decay parameters. The best fit values of the independent variables are listed in Tab. 3.3.

Since the variables used for the global analysis must be positive (see for example Eqs. (2.8)), the limit on each variable can be used to determine a limit on each term separately. This

Parameter	Fit result ($\times 10^{-3}$)
Q_{RR}	< 1.14
Q_{LR}	< 1.94
Q_{RL}	< 44
Q_{LL}	> 955
B_{LR}	< 1.27
B_{RL}	< 10.9
α/A	0.3 ± 2.1
β/A	2.0 ± 3.1
α'/A	-0.1 ± 2.2
β'/A	-0.8 ± 3.2

Table 3.3: Results of the 2005 global analysis of muon decay data [23]. Limits are 90% confidence. $P_\mu = 1$ is assumed.

method was used to obtain the limits on the weak coupling constants listed in Tab. 3.2.

TWIST is unable to make a sensitive direct measurement of the parameter η , but the global analysis indirectly determined $\eta = -0.0036 \pm 0.0069$ [23].

The global analysis results demonstrate the impact that the *TWIST* measurements have already had on the knowledge of the weak interaction. This analysis is repeated using the measurements of ρ and δ from the present work in Sec. 9.2.1.

Chapter 4

TWIST Experimental Setup

I cannot stress enough that nearly without exception one cannot randomly combine a collection of circuitry and power supplies into an object and 'accidentally discover' that it is a fully functioning death ray (I speak here from bitter experience).

—James Kakalios, “*The Physics of Superheroes*”

The core of the *TWIST* experiment is a symmetric stack of high-precision tracking chambers. Highly polarized muons are stopped in a target at the centre of the chamber stack; decay positrons are then tracked to determine their momenta and angles. The low mass of the detector reduces scattering and energy loss, and allows the muons to reach the central target. The strong 2 Tesla magnetic field focusses the incoming beam, maintains the polarization of the stopped muons in the face of muonium¹ formation, and allows the reconstruction of the momenta of decay positrons. A schematic of the *TWIST* detector can be found in Fig. 4.1.

This chapter describes the TRIUMF beamlines used to produce and transport the polarized muons (Sec. 4.1), details the components of the *TWIST* detector, magnet, etc. (Sec. 4.2), and discusses the calibration of the detector alignment and timing (Sec. 4.3).

4.1 TRIUMF Beamlines

4.1.1 Proton Beam and Production Target

The production of the muon beam used by *TWIST* begins with a 500 MeV proton beam provided by the TRIUMF cyclotron; a pulse of protons is produced every ~ 43 ns. *TWIST* muons are produced in beamline 1A. The protons collide with the T1 production target—here *TWIST* uses a shaped piece of graphite, or occasionally steel-jacketed beryllium—and produce pions in a wide range of energies.

¹Muonium is a bound state of a muon with an electron. Its formation rate depends on the material in which it is stopped. In the absence of an external magnetic field fixing the muon polarization, magnetic interactions between the muon and electron spins can cause the muon spin to change.




Figure 4.1: Schematic view of the *TWIST* detector. Muons travel down the beampipe in vacuum until they are well inside the magnetic field provided by the solenoid and steel yoke. They slow down as they pass through the detector chambers, stopping in the central target. See text for details.




Figure 4.2: M13 beamline layout [36]. Muons and other particles are produced at the 1AT1 target. B1 and B2 are dipole magnets, and Q1–Q7 are quadrupoles. F1–F3 are points where the beam is designed to be in focus; in the configuration used for *TWIST* there is no vertical focus at F2. The *TWIST* apparatus is situated just downstream of the final focus (F3). Slits and jaws control the shape of the beam at various points; in addition, after B1 disperses the beam, the horizontal slit at F1 is used to select the range of particle momenta.

4.1.2 Muon Beamline

TWIST uses the M13 secondary beamline for collection and delivery of muons from the 1AT1 production target. The beamline can be configured to deliver positive or negative particles by selecting the polarity of the currents in the beamline magnets. Positive particles were selected for this experiment, and as such all particles discussed in this section are positive.

Most pions leave the production target, but some come to a stop inside. As described in Sec. 2.2.1, pions decay with a lifetime of 26 ns into a muon and a neutrino. In the rest frame of the pion, the resulting muons are produced with negative helicity² (i.e. their spins are always opposite their momentum vectors). In the lab frame, if the pion decays in flight, its momentum provides a boost to the decay muon, breaking the relationship between momentum and spin. The polarization of these muons depends on the pion momentum distribution and the beamline geometry. Muons from pions decaying at the surface³ of the production target are commonly referred to as “surface muons,” while those from pions decaying in flight are called “cloud muons.” Note that surface muons, coming from pions decaying at rest, cannot have momentum greater than 29.8 MeV/c, as explained in Sec. 2.2.1.

The layout of the M13 beamline is shown in Fig. 4.2. The beamline collects muons

²This is true only within the Standard Model. Exotic weak currents can produce an admixture of positive helicity muons.

³If the pion decays more deeply inside the production target, the muon will scatter on its way out, reducing its polarization.

and other particles (pions, positrons, and protons) which leave the production target at an angle of 135° with respect to the proton beam. Momentum selection is done with the combination of a dipole magnet (B1) and a slit; since the curvature of a particle's path through a magnetic field depends on its momentum, the slit blocks particles with all but a specific, narrow range of momenta. The width of the slit determines the momentum "bite," $\Delta p/p$. During standard *TWIST* data acquisition, the beamline momentum was set to 29.6 MeV/c to largely avoid the muon momentum cutoff at 29.8 MeV/c, and the momentum bite was set to approximately 0.8%. To improve the stability of the muon beam, regulators were implemented beginning with the 2004 data taking to maintain the measured fields in the B1 and B2 dipole magnets at a given set point.

A 3 μm polyester window valve in the beamline stops radioactive beam gas (likely ^{11}C and ^{13}N for the most part), at the cost of a small scattering of the muons. This window was essential when using the graphite target since the radioactivity from this gas produced significant rates in our chambers.

The currents in the dipole and quadrupole magnets, the positions and sizes of the slits and jaws, and other properties of the beamline are controlled directly by the experimenters using EPICS (Experimental Physics and Industrial Control Software).

The beamline momentum must be calibrated, since the momentum selection depends on the setting of the B1 dipole magnet, the position of the selection slit, and even the precise location and geometry of the production target. The 29.8 MeV/c surface muon cutoff was used as a calibration point. The calibration can be obtained by plotting the muon flux as a function of B1 magnetic field strength, where the currents in the other beamline magnets are scaled as B1 is scaled; see Fig. 4.3. The position of the mid-point on the edge gives the B1 setting corresponding to 29.8 MeV/c, and the width of the edge shows the momentum bite. More details on the edge scan calculations are given in the author's 2001 technical note [37].

The muons travel in vacuum until they are well inside the magnetic field; this prevents scattering from broadening the beam. It is important to keep the beam as small as possible and as close to the z axis as possible in order to preserve its polarization [38].

At the end of the beampipe and well inside the magnetic field of the *TWIST* solenoid is a ~ 22 cm long chamber, called the gas degrader, with a tunable mixture of helium and CO_2 . The amount of energy lost by muons in the chamber depends on the He: CO_2 ratio. This allows the fine-tuning of the mean stopping position of the muons. A regulator was implemented beginning with the 2004 data-taking period, where slow control software reacted to shifts in the stopping distribution of muons within the DCs and PCs by adjusting the He: CO_2 ratio, greatly improving the stability of the stopping distribution. The stopping distribution in the chambers was calibrated by measuring the distribution of muon stops in

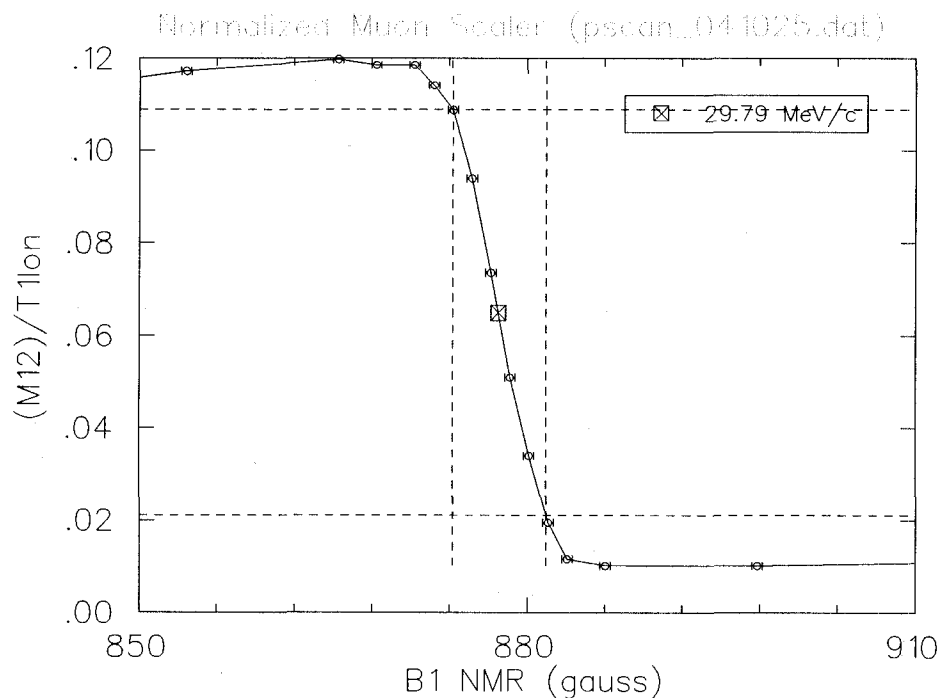


Figure 4.3: Muon flux vs. dipole magnetic field, for calibrating the M13 beamline momentum. The vertical axis is the number of muons detected, normalized to a counter proportional to the proton beam current. The horizontal axis is the magnetic field of the B1 dipole, which is used to select the beamline momentum; the field strength is measured with an NMR probe. The dashed lines show the 10% and 90% flux points, used to gauge the edge width, and the box shows the position of the edge centre. Here, the edge occurs at 877.75 Gauss, and the width is $\Delta p/p = 0.82\%$. Data were taken on 25 October, 2004, using the graphite target.

the chambers⁴ and comparing to MC, along with MC predictions of the relationship between muon stops in chambers and the mean stopping position of muons in the target itself.

The muon rate was approximately 2700 muons per second, as measured by the trigger rate (compensating for the 1% of triggers caused by positrons; see Sec. 4.2.2). The beam positron rate was approximately 7800 positrons per second, as determined from the mean number of beam positrons identified by the reconstruction software (Sec. 5.3.1) in the 16 μ s time window of each event.

Time of Flight Structure of the Beam

The difference in time of flight between different particles is useful in identification. The speed of a particle of mass m and momentum p is $v = pc/E$, where $E = \sqrt{p^2 + m^2}$, with mass, momentum, and energy measured in MeV/ c^2 , MeV/ c , and MeV. Then the amount of time t it takes to travel a distance d is

$$t = \frac{Ed}{pc} = \frac{d\sqrt{p^2 + m^2}}{pc}. \quad (4.1)$$

The distance through M13 from the 1AT1 production target to the *TWIST* trigger scintillator is 11.4 m. So, for example, at 29.6 MeV/ c , a muon takes 141 ns to reach the trigger, while a positron takes just 38 ns.

The difficulty is that we cannot know when a given particle left the production target. However, the proton bunches leave the cyclotron every 43 ns, and a capacitive probe on the proton beamline signals when this happens. At a given momentum, the time t_{cp} between the scintillator signal and the next capacitive probe (CP) signal will be different for each particle, according to its mass; this can be used for particle identification. Since t_{cp} is the time between the particle's arrival at the scintillator and the *next* CP signal, it will be smaller for later particles. t_{cp} is often called the **CP time of flight (CPTOF)**, or "TCAP" ("capacitive probe time"). Typical CPTOF spectra are shown in Fig. 4.4. (In practice, the next two capacitive probe signals are recorded, so that a continuous region of the CPTOF spectrum can be selected for surface muons.) Particles which are created in time with proton bunches will produce well-defined peaks in the CPTOF spectrum. Surface muons, however, are created from pions decaying at rest in the production target; the pion lifetime is about 26 ns, comparable to the CP period, so surface muons are not confined to a narrow CPTOF band. Cloud muons are the decays of pions in flight, and so must have decayed promptly; these are therefore produced in time with the proton bunches and produce a peak in the CPTOF spectrum. Since both surface and cloud muons have the same time of flight, the cloud muons are found directly under the peak of the surface muon distribution (see Fig. 4.4).

⁴Note that, because energy loss is a stochastic process, muons of a given energy will have a range of stopping positions. Since the *TWIST* target is very thin, many muons (around 25%) stop in the detector chambers instead.

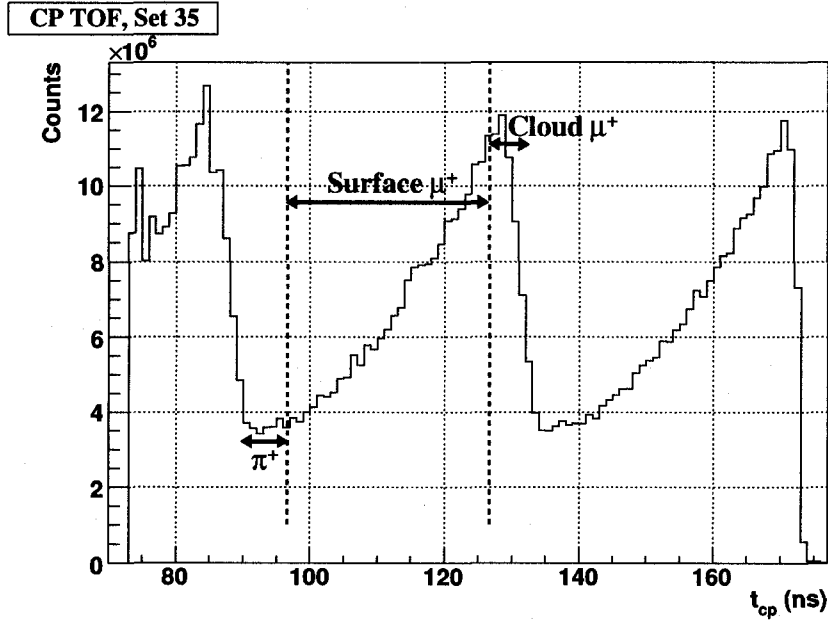


Figure 4.4: Time structure of the M13 beam (CPTOF) at 29.6 MeV/c. CPTOF is the time between a particle's arrival at the trigger scintillator and the *next* CP signal, so that later particles will have smaller CPTOF values; hence plots such as this showing the pion decay lifetime will appear backwards. The regions indicated by the arrows show the location of cloud muon and pion contamination, and the region of pure surface muons. The region $97 < t_{cp} < 127$ ns, indicated by the dashed lines, was selected for the final spectra. The bump at $t_{cp} \approx 110$ ns is believed to be an artifact of electronics, and should not affect the use of CPTOF for particle identification. Two capacitive probe signals are recorded for each trigger, giving two values of t_{cp} ; both are shown in this histogram so that a continuous region of the spectrum can be selected for surface muons.

Note that the signal threshold of the trigger scintillator was set to be relatively insensitive to positrons, while still efficiently detecting muons. As a result, few positrons appear in the CPTOF spectra. Most positrons seen by *TWIST* are produced by pion and muon decays in flight, so those that register in the scintillator produce a peak in the CPTOF spectrum. Positrons are also produced by muons stopped in various places around the production target and along the beamline, resulting in a constant background in the CPTOF spectrum.

A cut on the CPTOF of the trigger particle allows the selection of a clean sample of surface muons, free of contamination by cloud muons or pions [39].

4.1.3 Muon Beam Monitor

To monitor the position and angle (emittance) of the muon beam, for creating realistic beam input for the simulation, *TWIST* designed the Time Expansion Chamber, or TEC [38], which was installed in 2004. The device consists of two modules, one oriented in the x direction and the other in y . They measure the position as a function of z of the incoming muon, providing position and angle information of each particle's trajectory. The system is very low mass, and operates at low gas pressure (8 kPa, or about 0.08 atm) to minimize scattering. The TEC is situated inside a gas box in the beam pipe, near the nominal focus of the beamline, just entering the fringe of the *TWIST* solenoid's magnetic field; the effect of the field on the particle trajectories is small enough that the tracks through the TEC are straight to a very good approximation. Due to the low operating pressure, the TEC is only sensitive to particles such as muons which leave substantial amounts of ionization in the gas; positrons are not efficiently detected. (The positron beam was studied using the standard *TWIST* drift chambers with the solenoid magnet turned off, as described in Sec. 6.2.1).

The information from the x and y modules is used to map out the position and angle profiles of the incoming muon beam at the central plane of the TEC, including correlations between the profile parameters (angles vs. positions etc.). These profiles are corrected for the effects of multiple scattering in the TEC, and are then used to define the muon beam used in the simulation (Sec. 6.2.1).

4.2 *TWIST* Detector

4.2.1 Magnet

The *TWIST* magnetic field is supplied by a liquid-helium-cooled superconducting solenoid with a 1 m bore. The magnet is placed inside a ~ 3 m cube-shaped steel yoke, designed to increase the uniformity of the field within the detector region. The field was mapped using a custom-built apparatus involving a set of Hall probes on a rotating arm, with a movable NMR probe for calibration; the resulting magnetic field map is accurate to approximately ± 1 Gauss. The standard operating setting for *TWIST* is ~ 20 kGauss at the centre of the

solenoid; at this setting the field is uniform to within about 80 Gauss within the tracking region ($-60 < z < 60$ cm, $r < 16$ cm).

The Hall probes in the magnetic field mapping system measured only the z component of the field. The magnetic field was simulated using the OPERA-3d finite element software [40]. This provided the full three-dimensional magnetic field throughout the detector volume and well outside the measured region, as well as providing a smoother magnetic field map than would be possible from directly interpolating the field measurements. The current in the simulated coils, as well as the shape, structure, and material of the yoke steel, were adjusted in the OPERA-3d simulation to produce a magnetic field map as close as possible to the measured map. The final simulated field map agrees with measurements to within ± 2 Gauss [41].

The field strength is measured by an NMR probe located inside the *TWIST* detector stack, just outside the active region. A second probe was used periodically to confirm the stability of the relative readings. The measurements from these probes allow the determination of the absolute field strength to high precision (significantly better than 1 Gauss). The OPERA-3d field map was scaled globally to match the reading of the main NMR probe at its location.

4.2.2 Trigger and Other Scintillators

The trigger for data acquisition is supplied by a thin circular scintillator (195 μm thick, 3 cm diameter), placed just past the end of the M13 beampipe (at $z = -80$ cm). Positrons passing through this scintillator deposit little energy, while muons and heavier particles deposit much more. The threshold for the amount of energy in the scintillator required to produce a digital signal was set high enough to prevent most beam positrons from causing a trigger, while nearly all muons trigger. Part of the signal from the trigger scintillator is passed to an ADC (Analog to Digital Converter) to measure the amount of energy deposited; from this it was found that about 1% of the scintillator signals are due to positrons.

The trigger scintillator is mounted between a pair of plastic light guides. Each light guide covers half of the circumference of the trigger scintillator. The guides carry the scintillation light out to a pair of photomultiplier tubes, designated M1 and M2. These are used independently and as a linear sum, M12. The trigger is a coincidence of M1, M2, and M12. The scintillator provides an unbiased trigger which depends only on the passage of a muon.

A pair of positron scintillators is located just downstream of the muon scintillator; the muon trigger scintillator covers a 3 cm hole at their centre. The positron scintillators are read out by wavelength-shifting fibers wrapped around their circumferences, for a total of four photomultiplier signals. The coincidence of all four signals is designated PU. A similar

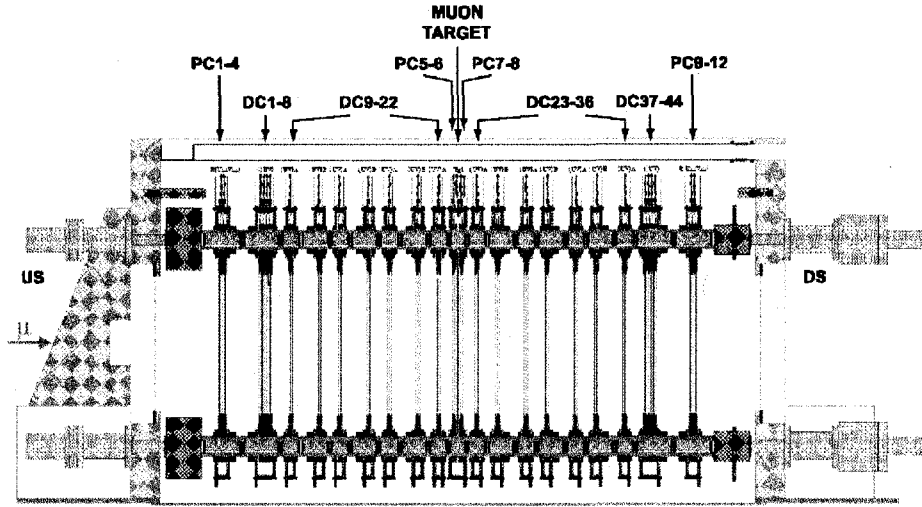


Figure 4.5: Side view of the *TWIST* spectrometer, showing the arrangement of drift chambers (DCs) and multi-wire proportional chambers (PCs). Muons enter from the left in the figure, slow down in the chambers, and stop in the aluminum foil (“Muon Target”) at the centre. The gas degrader is located inside the beampipe, just before the end. Drawing is approximate.

removable scintillator pair,⁵ designated DS, can be placed downstream of the detector, outside the steel yoke. These scintillators are used for some specialized studies, such as alignment.

4.2.3 Tracking Chambers

TWIST makes use of two types of tracking chambers: the Multi-Wire Proportional Chambers (MWPCs or simply PCs), which have fast response times (typically less than 20 ns) and are used mainly for timing, and the Drift Chambers (DCs), which have slow drift times (hundreds of nanoseconds) and are used for high-precision tracking. Figure 4.5 shows the arrangement of chambers within the spectrometer. The chamber design will be described here in brief; more detail can be found in ref. [42]. All chamber assembly, including the mounting of wires and foils, was done by hand at TRIUMF by *TWIST* personnel, including the author.

Both types of tracking chambers are of a planar design, perpendicular to the z axis. A single chamber consists of a plane of $15\text{ }\mu\text{m}$ diameter sense wires flanked in z by two cathode planes ($6.35\text{ }\mu\text{m}$ aluminized Mylar foils under high tension). Each plane is assembled on a 60 cm diameter glass annulus, with a thin printed circuit board laminated on. The glass plate has an inner diameter of $\sim 398\text{ mm}$. The DCs are generally assembled in two-chamber

⁵There were actually two versions of DS scintillators used in 2004 data-taking. A version read out with wavelength-shifting fibers, and including a hole at the centre similar to the PU scintillators, was used when the magnetic field was turned on. A simpler version, using a plastic wave guide and without a hole, was used when the field was turned off.

modules, consisting of three cathode planes and two wire planes. There are also two “dense stack” modules of DCs, consisting of eight wire planes; the set of modules not part of the dense stack are sometimes referred to as the “sparse stack”. DC planes run from $|z| = 4.4$ cm to $|z| = 50.0$ cm. There are modules of four PCs each at the upstream and downstream ends of the detector ($58.4 \text{ cm} \leq |z| \leq 60.0 \text{ cm}$), as well as a target module consisting of the stopping target with two PCs on each side; the stopping target also acts as a cathode foil for the two innermost PCs.

All 44 DC planes are identical, but become u or v planes depending on whether they are tilted by $+45^\circ$ or -45° (looking downstream) as they are mounted within the module. This is to keep the detector isotropic and reduce the effect of wire sag; horizontal wires would sag under gravity, by about $2.3 \mu\text{m}^6$, and vertical wires would not. (See Sec. 4.2.5 below for a complete description of the *TWIST* coordinate system.) The cathode planes are 4 mm apart. The wire planes are almost exactly in between the cathode planes, but due to the details of the cathode plane construction the wires are $\sim 15 \mu\text{m}$ off centre in z ; the present analysis now accounts for this offset. Wire and cathode foil positions are defined relative to 4 mm-thick spacers, made of thermally stable Sitall ceramic manufactured to be optically flat and parallel. (The same material is used for astronomical telescope mirrors.) Additional, thicker Sitall spacers (20 mm or 40 mm) define the gaps between modules. Module gaps are alternated to reduce the periodicity of the detector.

DC planes have 80 wires at a 4 mm pitch. (Due to limited availability of TDCs (Time to Digital Converters) for recording the wire signals, not all of the dense stack wires were instrumented during the 2002 data taking. This was corrected for the 2004 run, significantly improving the reconstruction.) Each plane is logically divided into drift cells, one per wire; a cell is the region of the chamber read out by a given wire. Drift chambers are filled with dimethylether (DME) gas, which has a low mean atomic number, a slow drift time (providing better resolution) and a low Lorentz angle (the angle by which the drifting electrons will be deflected by the magnetic field).

Each of the 12 PC planes have 160 wires at a 2 mm pitch. All wires in the outer PC planes are instrumented, while only the central 48 wires of each of the target module PCs are instrumented because of the construction of the stopping target (see Sec. 4.2.4). The PCs use a CF_4 /isobutane mixture, providing the fast response needed for timing.

The chambers are numbered starting at the upstream end. Thus, DCs 1–22 are upstream of the stopping target, and 23–44 are downstream; PCs 5–8 are in the target module. A “global” plane number can be defined, where both DCs and PCs are counted; this is useful for describing the muon stopping distribution, for example. In this scheme, the target PCs are given global plane numbers 27–30.

⁶Yuri Davidov, “TWIST Wire Sags”, *TWIST* General internal forum, 27 January, 2002. https://twist.phys.ualberta.ca/forum/view.php?bn=twist_general&key=1012182202

The gaps between chamber modules are filled with a 97/3 helium/nitrogen mixture, rather than air, to reduce the amount of material in the detector. Nitrogen is included to reduce sparking of the chambers or electronics. The helium mixture is kept at atmospheric pressure. The differential pressure between the tracking chambers and the interstitial helium is held stable during data taking; differences in the pressure will cause the outer cathode foils in each module to bulge out or in, changing the electric fields and drift time maps of the chambers. Online monitoring of the foil bulge was implemented for the first time partway through the 2004 data taking.

Pre-amplifiers mounted on the outside of each chamber amplify the signals for transmission outside the detector and magnet. The signals are enhanced again by more powerful post-amplifiers, then are passed through discriminators which produce rectangular pulses from the amplified signals. These pulses are sent for digitization to Lecroy 1877 FastBus TDCs (Time to Digital Converters).

TDCs are operated in common-stop mode. The signal from the muon trigger scintillator, delayed by approximately $10\ \mu\text{s}$, is used to stop timing; at this point the times of all hits received in the previous $16\ \mu\text{s}$ are recorded. In this way information is collected for $6\ \mu\text{s}$ before the trigger and $10\ \mu\text{s}$ after.

Each TDC can record up to eight signals per wire in each event; if additional signals are received, the most recent eight are kept. A fixed blanking time of $80\ \mu\text{s}$ was imposed after each trigger, preventing the recording of additional signals to give each TDC time to finish signal conversion for readout. Each TDC was required to produce a header listing the number of converted hits, even if no hits were received, to verify that all TDCs were able to complete hit conversion.

4.2.4 Stopping Target

A new muon stopping target was used for 2004 (and 2003) data taking. The target is a $71 \pm 1\ \mu\text{m}$ thick foil of high-purity aluminum ($> 99.999\%$), a material known to cause very little depolarization in stopped muons. Due to limitations in the size of foil that could be put under tension to keep it flat (necessary since the target is used as a cathode foil in the central PCs (6 and 7); flatness is also important for maintaining uniform target thickness), the foil was mounted over a cutout in a stretched $25\ \mu\text{m}$ aluminized Mylar foil (thicker than the usual cathode foils). The 150 mm aluminum target was glued to the edges of a 120 mm cutout, with a conductive silver epoxy to maintain conductivity. Kapton masks were applied to both sides of the target assembly to cover the edges of the aluminum and Mylar foils, to prevent high electric fields around the sharp edges; these masks left a 110 mm diameter opening to expose the stopping foil. Similar masks were applied to the neighbouring foils of PC6 and PC7 for symmetry. A schematic of the stopping target is shown in Fig. 4.6. As




Figure 4.6: Construction of the muon stopping target. The aluminum stopping target is attached to a cutout in an aluminized Mylar foil. Kapton masks cover the resulting sharp edges, preventing high electric fields in those regions. See text for further details. (Diagram by Mina Nozar.)

mentioned above, only the central 96 mm region is instrumented.

The aluminized Mylar foil holding the target is positioned at $z = +60 \pm 50 \text{ } \mu\text{m}$ in the *TWIST* coordinate system. Since the Mylar foil is $25 \text{ } \mu\text{m}$ thick, and the aluminum target is $71 \text{ } \mu\text{m}$ thick (both $\pm 1 \text{ } \mu\text{m}$ or less), the centre of the aluminum stopping target is at $z = 60 - (25/2 + 71/2) = +12 \pm 50 \text{ } \mu\text{m}$.

4.2.5 *TWIST* Coordinate Systems

The standard “lab frame” coordinate system is in (x, y, z) ; z is along the axis of the beamline, in the “downstream” direction (i.e. with the direction of the particles’ motion), y is up, and x is defined to complete the right-handed coordinate system.

The *TWIST* coordinate system used in the reconstruction matches that of the drift chambers; as mentioned in Sec. 4.2.3, the *TWIST* detector was designed with the wires of the drift chambers oriented at 45° to vertical. The system is in (u, v, z) ; z is again along the axis of the beamline, and u and v are found by a $+45^\circ$ rotation from x and y , about the z axis.

The z origin is defined to be half way between the wire planes of the central two PCs (PC6 and PC7).

4.2.6 *TWIST* Control Software

Many of the *TWIST* systems, including chamber high voltage, discriminator thresholds, etc., are controlled through a custom interface prepared using Tcl/Tk [43]. Data acquisition is handled with the Midas software [44], developed jointly at PSI and TRIUMF. Slow controls, such as regulating the M13 dipole magnets or the mixture in the gas degrader, were implemented through interfaces to Midas.

4.3 Calibration

4.3.1 Alignment

Once the detector is assembled, the chambers have some small random translational and rotational offsets in the u and v directions. (The Sitall spacers ensure precise alignment in z , independent of temperature; the z position of each Sitall is known to within a few microns.) These offsets are measured using wide-emittance 120 MeV/c pion beams, with the magnetic field turned off and using the PU and DS scintillators as triggers. Straight lines are fit to the pion tracks, and the mean tracking residuals (the distance between the fitted track and the measured particle signal position) are used to determine the translational and rotational offsets. The precision of the translational offsets was found to be about 10 μm , and the precision of the rotational offsets to be 0.02° , determined using Monte Carlo studies.

The detector is then aligned to the magnetic field, using the helical tracks from standard muon decay data.⁷ The wire centre data are fit using a modified helix formula:

$$u = a \sin(bz - c) + u_c + d_u z \quad (4.2a)$$

$$v = a \cos(bz - c) + v_c + d_v z \quad (4.2b)$$

$$(4.2c)$$

where

$$d_u = \sin(\theta_B) \cos(\phi_B) \quad (4.3a)$$

$$d_v = \sin(\theta_B) \sin(\phi_B) \quad (4.3b)$$

$$(4.3c)$$

and θ_B and ϕ_B are the polar and azimuthal angles describing the misalignment of the detector with respect to the magnetic field. Since θ_B is small, the correction can be applied by shifting the drift chambers in u or v by amounts which are linear in z . The error introduced by this approximation is small.⁸

⁷Roman Tacik, "Misalignment of detector wrt B", *TWIST Physics internal forum*, 1 October 2004. https://twist.phys.ualberta.ca/forum/view.php?bn=twist_physics&key=1096660263

⁸Robert MacDonald, "B-Field Alignment: Skew vs Tilt", *TWIST Alignments internal forum*, 7 July, 2006. https://twist.phys.ualberta.ca/forum/view.php?bn=twist_alignments&key=1152298619

The offsets are put into calibration files, which are used by the track reconstruction software (Sec. 5.3.1) to correct the assumed tracking chamber positions. The GEANT simulation software also reads these calibration files and positions the chambers accordingly (Sec. 6.2.4).

4.3.2 Wire Time Offsets (t_0)

Variations in cable lengths, variations in electronics, and other effects can offset the timing of signals recorded by the chambers and scintillators. These time offsets, or t_0 's, must be measured *in situ* and accounted for during the data analysis.

t_0 values are measured using wide-emittance 120 MeV/c pion beams, with the magnetic field turned off, and with the PU and DS scintillators as triggers. (These are the same data used for alignment.) Straight lines are fit to the pion tracks, accounting for time of flight through the detector, in order to determine when the ionization was deposited in each drift cell. From this the drift times can be calculated. A drift time distribution is assembled for each wire in the detector, and the position of the leading edge of this distribution is the t_0 for that wire.

As the effects which cause variations in the time offsets are not included in the simulation, t_0 's are determined only for data.

Chapter 5

TWIST Analysis

Real tracks are usually described by five parameters, but five-dimensional mesas are more difficult to visualize.

—F. James, NIM **211** (1983), 145–152

There are four main components of the *TWIST* analysis chain; Fig. 5.1 shows a diagram of the program flow. Given raw data from the detector (or the simulation), the Mofia analysis program (Sec. 5.3) classifies each event and reconstructs the positron tracks. Mofia produces files of processed data in the ROOT “tree” format [45]. A histogram of the decay spectrum is produced from the trees using spectrum summation (“tree summing”) software (Sec. 5.4); here the various cuts are applied to eliminate tracks or entire events, and the “best” decay track, if any, is identified for inclusion in the decay spectrum. The decay spectra from data and simulation are then calibrated against each other for consistent momentum reconstruction (Sec. 5.5), using the 52.8 MeV/c kinematic endpoint of the momentum spectrum. The spectrum summation is then repeated, applying the momentum calibration to each event before including it in or excluding it from the spectrum.

The final decay spectra produced from data and MC can then be compared using the spectrum fitter (Sec. 5.1). This determines the differences ($\Delta\rho, \Delta\delta, \Delta P_\mu\xi$) in the Michel parameters between data and MC. The true values of the parameters in data are found by adding the parameter values used to generate the MC to the differences found by the spectrum fitter: $\rho_{data} = \Delta\rho + \rho_{MC}$, etc. This allows *TWIST* to perform a blind analysis, by hiding the MC Michel parameter values. See Sec. 6.1 for details.

The same analysis is applied to both data and simulation. In this way, spectrum distortions due to detector response, positron energy loss and scattering, reconstruction biases, and other effects, which would otherwise lead to systematic errors in the measurement, are cancelled, at least to the level that the simulation accurately represents the data.

The spectrum fitter will be described first, to provide context for the rest of the analysis. The determination of the drift time maps used to convert TDC times into drift chamber locations is discussed, followed by the track reconstruction software and the spectrum summation. The technique used for energy calibration is presented. Finally, the choice of

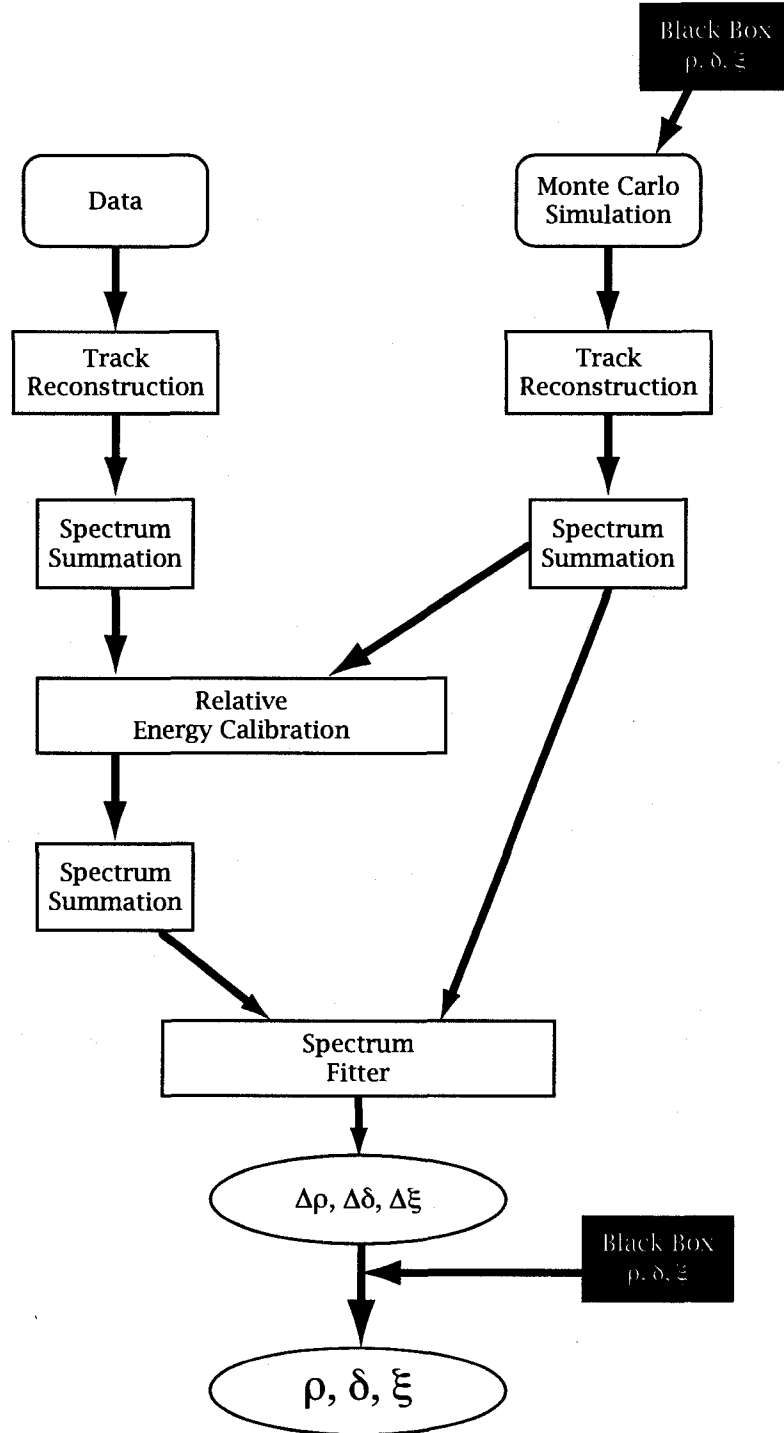


Figure 5.1: Flowchart of *TWIST* analysis software chain. The same analysis and summation software is applied in the same way to both data and simulation. See text for details.

region of the $(p, \cos \theta)$ spectrum to use for extracting the Michel parameters—i.e. the fiducial region—is discussed, influenced largely by considerations of the analysis software. Major improvements since the collaboration’s previous ρ and δ measurements are highlighted for each component.

5.1 Michel Spectrum Fitting

The goal of the spectrum fitting procedure is to determine the differences in the decay parameters describing two given decay spectra. The fit is performed using the ROOT rewrite of the standard MINUIT fitting libraries [46, 47].

The muon decay spectrum (Eq. (2.9) on p. 9) is linear in the decay parameters ρ and η , and in the parameter products $P_\mu \xi$ and $P_\mu \xi \delta$. Therefore, if we treat the products as an alternative parameterization, the derivatives of the spectrum with respect to these parameters are independent of the decay parameters. Let $S \equiv \frac{d^2 \Gamma}{dx d(\cos \theta)}$ represent the muon decay spectrum; then the linearity in these decay parameters allows us to express the data spectrum S_D in terms of the simulated spectrum S_M using the derivatives:

$$S_D = S_M + \frac{\partial S}{\partial \rho} \Delta \rho + \frac{\partial S}{\partial \eta} \Delta \eta + \frac{\partial S}{\partial P_\mu \xi} \Delta(P_\mu \xi) + \frac{\partial S}{\partial P_\mu \xi \delta} \Delta(P_\mu \xi \delta) \quad (5.1)$$

where

$$\frac{\partial S}{\partial \rho} = \left[\frac{m_\mu}{4\pi^3} E_{max}^4 G_F^2 \sqrt{x^2 - x_0^2} \right] \frac{2}{9} (4x^2 - 3x - x_0^2) \quad (5.2a)$$

$$\frac{\partial S}{\partial \eta} = \left[\frac{m_\mu}{4\pi^3} E_{max}^4 G_F^2 \sqrt{x^2 - x_0^2} \right] x_0 (1 - x) \quad (5.2b)$$

$$\frac{\partial S}{\partial P_\mu \xi} = \left[\frac{m_\mu}{4\pi^3} E_{max}^4 G_F^2 \sqrt{x^2 - x_0^2} \right] \frac{1}{3} \sqrt{x^2 - x_0^2} (1 - x) \cos \theta + (R.C.) \cos \theta \quad (5.2c)$$

$$\frac{\partial S}{\partial P_\mu \xi \delta} = \left[\frac{m_\mu}{4\pi^3} E_{max}^4 G_F^2 \sqrt{x^2 - x_0^2} \right] \frac{2}{9} \sqrt{x^2 - x_0^2} \left(4x - 3 + \left(\sqrt{1 - x_0^2} - 1 \right) \right) \cos \theta. \quad (5.2d)$$

Recall from Eq. (2.9) that $x = E/E_{max}$ is the positron’s reduced energy, and $x_0 = m_e/E_{max}$. In the *TWIST* formalism the anisotropic radiative corrections are multiplied by ξ , meaning we can treat $P_\mu \xi$ consistently as a single parameter. Since the radiative corrections are calculated within the Standard Model, and $P_\mu \xi|_{SM} = 1$, this introduces no new assumptions. In the event that non-Standard Model values for any of the Michel parameters are found, of course, everything related to the radiative corrections will be revisited.

