4 ANALYSIS

The differential cross section is defined as the number of interactions per target particle leading to an event with kinematics between Ω and $\Omega + \delta\Omega$ per number of incident particles per unit area. We express this as

$$\frac{d\sigma}{d\Omega} = \frac{N/\Delta\Omega}{N_T(\bar{N}_i/A)}. \tag{4.1}$$

where N is the number of dimuon events from the target, N_T the number of target particles in the path of the beam, $\bar{N}_i$ the mean number of protons incident on the target and A the cross sectional area of the beam.

The simple expression in equation 4.1 is somewhat misleading, since we did not measure the number of pp and pd interactions which led to a dimuon event directly. Instead, we measured the yields of oppositely charged muon pairs from the interaction of a proton beam with target flasks containing liquid H_2 and D_2. These yields were subject to contamination by background events originating from the interaction of the beam with the SWIC and the front and back faces of the target flasks, and from the random coincidence of two uncorrelated muons produced in the same beam bucket. The yields were also subject to losses due to the acceptance of the spectrometer and inefficiencies in the detection of the muons. Equation 4.1 also relies on quantities such as the number of target particles and incident protons, which were also not directly measured, and the cross sectional area of the beam, which was not measured with any great precision.

Fortunately, we can rewrite equation 4.1 entirely in terms of quantities which were directly measured, while eliminating terms which were not measured well. We define the number of dimuons, corrected for backgrounds, as $N_{\mu^+\mu^-}$. This can be related to the number of interactions in equation 4.1 by

$$N = \frac{N_{\mu^+\mu^-}}{\alpha\epsilon} = \frac{N^{\text{target}}_{\mu^+\mu^-} - N^{\text{background}}_{\mu^+\mu^-}}{\alpha\epsilon} \tag{4.2}$$

where α is the geometric acceptance of the spectrometer, defined as the probability that both tracks in a given event would traverse all of the active detector elements of the spectrometer and survive all data cuts, ϵ is the efficiency with which the given event was detected, and the "target" and "background" superscripts refer to events which originated in the liquid target and background events respectively. The number of target particles can be expressed in the usual way:

$$N_T = N_A \, \rho \, L \, A \tag{4.3}$$

where N_A is Avogadro's number, ρ is the density of cryogenic liquid, L is the length of the target, and A is the cross sectional area of the beam.

To get the mean number of incident protons, we need to average over the attenuated flux in the target. The number of protons at any point along the length z of the target can be expressed in terms of the incident number of protons and the hadronic absorption length λ, as

$$N_i(z) = N_p e^{-z/\lambda}. \tag{4.4}$$

Averaging over the total length of the target L yields

$$\bar{N}_i = \frac{1}{L}\int_0^L N_i(z)\,dz = \frac{N_p\lambda}{L}(1 - e^{-L/\lambda}). \tag{4.5}$$

Substituting equations 4.2-4.5 into equation 4.1 results in an expression for the differential cross section:

$$\frac{d\sigma}{d\Omega} = \frac{N_{\mu^+\mu^-}}{\mathcal{L}\alpha\epsilon\Delta\Omega}, \tag{4.6}$$

where we have grouped several of the factors in the denominator into a single quantity, called the luminosity, given by

$$\mathcal{L} = N_A\rho\lambda(1 - e^{-L/\lambda})N_p. \tag{4.7}$$

We will now explain in detail how each of the quantities in equation 4.6 was determined.

4.1 Event Yields

The total event sample produced by the second pass analysis contained several questionable events. These included backgrounds from uncorrelated muon pairs and correlated pairs originating outside of the target volume, events which were reconstructed by the analysis but were not responsible for the trigger which caused the DAQ to read out, and events which occupied regions of the spectrometer which complicated the reconstruction of their kinematics. These events were either filtered out of the second pass event sample, or independently binned and subtracted from it, to obtain the event yields used to calculate the cross sections. The same set of cuts were used on the Monte Carlo events used to calculate the acceptance. Simulations showed that contributions from the interactions of secondary hadrons produced in the target were negligible.

In order to avoid the J/ψ and Υ resonances, we selected events within the mass range $4.2 \leq M_{\mu^+\mu^-} \leq 8.7$ GeV, and events with masses $M > 10.85$ GeV. As the acceptance above the Υ resonance was quite small in the low and intermediate mass data sets, we only used the high mass data in that region. Due to uncertainties in the magnetic fields, data from the intermediate mass setting below 6.2 GeV were also discarded. Figure 4.1 shows the total E866 mass spectrum, with the contributions from the low-, intermediate- and high-mass data sets in their respective allowed mass ranges. The resonance regions have been retained, but all other cuts were applied. The number of events surviving all cuts is shown in table 4.1.

4.1.1 Target Flask Background

One of the larger sources of contamination was due to the beam interacting with the SWIC ($\approx$ 188 cm upstream of the target) and the walls of the target flasks which contained the cryogenic liquids. Figure 4.2 shows the yields from the cryogenic liquid and empty target flasks plotted versus ZUNIN. The SWIC is clearly visible in the empty target distributions, and much of the contamination was re-

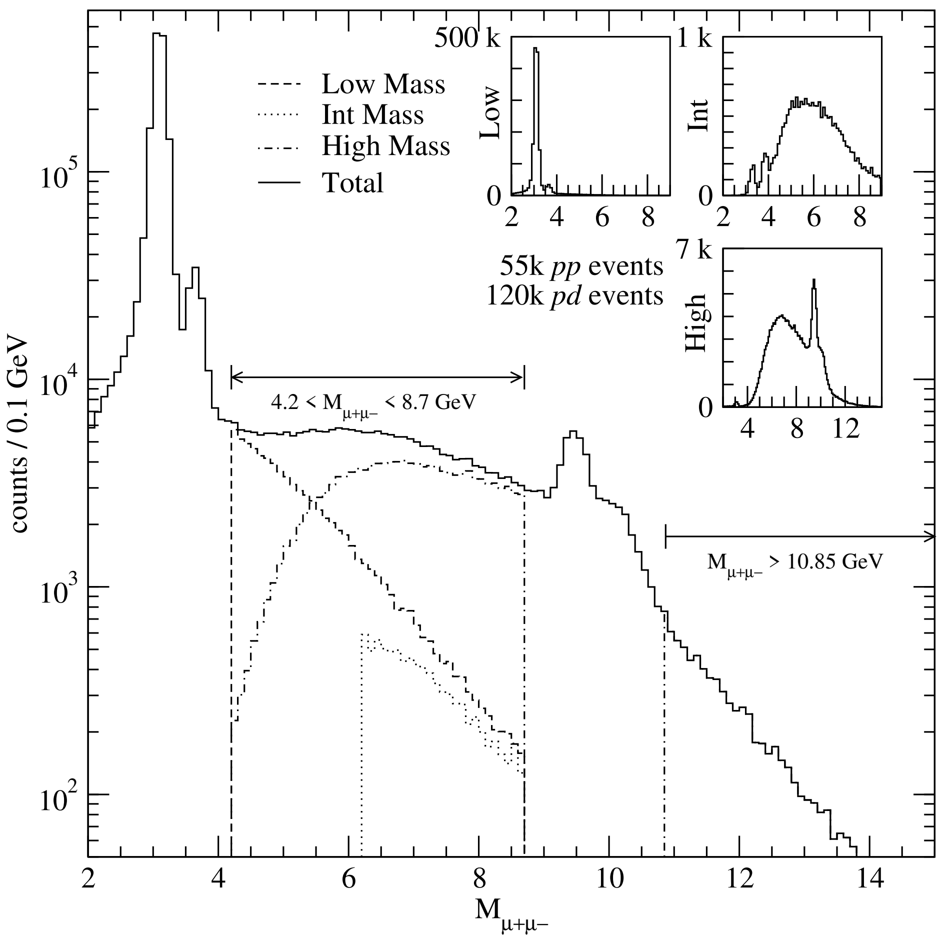

Figure 4.1: Total E866 mass spectrum surviving all cuts. The contribution of the low, intermediate and high mass data sets superimposed in their respective allowed mass ranges. Inset figures show the low, intermediate and high mass data sets on a linear scale over their entire mass range.

