3 EVENT RECONSTRUCTION

By the time the experiment ended, approximately 250 GB of data had been recorded to tape. The data on these tapes represented the hits on each chamber and hodoscope plane for each event, the general characteristics of the beam in any given spill and the overall state of the spectrometer. The reconstruction of the kinematics of the events from these raw data required several steps. After configuring the analysis for the specific data set, the tracks were reconstructed from the chamber hits. Once the tracks had been found, their paths through the SM0 and SM12 magnets had to be reconstructed while accounting for energy loss and multiple scattering in the absorbing wall and beam dump. These procedures are outlined in greater detail below.

3.1 Data Sets

During the nine months starting in September 1996 in which FNAL E866/NuSea took data, over 360,000 continuum dimuon events were written to tape. These events were grouped into several sequentially numbered data sets which were differentiated by the polarities and currents used to energize the three spectrometer magnets. Each data set fell into one of three categories, defined by the mass range for which the spectrometer was optimized. These categories were referred to as the low-, intermediate- and high-mass data, corresponding to spectrometer settings optimized for masses near the J/ψ, between the J/ψ and Υ masses, and near the Υ mass, respectively. Table 3.1 lists the data sets used in this analysis and the mass setting to which they belong. It also tabulates the currents used to power the spectrometer magnets, and whether the targets were filled with either the first or second sample of cryogenic liquids.

Table 3.1: Definition of the data sets. The SM3 magnet was operated at a single current of 4230 A, with the same polarity as the SM12 current.

mass setting	data set	SM0 current	SM12 current	target fill
low	5	−2100 A	2800 A	first
	10	2100 A	−2800 A	second
intermediate	9	0 A	2800 A	second
high	7	0 A	4000 A	first
	8	0 A	4000 A	second
	11	0 A	−4000 A	second

3.2 First Pass

Since only about $\sim 1\%$ of the events recorded to tape were due to a continuum dimuon event originating from the target, multiple analyses of the entire data sample were not feasible. It was therefore necessary to reduce this data sample to more manageable levels. A first pass analysis was performed which filtered out many of the bad (non-dimuon) events in the data sample, while preserving almost all of the events we were interested in. The computing resources for this analysis were provided by the Fermilab Computing Division. A number of IBM workstations were linked together into four parallel computing clusters called "farms", which were able to process the complete data sample in approximately two months.

3.2.1 Track Reconstruction

The analysis began by reading in the configuration of the spectrometer from various files, some specified by the user and some loaded automatically based on the data set being analyzed. Once configured, the analysis code examined the hits in the drift chambers at stations 2 and 3 to find a set of candidate tracks. If four of the six planes in a station registered a hit, the position of that hit was considered a possible track reconstruction point. Points which were inconsistent with the trigger which fired were dropped. The remaining reconstruction points were iteratively combined to form track segments, so long as those segments were consistent with a single particle forming both points.

The track segments were then compared with hits in station 1 to reduce further the number of track segments. Each segment was first extended to the SM3 bend plane, located between stations 1 and 2. The charge and momentum of the particle was not yet known, so a window-of-interest at station 1 which pointed back to the target in the X direction was searched. If none of the hits at station 1 met the search criteria, the track segment was discarded.

Once the track segment between stations 2 and 3 was connected with the station 1 reconstruction point (or points), the charge and momentum of the track were determined from the direction and size of the deflection of the track at SM3. The track segment was then extended back to station 4 to verify that it was indeed a muon. The concrete and lead shielding present at station 4 absorbed any hadrons and electrons which managed to penetrate the absorbing wall, but also caused any muons to multiple scatter. Since at this point the analysis knew the momentum of the track, a window consistent with the expected multiple scattering could be searched for hits in the five station 4 detectors. If the track fired less than three of the station 4 detectors, it was rejected on the grounds that it was not likely to be a muon.

3.2.2 Event Kinematics

At this point in the analysis, the charge, momenta and positions at the various detectors of the tracks in the event were known. It had also been determined at this point that each of the tracks was due to a muon. The tracks were then propagated back through the SM12 and SM0 magnetic fields, accounting for the energy lost in the hadronic absorbing wall and (if applicable) the dump. The paths of the tracks were reconstructed inside the magnetic field in 46-cm steps, using field maps which had been measured prior to the experiment. A correction[1] was also made to compensate for the energy lost by the muon where the track passed through either the absorbing wall or the beam dump.