The momentum dependences of Eqs. (5.2) are shown to scale in Fig. 5.2; full distributions in $(p, \cos \theta)$ are shown in Fig. 5.3 for all except $\partial S/\partial \eta$. The regions in $(p, \cos \theta)$ space where the derivatives differ by the largest amounts are those regions of the spectrum which are most sensitive to the values of the corresponding Michel parameters. For example, the ρ and $\xi \delta$ spectra both have peaks at the endpoint ($p = 52.8$ MeV/c) while the ξ spectrum does not, so a higher accepted momentum range will increase the sensitivity to those parameters.




Figure 5.2: Momentum dependences of the decay spectrum derivatives given in Eqs. (5.2). S , the standard Michel spectrum, is shown for $\cos\theta_s = 1$, and $P_\mu = 1$ has been assumed. Each derivative has been multiplied by its Standard Model value for scale, except for the η derivative which has been multiplied by 300 times its 1σ world limit of -0.011 [14]. Both the ρ and $\xi\delta$ derivatives are shown, but they differ only by terms in $x_0^2 = (m_e/E_{max})^2$. Vertical axis is proportional to decay probability.

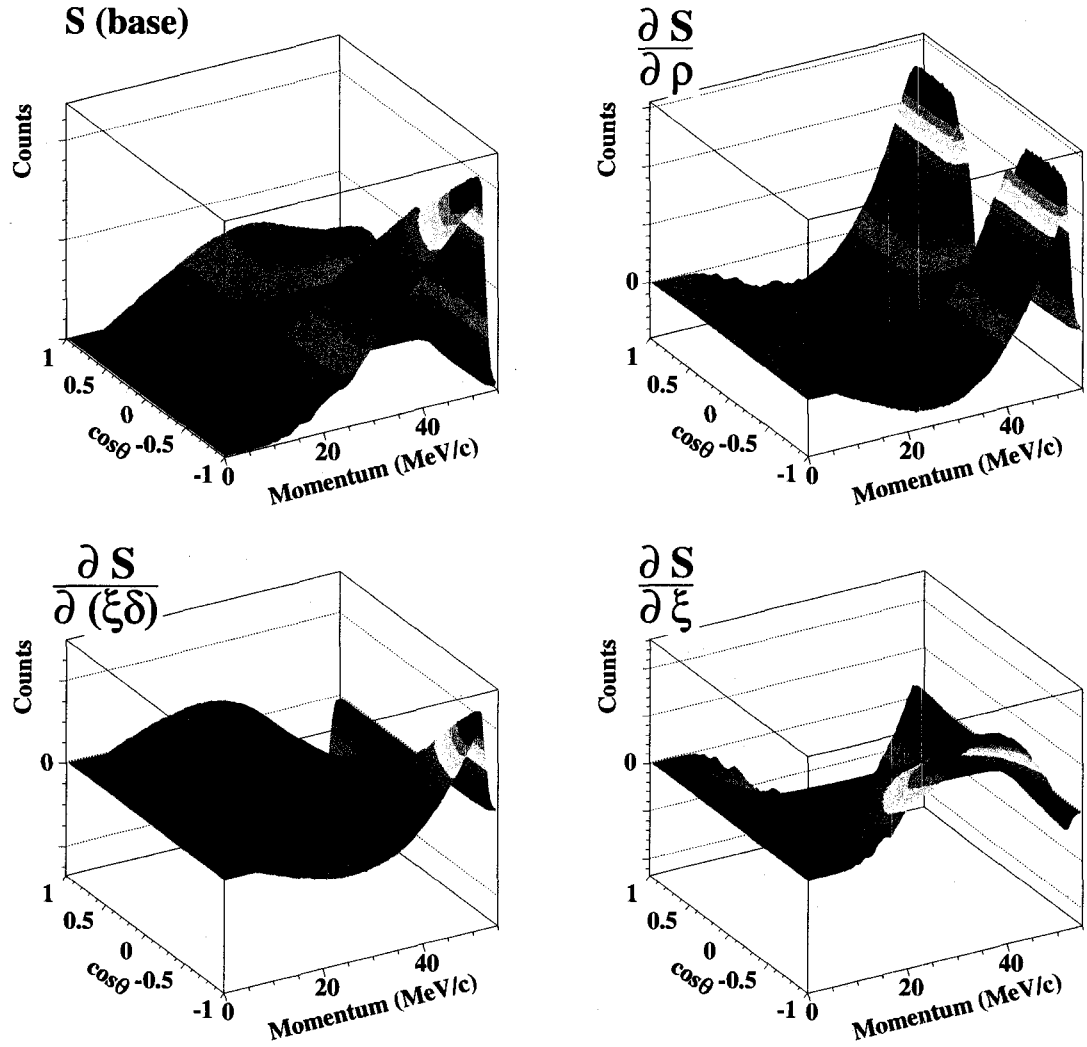


Figure 5.3: Two-dimensional distributions of the Michel spectrum S and three of its derivatives (η derivative not shown). Spectra are all fully simulated and reconstructed. $\cos\theta = 1$ is in the direction of the muon beam, opposite the direction of muon spin. The figures shown here are not normalized; normalization is done relative to the base spectrum at fitting time. Note that the derivative spectra can take on negative values.

Since the ξ and $\xi\delta$ spectra are both multiplied by $\cos\theta$, the small-angle regions of the spectrum (near $|\cos\theta| = 1$) both upstream and downstream will be most affected by those parameters, and the distinction from the shape of the ρ spectrum will be largest.

The coefficients $\Delta\rho$, $\Delta\eta$, $\Delta P_\mu\xi$, and $\Delta P_\mu\xi\delta$ in Eq. (5.1) represent the differences in the decay parameters between the data and MC spectra. These are the actual fit parameters used in the Michel spectrum fit.

Because of the small size of the η contribution to the decay spectrum (as seen in Fig. 5.2), *TWIST* lacks the sensitivity to measure η directly. Instead, an external measurement of η is assumed,¹ and only $\Delta\rho$, $\Delta P_\mu\xi$, and $\Delta P_\mu\xi\delta$ are fit. (While the “black box” remains closed, the value of η is fixed to whatever hidden value was chosen. Once the parameter values are revealed, the fits are repeated with the value of $\Delta\eta$ set to match the external measurement.) The value of η assumed by the spectrum fitter has a small but significant effect on the results, as described in Sec. 8.7.2.

Any two Michel spectra can be compared by the spectrum fitter—data vs. MC, data vs. data, MC vs. MC—to determine apparent differences in their Michel parameters. The two spectra to be fit, regardless of their origins, are designated **DATA** and **BASE**. Both are normalized to integrate to 1 in their fiducial regions. The “fitted” spectrum is a linear combination of **BASE** and the derivatives, and is also normalized to 1 in its fiducial region. To establish a correspondence between the fit coefficients in front of the derivative spectra and the differences in the associated Michel parameters, the **EFF** spectrum is used; this is a simulated spectrum generated under the same conditions as the derivative spectra. It is also normalized to 1 in its fiducial region, and the derivative spectra are normalized relative to **EFF** by comparing values of *nthrow*.² The **EFF** spectrum may be the same as the **DATA** or **BASE** spectra if either of those are simulated, or it may be a third independent spectrum.

Only spectrum bins whose centres lie within the fiducial region are included in the fit; see Sec. 5.6 for the definition of this region. This excludes untrusted regions of the spectrum, where the reconstruction is poor or other problems are identified.

The spectrum fitting software was tested by its author, Andrei Gaponenko [48], using simulated data sets with different, known Michel parameter values. The tests showed that the spectrum fitter successfully determined the differences in decay parameters between the two simulations, and that the fit results and uncertainties had the correct statistical properties: the fit probabilities were distributed uniformly, the parameter uncertainties scaled as $1/\sqrt{N}$, etc. For each analysis round, the spectrum fitter is further validated by

¹Currently *TWIST* is using the value $\eta = -0.0036 \pm 0.0069$ from the 2005 global analysis [23] of all available muon decay data by Gagliardi et al. The analysis included the *TWIST* ρ and δ measurements published in that year.

²*nthrow* is the number of random number trials required in the spectrum generator’s acceptance-rejection method in order to produce each spectrum; the more trials required to produce a spectrum with a given number of events, the lower the overall probability that spectrum represents. See Sec. 6.1 for more information.

so-called “white-box” tests. Once the black box is opened and the measurement results revealed, new simulations are run with the decay parameters set to the measured values; when these white-box spectra were analyzed and fit against the data, the spectrum fitter must find decay parameter differences consistent with zero, as expected [48–50]. These high-statistics tests are further verification that the spectrum fitter functions correctly.

5.2 Chamber Drift Time Calculations

The time for ionization electrons to drift to the sense wire of a drift chamber cell and produce a signal depends on the location where the ionization was produced and the shape of the electric field within the cell, as well as properties and density of the drift gas. The electric field, for a given applied voltage, is defined by the geometry of the cell—the locations of the cathode foils and the position and size of the sense wire. The map of drift times within a cell is known as the Space-Time Relations, or STRs. The GARFIELD drift chamber simulation software [51] was used to calculate the STRs corresponding to a grid of points within the ~ 4 mm DC cell; the grid spacing was $20\ \mu\text{m}$. These STRs are used in the GEANT simulation (Sec. 6.2.1) and in the reconstruction (Sec. 5.3). Since the details of the drift times were not important for the PCs, approximations are used in the simulation. (No PC STRs are used by the analysis software.)

5.3 Data Analysis

The *TWIST* data analysis software is called **Mofia**, and was largely developed by *TWIST* collaborators, including the author. This software performs particle identification and track reconstruction.

In order to analyze the large amount of data and simulation produced by the *TWIST* experiment, the Mofia analysis is run on the WestGrid computing facility [52]; one set of data or simulation can be analyzed there in 1–3 days.

5.3.1 Description of Analysis

The standard production analysis is described below, and applies to analyzing both data and simulation except where noted. Variations in the analysis can be made for some specialized studies.

Data files

Data are stored in individual data files, called “runs,” of about 2 GB each; they typically contain about 800,000 events in each run, where an event is defined as the information recorded during the $6\ \mu\text{s}$ before and $10\ \mu\text{s}$ after each muon trigger. A typical data set

consists of about 300 runs, or 2.4×10^8 triggers. Simulations are similarly stored in runs, of 500,000 events each. Simulation sets typically consist of 800–1600 runs, or $4\text{--}8 \times 10^8$ triggers. Each data file is analyzed by a separate instance of Mofia, which simplifies analysis logistics and allows processing of runs in parallel. Files from simulation are analyzed in the same way as those from real data, except for the headers.

An event is composed of a collection of TDC times at which signals were received from the various chamber wires, scintillators, etc.; these times are referred to as **hits**. Times are measured relative to a copy of the muon trigger scintillator signal which has been delayed by approximately $10\text{ }\mu\text{s}$; as the analysis software unpacks each event, the actual time of the trigger is subtracted from all other TDC times, and these corrected times are used subsequently. Leading and trailing edge times for each signal are recorded, from which the signal's width (duration) can be determined. (Signal width is approximately proportional to charge collected by the tracking chamber, or to light collected from the scintillator; this can be used as a simple measure of deposited energy. Signal width is also used to identify crosstalk.) The energy deposited in the muon trigger scintillator is also measured more directly using an ADC.³

Calibration files

When Mofia initially opens a data file, calibration files are selected based on the information in the file header. (Again, calibration files include detector geometry and alignment, drift time maps, the magnetic field map, etc.) For real data, the run number is looked up in the CFM (Calibration File Manager) database, where the list of calibration files for each data run is stored. For simulation the list of calibration files used to produce the data file is stored in a data bank in the header. In either case the choice of calibration files can be overridden.

Event unpacking

Once initialization is done and calibration files are read, Mofia processes each event in the file. TDC time offsets, or t_0 's, are subtracted from all hits at this stage, to account for differences in electronics delays etc. See Sec. 4.3.2 for more details.

Crosstalk removal

If the file being analyzed is not simulation, some of the TDC signals will be the result of crosstalk—false signals an electrical pulse induces in neighbouring wires or electronics. Mofia

³More precisely, a PACT (Pulse Amplitude Charge to Time converter) system is used, where the size of the muon scintillator signal (as measured by the amount of charge from the photomultiplier tubes) is measured, and a time signal is produced whose duration is directly proportional to the measured charge. This time signal's width is measured by the standard high-precision TDCs as normal.

identifies these by their widths (crosstalk signals are very narrow) and by their proximity in time to “real” signals on nearby wires. Hits identified as crosstalk are removed.

Time windows

Each event is then divided into time segments, or “windows.” Hits in the PCs define the start of a time window of width 1050 ns; parts of the event in between these windows are defined as empty windows. Windows are numbered consecutively. For example, a simple event consisting of a muon, which triggers the event acquisition, and its subsequent decay, would have a window structure as follows:

Window 1. From $-6\ \mu\text{s}$ to the start of Window 2.

Window 2. From 50 ns before the muon’s earliest PC hit to $1\ \mu\text{s}$ after.

Window 3. From the end of Window 2 to the start of Window 4.

Window 4. From 50 ns before the decay positron’s earliest PC hit to $1\ \mu\text{s}$ after.

Window 5. From the end of Window 4 to the end of the event.

Additional particles (more muons or beam positrons, or occasionally cosmic rays) would give rise to additional windows. If there are additional PC hits within an already-defined $1\ \mu\text{s}$ time window, e.g. due to a fast decay or a coincident beam particle, a new time window is begun.

Hits are grouped into their corresponding time windows. Within a window, only the earliest hit on a given wire is kept; any others are assumed to be due to additional ionization from the same particle’s passage through the drift cell.

Time window classification

The spatial distribution of hits within each window is used to determine what type of particle or particles are in that window, and classify the window accordingly. For example, beam positrons pass through the full length of the detector, while muons stop part way through. (The classification does not assume that the muons stop in the stopping target.) Decay positrons should originate near the muon stopping position, and travel to the end of the detector. Multiple, separated clusters of hits in a sequence of several planes suggests multiple particles, and comparing distributions on either side of the muon stopping position along with other spatial properties can help distinguish between the overlap of a beam positron with a decay positron and, say, a decay positron which scatters strongly enough to reverse direction.

PC hit widths are used as well, to help identify muons; muon hits are wider than positron hits since they deposit much more energy in a drift chamber, producing a larger and wider signal.

A full list of the window types used by Mofia is given in App. A. Some window types will be more likely to occur at certain angles and momenta—those windows involving delta rays or scatters, in particular.

Event classification

Each event is classified according to the types of windows it contains and their sequence. About 40% of events are classified as “Simple Clean,” consisting of a muon window containing the event trigger, and a subsequent decay positron window, separated in time by $1\ \mu\text{s}$ or more. The “Time Clean” event type, comprising 4% of events, is similar, including one or more beam positron windows in addition to the muon and decay windows described for “Simple Clean.” The beam positrons must also be separated in time from other particles by at least $1\ \mu\text{s}$. A full list of the event types used by Mofia is given in App. B, including the list of event types included in the final analysis.

Events where time windows overlap, or in which multiple muons are identified, can be excluded from the final analysis without introducing any bias in the decay spectrum. Event types where it is known that the decay positron (if any) cannot be reliably reconstructed are excluded. Events in which the decay positron was seen to scatter in the detector or produce a delta ray were included in the spectrum, since these events are strongly dependent on the decay angle and energy. Similarly, event types for which a beam positron apparently arrives within $\sim 200\ \text{ns}$ of the decay are also kept, because these can instead be the decay positron scattering off of outside material such as the yoke and returning to pass through the detector. Again, these are strongly dependent on energy and angle and must be included in the final spectrum to avoid introducing a bias.

Pattern recognition

The hits in windows involving decay positrons are passed to the track reconstruction by the window classification. The hits in each decay window are sorted according to which side of the muon stopping position they fall on, and a first guess at the parameters of the helix or helices they make is attempted. For this first guess, clusters of hit wires are used as the data points, rather than individual hits. Positions and widths of wire clusters are computed from groups of hit wires in a UV module [49]. Hits are required simultaneously in both the U and V planes of a UV module for the pattern recognition to form a wire cluster (which means that, if one plane of a module has no hits, any hits in the other plane of that module are ignored from this point onward). Various combinations of the hits are tested to see if a helix can be drawn through the clusters; the number of wires hit per plane, which is largely correlated with the track angle, is used as a guide. The helix is purely geometrical at this stage, described in terms of the helix centre, radius, wavelength, and pitch angle. Hits

which cannot be included in a single track are set aside (reducing the sensitivity to delta particles, beam positrons, backscatters, and other sources of extraneous hits), until finally a good track is found. The pattern recognition system defines a “good track” as one which passes through all wire clusters to be included in the track. Hits associated with this track are flagged, and the remaining hits are tested again to see if additional helices can be found.

To aid in finding tracks, the pattern recognition includes hits from all PCs except those in the target module (which are most vulnerable to the dead zone problem described in Sec. 6.2.1, caused by ionization left by the muon). PC hits are not passed to the helix fitter, however; the fast drift times in the PCs are designed to give good timing information, but the position resolution is essentially the size of the wire spacing (2 mm).

A major strength of this method is its ability to sort out multiple tracks, discard hits due to delta rays, etc. One drawback is that a hard scatter by the positron can result in hits dropped from the end of a track; a track can even be “split” by this process, where enough hits are removed from the track to form a second track. (It has been shown that the helix fitter, described below, is often able to fit through large scatters like this, if it is given the hits to work with, so one improvement to be made in the future is to soften the first guess’s track fit requirements to prevent tracks like this from splitting.)

The hits and parameters associated with each helix found by first guess are passed to the helix fitter. First guess does not use the magnetic field map in any way; a perfect helix is assumed, and its helix parameters are computed in terms of curvature, wavelength, etc.

Helix fitter

The helix fitter describes the track in terms of the position and momentum of the positron near the first hit used in the track, where the “first” hit depends on the assumed direction of travel of the positron. Under normal conditions all positrons are assumed to be traveling out of the target.

The track provided by the pattern recognition is translated into the three dimensional momentum of the positron at the start of the track, as well as the three dimensional location of the track start; these, plus the time at which the particle was at the track start position, are used as the helix fit parameters. The track is refined in two passes, using essentially the same method for both. First, starting at the head of the track (the end closest to the muon stopping position), the track is projected from hit to hit. Along the way, the track parameters are modified by the magnetic field map, and by the predicted energy loss of the particle along each step’s arc length through the materials described in the detector geometry.

The general formula for the most probable energy loss⁴ assumed by the helix fitter is

$$\Delta E(p) = \frac{1}{\cos \theta} \sum_i l_i (\varepsilon_i^{ion} + \varepsilon_i^{brem}(p)) \quad (5.3)$$

where ΔE is the energy lost by the positron, as a function of positron momentum p , as it passes through various layers of material with thickness l_i (measured along the z axis). ε_i^{ion} is the energy lost per unit thickness due to ionization in the material, and $\varepsilon_i^{brem}(p)$ is the energy lost per unit thickness due to bremsstrahlung: $\varepsilon_i^{brem}(p) = k_i p$, where k_i is a constant which depends on the material (radiation length, density, etc.). The author tested this model using simulated sets of mono-energetic positrons produced at several energies [53], and it was found that the momentum-dependent component $\varepsilon_i^{brem}(p)$ is effectively unseen by the helix fitter; furthermore, it was found that Eq. (5.3) provides a good description of the energy loss if $\varepsilon_i^{brem}(p)$ is removed. This is likely because bremsstrahlung production tends to be in the form of hard, discrete interactions, which are not fit through. (Probably the pattern recognition is splitting the tracks at these points, as discussed above.) The form implemented in the helix fitter is then reduced to

$$\Delta E = \frac{1}{\cos \theta} \sum_i l_i \varepsilon_i^{ion}. \quad (5.4)$$

The track is also modified by fitted kinks, based on a method described by G. Lutz [54]. A “kink plane” is defined at the centre of each DC module (UV pair of drift chambers) after the first, as well as inside each DC dense stack, and the track is allowed to change direction at each kink plane. The scattering of a track at these points is “penalized” by an increase in the fit χ^2 (see below), with the penalty related to the theoretical scattering probability based on the amount of material since the last kink plane.

At each wire cell with a hit, the track is projected through the cell and the point of shortest drift time to the wire is found, using the STR tables (Space-Time Relations, tables of drift times as a function of position within the wire cell); this is the “calculated” hit position for that cell. In the helix fitter’s first pass, the “true” position of the hit is assumed to be the centre of the cluster of wires hit in each plane (with the U and V planes treated separately, unlike in the pattern recognition stage). The residual, the difference between calculated and “true” hit position, is then assigned based on the track angle and the number of adjacent wires hit using the “narrow windows” method [55]; this helps to resolve left-right ambiguities.⁵ In the second pass, actual drift distances are used. The track is first traced through each drift cell, using interpolations on the STR tables to find the point with the shortest drift time. This point is the “track position” for that cell. This should be

⁴As pointed out in the Particle Data Book [14], the mean energy loss is a poorly defined quantity in almost all circumstances, as it is heavily influenced by statistical fluctuations in the tails of large energy loss, as well as other complications.

⁵The same drift time is found on both sides of the wire, and it can be difficult to distinguish on which side of the wire the track passed.

compared to the position of the actual hit in that cell (the “hit position”) for a drift-distance fit. To determine the hit position, a straight line is followed from the wire through the track position until a drift time is found which corresponds to the measured drift time. (This is an approximation, which assumes the isochrones—lines of constant drift time—are circular. In reality, they are asymmetric, and distort further as they approach the edges of the cell.)

An alternative helix fit method involves using the drift times to determine residuals, rather than the drift distances; it was found that the reconstruction was more accurate with drift distance fits.

Since the TDCs can only measure relative time offsets, the global time offset is calculated by the helix fitter to determine the drift times. The track’s start time is one of the fit parameters, acting as a global offset to the drift times. (For good tracks, the track’s start time is a good measure of the muon decay time.) The time of flight of the positron along the track is also accounted for. The start time and time of flight are subtracted from the recorded TDC time to determine the drift time of the electrons within a cell.

The hit residuals are normalized by the appropriate resolution function (which, for the final drift distance fits, is taken to be a flat 100 μm throughout the cell), and the normalized hit residuals are combined with the kink residuals as per Lutz [54] to determine the χ^2 for the track. The total χ^2 is given by

$$\chi^2 = \sum_{\text{hits}} \frac{(d_f - d_m)^2}{\sigma_d^2} + \sum_{\text{kinks}} \frac{\theta_k^2}{\sigma_\theta^2} \quad (5.5)$$

where d_f is the drift distance given by the fitted hit position, d_m is the measured drift distance, σ_d is the drift distance resolution, θ_k is the fitted kink angle, and σ_θ is the width of the theoretical scattering distribution [56] at each kink plane considering the amount of material around that plane. The first sum is over all hits included in the track, and the second sum is over all kink planes used in the fit.

The helix fit parameters which minimize the χ^2 are then found using the Gauss-Newton method for nonlinear least squares problems, iterating over the above steps of projecting the track and determining drift times.

It should be noted that the use of χ^2 for the helix fit is an approximation, as the χ^2 fit method assumes symmetric Gaussian uncertainties in the data points. The drift times vary in a non-linear way in a drift cell, and since these are the basis of the hit uncertainties the assumptions of the χ^2 method are violated. The mean of the χ^2/dof distribution is not expected to be near 1, as a result. However, as bias and resolution tests have shown acceptable performance by the helix fitter, this appears to be a reasonable approximation.

Once fitting is complete, each track is assigned an error code. An error code of 0 indicates a good fit; other values indicate problems encountered during the reconstruction, such as not having enough hits available to define the track properly, too many iterations required

to converge, etc.

The momentum resolution can be measured directly from the reconstructed shape of the 52.8 MeV/c positron momentum edge of the decay spectrum, using the method outlined in Sec. 5.5.1; the momentum resolution was found to be 70 keV/c in MC, and 75 keV/c in data, with uncertainties of less than 1 keV/c. As there is no well-defined angular feature in the decay spectrum, the angular resolution must be measured using simulation; the angle resolution was found to be $\sigma(\cos \theta) = 0.002$.

Processed output

Information about each event is written to a ROOT ntuple (“tree”) [45]. For each event, information about the window types and times, helix fit results and event information, and other details are stored for further analysis.

5.3.2 Improvements in Analysis Software Since Previous *TWIST* Measurements

In the previous *TWIST* analyses of ρ and δ , the drift time maps (STRs) used to analyze the data assumed symmetric drift cells, with the wires centred. For this analysis the STRs were generated using the more realistic asymmetric geometry, which improved the resolution of the reconstruction by 5–10 keV. The helix fitter is now able to compensate for energy loss in the positron tracks, improving the fits and reducing reconstruction biases. Another significant improvement was in the way the helix fitter determines the point of closest approach between a track and the wire; this point is where the drift time is referenced, and improving the algorithm has reduced the reconstruction bias, particularly at low track angles. Also helping the reconstruction of small-angle tracks is an improvement in the first guess technique, initially developed for fitting the muon beam but adapted for low-angle positron tracks.

5.3.3 Error Found Since Completion of Analysis

Since the analysis was completed, an error was found in the window classification software (Sec. 5.3.1) which caused a bias in the reconstructed spectrum. The bias depended on $|\cos \theta|$. A test analysis with the error corrected showed that both data and simulation were affected in the same way, and the Michel parameters extracted from the data-simulation fit (Sec. 5.1) were unaffected.

5.4 Positron Spectrum Summation

The *TWIST* program `tta` (for “*TWIST* Tree Analysis”) performs the positron spectrum summation, or “tree sum.” It reads in the processed event information from the trees

produced by Mofia, selects which events and tracks to use, and fills the two-dimensional decay spectrum histogram.

The spectrum summation is computationally simple, and is run locally on computers at TRIUMF. Summing over one data set or simulation typically takes 6–12 hours.

5.4.1 Spectrum Summation Method

All ROOT trees (one per run) produced by Mofia for a given data set or simulation are processed sequentially. First each event is compared against a series of criteria, or cuts, to decide whether to use it in the final spectrum. The reconstructed decay positron tracks in the selected events are then subjected to another set of cuts and selections, until either all tracks in the event have been rejected or a single decay track has been chosen. The event cuts and track selection process are detailed in Sec. 5.4.2.

The selected decay track is then included in two $(p, \cos \theta)$ histograms: a full decay spectrum covering full momentum range, and an endpoint spectrum covering the region around $p = 52.8$ MeV/c with finer binning, for use in energy calibration.

Each event from a GEANT run includes a weight of ± 1 , based on whether the original function used to define the probability distribution was positive or negative. The weight is used when an event is added to the final decay spectrum, so that each event is added to or subtracted from the spectrum; this allows the simulation of derivative spectra. See Sec. 6.1.

The spectrum summation is done in two passes. The first, or “raw,” pass uses the positron track parameters as they are found by Mofia. Both the full decay spectrum and the endpoint spectrum are filled during this pass. The endpoint spectrum is then used by the energy calibration (Sec. 5.5). The second, or “calibrated,” pass uses the energy calibration parameters to modify the momentum of each decay track, based on its angle, before track selection is performed. No new endpoint spectrum is produced on this pass.

5.4.2 Event Cuts and Track Selection

Cuts used to select good events are described below. Histograms of the number of events rejected by each cut are shown for a standard data set and standard simulation in Fig. 5.4; some of the criteria listed in the histograms are not used and so are not described here. About 25% of data events, and 35% of simulation events, are accepted in the final decay spectra. Cuts are applied successively, so that once an event has been rejected the subsequent cuts are not tested. The event cuts, in order of application, are:

TCAP. The time between the muon trigger and the signal from the proton beam capacitive probe, used to select surface muons. See Sec. 4.1.2 for how this time is determined. (This cut is only applied to data.)

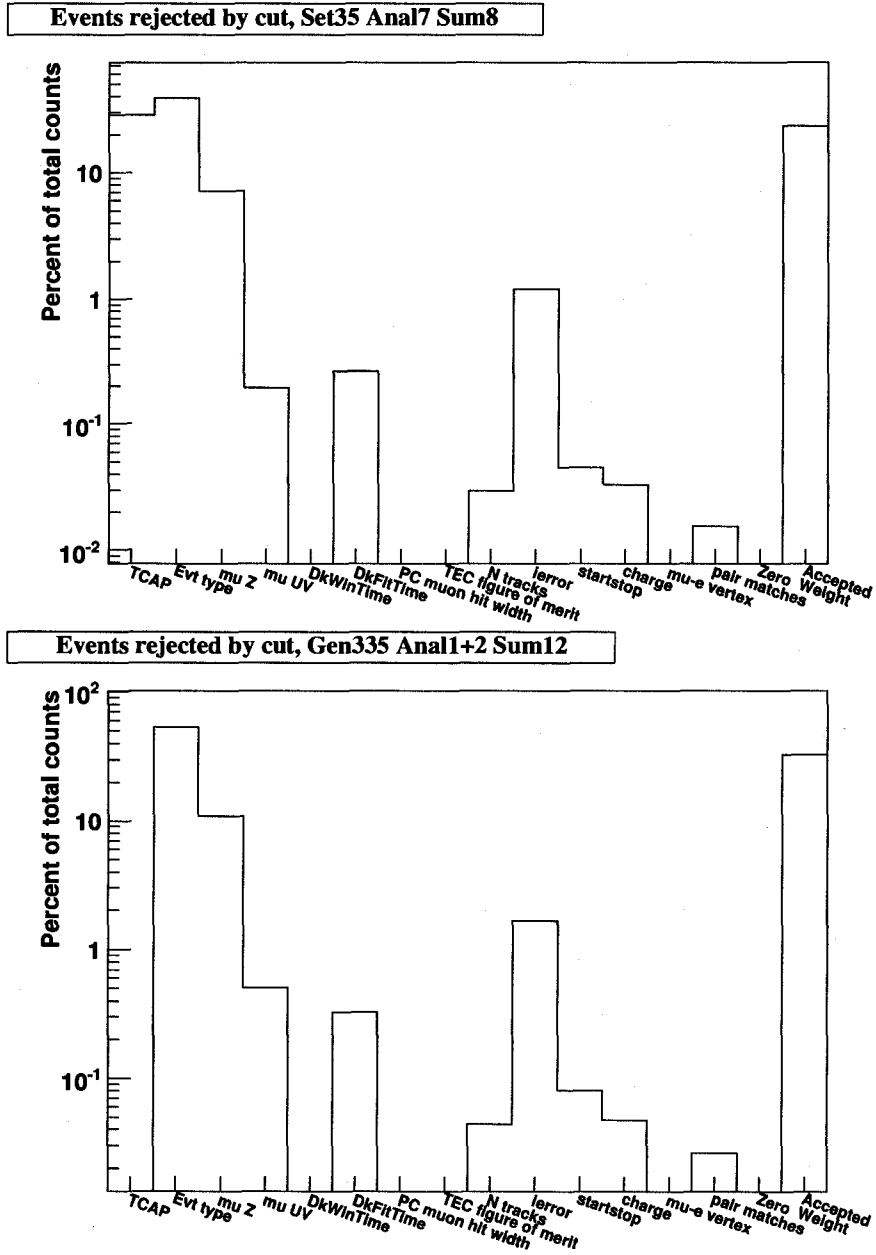


Figure 5.4: Events rejected by each cut during spectrum summation, for the standard raw treesum of data set 35 and simulation gen335. Note the excellent similarity between the distributions. (The first cut (TCAP) is not applied to MC because the cyclotron time structure is not included in the simulation of the beam.) See text for the sequence of cut application.

Event type. The following event types are accepted: 1, 2, 6, 7, 10, 11, 21, and 22. See Sec. 5.3.1 for details, and App. B for the list of event types.

Muon stopping z . The plane in which the downstream-most muon hit occurs must be PC 6, the PC immediately upstream of the stopping target.

Muon stopping (u, v) . The muon must stop within 2.5 cm of the centre of the stopping target, based on the positions of the muon hits in PCs 5 and 6. In conjunction with the cut on p_t (Sec. 5.6), this ensures that the decay positron is contained within the sensitive region of the detector and does not strike any of the support materials (glass frames etc.).

Number of tracks in decay window. Events which have no reconstructed tracks in the decay window are rejected here.

Decay time. The muon decay time must be at least $1\ \mu\text{s}$ after the muon trigger, to prevent muon hits from being confused with decay positron hits. This cut is based on the fitted track time from the helix fitter, not from the time of the decay window determined by classification as was used for previous analyses. See Sec. 5.3.1. (Note that, since events with overlapping windows are rejected in the Event Type cut, there is an implicit decay time cut of 1000 ns as well.)

Zero weight. If there was a problem during the simulation and the event was not completed correctly (see Sec. 6.2.2), GEANT sets the weight of the event (normally ± 1) to zero; these events are rejected. (This cut is tested after the track selection cuts. This cut is only applied to simulation.) Although the value of `nthrown` is not updated to reflect any unused decays, the error is small and this occurrence is rare.

Once the above event cuts are passed,⁶ selection criteria are applied to the individual tracks in the decay window. Figure 5.4 shows how many events were rejected by each track cut—i.e. from how many events each track cut rejected the last available track—for standard data and simulation. The list of track cuts is below; again, unused cuts are left out of the list.

iererror. Reconstructed decay tracks should be “good fits” according to the helix fitter’s error code.

Track start/stop. Decay tracks should not cross the target. Tracks with $\cos\theta < 0$ should be upstream of the target, and those with $\cos\theta > 0$ should be downstream.

⁶To be precise, since the decay time cut is based on the fitted time of the decay track, the track selections must be executed first. Thus, track selections are performed between the event cuts on number of tracks and on decay time.

Charge. Decay tracks should have positive charge.

Pair matches. If there are tracks on both sides of the muon stopping position, they are projected to see if they intersect at any z value between them. If they do, they are assumed to be part of the same track and both are rejected, since they must not have come from the muon.

In the rare case that an event passes all of the above cuts and there are still multiple decay tracks remaining (about 1 in 10,000 accepted events), all available decay tracks are projected back to the muon stopping z , and compared in (u, v) to the muon's stopping location. The track which projects closest to the muon is accepted.

Approximately 23% of data events and 33% of Monte Carlo events are accepted by the above cuts.

5.4.3 Improvements in Spectrum Summation Software Since Previous *TWIST* Measurements

Aside from the correction of a bug present in the initial versions of the tree sum program (which affected the first *TWIST* ρ measurement), there have been no substantial changes to the program behaviour compared to what was used for the first round of ρ and δ measurements. The most significant change is the use of the fitted track time instead of the decay window time to define the time of the muon decay. The fitted track time uses information from the entire track and determines the decay time much more precisely than the window time, which is defined as the time of the earliest of four PC hits.

5.5 Energy Calibration

As derived in Sec. 2.2.2, the maximum energy a positron can have from muon decay is 52.8 MeV. This is a sharp feature in the energy spectrum, and can be used to calibrate the measured energy of a positron.

There are a number of factors which can affect the energy calibration. Energy loss in the target module (the stopping target and the surrounding four PCs) will affect the energy calibration as well, since the positron's energy cannot be measured until the positron reaches the DCs; this effect has the largest impact on the position of the measured endpoint. To first order, energy loss depends linearly on the amount of material the particle passes through, which in the planar geometry of the detector varies with $1/\cos\theta$. Unless the average stopping z of the muons is exactly in the centre of the stopping target, the average amount of energy lost by the positrons will be different upstream and downstream, requiring separate energy calibrations in either direction. Biases in the track reconstruction arise from, among other things, the relationship of the STRs (drift time maps) used in track reconstruction and the

“true” STRs in the physical detector (or the simulation). Discrepancies between the length scales (in z , u , or v) of the physical detector and those assumed by the analysis will result in incorrectly reconstructed momenta, which the energy calibration can compensate for to some extent; this is a small effect (Sec. 8.5.2) thanks to the high-precision construction of the detector (Sec. 4.2.3). Errors in the magnetic field map used in track reconstruction will affect the endpoint as well, although this is also a small effect given the accuracy of the field map (Sec. 4.2.1). It is important to note that many of these effects, such as the stopping position or the STRs, are not necessarily precisely matched between the simulation and the real data.

If one has an analytical model for what the raw spectrum endpoint should look like, it is possible to fit that model to the measured endpoint to determine what energy calibration needs to be applied to move the endpoint to 52.8 MeV. Indeed, this is the approach used by previous *TWIST* analyses for energy calibration; see Sec. 5.5.3 for a brief description. However, this requires a good model of the endpoint shape, accounting for resolution and other effects, all with the correct angular dependence. The resulting changes in energy that need to be applied to calibrate the endpoint can be on the order of 100 keV or more.

A *relative* energy calibration technique was recently developed, and was used for the present analysis. In this approach, the endpoint of the data spectrum is compared to that of the MC spectrum, and the differences are used to calibrate one spectrum to the other. The simulation accounts for the various effects on the spectrum edge shape more accurately than an analytic model, so this technique should provide a better calibration between the two spectra. (The largest remaining vulnerability in this method is the 5 keV/c difference in momentum resolution between the simulation and the real detector, as mentioned in Sec. 5.3.1) An additional advantage to the relative energy calibration is that a smaller correction is applied, on the order of 10 keV/c instead of the ~ 100 keV/c absolute calibration; a smaller correction can tolerate a larger relative error. Furthermore, the implementation of the relative energy calibration no longer assumes that the projections of the upstream and downstream edges must match at 0, a requirement assumed by the previous absolute calibration technique.

There are two components to the energy calibration. First, the differences in positions of the endpoints of the two spectra (data and MC) are determined as a function of angle (more specifically, a function of $1/\cos\theta$, measured separately upstream and downstream). This is the energy calibration at the spectrum endpoint; the technique is described below. This must then be applied to all momenta; since the momentum dependence cannot be measured directly on the smooth momentum spectrum, some model must be assumed. The two logical extremes considered likely are a “shift” (the same *absolute* amount of energy is added at all momenta) and a “scaling” (the same *relative* amount of energy is added at all

momenta, so that the limits are no change at zero momentum and the measured change at the endpoint). Tests showed that the difference in results between the two models was negligible (see Sec. 8.2.3); a “shift” model was assumed, based on specialized comparisons of the positron interactions between simulation and data (Sec. 6.3.1).

5.5.1 Relative Energy Calibration at the Endpoint

The relative energy calibration is described in detail in a *TWIST* forum posting.⁷ As mentioned above, this method calibrates one reconstructed muon decay spectrum against another. Any two decay spectra can be calibrated under this procedure, of course; this discussion will assume that the spectra are data and simulation, for simplicity.

Before calibration, a spectrum fit is performed between the data and MC spectra, as per Sec. 5.1. This is primarily to compensate for differences in the Michel parameters between the two spectra. The spectrum fit uses the standard $(p, \cos \theta)$ range. Once the spectrum fit is complete, the data spectrum and the fitted MC spectrum are each normalized by the total counts in their respective fiducial regions (see Sec. 5.6). The relative energy calibration is then done between data and *fitted* MC, using $(p, \cos \theta)$ histograms of the endpoint region with the fit results applied. The endpoint histogram covers the full range of $\cos \theta$, and the momentum range $30 \leq p < 70$ MeV/c, much larger than is needed for the energy calibration.

The $\cos \theta$ bins of the endpoint histograms are of constant width in $1/\cos \theta$, to improve the number of bins near $1/\cos \theta = 1$. Each $\cos \theta$ bin of the endpoint histogram contains a momentum spectrum, and is examined separately. The subsection $52.3 \leq p < 53.4$ MeV/c of the data momentum spectrum is selected, containing the momentum edge. A subsection of the same size is selected from the MC spectrum, but offset from the data subset. Figure 5.5 shows an example of the momentum edges with the MC at zero offset.

A χ^2 goodness of fit measure is then calculated, based on the differences in the two subsection spectra. The selected subsection of the MC momentum is then incremented by one momentum bin and the process repeated, etc. This gives a χ^2 as a function of the MC spectrum offset in steps of the momentum bin widths; see Fig. 5.6 for an example. The offset with the minimum χ^2 is found, and a parabola is fit to a region around this point. (The region around the local minimum is approximately parabolic. In practice, the shift steps used to determine the minimum χ^2 are only about 10 bins on either side of the minimum, or 100 keV.) The position at which the parabola’s minimum occurs gives the difference in edge positions between the two spectra for this angle bin.

