Chapter 5

Distortions of the Spectrum

A detailed understanding of the spectrometer is crucial to correct distortions in the measured spectrum. At the energies studied, e^+ are vulnerable to significant scattering, annihilation and energy loss in the materials traversed, and inactive media degrade energy resolution. In addition, a class of problems arises because the number of e^+ produced in the low-energy regime is so small compared to the number at high energies. Thus, improbable processes by which high-energy particles can yield events at low-momentum tunes cause distortion. More quantitatively, if Comus is tuned to a central momentum of $P_e = 6.15$ MeV/c and there is a probability of 1.5×10^{-8} that an e^+ of higher energy will yield an event at this setting, the contamination is 1%, the same enhancement as if η had been increased by $+0.09$.

Thus, quite unlikely processes can cause substantial problems at low energies; for tunes below a few MeV/c, many effects are so large as to be unmanageable. For this reason, as well as the decreasing statistical sensitivity of this experiment to η at low energies, the lower limit to the energy range was about 5 MeV.

5.1 Target-area Effects

There are several distorting processes that occur in the μ^+-stop counter, $T1$, and the surrounding anti-counters. While it is not entirely possible to separate them in a calculation (one being forced to perform Monte Carlo simulations in order to handle the problem properly, since the target is not "thin" in the sense that all effects can be added linearly) it is useful to discuss the processes separately.

5.1.1 Multiple Scattering

In an ideal (but impractical) experiment, one would stop muons at the center of a tiny sphere of scintillator in the place of $T1$. In such a system there would be no angular anisotropy of the e^+ from unpolarized muons. However, the finite size of the muon beam spot requires that the $T1$ scintillator be fairly large in the directions transverse to the beam, while the need to minimize the amount of matter through which the decay e^+ must pass, or from which it might scatter, requires that it be thin. A moderately wide, flat target is then unavoidable.

When e^+ are emitted in directions nearly normal to the plane of $T1$, they experience relatively little scattering; for e^+ initially directed toward the spectrometer acceptance (in this case $\theta_T < 66°$, where θ_T is the angle of the e^+ trajectory from the normal to $T1$) the scattering is typically a few degrees or less and is not of major significance. However, those e^+ which are emitted with $\theta_T \approx 90°$ may pass through a very substantial thickness of scintillator before reaching an edge—up to 2.29 cm. For e^+ with energies of a few MeV, it is almost certain they will scatter out of one of the faces of $T1$ before they travel this far, while most 40 MeV e^+, for example, will not deviate far from their original direction. This effect is illustrated by Monte Carlo results in Figure 17, which show a substantial anisotropy for 6 MeV positrons. (The asymmetry with respect to $\theta_T = 90°$ derives from the average muon being deposited less than halfway through the scintillator.) Fortunately, the spectrometer does not accept particles at the angles most affected and the distortion is significant only for the lower-energy data points.

Much of this effect can be eliminated by rejecting events which have an unusually large energy deposition in $T1$. However, the effect of such a cut is complicated, and there is no alternative to Monte Carlo calculation to determine it.

Two scattering problems also occur in the light guide attached to $T1$:

- The light guide is twice the thickness of the $T1$ scintillator, so there is a small overhang which particles leaving $T1$ can strike. However, similar numbers of particles are scattered into, and out of, the spectrometer's angular acceptance, so the net effect is not large.
- The light guide, lying in the plane of $T1$, can scatter particles into the spectrometer acceptance in the same way as $T1$ itself tends to do. The problem is reduced by the anti-counters A1 and A2, which veto particles whose source is not very close to $T1$. The few remaining events cannot be rejected on the basis of their energy deposition (since they will seldom pass through enough Lucite to produce the required amount of Čerenkov light) and correction is made for them.

5.1.2 Continuous Energy Loss

While a distinction between Bhabha scattering and continuous energy loss for an e^+ is artificial, for calculational purposes one must choose some energy threshold above which to track scattered e^-, and below which to treat the energy loss as continuous. It is also helpful conceptually. Fortunately, the shape of the spectrum in the measured region is insensitive to the value of this threshold, as will be discussed in Section 5.1.3.

There are two parts of the spectrum which are most affected, relatively, by the continuous energy loss:

- The endpoint region: the effect here can be nearly 100%, since the whole spectrum shifts to lower energy and the measured intensities at those energies separated from the endpoint by less than the average energy loss are drastically reduced.