In addition to suffering energy loss inside the absorbing wall and dump, the muons also underwent multiple Coulomb scattering (MCS). The correction for this began by reconstructing the track back to an effective scattering bend plane (ZSCPLN) and recording its position and momentum there. The track was then completely analyzed, noting the deviations in X and Y from the nominal target center. A correction to the momentum vector at the ZSCPLN was calculated from these deviations, and the reconstruction of the track repeated from that point. This procedure was iterated until the track converged on the nominal center of the target.

The production of pions and other hadrons in the target contributed to a large rate of single muons in the spectrometer as they decayed between the target and dump. A significant fraction of these muons followed a beam-like trajectory through the dump and absorbing wall. Because there was no way to differentiate between one of these single muons and a muon which was a part of a valid dimuon event, these muons could cause ambiguities in the analysis that could not be resolved. To prevent this from happening, muons which passed through a 10.2 cm

[1] The most probable energy lost in a given section of absorbing wall or dump was parameterized versus the momentum of the track. This energy loss parameterization was used to estimate the energy lost by the muon.

horizontal band centered on the nominal beam position at the dump were eliminated from the event sample.

After the kinematics of the individual tracks were reconstructed at the target, they were combined to form dimuon pairs. Less than 0.08% of the reconstructed events contained more than two tracks. From the charge of the tracks and their momenta, the kinematics of the event were determined. Events which reconstructed to a mass $< 2.0\,\mathrm{GeV}$ were rejected on the grounds that the acceptance for such pairs arising from the target was zero[2]. Events whose uniterated vertex[3] indicated they were unlikely to have originated in the target were also cut. This cut was realized in terms of the z-coordinate of the uniterated vertex (ZUNIN), which was restricted to values less than 350 inches from the center of the target.

3.3 Second Pass

The goal of the first pass analysis was to reduce the data sample to a more manageable level. The purpose of the second-pass analysis was to reconstruct the kinematics of the events as accurately as possible, minimizing any systematic errors which would affect the cross sections. In the remainder of this chapter, we

[2] There was a problem in the application of this cut which applied to events with more than two tracks. If one of the pairs which could be combined from the tracks in the event was below the mass cut, the entire event was cut rather than the particular combination of tracks. Studies indicate that the error introduced by this is on the order of 0.1%, which is negligible compared to other uncertainties in the experiment.

[3] The position at which the tracks made their closest approach before the MCS corrections were made.

will discuss the differences between the algorithms the first- and second-pass analyses used to correct for multiple scattering and energy loss. We will also outline the procedures used to determine the strengths of the fields in the spectrometer magnets, and the position and angles of the beam at the target.

3.3.1 Multiple Scattering

The ZSCPLN was used to correct for the multiple scattering that the muons underwent in traversing the absorbing wall and beam dump as described above. The first pass-analysis used a single position for the ZSCPLN, which was empirically determined by minimizing the correlation between the mass of the dimuon pair and the mean ZUNIN. This approach ignored the additional scattering that events which passed through the dump were subject to. The second-pass analysis based the position of the ZSCPLN on the length of the dump which the muon penetrated. It was found that for every 1 cm the track traveled through the dump, the ZSCPLN should be moved upstream 0.7 cm.

The second pass analysis also employed a different technique for finding the optimal position of the ZSCPLN. Monte Carlo studies indicated that the correlation between the mass of the events and ZUNIN was due to energy loss fluctuations in the beam dump and hadronic absorber, rather than difficulties in reconstructing the track angles as had previously been believed. The ZSCPLN in the second-pass analysis was placed at the position which optimized the reconstruction of the track angles at the target for Monte Carlo events. This resulted in a ZSCPLN located 12 m downstream of the target, near the end of the carbon section of the hadron absorber. The method used in the first pass analysis placed the ZSCPLN 22 m downstream of the target, well outside of the absorbing wall.

3.3.2 Energy Loss

As will be discussed below, determination of the strength of the magnetic fields from the data relied on an accurate reconstruction of the masses of the J/ψ and Υ resonances, and the ZUNIN distribution of the events. Deviations between the average energy lost by the muons and the parameterization used to correct for the energy loss in the analysis could cause systematic shifts in both mass and z-vertex reconstruction, and thus the measurement of the magnetic fields.