Table 4.1: Number of events in each data set which survived all data cuts.

Set	H_2 target	D_2 target	Empty target
5	17217	36962	699
7	9548	21030	425
8	20049	44129	1007
9	1768	4966	63
10	1658	3509	60
11	4850	10545	230
total	55090	121241	2484

moved by simply applying cuts (ZUNIN $> -50''$ in the low mass data, and ZUNIN $> -60''$ in the intermediate and high mass data). Events with ZUNIN $> 60''$ in the low mass data, and ZUNIN $> 90''$ in the intermediate and high mass data were also rejected, as they were likely to have originated in material downstream of the target.

The remaining background from the target flask was measured during the experiment, allowing us to subtract it from our event yields. A correction was made to the empty target yields based on the uniterated z-vertex of the event to account for the $\approx 7\%$ and $\approx 14\%$ attenuation of the beam between the front and back end caps of the flasks containing liquid hydrogen and liquid deuterium. This

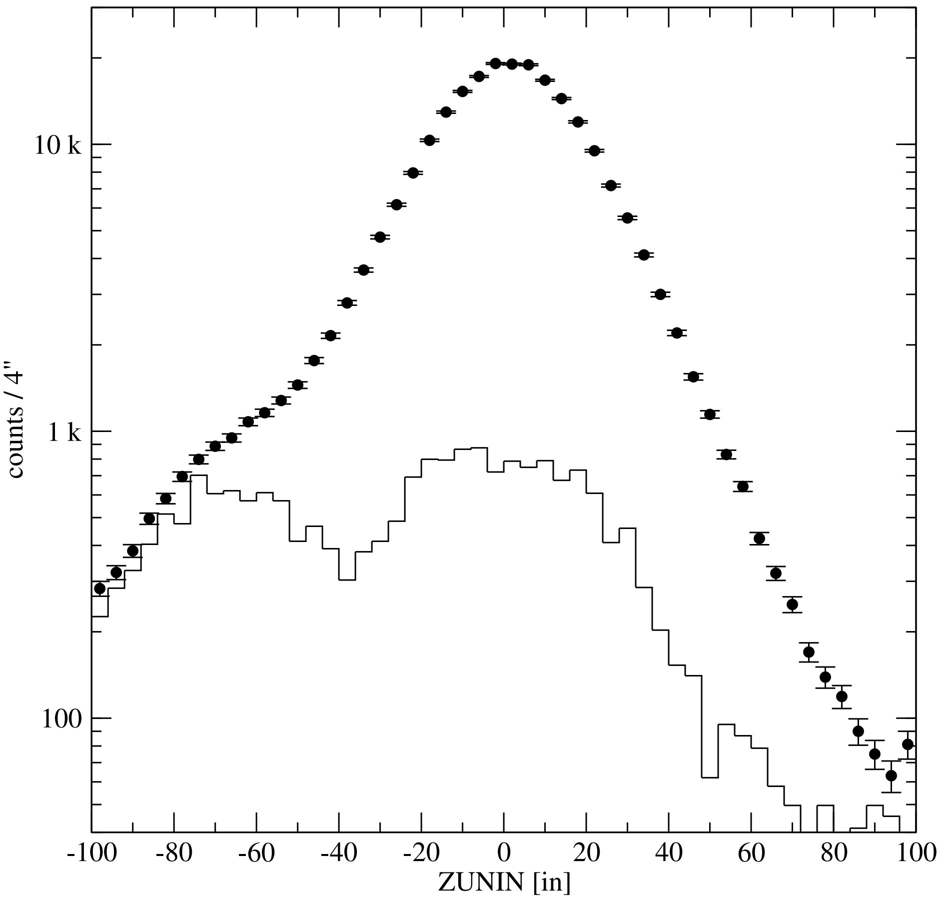

Figure 4.2: Comparison of ZUNIN distributions from the sum of the liquid hydrogen and deuterium targets (circles) to the contribution from the empty target (histogram), as calculated for deuterium.

correction was given by

$$N_{\mu^+\mu^-}^{empty} = \begin{cases} N_{empty} \times e^{-L/\lambda}, & \text{if ZUNIN} > 0 \\ \\ N_{empty}, & \text{if ZUNIN} < 0 \end{cases} \tag{4.8}$$

where we corrected events which appeared to come from downstream of the nominal target position $z > 0$ by the attenuation of the cryogenic liquid appropriate to the target.

4.1.2 Combinatoric Background

Another important background source was the random coincidence of oppositely charged, uncorrelated muons produced in the target in the same beam bucket. The numbers and kinematic distributions of these random events were estimated using event samples triggered on the the like-sign (PhysA2) and single-muon (PhysB2) triggers. These random events were typically confined to masses below 5 to 6 GeV and were only important in the low- and intermediate-mass data sets, the latter being a non-issue due to the 6-GeV cut in the intermediate mass.

The single-muon trigger was used to trigger on what appeared to be single muons from the target traversing one of the matrix elements. Due to the high rate of such events from the target, these triggers were prescaled when in use during normal run conditions. Single-muon event samples were taken either during the normal course of the data collection or in special low-intensity runs. The resulting event samples were subjected to the same analysis chain as the opposite-sign data.

The single muons were then randomly combined with other single muons to form an uncorrelated dimuon event. This produced a sample of random dimuon events which, when properly normalized, could be subtracted from the data sample. Figure 4.3 compares the total dimuon yield in the data (summed over both targets and the high and low mass data sets) to the combinatoric background. The low mass data set contributes the largest contamination – approximately 5% of its total yield was due to randoms.

The normalization of the randoms was accomplished by comparing a subset of the randoms (those with two muons with the same charge), to events which satisfied the like-sign trigger in the data. These like-sign events were analyzed in the same way as any other event, except that one of the tracks was reflected about $y = 0$ when the kinematics of the events were calculated. Figure 4.4 shows the real like-sign events compared to the combinatoric like-sign events, scaled to match the real like-signs.

4.1.3 Trigger Cuts

For any given event, the possibility existed that the muons which fired the trigger and resulted in the event being read out by the DAQ were not the same muons which were reconstructed by the analysis. The easiest way for this to occur was to have a single, beam-like muon fire a set of hodoscopes, one of which was outside of the matrix, while simultaneously a second, low-momentum muon fired the missing

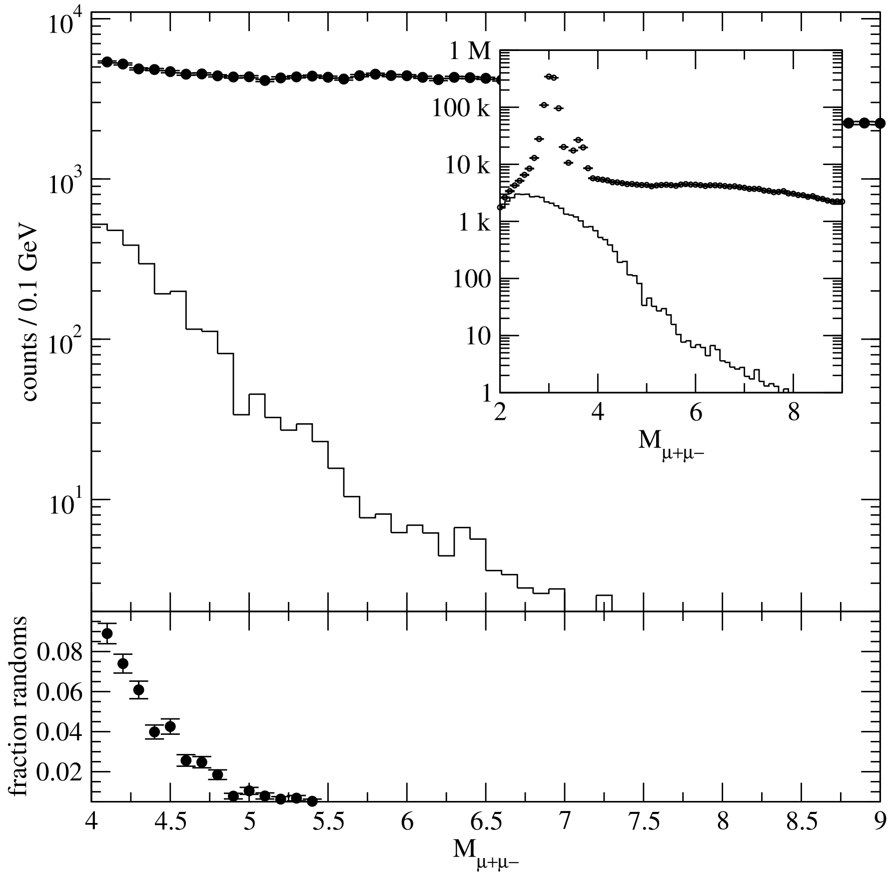

Figure 4.3: Comparison of the total E866 dimuon yield (dots) to the combinatoric background (histogram). The top panel shows the unsubtracted mass distributions, with the inset showing the same distributions down to 2 GeV. The bottom panel shows the fractional contribution of randoms vs. mass.