- The low-energy region of the spectrum: this can be understood by considering the Michel spectrum without radiative corrections, $F(x) \equiv d\Gamma^0(x)/dx = -2x^2(3-2x)$. In this approximation,

  $$\frac{1}{F(x)}\frac{dF(x)}{dx} = \frac{6(1-x)}{x(3-2x)} .$$

  It is clear that a fixed energy loss causes a fractional change in the spectrum that is divergent at $x = 0$. While the effect is modified by both radiative corrections and an accurate energy-loss treatment, this illustrates the basic trend; the continuous energy loss for an e^+ is, in fact, nearly constant over the energy range of interest. The results of a more careful analytic calculation are shown in Figure 18.

5.1.3 Bhabha Scattering

The Comus spectrometer did not distinguish between e^+ and e^-, which merely spiral in opposite directions around the spectrometer axis. In principle, energy measurement in the detector could distinguish between them, but our energy resolution was not sufficient for event-by-event determination. Thus, an e^- which was Bhabha scattered out of $T1$ was easily mistaken for an e^+ of equal momentum coming directly from μ^+ decay. The probability for Bhabha scattering in a thickness dx is given by[1]

$$d^2\phi_{Bhabha} = \frac{2\pi r_e^2 n_e}{\beta_0^2}\frac{dT_1}{T_1^2}\left\{1-\beta_0^2\left[f_1(y)\epsilon - f_2(y)\epsilon^2 + f_3(y)\epsilon^3 - f_4(y)\epsilon^4\right]\right\} dx ,$$

where we have defined

$$f_1(y) = 2 - y^2 ,$$
$$f_2(y) = 3 - 6y + y^2 - 2y^3 ,$$
$$f_3(y) = 2 - 10y + 16y^2 - 8y^3 ,$$
$$f_4(y) = 1 - 6y + 12y^2 - 8y^3 ,$$
$$y = \tfrac{1}{T_0+2} \text{ and } \epsilon = \tfrac{T_1}{T_0}.$$

T_0 is the kinetic energy of the initial e^+, T_1 is the kinetic energy of the scattered e^-, r_e is the classical electron radius and n_e is the electron density of the material. Both T_0 and T_1

---

[1] H. J. Bhabha, Proc. R. Soc. London, Sect. A **154**, 195 (1936).

are given in units of m_e.

The first point to notice is that $d^2\phi_{Bhabha}$ is divergent as $T_1 \rightarrow 0$. Thus, one might *a priori* expect a very large amount of e^- contamination in low-energy measurements. (Actually, the divergence is prevented by atomic effects, but this has little practical importance here.) The contamination is drastically reduced in the Comus spectrometer, however, by the target-area veto counters which will usually be struck by the coincident e^+ and veto the event. Because Bhabha scattering (ignoring soft photons) is a two-body process, the scattering angle of the e^- is fixed and given by

$$\cos\theta_{scat} = \sqrt{\frac{1+\frac{2}{T_0}}{1+\frac{2}{T_1}}} .$$

A little trigonometry gives the more directly useful quantity, θ_{sep}, the angle separating the final e^+ and e^- directions:

$$\cos\theta_{sep} = \frac{1}{\sqrt{\left(1+\frac{2}{T_1}\right)\left(1+\frac{2}{T_0-T_1}\right)}} .$$

For scattered e^- with energies of a few MeV, θ_{scat} tends to be fairly large. For example, an e^- of 6.15 MeV/c momentum will be separated by at least 24.4° from the correlated e^+. This makes a veto in the $A1$, $A2$ counters likely, and the effect is reduced by an order of magnitude, to a manageable level.

Another concern is that the e^+ might be accepted, with the e^- causing vetoes in an energy-dependent way. Such energy dependence would exist if $d^2\phi_{Bhabha}$ significantly varied with T_0 in the region where the cross section is large—for small ϵ. It would also exist if θ_{sep} were a strong function of T_0. Accurate calculation would then be difficult, complicated by multiple scattering and energy loss in $T1$, and the absolute calibration and position-dependent light collection efficiency for the veto counters. Fortunately, both $d^2\phi_{Bhabha}$ and θ_{sep} are insensitive to T_0 when ϵ is small. Another way in which energy-dependent veto probability might arise is through deflection of the e^- by the spectrometer's fringe field. This problem is circumvented by the choice of cuts on $A1$ and $A2$: they must be set high enough to avoid vetoes from the extremely low-momenta e^- subject to deflection at the higher spectrometer settings.