Every muon which passed through the spectrometer was subject to a large amount of energy loss in the absorbing wall and possibly the beam dump. A typical 200 GeV muon lost on average ~ 2 GeV in the absorber. Figures 3.1 and 3.2 show the energy lost in traversing the beam dump and absorbing wall for muons with different incident momenta. The calculations were performed using the TRAMU muon transport program [Gin 86]. The large, non-Gaussian tail skews the mean to much larger energy losses than the most probable value. It was important to parameterize the energy loss in the reconstruction of the data with emphasis on accurate (on average) reconstruction of the kinematics of the events rather than optimal resolution. Thus, the mean energy loss was used in this analysis.

Given the incident momentum of the muons on a specified material, the parameterization of the mean energy loss is a well defined problem. However, the incident momentum of the muons on the absorbing wall or dump was the quantity which we wished to estimate from our measurement of the momentum downstream of SM12 and some parameterization of the random distribution of energy loss inside the absorbing wall and dump. To further complicate the problem, muons which lost a large amount of momentum in the dump/absorber could be swept out of the acceptance. Muons with lower incident momenta would be more susceptible to this, introducing an acceptance effect into the parameterization of the energy loss.

Figure 3.3 shows the mean energy lost by muons in the 24″ section of the absorbing wall according to the Monte Carlo (where we knew the energy lost by each muon thrown). We plot the mean energy loss versus the momentum incident upon the 24″ of copper, and versus the momentum after the 24″ of copper. The "after" parameterization is generally flatter than the "incident" parameterization. This is due to the fact that larger mean energy losses in the "incident" parameterization will be shifted to smaller values of p_{after}. It should be clear from the figure that this effect is important, especially at lower momenta where we have most of our statistics in the data.

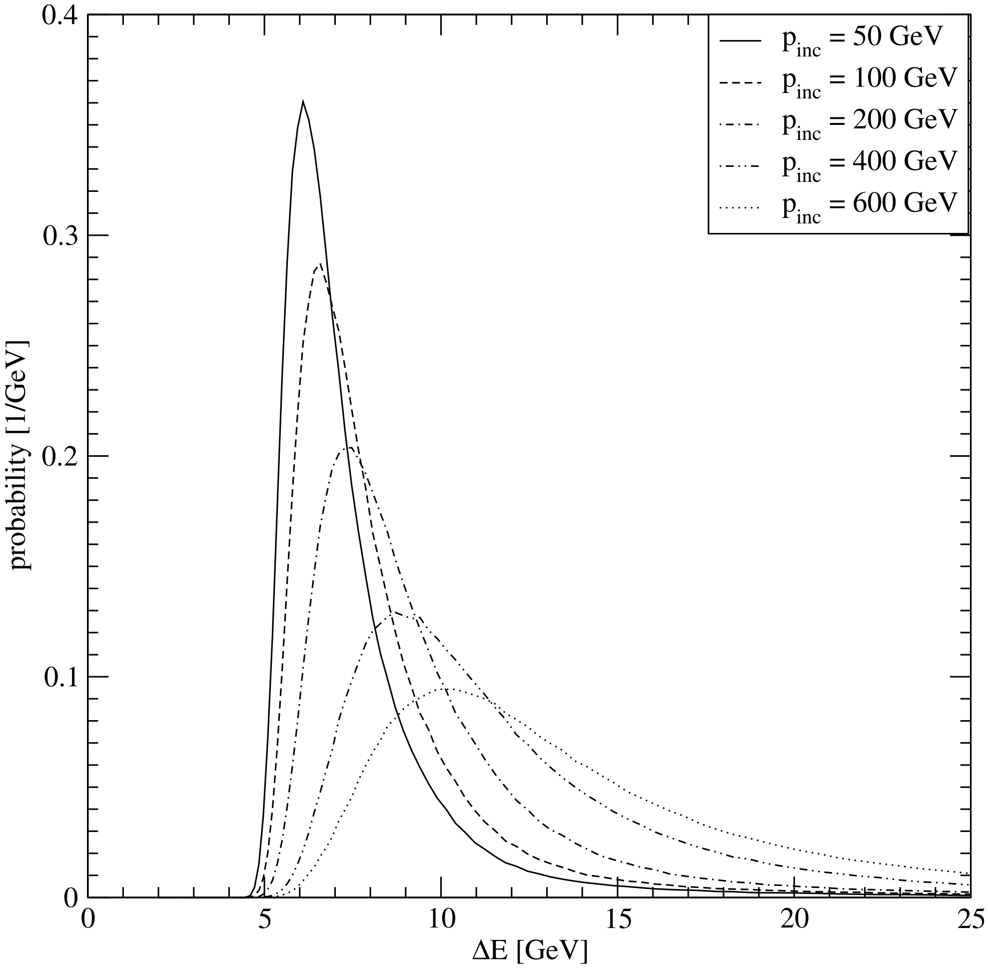

Figure 3.1: Energy loss distributions for muons passing through 144” of Cu (beam dump) for various incident energies.