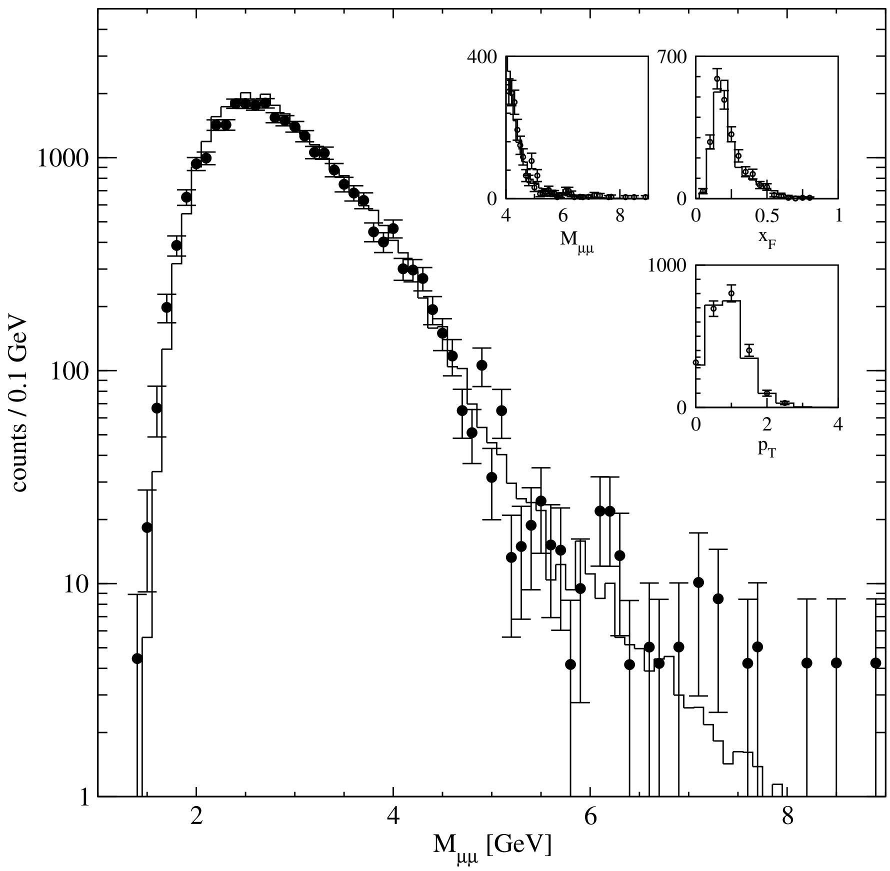

Figure 4.4: Comparison of like-sign distributions measured in the data (dots) to the like-sign distributions measured using the specialized trigger discussed in the text (histogram). The data and combinatoric like-signs are summed over the low and high mass data sets and both targets. The inset figures show the mass, x_F and p_T like-sign distributions over the mass range used in this analysis.

matrix element. When combined with a second, uncorrelated, oppositely charged muon which fired a valid matrix element, this would appear to be a high p_T continuum dimuon event. We therefore had to check that the muons in each dimuon event satisfied the trigger condition which had caused the reading out of that event. This was done by comparing the reconstructed positions of the muons at each hodoscope plane to the aperture of the hodoscopes in the trigger which fired. If the track was outside that aperture (which was widened slightly so that we would not lose events due to chamber resolution effects), it was discarded. Events which fired PhysA2 (the like-sign trigger) were binned separately from events which fired the other PhysA triggers.

In addition to the trigger cuts outlined above, comparisons of the data and Monte Carlo suggested that there was a problem with the trigger hardware corresponding to certain matrix elements. To first order, the acceptance of the spectrometer should be left-right symmetric. A separate analysis [Bro 01], which studied events from the dump in the high mass data, noticed that several matrix elements yielded much more data relative to Monte Carlo on one side of the spectrometer than the other. This discrepancy was much larger than could be explained by left-right asymmetries in the spectrometer, and thus these matrix elements were suspect. Similar studies were performed for the target data, confirming that there was a problem with several matrix-elements. These matrix elements, which tended to be at the edge of the acceptance where the statistics were generally poor, were eliminated from study.

4.1.4 Event Geometry

Several events which passed through all of the second pass cuts occupied regions of the spectrometer which complicated the reconstruction of their kinematics, or were judged to raise questions about their origins. Events which had a muon which appeared beam-like (tracks with $|\theta_y^{target}| < 0.006$ or tracks which were on the wrong side of the dump given their charge) were discarded. Such tracks could have originated from the decay of a hadron produced in the dump, and also followed a trajectory which skimmed the edge of the dump, calling into question the amount of energy lost by the track.

Another problematic class of events was found by studying the Monte Carlo. The energy lost by tracks in the Monte Carlo was plotted versus the x position of the tracks at SM3. Tracks near the edge of the acceptance, $|x_{SM3}| > 24''$, had significantly more energy loss in them than tracks with $|x_{SM3}| < 24''$. The source of this excess energy loss was the walls and teeth of the SM12 magnet, indicating that this class of tracks passed through a region of the magnetic field which was highly uncertain.

As discussed in the previous chapter, energy-loss fluctuations and multiple coulomb scattering cause the events to appear to reconstruct away from the actual interaction point in the target. To limit the fluctuation of energy loss and the

Table 4.2: Spill quality cuts. Events were kept if they obeyed the inequality.

	mass setting		
Cut	low	intermediate	high
SEM	> 150	> 100	> 100
IC3	> 5000	> 5000	> 5000
AMONSB / AMON	> 0.9	> 0.9	> 0.9
DUTY	> 25	> 50	> 25

effects of multiple scattering in the data, and to provide additional discrimination against events which originated outside the target volume, we required the event to reconstruct to within $2.5''$ of the nominal transverse position of the beam.

4.1.5 Spill Quality Cuts

During each beam spill, a number of quantities were monitored which told us the overall quality of the data being collected. These included the intensity of the beam (SEM and IC3), the luminosity of the target (AMON), the live-time of the spectrometer (AMONSB/AMON) and the beam duty factor (DUTY). Table 4.2 lists the spill quality cuts which were required for each spectrometer setting.

4.2 Luminosity

The intensity of the beam in each spill was measured with the SEM6 counter on the Meson East beam line. This recorded a SEM count which was proportional to the number of protons in the spill. The SEM6 counter was calibrated during special calibration runs where a thin copper foil was inserted into the beam and a SEM count accumulated during a number of spills. By studying the rate of the 1368 keV γ's emitted by the ^{24}Na produced in the interaction, the number of protons incident on the foil could be deduced, and the SEM response calibrated. These calibrations rely upon a measurement of the ^{24}Na production cross section [Bak 84] of 3.90 ± 0.11 mb[1].