The primary effect of the very low-energy e^- is then to reduce the number of good events uniformly. Even so, one might worry that too many events will be lost in this manner, given the nearly divergent cross section for low-energy Bhabha scatters; this does not occur because the *very* low-energy e^- are stopped in $T1$ and there is a finite threshold for event rejection in the veto counters. An estimate of the overall Bhabha scattering effect is shown in Figure 19.

5.1.4 External Bremsstrahlung

Bremsstrahlung, like most other energy-loss processes, has a nearly divergent cross section for producing low-energy secondaries—photons, in this case. The divergence is $1/k$, where k is the photon energy. However, the situation is simpler than for Bhabha scattering; there is little chance that the secondary photons will either cause or veto events. The $T1$, Tb, $A1$, $A2$, $A3$, A, B, C and D scintillators are all fairly thin and carbon ($Z = 6$) is their heaviest component element, so interactions in them are unlikely for photons of high enough energy to cause vetoes.

With regard to vetoing an event, another important thing to consider is the angular separation between the secondary e^+ and the photon. Naturally, being a three-body process, there is a continuous distribution of angles for bremsstrahlung, but the average rms angle between the photon and secondary electron for $Z = 6$ is roughly $0.6\,\frac{\ln\gamma_0}{\gamma_0}$ radians.[2] This rms angle is only a weak function of the energy division between secondaries, with the coefficient varying from about 0.50 to 0.75 as the photon's share of energy varies from 100% to 0%. As an example, if a 40 MeV e^+ emits a photon and is reduced in energy to 5 MeV, the angle separating the secondaries will typically be 1.8 degrees. When the e^+ is accepted, it is unlikely the photon will strike $A1$ or $A2$, although it may hit A, B, C or D. About 0.05% of all events, otherwise good, are vetoed in this way.

---

[2] M. Stearns, Phys.Rev 76, p.836 (1949).

Still, bremsstrahlung has an important effect on the muon decay spectrum. It is possible for an e^+ to transfer most of its energy to the photon, so that a high-energy e^+ may spuriously appear in the low-energy part of the spectrum. Figure 20 shows the approximate effect on the spectrum shape.

One other comment on bremsstrahlung in the $T1$ counter: carbon and hydrogen are of low enough Z that bremsstrahlung in the field of electrons cannot be ignored in comparison to nuclear bremsstrahlung. Further, the energy dependence of this process is somewhat different from nuclear bremsstrahlung. This will be discussed further in Section C.5.

5.1.5 In-Target Annihilation

The e^+, as they pass through $T1$, risk annihilation with the e^- in the scintillator. The energy dependence of this process can be seen from the following approximate formula[3] for the probability of an e^+ of energy γ_e (in units of m_e) annihilating in a thickness dx of material with electron density n_e:

$$d\phi_{annih} = 1.60\pi r_e^2 n_e \gamma_e^{-\frac{7}{9}} dx \ .$$

The effect on the spectrum is fairly small, around 0.1%, for the energies used in this experiment. The explanation is that the target is thin and the process is quite unlikely to produce spurious, low-energy events from high-energy decays.

5.1.6 Internal Bremsstrahlung

Internal bremsstrahlung, i.e. radiative decay, results in very large corrections to the spectrum, but the radiative corrections already discussed include them as a subset, and they are known analytically to adequate accuracy. However, the implicit assumption was made in the previous discussion of radiative corrections that the real photons were not detected, an assumption which is not entirely valid in this experiment: the photons may cause vetoes in

---

[3] H. Messel and D. F. Crawford, *Electron-Photon Shower Distribution Function* (Pergamon, Oxford, 1970), p. 17.

the A, B, C, D, $A1$, $A2$ or Tb anti-counters. Section 5.1.4 stated that vetoes from external bremsstrahlung photons were unlikely, but that photon flux was an order of magnitude less than from internal bremsstrahlung in this experiment; the effect cannot be dismissed without study.

To be more quantitative, we use the formula of Fronsdal and Überall,[4] in which $y = E_\gamma / E_e(max)$ and y_0 is the minimum value of y considered, and create the following table of branching ratios.

x	Γ_γ/Γ			
	$y_0 = 10^{-3}$	$y_0 = 10^{-2}$	$y_0 = 10^{-1}$	$y_0 = 0.5$
0.105	29.0%	26.2%	22.2%	11.3%
0.205	17.0%	13.5%	9.4%	3.4%
0.305	14.3%	10.4%	6.1%	1.4%
0.405	13.3%	9.1%	4.6%	0.5%
0.505	12.8%	8.3%	3.7%	0.0%
0.605	12.3%	7.7%	3.0%	0.0%
0.705	11.9%	7.1%	2.2%	0.0%
0.805	11.2%	6.3%	1.4%	0.0%
0.905	9.9%	4.8%	0.0%	0.0%