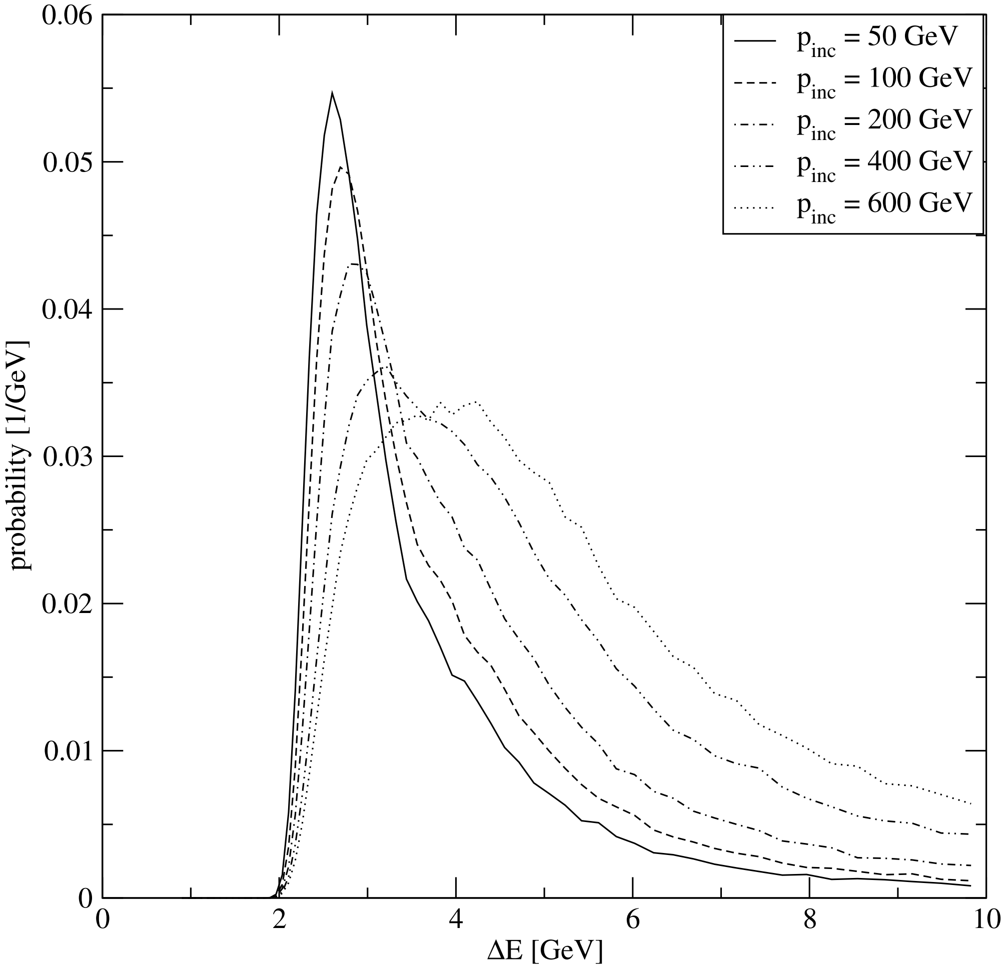

Figure 3.2: Energy loss distributions for muons passing through the entire absorbing wall – including the copper, graphite and polyethylene sections – for various incident energies.