Table 4.3 shows the results of several calibration measurements dating back to the original E605 experiment. Based on these measurements (taking the more recent measurements with greater weight), we estimate a SEM response of $(0.79 \pm 0.051) \times 10^8$ protons / SEM[Bro 97]. The uncertainty in the normalization was determined by adding the uncertainty due to the ^{24}Na cross section in quadrature with the standard deviation of the measurements in table 4.3.

Studies of the linearity of the SEM response were also performed by comparing the SEM counts to the luminosity as measured by the AMON and WMON detectors, and comparing to a redundant measurement of the beam intensity using the

[1] Although reference [Bak 84] measures the ^{24}Na cross section using 400-GeV protons, it is generally believed that the cross section is independent of energy.

Table 4.3: Measured SEM calibrations on the Meson East beam dating back to 1984.

Date	Protons	Protons/SEM $\times 10^8$
01/26/84	2.8E13	.83
02/11/85	6.9E12	.78
03/25/85	2.0E12	.85
05/06/85	2.0E12	.90
07/25/85	9.9E13	.80
08/29/85	9.6E13	.80
11/10/87	3.4E13	.79
11/16/87	2.4E13	.77
02/11/88	3.6E13	.78
10/29/96	1.7E12	.885
10/29/96	8.5E11	.915
12/20/96	5.1E13	.76
08/15/97	5.0E11	.86
08/29/97	2.0E11	.77

ion chambers. These studies showed the SEM response was linear with increasing beam intensity, but subject to a constant offset which varied from run to run. These offsets, shown in table 4.4, were subtracted from the SEM counts before calculating how many protons were incident on the target in each spill.

Both the density and attenuation lengths of the targets depended upon the exact composition of the targets. Therefore, assays of the cryogenic liquids were performed using samples taken when the target flasks were emptied for the Christmas shutdown (end of data set 7) and at the end of the planned run (end of data set 11). The cryogenic liquids used in the first fill had negligible amounts of contaminants. The same was true for the liquid hydrogen in the second fill. However, two separate analyses of the second fill's liquid deuterium showed a significant HD contamination. Table 4.5 shows the measured contamination and the estimated nucleon content in the target.

In order to determine the target densities, the vapor pressures of the cryogenic targets were measured and periodically recorded. The high circulation rate of the liquids, combined with the comparatively low beam current ensured that the cryogenic liquids did not boil during the spill which would have reduced the densities of the targets. Table 4.6 shows the mean target pressure for selected data sets in the experiment.

Cryogenic data tables [Rod 73] were used to calculate the target densities from the average vapor pressures. Taking the pressure P to be measured in psi, and

Table 4.4: Offsets to the SEM counters which were subtracted from each spill. I_{run} is the run number.

Starting run	Ending run	SEM offset
1551	1588	27
1589	1620	$1.479 \times I_{\mathrm{run}} - 2436$
1631	1690	28
1879	1884	28
1886	1887	155
1888	——	275
1889	1944	375
1945	2020	130
2021	2046	0
2049	2053	98
2056	2078	100
2079	2115	30

Table 4.5: Percent molecular and atomic abundance of second deuterium fill.

Molecule	percent abundance (first analysis)	percent abundance (second analysis)
D_2	94.13 ± 0.58	92.7 ± 0.8
HD	5.82 ± 0.58	6.89 ± 0.69
H_2	0.05 ± 0.01	0.147 ± 0.015
other	–	–
Atom	**estimated abundance**	
D	97.00 ± 0.6	
H	3.00 ± 0.6	

Table 4.6: Target pressures (in psi) for selected data sets.

Data set	hydrogen	deuterium
5	14.97	14.92
7	15.04	14.96
8	15.11	15.17
11	15.15	15.21

density ρ in g/cm^3, the expression for the density of hydrogen is given by

$$\frac{1}{\rho_h} = 62.473\,\{0.2115 + 1.171 \times 10^{-3}P - 1.109 \times 10^{-5}P^2\}, \tag{4.9}$$

and that for deuterium by

$$\rho_d = 4.028 \times 10^{-3}\{43.291 - 3.4176\frac{P}{14.6959} + 0.5783(\frac{P}{14.6959})^2\} \tag{4.10}$$

These result in densities which vary from data set to data set by no more than $\sim 2\%$, which is negligible compared to the 6.5% uncertainty in the beam normalization. From the measured pressures, the densities of the hydrogen and deuterium targets were determined to be

$$\rho_h = 0.0706\,\mathrm{g/cm}^3$$

$$\rho_d = 0.1630\,\mathrm{g/cm}^3.$$

As reported in the particle data book [Gro 00], the hadronic interaction lengths for hydrogen and deuterium are 50.8 g/cm^2 and 54.7 g/cm^2 respectively. To express these lengths in terms of the physical length, as opposed to the "thickness" of the target, we need to divide by the density of the targets. For the cryogenics used in the first fill, and the liquid hydrogen from the second fill, this is all we needed to do. For the liquid deuterium from the second fill, however, we needed to correct for the hydrogen contamination by taking the weighted average. The resulting hadronic absorption lengths are

$$\begin{aligned}
\lambda_{lh2} &= 719.5\,\text{cm} \\
\lambda_{ld2}^{(\text{1st})} &= 335.6\,\text{cm} \\
\lambda_{ld2}^{(\text{2nd})} &= 340.6\,\text{cm}
\end{aligned}$$

where the superscripts indicate either the first or second fill. The total integrated luminosities for each data set and the entire data sample are shown in table 4.7.

4.3 Trigger Efficiencies

In order to properly measure the efficiencies of the hodoscopes, it was necessary to use a sample of events which were triggered independently of the hodoscopes being measured. Because we only use the information from the Y hodoscopes in the experiment, events which triggered on the X hodoscopes could provide this independent measurement. Unfortunately, the rate of events which satisfied the X134L·X134R trigger (a muon on either side of the spectrometer) was prohibitively high to take along with the physics data. Therefore, several specialized runs were taken periodically to determine the efficiencies of the Y hodoscopes. A specialized

Table 4.7: Integrated luminosities in the E866 data sets.

Data Set	Integrated H_2 Luminosity (nucleon/nb)	Integrated D_2 Luminosity (nucleon/nb)
5	4.74×10^6	1.04×10^7
7	1.97×10^7	4.34×10^7
8	4.04×10^7	8.95×10^7
9	2.95×10^6	6.65×10^6
10	4.72×10^5	1.03×10^6
11	1.03×10^7	2.28×10^7
Total	7.86×10^7	1.74×10^8

trigger required that five hodoscope planes on one side of the spectrometer were hit. Reconstructed single muon tracks which fired at least six hodoscopes (one in each plane) were then studied to determine the efficiency of the hodoscope in the seventh plane. The typical efficiencies measured were $\sim 97\%$. However, these specialized runs did not measure the efficiencies near the edge of the acceptance with sufficient precision. To increase the statistics in these regions, we used events from the full data set. Although these events introduced a bias from the trigger to the efficiency measurement, they did probe the edge of the acceptance with far more statistics. The efficiencies measured here were checked against those of the specialized runs and found to be largely consistent.

4.4 Chamber Efficiencies

In order to study the efficiency of each detector plane, we began with a sample of tracks which would have been reconstructed regardless of whether or not the particular detector plane being studied had fired or not. This was possible due to the redundant tracking information at each station. Each plane was divided into quadrants, and the ratio of the number of tracks which fired in a given quadrant to the total number of tracks passing through the quadrant was used to estimate the efficiency of the drift chamber.

The use of single-hit TDC's in the DAQ electronics introduced an additional source of inefficiency. A second charged particle which passed through a drift cell

Table 4.8: Rate-dependence correction factors for the E866 data. Each event is weighted by $1 - SEM \times R_i$, where i is the target.

Mass Setting	R_d	F	$R_h = R_d/F$
Low	-0.2035×10^{-4}	1.51	$-.1348 \times 10^{-4}$
Intermediate	-0.4976×10^{-5}	1.18	$-.4152 \times 10^{-5}$
High	-0.2696×10^{-5}	1.78	$-.1515 \times 10^{-5}$

following the first particle in an event would not have its position recorded to tape, resulting in the event being more difficult to reconstruct. This effect would have a definite dependence on the event rate, and therefore the beam intensity.