Thus, an average γ-detection efficiency of a few percent could result in substantial spectrum shape changes. For the geometry and data cuts used in this experiment, 0.5% of events are lost from this source, and the shape distortion between $x = 0.10$ and $x = 0.95$ is no more than 0.11%. Looser cuts on the veto counters would reduce the loss of events, but *increase* the distortion. The near-cancelation in the energy dependence of several effects in this experiment (the cross section, the photon-positron angular separation and photon detection efficiency) is fortuitous; a different experiment could have larger (or smaller) distortion.

---

[4] C. Fronsdal and H. Überall, Phys. Rev. **113**, 654 (1959).

For finite cuts, however, the spectrum near the endpoint will show more distortion. As the phase space for photons of adequate energy to generate a veto vanishes, the event rate depression also vanishes. Thus, there is a 0.5% correction in the spectrum region used to confirm the spectrometer line shape.

5.2 Momentum-Dependent Detection Efficiency

The particle-detection efficiency in our spectrometer shows a certain amount of momentum dependence, especially at tunes of a few MeV/c. The physical processes responsible will be discussed in the subsections below. Together, these processes make the detector efficiency complicated enough that there is no alternative to a Monte Carlo calculation, which is done with the EGS code. Results are shown in Figure 21 for the positron-detection efficiency. In order to eliminate energy-dependent effects discussed elsewhere, these events were required to not impact any material between $T1$ and the downstream edge of the $C2$ collimator. The figure does include effects due to the last set of collimator-suspension cables and the vacuum window, as well as the detection counters themselves.

5.2.1 Scattering Effects

The optics of Comus are designed so that almost all unscattered tracks will pass within the bounds of the $T2$ scintillator. However, particles must cross the vacuum window and 0.7% of them strike the cables centering the $K4$ collimator. The vacuum window is 0.025 cm thick aluminum, while there are four stainless steel cables: one of 0.16 cm diameter and three of 0.12 cm diameter. Some particles are significantly scattered, especially those which would have passed near the outer radius of $T2$, and either miss $T2$ or only graze a corner. This problem is linked to the 3.5 cm gap which separates the vacuum window from $T2$ and the even larger gap between the cables and $T2$, making scattering angles of only a few degrees significant. The rms scattering angle for relativistic particles varies as P_e^{-1}.

Another scattering effect which reduces detection efficiency is backscattering from the counter wrappings and scintillator. Although this material is thin and low Z, this probability is not negligible for low-energy particles. The $T2 \cdot T3$ requirement will usually reject the event if backscattering occurs upstream of some plane in $T3$, the position of the plane being determined by the cut on $T3$. The total observed-energy cut may reject events in which the backscatter occurs even later; this is especially true at higher P_{tune} settings, where this cut essentially requires energy deposit in the NaI.

5.2.2 Annihilation

Positrons annihilate in the collimator-suspension cables near the detector, the vacuum window, $T2$, $T3$ and their wrapping material; once annihilation γ's are formed, charged particles do not usually reappear until the NaI(Tl) crystal is reached. Thus, $T2$ and/or $T3$ may have little, or no, energy deposited in them, and detector inefficiency results. As mentioned in Section 5.1.5, annihilation is energy-dependent. The energy loss by a low-energy e^+ in the early part of the detector can be a significant fraction of its energy, and the energy-loss tail increases the annihilation probability, moving particles closer to the divergence at $E_e = 0$. There is the also the energy-dependent effect of backward-traveling secondaries from the shower in the NaI that may trigger $T2$ or $T3$, despite an annihilation.

In-flight annihilations can also occur after $T3$, of course. In the NaI crystal itself, they have little effect; in its window, the effect is to reduce losses in the remainder of the window, increasing the energy deposit in the NaI.

5.2.3 Unintended Vetoes

Since $C1$, $C2$ are relatively close to the detector, there is a small probability that particles in the detection shower will enter the spectrometer and strike them. Very low-energy $e^{\pm}$ are reflected by the magnetic field, of course, but high-energy ones and γ's are not. The TDC's on these counters cannot separate this class of vetoes: the extra delay is not enough to be clearly resolved (the counters are about a meter long, being wrapped on the collimators) and it is necessary to cut so tightly on the ADC for these counters that the TDC has often not been triggered anyway. Typically, these vetoes reduced the detection efficiency by 0.5%.