This procedure is repeated for each angle bin in the endpoint histogram, and the relative edge position plotted as a function of $1/\cos \theta$, since it has been found that this gives a linear relationship for a reasonable range of $1/\cos \theta$. Figure 5.7 shows an example for the

⁷Ryan Bayes, “Energy calibration status report”, *TWIST* Software internal forum, 18 June, 2007. https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_software&key=1182195189




Figure 5.5: Normalized momentum edges used for energy calibration. Shown are standard data (set35) and simulation (gen335) for the angle bin at $\cos \theta = 0.51$, with the simulation at zero offset compared to the data.

calibration of a standard data set against its corresponding Monte Carlo. Linear fits are performed for some region of the plot of endpoint difference vs $1/\cos \theta$, for positive and negative $\cos \theta$ separately; the four fit parameters (slopes A and intercepts B for upstream and downstream fits) represent the energy calibration between the given data and MC decay spectra. The range $0.5 \leq |\cos \theta| \leq 0.9$ (i.e. $1.11 \leq \frac{1}{|\cos \theta|} \leq 2.0$) was considered the “fiducial” region of the endpoint fit. This is similar to the $\cos \theta$ region used in the spectrum fit (Sec. 5.6).

5.5.2 Application of the Energy Calibration

Once the relative energy calibration has been performed, the data spectrum is re-summed, this time applying the energy calibration to each decay positron’s momentum as it is read in. For a reconstructed momentum p_{rec} , the calibrated momentum p_{cal} is given by

$$p_{cal} = p_{rec} - B_i + \frac{A_i}{|\cos \theta|} \quad \text{where } i = \text{US, DS.} \quad (5.6)$$

Again, A_i and B_i are the slopes and intercepts of the linear fits from the energy calibration.

5.5.3 Absolute Energy Calibration

Previous *TWIST* analyses used an absolute energy calibration technique instead of the one described above. In addition to determining the absolute location of the momentum edge in a given spectrum, this technique determined the width of the edge, providing a measure




Figure 5.6: χ^2 of comparative energy calibration as MC is shifted, for the calibration of standard data (set35) against simulation (gen335), for the angle bin at $\cos\theta = -0.85$ or $1/\cos\theta = -1.18$. The full range of shifts is shown in the upper plot; the ~ 100 keV/c region around the minimum is shown in the lower plot. A quadratic fit similar to the one shown is used to find the optimal MC shift.

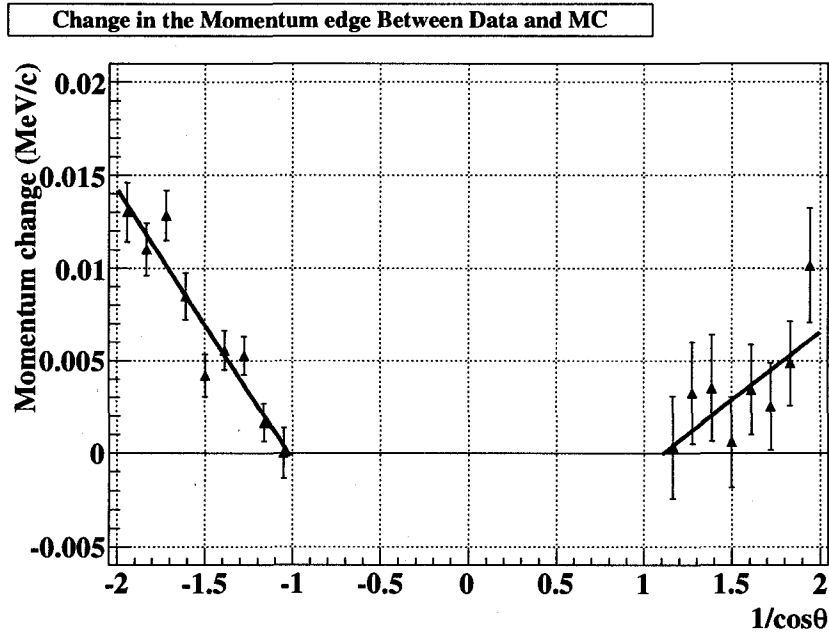


Figure 5.7: Comparative energy calibration result, for the calibration between standard data (set35) and simulation (gen335). The difference, which is of the form (MC – Data), shows that the MC edge is at a higher momentum than the data edge.

of the momentum resolution. This method is still useful for determining resolution, and so will be described here in brief.

In this method an analytic function is fit to the endpoint. The function consists of a sloped line with a step to zero at the edge, convoluted with a Gaussian curve. The slope of the line, the position of the step, and the width of the Gaussian are all fit parameters. Rather than fitting each angle bin independently, certain angular dependences are assumed; in particular, the edge position is proportional to $1/|\cos\theta|$, since this is proportional to the amount of material encountered by the positron, and the Gaussian width is approximately proportional to $1/|\sin\theta|$. The space-point resolution is essentially independent of $\cos\theta$, which makes the resolution of the measured helix radius independent of the radius. This implies a track reconstruction resolution which depends on $1/|\sin\theta|$ for constant momentum tracks. This was confirmed empirically [48]. The $1/|\sin\theta|$ approximation is only valid over a limited range of angles, which limited the range of angles which could be used for the absolute energy calibration. The momentum resolution determined by this procedure represents an extrapolation to $1/|\sin\theta| = 1$.

2005 <i>TWIST</i> Measurement [24, 25]	This Analysis
$0.50 < \cos \theta < 0.85$	$0.50 < \cos \theta < 0.92$
$p < 50.0 \text{ MeV/c}$	$p < 51.5 \text{ MeV/c}$
$p_t < 38.5 \text{ MeV/c}$	$10.0 < p_t < 39.7 \text{ MeV/c}$
$ p_z > 13.7 \text{ MeV/c}$	$ p_z > 13.7 \text{ MeV/c}$

Table 5.1: Fiducial cuts used during spectrum fitting. During the spectrum fit, only those histogram bins which satisfy all listed cuts are considered. The first column describes the fiducial region used for the ρ and δ measurements previously published by *TWIST* [24, 25]; the second column describes the fiducial region used by the current analysis. Here, p is the total decay positron momentum, and p_t and p_z are the transverse and longitudinal components.

5.6 Fiducial Region Considerations

The fiducial⁸ region of the decay spectrum is the region in $(p, \cos \theta)$ used during spectrum fitting (Sec. 5.1) to determine the difference in the muon decay parameters between the two spectra. The cuts defining the fiducial region are determined primarily by track reconstruction considerations; reconstruction of tracks with $(p, \cos \theta)$ outside this region may have significant biases, reduced resolution, or other problems.

The fiducial region should be as large as possible. The most obvious reason is that using a larger region of the decay spectrum reduces the statistical uncertainty. Increasing the fiducial region can also improve the sensitivity to changes in the decay parameters, as discussed in Sec. 5.1.

The fiducial region used for the present analysis is described in Tab. 5.1, in terms of a series of cuts on $(p, \cos \theta)$; during the spectrum fit, only those histogram bins which satisfy all the listed cuts are considered. The individual cuts are discussed below, and are compared to the fiducial cuts used for the ρ and δ measurements previously published by *TWIST* [24, 25].

It is useful to remember in the following discussion that p_z and p_t are proportional to the wavelength and radius of the helix.

Although the shape of the smooth spectrum is not strongly influenced by the momentum reconstruction resolution, the resolution defines the shape of the endpoint. One purpose of the maximum momentum cut is to exclude the endpoint region, greatly reducing the influence of the reconstruction resolution. This cut was previously set conservatively to 50.0 MeV/c. Since the resolution was found to be less than 100–300 keV/c over the full range of angles (Sec. 5.3.1), and no reconstruction problems were identified at higher momenta under the current analysis, this cut was raised to 51.5 MeV/c—still low enough to stay well back from the endpoint region used in the energy calibration ($52.3 < p < 53.4 \text{ MeV/c}$).

Tracks with very high angles (small $|\cos \theta|$) undergo large amounts of multiple scattering,

⁸From a late Latin word meaning “trust.”




Figure 5.8: Diagram of the fiducial region overlaid on a reconstructed decay spectrum. The region outlined in black shows the fiducial region for the current analysis; the region outlined in white shows the fiducial region for the previous *TWIST* ρ and δ measurements [24, 25].

interfering with track reconstruction, so a minimum $\cos \theta$ cut is required. Requiring $|\cos \theta| > 0.5$ keeps well away from the problem regions, and little is gained in statistics or sensitivity by pushing the minimum $|\cos \theta|$ cut lower.

The maximum $|\cos \theta|$ cut is limited by the reconstruction bias when fitting small-angle tracks. Recent improvements to the helix fitter (see Sec. 5.3.2) allowed the increase of this cut to $|\cos \theta| < 0.92$.

The increase in the $|\cos \theta|$ cut requires the imposition of a minimum p_t cut as well, because of the reconstruction efficiency; see Sec. 6.3.1 for details. Tracks with radii comparable to the 4 mm wire spacing cannot be reconstructed. The previous cut of $|\cos \theta| < 0.85$ avoided this region, making the minimum p_t cut unnecessary.

The maximum p_t cut eliminates large-radius positron tracks which can travel outside of the instrumented region of the detector and scatter on the support materials (glass frames etc.). The relationship between a track's transverse momentum and its radius of curvature r is given by $p_t = 300Br$ [14], where p_t is in MeV/c, B is the magnetic field in Tesla, and r is in metres. The smallest support structure is the ring supporting each cathode foil, which has an inner diameter of 34.0 cm. An event cut is applied which requires that each muon stops within 2.5 cm of the z axis (Sec. 5.4.2), so the largest helix guaranteed not to hit support material has a radius of $(17 - 2.5)/2 = 6.75$ cm. In 2002 data were taken with magnetic fields of 1.96 T and 2.04 T; the p_t cut was chosen to be consistent with these fields

as well, in case more data are to be taken in the future at either of these magnetic field strengths. Weaker magnetic fields result in larger helix radii, so $B = 1.96$ T was used to calculate the maximum p_t cut; $r = 6.75$ cm and $B = 1.96$ T give a cut of $p_t < 39.7$ MeV/c. The previously used cut of $p_t < 38.5$ MeV/c was conservative.

The detector construction is periodic in z ; the chamber modules are spaced with two alternating gap sizes, so that the base-to-base distance between chambers is 5.2 cm or 7.2 cm, resulting in a 12.4 cm periodicity and providing only two independent (u, v) points in the sparse stack⁹. This severely weakens the reconstruction of tracks with this wavelength which are integer multiples of 12.4 cm. It can be shown that $p_t = 300Br$ implies $p_z = 300B\lambda/(2\pi)$, so a wavelength of $\lambda = 12.4$ cm corresponds to $p_z = 11.8$ MeV/c. The problem is not confined to this precise wavelength, of course; tracks near this wavelength will be affected to some degree as well. Tests of reconstruction resolution on simulation suggest a cut of $|p_z| > 13.7$ MeV/c to remove this problem.¹⁰

About 16% of events accepted by the cuts in Sec. 5.4.2 are within the fiducial region described above, for either data or simulation.

⁹Recall that the “sparse stack” is composed of those chambers not in the dense stack.

¹⁰See these *TWIST* Software internal forum postings by Andrei Gaponenko for details: “Detector granularity”, 16 April, 2003.

https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_software&key=1050476690

and “Drift distance helix fitter, iteration 4: magic wavelength”, 23 September, 2003.

https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_software&key=1064375351

Chapter 6

TWIST Simulation

*In the face of entropy and nothingness, you
kind of have to pretend it's not there if you
want to keep writing good code.
—Karl Lehenbauer*

The *TWIST* experiment determines the Michel parameters by comparing the shape of a measured decay spectrum to the shape of a simulated spectrum. The simulation models the transportation of particles such as muons and positrons through the *TWIST* detector, and their interactions with the detector materials. When a muon decays, the momentum and angle of the positron is provided by an external program; this allows the simulation of Michel parameters which differ from the Standard Model.

From the start, the values of the Michel parameters used in the simulation are hidden; during analysis, the experiment measures *differences* between the data and MC values. Only when the hidden values of the MC Michel parameters are revealed are the measured Michel parameters in data known. This blind analysis technique allows *TWIST* to study systematic uncertainties and set-to-set consistency while avoiding “human bias”—the tendency to keep looking for problems when an unexpected answer is obtained, and stop once the answer is “reasonable.”

Due to the large amount of simulation required by *TWIST* (2–3 times the amount of data taken), the Westgrid computing cluster [52] was used. (Approximately 10,000 CPU-days on Westgrid for this simulation and analysis.)

Section 6.1 describes the *TWIST* blind analysis technique, and the spectrum generation software. Section 6.2 discusses the *TWIST* Monte Carlo simulation. Section 6.3 lists the methods used to validate important aspects of the simulation.

6.1 Michel Spectrum Generation and Blind Analysis

TWIST uses a stand-alone program, called `micheld`, to generate the $(p, \cos\theta)$ pairs that make up the Michel spectrum or one of the derivative spectra. (See Sec. 5.1 for an explanation of the derivative spectra; these are the derivatives of the muon decay spectrum with respect to the Michel parameters.) The `micheld` server resides on a dedicated computer,

and is communicated with remotely through a client program. The server generates the spectrum samples (collections of $(p, \cos \theta)$ pairs, typically 100,000 pairs per sample), and supplies samples to GEANT.

In order to generate a decay spectrum, some set of Michel parameters must be assumed. This is where the *TWIST* Blind Analysis is implemented. Upon request, `micheld` will randomly choose values for ρ, δ, ξ , and η within some specified tolerances, which are generally set to be about five times the current uncertainty in the parameters. The chosen Michel parameter values are stored in a database using public key encryption, and the private key needed to reveal their values is stored on a CD in a safe. In this way the Michel parameters are kept hidden, as a so-called “black box.” Since the spectrum fitting technique described in Sec. 5.1 only determines the *differences* in decay parameters between two spectra, the decay parameters in the measured data cannot be determined until the parameters used in the simulation are revealed by supplying the decryption key and “opening the box.”

The `micheld` program generates $(p, \cos \theta)$ pairs using an acceptance-rejection method, where a $(p, \cos \theta, y)$ triplet is selected at random, where y is a value between 0 and 1 with uniform probability. The value of y is compared against the decay probability (Eq. (2.9) on p. 9, with radiative corrections) for the selected $(p, \cos \theta)$; if y is too large, a new triplet is tried. The program keeps track of the number of trials (`nthrown`) used to generate the `micheld` samples, for later normalization of the GEANT spectrum generated using these samples.

In practise, y is thrown between 0 and some upper bound, chosen to be larger than any value the probability spectrum can take on. This prevents selecting a large number of y values which will never be accepted, saving a large amount of computation time. Then an effective value of `nthrown` is calculated by dividing the actual number of trials by the upper bound; this effective `nthrown` is the value reported by `micheld` and used by the spectrum fitter for normalization.

Radiative corrections are included in the decay probability function, up to $\mathcal{O}(\alpha^2 L)$, where α is the fine structure constant and $L = \log(m_\mu^2/m_e^2)$. (See Sec. 2.4 for more information.) $\mathcal{O}(\alpha^2 L^0)$ terms have been calculated [21], but were not included in the probability function used by `micheld`; the effect of neglecting these terms is evaluated in Sec. 8.7.1 and was found to be negligible. Note that all radiative decay modes are treated as radiative corrections to the $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ decay, i.e. secondary particles such as photons are not tracked by the simulation.

Essentially the same procedure is used to generate derivative spectra (Eqs. (5.1)). Since the derivatives can have negative values, the absolute values $|\partial S/\partial \rho|$ etc. are used as the probability functions. For $(p, \cos \theta)$ values where the derivative is negative, a weight of -1 is assigned to that event; otherwise the weight is $+1$. This weight is carried through the

simulation and analysis to the spectrum summation (see Sec. 5.4), so that the event using this $(p, \cos \theta)$ pair is subtracted from the spectrum rather than added. Derivative spectra are otherwise treated in exactly the same fashion as the standard simulation.

6.2 Monte Carlo Simulation

GEANT3.21 was used for the *TWIST* simulation [57]. The simulation tracks particles through the detector materials in the modeled magnetic field, applies interaction processes such as energy loss and scattering, and produces output files in the same format as the real data files. Description of the *TWIST* detector in GEANT, and local modifications to the GEANT software, were coded by Peter Gumplinger, Dennis Wright, and others.

6.2.1 Simulation Technique

The geometry and properties of the detector are described in terms of volumes and materials. Most of this is defined within the program code, but many parameters, such as chamber foil positions and stopping target geometry, are read in from the same geometry file used for analysis.

Aspects of the detector which need to be frequently modified to match data conditions (or for specialized studies), such as atmospheric pressure or temperature, fraction of CO_2 in the gas degrader, beam momentum, etc., are set in a configuration file. These are generally set to match measured conditions in the data, although in the case of the gas degrader the CO_2 fraction was set by matching the muon stopping distribution. (The fact that the simulation required a different CO_2 fraction to achieve the same stopping distribution suggests the amount of material in the muon's path is inconsistent; the amount of difference is equivalent to approximately $60 \mu\text{m}$ of Mylar, very small compared to the total amount of material in the detector.)

Particles are tracked through the magnetic field (described using the Opera 3D field simulation software; see Sec. 4.2.1 for description, and Sec. 8.2.1 for validation) using a standard fourth-order Runge-Kutta integration technique. Step sizes are determined from the amount of variation in the magnetic field, the material being passed through, volume boundaries, the occurrence of "physics processes" such as bremsstrahlung or scattering, etc.

The spin of the muon changes direction as the muon travels through a magnetic field. In the absence of scattering, the directions of the spin of polarized muons will be parallel to their momentum vectors. Scattering breaks this connection, so the simulation tracks the spin separately, integrating the Bargmann-Michel Telegdi equation for spin propagation [58] using the same technique used for tracking a particle's trajectory.

Energy loss and multiple scattering which take place at levels below certain cutoffs are treated as continuous and included in the particle tracking as random deviations in

the particle's direction and momentum. "Discrete" processes, such as secondary particle production (e.g. bremsstrahlung, delta rays, etc.) are performed at appropriate random intervals.

When a charged particle moves through a tracking chamber, GEANT randomly determines the locations along the particle track where ion clusters are produced, using a mean cluster spacing set in the configuration file. STR (Space-Time Relation) tables provide drift times corresponding to the ion cluster locations. A Gaussian smearing was applied to simulate effects of the electronics as well as diffusion in the gas; the smearing width was 1.5 ns near the wire, increasing exponentially with drift time. STR tables for the drift chambers were calculated using the GARFIELD software [51], as described in Sec. 5.2; crude hand-made tables are used for the proportional chambers, since the details of the drift times there are not important. From these cluster drift times, and the signal durations set in the configuration file, GEANT determines when the ion cluster signals overlap. The threshold number of simultaneously overlapping signals required to produce a signal (a "wire hit") in the electronics is also set in the configuration file. Details of the ion cluster spacing, threshold, and time smearing were tuned using the leading edge of the TDC time spectra (see Sec. 6.3.4, below).

In the physical detector, the passage of the muon leaves a large amount of ionization in the chamber gas, which does not clear immediately. This "deadens" that region of the chamber by masking the electric field, to the point where the ionization from a subsequent positron cannot be amplified into a measurable signal. This occurs in both the drift chambers¹ and the multi-wire proportional chambers [59]. This dead zone process is included in the simulation, with the size and duration of the dead zone matched against studies of the physical detector. The effect of the dead zone is most pronounced in the PCs around the stopping target, as the decay track starts where the muon passage ends; however, the current analysis is almost insensitive to this effect because the target PCs are no longer used in track reconstruction.

Wire hits are written out in data files using the same format as the real data acquisition produces (see Sec. 4.2.3). GEANT includes additional information in this output file, such as the list of configuration files used (see Sec. 6.2.4) and optional "truth banks," banks of information on the true information about the event (see Sec. 6.2.3).

Muon and Positron Beam Inputs

Muons and positrons are the only beam particles included in the simulation. The simulation includes pile-up of both types of particles. The beam rates for the two types of particles are specified in the configuration file. The numbers of pile-up muons and positrons in each

¹Art Olin, "dead zone modeling", TWIST Monte Carlo internal forum, 4 December, 2005.
https://twist.phys.ualberta.ca/forum/view.php?bn=twist_montecarlo&key=1133750300

event are chosen from Poisson distributions corresponding to the specified rates; the pile-up particles (if any) are then simulated at uniformly random times within the event.

The initial (x, y) position and direction of each muon are sampled from correlated position and angle distributions, which were produced from TEC measurements with the *TWIST* solenoid turned on (see section 4.1.3). These distributions varied between data sets. The initial total momentum of the muon was chosen using a uniform random distribution, which should correspond to the real distribution obtained when momentum selection is done with the combination of a bending magnet and a slit (Sec. 4.1.2). No muons were generated with initial momentum more than the cutoff of ~ 29.79 MeV/c, in keeping with the kinematics of pion decay (see Sec. 2.2.1). The centroid of the momentum distribution was set to match the estimated beam momentum of 29.4 MeV/c, and the width, or momentum bite, was set to 0.7% (see Sec. 4.1.2).

The muon beam profile was measured before the magnetic field's fringe region with the solenoid turned on (see Sec. 4.1.3); the simulated muons were started at the centre of the TEC ($z = -191.944$ cm), before the fringe of the magnetic field, and projected into the magnetic field and detector from there.

The initial positions and directions of the beam positrons are sampled from correlated distributions in a similar way to the muons. Distribution profiles for the beam positrons were measured using the main *TWIST* drift chambers with the magnetic field turned off, and projected linearly back to the TEC position. From there, they were allowed to travel into the magnetic field and detector normally. The momentum distribution, defined in the configuration file, was set to be the same as that of the muon.

6.2.2 Muon Decay Spectrum

When a muon decays, the momentum and angle of the decay positron is read from a decay spectrum sample supplied by the `micheld` server (Sec. 6.1); this is true for the primary muon and for secondary (pile-up) muons. For each GEANT run, one or more samples are read from the `micheld` server. (Five samples of 100,000 events each are used per GEANT run under standard production.) Several items of meta-data are stored in the output data banks related to the spectrum samples: the sample numbers are stored for reproducibility; the values of `nthrown`, the number of random trials used to generate the spectrum, are stored for normalization; and a weight of ± 1 is stored for each muon decay, based on whether the original function used to define the probability distribution was positive or negative (needed for derivative spectra; see Sec. 6.1). In the case where, in the last event of a run, there are more muons generated than there are $(p, \cos \theta)$ pairs remaining in the sample, a weight of 0 is assigned to that event, and it is not used in the final spectrum.

6.2.3 Optional Output Data Banks: “Truth Banks”

There are several banks of optional information available for inclusion in the simulated data files. The two most commonly used are the **vertex bank** and the **space points bank**. The vertex bank stores information about each particle’s position and momentum at up to four points: particle creation, the first drift chamber encountered, the last drift chamber encountered, and the end of the particle track (because it stopped, decayed, or left the “world” volume). The vertex bank, being relatively small, is included in all standard production simulations, although it is not often used in the analysis. The space points bank includes each particle’s position and momentum each time an ionization cluster is generated in a drift chamber (see above, Sec. 6.2.1). The space points bank is used in some specialized studies, but is not normally included due to being huge.

6.2.4 Calibration Files in Simulation

GEANT uses the same calibration files used in the analysis of data, including magnetic field map, muon and positron beam properties, drift time maps, etc. The geometry and alignment files are used to define the layout of the drift chambers, stopping target, etc, with additional information supplied within the *TWIST* GEANT code. The supplementary information defined in the code describes materials not needed by the analysis, such as support structures, beam pipe, etc. Alignment files are used rather than assuming perfectly aligned chambers in order to be as realistic as possible, and to reduce the unphysical regularity which would result from perfect alignment.

The list of calibration files used by the simulation is stored in a data bank in the output file. This data bank is read in by the analysis software so that the correct calibration files are used for the analysis.

6.2.5 Improvements in Simulation Software over Previous Analysis

Possibly the largest improvement in the *TWIST* simulation software since the previous measurements of the ρ and δ parameters is the ability to simulate asymmetric drift chamber cell geometry and STRs, so that the wire planes need not be centered between the cathode foils. In addition, improved studies of the size and duration of the dead zone left by the muon allowed it to be better simulated; and beam profiles from the TEC and other beam studies greatly improved both the muon and positron beams used in GEANT. The switch to a stopping target with precisely known geometry and materials allowed its accurate simulation, as well.

6.3 Validation of *TWIST* Simulation

Since *TWIST* determines the Michel parameters by comparing the shape of a measured decay spectrum to the shape of a simulated spectrum, validation of the simulation is vital. In order to avoid biasing the results, the validation must use methods that are independent of the Michel parameters themselves.

The validation process determines how well the *TWIST* simulation matches reality. The results of this process are used to estimate the corresponding systematic uncertainties in the final Michel parameter measurements, as detailed in Chap. 8.

6.3.1 Positron Interactions


The interactions between the decay positrons and the detector materials will influence the track reconstruction, potentially affecting the resulting Michel spectrum. The degree to which these interactions were correctly reproduced by the simulation was studied, using a specialized data set and its corresponding simulation. This study was performed by the author. The study is sensitive to interactions throughout the detector, but particularly to interactions in the target module. For muons decaying in the target, the tracking only begins after the positron has left the target, so it is especially important that the simulation correctly reproduce spectrum-distorting effects such as scattering and energy loss in the target module.

Specialized Data: Upstream Stops

A set of data (Set 33) was taken with the beam momentum lowered to about 24.6 MeV/c; at this momentum most muons stopped in the trigger scintillator or the first PCs (1–4). This is referred to as “upstream stops” data. A simulation (Gen 333) was generated to match. Decay positrons produced in the downstream direction would then pass through the entire length of the detector. The upstream and downstream halves of the positron’s path are reconstructed separately, and their momentum and track angle compared. Energy loss, scattering, helix fitter biases, and reconstruction resolution can all result in differences in the properties of the two tracks. Distributions of $\Delta p = p_{DS} - p_{US}$, $\Delta\theta = \theta_{DS} - \theta_{US}$, and other differences then allow the direct examination of positron interactions in the detector—and hence the comparison of the positron interactions in the simulation to those in the real detector—independent of the shape of the Michel spectrum.

68 runs of upstream stops data were taken in 2004, with 1.48×10^6 events accepted after event and track selection cuts were applied. 1,000 runs of matching simulation were generated, with 1.27×10^7 events accepted. Figure 6.1 shows the $(p, \cos\theta)$ distributions of events accepted, for data and simulation. Criteria required to accept the event include:

US Spectrum (after cuts), Data



US Spectrum (after cuts), MC

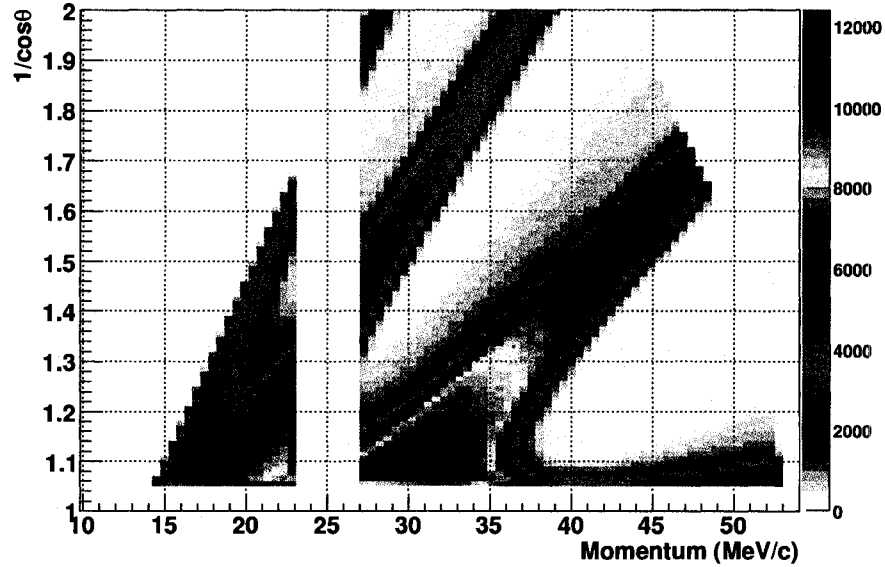


Figure 6.1: Distribution of events accepted in upstream stops studies, for data (top) and simulation (bottom). As analysis of upstream stops data is typically in terms of $1/\cos\theta$ instead of $\cos\theta$, the distributions are plotted in those terms. The band structure is due to the requirement that the positron pass through aluminum disk of the stopping target, i.e. within 4 cm of the z axis at $z = 0$. Events with momentum within 2 MeV/c of 25 MeV/c were removed; this eliminated the beam positrons, which would otherwise skew the measurement.

- The event must include some type of decay or beam positron window. (Positrons from muon decays this far upstream often appear to be beam positrons to the window classification.)
- If the muon entered the detector it must have stopped in PCs 1–4.
- The decay or beam positron window must occur at least $1\ \mu\text{s}$ after the trigger (eliminating beam positron trigger events).
- There must be exactly one good track reconstructed upstream of the stopping target, and one good track downstream.
- To ensure the positrons pass through the stopping target foil (rather than through the various masks and other structures (Fig. 4.6), which may not be as well simulated and are not of particular interest), the tracks are projected to $z = 0$ and the resulting (u, v) positions must be within 4 cm of the centre of the target.
- Fiducial cuts, as described in Sec. 5.6, are applied based on the upstream track. (No fiducial requirements are made of the downstream track.) Note that, for histograms which are functions of momentum only, no momentum cuts are applied, and similar for angle-based histograms.
- The momentum of the upstream track must be below 23 MeV/c or above 27 MeV/c. This eliminates beam positrons; since they have a narrow phase space and large statistics they tend to skew the analysis.

Integrated Energy Loss and Scattering Distributions

Due to the planar geometry of the detector, the amount of material a particle passes through and hence the amount of energy loss it experiences is proportional to $1/\cos\theta$. This means that any differences in the angular distributions of the simulation and the data (caused by, for example, different polarizations or different values of ξ) will affect the integrated energy loss distribution. This dependence on the angular distribution can be eliminated by studying the distribution of $\Delta p \cos\theta$ rather than Δp . The $\Delta p \cos\theta$ distributions for data and simulation are shown in Fig. 6.2.

Scattering is expected to be less important to the measurement of the Michel parameters, since its effects should be symmetric upstream and downstream, so simple $\Delta\theta$ distributions were used to study scattering. The $\Delta\theta$ distributions for data and simulation are shown in Fig. 6.3.

To estimate the peak positions and widths of the various distributions, a function con-

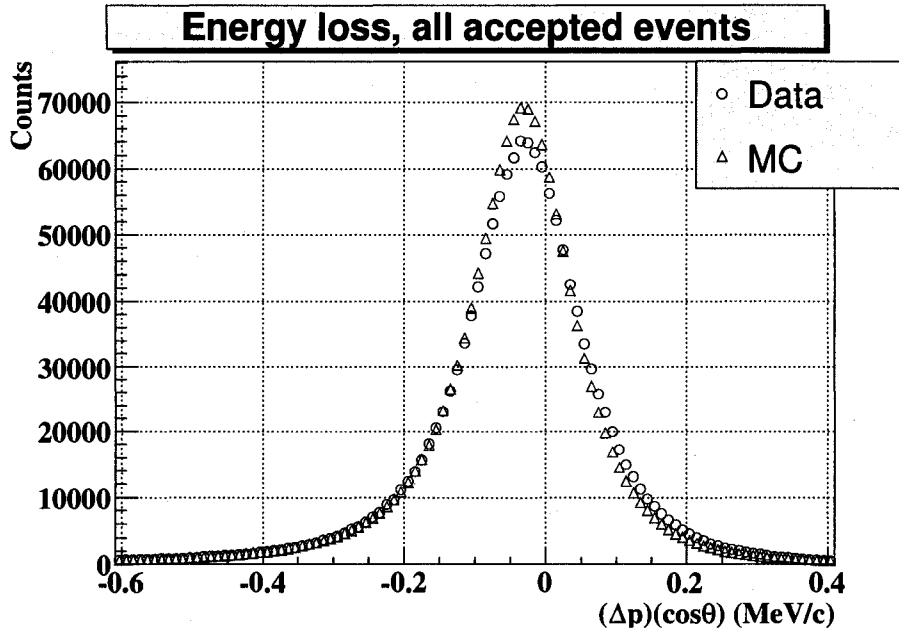


Figure 6.2: Integrated energy loss distribution, near the peak energy loss region, for data and MC, for all events shown in Fig. 6.1. The histogram from simulation has been normalized to the data. Energy loss is multiplied by $\cos\theta$ to account for first-order angle dependence, reducing the dependence of the distribution on the phase space of the accepted events. Error bars are shown but are vanishingly small.

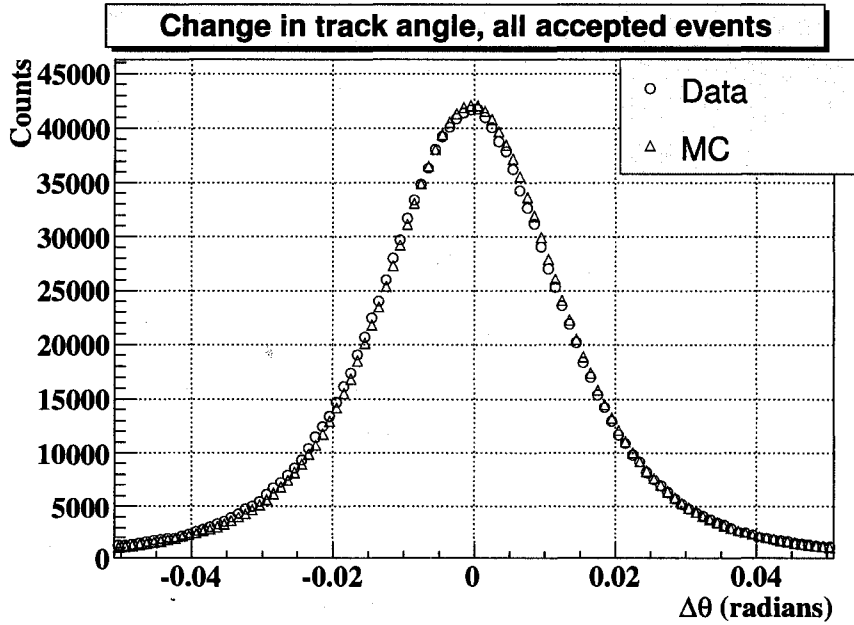


Figure 6.3: Integrated scattering distribution, near the peak scattering region, for data and MC, for all events shown in Fig. 6.1. The histogram from simulation has been normalized to the data. Error bars are shown but are vanishingly small.

	Energy Loss	
	Data	MC
Peak	-28.4 ± 0.1 keV	-29.65 ± 0.04 keV
FWHM	155.9 ± 0.1 keV	141.64 ± 0.04 keV

	Scattering	
	Data	MC
Peak	-0.97 ± 0.02 mrad	-0.581 ± 0.007 mrad
FWHM	29.75 ± 0.02 mrad	29.159 ± 0.007 mrad

Table 6.1: Peak positions and widths (full width at half maximum) of the energy loss and scattering distributions shown in Figs. 6.2 and 6.3.

sisting of a Gaussian peak with asymmetric exponential tails was used:

$$f(x) = \begin{cases} Ne^{\lambda_-^2/2} e^{\lambda_-(x-\bar{x})/\sigma} & (x - \bar{x}) < -\sigma\lambda_- \\ Ne^{-(x-\bar{x})^2/2\sigma^2} & -\sigma\lambda_- \leq (x - \bar{x}) \leq \sigma\lambda_+ \\ Ne^{\lambda_+^2/2} e^{-\lambda_+(x-\bar{x})/\sigma} & (x - \bar{x}) > \sigma\lambda_+ \end{cases} \quad (6.1)$$

where N is an overall normalization, λ_- and λ_+ control the shapes of the left and right exponential tails, σ is the Gaussian width parameter, and \bar{x} is the peak position. All of these are fitted parameters. The only parameter of interest is \bar{x} . The width of the distribution is estimated by measuring the full width at half maximum (FWHM) of the fit function (to prevent fluctuations in FWHM due to statistical variations of individual bins). The peak positions and widths of the $\Delta p \cos \theta$ and $\Delta \theta$ distributions are listed in Tab. 6.1. The simulation reproduces the scattering very well, to within half a milliradian. The energy loss distribution shows some slight difference: the most probable energy loss is about 1 keV larger in simulation, and the width is roughly 10% narrower.

The “hard” interactions—events where the energy or angle of the positron changes by a large amount in the target—are very well reproduced by the simulation. Figs. 6.4 and 6.5 show the distributions of Δp and $\Delta \theta$ on logarithmic scales; the simulated distribution matches that measured from data over several orders of magnitude. The energy loss distribution is well reproduced by the simulation for losses as large as 10 MeV/c. The scattering distribution is reproduced over the full range of possible $\Delta \theta$. (Almost no events were seen with $\Delta \theta \gtrsim 1$ rad, which corresponds to $\cos \theta = 0.54$.) The range of $\Delta \theta$, and the asymmetry between positive and negative values, is also matched by the simulation. The asymmetry is likely due to phase space; the fiducial cuts limit the size of θ_{US} , so that there are generally more ways to scatter to larger angles than to smaller angles. This may be related to the non-centered $\Delta \theta$ peaks in Tab. 6.1, but since the peak offsets were very small this was not investigated.

A quantitative measure of the simulation of hard interactions can be made by comparing the number of events in the “tails” of the Δp and $\Delta \theta$ distributions to the total number of events in each, giving a measure of the probability of a hard interaction occurring; see Tab. 6.2. The Δp tail was defined as those events with $(\Delta p - \overline{\Delta p}) < -1$ MeV/c, where $\overline{\Delta p}$ is

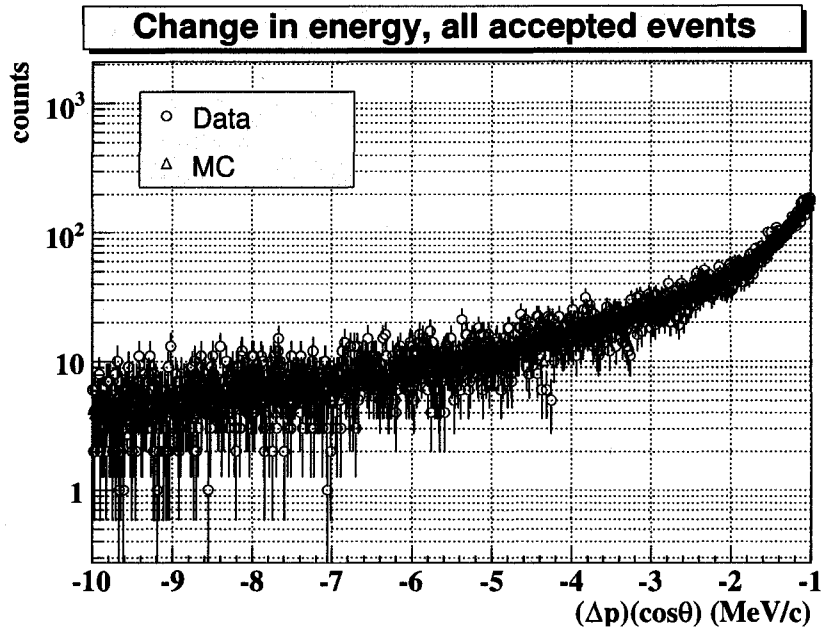


Figure 6.4: Integrated energy loss distribution, for large energy loss events, for data and MC, for all events shown in Fig. 6.1. The histogram from simulation has been normalized to the data, with the same normalization as in Fig. 6.2. Energy loss is multiplied by $\cos\theta$ to account for first-order angle dependence, reducing the dependence of the distribution on the phase space of the accepted events.

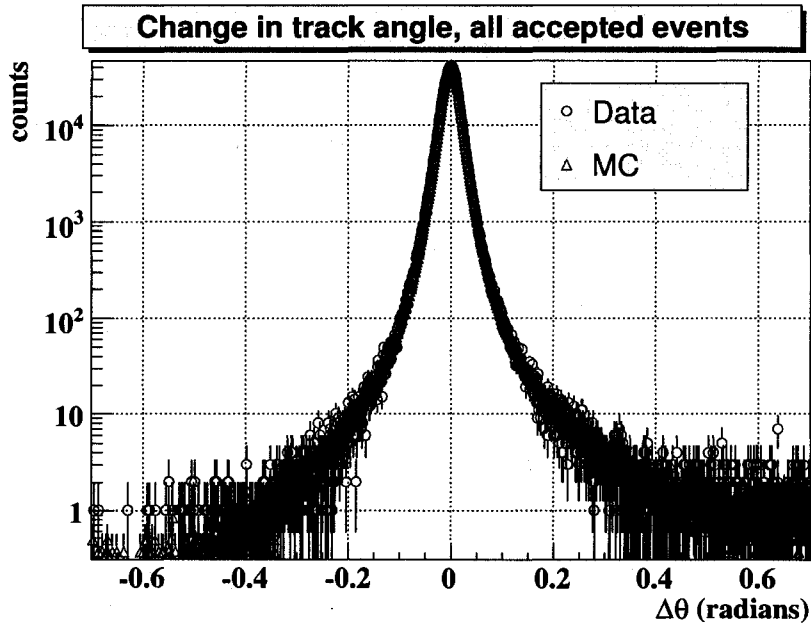


Figure 6.5: Integrated scattering distribution, for large scattering events, for data and MC, for all events shown in Fig. 6.1. The histogram from simulation has been normalized to the data, with the same normalization as in Fig. 6.3.