To correct for the rate dependence, the number of reconstructed events per unit beam intensity (measured in terms of SEM counts) was plotted versus the beam intensity and fit to a straight line. The fits for hydrogen and deuterium were constrained by the relative number of extra drift chamber hits in an average event $F = \frac{N_d^{\text{extra}}}{N_h^{\text{extra}}}$ for each target [Tow 99]. This meant that only one parameter, the slope R_d of the efficiency-versus-intensity plot for deuterium, needed to be determined for each mass setting. Figure 4.5 shows the fit for the high mass data, and table 4.8 shows the corrections obtained by reference [Tow 99].

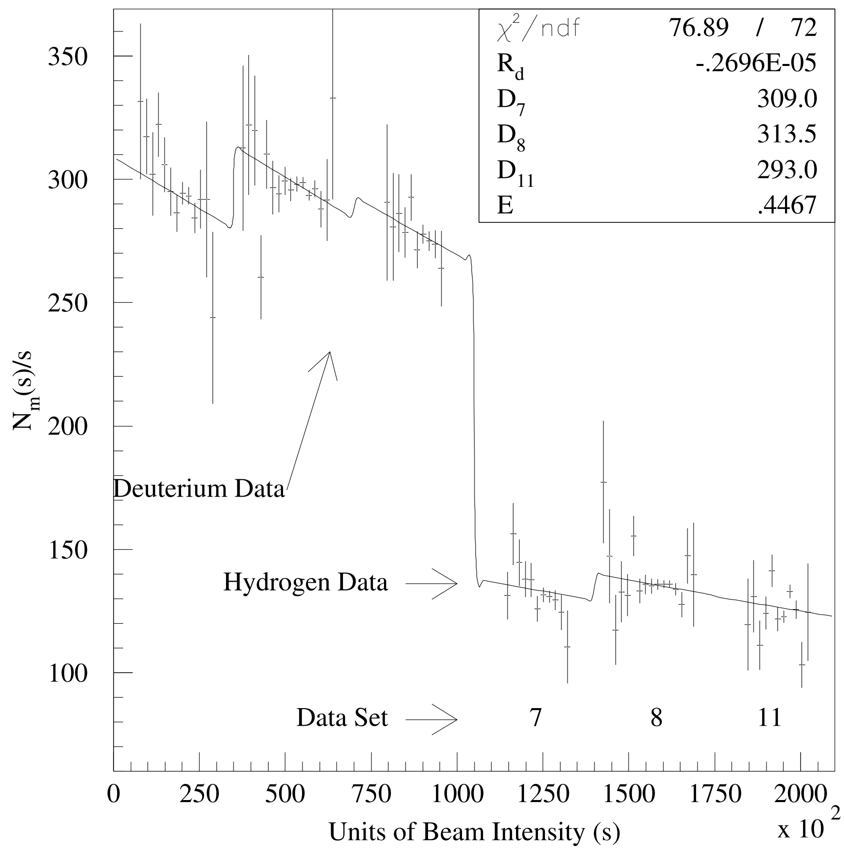

Figure 4.5: Dimuon yield (N_m) per unit beam intensity (s) versus beam intensity for the high mass data, as given in reference [Tow 99].

4.5 Monte Carlo and Acceptance

The acceptance of the spectrometer was a complicated function involving the detector geometry, the magnitudes of the magnetic fields in each of the three spectrometer magnets, the efficiencies of the detectors and trigger, and the various physical processes which muons traversing the spectrometer were subject to. This level of complexity required the use of a Monte Carlo simulation to determine the acceptance of the spectrometer.

The E866 Monte Carlo consisted of custom routines designed to simulate the production of dimuon events, the passage of the resulting muons through the spectrometer, and the detector response to their passage. The output of the Monte Carlo was a file whose format was identical to that produced by the DAQ during the experiment. This enabled us to analyze the Monte Carlo events with the exact same analysis chain as the real data, minimizing any systematic differences between the data and Monte Carlo.

To further minimize these differences, the Monte Carlo was configured to simulate as closely as possible the state of the spectrometer while it was taking data, the physical processes the muons were subject to in passing through the spectrometer, and the physical processes giving rise to the dimuon events. In chapter 3 we discussed how we obtain information about the state of the spectrometer (magnetic fields, detector alignments and efficiencies, etc...). We will discuss how events were generated in the Monte Carlo, the simulation of the physical processes the resulting muons were subject to, and the final set of calculations needed to compute the acceptance of the spectrometer.

4.5.1 Event Generation

In order to generate a continuum dimuon event, we had to sample the six kinematic variables $M_{\mu^+\mu^-}$, x_F, p_T, ϕ_p, ϕ_d, and θ_d from distributions which closely approximate the actual differential cross section $d^6\sigma/dMdx_Fdp_Td\theta_dd\phi_pd\phi_d$. *A priori* we know the cross section differential in the angular variables goes like,

$$\frac{d^3\sigma}{d\theta_d d\phi_d d\phi_p} \propto 1 + \cos^2(\theta_d) \tag{4.11}$$

since there is no preferred direction in the experiment to give a ϕ dependence, and the virtual photon in a $q\bar{q}$ annihilation is produced transversely polarized as predicted by QED and confirmed by experiment [Con 89, McG 99, Cha 99, Bro 01].

Next-to-leading order calculations of the doubly-differential cross section $d^2\sigma/dMdx_F$ agree well with existing measurements of the deuterium and nuclear cross sections, but problems arise in QCD in calculating the full triple differential cross section $d^3\sigma/dMdx_Fdp_T$. Therefore, an empirical form [Kap 78] for the p_T dependence of the cross section was used

$$\frac{d^3\sigma_{\text{empirical}}}{dMdx_Fdp_T} \propto \frac{1}{p_T^0}\frac{10p_T/p_T^0}{(1+(p_T/p_T^0)^2)^6} \times \frac{d^2\sigma_{\text{NLO}}}{dMdx_F} \tag{4.12}$$

where $p_T^0 = 2.8$ GeV was fit to both the hydrogen and deuterium data, and the MRST 98 [MRST 98] partons were used in the NLO calculations of $d^2\sigma/dMdx_F$.

4.5.2 Physics Simulation and Detector Response

Once an event had been sampled from the above distributions, the kinematics of each muon were calculated and boosted into the lab frame, and the event vertex sampled from

$$P(z) = \frac{1}{\lambda}\frac{e^{-z/\lambda}}{1 - e^{-L/\lambda}}. \tag{4.13}$$

Inside the target, the very small amount of multiple scattering the muons were subject to was simulated. This multiple scattering was found to be negligible.

After the effects on the muons in the target had been simulated, the muon's positions were extrapolated to the point at which the SM0/SM12 field began. The same magnetic field maps which were used in the analysis as well as the overall field strengths as computed from the data were used here. The muons were propagated through the magnetic field one at a time in 5 cm steps. At each step in the field, the magnetic field kick was calculated, using

$$dp_x = -q\ dz\, B_y \tag{4.14}$$

$$dp_y = q\ dz\, B_x \tag{4.15}$$

which neglected only the small z component of the force[2].

[2]To be precise, we computed the changes in the x and y components of the muon's momentum,

In addition to the simulation of the magnetic fields, the effects of the beam dump and absorbing wall had to be included. Multiple scattering was simulated in each 5 cm step of the dump and absorbing wall, using the *gmols* routine from the standard GEANT [GEANT] simulation package. This was the only case where a GEANT routine was used in the Monte Carlo. Multiple scattering was also simulated in the materials found in the helium bags which filled SM0 and SM12 downstream of the absorber, the air and helium the muons had to pass through, the hodoscope and drift chamber planes, and the bulk material at station 4.