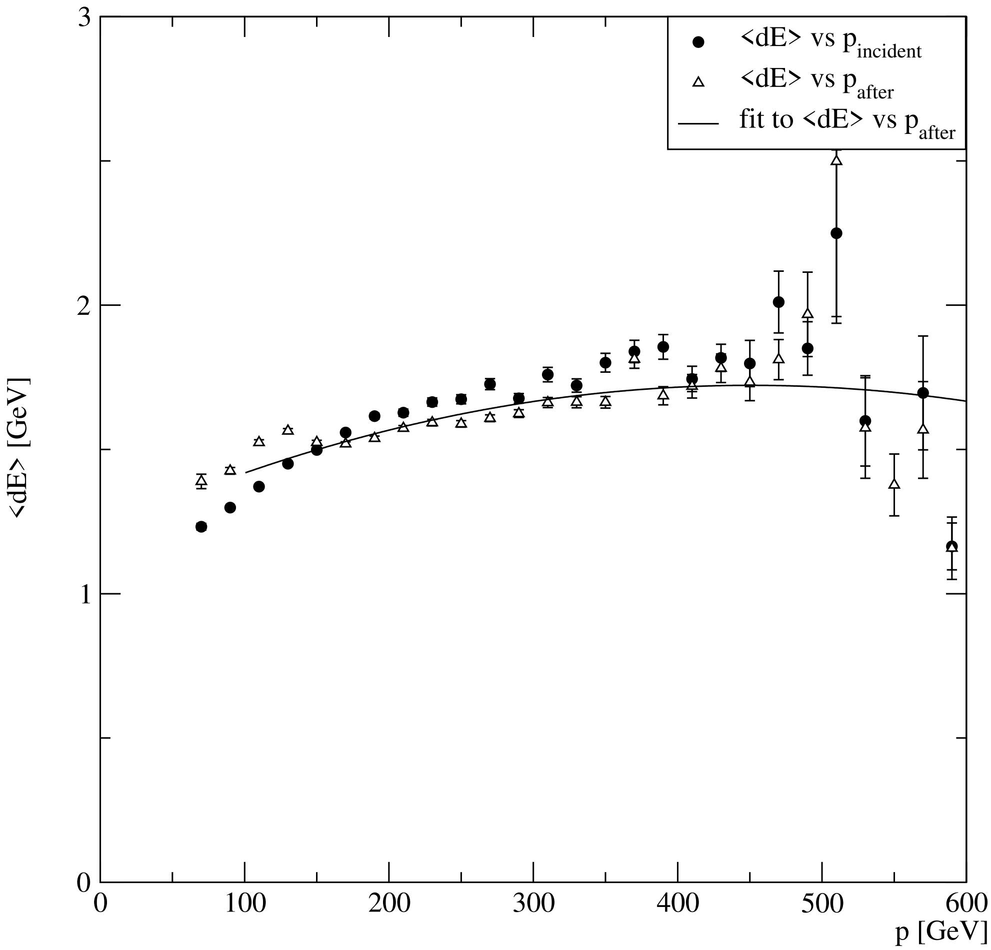

Figure 3.3: Mean energy loss for muons in the 24″ Cu section of the absorbing wall. Shown is the energy loss versus the momentum of the muon entering the section (circles) and the energy loss versus the momentum of the muon exiting the section (triangles). The solid line is the fit to the second set of points which was used in reconstructing the events.

Figure 3.3 also shows the fit to the energy loss in the 24″ copper section of the absorbing wall. The parameterizations which we used for all of the absorber sections in the high-mass data are given by

$$
\begin{aligned}
dE_{Cu} &= 1.2202 + 0.2228 \times 10^{-2} p_{\text{after}} - 0.2472 \times 10^{-5} p_{\text{after}}^2 \\
dE_{C} &= 0.2461 + 0.3135 \times 10^{-3} p_{\text{after}} - 0.4362 \times 10^{-6} p_{\text{after}}^2 \\
dE_{CH_2} &= 0.1857 + 0.2073 \times 10^{-3} p_{\text{after}} - 0.3501 \times 10^{-6} p_{\text{after}}^2
\end{aligned}
$$

where p_{after} is the momentum immediately downstream of the absorbing wall section being traced through. Deviations in the shape of this parameterization, but not intercept, were found for the low- and intermediate-mass data sets. When used to reconstruct Monte Carlo events, systematic shifts in the masses and ZUNIN of the events are held to acceptable levels.

3.3.3 Magnetic Fields

The magnetic fields in the three spectrometer magnets were measured prior to the installation of the absorbing wall. Measurements were made at each of the planned operating currents, except for the 4000 A setting of the SM12 magnet. This setting was not measured due to concerns over the structural integrity of certain parts of the magnet under the increased strain of higher magnetic forces. At the time, it was anticipated that the 2800 A setting of the SM12 field would be the largest current at which that magnet would be operated. Later during the experiment, the SM12 magnet was successfully operated at the 4000 A setting, providing the largest and cleanest sample of continuum dimuon events in the E866 data sample.

The magnet map used in the high-mass data utilized a measurement of the SM12 field from a previous experiment. However, the upstream portion of the magnet was reconfigured for E866. This region of the magnet map relied on an OPERA [Ope 94] calculation of the magnetic field, matched to the field map in the known region of the magnet. A similar procedure had to be performed to determine the magnetic fields in the beam dump, where it was not possible to measure the fields.