	Energy Loss	
	Data	MC
Tail/Total	0.0142 ± 0.0001	0.01420 ± 0.00004
	Scattering	
	Data	MC
Tail/Total	0.00513 ± 0.00003	0.00497 ± 0.00007

Table 6.2: Hard interaction probabilities, in data and simulation. These are calculated as the ratio of the number of counts in the distribution’s tail(s) to the total counts in the distribution. See text for details.

the position of the peak of the Δp distribution. The cutoff was chosen somewhat arbitrarily, with the requirement that it be well away from the peak, to reduce any confusion between effects of actual physical processes and of resolution moving events into or out of the tail. The agreement between data and MC was consistent for several tested values of the cutoff. Note that the tail is defined relative to the position of the peak, so that a shift in the peak will not translate into a false change in tail counts. Since events with $\Delta p > 0$ were due to resolution and other reconstruction effects (physically a positron cannot gain energy as it crosses the detector), only the negative tail was examined. The $\Delta\theta$ tail was defined as $|\Delta\theta - \overline{\Delta\theta}| > 0.1$ rad, where $\overline{\Delta\theta}$ is the position of the peak of the $\Delta\theta$ distribution, this time looking at the combined sum of the negative and positive tails. In both cases, varying the cutoff points defining the tails changed the measured tail probabilities, but did not affect the comparison between data and simulation. As can be seen from Tab. 6.2, the relative tail counts in the simulation are an excellent match to those measured in data.

Energy and Angle Dependence of Energy Loss

Energy loss depends linearly on the amount of material the particle passes through, which in the planar geometry of *TWIST* varies with $1/\cos\theta$. The energy calibration (Sec. 5.5) provides a direct comparison of the energy loss behaviour in simulation to that in data, but only at the spectrum endpoint. The upstream stops data allow the energy loss behaviour to be compared at other momenta away from the edge. Figure 6.6 shows the peak energy loss as a function of $1/\cos\theta$ for several 5 MeV/c bins of momentum. (Other values of p were missing too much data to be useful, because of the effect of requiring that all positrons cross the target within 4 cm of the z axis to ensure they passed through aluminum; see Fig. 6.1.) As the plots show, the simulation reproduces the peak energy loss seen in data to within less than 10 keV/c, usually much less, in all bins shown. The results of linear fits to the plots are listed in Tab. 6.3, showing that there is very little momentum dependence, if any, to the $1/\cos\theta$ behaviour of the peak energy loss. This result motivated the momentum-independent application of the energy calibration to the bulk of the spectrum (Sec. 5.5). The constant terms in the linear fits were expected to be zero. The fact that they are not

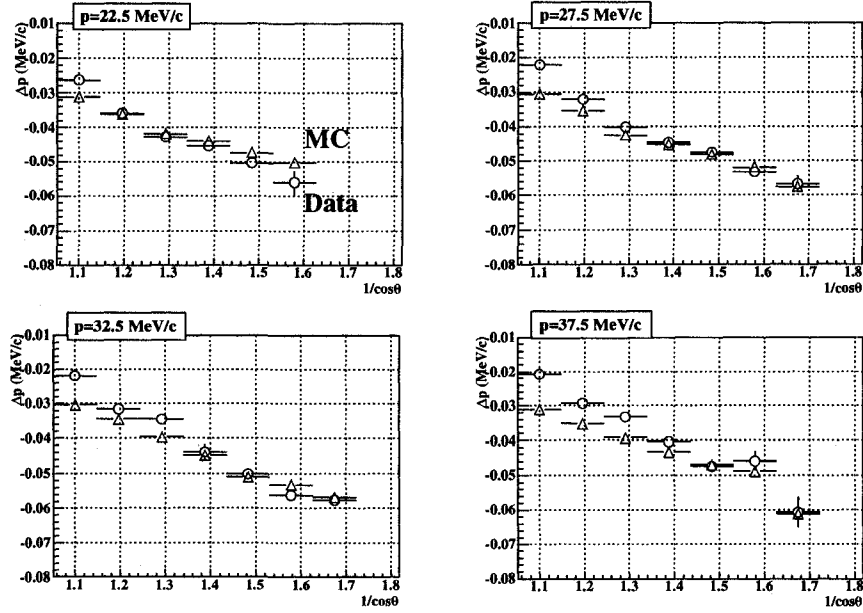


Figure 6.6: Peak energy loss as a function of $1/\cos\theta$ for several 5 MeV/c bins of momentum.

p (MeV/c)	Data		Simulation	
	Constant (keV/c)	Slope (keV/c)	Constant (keV/c)	Slope (keV/c)
22.5	43 ± 4	-64 ± 3	16 ± 1	-43.1 ± 0.9
27.5	47 ± 4	-65 ± 3	18 ± 1	-45 ± 1
32.5	56 ± 3	-71 ± 2	23.9 ± 0.8	-49.2 ± 0.7
37.5	45 ± 4	-61 ± 3	14 ± 2	-41 ± 1

Table 6.3: Linear fits to peak energy loss vs. $1/\cos\theta$ for bins of momentum, over the range $1.05 < 1/\cos\theta < 1.62$; the last bin tended to deviate significantly from linearity, if it was available at all, and so was not included in the fits. The constant value is the projection to $1/\cos\theta = 0$.

may indicate higher-order behaviour of the energy loss beyond a simple $1/\cos\theta$ dependence; it may also be related to the varying occupancies of the bins used in Fig. 6.1, or to biases in the track reconstruction which treat tracks traveling towards the target differently from tracks leaving the target. Note that differences in slopes between simulation and data can change the values of the constant terms, as well.

6.3.2 Reconstruction Efficiency

The upstream stops data described in Sec. 6.3.1 were also used to study the reconstruction efficiency in data in comparison to that from simulation, as a function of $(p, \cos\theta)$. The method to measure e.g. the downstream reconstruction efficiency is to ask whether a good track is reconstructed downstream, given a good track upstream at some $(p, \cos\theta)$ value; upstream reconstruction efficiency is defined similarly. Only events with hits in both the far upstream and far downstream PCs (1–4 and 9–12) were considered, ensuring that the positron crossed the entire detector. To reduce sensitivity to the effects of scattering, energy loss, reconstruction biases, and fiducial edge effects the only requirement that was placed on the nature of the downstream track was that it is a good track according to the helix fitter. This is a somewhat oversimplified definition of reconstruction efficiency, but it is sufficient for comparing data to simulation.

In practice the distribution of reconstruction efficiency is determined using the ratio of two $(p, \cos\theta)$ histograms. To study downstream efficiency, a histogram is made of all good *upstream* reconstructed tracks for which any good downstream track was reconstructed (r), and this is divided by the histogram of *all* good upstream tracks (T). Again, upstream reconstruction efficiency is defined similarly. If $\epsilon \equiv r/T$ is the reconstruction efficiency in a given $(p, \cos\theta)$ bin, then its uncertainty is given by

$$\sigma_\epsilon = \frac{1}{T} \sqrt{\frac{r(T-r)}{T}}. \quad (6.2)$$

Figures 6.7 and 6.8 show the reconstruction inefficiency ($1-\epsilon$, to make variations easier to see on a logarithmic scale) for data and simulation, for upstream and downstream respectively.

(Note that the statistics available in the data are significantly lower than that in the simulation.) There is a feature in the upstream inefficiency plots at the beam momentum (24.6 MeV/c) caused by beam positron contamination: beam positrons are nearly parallel to the z axis and are often not reconstructed upstream, but they can scatter (particularly in the target) and gain enough transverse momentum that they can be reconstructed downstream. This appears as an inefficiency upstream. This feature is reproduced in the simulation, although in a different momentum bin since the beam momentum was not set exactly the same.

To see the distributions more clearly, projections of the fiducial regions onto the p and




Figure 6.7: Upstream reconstruction inefficiency $(1 - \epsilon)$ as a function of $(p, \cos \theta)$, for data (top) and simulation. The black outline shows the fiducial region.




Figure 6.8: Downstream reconstruction inefficiency ($1 - \epsilon$) as a function of $(p, \cos\theta)$, for data (top) and simulation. The black outline shows the fiducial region.

	Upstream	Downstream
Data	$(5.7 \pm 0.2) \times 10^{-4}$	$(8.2 \pm 0.2) \times 10^{-4}$
MC	$(5.12 \pm 0.05) \times 10^{-4}$	$(7.87 \pm 0.06) \times 10^{-4}$
(MC - Data)	$(0.6 \pm 0.2) \times 10^{-4}$	$(0.3 \pm 0.2) \times 10^{-4}$

Table 6.4: Integrated inefficiencies over the fiducial regions, for data and simulation.

$\cos\theta$ axes were made. (This was done by projecting the fiducial regions of the original histograms onto these axes before calculating efficiencies.) See Fig. 6.9 for the upstream projections; downstream projections are similar. (The beam momentum feature is visible in the momentum projections, of course, and shows that the beam momentum in the simulation was in fact slightly lower than in data.) The differences between reconstruction inefficiencies in data and simulation are shown in Fig. 6.10, with p and $\cos\theta$ projections shown in Fig. 6.11; again, upstream inefficiency results are shown, and downstream differences are similar. The reconstruction efficiency is consistently higher in the simulation, but by less than 1×10^{-4} ; linear fits to the p and $\cos\theta$ projections showed that the difference is constant over the fiducial region.

The difference in efficiencies between downstream and upstream introduces an asymmetry that could affect the measurement of the Michel parameters δ and $P_\mu\xi$, if the difference is not reproduced in simulation. Figure 6.12 shows the downstream-upstream differences in efficiency for data and simulation, and Fig. 6.13 shows the p and $\cos\theta$ projections. Although an asymmetry in the reconstruction efficiency exists, it is well reproduced by the simulation over the fiducial region, and is therefore a feature of reconstruction technique and not a problem in the simulation.

The integrated inefficiencies over the entire US and DS fiducial regions are listed in Tab. 6.4, along with the differences between data and simulation. There is a suggestion of an inconsistency between data and simulation, particularly upstream; see Sec. 8.1.5 for an assessment of possible implications for the measurement of the Michel parameters.

6.3.3 Outside Material

Although material outside of the active volume of the detector, such as the beam pipe and related objects, do not directly influence positron track reconstruction or other aspects of the analysis, a particle striking such material can scatter back into the active volume or generate secondary particles which do so. These leave additional hits, or even entire tracks, which may weaken the reconstruction. Some material, such as the steel yoke, is not simulated at all, and the geometry and other properties of the materials which are simulated, such as the beam pipe, may have errors.

A simple test of the accuracy of the simulation of this outside material is to count the number of times a particle (either the decay positron or a secondary particle) backscatters

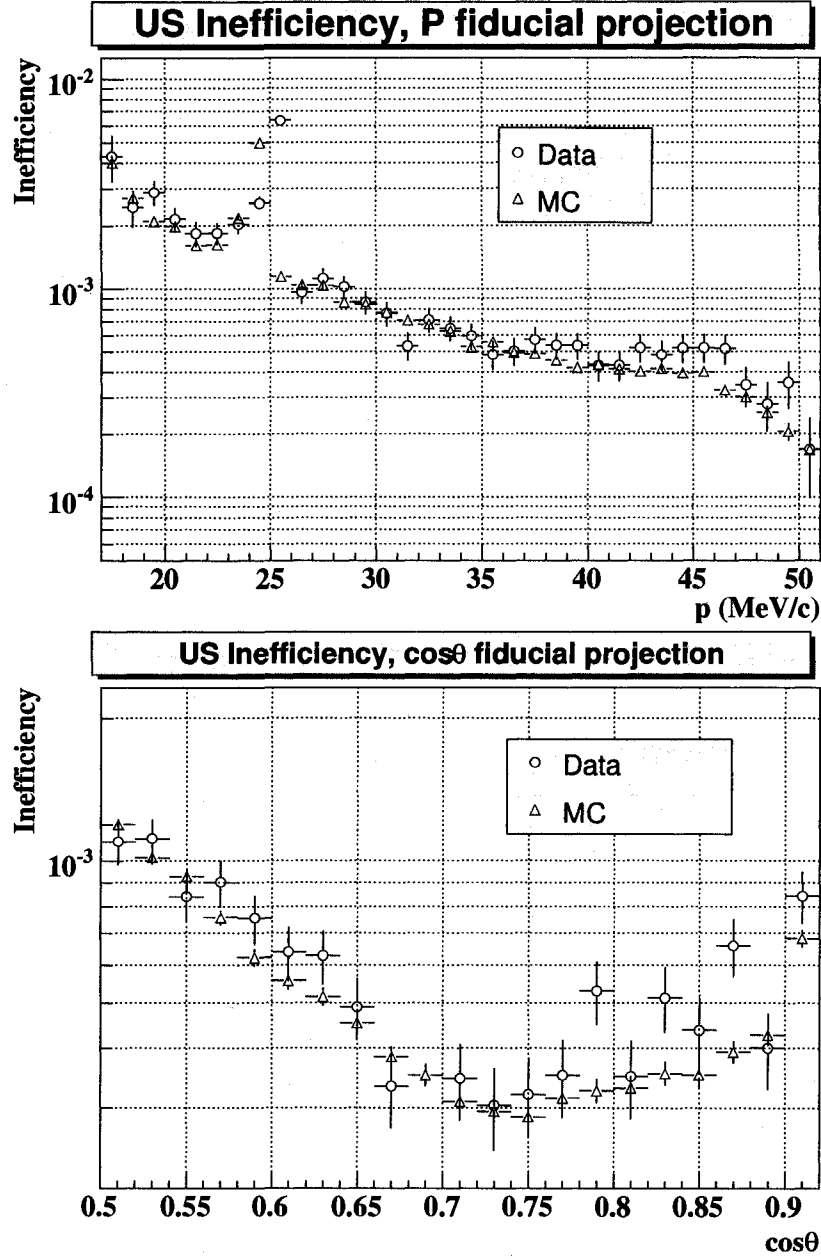


Figure 6.9: Projected upstream inefficiencies as functions of p (top) and of $\cos\theta$, for data and simulation. Each bin represents an average inefficiency across the fiducial region.




Figure 6.10: Difference between data and MC inefficiencies, as a function of $(p, \cos \theta)$. A negative value means the simulation is less inefficient (more efficient) than the data. Yellow represents differences $> +0.002$, and white represents differences < -0.002 . The beam momenta were not exactly the same in data and simulation, so the resulting features in their inefficiency measurements do not line up, leaving the yellow and white stripe seen here near $p = 25$ MeV/c.

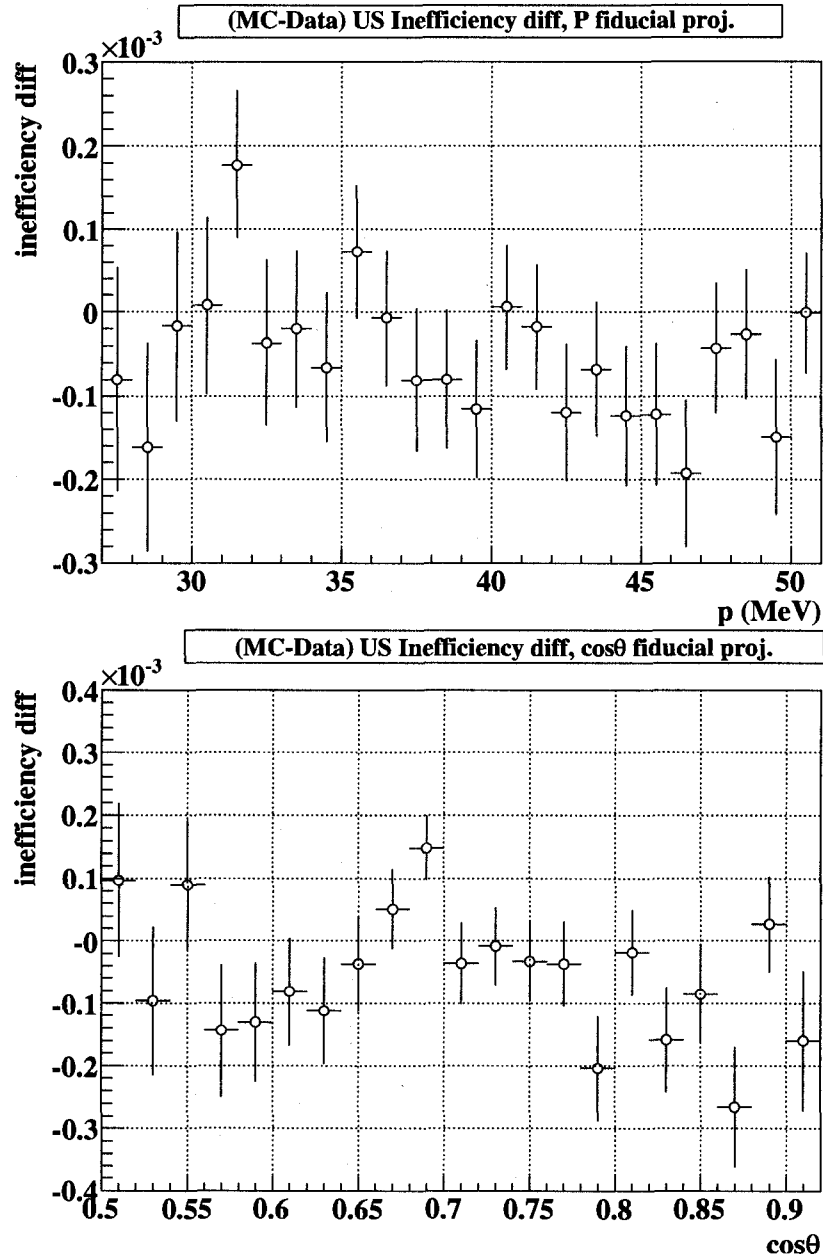



Figure 6.11: Projected differences between data and MC inefficiencies, as functions of p and $\cos\theta$, for the fiducial region. A negative value means the simulation is less inefficient (more efficient) than the data.

(DS-US) Difference in Inefficiency, Data



(DS-US) Difference in Inefficiency, MC




Figure 6.12: Difference between downstream and upstream reconstruction inefficiencies, for data (top) and simulation. The black outline shows the fiducial region. Negative values mean the downstream reconstruction inefficiency is lower, or the efficiency is higher. The stripe near $p = 25$ MeV/c is an artifact of the beam positrons (see text).

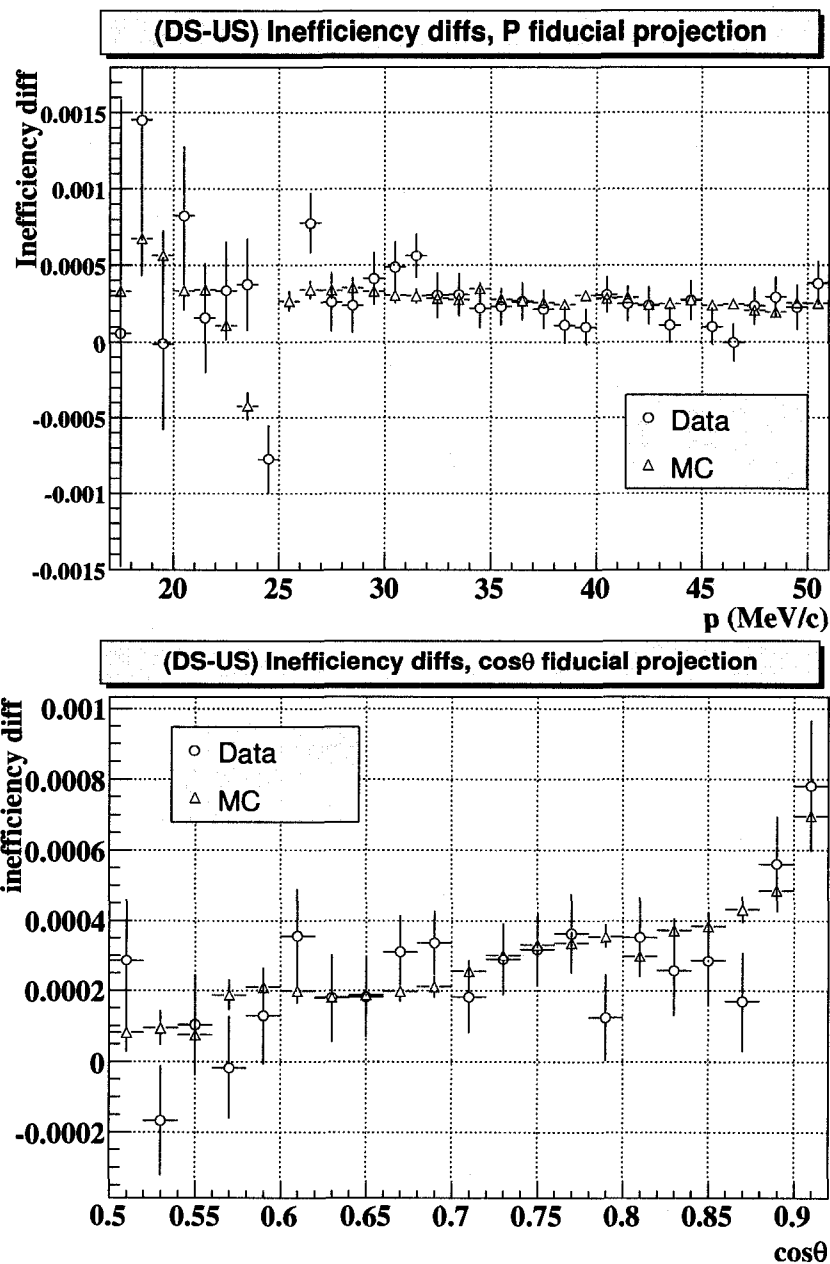


Figure 6.13: Projected differences between downstream and upstream reconstruction inefficiencies as functions of p (top) and $\cos\theta$, for the fiducial regions of data and simulation. Negative values mean the downstream reconstruction inefficiency is lower, or the efficiency is higher.

	Upstream	Downstream
MC	$(1.641 \pm 0.001) \times 10^6$	—
Data	$(1.566 \pm 0.002) \times 10^6$	$(0.178 \pm 0.001) \times 10^6$
(MC – Data)	$(0.075 \pm 0.002) \times 10^6$	

Table 6.5: Normalized backscatter counts for data and simulation, with background removed. The MC PC TOF histogram for downstream decays was used for the background.

into the detector. Backscatters were counted using “PC Time-of-Flight” or “PC TOF” histograms. For each event containing a window of types 7, 8, 14, or 15, and in which there were hits in both the upstream and downstream PCs (1–4 and 9–12), the difference in the average hit times between the two PC groups gave a “time of flight” measure for whatever particle traversed the detector. Window types 7 and 8 were upstream or downstream decay positron windows for which the classification identified a delta ray, and types 14 and 15 were upstream or downstream decay positron windows for which the classification identified a beam positron overlapping the decay; a backscatter could be mistaken for either of these, depending on its angle etc. PC TOF histograms for data and simulation are shown in Fig. 6.14. Each histogram is normalized using the number of muon windows identified by the classification (window type 1). Each PC TOF histogram is a combination of an asymmetric bell-shaped background plus a roughly Gaussian peak at ± 15 ns; the peak contains the backscatters. Since there are no materials simulated downstream of the detector in the Monte Carlo (the only material that could be there is the yoke, which is not simulated), the MC downstream backscatter histogram was used as the background, and subtracted from the other PC TOF histograms. The PC TOF histograms with background removed are shown in Fig. 6.15. The slight undershoot visible in the data after background removal suggests the background was somewhat overestimated, but the effect is much smaller than the peak sizes.

To count backscatters, each normalized PC TOF histogram (*before* background removal) was integrated, and the integral of the background histogram was subtracted. (This was done rather than integrating the histograms after background removal to simplify the error calculations.) The results are shown in Tab. 6.5. The largest discrepancy is due to backscatters from decay positrons striking the steel yoke, which is not simulated. Backscatters from the upstream material are reproduced correctly to about 5%. The impact of the mismatches on the reconstructed Michel parameters is examined in Sec. 8.3.2.

6.3.4 TDC Spectra

The details of the drift chamber signal simulation (ion cluster spacing, signal threshold, time smearing) were chosen based on theoretical considerations² and tuned to match data

²Vladimir Selivanov, “Re: Tuning of geant ionization clustering to data”, TWIST General internal forum, 4 April, 2005.

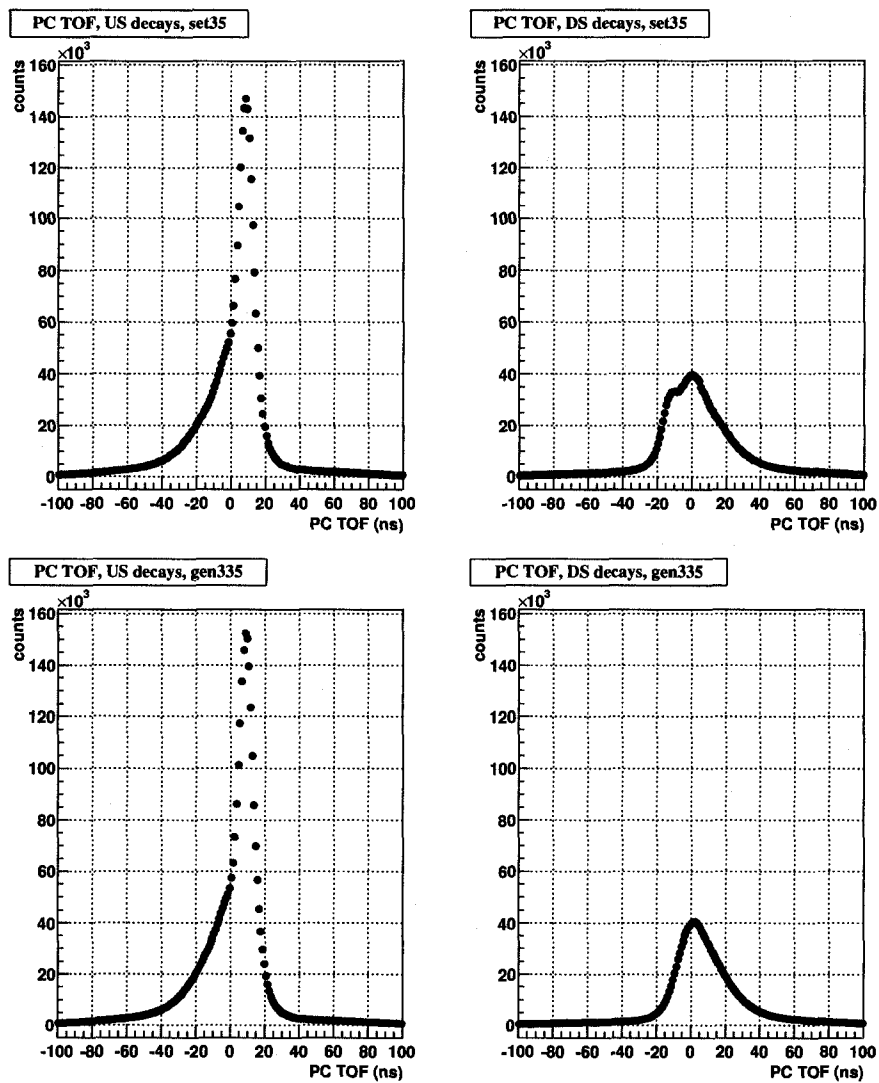


Figure 6.14: PC Time-of-Flight histograms for data (top) and simulation. Upstream decays (left) are from window types 7 and 14; downstream decays (right) are from window types 8 and 15.

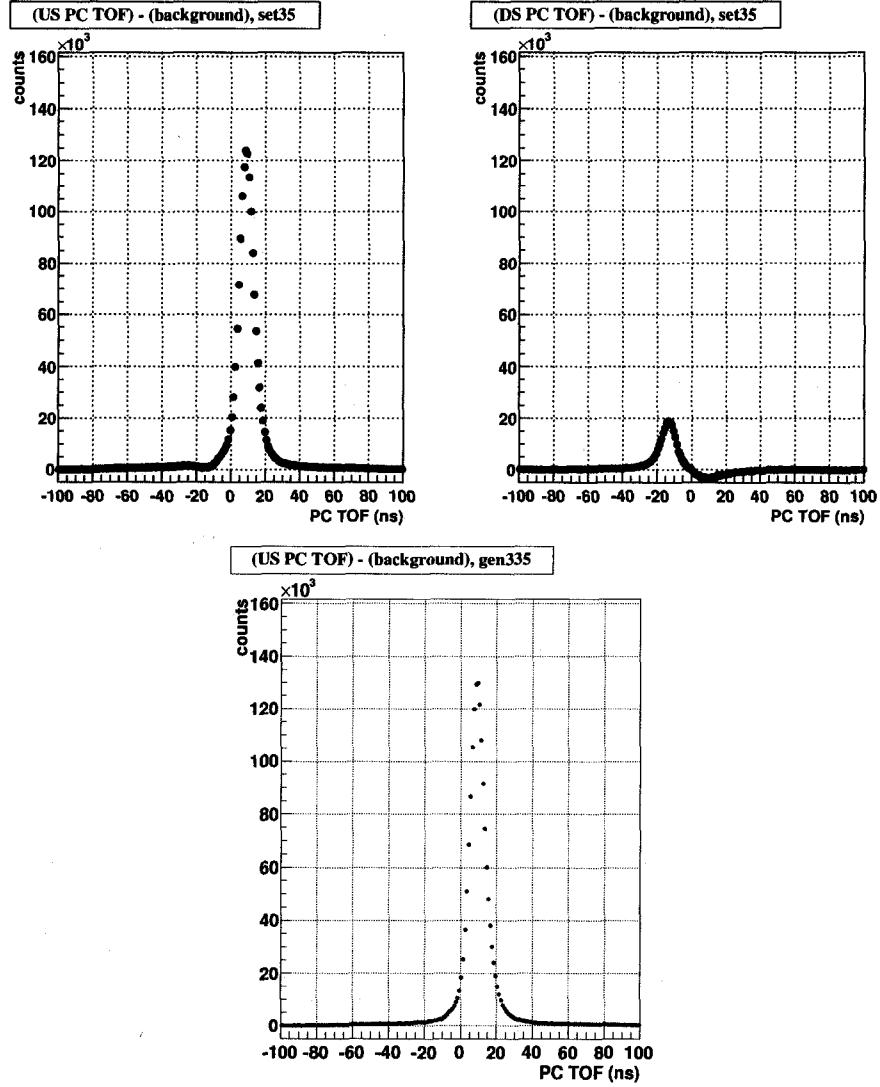


Figure 6.15: PC Time-of-Flight histograms for data and simulation, with background removed. The MC PC TOF histogram for downstream decays was used for the background.




Figure 6.16: Leading edge of TDC signal time distributions for data (closed squares) and simulation (open circles). Distributions are normalized to total counts. Error bars are shown, but are vanishingly small. Figure by Jingliang Hu.

using the leading edge of the TDC signal time distribution³ after t_0 correction. Figure 6.16 shows the final agreement between data and simulation. The simulation reproduces the TDC distribution seen in data fairly well, although it is slightly wider. The simulation of the ion clustering and related signal processes, along with the STRs, can affect the chamber efficiency, reconstruction resolution, etc.

https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_general&key=1049482467

³Konstantin Olchanski, "Re: Tuning of geant ionization clustering to data", TWIST General internal forum, 25 May, 2005.

https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_general&key=1053913884

Chapter 7

TWIST Data

We are never in a position to say what really is or what really happens, but we can only say what will be observed in any concrete individual case.

—Erwin Schrödinger, “The Fundamental Idea of Wave Mechanics”

The data studied in this analysis were taken during the *TWIST* 2004 run period, from 15 October through 7 December. This included a number of tests and studies, some with the solenoid turned off and some with it on, as well as the seven data sets used for the measurement of the Michel parameters; these data sets totalled nearly 4 TB of data, including over 1.5×10^9 muon decays, nearly twice the amount of data used in the previous ρ and δ measurements. The data sets, which are summarized in Tab. 7.1, were taken under a variety of conditions (beam steering, rate, muon stopping position, etc.), with the philosophy that, provided the simulation reproduces these conditions correctly, the Michel parameters extracted from the data should be independent of the run conditions. Data sets are numbered chronologically.

This is the same data analyzed by Blair Jamieson et al. for the measurement of the Michel parameter $P_\mu\xi$, published in 2005 [50]. However, the $P_\mu\xi$ analysis did not include the study of the ρ or δ systematic uncertainties and corrections. Furthermore, the present measurement uses more advanced analysis and simulation software than was available for the 2005 analysis. Most importantly, a new, independent set of hidden Michel parameters was generated for use in the simulation (see Sec. 6.1), and this new set was used exclusively for this analysis; thus, measurements and systematics could not be compared to the previous analysis or to the Standard Model until the the black box was opened. Thus, blindness is preserved for the present measurement.

In addition to the standard data sets, a number of specialized studies were undertaken. In particular, in Set 33 the muons were stopped in the trigger scintillator or in PCs 1–4 for the purpose of studying the simulation of the detector response and the details of the positron interactions; see Sec. 6.3.1 for the results of this study. Set 33 contained 68 runs

Set	Description	Runs	Events ($\times 10^6$)	Fiducial Events ($\times 10^6$)
31	B2 @ 949 G, centred stops	249	209	8
32	B2 @ 949 G, PC5 stops	118	94	2
35	centred stops	342	287	11
36	3/4 stops	384	323	12
37	high rate	245	198	7
38	aperture in beam	301	263	9
39	3/4 stops	182	157	6
Total:		1821	1531	55

Table 7.1: Description of data sets used for Michel parameter measurements. Each “run” is a 2 GB data file. Approximately 4% of events pass both the track selection and fiducial cuts to be used in the final analysis (see Secs. 5.4 and 5.6); these events are listed in the last column. Sets 33 and 34 are specialized studies; Set 33 is described in Sec. 6.3.1, but Set 34 contained too little data to be useful.

of 2 GB each, providing just under 10^7 fiducial events after cuts. Another set, 34, was taken with the muons stopped mainly in a plastic disk placed downstream of the detector, to mirror Set 33 and check for biases etc. This was a very small set, however, with only 11 runs, most of which were not useful events. To allow the muons to reach the downstream plastic disk, it was necessary to increase the beam momentum to 36 MeV/c. As only cloud muons¹ are available at this momentum, the muon rate was exceedingly low, with heavy beam positron contamination; moreover, range straggling² meant that many of the muons were stopping well inside the detector, and were not useful for this study. Thus, Set 34 was not used.

A description of the individual data sets used for this measurement follows. Data quality monitors and consistency checks are discussed as well.

7.1 Description of Data Sets

An overview of the primary data sets and the conditions under which they were taken is provided in Tab. 7.1. Many of the condition changes were for the purposes of studying their effects on the muon polarization. Unless otherwise specified, the beamline elements were set according to the values in Tab. 7.2.

For **Set 31**, the beam was steered off axis by increasing the magnetic field in the B2 dipole by 5×10^{-4} T to 0.0949 T. The muon stopping distribution was adjusted to set the peak of the distribution in the centre of the stopping target (see Sec. 4.1.2).

Set 32 was similar to Set 31, but the gas degrader was used to move the stopping distribution well upstream, so that more muons were stopping in PC5. This increased the number of accepted events where muons stopped in the PC6 gas instead of the target, in

¹Cloud muons are muons from pions decaying in flight; see Sec. 4.1.2.

²Since energy loss is a stochastic process, the range of a stopping particle has some distribution, which grows wider as the momentum of the particle increases. This process is known as range straggling.

Magnet	Current (A)	Field (T)
B1	39.8	0.087137
B2	43.8	0.094400
Q1	149.8	
Q2	57.4	
Q3	113.3	
Q4	89.5	
Q5	113.4	
Q6	52.1	
Q7	59.6	

Slit/Jaw	Width (mm)	Position (mm)
F0 Hori. Slit	60.0	0.0
F0 Vert. Jaws	18.0	—
F1 Hori. Slit	10.0	0.0
F2 Vert. Jaws	20.0	—

Table 7.2: Standard M13 beamline element settings used for 2004 data; see Fig. 4.2 on page 22 for element descriptions. Magnet currents are read from a direct current transformer (DCT); in addition the dipole magnets B1 and B2 were equipped with NMR probes for precise magnetic field readings. Slit and jaw settings were precise to ± 0.2 mm. Slits can change both width and position, while jaws can only change width.

order to test the effect on the polarization.

Set 35 was a nominal data set, with all beam settings at standard values and the stopping distribution centred in the target. **Sets 36 and 39** also used the standard beam settings, but the amount of CO_2 in the gas degrader was reduced to move the stopping distribution 3/4 of the way through the target.

For **Set 37** the horizontal slit and vertical jaws at F0, near the production target, were widened to 24.0 ± 0.1 mm and 80.0 ± 0.1 mm, respectively, increasing the beam rate to about 5000 muons per second (and changing the beam profile in the process).

Set 38 was taken using the same “high-rate” beamline settings as Set 37, but a 1 cm diameter aluminum aperture was inserted at $z = -177$ cm (on the downstream face of the TEC enclosure) to reduce the radial size of the beam. This also reduced the beam rate to about 1200 muons per second, roughly half of the standard rate.

Since this was not the first time these data sets have been analyzed [27], analysis numbers were assigned; this was necessary to track subsequent analyses for systematic studies as well. For example, the standard analysis of Set 31 was Analysis 2, while the standard analysis of Set 35 was Analysis 7.

7.1.1 Matching Simulation to Data

A separate simulation was run for each data set, taking care to match all conditions (rate, temperature, pressure, etc.) under which the original data was taken. Simulation labels are of the form “Gen 3xx”, where xx is the corresponding data set number (e.g. Gen 335 corresponds to Set 35). Each simulation used an appropriate muon beam profile, measured using the TEC (Sec. 4.1.3). The same STR (drift time map) file was used for all simulations, after determining the variations in gas density between sets was small enough to permit this.

Muons in the simulation do not experience as much energy loss travelling through the

Set	P/T (Pa/K)	μ rate (Hz)	μ Stop \bar{z} (cm)	Foil bulge (μm)
31	335.8 [1.9]	2540 [60]	-0.890 [0.042]	-50 [55]
32	334.9 [0.4]	2560 [70]	-2.184 [0.062]	-75 [45]
35	335.0 [0.9]	2740 [200]	-0.745 [0.034]	-30 [45]
36	332.1 [1.2]	2790 [90]	-0.300 [0.024]	-25 [55]
37	334.1 [0.5]	5970 [210]	-0.287 [0.031]	-25 [50]
38	332.2 [2.3]	1400 [50]	-0.301 [0.029]	-20 [65]
39	332.9 [0.8]	3340 [60]	-0.293 [0.026]	-15 [40]

Table 7.3: Mean values of several experimental variables, with RMS variations in square brackets. P/T is proportional to the atmospheric density. Muon rate is measured as the rate of M scintillator signals. For the muon stopping position, $\bar{z} = -0.7$ cm represents a distribution centred in the target, and $\bar{z} = -0.3$ cm stopping at a depth of 3/4 of the total stopping target thickness. A negative foil bulge means foils are deflected towards the wires; foil bulge uncertainty is about 25 μm .

detector as they do in the real data; the difference corresponds to apparently “missing” material equivalent to approximately 63 μm of Mylar. This is believed to be due to errors in the GEANT integration of energy loss in thin films, of which the *TWIST* detector has many; the thickness of the trigger scintillator and its wrappings could also be different from what is set in GEANT. Thus, rather than setting the He:CO₂ mixture in the gas absorber in the simulation to match that of data, it was adjusted until the stopping distribution as seen in the tracking chambers matched that of the corresponding data set. This meant a significantly different helium fraction in the simulation; for example, 63% He was used for Set 35, while Gen 335 used just 33% He.

Each simulation was analyzed in two sub-groups. For the first 500 runs of each set, the vertex bank (see Sec. 6.2.3) was included in the ROOT trees produced by Mofia, to facilitate future systematic studies. (In practise, only the vertex banks in Gen 335 were used; future simulations will likely not include vertex banks for all gens.) The sub-groups were combined for analysis that did not involve the vertex bank.

Significantly more events were simulated than were included in the corresponding data sets, typically by factors of 2–6, such that the simulation did not contribute significantly to the statistical uncertainty in the measurements of the Michel parameters. (As will be shown, however, the uncertainty in the final measurements is entirely dominated by systematic uncertainties, not statistical uncertainties.)

7.2 Data Quality

7.2.1 Online Monitoring and Regulation

Many aspects of the experiment were monitored and recorded in a database during the course of the data taking. Table 7.3 shows the values and variations of several important experimental variables. The steering of the proton beam on the production target, the

settings and readings of the M13 beamline elements, the muon trigger rate, pressure and temperature inside and outside the detector, and other variables were recorded. The magnetic field strengths in the B1 and B2 dipole magnets were monitored using NMR probes, and the dipoles' currents adjusted automatically to maintain the fields at set strengths; this kept the fields stable to better than $\pm 0.01 \times 10^{-4}$ Tesla.

The muon stopping position was estimated using the mean z value of all muon stops within 12 cm of the target, as measured by the nearby DCs and PCs. (A muon stopping within a given DC or PC was said to have stopped at the centre of that chamber, for the purpose of this calculation.) $\bar{z} = -0.7$ cm corresponded to a centred stopping distribution, while for $\bar{z} = -0.3$ cm the most probable stopping position is at a depth of 3/4 of the total stopping target thickness, according to Monte Carlo simulations. Once set, the stopping position was held stable by automatic adjustments to the He:CO₂ mixture in the gas degrader. (The regulator system was implemented early in the taking of Set 35, and was not available for earlier sets.)