The energy lost by muons traversing the spectrometer was sampled from a set of fast lookup tables generated using Von Ginniken's TRAMU program [Gin 86]. Twelve tables for muons with incident momenta between 10 and 600 GeV were interpolated to obtain the energy loss for muons of arbitraty momenta. This number was then scaled from the length of material assumed in the table to the step size of 5 cm, and subtracted from the muon's momenta in each step. A new energy loss was sampled once the muon had traversed the length of material specified by the energy loss routines. This greatly simplified the energy-loss simulation in the dump, where the total length of material the muon would pass through was not known beforehand. The only other places where energy loss was simulated was in the teeth and walls of the SM12 magnet, and the hadronic and EM calorimeters

then applied momentum conservation to calculate the final z component of the momentum. It can be shown that this causes a systematic error in the momentum at each step which is negligible over the entire magnetic field.

at station 4. A precise simulation of the energy loss was unnecessary here, so a constant energy loss was subtracted in these materials.

Equation 4.6 indicates seperate corrections for the acceptance of the spectrometer and detector efficiencies. In reality, though, the Monte Carlo simulated detector inefficiencies based on the measurements described above. Thus, the acceptance calculations contained the efficiency corrections. Only the rate-dependent inefficiencies were corrected for seperately.

4.5.3 Acceptance Calculation

Once a sufficient number of Monte Carlo events had been recorded, they were subjected to the same analysis chain and cuts described in chapters 3 and 4. The only difference was that spill quality cuts were not applied to the Monte Carlo. Once we had filtered the Monte Carlo event sample, the acceptance in any given bin could be calculated as the ratio of the number of surviving Monte Carlo events in that bin divided by the total number of events generated in that bin

$$\alpha(M, x_F, p_T) = \frac{N^{\mathrm{acc}}_{\mathrm{MC}}(M, x_F, p_T)}{N^{\mathrm{gen}}_{\mathrm{MC}}(M, x_F, p_T)}. \tag{4.16}$$

Calculating the number of accepted and generated Monte Carlo events in each bin would have been a straightforward counting exercise, save for two reasons. First, comparisons of the data to the Monte Carlo revealed differences on the order of 10 to 15% which we wished to eliminate[3]. Second, in order to eliminate the

[3]These differences were consistent with the observed differences between the measured cross

systematic error between the hydrogen and deuterium cross sections introduced by finite Monte Carlo statistics in the acceptance calculation, the same sample of Monte Carlo events were used to calculate the acceptance for hydrogen and deuterium. Since the event sample had been generated using either a hydrogen or deuterium target, the events had to be reweighted by the cross section appropriate to the target being calculated.

Differences between the data and Monte Carlo were reduced by reweighting the Monte Carlo events as a function of M, x_F and p_T. These functions were determined by polynomial fits to the ratios of the data divided by the Monte Carlo. Furthermore, the p_T cross section was seen to vary as a function of M and x_F. This correlation was dealt with by fitting p_T^0 in each bin, resulting in the empirical form

$$p_T^0(M, x_F) = 1.21 + 0.350M - 0.0182M^2 + 1.37x_F + 2.60x_F^2 \tag{4.17}$$

which was the same for both targets, within the uncertainties in the fit parameters. The resulting weighting functions for the Monte Carlo were consistent with the differences observed between the measured cross sections and the double-differential cross section tables used to generate the Monte Carlo.

Monte Carlo events which had been generated using hydrogen as the target were reweighted with the deuterium cross section and vertex distribution when

sections, and the NLO calculations upon which the Monte Carlo generator was based.

calculating the deuterium cross section. Thus, the deuterium yield calculated from a hydrogen Monte Carlo event sample was given by

$$N_{\mathrm{D2MC}} = \sum_{i=1}^{N_{\mathrm{H2MC}}} \frac{d^3\sigma_{pd}(M^i, x_F^i, p_T^i)}{d^3\sigma_{pp}(M^i, x_F^i, p_T^i)} \times \left(\frac{e^{-z/\lambda_{\mathrm{D2}}}}{e^{-z/\lambda_{\mathrm{H2}}}} \right) \tag{4.18}$$

Hydrogen yields were calculated from deuterium Monte Carlo samples in a similar manner:

$$N_{\mathrm{H2MC}} = \sum_{i=1}^{N_{\mathrm{D2MC}}} \frac{d^3\sigma_{pp}(M^i, x_F^i, p_T^i)}{d^3\sigma_{pd}(M^i, x_F^i, p_T^i)} \times \left(\frac{e^{-z/\lambda_{\mathrm{H2}}}}{e^{-z/\lambda_{\mathrm{D2}}}} \right). \tag{4.19}$$

This procedure was carried out for both the accepted and generated events, maintaining the relative normalization of the thrown and accepted Monte Carlo, and equation 4.16 was used to compute the acceptance.

Figures 4.6 and 4.7 compare the data and Monte Carlo for several reconstructed pair- and detector-quantities, summed over both targets and averaged over all data sets. The agreement between the data and Monte Carlo is quite good for most physics and detector distributions. The discrepancies in the single track momentum distributions, which occur at larger momenta where the statistics are poor, are thought to be due to beam-alignment problems. The other major discrepancy is in the ZUNIN distributions. This is partly due to a slightly wider ZUNIN distribution in the Monte Carlo than is observed in the data, but mainly due to the centroids of the data and Monte Carlo distributions not being precisely aligned[4].

[4]This is a major component of the estimated systematic errors discussed below.

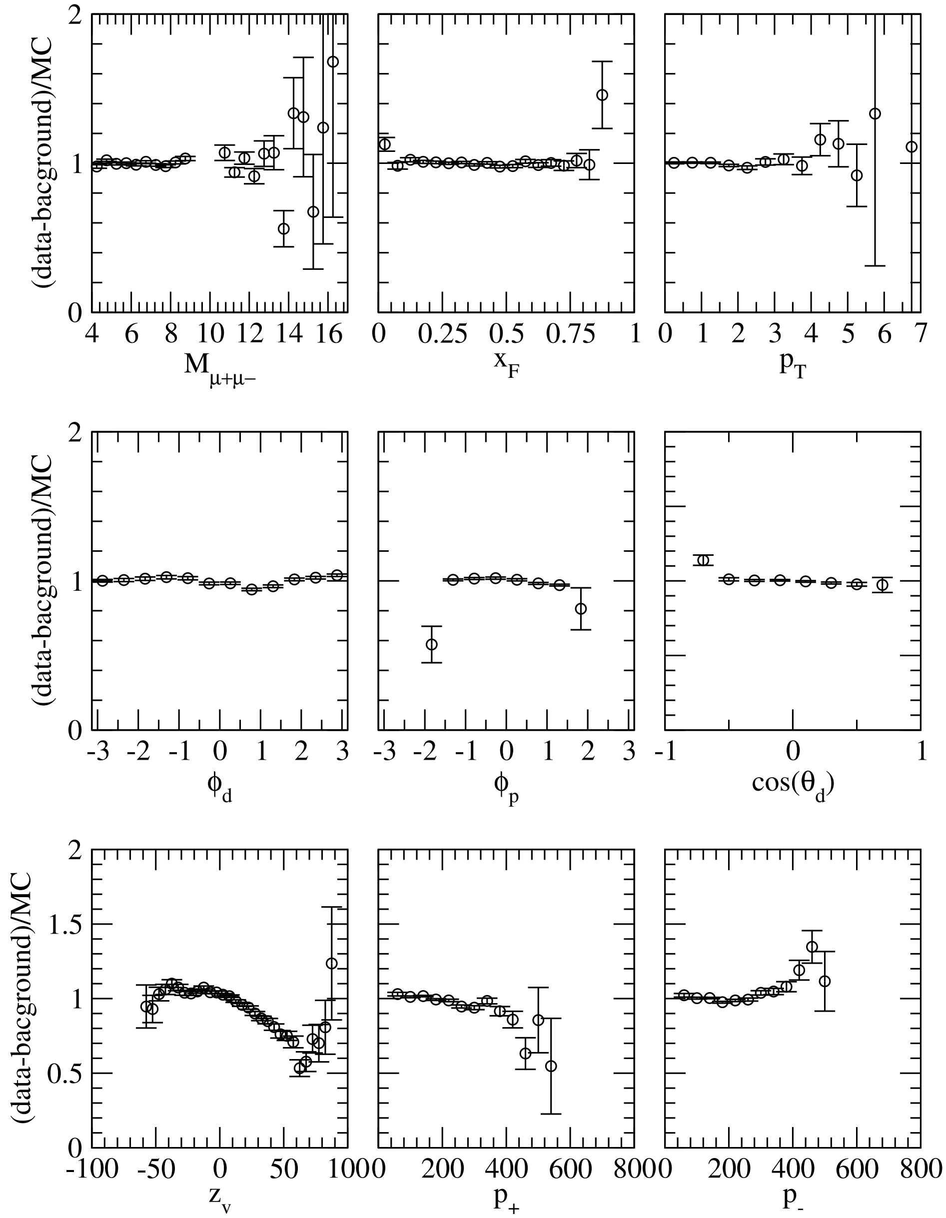

Figure 4.6: Ratios of data and Monte Carlo yields for several "pair" quantities (mass, x_F, p_T, ϕ_d, ϕ_p, $cos(\theta_d)$ and ZUNIN) and the positive and negative track momenta. Data averaged over all targets and data sets.