Although the shapes of the magnetic fields were well determined, the absolute strengths of the fields were uncertain at the $\sim 1\%$ level. Monte Carlo studies indicated that this level of uncertainty in the SM0 and SM12 magnets could be a significant source of uncertainty in the acceptance of the spectrometer, especially near its low mass edge. Figure 3.4 shows the ratio of Monte Carlo events generated with magnetic fields differing by 5% using a fast version of the E866 Monte Carlo. The relative acceptance falls by nearly 40% over the given mass range. A systematic uncertainty of $\sim 1\%$ in the SM12 field would correspond to a $\sim 8\%$ point-to-point systematic uncertainty in the acceptance, which would have been unacceptable given the precision of our data. Furthermore, systematic shifts in the SM3 field would cause systematic shifts in the reconstructed kinematics of the events. Such shifts were a problem, considering the steeply-falling nature of the

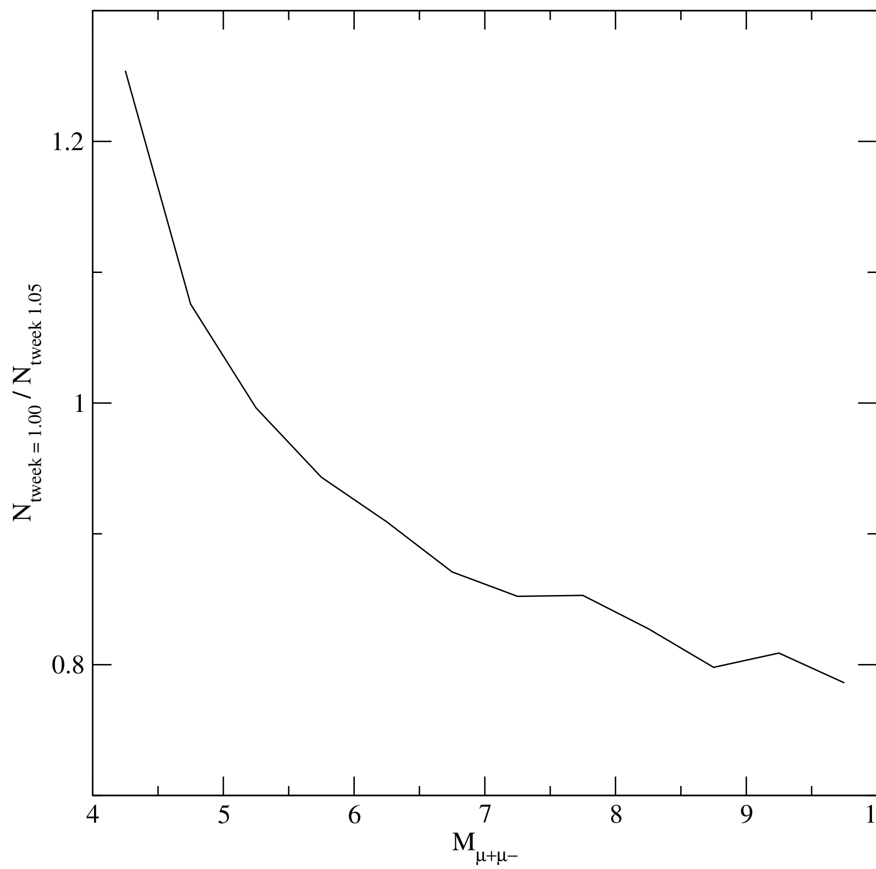

Figure 3.4: Ratio of Monte Carlo events generated with the nominal SM12 field to those generated with the SM12 field augmented by 5%. The ratio is shown for $0.5 \leq x_F < 0.55$.

cross section we were trying to measure.

The general procedure for determining the absolute field strengths from the data, once the impact of the energy loss parameterization had been limited, was as follows. We began by fitting the masses of any resonances (J/ψ, Υ) in the mass spectrum to a Gaussian plus a quadratic background, noting the reconstructed mass of the peaks. We also estimated the mean of the ZUNIN distribution. Deviations of these quantities from the nominal position of the target and the published resonance masses indicated an incorrect value for the SM12 and SM3 fields respectively. The data set was then re-analyzed with a multiplicative constant – referred to as a tweek (*sic.*) – applied to each of the magnetic fields to try to compensate. This was performed iteratively until the resonance masses converged to their published values, and ZUNIN was centered on the nominal target position.