The drift times in the DCs depend on the shape of the electric field within the chamber, which is governed by the positions of the cathode foils. To monitor the amount of foil bulge, the mean drift times were compared between two regions of the drift chambers: $r_A < 5.0$ cm, and $9.5 < r_B < 12.5$ cm. Drift times were combined from all upstream DCs in the sparse stack to increase statistical precision. The difference $\Delta t = \bar{t}_B - \bar{t}_A$ is a measure of the amount of foil bulge; a positive Δt signifies an inward foil bulge. Data were taken with the foils bulged by known amounts (as calculated from differential pressure, foil tensions and chamber capacities) for calibration.

7.2.2 Consistency of Measurements

As mentioned previously, the differences in Michel parameters between data and simulation should be consistent between data sets, independent of the conditions under which the data were taken. This increases confidence that the simulation is correctly accounting for these variations between sets. Figure 7.1 shows the measured values of $\Delta\rho$, $\Delta\delta$, and $\Delta P_\mu\xi$ for each set, with constants fit to obtain weighted averages. The χ^2/dof and confidence level values of the constant fits are reasonable for statistically independent measurements, confirming that the Michel parameter measurements will be consistent among the data sets. Note that the true values of δ and ξ are needed to convert the fit parameters from the $(\xi|\delta, \xi\delta)$ notation used in the fit into the standard (ξ, δ) notation, as described in Sec. 5.1. For the blind fits performed here, Standard Model values are assumed for the conversion. (Since the hidden values of these parameters were restricted to be close to the Standard Model values (Sec. 6.1), this did not affect the evaluation of systematic uncertainties in Chap. 8.)




Figure 7.1: Fitted values of $\Delta\rho$, $\Delta\delta$, and $\Delta P_\mu\xi$ for each data set. A constant is fit for each parameter to obtain the weighted average fit result. The χ^2/dof and confidence level ("Prob") values for the constant fits show that the Michel parameter measurements are consistent between data sets.

Chapter 8

Corrections and Systematic Uncertainties

Awwugh!
—Charles Schultz, “Charlie Brown”

The precision of the *TWIST* experiment is limited by its systematic uncertainties rather than by statistics. Once a possible source of error is identified and its uncertainty determined, its impact on the Michel parameter measurement must be assessed. The relationship between the error σ_q in some parameter q of the experiment (e.g. the simulation of energy loss, the magnetic field map, etc.) and the resulting uncertainty in a Michel parameter can be expressed as

$$\sigma_\lambda = \frac{\partial \lambda}{\partial q} \sigma_q \quad (8.1)$$

where $\lambda \in \{\rho, \delta, \xi\}$ represents the Michel parameters. This description assumes that the q parameter being studied is not correlated to other parameters of the experiment, which is generally true here. If the effect on the Michel parameters is linear in the variation of q , Eq. (8.1) can be approximated as

$$\sigma_\lambda = \frac{\Delta \lambda}{\Delta q} \sigma_q \quad (8.2)$$

which leads to the method *TWIST* uses to estimate most systematic uncertainties. In this method the source of error is exaggerated in some way by an amount Δq , usually by modifying the simulation or by making changes to the input files in the analysis; the scale factor by which the uncertainty has been exaggerated is $\Delta q / \sigma_q$. For example, for the study of one systematic uncertainty, simulations were run with the production rates of various secondary particles increased (Sec. 8.3.1), to examine the effects of errors in these rates in the standard simulation; for another study, a standard simulation was reanalyzed using a modified magnetic field map (Sec. 8.2.1). Each exaggerated study produces a decay spectrum which can be fit against standard simulation in the usual way, to measure directly the impact the exaggeration has on the measured Michel parameters, $\Delta \lambda$. The scale factors used for these studies need to be as large as possible, so that the effect is well measured, while being small enough that the analysis software is not impacted in an unrealistic way.

Finally this Michel parameter measurement is scaled by the exaggeration factor, giving the systematic uncertainty in the Michel parameters due to the source of error in question as expressed in Eq. (8.2).

In a few cases the uncertainty in the systematic measurement is larger than the sensitivity itself. In these cases the central value of the measurement was still used in the same way as for other systematic uncertainties. The contributions in these cases were small, with the exception of the measurement of the systematic uncertainties due to errors in the drift time maps (Sec. 8.1.1); in that case, however, since the systematic measurement resulted in a correction, the uncertainty in the correction is the correct systematic uncertainty to use.

The statistical power of the systematic studies can be improved by reanalyzing the same simulation under new conditions; since the same events are being considered in both cases, the resulting spectra can be highly correlated. The correlation is reflected in the χ^2/dof of the Michel parameter fit. The uncertainties from the fit can then be reduced by a factor of

$$R = \sqrt{\chi^2/dof}. \quad (8.3)$$

All systematic uncertainties were assumed to be common to all data sets. In some cases there is some variation between sets, but the contributions of these variations to the total systematic uncertainty is negligible.

Improvements to the *TWIST* detector, software, and methodology have significantly reduced the systematic uncertainties since the previous measurements of ρ [24] and δ [25]. The methodology for measuring many of the systematic uncertainties has been improved, and much greater statistics are available for the studies. The precision stopping target geometry now in use (a 71 μm high-purity aluminum foil, replacing a 125 μm Mylar foil spray-coated with $\sim 10 \mu\text{m}$ of graphite) has reduced the positron interaction uncertainties (Sec. 8.3.1), and eliminated the systematic uncertainty due to the target thickness (Sec. 8.8). More precise and detailed testing of the spectrometer alignment procedures have reduced the systematic uncertainty from alignment to almost nothing (Sec. 8.5.1). Online monitoring of properties of the experiment such as differential pressures, and increased instrumentation of the drift chambers, have reduced one source of chamber response uncertainty, and the development of a technique to determine chamber drift times directly from data has allowed a more direct evaluation of the remaining uncertainty (Sec. 8.1.1). New energy calibration techniques have been developed, which both reduced the associated systematic uncertainty and improved our understanding of it (Sec. 8.2.2).

Leading systematics for the present measurement include the chamber drift times (Sec. 8.1.1), energy calibration (Sec. 8.2.2), and the simulated delta ray production rate (Sec. 8.3.1).

Table 8.1 summarizes the current systematic uncertainties of the *TWIST* experiment, organized by categories; the remainder of this chapter discusses how the systematic uncertainties were evaluated. Each section discusses one of the systematics categories. In most

Category	ρ	δ
Chamber Response (8.1)	2.94	5.22
Energy Scale (8.2)	2.86	4.07
Positron Interactions (8.3)	1.60	0.93
Resolution (8.4)	1.22	1.37
Alignment and Lengths (8.5)	0.30	0.29
Beam Intensity (8.6)	0.06	0.15
External Uncertainties (8.7)	1.16	0.15
Total	4.57	6.71

Table 8.1: Summary of *TWIST* systematic uncertainties, in units of 10^{-4} . Sections detailing each systematic category are given in parentheses.

Correction	ρ	δ
Drift time maps (8.1.1)	-1.99 ± 2.92	$+1.61 \pm 5.17$
Momentum resolution (8.4.1)	$+1.22 \pm 0.24$	$+1.31 \pm 0.26$
Total	-0.77 ± 3.16	$+2.92 \pm 5.43$

Table 8.2: List of corrections to final Michel parameter measurement, in units of 10^{-4} . Sections detailing each correction are given in parentheses. Uncertainties in the corrections are included in the systematic uncertainties in Tab. 8.1.

cases there are separate subsections devoted to the individual systematic uncertainties, but in some cases multiple uncertainties were evaluated in the same study, so their descriptions are combined. All studies of systematic uncertainties described here were designed and performed by the author, except where noted.

In some cases, what is under study is a systematic error, rather than an uncertainty—something is being done incorrectly by a known amount (with a known sign). Measuring the effects of these errors on the final Michel parameters allows corrections to be applied to account for this. In most cases the error was small enough that it was simply treated as a systematic uncertainty, but errors due to the drift time maps (Sec. 8.1.1) and the simulated momentum resolution (Sec. 8.4) were significant and were treated as corrections. These are listed in Tab. 8.2.

Section 8.8 briefly describes systematic uncertainties which appeared in prior *TWIST* measurements of ρ and δ , but which are not included here.

Note that, although the present work is concerned specifically with the decay parameters ρ and δ , these systematic uncertainty studies generally apply to the $P_\mu\xi$ parameter as well. Since each study’s uncertainties are determined for all three parameters simultaneously, the $P_\mu\xi$ results are recorded here for future reference, though they are not included in the summary table (Tab. 8.1).

8.1 Chamber Response

The Chamber Response category includes effects related to the way the detector itself reacts to the passage of charged particles. The systematic uncertainties due to these effects are summarized in Tab. 8.3. The time-independent inaccuracies in the STRs used by the analysis also led to a correction to be applied to the measured Michel parameters.

Chamber Response	ρ	δ
STRs (time-independent) (8.1.1)	2.9	5.2
STRs (gas-density variations) (8.1.2)	0.07	0.12
STRs (foil bulge variations) (8.1.3)	0.09	0.15
Foil Bulge (geometric effects) (8.1.4)	0.12	0.25
Efficiency Asymmetry (8.1.5)	0.11	0.60
Wire Time Offsets (t_0 's) (8.1.6)	0.22	0.24
Total	2.94	5.22

Table 8.3: Summary of systematic uncertainties related to chamber response, in units of 10^{-4} .

8.1.1 Chamber Drift Times (STRs): Time-Independent Component

For this simulation and analysis, the same STRs were used for all planes of the detector, and for all sets. The systematic uncertainty associated with this can be described with separate time-independent and time-dependent components.

To the degree that the Garfield-generated STRs do not perfectly describe the true STRs in the physical detector, the data was analyzed with the “wrong” STRs. Due to effects such as ion clusterization, and perhaps helix fitter biases etc., the Garfield-generated STRs are not really the correct STRs to use for analyzing Monte Carlo spectra either. The relationship between the STRs used for analysis and the STRs inherent to the drift chambers is different between data and MC, possibly introducing different biases.

Jingliang Hu and Alexander Grossheim have recently developed¹ a method for directly measuring STRs from data (or from MC). This is an iterative method where the measured drift times are averaged for bins of hit positions in the drift cell (as reported by the helix fitter; see Sec. 5.3.1), and the average drift time at each point is used to correct the original STR map. This is repeated until the corrections converge to zero. These measured STRs include all of the analysis effects as well as the original chamber STRs themselves, and are the “correct” STRs to use for analysis (using data-driven STRs for analyzing data, and MC-driven STRs for analyzing MC). However, since these measured STRs were not available at the time of production analysis, the Garfield STRs were used.

¹See for example Alexander Grossheim, “News on data-driven STRs”, TWIST Chambers internal forum, 5 December 2006.
https://twist.phys.ualberta.ca/forum/view.php?bn=twist_chambers&key=1165358242




Figure 8.1: STRs generated by Garfield software. The top plot uses a linear contour scale; the bottom plot uses a logarithmic contour scale. Contour units are in nanoseconds. The wire is at $(0,0)$, and the cathode foils are at $z = -0.185$ cm and $z = +0.215$ cm.

For the following study, an early version of the measured STRs was used. These were “global” STRs, where the same STRs were used for all planes. Their accuracy was sufficient for this study.

Measured vs Garfield STRs

The STRs produced by the Garfield drift chamber simulation software are shown in Fig. 8.1, and those measured from data and MC are shown in Fig. 8.2. Figure 8.3 shows the changes between measured and Garfield STRs. The changes can be quite significant, especially in the corner regions of the cell. However, as Fig. 8.4 shows, most of the hits used by the analysis software are near the centre of the cell.




Figure 8.2: STRs measured from data (left), and from Monte Carlo (right). The structures near $z = \pm 0.15$ cm in data are artifacts of the smoothing procedure, and are in regions of the drift cell where no hits are used (see Fig. 8.4); they do not affect the reconstruction. The top plots use a linear contour scale; the bottom plots use a logarithmic contour scale. Contour units are in nanoseconds.

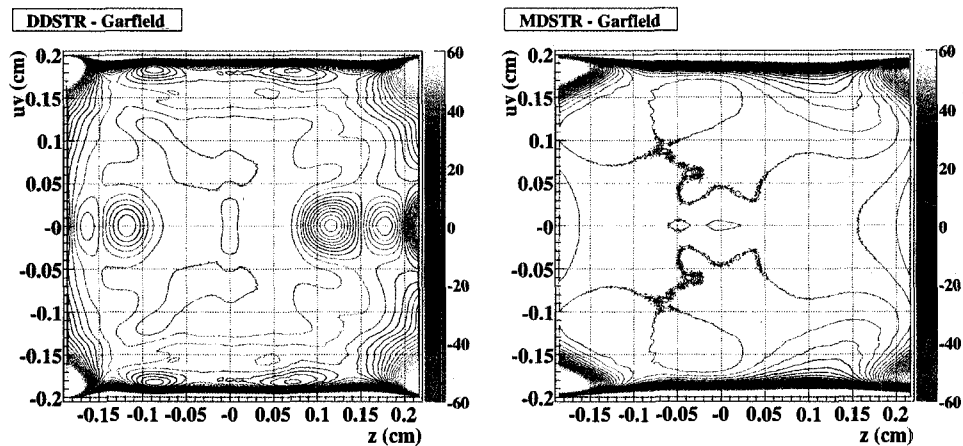


Figure 8.3: Differences between data-driven (left) or Monte Carlo-driven (right) STRs and Garfield STRs. The main differences between measured and Garfield STRs are near the cell corners and along the edges; the central region of the cell, where the majority of hits are found (Fig. 8.4), generally changes by 5 ns or less. The contour scale is in nanoseconds.

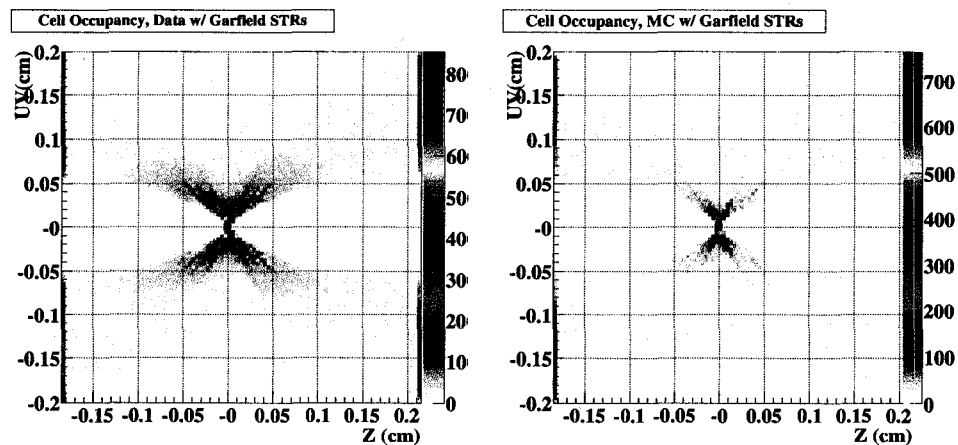


Figure 8.4: Hit occupancies within a drift cell, for data and MC, showing that the majority of hits used by Mofia are near the centre of the cell. The colour scales are in arbitrary units.




Figure 8.5: Effects of analyzing data and MC using measured STRs, vs using Garfield STRs. The plots show how the measured spectra change, for data (top left) and MC (top right); the differences between data and MC are shown in the bottom plot. Colour scales are in arbitrary units; the important information in these plots are the shapes of the spectrum differences. Note that the band structure apparent in the spectrum differences for data and MC are apparently gone in the (MC – data) difference plot.

Analysis Bias Due to Garfield STRs

The use of the STRs measured from data or MC changes the reconstructed decay spectrum slightly, and the changes are broadly similar between data and MC. If data and MC spectra are changed differently, this will change the decay parameters measured from the MC-data comparison. To measure the change of the spectrum shape, data (set35) and MC (gen335) were re-analyzed with their corresponding measured STRs, and the reconstructed spectra were compared to the set35 and gen335 spectra reconstructed using Garfield STRs. The differences in the normalized spectra produced using the different STRs are shown in Fig. 8.5 for data and MC, as is the MC-data difference of the shape changes. The differences in shape are relatively small—the reconstructed spectra still “look” like Michel spectra.

The data spectrum reconstructed using measured STRs was fit against the data spectrum reconstructed using Garfield STRs, and the same was done for MC. These fits characterize the distortions of the reconstructed spectrum, but it is their *difference* that changes the decay parameters that *TWIST* measures. The fits and their differences are listed in Tab. 8.4. These differences can be used to correct the standard decay parameter measurements; in this case, the uncertainty in the differences represents a systematic uncertainty in the final measurements.

	Data	MC	(MC-Data)
$\Delta\rho$	$(-3.7 \pm 2.1) \times 10^{-4}$	$(-1.27 \pm 2.0) \times 10^{-4}$	$(-2.0 \pm 2.9) \times 10^{-4}$
$\Delta\delta$	$(4.67 \pm 3.8) \times 10^{-4}$	$(3.06 \pm 3.5) \times 10^{-4}$	$(+1.6 \pm 5.2) \times 10^{-4}$
$\Delta\xi$	$(-11.53 \pm 4.7) \times 10^{-4}$	$(-7.54 \pm 4.4) \times 10^{-4}$	$(+4.0 \pm 6.4) \times 10^{-4}$

Table 8.4: Fits between spectra reconstructed using measured STRs to those using Garfield STRs, for data and MC. Differences are (Simulation – Data), and are used to correct the measured decay parameters as proscribed in the text. Uncertainties in the differences then represent systematic uncertainties in the measured Michel parameters.

Correction of the Decay Parameters due to Garfield STRs

To determine the correct sign with which to apply the corrections in Tab. 8.4, recall that the standard decay parameter fit is of the form

$$\text{Data} = \text{MC} + \text{derivatives.}$$

The fit parameters are coefficients on the derivatives, describing Data – MC. The fits done to determine the STR corrections in Tab. 8.4 were of the form

$$\text{DDSTR} = \text{Garfield} + \text{derivatives}$$

where “DDSTR” refers to the analysis using data-driven (or MC-driven) STRs. So the fit parameters in this case describe DDSTR – Garfield. Then the MC – Data differences listed in Tab. 8.4 are

$$(\text{MDSTR} - \text{Garfield}_{\text{MC}}) - (\text{DDSTR} - \text{Garfield}_{\text{Data}}) \quad (8.4)$$

or

$$(\text{MDSTR} - \text{DDSTR}) - (\text{Garfield}_{\text{MC}} - \text{Garfield}_{\text{Data}}) \quad (8.5)$$

Since fit parameters from a standard decay parameter fit are of the form $\text{Garfield}_{\text{Data}} - \text{Garfield}_{\text{MC}}$, the correction to add is

$$(\text{DDSTR} - \text{MDSTR}) - (\text{Garfield}_{\text{Data}} - \text{Garfield}_{\text{MC}}) \quad (8.6)$$

which is the negative of Eqs. 8.5 and 8.4. So the correction to be applied to the measured decay parameters is of the opposite sign of the MC – Data differences listed in Tab. 8.4.

The uncertainties in these corrections then results in a systematic uncertainty in the measured decay parameters.

8.1.2 Chamber Drift Times (STRs): Gas Density Variations

The time for an electron to drift to the sense wire in a drift chamber cell depends on the mean free path of the electron, and hence on the gas density. The same STRs were used for the analysis of all data sets, so variations in the gas density represent a systematic error. The gas density varied by about 0.4% (RMS/mean) over the course of the 2004 data taking, with no set more than 0.7% away from the mean.

	Slope (keV/c)	Intercept (keV/c)
Upstream	13 ± 1	26 ± 2
Downstream	16 ± 3	30 ± 5

Table 8.5: Results of comparative energy calibration between simulation with low-density STR file and standard. By comparison, the slopes and intercepts energy calibration between data and simulation are in the range of 5–20 keV/c.

	Gas Density Variation
$\Delta\rho$	$(1.75 \pm 3.00) \times 10^{-4}$
$\Delta\delta$	$(-2.99 \pm 5.30) \times 10^{-4}$
$\Delta\xi$	$(3.83 \pm 6.54) \times 10^{-4}$
χ^2/dof	1157/2463
Renormalization	0.685
Scale factor	25

Table 8.6: Unscaled Michel parameter fit of simulation analyzed with low-density STRs against standard simulation.

Note that the variations in the gas density will also affect the centroid and width of the stopping distribution; however, these systematic effects are accounted for in the studies of the energy calibration (Sec. 8.2.2) and the momentum resolution (Sec. 8.4.1).

Systematic Effects

A new STR file was created using Garfield, with conditions to match the STRs used for standard analysis but with a pressure of 90.7 kPa instead of the usual 100.7 kPa (the mean pressure for all 2004 data)—a reduction in the pressure, and hence the gas density (which is proportional to P/T) of 10%. The scale factor is then $10\%/0.4\% = 25$. 800 runs of standard simulation were analyzed with this new STR file, and the result was compared to the standard analysis.

The results of energy calibration between the two analyses are given in Tab. 8.5; changing the STR file in this way results in a significant change in the energy calibration. The results of the spectrum fit of the two analyses are listed in Tab. 8.6, before rescaling, and the results with scaling applied are listed in Tab. 8.7; the latter represent the systematic uncertainty due to the variations in the chamber gas density.

	Gas Density Variation
$\Delta\rho$	$(0.07 \pm 0.08) \times 10^{-4}$
$\Delta\delta$	$(-0.12 \pm 0.15) \times 10^{-4}$
$\Delta\xi$	$(0.15 \pm 0.18) \times 10^{-4}$

Table 8.7: Michel parameter fit of simulation analyzed with low-density STRs against standard simulation, after applying scale factors. This represents the systematic uncertainty due to variations in the drift chamber gas density.

8.1.3 Chamber Drift Times (STRs): Foil Bulge Variations

The time for an electron to drift to the sense wire in a drift chamber cell depends on the shape and strength of the electric field within the drift cell; the electric field depends on the positions of the cathode foils and anode wire. As explained in Sec. 4.2.3, variations in the differential pressure between the drift chambers and the interstitial helium will change the positions of the outer cathode foils of each module. The typical RMS variation in the foil positions over the course of the 2004 data taking, according to measurements by Jingliang Hu, was about $50\text{ }\mu\text{m}$ (see Tab. 7.3). Foil positions were measured using changes in the drift times of hits in each chamber; the foil positions affect the chambers' electric fields, which change the drift times. Data were taken with the chambers' differential pressure changed by known amounts, to calibrate the changes in drift times. The relationship between differential pressure and foil bulge was calculated using the tension of the foils.

Systematic Effects

Dick Mischke² determined that it was possible to approximately relate the ways in which the STRs changed due to several possible effects—the effect of a certain shift in the foil positions was of a similar order as certain changes in the gas density, etc. This can be used to obtain an order of magnitude estimate in the effects of the foil bulge variations instead of running additional simulations. Renee Poutissou³ studied the ways in which various condition changes affected the STRs. She showed that the changes in STRs due to a $200\text{ }\mu\text{m}$ cathode foil shift was approximately 0–8%, depending on the region of the cell, with the typical change being about 3% for the parts of the cell in which significant numbers of hits were recorded. The changes in STRs due to a 10% reduction in gas density was around 15% over most of the cell. Although the shapes of the changes are quite different, this suggests that in magnitude the effect of a $200\text{ }\mu\text{m}$ foil shift is roughly 4–5 times smaller than the effect of a 10% reduction in gas density.

By these arguments, the expected changes in the Michel parameters due to a $200\text{ }\mu\text{m}$ foil shift are about 5 times smaller than those shown in Tab. 8.6 for the gas density variations study above. Since the variation in foil position was approximately $50\text{ }\mu\text{m}$, an order of magnitude estimate of the systematic uncertainties in the Michel parameters due to variations in the foil bulge can be obtained from Tab. 8.6 using a scale factor of $(5)(4) = 20$. The results are given in Tab. 8.8. These results are small enough that this estimate is sufficient.

²Dick Mischke, “STR files”, TWIST Physics internal forum, 14 August, 2004.

https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_physics&key=1092509786

³Renee Poutissou, “Study of str's produced under various conditions”, TWIST General internal forum, 8 November, 2003.

https://twist.phys.ualberta.ca/forum/view.php?bn=twist_general&key=1068329550

	Foil Bulge Variation
$\Delta\rho$	$(0.09 \pm 0.10) \times 10^{-4}$
$\Delta\delta$	$(-0.15 \pm 0.18) \times 10^{-4}$
$\Delta\xi$	$(0.19 \pm 0.22) \times 10^{-4}$

Table 8.8: Estimate of the systematic uncertainty due to variations in the drift chamber foil bulge, based on an approximate scale factor of 20 applied to the Michel parameter fit in Tab. 8.6.

8.1.4 Foil Bulge: Geometric Effects

In addition to affecting the STRs, changes in the foil positions affect the size of the drift cell, changing the probability that a track will pass through it or not and hence changing the probability of a hit being recorded for a given track. Since the simulation used a constant foil geometry, the number of hits per track available for analysis will not be the same as for data, possibly influencing reconstruction biases.

Measurement of Systematic Uncertainties

The standard geometry file (dt_geo.00053) was edited to move the outer DC foils of each module inward toward the central foil by $500\text{ }\mu\text{m}$, 10 times more than the $50\text{ }\mu\text{m}$ foil position variation measured by Jingliang Hu (see Sec. 8.1.3). This altered geometry file was then used for a simulation which was otherwise identical to the standard, including using the same STR file.

The expected impact of this change to the geometry file is to reduce the number of hits available to the analysis; in particular, tracks which would have just crossed the corner of a 4 mm cell and left a hit would not cross the corner of a 3.5 mm cell. Not only does this reduce the number of hits available to the helix fitter, but it decreases the size of hit clusters used by FirstGuess to estimate the track angle. This change appears in the number of degrees of freedom reported by the helix fitter, as seen in Fig. 8.6; the average number of degrees of freedom dropped from 51.2 to 49.4.

Systematic Effects

The simulation with exaggerated foil bulge was fit against the complete standard spectrum in the usual way, to measure the Michel parameters. Since this was a new simulation, the correlations to the standard simulation were low; no renormalization of the uncertainties will be performed. The raw fits are shown in Tab. 8.9. Table 8.10 shows the fits after applying the scale factors. These rescaled results represent the systematic uncertainties from the geometric effects of the incorrect simulation of the foil bulge.




Figure 8.6: Number of degrees of freedom available to the helix fitter, with standard geometry and with a 500 μm foil bulge.

	Foil Bulge
$\Delta\rho$	$(-1.15 \pm 3.02) \times 10^{-4}$
$\Delta\delta$	$(-2.52 \pm 5.33) \times 10^{-4}$
$\Delta\xi$	$(-7.13 \pm 6.58) \times 10^{-4}$
χ^2/dof	2404/2463
Renormalization	1
Scale factor	10

Table 8.9: Unscaled Michel parameter fits for the geometric foil bulge systematic: simulation with exaggerated foil bulge, fit against standard.

	Foil Bulge
$\Delta\rho$	$(-0.12 \pm 0.30) \times 10^{-4}$
$\Delta\delta$	$(-0.25 \pm 0.53) \times 10^{-4}$
$\Delta\xi$	$(-0.71 \pm 0.66) \times 10^{-4}$

Table 8.10: Michel parameter fits for the geometric foil bulge systematic, after applying the scale factor. This represents the systematic uncertainties due to the geometric effects of the incorrect simulation of the foil bulge.

	Asymmetric Reconstruction Efficiency
$\Delta\rho$	$(-1.06 \pm 2.46) \times 10^{-4}$
$\Delta\delta$	$(6.00 \pm 4.34) \times 10^{-4}$
$\Delta\xi$	$(-11.36 \pm 5.35) \times 10^{-4}$
χ^2/dof	0.178
Renormalization	0.00850
Scale factor	10

Table 8.11: Unscaled Michel parameter fit of standard simulation with an upstream reconstruction efficiency imposed against the unaltered simulation.

	Asymmetric Reconstruction Efficiency
$\Delta\rho$	$(-0.106 \pm 0.002) \times 10^{-4}$
$\Delta\delta$	$(0.600 \pm 0.004) \times 10^{-4}$
$\Delta\xi$	$(-1.136 \pm 0.005) \times 10^{-4}$

Table 8.12: Michel parameter fit of standard simulation with an upstream reconstruction efficiency imposed against the unaltered simulation, after applying scale factors. This represents the systematic uncertainty due to asymmetric reconstruction efficiency.

8.1.5 Efficiency Asymmetry

As shown in Sec. 6.3.2, there is no evidence of any momentum or angle dependence to the difference in reconstruction efficiencies between data and simulation. For an estimate of possible systematic errors in the Michel parameters due to differences in reconstruction efficiency, the effect of a constant MC-data difference in efficiencies for upstream and a different constant difference downstream is studied. Recall that the measured MC-data difference was $(6 \pm 2) \times 10^{-5}$ upstream, and $(3 \pm 2) \times 10^{-5}$ downstream. As a conservative estimate, the differences were taken as 6×10^{-5} upstream and 0 downstream.

Systematic Effects

A standard simulation was altered by multiplying all upstream bins by $1 - (6 \times 10^{-4})$, to reproduce the effect of an inefficiency with an exaggeration factor of 10. This altered spectrum was fit against the original. Table 8.11 lists the spectrum fit results; the results with the scale factor applied, shown in Tab. 8.12, represent the systematic uncertainty due to an asymmetry in the reconstruction efficiency.

8.1.6 Wire Time Offsets (t_0 's)

Recall that wire time offsets, or t_0 's, are wire-dependent time offsets added to the times of received signals from the wire chambers (see Sec. 4.3.2 for details). These are to account for differences in signal cable lengths, variations in electronics, and other effects, and must be measured *in situ*.

There are two possible sources of systematic error from the measurement of the t_0 values: measurement bias, and time variation. These will be considered separately.

t_0 Measurement Bias

t_0 values are measured from the data and used as corrections during data analysis. If the same procedure were applied during MC analysis, any biases in the measurement of t_0 values might be cancelled in the comparison between data and MC, but this is not done. Therefore, biases in the measurement of t_0 's may result in systematic errors.

t_0 values are measured from the leading edge of the time spectrum. The shape of this leading edge varies with track angle, and this produces a variation with wire number. This effect is due to the way the spatial distribution of hits in the drift cell varies with track angle and is not a real time offset, so it represents a measurement bias. Because there is nothing in the MC to cause a true t_0 offset, measuring the t_0 's from MC will reveal the measurement bias. The "correct" t_0 value depends on settings of the t_0 analysis software, including assumed scintillator time offsets, but this should be the same for all wires on all planes, and a global shift in the t_0 's does not affect the normal analysis. Furthermore, the effect of random t_0 errors should only be to degrade the reconstruction resolution, and the systematic error due to resolution is measured separately. Thus trends in the measured t_0 's are the measures of interest.

t_0 values were measured using broad beams of 120 MeV/c pions. A similar beam was simulated for this study, and t_0 values measured. To look for trends, averages were taken for each wire in a set of DC planes, considering upstream and downstream U and V planes separately, and fourth-degree polynomials were fit to parameterize any trends; see Fig. 8.7. Wire-to-wire variations were seen on the order of 0.1–0.2 ns, and the shapes are broadly similar among the four sets of planes. Averages were also taken for all wires in each plane (Fig. 8.8). Here more scatter is seen, with a range of up to 0.5 ns, but there are no obvious trends, and the difference between upstream and downstream is about 0.04 ns.

t_0 Time Variations

For 2004 data taking, t_0 values derived from data taken at the start of the run period were used for analysis of all data. t_0 data was also taken at the end of the run period to confirm that the values did not change over the course of data taking. This end-of-run data was complicated by the fact that several DC planes were by that point not working, and the signal threshold on DC 36 had been changed due to signal oscillations. (All of these changes occurred after primary data taking was finished, and so should not affect the data analysis.) Still, while the end-of-run data is not useful for directly determining t_0 values for use in analysis, by averaging over good planes comparisons can still be made between start-of-run

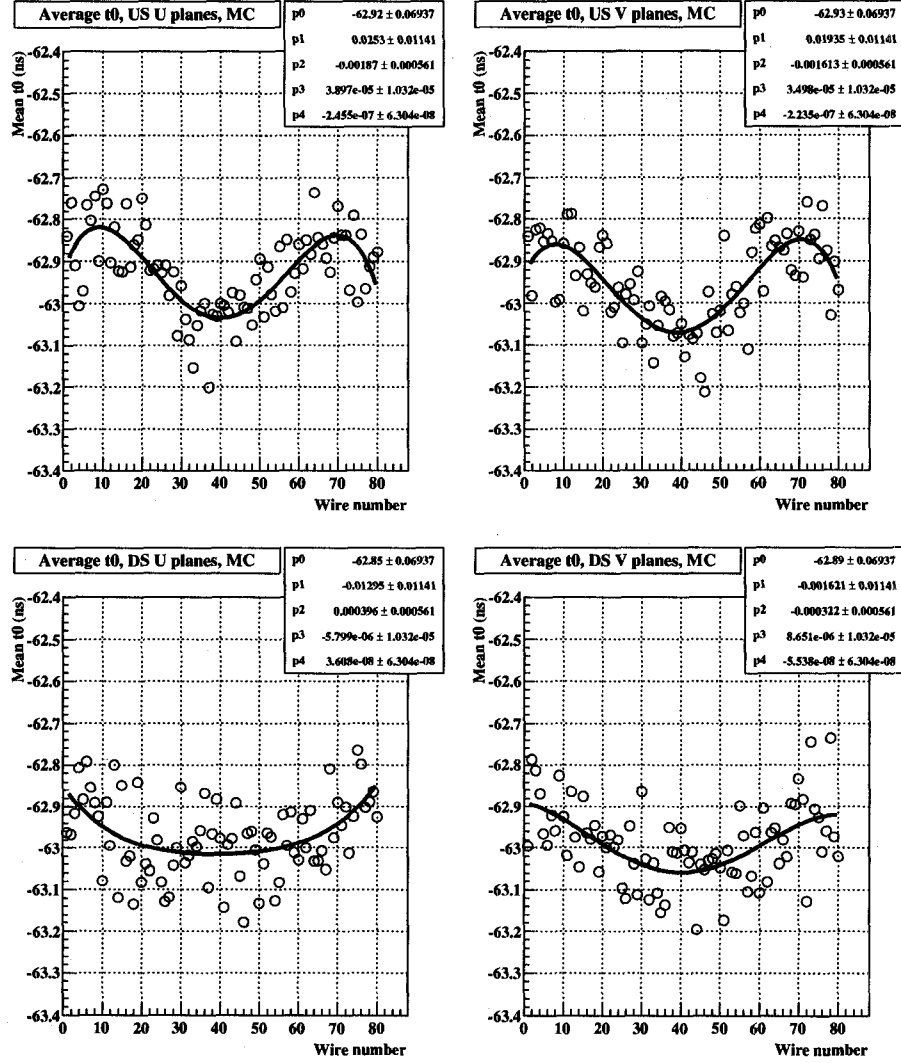


Figure 8.7: t_0 values for each wire in a DC plane, averaged over all upstream and downstream U and V planes, as measured from MC. Fourth-order polynomials have been fit to each set of t_0 's.




Figure 8.8: Average t_0 values for each DC plane, as measured from MC. Because there are no obvious trends, no polynomial fits were performed. The average t_0 value is -62.94 ns upstream, and -62.98 ns downstream.

and end-of-run t_0 's to check for time variations.

Differences were calculated between end- and start-of-run t_0 values, and as with MC averages were taken for corresponding wires in a set of DC planes, considering upstream and downstream U and V planes separately, and fourth-degree polynomials were fit to parameterize any trends; see Fig. 8.9. DC 36 was excluded from these averages. Averages were also taken for all wires in each plane (Fig. 8.10); the t_0 value for DC 36 was found to have changed by more than 2 ns. Ignoring this plane, the t_0 changes are seen to roughly follow a quadratic shape as a function of plane number, covering mainly a ~ 0.2 ns range.

Systematic Effects

Since the t_0 measurement bias and time variation are of similar magnitude, the effects of each were measured in separate systematic studies.

For the case of the measurement bias test, a t_0 file was created using quartic functions with coefficients some factor larger than the fit results shown in Fig. 8.7; the exaggeration was done after subtracting the upstream or downstream average t_0 value found in Fig. 8.8. No other z -dependent offset was applied, since the upstream and downstream averages were consistent with each other.

For the case of the time dependent variation test, a t_0 file was created by adding quartic functions to the standard DC t_0 file. Again, the quartic functions were produced using

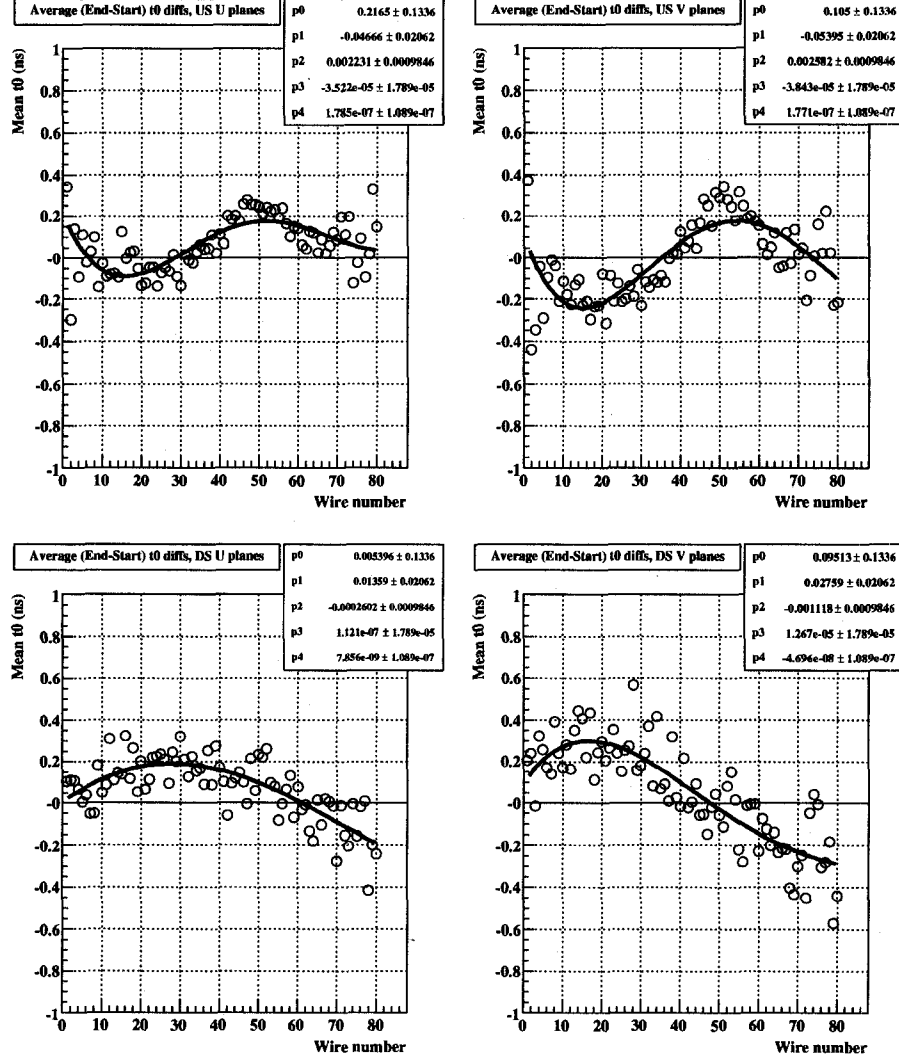


Figure 8.9: Average differences in end-of-run and start-of-run t_0 values for each wire in a DC plane, averaged over all upstream and downstream U and V planes. Fourth-order polynomials have been fit to each set of t_0 's.




Figure 8.10: Average differences in end-of-run and start-of-run t_0 values for each DC plane. The top and bottom graphs differ only in vertical scale; the top graph shows the aberration of DC 36 t_0 due to its raised threshold. A second-degree polynomial was fit to the average t_0 differences from good planes (bottom graph).

	Scale Factor	Mean χ^2/dof
Standard	—	2.508 ± 0.004
Measurement Bias	10	2.54 ± 0.03
	20	2.57 ± 0.03
Time variation	10	2.55 ± 0.03
	20	2.67 ± 0.03

Table 8.13: Effects of doctored t_0 files on the χ^2/dof results from the helix fitter, for measurement bias and time variation tests, for different scale factors.

	Measurement Bias	Time Variation
$\Delta\rho$	$(-1.46 \pm 3.01) \times 10^{-4}$	$(1.67 \pm 3.01) \times 10^{-4}$
$\Delta\delta$	$(1.68 \pm 5.30) \times 10^{-4}$	$(1.64 \pm 5.30) \times 10^{-4}$
$\Delta\xi$	$(-3.98 \pm 6.54) \times 10^{-4}$	$(0.12 \pm 6.54) \times 10^{-4}$
χ^2/dof	899/2463	924/2463
Renormalization	0.604	0.613
Scale factor	10	10

Table 8.14: Unscaled Michel parameter fits for t_0 systematics, based on measurement bias and on time variations: simulations with doctored t_0 files, fit by standard.

coefficients some factor larger than those from the end-start comparison shown in Fig. 8.9, after the upstream and downstream averages were subtracted off. On top of this, the z -dependence of the time variations was added using an exaggeration of the quadratic function shown in Fig. 8.10.