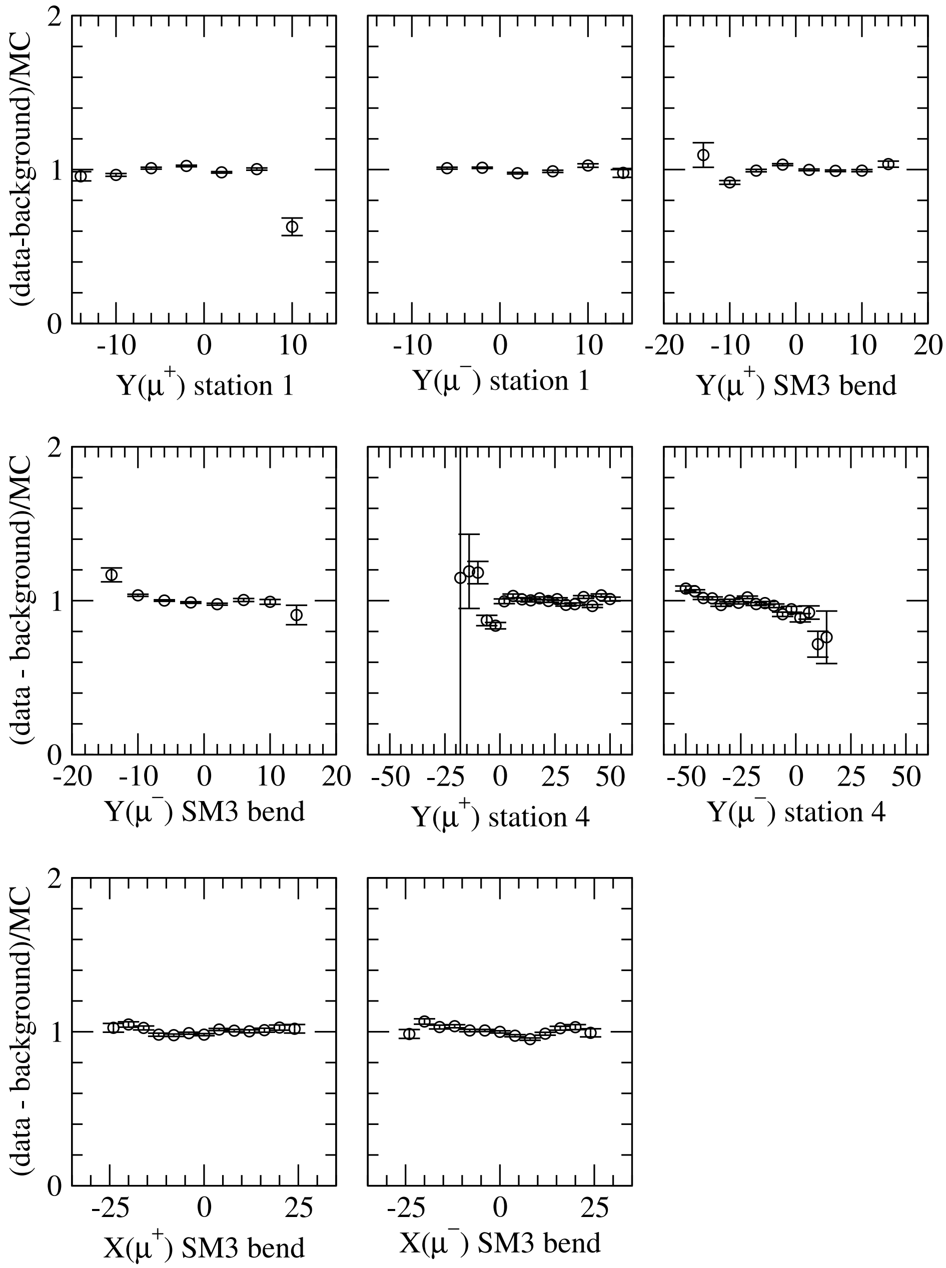

Figure 4.7: Ratios of data and Monte Carlo yields for the distributions of positive and negative tracks projected onto the station 1 hodoscope-, SM3 bend- and station 4 hodoscope-planes. Data averaged over all targets and data sets.

4.6 Combining the Data

Once all of the data were binned and subjected to all of the corrections described above, we had a measurement of the triple-differential cross sections $d^3\sigma(pp)/dMdx_Fdp_T$ and $d^3\sigma(pd)/dMdx_Fdp_T$ for each of the six data sets used in the analysis. For the purposes of the following discussion, we denote the measured cross section in a given M, x_F, p_T bin as $\sigma_i \pm \delta\sigma_i \pm \Delta\sigma_i$, where i represents the data set, $\delta\sigma_i$ is the statistical uncertainty and $\Delta\sigma_i$ is the systematic uncertainty.

The results from the six data sets were first combined into three results corresponding to each mass setting. We used the weighted average

$$\sigma \pm \delta\sigma = \frac{\sum_i \sigma_i \omega_i}{\sum_i \omega_i} \pm \sqrt{\frac{1}{\sum_i \omega_i}} \tag{4.20}$$

where the weights were given by $\omega_i = \frac{1}{(\delta\sigma_i)^2}$. After the average cross section for each mass setting was calculated, the three mass settings were combined into a single result, again using equation 4.20. This procedure neglected the systematic uncertainties between data sets and mass settings, which were sufficiently small so that the average was not affected much by the omission. Also, it would not have been appropriate to use the systematic uncertainties in computing the averages, as most of the systematic uncertainty in the cross sections were (by design) common to the hydrogen and deuterium measurements.

4.6.1 Systematic Uncertainties

Computing the systematic uncertainty in the combined cross section was a fairly straightforward exercise in error propagation. The standard error propagation formula for an expression $f(x_1, x_2, ..., x_N)$ is given by

$$(\Delta f)^2 = \sum_{i=1}^{N} \left(\frac{\delta f}{\delta x_i}\right)^2 (\Delta x_i)^2 . \tag{4.21}$$

From this formula, it can be shown that the systematic uncertainty in equation 4.20 is given by

$$(\Delta\sigma)^2 = \frac{\sum_i \omega_i^2 (\Delta\sigma_i)^2}{\left(\sum_i \omega_i\right)^2} \tag{4.22}$$

where the ω_i are the statistical weights used in the weighted average. One need only compute the systematic uncertainty in each data set (mass setting) $\Delta\sigma_i$.