The procedure outlined above was essentially the procedure used to determine the tweeks on the SM12 and SM3 magnets in the high-mass data. It is a slight simplification, since a shift in the SM3 magnet would also have caused a shift in the ZUNIN of the events due to the increased momentum and resulting decrease in the deflection of the muon by SM12. The one important change to note was that only events which missed the dump were studied. Muons which hit the dump went through a somewhat uncertain amount of material, contributing an uncertainty to the amount of energy loss which should have been corrected for. Additionally, the ZUNIN of the events with one track in the dump was largely determined by the other track in the event. This reduced the sensitivity of ZUNIN to the magnetic field in SM12.

Additional problems arose in the analysis of the low and intermediate mass data sets. In the intermediate-mass data set, we had only a limited number of events to work with. The reconstructed Υ mass, the only resonance available in that data set, was too uncertain to provide an accurate assessment of the SM3 tweek. Therefore, the value found in the high mass analysis was used for SM3, and the usual procedure of finding the SM12 tweek using the ZUNIN distribution performed.

A similar problem occurred in the low-mass data set, where we had an extra magnetic field (SM0) to worry about. In principle, since we could resolve both the J/ψ and ψ' resonances, we had the extra information needed to tune the SM0 field. In practice, the ψ' mass was relatively insensitive to the SM3 field, for a reasonable set of tweeks. Once again, the SM3 field tweek from the high-mass data was used in this data set. The SM0 tweek was found from the fits to the J/ψ mass, and the usual ZUNIN fitting procedure used to extract the SM12 tweek.

The procedures described above provided an estimate of the average field strength in a given data set. It also provided a way to estimate the systematic uncertainties on the tweeks by studying the impact of a given change in tweek on the ZUNIN distribution and the resonance masses. Uncertainties in the mass and vertex fits could then be translated into uncertainties in the magnetic fields. Once the uncertainties in the magnetic fields were determined, the corresponding acceptance uncertainty could be calculated by running the Monte Carlo using different input fields. We will describe this in greater detail in the next chapter.

Figure 3.5 shows the J/ψ and Υ mass regions from the low- and high-mass data, when reconstructed using the optimal magnetic field tweeks. The reconstructed J/ψ and Υ masses of 3.095 and 9.461 GeV compare quite well with the published values of 3.097 and 9.46 GeV [Gro 00]. The intrinsic widths of the res-

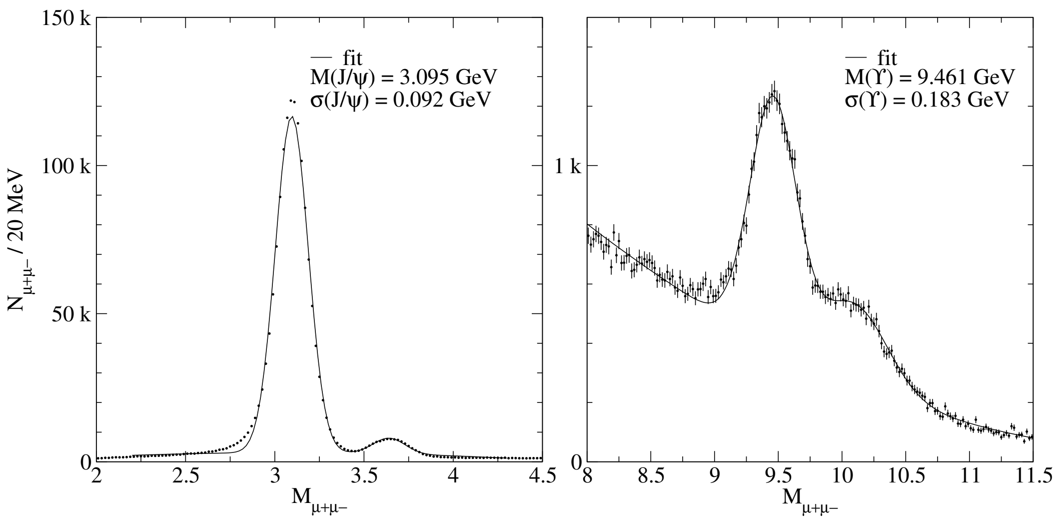

Figure 3.5: Reconstructed masses of the J/ψ and Υ resonance for the low- and high-mass data sets. Data from both targets are included.

onances (which are best expressed in keV), are smeared out by energy loss and multiple scattering in the absorbing wall and dump – our mass resolution is 92 MeV at the J/ψ mass and 183 MeV near the Υ. Figure 3.6 shows a fit to the ZUNIN distribution of the continuum dimuon regions, where the data has been summed over all data sets. In practice, fits to the ZUNIN distributions and resonance regions were performed separately in each data set. The deviations of the fits from their expected values amount to uncertainties of only $\pm 0.10\%$ and $\pm 0.15\%$ in the SM0/12 and SM3 tweeks, respectively.