The scale factors were chosen after performing small test analyses (20 000 events per test) of MC with t_0 files doctored as described above, using scale factors of 10 and 20, and checking the results on the mean χ^2/dof . The results are shown in Tab. 8.13; the scale factor of 20 is not significantly worse than 10. However, a scale factor of 10 was chosen to be conservative, in both cases. This scale factor was applied to the coefficients of the fit functions and new t_0 files were created, as described above.

The exaggerated measurement bias (anal8) and time variation (anal9) analyses were each fit against the complete standard spectrum to measure the effects on the Michel parameters. Table 8.14 shows the raw fit results. Table 8.15 shows the fit results after applying the scale factors and renormalizing the uncertainties. These rescaled results, combined in quadrature, represent the systematic uncertainties from the t_0 measurements.

	Measurement Bias	Time Variation	Total
$\Delta\rho$	$(-0.146 \pm 0.182) \times 10^{-4}$	$(0.167 \pm 0.184) \times 10^{-4}$	$(0.222 \pm 0.259) \times 10^{-4}$
$\Delta\delta$	$(0.168 \pm 0.320) \times 10^{-4}$	$(0.164 \pm 0.325) \times 10^{-4}$	$(0.235 \pm 0.456) \times 10^{-4}$
$\Delta\xi$	$(-0.398 \pm 0.395) \times 10^{-4}$	$(0.012 \pm 0.401) \times 10^{-4}$	$(0.398 \pm 0.563) \times 10^{-4}$

Table 8.15: Michel parameter fits for t_0 systematics, based on measurement bias and on time variations, after applying scale factors and renormalizing the errors. The “Total” column is the quadrature sum of the other two; this represents the systematic uncertainties due to uncertainties in the measurement bias and time variation of the t_0 ’s.

8.2 Energy Scale

The Energy Scale category includes effects which influence the measurement of the positron energy. The systematic uncertainties due to these effects are summarized in Tab. 8.16.

Energy Scale	ρ	δ
Magnetic Field Map Shape (8.2.1)	0.60	0.74
Energy Calibration (8.2.2)	2.74	3.86
Momentum Dependence of Energy Calibration (8.2.3)	0.79	0.73
Total	2.86	4.07

Table 8.16: Summary of systematic uncertainties related to energy scale, in units of 10^{-4} .

8.2.1 Magnetic Field Map Shape

The magnetic field within the *TWIST* detector needs to be known in two ways: the absolute strength of the field is needed for determining the momentum of a positron from the shape of its helical track; and the shape of the field is needed for properly reconstructing the positron track within the detector.

The magnetic field strength is measured by an NMR probe located inside the *TWIST* detector stack (Sec. 4.2.1). The magnetic field map is scaled to match the field strength at the NMR probe position for analysis; the uncertainty in this scaling is essentially the degree to which we trust the field map, which is at the $\pm 1 \times 10^{-4}$ T level. However, part of the role of the energy calibration is to compensate for mis-scaling the magnetic field, making the uncertainty in the field strength a part of the energy calibration systematic (section 8.2.2).

The shape of the simulated field map used for the analysis of the data does not perfectly match the real shape of the magnetic field (Sec. 4.2.1). The differences mean that the positron tracks will not be reconstructed correctly. (Note that this is not true for simulation, as the same field map is used in simulation and analysis.) A systematic error therefore arises from using a field map with the wrong shape.

Systematic Effects

To test the effect of the field map shape, the difference ΔB_z between the measured and simulated field maps was taken, and parameterized with the function

$$\Delta B_z = c_2 z^2 + c_3 z^3 + c_r r \quad (8.7)$$

where c_2 , c_3 , and c_r were the fit parameters; the best fit values were found to be

$$c_2 = -6 \times 10^{-8} \text{ T/cm}^2, \quad c_3 = -4 \times 10^{-10} \text{ T/cm}^3, \quad c_r = -1.25 \times 10^{-5} \text{ T/cm}. \quad (8.8)$$

	Distorted Field Map
$\Delta\rho$	$(5.96 \pm 2.46) \times 10^{-4}$
$\Delta\delta$	$(7.42 \pm 4.34) \times 10^{-4}$
$\Delta\xi$	$(-7.69 \pm 5.35) \times 10^{-4}$
χ^2/dof	165/2463
Renormalization	0.259
Scale factor	10

Table 8.17: Unscaled Michel parameter fits for field map systematic.

	Distorted Field Map
$\Delta\rho$	$(0.60 \pm 0.06) \times 10^{-4}$
$\Delta\delta$	$(0.74 \pm 0.11) \times 10^{-4}$
$\Delta\xi$	$(-0.77 \pm 0.14) \times 10^{-4}$

Table 8.18: Michel parameter fits for field map systematic, after applying the scale factor and renormalizing the errors. These represent the systematic uncertainties due to uncertainties in the shape of the magnetic field map.

For the systematic exaggeration, a new field map was produced by adding to the standard field map an exaggerated ΔB_z from Eq. 8.7 with coefficients

$$c_2 = -6 \times 10^{-7} \text{ T/cm}^2, \quad c_3 = -4 \times 10^{-9} \text{ T/cm}^3, \quad c_r = -1.25 \times 10^{-4} \text{ T/cm}, \quad (8.9)$$

for a scale factor of 10. A standard surface muon simulation was analyzed using this modified field map, and the result was fit against the standard analysis using the Michel spectrum fitter. The results of the spectrum fit are shown in Tab. 8.17; and the systematic errors obtained after scaling the results and renormalizing the errors are listed in Tab. 8.18.

8.2.2 Energy Calibration

Variations in the energy calibration affect the measured Michel parameters, leading to a systematic uncertainty in the result. This uncertainty is statistical, but since the energy calibration is determined separately from the Michel parameter measurement its uncertainty will be treated separately as well.

Systematic Error in the Energy Calibration

This systematic uncertainty was studied by Ryan Bayes.⁴ Recall from Sec. 5.5.1 that the energy calibration is parameterized using upstream and downstream slopes ($\Delta\varsigma_{US}$ and $\Delta\varsigma_{DS}$) and intercepts ($\Delta\varrho_{US}$ and $\Delta\varrho_{DS}$) of the differences in endpoint positions between the two spectra of interest. The result of the energy calibration between standard data and the corresponding simulation is:

$$\begin{aligned} \Delta\varsigma_{US} &= 0.0143 \pm 0.0017 \text{ MeV/c} & \Delta\varrho_{US} &= 0.0145 \pm 0.0026 \text{ MeV/c} \\ \Delta\varsigma_{DS} &= 0.0073 \pm 0.0038 \text{ MeV/c} & \Delta\varrho_{DS} &= 0.0082 \pm 0.0060 \text{ MeV/c}. \end{aligned} \quad (8.10)$$

⁴Ryan Bayes, "statistical part of endpoint calibration systematic", TWIST Physics internal forum, 1 August, 2007.
https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_physics&key=1186006616

	US Slope + 100 keV/c	US Intercept + 100 keV/c
$\Delta\rho$	$(-17.2 \pm 3.7) \times 10^{-4}$	$(+6.3 \pm 3.7) \times 10^{-4}$
$\Delta\delta$	$(-93.0 \pm 6.5) \times 10^{-4}$	$(+67.2 \pm 6.7) \times 10^{-4}$
$\Delta\xi$	$(+116.3 \pm 8.2) \times 10^{-4}$	$(-83.4 \pm 8.2) \times 10^{-4}$
χ^2/dof	2535/2463	2496/2463
	DS Slope + 100 keV/c	DS Intercept + 100 keV/c
$\Delta\rho$	$(-37.8 \pm 3.7) \times 10^{-4}$	$(+22.1 \pm 3.7) \times 10^{-4}$
$\Delta\delta$	$(+45.4 \pm 6.6) \times 10^{-4}$	$(-35.9 \pm 6.7) \times 10^{-4}$
$\Delta\xi$	$(-3.0 \pm 8.2) \times 10^{-4}$	$(+2.3 \pm 8.2) \times 10^{-4}$
χ^2/dof	2490/2463	2541/2463

Table 8.19: Unscaled Michel parameter fits for the energy calibration systematic: simulation with energy calibration parameters exaggerated, fitted against simulation with standard calibration. The χ^2/dof suggests the fits are essentially uncorrelated, probably due to the fact that the treesums with exaggerated calibrations were only subsets of the standard treesum. The exaggerated calibrations also had substantial impacts on the spectrum shape, as seen in the fit results, and this likely contributed to the large χ^2/dof as well.

	$\Delta\varsigma_{US}$	$\Delta\varrho_{US}$	$\Delta\varsigma_{DS}$	$\Delta\varrho_{DS}$
$\Delta\varsigma_{US}$	2.98×10^{-6}	4.42×10^{-6}	0	0
$\Delta\varrho_{US}$		6.72×10^{-6}	0	0
$\Delta\varsigma_{DS}$			1.42×10^{-5}	-2.22×10^{-5}
$\Delta\varrho_{DS}$				3.56×10^{-5}

Table 8.20: The error matrix for a standard energy calibration. Diagonal elements are the squares of the uncertainties in the fit parameters; off-diagonal elements show correlations between parameters. The matrix is symmetric; only the upper triangle is shown.

The uncertainties in the downstream parameters are larger because the momentum edge is less well defined and the statistics are much poorer, due to the shape of the decay spectrum.

Systematic Effects

To test the effects of the uncertainty in the energy calibration on the measured Michel parameters, the above calibration (8.10) was applied to the standard simulation in a spectrum summation. Then four new spectra were created, each with one of the parameters exaggerated by 100 keV/c, and were fitted against the spectrum made with standard calibration parameters. (This was repeated with exaggerations of 75 keV/c, 125 keV/c, and 150 keV/c, to test for linearity; the effect was found to be linear in the exaggeration used. The 100 keV/c results are presented here.) The new spectra used 1100 runs. The results of the spectrum fits are shown in Tab. 8.19.

There are significant correlations between the energy calibration parameters, as shown in Tab. 8.20; a slightly different technique from usual was used to determine the final systematic uncertainty to take these correlations into account. The relationship between the error σ_q in some parameter q of the experiment and the resulting uncertainty in a Michel parameter, as expressed in Eq. (8.1), assumes that the parameters being considered are uncorrelated. In the case of energy calibration, however, the correlations between parameters must be

	Energy calibration
$\Delta\rho$	2.74×10^{-4}
$\Delta\delta$	3.86×10^{-4}
$\Delta\xi$	0.45×10^{-4}

Table 8.21: Systematic uncertainties in Michel parameters due to uncertainties in the energy calibration, determined using Eq. (8.11).

considered, and a more general expression must be used:

$$\sigma_{\lambda\alpha}^2 = \frac{\partial\lambda_\alpha}{\partial q_\gamma} \frac{\partial\lambda_\beta}{\partial q_\delta} \sigma_{\gamma\delta}^2 \quad (8.11)$$

where repeated indices denote summation. Here, $\sigma_{\lambda\alpha}^2$ is the error matrix (or covariance matrix) for the Michel parameters under the systematic uncertainty being studied—in this case, the energy calibration parameters; and $\sigma_{\gamma\delta}^2$ is the error matrix for the parameters being studied. The diagonal elements of $\sigma_{\lambda\alpha}^2$ represent the systematic uncertainties in the Michel parameters. Note that Eq. (8.11) simplifies to Eq. (8.1) when $\sigma_{\gamma\delta}^2$ is diagonal. The covariance matrix for the energy calibration parameters is calculated from a least squares method used in the linear fit component of the calibration:

$$\sigma_{\gamma\delta}^2 = \left(\frac{\partial^2 \chi^2}{\partial q_\gamma \partial q_\delta} \right)^{-1}. \quad (8.12)$$

The uncertainties in the energy calibration parameters are the diagonal elements of Eq. (8.12). $\sigma_{\gamma\delta}^2$ depends both on the correlations between fit parameters and on the statistics in the data and simulation spectra being correlated. For this study, $\sigma_{\gamma\delta}^2$ is taken from Tab. 8.20; $\frac{\partial\lambda_\alpha}{\partial q_\gamma}$ is calculated using Tab. 8.19 for the values of $\partial\lambda_\alpha$, with $\partial q_\gamma = 100$ keV/c. Using these values in Eq. 8.11 gives the systematic uncertainty due to uncertainties in the energy calibration, as listed in Tab. 8.21.

8.2.3 Momentum Dependence of Energy Calibration

Once the energy calibration is determined at the 52.8 MeV/c momentum edge, it must be applied to the bulk of the spectrum with some assumed momentum dependence, as discussed in Sec. 5.5.2. The two likely “extremes” identified were linear or constant momentum dependence. In standard analysis, recall that the calibration is applied independent of momentum, using

$$p_{cal} = p_{rec} - \Delta q_{US,DS} + \frac{\Delta\zeta_{US,DS}}{|\cos\theta|} \quad (8.13)$$

where p_{rec} is the momentum as determined by the track reconstruction, and p_{cal} is the calibrated momentum.

	Momentum Dependence of Energy Calibration
$\Delta\rho$	$(0.79 \pm 2.55) \times 10^{-4}$
$\Delta\delta$	$(0.73 \pm 4.34) \times 10^{-4}$
$\Delta\xi$	$(0.09 \pm 5.35) \times 10^{-4}$
χ^2/dof	6.5/2463
Renormalization	0.051
Scale factor	1

Table 8.22: Michel parameter fits for systematic uncertainty due to the momentum dependence of the energy calibration. No scale factor was applied in the study.

Systematic Effects

The effect of assuming a linear dependence instead of Eq. (8.13) was studied by Ryan Bayes.⁵ The difference in the calibrated momentum p_{cal} between the momentum-independent calibration of Eq. (8.13) and a momentum-dependent calibration can be expressed to first order as

$$\Delta p_{tot} = \left(\frac{\Delta_{SUS,DS}}{|\cos\theta|} - \Delta_{QUS,DS} \right) \left(\frac{p_0 - p_{rec}}{p_0} \right) \quad (8.14)$$

where $p_0 \approx 52.83$ MeV/c is the kinematic cutoff momentum. This was applied to the standard surface muon simulation, using the parameters from the calibration against standard data,, listed in Eq.(8.10), and the resulting spectrum was fit against the uncalibrated simulation. The results of this fit, shown in Tab. 8.22, represent the systematic uncertainty due to the unknown momentum dependence of the energy calibration.

8.3 Positron Interactions

The Positron Interactions category includes effects related to processes undergone by the positrons which affect their energy and direction. The systematic uncertainties due to these effects are summarized in Tab. 8.23.

Positron Interactions	ρ	δ
Simulation of delta ray production (8.3.1)	-1.5	-0.9
Simulation of bremsstrahlung production (8.3.1)	-0.04	-0.03
Outside Materials (8.3.2)	0.39	-0.27
Total	1.60	0.93

Table 8.23: Summary of systematic uncertainties related to positron interactions, in units of 10^{-4} .

⁵Ryan Bayes, "Re: energy calibration: test of momentum dependence", TWIST Software internal forum, 5 October, 2007. https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_software&key=1191626571

8.3.1 Simulated Bremsstrahlung and Delta Ray Production Rates

The Positron Interactions category of systematic uncertainties covers the simulation of the way positrons behave as they pass through the *TWIST* detector, primarily regarding the way they scatter and lose energy. “Softer” interaction processes are covered by other systematic uncertainties: the effects of the continuous energy loss model is corrected as part of the energy calibration (Sec. 8.2.2), and the multiple scattering processes affect the reconstruction resolution (Sec. 8.4). This section looks at the effect of discrete processes, of which bremsstrahlung and delta rays are the most important.

Delta rays are electrons which are knocked out of their atomic orbits by passing energetic particles (mainly positrons, in our case), and which are given enough energy to leave their own tracks in the detector. Bremsstrahlung is the radiation of one or more photons emitted as the passing particles (positrons in particular, in our case) decelerate. These processes can significantly change the energy and (in the case of delta ray production) angle of the incident positron, which will distort the reconstructed muon decay spectrum; the sudden change can also affect track reconstruction. These effects are accounted for in the *TWIST* standard analysis by comparing data to simulation, which includes both processes in detail; an error in the simulated rates of these processes, then, can affect the data-MC comparison and the reconstructed decay parameters.

Systematic Error in the Delta Ray Rate

A simple method of evaluating the relative delta ray production rates between data and MC involves the “upstream stops” measurements described in Sec. 6.3.1. An event in which a delta ray is produced in or near the stopping target will have a single track reconstructed upstream, and two tracks reconstructed downstream, provided the delta ray has enough energy and angle to be reconstructed. So, if N_{11} is the number of events in a given simulation with one track upstream and one track downstream of the stopping target, and N_{12} is the number of events with one track upstream and two tracks downstream, then

$$R_\delta = \frac{N_{12}}{N_{11}} \quad (8.15)$$

is a measure of the rate of delta ray production.

This is not the only type of event which can produce an event topology like this, however, especially if no fiducial cuts are imposed on the tracks. For example, the hits from a high-angle track could be poorly reconstructed as two separate tracks, a positron could backscatter downstream and return to the detector, etc. Therefore R_δ should not be used directly. To determine the contribution from other processes, several GEANT executables were produced with the delta ray production rates multiplied⁶ by 10, 5, 3, and zero, for

⁶Production rates were increased by modifying the tracking step sizes in the simulation, as described in Sec. 6.2.1.




Figure 8.11: Measured delta rate $R_\delta = N_{12}/N_{11}$ for different delta rate boost factors. The vertical axis is in units of 10^{-2} .

electrons and positrons only, to avoid impacting the stopping distribution. These executables were used to run “upstream stops” simulations of 500,000 events each, using the same parameters as the standard upstream stops simulation, and these were analyzed in the same way. R_δ was measured for each simulation, as well as for the standard upstream stops simulation (which has a “boost factor” of 1). These are shown in Fig. 8.11, along with a linear fit; R_δ is clearly responding directly to the increase in delta rate in the simulation. The slope p_1 of the fit indicates the relationship between R_δ and the actual delta production rate, and the intercept p_0 indicates the other contributions to R_δ .

R_δ was also measured for data (0.01432 ± 0.00003). If we make the conservative assumption that any difference in R_δ between data and MC is due to an actual difference in the delta ray production rate, then the results in Fig. 8.11 can be used to estimate the mismatch in delta rate between data and MC, by determining the “boost factor” needed to reproduce the R_δ seen in data: $\text{boost} = (R_\delta - p_0)/p_1$. This gives a boost factor for data of 1.18 ± 0.04 ; the delta rate in MC is about 18% lower than that in data.

It is likely that this 18% discrepancy is not due entirely to the delta ray production physics. For example, if the stopping target thickness is different between MC and data, that will affect all interactions in the target, not just the delta rays; in that case, the contribution from the p_0 term would change as well, and the difference between data and MC would be much smaller than 18%. In the extreme case where a direct comparison is possible, which might be true if the target thickness is the cause, the difference would be

less than 4%. The conservative 18% difference will be used here.

Systematic Error in the Bremsstrahlung Production Rate

The primary effect of bremsstrahlung on the positron is to carry away some of its energy. Low energy bremsstrahlung contributes to the approximately continuous slowing down of the positron, and so is accounted for by the energy calibration; this study will be concerned with “hard” bremsstrahlung, essentially a discrete process which results in a significant and sudden change in the positron’s momentum.

The (hard) bremsstrahlung production rate can be studied for MC-data comparison using the “upstream stops” measurements (again, see Sec. 6.3.1), looking for positrons which experience a large change in momentum as they pass through the stopping target. A “large change” is defined here as $\Delta p \equiv p_{ds} - p_{us} < \langle \Delta p \rangle - 1 \text{ MeV}/c$, where $\langle \Delta p \rangle$ is the *peak* or most probable energy loss (this form is used to reduce sensitivity to the “soft” energy loss processes which shift the distribution); essentially the study is concerned with positrons which lose more than about 1 MeV/c through the stopping target. If N_h is the number of hard energy loss events, and N_{tot} is the total number of reconstructed events, then

$$R_B = \frac{N_h}{N_{tot}} \quad (8.16)$$

is a measure of the rate of hard energy loss events. The hard energy loss rates in data and simulation are in good agreement qualitatively, as shown in Fig. 6.4, so the R_B measurements should be similar.

Bremsstrahlung is not the only process that can cause this type of hard energy loss, of course; delta rays and even nuclear interactions can play some role. Using the same method as for the delta ray production rate measurement, above, several GEANT executables were produced with the bremsstrahlung production rates multiplied by 10, 5, 3, and zero, for electrons and positrons only. These executables were used to run “upstream stops” simulations of 4,500,000 events each—the effect of bremsstrahlung on the spectrum was found to be much smaller than that of delta rays—using the same parameters as the standard upstream stops simulation, and these were analyzed in the same way. R_B was measured for each simulation, as well as for the standard upstream stops simulation (which has a “boost factor” of 1); these are shown in Fig. 8.12, along with a linear fit; R_B is clearly responding directly to the increase in bremsstrahlung production in the simulation. The slope p_1 of the fit indicates the relationship between R_B and the actual bremsstrahlung production rate, and the intercept p_0 indicates other contributions to R_B .

R_B was also measured for data (0.0142 ± 0.0001). Again making the conservative assumption that any difference in R_B between data and MC is due to an actual difference in the bremsstrahlung production rate, then the results in Fig. 8.12 can be used to estimate the mismatch in bremsstrahlung rate between data and MC, by determining the “boost factor”




Figure 8.12: Measured hard energy loss rate $R_B = N_h/N_{tot}$ for different bremsstrahlung rate boost factors.

needed to reproduce the R_B seen in data: $\text{boost} = (R_B - p_0)/p_1$. This gives a boost factor for data of 1.001 ± 0.019 ; the bremsstrahlung rate is consistent with the data to about 2%.

Because the uncertainty on the match of the bremsstrahlung rate between data and MC is so large compared to its measurement—much larger than the relative uncertainties for other systematic measurements—that uncertainty will need to be carried through to the scale factor as well, and ultimately through to the systematic uncertainty. This will be done by determining the scale factors for data boost factors of $1.001 - 0.019 = 0.981$ and $1.001 + 0.019 = 1.020$, and comparing the scaled systematic uncertainties to those using the central value of 1.001.

Systematic Effects

To determine the sensitivity of the muon decay parameters to the delta ray and bremsstrahlung production rate, GEANT executables with the delta ray and bremsstrahlung rates increased by factors of 3 and 10 were used for simulations with the same parameters as a standard surface muon simulation. These were analyzed in the usual way, and relative energy calibration applied using the standard simulation as the base. The scale factors for these exaggerations are

$$\left| \frac{(\text{Exaggerated MC}) - (\text{Standard MC})}{(\text{Data}) - (\text{Standard MC})} \right| \quad (8.17)$$

giving 11 and 50 for the delta ray exaggeration, and 2000 and 9000 for the bremsstrahlung exaggeration.

Boost factor	Delta ray		Bremsstrahlung	
	3	10	3	10
US slope	19 ± 2 keV	74 ± 1	1 ± 1 keV	6 ± 1
US intercept	8 ± 4 keV	26 ± 2	-2 ± 2 keV	-4 ± 2
DS slope	20 ± 5 keV	111 ± 3	11 ± 3 keV	30 ± 3
DS intercept	3 ± 8 keV	90 ± 5	15 ± 5 keV	41 ± 5

Table 8.24: Results of the relative energy calibration between simulation with increased delta ray or bremsstrahlung production and standard surface muon simulation.

Boost factor	Delta ray	
	3	10
$\Delta\rho$	$(-17.1 \pm 3.0) \times 10^{-4}$	$(-79.7 \pm 3.0) \times 10^{-4}$
$\Delta\delta$	$(-9.8 \pm 5.3) \times 10^{-4}$	$(-76.3 \pm 5.2) \times 10^{-4}$
$\Delta\xi$	$(15.1 \pm 6.5) \times 10^{-4}$	$(103.5 \pm 6.5) \times 10^{-4}$
χ^2/dof	2423/2463	3889/2463
Renormalization	1	1
Scale factor	11	50

Table 8.25: Unscaled Michel parameter fits for the delta ray production rate systematic tests, fit by standard simulation.

Results of the energy calibration are shown in Tab. 8.24. Not surprisingly, increasing the delta ray or bremsstrahlung rates affects the energy loss seen by the positrons. The unscaled Michel parameter fit results are shown in tables 8.25 and 8.26, and plots of the fit residuals

$$\frac{\text{Fit} - \text{Data}}{\sigma}$$

are shown in Figs. 8.13 and 8.14. The plots illustrate that the Michel fitter is able to fit the boost factor 3 simulations reasonably well in both cases, but the residual plots of the boost factor 10 fits have a lot of structure, indicating very poor fits; this is reflected in the χ^2/dof values in tables 8.25 and 8.26. The fit results with the scale factor applied are in tables 8.27 and 8.28. These last represent the systematic uncertainties in the reconstructed Michel parameters due to errors in the MC delta ray or bremsstrahlung production rates. Notice that, in spite of the poor quality of the boost factor 10 fits, after scaling they result in systematic uncertainties consistent with those taken from the boost factor 3 fits, indicating

Boost factor	Bremsstrahlung	
	3	10
$\Delta\rho$	$(-81.7 \pm 3.0) \times 10^{-4}$	$(-349.1 \pm 3.0) \times 10^{-4}$
$\Delta\delta$	$(-62.5 \pm 5.3) \times 10^{-4}$	$(-268.6 \pm 5.2) \times 10^{-4}$
$\Delta\xi$	$(28.6 \pm 6.5) \times 10^{-4}$	$(76.0 \pm 6.6) \times 10^{-4}$
χ^2/dof	2539/2463	3452/2463
Renormalization	1	1
Scale factor	2000	9000

Table 8.26: Unscaled Michel parameter fits for the bremsstrahlung production rate systematic tests, fit by standard simulation.




Figure 8.13: Residuals of Michel parameter fits for increased delta ray production, against standard surface muon simulation. The top plot is for a delta ray boost factor of 3, and the bottom plot is for a boost factor of 10.

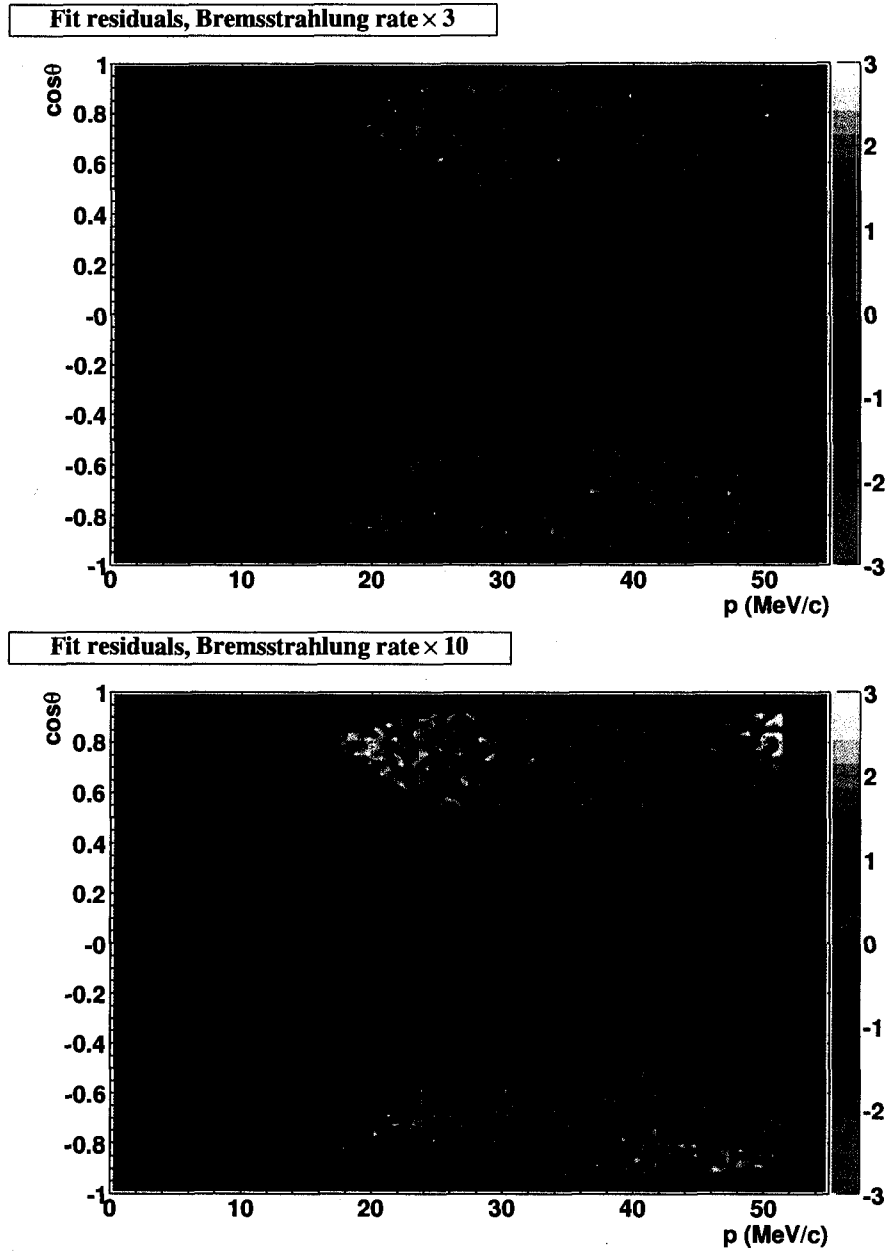


Figure 8.14: Residuals of Michel parameter fits for increased bremsstrahlung production, against standard surface muon simulation. The top plot is for a bremsstrahlung boost factor of 3, and the bottom plot is for a boost factor of 10.

Delta ray		
Boost factor	3	10
$\Delta\rho$	$(-1.5 \pm 0.3) \times 10^{-4}$	$(-1.59 \pm 0.06) \times 10^{-4}$
$\Delta\delta$	$(-0.9 \pm 0.5) \times 10^{-4}$	$(-1.5 \pm 0.1) \times 10^{-4}$
$\Delta\xi$	$(1.4 \pm 0.6) \times 10^{-4}$	$(2.1 \pm 0.1) \times 10^{-4}$

Table 8.27: Michel parameter fits for the delta ray production systematic, after applying the scale factors. These represent the systematic uncertainty due to the error in the MC delta ray production rate; the systematics for the boost factor of 3 will be used (see text).

Bremsstrahlung		
Boost factor	3	10
$\Delta\rho$	$(-0.04 \pm 0.74) \times 10^{-4}$	$(-0.04 \pm 0.70) \times 10^{-4}$
$\Delta\delta$	$(-0.03 \pm 0.56) \times 10^{-4}$	$(-0.03 \pm 0.54) \times 10^{-4}$
$\Delta\xi$	$(0.02 \pm 0.25) \times 10^{-4}$	$(0.01 \pm 0.15) \times 10^{-4}$

Table 8.28: Michel parameter fits for the bremsstrahlung production systematic, after applying the scale factors. These represent the systematic uncertainty due to the error in the MC delta ray production rate; the systematics for the boost factor of 3 will be used (see text). Uncertainties listed here take into account the uncertainty on the scale factor.

that these measurements of systematic uncertainties are robust in spite of the large impact of the exaggerations. The results from the higher quality boost factor 3 fits will be used for the final systematic uncertainties. Uncertainties in the bremsstrahlung systematic uncertainty were determined using the uncertainty in the data's bremsstrahlung measurement.

8.3.2 Outside Materials

The material outside the sensitive region of the detector (the beam pipe, the yoke, etc.) could influence the reconstructed decay spectrum by scattering particles back into the tracking region, adding hits and making reconstruction more difficult. This is compensated for by the inclusion of these materials in the simulation; the degree to which the simulation does not reproduce these represents a systematic uncertainty.

The probability of particles backscattering through the tracking region was determined using PC Time-of-Flight measurements, as described in Sec. 6.3.3. Backscatters for up-stream decays (scattering mainly on the beam pipe and support structures) and for down-stream decays (scattering from the steel yoke, which is not simulated) were studied separately.

Systematic Effects

To exaggerate the number of backscatters, a simulation was run with an aluminum disk placed downstream of the detector. The number of particles backscattering through the tracking region due to the added aluminum disk was again determined using PC Time-of-Flight measurements. As with the study of standard data and simulation, the PC TOF histogram for downstream decays was used as the background and was subtracted from the

	Outside Material
$\Delta\rho$	$(6.8 \pm 3.0) \times 10^{-4}$
$\Delta\delta$	$(-4.6 \pm 5.3) \times 10^{-4}$
$\Delta\xi$	$(10.9 \pm 6.5) \times 10^{-4}$
χ^2/dof	2423/2463
Renormalization	1
Scale factor	17

Table 8.29: Unscaled Michel parameter fits for the effect of outside materials: a standard simulation with the addition of an aluminum plate downstream fit to standard simulation.

	Outside Material
$\Delta\rho$	$(0.39 \pm 0.17) \times 10^{-4}$
$\Delta\delta$	$(-0.27 \pm 0.31) \times 10^{-4}$
$\Delta\xi$	$(0.63 \pm 0.38) \times 10^{-4}$

Table 8.30: Michel parameter fit of a simulation with an aluminum plate downstream against standard simulation, after applying the scale factor. These represent the systematic uncertainties due to errors in the simulation in the materials outside the detector.

normalized downstream PC TOF histogram of the simulation with the aluminum in place; the latter is shown in Fig. 8.15, before and after background subtraction. All PC TOF histograms are normalized by the number of muon windows compared to the number in the simulation.

The normalized PC TOF histogram with the aluminum plate added was integrated, and the integral of the background histogram was subtracted, as before. The result was a normalized backscatter count of $(3.083 \pm 0.002) \times 10^6$ from the aluminum plate. The difference in backscatters between standard simulation and data is due mainly to the yoke; the normalized count of yoke backscatters in data (Tab. 6.5) was $(0.178 \pm 0.001) \times 10^6$. Thus the scale factor for this study is $3.083/0.178 = 17$.

The simulation with the added aluminum plate was analyzed and the spectrum fit against standard simulation; Tab. 8.29 shows the fit results. The systematic uncertainty in the measured Michel parameters due to errors in the simulation in the materials outside the detector is then found by applying the scale factor; Tab. 8.30 gives the final systematic uncertainty.

8.4 Reconstruction Resolution

The Reconstruction Resolution category describes the way the measured Michel parameters are affected by differences in angle and momentum resolution between data and simulation. The systematic uncertainties due to these effects are summarized in Tab. 8.31. The difference in momentum resolution between data and simulation also led to a correction to be applied to the measured Michel parameters.

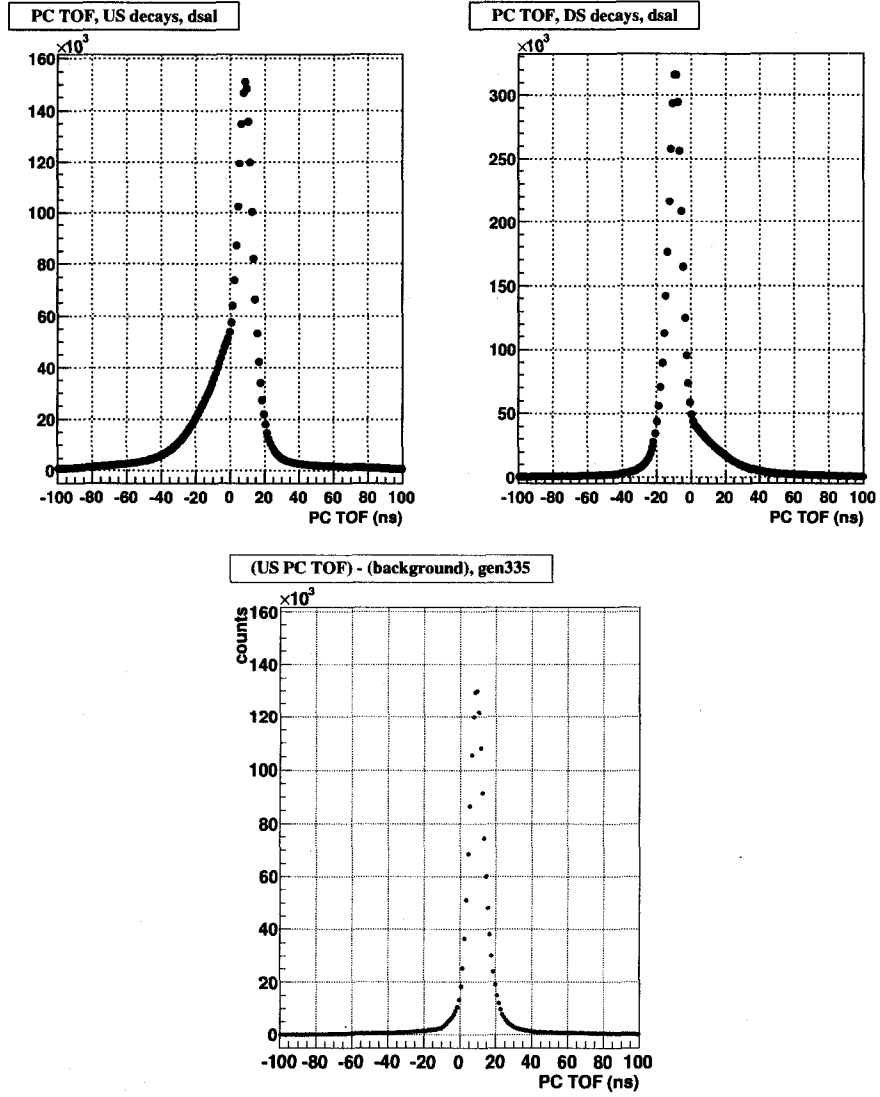


Figure 8.15: Normalized backscatter counts for simulation with an aluminum plate added downstream. The PC TOF histogram for downstream decays from standard simulation was used as the background.

Reconstruction Resolution	ρ	δ
Angle Resolution (8.4.1)	0.09	0.41
Momentum Resolution (8.4.1)	0.24	0.26
Total	0.26	0.49

Table 8.31: Summary of systematic uncertainties related to reconstruction resolution, in units of 10^{-4} .

8.4.1 Momentum and Angle Resolution

The muon decay distribution is smooth, and varies slowly compared to the reconstruction resolution. As a result, resolution in energy and angle should not significantly affect the shape of the spectrum within the fiducial region. However, the momentum resolution does significantly change the shape of the high-momentum edge of the spectrum, near 52.8 MeV/c; since the energy calibration uses this edge to compare the energy scales in data and simulation, it could well be vulnerable to a data-MC mismatch in the momentum resolution, and this is what is seen. Since there are no sharp features in the angle distribution, angle resolution is not expected to be a major source of error in the *TWIST* experiment, and again this is what is observed.

Systematic Uncertainty in the Momentum and Angle Resolutions

The “upstream stops” data (Sec. 6.3.1) were analyzed in bins of $(p, \cos \theta)$. The widths of $\Delta\theta = \theta_{DS} - \theta_{US}$ and $\Delta p = p_{DS} - p_{US}$ distributions include effects of scattering, energy straggling, and reconstruction resolution. Conservatively, we take the difference in widths measured for data and MC to be due entirely to differences in resolution. (Even if it is not, the end effects on the Michel parameter reconstruction should be similar to those of resolution.)

Resolution is defined in terms of a Gaussian σ parameter; however, the $\Delta\theta$ and Δp distributions are not Gaussian. To estimate the resolution, the full width at half maximum (FWHM) of each $\Delta\theta$ and Δp distribution was measured for each $(p, \cos \theta)$ bin, and divided by $2\sqrt{2\ln 2} \approx 2.35$, which is the relationship between σ and FWHM in a Gaussian distribution. This approximation should be good enough for the purposes of this measurement. Uncertainties in FWHM were estimated using the same formula as for arithmetic RMS: $\text{FWHM}/\sqrt{2N}$. Since reconstruction resolution is expected to follow a $1/\sin \theta$ dependence (see Sec. 5.5.3), resolutions were plotted against $1/\sin \theta$ for each of several bins of momentum; see Figs. 8.16 and 8.17. Since the distributions were binned in $1/\cos \theta$, the $1/\sin \theta$ bins are of irregular size.