Recalling equation 4.6, the systematic uncertainty on a given cross section measurement will, by equation 4.21, be given by

$$\left(\frac{\Delta\sigma_i}{\sigma_i}\right)^2 = \left(\frac{\Delta\alpha_i}{\alpha_i}\right)^2 + \left(\frac{\Delta\epsilon_i}{\epsilon_i}\right)^2 + \left(\frac{\Delta\mathcal{L}_i}{\mathcal{L}_i}\right)^2 + \left(\frac{\Delta N_{\mu^+\mu^-}}{N_{\mu^+\mu^-}}\right)^2 \tag{4.23}$$

We must be careful here, however. The uncertainty in the luminosity $\Delta\mathcal{L}$, for instance, is dominated by a $\pm 6.5\%$ uncertainty in the beam normalization. This uncertainty is common to all data sets and mass settings. It would not be appropriate to treat it as an independent 6.5% uncertainty in each measurement. We also assume that the uncertainties in detector efficiencies $\Delta\epsilon$ are common to each data set in a given mass setting, since events with similar kinematics will occupy similar regions of the spectrometer. We discuss how the systematic uncertainties were treated below.

4.6.2 Uncorrelated Systematic Uncertainties

The systematic uncertainties, not including the $\pm 6.5\%$ normalization uncertainty, fell into three categories: uncorrelated uncertainties, uncorrelated uncertainties which are common to a given mass setting, and correlated uncertainties which are common to a given mass setting. In general, we treated these systematic uncertainties separately. In the end they were combined in quadrature to estimate the total systematic uncertainty on the averaged cross sections.

The only uncorrelated systematic uncertainty which was not shared between the data sets in a given mass setting was the uncertainty in the acceptance due to the statistical uncertainty in the Monte Carlo. The uncertainty in each mass setting due to this statistical uncertainty was therefore estimated to be, according to the prescription of equation 4.22

$$(\Delta\sigma_{\alpha,stat})^2 = \frac{\sum_i \omega_i^2 \left(\Delta\sigma_i^{\alpha,\mathrm{stat}}\right)^2}{\left(\sum_i \omega_i\right)^2} \tag{4.24}$$

where $\Delta\sigma_i^{\alpha,\mathrm{stat}} = \frac{\Delta\alpha_i^{\mathrm{stat}}}{\alpha_i}\sigma_i$ is the uncertainty in the i^{th} measurement due to the statistical uncertainty in the acceptance $\Delta\alpha_i$.

The next step combined all of the uncorrelated uncertainties which were common to a given mass setting. These included the uncertainties in the acceptance in each mass setting, calculated above, and the uncertainties in the detector efficiencies. Again we use equation 4.22

$$(\Delta\sigma_{\text{uncorrelated}})^2 = \frac{\sum_i \omega_i^2 \left\{(\Delta\sigma_\alpha)^2 + (\Delta\epsilon_{\text{hodo}})^2 + (\Delta\epsilon_{\text{chamber}})^2 + (\Delta\epsilon_{\text{trig}})^2\right\}}{\left(\sum_i \omega_i\right)^2} \quad (4.25)$$

where the sum is now over all mass settings instead of all data sets in a given mass setting.

The remaining source of systematic uncertainty was the uncertainty in the acceptance due to uncertainties in the magnetic-field tweeks as measured in chapter 3. Unlike the previous sources of uncertainties we have discussed, we expect this uncertainty to be at least partially correlated from mass setting to mass setting – the SM3 tweek was wholly determined by measurements from the high mass data.

An estimation of the size of this uncertainty in any given mass setting was performed using a Monte Carlo based on the code described above, modified to propagate a given event through the spectrometer multiple times with different settings of the magnetic fields. In this modified Monte Carlo, all random processes, such as energy loss, multiple scattering and detector resolutions and efficiencies, were turned off. The result was a highly correlated sample of Monte Carlo events where any differences in reconstructed yields were due solely to acceptance effects.

4.6.3 Correlated Systematic Uncertainties

The systematic uncertainty in the acceptance due to the uncertainties in the magnetic field in any given mass setting could then be estimated by

$$\left(\frac{\Delta\alpha_{\text{tweek}}}{\alpha}\right) \approx 2\,|\frac{N_{\text{nominal}+\delta} - N_{\text{nominal}}}{N_{\text{nominal}+\delta} + N_{\text{nominal}}}|. \tag{4.26}$$

Here, N_{nominal} represented the number of Monte Carlo events subjected to the nominal magnetic field tweeks in traversing the spectrometer which survived the standard analysis chain. Similarly, $N_{\text{nominal}+\delta}$ were the number of Monte Carlo events subjected to a set of augmented magnetic fields in traversing the spectrometer, which were again subjected to the standard analysis chain. The field tweeks in this case were all augmented by the estimates of their uncertainties[5], which should result in the maximum possible uncertainty.

This results in a systematic uncertainty in the cross section for each mass setting given by

$$\Delta\sigma_i^{\alpha,\text{tweek}} = \left(\frac{\Delta\alpha_{\text{tweek}}}{\alpha}\right)\sigma_i. \tag{4.27}$$

To estimate the uncertainty in the total cross section, we depart from the prescription given by equation 4.22. We do this, because equation 4.22 is based on the assumption that the systematic uncertainties obey Gaussian statistics. This

[5] In the case of the low mass data, the SM12 and SM3 field tweeks were augmented by their systematic uncertainties while the SM0 field tweek was reduced by its uncertainty. This resulted in the maximum possible systematic uncertainty in the acceptance for that spectrometer configuration.

is a good assumption for many of the sources of uncorrelated systematic uncertainties. Here it is not. Therefore we treat the systematic uncertainty $\Delta\sigma_i^{\alpha,\mathrm{tweek}}$ as a potentially "missing" or "excess" cross section for each mass setting, for which we wish to obtain an average estimate for the total cross section. We take

$$\Delta\sigma_{\mathrm{correlated}} = \frac{\sum_i \omega_i \Delta\sigma_i^{\alpha,\mathrm{tweek}}}{\sum_i \omega_i} \tag{4.28}$$

where once again the ω_i are the statistical weights of the given mass setting. Finally, we add the correlated and uncorrelated systematic uncertainties in quadrature to obtain an estimate for the total systematic uncertainty

$$\Delta\sigma = \sqrt{\left(\Delta\sigma_{\mathrm{uncorrelated}}\right)^2 + \left(\Delta\sigma_{\mathrm{correlated}}\right)^2}. \tag{4.29}$$

4.6.4 Integrating the Data

After the data were combined using equation 4.20, we integrated the triple-differential cross sections over p_T to obtain the double-differential cross sections $d^2\sigma(pp)/dMdx_F$ and $d^2\sigma(pd)/dMdx_F$. The integrated cross sections were given by

$$M^3 \frac{d^2\sigma}{dMdx_F} = M^3 \sum_i^N \left(\frac{d^3\sigma}{dMdx_F dp_T}\right)_i \Delta p_{T,i} + f(M, x_F; \alpha) \tag{4.30}$$

where the sum is over all N p_T bins whose widths are denoted $\Delta p_{T,i}$ and $f(M, xF; \alpha)$ represents a correction explained below. The statistical and systematic uncertainties[6] on the integrated cross sections were calculated using the usual error prop-

[6] Correlated and uncorrelated systematic uncertainties were first treated separately, then combined in quadrature to obtain the total systematic uncertainty.

agation in equation 4.23 – in other words, the uncertainties (multiplied by Δp_T) were added in quadrature.

The correction in the integral $f(M, x_F; \alpha)$ accounted for the fact that the acceptance of the spectrometer could not reach the full kinematic limit in p_T. For each bin in mass and x_F, the triple-differential cross sections were integrated only over the range in p_T where the acceptance was nonzero. Where the acceptance was nonzero, and all data sets lacked events, we made an estimate of the missing cross section $f(M, x_F; \alpha)$. The estimate was based on the empirical p_T distribution given in equation 4.12 and the parameterization of p_T^0 given in equation 4.17, extrapolated from the largest nonzero p_T bin to infinity.

The resulting corrections were well below 1% over much of the kinematic range covered. Only at larger masses did the correction to the cross sections approach 5%. While the empirical form for the p_T dependence of the cross section does a fairly good job of reproducing the data, it tends to underestimate the tail of the distribution. Given this, and the uncertainties in the parameterization of p_T^0, we estimate that the correction $f(M, x_F; \alpha)$ is uncertain by about $\pm 15\%$. This uncertainty was added in quadrature to the systematic uncertainties in the final result.