One other source of systematic error in the tweeks was uncovered using the Monte Carlo. When Monte Carlo events were retracked through the reconstruction algorithm, the momentum found at the SM3 bend plane was $\sim 0.06\%$ larger than that which was actually generated by the Monte Carlo. Thus, the correct

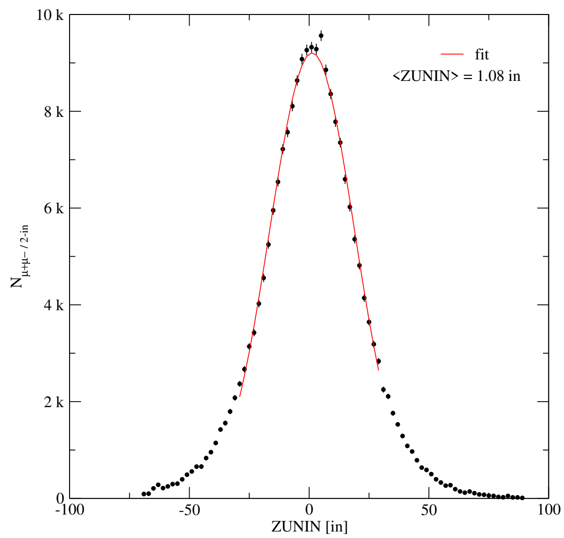

Figure 3.6: Distribution of empty-target-background subtracted events vs the uniterated z-vertex (ZUNIN), summed over both targets and all data sets.

field tweeks were subject to a small offset which had to be accounted for to obtain a more accurate simulation of the spectrometer. The origin of this effect was unclear. It may have been due to slight differences between the actual and simulated geometry of the apparatus, resolution effects coupled to the single bend plane approximation used for SM3, or some combination of these and other issues. This additional source of systematic uncertainty in the SM3 field does not contribute significantly to the overall size of the systematic uncertainties in the acceptance.

In table 3.2 we list the tweeks found for each mass setting using the previously described procedures. Studies of the stability of the mean ZUNIN and Υ masses versus data set and run number in the high-mass data suggest that a single value for the SM3 field tweek is reasonable. These studies also indicated that the SM12 field was reliably stable, and a single value was used for each data set within a

Table 3.2: Field strengths for each data set.

Mass setting	Data set	SM0 tweek	SM12 tweek	SM3 tweek
low	5	1.0140 ± 0.0010	1.0068 ± 0.0010	1.0049 ± 0.0015
	10	1.0140 ± 0.0010	1.0068 ± 0.0010	1.0049 ± 0.0015
int	9	N/A	1.0093 ± 0.0010	1.0049 ± 0.0015
high	7	N/A	0.9865 ± 0.0010	1.0049 ± 0.0015
	8	N/A	0.9865 ± 0.0010	1.0049 ± 0.0015
	11	N/A	0.9865 ± 0.0010	1.0049 ± 0.0015

given mass setting.

3.3.4 Beam Position and Angles

The position of the beam was determined by fitting the distribution of the x and y vertices of the raw data at the target position. The uniterated x and y positions of the two tracks in the event were averaged together to yield the x and y vertices. Run-to-run variations in beam position were found to be negligible for our purposes.

The beam angle was found through an iterative process involving the Monte Carlo. In principle, the Y component (and similarly the X component) of the beam angle could be found from just the data by looking at the muon momentum vectors $\theta_y^{dimuon} = (p_y^+ + p_y^-)/(p^+ + p^-)$. However, asymmetries in the dump cut in the first pass analysis limited the accuracy of this technique. An independent test of the validity of the beam angle extracted from the data was to check the ϕ_p distribution in the data against that of the Monte Carlo. Since we know that the ϕ_p distribution must be isotropic, and having generated an isotropic distribution in the Monte Carlo, any difference between the distribution in the data versus the Monte Carlo could only result from a difference in the beam angle between the experiment and the Monte Carlo. The beam angles were improved by minimizing these differences, and we estimate that the uncertainty in the beam angle contributes no more than 5% ($\sim$ 0.7 GeV) of the transverse momentum of an average event.