Ultimately, in order to test the sensitivity of the Michel parameters to the energy and angle resolution, the reconstructed track parameters were “smeared” before inclusion in the spectra (see below). Any smearing applied adds to the measured resolution in quadrature. To determine what width of smearing needs to be applied to the simulation in order to make its resolution similar to data (or vice versa, if the simulation is wider), the quadrature differences $\Delta\sigma_Q$ of the plots in Figs. 8.16 and 8.17 were taken for each $(p, \cos \theta)$ point, according to

$$\Delta\sigma_Q = \begin{cases} \sqrt{\sigma_D^2 - \sigma_M^2} & \sigma_D > \sigma_M \\ -\sqrt{\sigma_M^2 - \sigma_D^2} & \sigma_M > \sigma_D \end{cases} \quad (8.18)$$

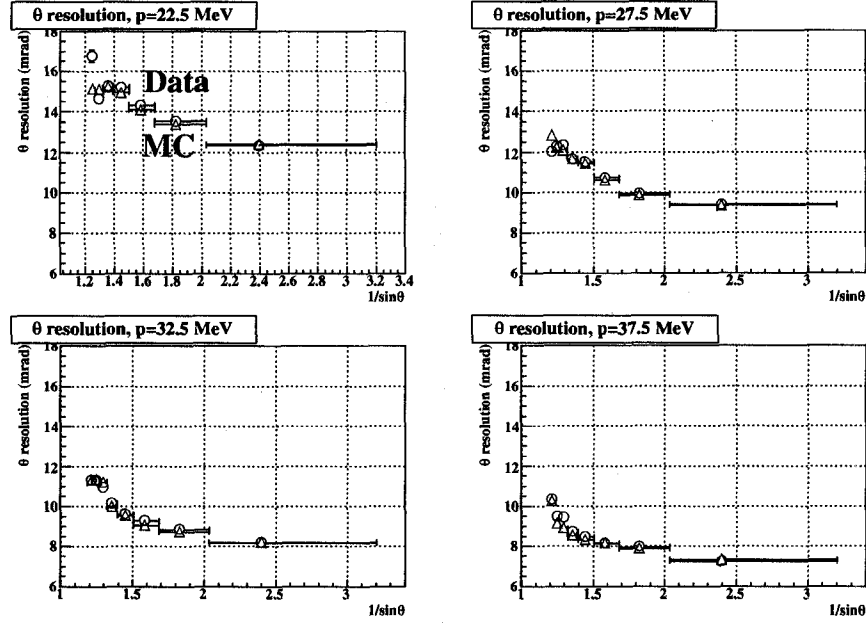


Figure 8.16: Angle resolution vs $1/\sin \theta$, for four 5 MeV/c momentum bins, for data (red circles) and simulation (blue triangles). Vertical error bars are included but are very small.

where σ_D and σ_M are the resolutions for data and MC; these quadrature differences are shown in Figs. 8.18 and 8.19. (Herein the notations $\Delta\sigma_{Q\theta}$ and $\Delta\sigma_{Qp}$ will denote the quadrature differences in angle and momentum resolution.) An attempt was made to parameterize the behaviour of $\Delta\sigma_Q$ vs $1/\sin \theta$, using linear fits of the form

$$\Delta\sigma_Q = m_s \left(\frac{1}{\sin \theta} - \frac{1}{\sin \theta_0} \right) + k_s \quad (8.19)$$

where $1/\sin \theta_0$ is a location near the middle of the fit domain; using this instead of the standard $y = mx + b$ form reduces the correlations between the fitted slope and the constant term. The fit results are included in Figs. 8.18 and 8.19; clearly they are a crude approximation, but they will suffice as a characterization for this analysis. The fitted slopes $m_s(p)$ for each momentum slice are shown in Figs. 8.20 and 8.21, with averages \bar{m}_s taken by fitting constants to each figure; the error bars in the figures are underestimates as they do not take the scatter of the original measurements into account. Again, the averages are sufficient for this analysis. Similarly, the fitted offsets $k_s(p)$ for each momentum slice are shown in Figs. 8.22 and 8.23. Here linear fits similar to Eq. (8.19) were applied, of the form

$$k_s = m_p(p - p_0) + k_p \quad (8.20)$$

where again p_0 is a location near the middle of the fit domain.

Although this is the best information we have available on the angle resolution in data and simulation, the momentum resolution was also measured using the absolute energy calibration technique used for the previous *TWIST* measurements, as described in Sec. 5.5.3.

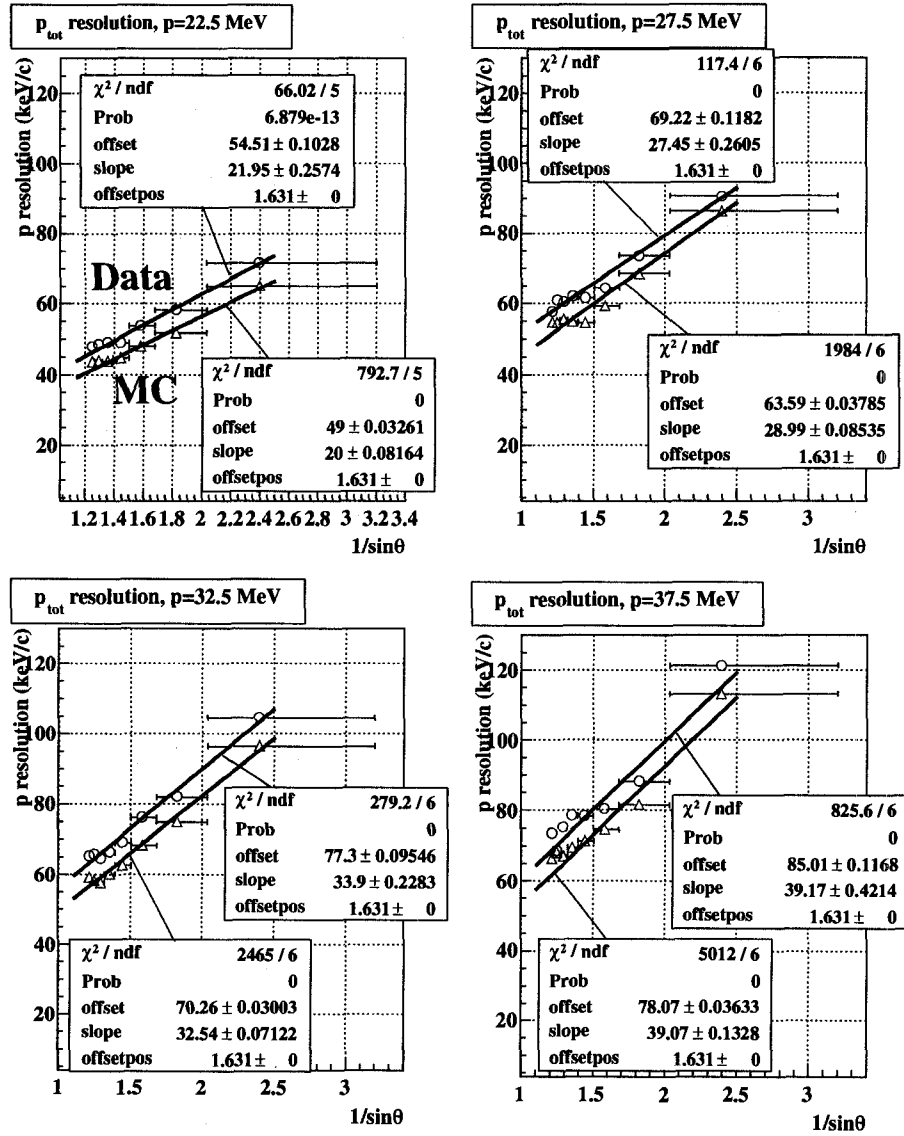


Figure 8.17: Momentum resolution vs $1/\sin\theta$, for four 5 MeV/c momentum bins, for data (red circles) and simulation (blue triangles). Vertical error bars are included but are very small.




Figure 8.18: Data-MC quadrature differences in angle resolution, vs $1/\sin\theta$, for four momentum bins.

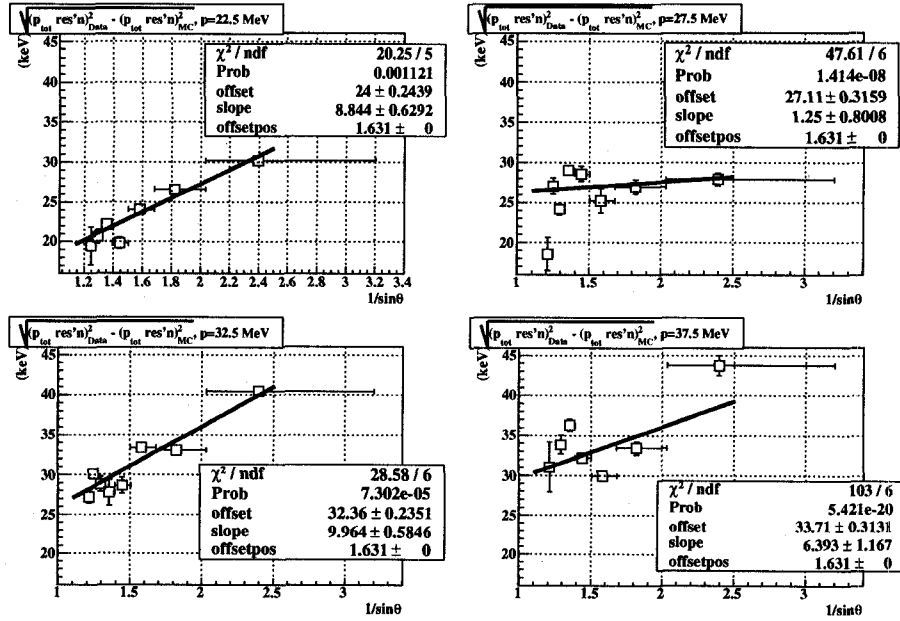


Figure 8.19: Data-MC quadrature differences in momentum resolution, vs $1/\sin\theta$, for four momentum bins.




Figure 8.20: Slopes of $\Delta\sigma_{Q\theta}$ vs $1/\sin\theta$, from Fig. 8.18, with constant fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Q\theta}$ values, and are underestimates.




Figure 8.21: Slopes of $\Delta\sigma_{Qp}$ vs $1/\sin\theta$, from Fig. 8.19, with constant fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Qp}$ values, and are underestimates.




Figure 8.22: Offsets of $\Delta\sigma_{Q\theta}$ vs $1/\sin\theta$, from Fig. 8.18, with fit. The vertical error bars do not account for the scatter in the $\Delta\sigma_{Q\theta}$ values, and are underestimates.




Figure 8.23: Slopes of $\Delta\sigma_{Qp}$ vs $1/\sin\theta$, from Fig. 8.19, with constant fit.

	Angle	Momentum
\bar{m}_s	-1.4 ± 0.1 mrad	7.5 ± 0.4 keV
m_p	0.02 ± 0.01 mrad/MeV	0.71 ± 0.02 keV/MeV
k	1.10 ± 0.06 mrad	29.5 ± 0.1 keV
$1/\sin \theta_0$	1.631	
p_0	30 MeV/c	

Table 8.32: Constants for parameterization of data-MC resolution differences Eq.(8.21).

This technique measured the width of the high momentum edge, and fit the widths along the edge to a $1/\sin \theta$ function to obtain the momentum resolution at $1/\sin \theta = 1$. While this has nothing to say about the behaviour of the resolution in the “bulk” spectrum, it is considered a robust measurement of the resolution at the momentum edge. It can therefore be used as a measure of the overall difference between the data and MC resolution.

Systematic Effects

Under the assumption that the energy and angle behaviour of the resolution difference are independent of each other, Eqs. (8.19) and (8.20) can be combined to parameterize the overall behaviour of the data-MC quadrature difference in resolution:

$$\Delta\sigma_Q = \bar{m}_s \left(\frac{1}{\sin \theta} - \frac{1}{\sin \theta_0} \right) + m_p(p - p_0) + k \quad (8.21)$$

where $1/\sin \theta_0$ and p_0 are the locations used for the associated fits, \bar{m}_s is the average slope from Fig. 8.20 (angle) or 8.21 (momentum), and m_p and k are the slopes and constants of the fits in Fig. 8.22 (angle) or 8.23 (momentum). These constants are listed in Tab. 8.32. It is important to keep in mind that, contrary to what the uncertainties in Tab. 8.32 suggest, this parameterization is very rough, but will suffice for estimating the systematic uncertainties due to the resolution mismatch.

The sensitivity of the Michel parameters to the data-MC difference in resolution can be measured by smearing the MC, using Eq. (8.21), with the coefficients scaled by various factors.

For the angle resolution, this was applied directly. For each event in the standard simulation, the angle θ of each reconstructed track was smeared using a random Gaussian distribution with the width given by Eq. (8.21), with the coefficients \bar{m}_s , m_p , and k all multiplied by 10. The resulting smeared simulation was run through the standard treesum and energy calibration, and the final decay spectrum was fit against the unmodified spectrum. Raw Michel parameter fit results are listed in Tab. 8.33, and the effect on the Michel parameters after applying scale factor and renormalization are in Tab. 8.34.

The comparative energy calibration turned out to be very sensitive to the momentum resolution. Figure 8.24 shows the change in momentum edge as measured by the comparative energy calibration, for a momentum smearing using the coefficients in Tab. 8.32 without

Angle Resolution	
$\Delta\rho$	$(-0.91 \pm 3.01) \times 10^{-4}$
$\Delta\delta$	$(-4.09 \pm 5.31) \times 10^{-4}$
$\Delta\xi$	$(-3.31 \pm 6.55) \times 10^{-4}$
χ^2/dof	1575/2463
Renormalization	0.800
Scale factor	10

Table 8.33: Unscaled Michel parameter fits for the angle resolution systematic test, fit by standard simulation.

Angle Resolution	
$\Delta\rho$	$(-0.09 \pm 0.24) \times 10^{-4}$
$\Delta\delta$	$(-0.41 \pm 0.42) \times 10^{-4}$
$\Delta\xi$	$(-0.33 \pm 0.52) \times 10^{-4}$

Table 8.34: Michel parameter fits for the angle resolution systematic, after applying the scale factor and renormalizing the uncertainties. These represent the systematic uncertainty due to the difference in resolution between data and simulation.

scaling them; this is the same order as the difference in resolution between data and MC. Already the behaviour of the edge is nonlinear, curling up towards $1/|\cos\theta| = 1$. This effect on the energy calibration causes significant changes to the reconstructed Michel parameters, as shown in Tab. 8.35.

Because of concerns over nonlinearity of the effect of momentum resolution on the Michel parameters, and evidence that the sensitivity of the Michel parameters to the momentum resolution is quite significant, standard simulation was smeared several times with different scalings of the coefficients in Eq. (8.21). As described above, full treesums and energy calibration were done for each smearing. In addition, the endpoint resolution was measured for each smearing using the absolute energy calibration software. This measure of resolution was used for the scaling. Figure 8.25 shows the changes in the Michel parameters, from Michel fits vs standard simulation, for each smearing, plotted against the endpoint resolution. The Michel parameters turn out to depend linearly on endpoint resolution, which provides confidence in this procedure.

The systematic effects due to the data-MC difference in momentum resolution can be determined by reading the Michel parameter shifts from the plots in Fig. 8.25 at the appro-

Momentum Resolution		
	Uncalibrated	Calibrated
$\Delta\rho$	$(0.02 \pm 2.46) \times 10^{-4}$	$(-1.50 \pm 2.46) \times 10^{-4}$
$\Delta\delta$	$(0.02 \pm 4.34) \times 10^{-4}$	$(-1.76 \pm 4.33) \times 10^{-4}$
$\Delta\xi$	$(0.16 \pm 5.35) \times 10^{-4}$	$(3.99 \pm 5.35) \times 10^{-4}$
χ^2/dof	126/2463	126/2463

Table 8.35: Unscaled Michel parameter fits for $\times 1$ momentum smearing, before and after applying the comparative energy calibration.




Figure 8.24: Change in momentum edge measured by comparative energy calibration, for $\times 1$ momentum smearing (i.e. using the coefficients in Tab. 8.32 for smearing).

Momentum Resolution	
$\Delta\rho$	$(-1.22 \pm 0.09) \times 10^{-4}$
$\Delta\delta$	$(-1.31 \pm 0.11) \times 10^{-4}$
$\Delta\xi$	$(3.20 \pm 0.33) \times 10^{-4}$

Table 8.36: Systematic corrections and uncertainties due to the data-MC momentum resolution mismatch, determined from fits in Fig. 8.25, with renormalized errors. The central values will be used as corrections to be subtracted from the measured Michel parameters, since the resolution is known to be smaller in the simulation than in the data; the uncertainties in the corrections then represent the systematic uncertainties due to the resolution mismatch.

appropriate resolution. Since this is a smearing applied to the standard simulation, the appropriate resolution is the endpoint resolution for standard simulation plus the data-MC difference in endpoint resolution:

$$69.63 \text{ keV/c} + 5.07 \text{ keV/c} = 74.70 \text{ keV/c} \quad (8.22)$$

To determine the shifts at this point, formulas such as

$$\Delta\rho = m_\sigma(\sigma - (74.70 \text{ keV/c})) + k_\sigma, \quad (8.23)$$

and similar for δ and ξ , were fit to the plots. The values of k_σ from these fits then represent the systematic uncertainties for the data-MC mismatch in momentum resolution. Since the data in Fig. 8.25 are highly correlated, the errors on k_σ can be renormalized by $\sqrt{\chi^2/\text{dof}}$ as usual. The resulting systematic uncertainties are listed in Tab. 8.36.

Since the momentum resolution is systematically smaller in simulation than in data, these systematic effects can be used to apply corrections to the measured Michel parameters

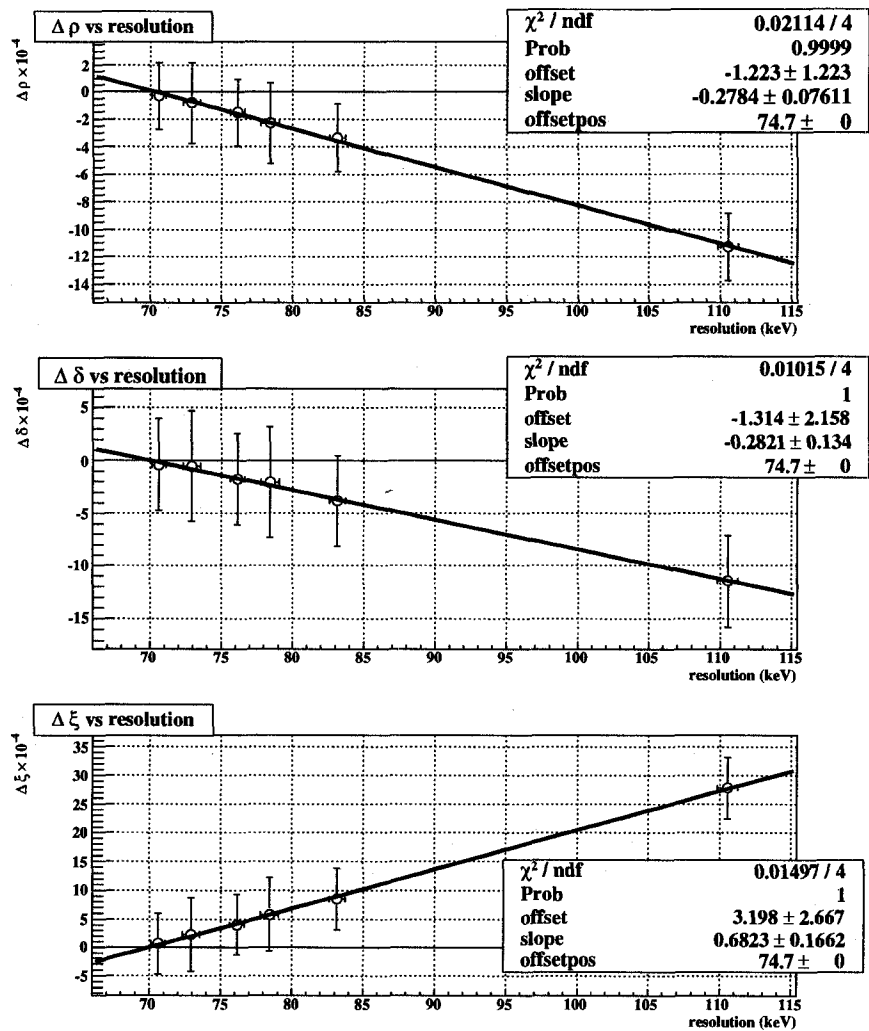


Figure 8.25: Changes in Michel parameters vs endpoint resolution, for several momentum smearings. Error bars are statistical, but do not account for the highly correlated nature of the fits.

to account for the resolution difference. The measured shifts in Michel parameters are from the differences between smeared and unsmeared spectra; here the smeared spectrum corresponds to the data, and the unsmeared spectrum to the simulation. The spectrum fit measured Data – MC, so these amounts need to be subtracted from the final measurements to account for the resolution mismatch. In this case, the central values of the shifts represent the correction to be subtracted, and the uncertainties of the shifts represent the remaining systematic uncertainties due to the difference in resolution.

As the systematic error due to the mismatch in angle resolution is small and consistent with zero, it will not be used for a correction.

8.5 Spectrometer Geometry

The Spectrometer Geometry section includes effects related to the alignment of the detector components, and to the knowledge of the overall length scales of the detector. The systematic uncertainties due to these effects are summarized in Tab. 8.37.

Spectrometer Geometry	ρ	δ
Translation Alignment (8.5.1)	0.03	0.01
Rotational Alignment (8.5.1)	0.01	0.01
U/V Width Scale (8.5.2)	0.16	0.17
Z Length Scale (8.5.2)	0.25	0.24
Total	0.30	0.29

Table 8.37: Summary of systematic uncertainties related to spectrometer geometry, in units of 10^{-4} .

8.5.1 Alignment

The effect of truly random misalignments, at least on the scale of our alignment uncertainties, should merely be to degrade the resolution of the detector; the tracking residuals should simply be smeared, without being biased. Systematic uncertainties related to resolution are discussed in Sec. 8.4. Thus, the only alignment systematics we should be considering are *correlated* misalignments: translational “shear”, and rotational “corkscrew.” A description of the alignment procedure can be found in Sec. 4.3.1.

Translational “shear” means the systematic translational offset of the detector planes, with the offset changing linearly from one end of the detector to the other. It is approximately equivalent to a rotation of the entire detector with respect to the magnetic field. The detector is explicitly aligned to the magnetic field as part of the standard alignment process, to a precision of 0.03 mrad, and the alignment is implemented in the analysis by applying a shearing offset to the detector plane positions.

Rotational “corkscrew” means the systematic rotation of the detector planes about the Z axis, with the rotation changing linearly from one end of the detector to the other. Unlike

translational shear, the amount of corkscrew in the aligned detector has not been measured previously. The typical difference between the largest negative and largest positive rotational corrections⁷ is about 1.7 mrad, and this can be used as an estimate of the possible corkscrew introduced by the alignment procedure; this is similar to an estimate from Wayne Faszer⁸ on the maximum corkscrew physically possible in the detector, based on the amount of clearance between the alignment rods and the citals.

Measurement of the Systematic Uncertainties

Several possible skew and corkscrew exaggerations were tested, by applying the exaggerated misalignments to the appropriate alignment files and analyzing 10,000 simulated standard surface muon events; the mean and RMS of the χ^2/dof reported by the helix fitter was then checked to determine how much the applied misalignments degraded the fit quality. Since the same 10,000 events were used for each test, changes in the mean χ^2/dof are significant at much smaller levels than the uncertainties would indicate. For simplicity, the same effect was applied to both U and V planes, so the shear or corkscrew depends only on the Z of the plane.

The translational alignment file was modified using

$$\begin{aligned}\phi_{shear} &= (0.03 \times 10^{-3} \text{ rad})S \\ dU_i &= z_i \tan(\phi_{shear})\end{aligned}\tag{8.24}$$

where ϕ_{shear} is the angle by which the detector was sheared (in mrad), S is the scale factor used for the exaggeration, dU_i is the amount added to the translational alignment offset of plane i in the alignment file (for both U and V planes), and z_i is the z position of plane i . The uncertainty in the alignment of the detector to the magnetic field, 0.03 mrad, is the quantity being exaggerated in this test.

The rotational alignment file was modified using

$$\begin{aligned}\phi_{twist} &= (1.7 \times 10^{-3} \text{ rad})S \\ d\phi_i &= (z_i/L)\phi_{twist}\end{aligned}\tag{8.25}$$

where ϕ_{twist} is total twist (or ‘‘corkscrew’’) applied to the detector (the effect is that the last plane has been rotated by an angle ϕ_{twist} with respect to the first), $d\phi_i$ is the amount added to the rotational alignment offset of plane i in the alignment file, L is the total distance from the first DC plane to the last, and other variables are as above. The conservative upper

⁷Shirvel Stanislaus, ‘‘Alignment for 2004 Data’’, TWIST Alignment internal forum, 20 May, 2005.
https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_alignments&key=1116622372

⁸Robert MacDonald, ‘‘Mechanical limits on rotational misalignment’’, TWIST Alignments internal forum, 7 July, 2006.
https://twist.phys.ualberta.ca/forum/view.php?bn=twist_alignments&key=1152316434

	Scale Factor S	Mean χ^2/dof	RMS χ^2/dof
Unmodified	—	2.507 ± 0.027	2.370 ± 0.019
Translation			
0.15 mrad	5	2.512 ± 0.027	2.369 ± 0.019
0.21 mrad	7	2.528	2.381
0.30 mrad	10	2.542	2.502
0.60 mrad	20	2.541	2.380
Rotation			
8.5 mrad	5	2.517	2.385
17 mrad	10	2.524	2.392
68 mrad	40	2.527	2.397

Table 8.38: Effects of translational (shear) and rotational (corkscrew) exaggerations on helix fit χ^2/dof . See text for test details. Uncertainties are the same for each test. However, since the same events are being analyzed in each test, the mean and RMS χ^2/dof can be compared more or less directly. The scale factors chosen for the shear translational and rotational systematic studies are shown in bold.

limit on the possible corkscrew of the detector, 1.7 mrad over the length of the detector, is the quantity being exaggerated.

Table 8.38 shows the effects of shear and corkscrew exaggerations with various scale factors on the mean and RMS χ^2/dof returned by the helix fitter. In summary, the shear has a relatively large effect on the χ^2/dof , which limits the exaggeration we can use. (This is not a surprise, as the effect on the χ^2/dof is how the magnetic field alignment is measured.) The corkscrew has a much smaller impact on the χ^2/dof ; it takes a fairly large exaggeration before the χ^2/dof notices. In both cases, reasonable scale factors can be chosen which still have “small” χ^2/dof effects. Scale factors of 7 for the translational systematic and 40 for the rotational systematic were chosen.

For the translational alignment systematic, a new translational correction file was created by applying a shear of 0.21 mrad, according to Eq. 8.24. For the rotational alignment systematic, a new rotational correction file was created by applying a corkscrew of 68 mrad (end to end), according to Eq. 8.25. These were used to reanalyze the standard simulation.

Systematic Effects

The shear-applied analysis, representing the translational alignment exaggeration, and the corkscrew-applied analysis, representing the rotational alignment exaggeration, were each fit against the standard simulated spectrum to measure the effects on the Michel parameters. The raw fit results are shown in Tab. 8.39. Table 8.40 shows the fit results after applying the scale factors; these are the systematic uncertainties for translational and rotational alignment.

	Translational	Rotational
$\Delta\rho$	$(-0.021 \pm 0.246) \times 10^{-4}$	$(0.036 \pm 0.246) \times 10^{-4}$
$\Delta\delta$	$(-0.006 \pm 0.433) \times 10^{-4}$	$(0.052 \pm 0.433) \times 10^{-4}$
$\Delta\xi$	$(0.006 \pm 0.534) \times 10^{-4}$	$(-0.081 \pm 0.534) \times 10^{-4}$
χ^2/dof	150/2463	83.2/2463
Renormalization	0.247	0.184
Scale factor	7	40

Table 8.39: Unscaled Michel parameter fits for translational and rotational alignment systematics.

	Translational	Rotational
$\Delta\rho$	$(-0.030 \pm 0.087) \times 10^{-4}$	$(0.009 \pm 0.011) \times 10^{-4}$
$\Delta\delta$	$(-0.009 \pm 0.153) \times 10^{-4}$	$(0.013 \pm 0.020) \times 10^{-4}$
$\Delta\xi$	$(0.009 \pm 0.189) \times 10^{-4}$	$(-0.020 \pm 0.025) \times 10^{-4}$

Table 8.40: Michel parameter fits for translational and rotational alignment systematics, after applying scale factors and renormalizing the errors. These represent the systematic uncertainties due to uncertainties in the translational and rotational alignments.

8.5.2 Spectrometer Length Scales

The length (z) and width (u or v) of the *TWIST* detector are the scales used to transform the shape of a reconstructed helix into the transverse and longitudinal momenta of a positron track. Therefore errors in these length scales will change the shape of the reconstructed decay spectrum, and represent systematic uncertainties. (Random displacements of wires or wire planes should not represent a systematic uncertainty but merely an increase in the reconstruction resolution; moreover, as these random displacements are very small, their impact should be negligible.)

The z length of the DC stack (from one end of the upstream dense stack to the opposite end of the downstream dense stack) was measured using a gauge block and a dial gauge, and was found to be $215 \pm 25 \mu\text{m}$ short of nominal, consistent with the sum of the measurements of the thicknesses of the Citals in the top column.⁹ The total DC stack length was nominally 100 cm, so the fractional uncertainty in the z length scale is $(25 \mu\text{m})/(100 \text{ cm}) = 2.5 \times 10^{-5}$.

The wires were positioned in u or v by hand to an accuracy of 3–5 μm . However, these positioning errors cannot accumulate across a wire plane, because that accuracy was determined by measuring wire positions relative to their nominal positions using a traveling microscope; the error of the microscope in the total width of the wire plane is less than 2 μm , according to the manufacturer. So we can conservatively estimate the uncertainty in the u and v length scales to be 5 μm . The total width of a wire plane is 320 mm, so the fractional uncertainty in the u and v scales is $(5 \mu\text{m})/(320 \text{ mm}) = 1.6 \times 10^{-5}$.

⁹Wayne Faszer, “Detector Stack Length Measurement”, *TWIST Chambers internal forum*, 11 July, 2003. https://twist.phys.ualberta.ca/forum/view.php?site=twist&bn=twist_chambers&key=1057971936

	Length (p_z)	
$\Delta\rho$	$(10.33 \pm 2.46) \times 10^{-4}$	$(2.49 \pm 2.46) \times 10^{-4}$
$\Delta\delta$	$(10.87 \pm 4.35) \times 10^{-4}$	$(2.71 \pm 4.34) \times 10^{-4}$
$\Delta\xi$	$(-12.13 \pm 5.35) \times 10^{-4}$	$(-3.06 \pm 5.35) \times 10^{-4}$
χ^2/dof	150/2463	35/2463
Renormalization	0.247	0.119
Scale factor	40	10

	Width (p_t)	
$\Delta\rho$	$(-8.12 \pm 2.46) \times 10^{-4}$	$(-1.58 \pm 2.46) \times 10^{-4}$
$\Delta\delta$	$(-8.67 \pm 4.33) \times 10^{-4}$	$(-1.65 \pm 4.34) \times 10^{-4}$
$\Delta\xi$	$(9.60 \pm 5.35) \times 10^{-4}$	$(1.81 \pm 5.35) \times 10^{-4}$
χ^2/dof	116/2463	23/2463
Renormalization	0.217	0.097
Scale factor	50	10

Table 8.41: Unscaled Michel parameter fits for length and width scale systematics: standard simulation with distorted p_z or p_t fit by standard simulation. Two different scale factors are shown.

Measurement of Systematic Uncertainties

Since the only effect of changing the length scales of the detector is to change the conversion between helix parameters and track momentum, this systematic error can be studied using the decay spectra directly, without repeating the full analysis. New spectra were created by multiplying p_z or p_t of each event by a distortion factor immediately before including them in the decay spectrum and the energy calibration histogram.

Because the effects of distorting p_z and p_t are inherently non-linear in $(p, \cos\theta)$, each distortion was done using two different scale factors to test for approximate linearity: 1.00025 and 1.001 for p_z (scale factors of 40 and 10), and 1.00016 and 1.0008 for p_t (scale factors of 50 and 10). Energy calibration was applied using the distorted spectra. The resulting Michel spectra were fit against the standard simulation spectrum.

Systematic Effects

The raw measured changes in the Michel parameters are shown in Tab. 8.41. Table 8.42 shows the fit results after applying the scale factors; note that the rescaled results are consistent between scale factors, confirming that the systematic effects are approximately linear. These rescaled results are the systematic uncertainties for the length and width scales of the *TWIST* detector.

The energy calibration keeps the total momentum approximately the same after p_z or p_t distortions. Since $p_{tot}^2 = p_z^2 + p_t^2$, the effects of distortions in p_z and p_t should have similar magnitudes and opposite signs, and this is seen in Tab. 8.41.

		Length (p_z)	
Scale factor		40	10
$\Delta\rho$		$(0.258 \pm 0.015) \times 10^{-4}$	$(0.249 \pm 0.029) \times 10^{-4}$
$\Delta\delta$		$(0.272 \pm 0.027) \times 10^{-4}$	$(0.271 \pm 0.052) \times 10^{-4}$
$\Delta\xi$		$(-0.303 \pm 0.033) \times 10^{-4}$	$(-0.306 \pm 0.064) \times 10^{-4}$
		Width (p_t)	
Scale factor		50	10
$\Delta\rho$		$(-0.162 \pm 0.011) \times 10^{-4}$	$(-0.158 \pm 0.024) \times 10^{-4}$
$\Delta\delta$		$(-0.173 \pm 0.019) \times 10^{-4}$	$(-0.165 \pm 0.042) \times 10^{-4}$
$\Delta\xi$		$(0.192 \pm 0.023) \times 10^{-4}$	$(0.181 \pm 0.051) \times 10^{-4}$

Table 8.42: Michel parameter fits for length and width scale systematics, after applying scale factors and renormalizing the errors. These represent the systematic uncertainties due to uncertainties in the length and width scales of the *TWIST* detector; the results with the larger scale factors will be taken.

8.6 Beam Intensity

The Beam Intensity category includes effects related to the accuracy of the simulation of the muon and positron beam rates. The systematic uncertainties due to these effects are summarized in Tab. 8.43.

Beam Intensity	ρ	δ
Muon Beam Intensity (8.6.1)	0.06	0.10
Positron Beam Intensity (8.6.1)	0.00	0.11
Total	0.06	0.15

Table 8.43: Summary of systematic uncertainties related to beam intensity, in units of 10^{-4} .

8.6.1 Muon and Positron Rates

Because of the relatively low rate of particles entering the *TWIST* detector—on the order of 10^3 – 10^4 per second, slow compared to the 16 μ s event gate—most events consist of a single muon entering the detector and decaying.

In some cases, however, the hits from pile-up beam particles can interfere with the reconstruction. For example, when the pile-up particle arrives within about 200 ns of either the arrival or (especially) the decay of the muon, it can sometimes be problematic for the analysis software to recognize this and correctly discard the event. Hits from the pile-up particle may be confused with hits from the decay positron, degrading or biasing the track reconstruction. This and other effects can lead to a distortion in the reconstructed decay spectrum, and the magnitude of the distortion depends on the intensity of the particle beam.

The position and angle distributions of muons and positrons in the particle beam were both measured in data, and those measurements were used to reproduce the beam properties in simulation. The simulation correctly reproduces the effects of changing the beam rates.¹⁰

¹⁰Alexander Grossheim, “Re: input for 2007 run plan”, *TWIST* General internal forum. 12 February

	Data	MC	MC – Data	Relative Diff	Error (Hz)
Set 31	0.00439	0.00458	+0.00019	+0.043	+109
Set 32	0.00439	0.00454	+0.00015	+0.034	+87
Set 35	0.00478	0.00488	+0.00010	+0.022	+90
Set 36	0.00485	0.00497	+0.00012	+0.026	+73
Set 37	0.01053	0.01111	+0.00058	+0.055	+328
Set 38	0.00220	0.00243	+0.00023	+0.103	+144
Set 39	0.00559	0.00597	+0.00038	+0.068	+227

Table 8.44: Multiple muon fractions R_μ for each data set and its corresponding simulation. “Relative diff” is the difference between MC and data divided by the data ratio.

The effects of pile-up particles on the classification and reconstruction should therefore be similar in MC and data. The systematic uncertainty, then, will be dominated by differences between the beam rates used in the simulation and the actual rates in data.

To compare the muon rate seen in data to that in the simulation, a measure of the rate which could be applied to both was needed. A relative measure of the muon rate can be obtained using the ratio

$$R_\mu = \frac{(\text{Multiple Muon Events})}{(\text{Multiple Muon Events}) + (\text{Simple Clean Events})}. \quad (8.26)$$

A “Simple Clean” event (event type 1) consists of a muon and a decay positron well separated in time with no other particles. A “Multiple Muon” event (event type 15) consists of two or more muons and their decay positrons, in order, all well separated in time. The ratio R_μ in Eq. 8.26 is then an estimate of the fraction of events with more than one muon. The values of this ratio seen in each data set and its corresponding simulation are listed in Tab. 8.44. The systematic error in the muon rate for each simulation can be determined by multiplying the *relative* difference between data and MC by the actual muon rate measured in data using the muon trigger scintillator. This error is listed in the last column of Tab. 8.44.

The beam positron rates in data and simulation can be compared using the same method described for the muons, above, using “Time Clean” events (event type 2). As with R_μ , the ratio R_e is an estimate of the fraction of events with one or more beam positrons. The values of this ratio seen in each data set and its corresponding simulation are listed in Tab. 8.45. The systematic error in the beam positron rate for each simulation can be determined by multiplying the *relative* difference between data and MC by the actual beam positron rate measured in data using Mofia’s count of the number of beam positrons per event. This error is listed in the last column of Tab. 8.45.

The large difference between the beam positron rate in the set 38 simulation compared to the data is understood. Set 38 was taken with an aperture in place to limit the phase space of the beam. This aperture limits beam positrons as well as muons. The beam positron rate

	Data	MC	MC – Data	Relative Diff	Error (Hz)
Set 31	0.119	0.105	–0.014	–0.115	–1122
Set 32	0.117	0.103	–0.014	–0.116	–1106
Set 35	0.096	0.085	–0.011	–0.115	–894
Set 36	0.095	0.085	–0.010	–0.102	–778
Set 37	0.161	0.143	–0.019	–0.117	–1587
Set 38	0.127	0.071	–0.051	–0.399	–4224
Set 39	0.083	0.071	–0.012	–0.141	–892

Table 8.45: Beam positron fractions R_e for each data set and its corresponding simulation. “Relative diff” is the difference between MC and data divided by the data ratio.

is of course measured downstream of the aperture, but in the simulation the beam positrons are started *upstream* of the aperture. This was not accounted for when the beam positron rate was set in the simulation, and since the same beam positron profile was used for all simulations the aperture reduced the rate of beam positrons reaching the detector. The reason this did not affect the muon rate has not been studied at this time.

Measurement of Systematic Uncertainties

Two new simulations were produced to test the effect of changing the muon and beam positron rates. These were both identical to the standard simulation, except for the beam rates. The first used a muon rate of 27310 Hz, ten times larger than the rate of 2731 Hz used in the standard simulation. The second used a beam positron rate of 77690 Hz, ten times larger than the standard simulation’s rate of 7769 Hz.

To confirm that raising the beam rates did not have unexpected side effects, the event type distributions of the high-rate simulations were compared to those of the standard simulation; see Figs. 8.26 and 8.27. The changes in both cases were as expected: mainly, events were moved from the “clean” types to the types with extra particles (“multiple muon” or “time clean” types, respectively). A description of the event types is given in App. B.

The exaggerations for these simulations are based on the changes in beam rates: $(27310 - 2731) = 24417$ Hz for the high muon rate, and $(77690 - 7769) = 69921$ Hz for the high beam positron rate. The scale factors will then be the ratio of these rate exaggerations to the errors listed in tables 8.44 and 8.45. Conservative scale factors of 100 and 60 can be safely assumed for the muon and beam positron rate systematics for all sets, with the exception of the beam positron rate systematic for set 38; the scale factor there is only 16.

Systematic Effects

The high muon rate and high beam positron rate simulations were each fit against the complete standard simulation spectrum in the usual way, to measure the effects on the Michel parameters. Since these were new simulations, the correlations to the standard simulation were low; no renormalization of the uncertainties will be performed. The raw fit results are




Figure 8.26: Event type distributions for standard and high muon rate simulations. The high rate distribution was normalized to that of the standard. The top plot shows the number of events with each event type; the bottom plot shows the difference between standard and high muon rate.




Figure 8.27: Event type distributions for standard and high beam positron rate simulations. The high rate distribution was normalized to that of the standard. The top plot shows the number of events with each event type; the bottom plot shows the difference between standard and high beam positron rate.

	High Muon Rate	High Beam Positron Rate
$\Delta\rho$	$(5.51 \pm 3.64) \times 10^{-4}$	$(-0.049 \pm 3.33) \times 10^{-4}$
$\Delta\delta$	$(-9.99 \pm 6.40) \times 10^{-4}$	$(6.31 \pm 5.88) \times 10^{-4}$
$\Delta\xi$	$(22.85 \pm 7.91) \times 10^{-4}$	$(-4.68 \pm 7.25) \times 10^{-4}$
χ^2/dof	2349/2463	2339/2463
Renormalization	1	1
Scale factor	100	60

Table 8.46: Unscaled Michel parameter fits for the rate systematics, based on high muon and beam positron rates.

	Muon Intensity	Positron Intensity	Total
$\Delta\rho$	$(0.05 \pm 0.04) \times 10^{-4}$	$(-0.00 \pm 0.06) \times 10^{-4}$	$(0.05 \pm 0.07) \times 10^{-4}$
$\Delta\delta$	$(-0.10 \pm 0.06) \times 10^{-4}$	$(0.11 \pm 0.10) \times 10^{-4}$	$(0.21 \pm 0.11) \times 10^{-4}$
$\Delta\xi$	$(0.23 \pm 0.08) \times 10^{-4}$	$(-0.08 \pm 0.12) \times 10^{-4}$	$(0.31 \pm 0.14) \times 10^{-4}$

Table 8.47: Michel parameter fits for the rate systematics, based on high muon and beam positron rates, after applying scale factors. The “Total” column is the quadrature sum of the other two; this represents the systematic uncertainties due to uncertainties in the muon and positron beam intensities.

shown in Tab. 8.46. Table 8.47 shows the fit results after applying the scale factors. These rescaled results, combined in quadrature, represent the systematic uncertainties from the incorrect simulation of the beam intensities.

The scale factor for the positron beam intensity systematic is five times higher for set 38 than for the other sets. When added in quadrature to the muon beam intensity systematic, the total errors become $(0.05 \pm 0.24) \times 10^{-4}$ for ρ and $(0.52 \pm 0.40) \times 10^{-4}$ for δ . This does not make a significant impact on the final results.

8.7 External Uncertainties

The External Uncertainties category includes effects which require calculations or measurements from outside the *TWIST* experiment. The systematic uncertainties due to these effects are summarized in Tab. 8.48.

External Uncertainties	ρ	δ
Radiative Corrections (8.7.1)	0.29	0.11
Correlations with η (8.7.2)	1.12	0.10
Total	1.16	0.15

Table 8.48: Summary of systematic uncertainties external to the *TWIST* experiment, in units of 10^{-4} .

8.7.1 Radiative Corrections

The highest order radiative corrections used in the *TWIST* simulation are $\mathcal{O}(\alpha^2 L)$ (Sec. 6.1), where $L = \log(m_\mu^2/m_e^2)$. However, radiative corrections have been calculated to $\mathcal{O}(\alpha^2)$




Figure 8.28: Reconstructed radiative correction spectrum, generated using only the $\mathcal{O}(\alpha^2 L)$ radiative correction.

	Radiative Corrections
$\Delta\rho$	$(-1.59 \pm 3.09) \times 10^{-4}$
$\Delta\delta$	$(-0.62 \pm 5.47) \times 10^{-4}$
$\Delta\xi$	$(-2.86 \pm 6.74) \times 10^{-4}$
χ^2/dof	0.639/2463
Renormalization	0.016
Scale factor	5.55

Table 8.49: Unscaled Michel parameter fit between a simulated spectrum with exaggerated radiative corrections against standard simulation.

(Sec. 2.4); thus the $\mathcal{O}(\alpha^2)$ can be treated as the theoretical uncertainty.

Systematic Effects

A decay spectrum was generated by `micheld` using the $\mathcal{O}(\alpha^2 L)$ radiative correction as the entire probability distribution (similar to how the derivative spectra are generated). This “RC-only” spectrum was then passed to GEANT and simulated and analyzed normally; the reconstructed spectrum is shown in Fig. 8.28. The resulting reconstructed spectrum, scaled down by a factor of 200, was added to a standard simulated spectrum. The pure RC spectrum had to be scaled down before combining it with the standard spectrum so that the combination would not be unfittably distorted. Energy calibration was applied between RC+standard and the standard spectrum, and a spectrum fit performed between the calibrated RC+standard and standard. The fit results are listed in Tab. 8.49.

The scale factor can be determined from two parts: the amount of RC-only spectrum

	Radiative Corrections
$\Delta\rho$	$(-0.29 \pm 0.07) \times 10^{-4}$
$\Delta\delta$	$(-0.11 \pm 0.12) \times 10^{-4}$
$\Delta\xi$	$(-0.51 \pm 0.15) \times 10^{-4}$

Table 8.50: Michel parameter fits of a decay spectrum with exaggerated radiative corrections (“RC+base”; see text) against a standard simulation, after applying the scale factor. This represents the systematic uncertainties due to the radiative corrections.

which has been added to the standard to produce the RC+base spectrum, and the comparison between the $\mathcal{O}(\alpha^2 L)$ and $\mathcal{O}(\alpha^2)$ corrections.

To determine how much RC-only spectrum was added, one can compare the `nthrown` values reported by the `micheld` spectrum generator, correcting the RC-only `nthrown` value by the factor of 200 which was applied before combining spectra. The value of `nthrown` for the RC-only spectrum was 2.2×10^{11} , and for the standard spectrum it was 9.9×10^8 , for a ratio of

$$\frac{2.2 \times 10^{11}/200}{9.9 \times 10^8} = 1.11.$$

The $\mathcal{O}(\alpha^2)$ radiative corrections are on the order of 20% of the $\mathcal{O}(\alpha^2 L)$ corrections [21], providing an additional scale factor of 5 to this study. The final scale factor is then 5.55.

Table 8.50 gives the final systematic uncertainty due to the radiative corrections of the muon decay spectrum.

8.7.2 Correlations with η

Recall that the *TWIST* experiment is not sensitive to the Michel parameter η , as explained in Sec. 5.1. However, the value of η assumed in the Michel parameter fit affects the fit results to some degree, particularly the value of ρ . To measure the correlations, a standard data-MC fit was repeated with three different values of η : the world average value of $\eta = -0.0036$, from [14], and one standard error ($\delta\eta = \pm 0.0069$) to either side ($\eta = -0.0105$ and $\eta = +0.0033$). The changes in the fit parameters $\Delta\rho$ etc. give the correlations

$$\frac{d\rho}{d\eta} = 0.0162 \qquad \frac{d\delta}{d\eta} = 0.0015 \qquad \frac{d\xi}{d\eta} = 0.0155 \qquad (8.27)$$

which can be used to determine the systematic uncertainty in each parameter due to the uncertainty in η . In this way, external measurements which improve the knowledge of η can be incorporated directly into the systematic uncertainty. The present uncertainty in η results in the systematic uncertainties in the other Michel parameters listed in Tab. 8.51. Note that, as the uncertainty in η is reduced by future experiments, the correlations in Eq. (8.27) can be used to reduce this systematic uncertainty; for this reason, the systematic uncertainty due to the uncertainty in η is quoted separately from the other uncertainties.

	Radiative Corrections
$\Delta\rho$	1.12×10^{-4}
$\Delta\delta$	0.10×10^{-4}
$\Delta\xi$	1.07×10^{-4}

Table 8.51: Systematic uncertainties due to correlations between η and the other Michel parameters, based on an uncertainty in η of 0.0069 and fit correlations given in Eq. (8.27).

8.8 Negligible Systematic Uncertainties

Several systematic uncertainties reported in previous rounds [48–50] are not included in the above lists. Some of these are briefly discussed below.

The stopping target used for the first ρ and δ measurements was Mylar, sprayed with graphite as a conduction layer. The thickness and uniformity of the graphite layer was largely unknown. The thickness and other properties of the current stopping target, a high-purity aluminum foil, are known to high precision.

The data-MC mismatch in the mean stopping position of muons within the target is accounted for by the energy calibration. Differences in the shape of the stopping distribution could smear the momentum (degrading the apparent momentum resolution), asymmetrically between upstream and downstream. A positron passing entirely through the target will lose on the order of 40 keV of energy, so a positron leaving from a decay within the target will lose some fraction of this. Simulation studies show that the width of the stopping distribution within the target is equivalent to a roughly 10 keV variation in energy loss. This smearing in the positron energy is added in quadrature to the width of the Landau distribution of energy loss the positrons experience as they travel to the first DC where they are measured, and the reconstruction resolution of the helix fitter, a combined width of around 100 keV or more. Thus any effect from a mismatch of stopping distribution widths between data and simulation is completely negligible.

The ionization left by the passage of a muon through a tracking chamber results in a “dead zone” in which no further hits can be measured for some time (see Sec. 6.2.1). The probability of a positron passing through this dead zone depends on the energy and angle of the positron; moreover, since the muon only passes through the upstream half of the detector this can introduce an upstream-downstream asymmetry. This is primarily a problem in the target PCs, where a positron which decays upstream is almost guaranteed to pass through the dead zone produced by the stopping muon. For this reason, the current analysis does not use the target PCs for pattern recognition. With this change, the systematic uncertainty due to errors in the simulation of the dead zone became negligible.

In addition to the above, systematic uncertainties due to hard scattering, chamber hit efficiencies, the stability of the chamber high voltage, and chamber crosstalk were found in previous rounds of analysis to be negligible, and are therefore not included here.

Chapter 9

Results

If you haven't found something strange during the day, it hasn't been much of a day.
—J.A. Wheeler

9.1 Results Revealed

Once all systematic uncertainties had been evaluated,¹ the private encryption key was applied to the database and the hidden values of the Michel parameters extracted. The fits between each data set and its corresponding simulation were then repeated (c.f. Sec. 7.2.2), with the value of η fixed in the fit to the world average value of $\eta = 0.0036$ from ref. [14], and the true black box values of δ and ξ were used for the conversion of the parameters to standard (ξ, δ) notation.

The parameter fits of each data set to its corresponding simulation are listed in Tab. 9.1, along with the χ^2 value of each fit. For each parameter a weighted mean measurement was obtained. Figs. 9.1 and 9.2 show the fit results graphically. The consistency of the seven independent measurements, taken under varying conditions, can be seen visually and

¹The exception is the fit correlation between η and the other parameters (Sec. 8.7.2), which was evaluated once the hidden value of η was known; the procedure was decided upon before the hidden values were revealed.

Set	$\Delta\rho (\times 10^{-4})$	$\Delta\delta (\times 10^{-4})$	$\Delta(P_\mu\xi) (\times 10^{-4})$	χ^2
31	-75.1 ± 4.4	$+91.1 \pm 7.7$	$+95.2 \pm 9.6$	2505
32	-69.3 ± 7.9	$+60.2 \pm 14.0$	$+66.3 \pm 17.3$	2434
35	-82.8 ± 3.8	$+92.5 \pm 6.7$	$+75.2 \pm 8.2$	2458
36	-78.1 ± 3.7	$+91.7 \pm 6.6$	$+73.6 \pm 8.1$	2483
37	-80.2 ± 5.0	$+89.1 \pm 8.8$	$+66.5 \pm 10.9$	2392
38	-78.6 ± 4.5	$+96.0 \pm 8.0$	$+72.3 \pm 9.9$	2555
39	-76.3 ± 4.9	$+72.7 \pm 8.6$	$+90.2 \pm 10.7$	2468
Average	-78.3 ± 1.7	$+88.3 \pm 3.0$	$+77.9 \pm 3.7$	

Table 9.1: Measured differences in Michel parameters between data and simulation, for each data set. All fits were performed with η fixed to the world average value ($\eta = 0.0036$) from [14]. Uncertainties are statistical. χ^2 values are shown for each parameter fit, for 2463 degrees of freedom. Data sets are described in detail in Chap. 7.





Figure 9.1: Measured differences in ρ and δ between data and simulation, for each data set. All fits were performed with η fixed to the world average value ($\eta = 0.0036$) from ref. [14]. A constant is fit for each parameter to determine the average measurement. Error bars are statistical. Data sets are described in detail in Chap. 7.




Figure 9.2: Measured differences in $P_\mu\xi$ between data and simulation, for each data set. All fits were performed with η fixed to the world average value ($\eta = 0.0036$) from [14]. A constant is fit to determine the average measurement. Error bars are statistical. Data sets are described in detail in Chap. 7.

from the χ^2/dof of the constant fits. In Tab. 9.2, the “black box” values of the Michel parameters used in the simulation are added to the average Data-MC differences, along with the corrections determined in Chap. 8. For the current analysis, we find

$$\rho = 0.75014 \pm 0.00017(\text{stat}) \pm 0.00044(\text{syst}) \pm 0.00011(\eta)$$

and

$$\delta = 0.75067 \pm 0.00030(\text{stat}) \pm 0.00067(\text{syst}).$$

The third uncertainty listed for ρ is due to the uncertainty in η ($\eta = -0.0036 \pm 0.0069$ [14]) and its correlation with ρ ($\partial\rho/\partial\eta = 0.016$), as described in Sec. 8.7.2.

	Data – MC	Black Box	Corrections	Final Measurement
ρ	-0.00783	0.75805	-0.00008	$0.75014 \pm 0.00017 \pm 0.00044$ ± 0.00011
δ	+0.00883	0.74155	+0.00029	$0.75067 \pm 0.00030 \pm 0.00067$
$P_\mu\xi$	+0.00779	0.99866	-0.00392	1.00253 ± 0.00037

Table 9.2: Measured values of the Michel parameters, with black box values and corrections applied. The first uncertainty in the total is statistical, the second systematic. The third uncertainty in ρ is due to the uncertainty in the parameter η . Note that no systematic uncertainty is quoted for $P_\mu\xi$.

The uncertainty on the new measurement of ρ is a factor of 2.1 improvement over the previous *TWIST* measurement [24], and a factor of 5.2 improvement over the pre-*TWIST* world average [14]. The uncertainty on the new measurement of δ is a factor of 1.8 improvement over the previous *TWIST* measurement [25], and a factor of 5.2 improvement over the pre-*TWIST* world average [14]. These improvements meet the goals of the *TWIST* collaboration for this round of analysis.

The systematic uncertainties in $P_\mu\xi$ were not studied explicitly, since $P_\mu\xi$ has already been measured from this data [27]. In particular, major $P_\mu\xi$ uncertainties related to the muon beam profile, muon depolarization in the stopping target, and other effects were unaddressed under the new analysis, although they should be largely unchanged from the previous study and continue to dominate the systematic uncertainty. Therefore the value of $P_\mu\xi$ obtained in this analysis should not be considered a new measurement.

Corrections corresponding to those listed in Tab. 8.2 were applied to the value of $P_\mu\xi$ obtained in this study. In particular, the correction for the choice of drift time maps used (see Sec. 8.1.1) was applied here (a correction of $\Delta P_\mu\xi = (4.0 \pm 6.4) \times 10^{-4}$), and this correction was not made for the previous $P_\mu\xi$ analysis. Taking this correction and other changes in the analysis procedure into account, the value of $P_\mu\xi$ obtained in the present analysis is consistent with the previous *TWIST* measurement [27].

9.1.1 “White Box” Consistency Check

The “white box” consistency check is a test of the *TWIST* blind analysis procedure, where a new simulation generated using the measured Michel parameters should give results consistent with zero when fit against the same data.

The nominal data set 35 was used for this test. Given the “black box” values listed in Tab. 9.2, the Michel parameters measured from set 35 were $\rho = 0.7498 \pm 0.0004$, $\delta = 0.7508 \pm 0.0007$, and $P_\mu\xi = 1.0062 \pm 0.0008$. Unfortunately this combination of Michel parameters cannot be used to generate a decay spectrum, as $P_\mu\xi\delta/\rho = 1.0076 > 1$, which is disallowed because it results in portions of the decay spectrum with negative probability. (This constraint is not imposed during the spectrum fit (Sec. 5.1), so it is free to find the best fit values for all three parameters separately.) Since ρ and δ are of particular interest for this analysis, the value $\xi = (0.9999)\rho/\delta = 0.998518$ was used for the white box spectrum. With these values of the Michel parameters, the white box test should give results consistent with $\Delta\rho = \Delta\delta = 0$, and $\Delta P_\mu\xi = +0.007659$. The white box spectrum was simulated and analyzed in the usual way, and the spectrum fit results were

$$\Delta\rho = -0.00043 \pm 0.00039 \quad \Delta\delta = -0.00055 \pm 0.00068 \quad \Delta P_\mu\xi = +0.00778 \pm 0.00085$$

which is consistent with the prediction, validating the *TWIST* blind analysis procedure.

	2005 Global Analysis [23] ($\times 10^{-3}$)	Present Analysis ($\times 10^{-3}$)
Q_{RR}	< 1.14	< 0.998
Q_{LR}	< 1.94	< 1.45
Q_{RL}	< 44	< 43
Q_{LL}	> 955	> 955
B_{LR}	< 1.27	< 0.722
B_{RL}	< 10.9	< 10.9
α/A	0.3 ± 2.1	0.1 ± 1.6
β/A	2.0 ± 3.1	2.1 ± 3.0
α'/A	-0.1 ± 2.2	-0.1 ± 1.6
β'/A	-0.8 ± 3.2	-0.8 ± 3.1
ρ	0.74959 ± 0.00063	0.74964 ± 0.00035
δ	0.74870 ± 0.00114	0.74997 ± 0.00065
ξ	0.99881 ± 0.00212	0.99726 ± 0.00132
$\xi\delta/\rho$	0.99764 ± 0.00085	0.99771 ± 0.00083
η	-0.0034 ± 0.0069	-0.0042 ± 0.0064

Table 9.3: Results of a new global analysis of muon decay data, including the present measurements (90% C.L.). $P_\mu = 1$ is assumed. Best fit values of selected Michel parameters are also listed.

9.2 Results in Context

9.2.1 Global Analysis of Muon Decay

The global analysis of all available muon decay data performed by Gagliardi *et al.* [23], described in Sec. 3.5, has been repeated by the author using their code. This reanalysis includes the new *TWIST* measurements of ρ and δ , as well as the *TWIST* measurement of $P_\mu\xi$ [27], which was not yet available at the time of the 2005 analysis. All other input values used in [23] were used in the present global analysis as well.

Recall from Sec. 3.5 that the global analysis used a Monte Carlo method to map out the joint probability distributions for 9 independent variables, Q_{RR} , Q_{LR} , Q_{RL} , B_{LR} , B_{RL} , α/A , β/A , α'/A , and β'/A , each a bilinear combination of the weak coupling constants g_{ij}^κ . The Michel parameters can then be written in terms of these independent variables. Table 9.3 shows the results of this global analysis, as well as the results of the 2005 analysis. The 90% confidence limits are given for the independent variables listed above, and global best-fit values of the Michel parameters ρ , δ , ξ , and η are given. The present analysis represents significant improvements in the limits on Q_{LR} and B_{LR} , and tightens several of the other limits. It is interesting to note that the global analysis significantly reduces the uncertainty in the value of ρ , from a total of 0.00050 to 0.00035.

This global analysis also provides a new indirect limit on the value of $P_\mu\xi$, due to the new values of ρ , and δ , in combination with the measurement of $P_\mu\xi\delta/\rho$ by Jodidio [28,29]: $P_\mu\xi = 0.99726 \pm 0.00132$, an improvement by a factor of 1.6.

The values of the Q_{ij} from this global analysis can be used in Eqs. (2.8) on p. 8 to

	Standard Model	Previous limits	2005 Global analysis [23]	Present Analysis
$ g_{RR}^S $	0	< 0.066	< 0.067	< 0.063
$ g_{RR}^V $	0	< 0.033	< 0.034	< 0.032
$ g_{LR}^S $	0	< 0.125	< 0.088	< 0.076
$ g_{LR}^V $	0	< 0.060	< 0.036	< 0.027
$ g_{LR}^T $	0	< 0.036	< 0.025	< 0.022
$ g_{RL}^S $	0	< 0.424	< 0.417	< 0.415
$ g_{RL}^V $	0	< 0.110	< 0.104	< 0.105
$ g_{RL}^T $	0	< 0.122	< 0.104	< 0.104
$ g_{LL}^S $	0	< 0.550	< 0.550	< 0.550
$ g_{LL}^V $	1	> 0.960	> 0.960	< 0.960

Table 9.4: New limits on the weak coupling constants from this analysis. (Limits on $|g_{LL}^S|$ and $|g_{LL}^V|$ are from [14].) The present analysis represents an improvement of approximately 10% in the limits on $|g_{LR}^S|$, $|g_{LR}^V|$, and $|g_{LR}^T|$.

place limits on the magnitudes of the weak coupling constants $|g_{ij}^\kappa|$. Since each Q_{ij} is a combination of terms of the form $|g_{ij}^\kappa|^2$, which must be positive, a limit on each $|g_{ij}^\kappa|$ can be determined by assuming that each coupling in turn is the only one to differ from the Standard Model. (The exceptions are $|g_{LL}^V|$ and $|g_{LL}^S|$, which are determined more sensitively from inverse muon decay, $e^- \nu_\mu \rightarrow \mu^- \nu_e$.) The limits determined with this method are listed in Tab. 9.4, along with the values from the 2005 global analysis and the previous limits.

9.2.2 Implications for Physics Beyond the Standard Model

Right-Handed Muon Decay

As discussed in Sec. 2.5.1, the quantity $Q_R^\mu = Q_{RR} + Q_{LR}$ represents the total probability for a right-handed muon to decay into any type of electron, a process forbidden under the Standard Model weak interaction. The new measurements of ρ and δ lead to the new limits on Q_{RR} and Q_{LR} shown in Tab. 9.3, and hence to a new 90% confidence limit upper bound on the combined probability $Q_R^\mu < 0.0024$, a significant improvement over the limit of $Q_R^\mu < 0.007$ from the 2005 global analysis [23].

Left-Right Symmetric Models

The left-right symmetric models, discussed in Sec. 2.5.2, are a class of extensions to the Standard Model where right-handed particles also interact weakly, through a “(V + A)” interaction which is suppressed but non-zero [26]. The right-handed interaction in this model is mediated by a second charged boson, the W_R , which mixes with the usual W_L :

$$\begin{aligned}
W_L &= W_1 \cos \zeta + W_2 \sin \zeta \\
W_R &= e^{i\omega} (-W_1 \sin \zeta + W_2 \cos \zeta).
\end{aligned}$$

See Sec. 2.5.2 for details. In general, the right-handed and left-handed interactions are governed by separate coupling constants, g_R and g_L . Recall from Eq. (2.21) that the right-handed current modifies the Michel parameter ρ :

$$\rho \simeq \frac{3}{4}(1 - 2\zeta_g^2)$$

where $\zeta_g \equiv \frac{g_R}{g_L}\zeta$. Under the left-right symmetric models, $\rho > 0.75$ is forbidden, so the general measurement of ρ can be converted into a 90% confidence limit lower bound: $\rho > 0.7493$ (compared with $\rho > 0.7487$ from the ρ measurement previously published by *TWIST*). This gives a limit on the left-right mixing angle of $|\zeta_g| < 0.022$, a significant improvement over the limit of $|\zeta_g| < 0.030$ for the published *TWIST* value of ρ .

9.2.3 Future Work

Many improvements to the simulation and analysis are already underway by the *TWIST* collaboration, and a large amount of new data have been taken since the data sets discussed in this work. Particular improvements suggested by the systematic uncertainties studied in Chap. 8 are mentioned below.

The use of STRs derived directly from the data and the simulation (Sec. 8.1.1), rather than the use of STRs simulated by GARFIELD, will greatly reduce the systematic uncertainties due to chamber response, and will eliminate the need for one of the corrections applied here.

The systematic uncertainty due to the uncertainty in the energy calibration (Sec. 8.2.2) will be reduced simply by the planned increase in statistics available for the final *TWIST* measurements. The algorithm used to perform the energy calibration can be improved to take into account differences in momentum resolution between the two spectra being calibrated, thus reducing the sensitivity to resolution (Sec. 8.4.1).

The systematic uncertainty due to errors in the production of delta rays (Sec. 8.3.1) was studied by adjusting the production rate in the simulation, and comparing the results to the production rate seen in data. This same method could be used to tune the production rate in the simulation. This would require a more detailed study of the delta ray production rate, however, to confirm that the mismatch with data is indeed due to delta rays.

9.3 Conclusion

New high-precision measurements of the Michel parameters ρ and δ have been made. These results are consistent with the Standard Model, and are approximately twice as precise as the previous *TWIST* measurements, and more than five times as precise as the previously accepted values. These measurements put significant new constraints on possible extensions to the Standard Model weak interaction. In particular, the limit on the probability that

a right-handed muon will decay via the weak interaction has been reduced by more than a factor of two. Limitations of these measurements have been identified, and work is in progress to greatly reduce them for future measurements. The prospects for the *TWIST* experiment to achieve its ultimate goals, to improve the knowledge of the Michel parameters by an order of magnitude, are excellent.

The Standard Model of the weak interaction is safe—for now...

Bibliography

- [1] T.D. Lee and C.N. Yang. Question of parity conservation in weak interactions. *Physical Review*, 104(1):254–258, Oct 1956.
- [2] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Physical Review*, 105(4):1413–1415, Feb 1957.
- [3] Richard L. Garwin, Leon M. Lederman, and Marcel Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Physical Review*, 105(4):1415–1417, Feb 1957.
- [4] J.C. Street and E.C. Stevenson. New evidence for the existence of a particle of mass intermediate between the proton and electron. *Physical Review*, 52(9):1003–1004, Nov 1937.
- [5] E.P. Hincks and B. Pontecorvo. The absorption of charged particles from the 2.2- μ sec. meson decay. *Physical Review*, 74(6):697–698, Sep 1948.
- [6] J. Steinberger. On the range of the electrons in meson decay. *Physical Review*, 75(8):1136–1143, Apr 1949.
- [7] L. Michel. Interaction between four half-spin particles and the decay of the μ -meson. *Proceedings of the Physical Society, A*, 63:514, 1371, 1950.
- [8] T. Kinoshita and A. Sirlin. Polarization of electrons in muon decay with general parity-nonconserving interactions. *Physical Review*, 108(3):844–850, November 1957.
- [9] D.J. Griffiths. *Introduction to Elementary Particles*. John Wiley & Sons, Inc., 1987.
- [10] J.H. Christenson *et al.* Evidence for the 2π decay of the K_2^0 meson. *Physical Review Letters*, 13(4):138–140, July 1964.
- [11] J.W. Cronin. CP symmetry violation—the search for its origin. *Review of Modern Physics*, 53(3):373–383, July 1981.
- [12] D.N. Spergel *et al.* Three-year wilkinson microwave anisotropy probe (WMAP) observations: Implications for cosmology. *The Astrophysical Journal Supplement Series*, 170(2):377–408, 2007.
- [13] P. Depommier. Muon polarization in the $\pi^+ \rightarrow \mu^+ \nu \gamma$ decay. Technical Note TN-82, TRIUMF TWIST experiment, 25 June, 2003. Supplmented in TN-89 and TN-96.
- [14] W.-M. Yao *et al.* Review of Particle Physics. *Journal of Physics G*, 33, 2006.
- [15] H.-J. Gerber W. Fetscher and K.F. Johnson. Muon decay: complete determination of the interaction and comparison with the standard model. *Physics Letters B*, 173(1):102–106, May 1986.
- [16] K. Schilcher A.I. Davydychev and H. Spiesberger. Hadronic corrections at $\mathcal{O}(\alpha^2)$ to the energy spectrum of muon decay. *European Journal of Physics*, C19:99–104, 2001.
- [17] W. Fetscher and H.-J. Gerber. Muon decay parameters. In *Review of Particle Physics* [14].
- [18] A. Arbuzov. First-order radiative corrections to polarized muon decay spectrum. *Physics Letters B*, 524:99, January 2002.

- [19] A. Czarnecki A. Arbuzov and A. Gaponenko. Muon decay spectrum: Leading logarithmic approximation. *Physical Review D*, 65:113006, 2002.
- [20] A. Arbuzov and K. Melnikov. $\mathcal{O}(\alpha^2 \ln(m_\mu/m_e))$ corrections to electron energy spectrum in muon decay. *Physical Review D*, 66:093003, 2002.
- [21] K. Melnikov C. Anastasiou and F. Petriello. The electron energy spcctrum in muon decay through $\mathcal{O}(\alpha^2)$. *Journal of High Energy Physics*, (09):014, 2007.
- [22] W. Fetscher and H.-J. Gerber. Precision measurements in muon and tau decays. In P. Langacker, editor, *Precision Tests of the Standard Electroweak Model*. World Scientific, Singapore, 1995.
- [23] R.E. Tribble C.A. Gagliardi and N.J. Williams. Global analysis of muon decay measurements. *Physical Review D*, 72(7):073002, 2005.
- [24] J.R. Musser et al. Measurement of the michel parameter ρ in muon decay. *Physical Review Letters*, 94(10):101805, 2005.
- [25] A. Gaponenko et al. Measurement of the muon decay parameter δ . *Physical Review D*, 71(7):071101, 2005.
- [26] B. Herczeg. On muon decay in left-right-symmetric electroweak models. *Physical Review*, D34(11):3449–3456, Dec 1986.
- [27] B. Jamieson et al. Measurement of $P_\mu \xi$ in polarized muon decay. *Physical Review D*, 74(7):072007, 2006.
- [28] A. Jodidio et al. Search for right-handed currents in muon decay. *Physical Review*, D34:1967, 1986.
- [29] A. Jodidio et al. Erratum: Search for right-handed currents in muon decay. *Physical Review*, D37:237–238, January 1988. Erratum to [28].
- [30] J. Peoples. *Positron Spectrum from Muon Decay*. PhD thesis, Columbia University, 1966. Report No. NEVIS-147 (unpublished).
- [31] M. Bardon, P. Norton, J. Peoples, A.M. Sachs, and J. Lee-Franzini. Measurement of the momentum spectrum of positrons from muon decay. *Physical Review Letters*, 14(12):449–453, Mar 1965.
- [32] S.E. Derenzo. Measurement of the low-energy end of the mu-plus decay spectrum. *Physical Review*, 181:1854–1866, 1969.
- [33] B. Balke et al. Precise measurement of the asymmetry parameter delta in muon decay. *Physical Review*, D37:587–617, 1988.
- [34] H. Burkard et al. Muon decay: Measurement of the transverse positron polarization and general analysis. *Physics Letters B*, 160:242–246, October 1985.
- [35] N. Danneberg et al. Muon decay: Measurement of the transverse polarization of the decay positrons and its implications for the fermi coupling constant and time reversal invariance. *Physical Review Letters*, 94:021802, 2005.
- [36] G.M. Marshall C.J. Oram, J.B. Warren and J. Doornbos. Commissioning of a new low energy pi-mu channel at TRIUMF. *Nuclear Instruments and Methods*, 179:99–103, 1981.
- [37] R.P. MacDonald. Summary of m13 momentum scans of surface muon edge. Technical Note TN-52, TRIUMF TWIST Experiment, 8 March 2001.
- [38] J. Hu et al. Time expansion chamber system for characterization of twist low-energy muon beams. *Nuclear Instruments and Methods A*, 566:563–574, 2006.
- [39] R.P. MacDonald. Examination and removal of backgrounds for twist at triumph. Master’s thesis, University of Alberta, 2001.
- [40] Vector Fields Ltd. *OPERA-3d*, 2002.

- [41] R. Armenta. Twist magnet: Field map study. Technical Note TN-68, TRIUMF TWIST experiment, August 2002.
- [42] R.S. Henderson et al. Precision planar drift chambers and cradle for the twist muon decay spectrometer. *Nuclear Instruments and Methods A*, 548:306–335, 2005.
- [43] Tcl programming language and tk graphical user interface toolkit. <http://www.tcl.tk/>.
- [44] S. Ritt et al. Midas data acquisition system. <http://midas.psi.ch/>, 2004.
- [45] R. Brun and F. Rademakers. Root—an object oriented data analysis framework. *Nuclear Instruments and Methods A*, 389:81–86, April 1997. See also <http://root.cern.ch/>.
- [46] CN/ASD Group. MINUIT—*Users Guide, nProgram Library D506*. CERN, 1993.
- [47] F. James and M. Winkler. MINUIT. <http://www.cern.ch/minuit>, November 2005.
- [48] A. Gaponenko. *A precision measurement of the muon decay parameter delta*. PhD thesis, University of Alberta, 2005.
- [49] J.R. Musser. *Measurement of the Michel parameter ρ in muon decay*. PhD thesis, Texas A&M University, December 2005.
- [50] B. Jamieson. *Measurement of the Muon Decay Asymmetry Parameter with the TWIST Spectrometer*. PhD thesis, University of British Columbia, November 2005.
- [51] R. Veenhof. *GARFIELD*. CERN. Version 7.10.
- [52] Westgrid: Western canada research grid. <http://westgrid.ca>, 2007.
- [53] R.P. MacDonald. Tests of the energy loss correction in the helix fitter. Technical Note TN-97, University of Alberta, May 2005.
- [54] G. Lutz. Optimum track fitting in the presence of multiple scattering. *Nuclear Instruments and Methods A*, 273:349–361, May 1988.
- [55] F. James. Fitting tracks in wire chambers using the chebyshev norm instead of least squares. *Nuclear Instruments and Methods*, 211(1):145–152, June 1983.
- [56] H. Bichsel and D.E. Groom. Passage of particles through matter. In *Review of Particle Physics* [14].
- [57] F. Carminati R. Brun and S. Giani. *GEANT Detector Description and Simulation Tool*. CERN Application Software Group, Geneva, Switzerland, March 1994. Version 3.21.
- [58] V. Bargmann, L. Michel, and V.L. Telegdi. Precession of the polarization of particles moving in a homogeneous electromagnetic field. *Physical Review Letters*, 2(10):435–436, May 1959.
- [59] B. Barrett. Analysis of space charge and dead zone effects in the twist spectrometer. Technical Note TN-95, St. Mary’s University, April 2005.

Appendix A

Window Type Definitions

The window types identified by the Mofia classification, as described in `mainf90/evalwin_mod.f90`, are listed below. See section 5.3.1 for information. The distributions of window types found in simulation and data are shown in figure A.1.

- 1 Muon
- 2 Upstream Decay
- 3 Downstream Decay
- 4 Beam Positron
- 5 Empty
- 6 PC Overlap involved
- 7 Upstream Decay (DS delta identified)
- 8 Downstream Decay (US delta identified)
- 9 Upstream Decay (Scatter back downstream)
- 10 Downstream Decay (Scatter back upstream)
- 11 Muon and Fast Decay Downstream
- 12 Muon and Fast Decay Upstream
- 13 Muon and Beam Positron
- 14 Upstream Decay with Beam Positron
- 15 Downstream Decay with Beam Positron
- 16 Soft, uncorrelated track
- 17 Cosmic ray, noise, or beam gas
- 18 Too early
- 19 Too late
- 20 Don't know
- 21 Pion




Figure A.1: Window type distributions from data (Set35) and simulation (Gen335). The upper plot shows all window types, and the lower plot is a close-up to show the distribution of the less common window classifications.

Appendix B

Event Type Definitions

The event types identified by the Mofia classification, as described in `mainf90/classifybj_mod.f90`, are listed below. See section 5.3.1 for information. Standard production spectrum summation, for either simulation or data, uses the following event types: 1, 2, 6, 7, 10, 11, 21, and 22. The distributions of event types found in simulation and data are shown in figure B.1.

- 1 (**Simple Clean**) Muon and Decay, DC and PC hits well separated in time
- 2 (**Time Clean**) Muon, Decay, and Beam Positrons, DC and PC hits well separated in time
- 3 (**Simple DC Overlap**) Muon and Decay, PC times separated in time, DC times overlap
- 4 (**Time DC Overlap**) Muon, Decay, and Beam Positron PC times separated in time, DC times overlap
- 5 (**PC Overlap**) Overlap of PC hits
- 6 (**Simple Delta Cleaned**) Muon and Decay, DC and PC hits well separated in time, delta removed
- 7 (**Time Delta Cleaned**) Muon, Decay, and Beam Positrons, DC and PC hits well separated in time, delta removed
- 8 (**Simple DC Overlap Delta Cleaned**) Muon and Decay, PC times separated in time, DC times overlap, delta removed
- 9 (**Time DC Overlap Delta Cleaned**) Muon, Decay, and Beam Positrons, PC times separated in time, DC times overlap, delta removed
- 10 (**Simple Scatter Cleaned**) Muon and Decay, well separated in time, scatter removed
- 11 (**Time Scatter Cleaned**) Muon, Decay, and Beam Positrons, well separated in time, scatter removed
- 12 (**Simple DC Overlap Scatter Cleaned**) Muon and Decay, PC times separated in time, DC times overlap, scatter removed
- 13 (**Time DC Overlap Scatter Cleaned**) Muon, Decay, and Beam Positrons, DC times overlap, PC times separated in time, scatter removed
- 14 (**Beam Positron Trigger**) Event triggered by Beam Positron
- 15 (**Multiple Muon Decays Simple Clean**) Two or more Muon Decays in order, well separated in time
- 16 (**Multiple Muon Decays Time Clean**) Two or more Muon Decays in order, with Beam Positrons, well separated in time
- 17 (**Multiple Muon Decays DC Overlap Simple Clean**) Two or more Muon Decays in order, PC times separated in time, DC times overlap





Figure B.1: Event type distributions from data (Set35) and simulation (Gen335). The upper plot shows all event types, and the lower plot is a close-up to show the distribution of the less common event classifications.

- 18 (Multiple Muon Decays DC Overlap Time Clean)** Two or more Muon Decays in order, with Beam Positrons, PC times separated in time, DC times overlap
- 19 (Multiple Muon Decays Dirty)** Two or more Muon Decays in wrong order
- 20 (Just a Muon and Beam Positrons)** No Decay Positron identified
- 21 (Simple Beam Positron Cleaned)** Muon and Decay, well separated in time, but with Beam Positron PC overlap in any
- 22 (Time Beam Positron Cleaned)** Muon, Decay, and Beam Positrons well separated in time, but with Beam Positron PC overlap in any
- 23 (Simple DC Overlap Beam Positron Cleaned)** Muon and Decay, PC times separated in time, but with Beam Positron PC overlap in either
- 24 (Time DC Overlap Beam Positron Cleaned)** Muon, Decay, and Beam Positrons, PC times separated in time, but with Beam Positron PC overlap in any
- 25 (Just Beam Positrons, Unknown Trigger)**
- 26 (Other, Unknown Trigger)**
- 27 (Other)** Whatever is left
- 28 (Simple Clean, Too Few Planes Hit)** for types 1, 6, 10, 15
- 29 (Time Clean, Too Few Planes Hit)** for types 2, 7, 11, 16
- 30 (Simple Clean, Too High Angle)** for types 1, 6, 10, 15
- 31 (Time Clean, Too High Angle)** for types 2, 7, 11, 16

Index

- t_0
 - systematic, 107
- GEANT
 - calibration files, 66
 - data banks, 66
 - dead zone, 64
 - digitization, 64
 - energy loss, 63
 - ion clusterization, 64
 - multiple scattering, 63
 - tracking, 63
- GEANT, 63
- alignment, 33
 - systematic, 139
- beam
 - in GEANT, 64
 - intensity
 - systematic, 144
 - pile-up, 64
- beamline
 - BL1A, 20
 - M13, 22
 - calibration, 23
 - momentum, 23
 - momentum bite, 23
- blind analysis, 61
- bremsstrahlung, 69
 - production rate, 121
 - systematic, 119
- calibration files
 - in GEANT, 66
 - in Mofia, 42
- CFM, 42
- chambers
 - response
 - systematic, 97
- classification
 - event, 44
 - time windows, 43
- coordinate systems, 32
- cptof, 25
- crosstalk
 - removal, 42
- data, 89
- data files, 41
- data sets, 89
- decay spectrum, 9
 - derivatives, 37
- delta ray
 - production rate, 119
 - systematic, 119
- drift chambers
 - gas density
 - systematic, 102
- drift chambers, 29
 - dense stack, 30
 - drift times, 41
 - numbering, 30
 - sparse stack, 30
- energy calibration
 - application, 55, 73
 - momentum dependence, 117
 - systematic, 115
- energy loss, 69
- η
 - correlation systematic, 151
- η
 - global analysis measurement
 - previous, 19
- event type
 - List of event types, 166
- foil bulge
 - systematic, 104
- foil bulge, 92
- gas degrader
 - simulation, 63
- gas absorber
 - simulation, 90
- gas degrader, 23
- hard scattering, 69
- helix fitter, 45
- kinematics
 - muon decay, 7
 - pion decay, 6
- m13
 - settings, 90
- magnetic field, 27
 - systematic, 114
 - tracking, 63
- mcfitter, 37
- Michel parameters, 9
- micheld, 61
 - in GEANT, 65
- Mofia, 41
 - calibration files, 42
 - multiple scattering, 69
- muon
 - beam

- intensity, **144**
 - cloud, 22
 - decay parameters, 14
 - lifetime, 7
 - polarization, 6
 - production, 22
 - surface, 22
- pattern recognition, 44
- pion
 - lifetime, 6
- positron
 - beam
 - intensity, **144**
- proportional chambers
 - numbering, 30
- proportional chambers, **29**
- radiative corrections
 - in GEANT, 62
 - systematic, **149**
- resolution
 - Data-MC difference, **129**
 - systematic, **127**
- scintillators, **28**
- solenoid, **27**
- spectrometer
 - dimensions
 - systematic, **142**
- spin
 - in GEANT, 63
- Standard Model
 - extensions, **12**
 - predictions, 7
- stopping distribution
 - simulation, 90
- STR
 - data-driven, 97, 98
 - decay parameter correction, 102
 - in simulation, 64
 - systematic
 - time-independent, **97**
- STRs, **41**
 - foil bulge
 - systematic, 104
 - gas density
 - systematic, 102
- TDC
 - simulation, 64
- TEC, **27**, 65
- Time Expansion Chamber, **27**
- time windows
 - classification, 43
 - definition, 43
- track reconstruction
 - helix fitter, 45
 - pattern recognition, 44
- trigger, **28**
- upstream stops, **67**
- weak coupling constants, 7
- window type
 - List of window types, **164**