5.3 Effect of Suspension Cables

Due to the mass of the collimators suspended within Comus (about 255 kilograms) and the need for stability when the spectrometer is moved, several sets of cables are necessary for support and centering. The set furthest downstream has been mentioned for its contribution to detector inefficiency, but the upstream cables are qualitatively different in their effect, since they are struck by a less-filtered distribution of particles. The $e^{\pm}$ leaving $T1$ risk encountering these cables and interacting, so that an event can be either added or lost.

There are two sets of stainless steel cables which support the $K2$ collimator; each set has four cables, 0.23 cm in diameter. Two of these cables are secured to $K2$ by small brackets which are exposed to the incident particles. The weight of the $K4$ collimator is offset by a third set of cables, two pairs, 0.12 cm in diameter.

These cables are rather opaque, in the sense that events accepted without the cables are seldom accepted when they are hit. At 6.15 MeV/c only 0.11% of the impacting particles are still accepted; at 50 MeV/c the fraction accepted rises to 6%. With the three cable sets blocking 2.5% of the spectrometer aperture, the net distortion of the spectrum is 0.15%.

The larger effect of these cables is to introduce spurious events, either by elastic scattering or by bremsstrahlung and Bhabha scattering. The effect of elastic scatters persists throughout most of the spectrum at a non-negligible level, because of the small angular deflections needed to change the fate of a particle (of order 2 degrees). The bremsstrahlung and Bhabha scattering, on the other hand, affect mostly the low-energy portion of the spectrum. These effects are modified by the presence of veto counters, analogous to the situation with $T1$. A calculation of the net effect of the cable sets as a function of energy is shown in Figure 22.

Included in this figure are "second-order" events in which particles, after scattering from the cables, shower in the spectrometer back plate or collimators to produce spurious events. These are not totally negligible because scattering allows high-energy particles to reach locations at which they constitute a significant fraction of the incident particles. Explicitly: at low P_{tune}, the magnetic field does not deflect high-energy particles enough to hit the vulnerable parts of the spectrometer back plate; scattering does. The fractional size of the effect increases at low energies; at $P_{tune} = 6.15$ MeV/c the contribution is 0.17%.

5.4 Effect of the $K1$ Collimator

It is possible for $e^{\pm}$ tracks originating on the back or inner bevel of the $K1$ collimator to be accepted. Thus, in general, one would expect that showering in the $K1$ collimator, or scattering from its inner bevel, would lead to spurious events. Muon decays in the beam collimator, etc. undoubtedly lead to event triggers in this way, but these do not have an in-time signal in $T1$ and are background that can be subtracted in time-of-flight spectra. Further, $A1$ and $A2$ veto nearly all charged particles which might strike $K1$ from the target.

The remaining possibility is that the γ ray from a radiative μ^+ decay in $T1$ will cause showering in $K1$, while the e^+ avoids hitting the various veto counters. Because of the strong angular correlation between the γ and e^+ in radiative decay, this is unlikely: only about 14% of such events escape vetoes from $A1$ or $A2$. This residual fraction is not completely negligible, however, and is not reduced much by energy-deposition cuts, since the particles accepted from this source are not much lower in momentum than the desired ones. The contamination is 0.21% at $P_{tune} = 6.15$ MeV/c, but falls rapidly at higher settings (because of the increasing rate of normal events, more than any reduction in the rate of events from $K1$ scatters).

5.5 Effect of the $K2$ Collimator

Showers directly through this collimator produce only a small number of events. The thickness of $K2$, combined with the shielding effect of $K4$ and the magnetic field, was adequate to reduce the probability of a spurious event from this source to less than 0.07% at $x = 0.1$. The contamination drops rapidly at higher tunes, since the absolute trigger rate is nearly constant and total-energy cuts in that case are effective in rejecting these low-energy particles.

A more significant source of events is showering by high-energy particles on the edges of $K2$. These events are suppressed by the fact that $K2$ is not a boundary of the acceptance for particles originating at $T1$, but there is a region of phase space in which particles leaving $K2$ can be accepted. On the bevel, this is centered at $0.89 \cdot P_{tune}$ and an angle 16 degrees away from typical incident trajectories. The effect on the spectrum is of order 1%, including a small second-order effect in which events are generated by rescattering from the back plate; this is shown in Figure 23.

5.6 Effect of the K_3 Collimator

Particles which strike the beveled inner surface of this ring at grazing incidence have a significant chance of scattering into the acceptance. Because small scattering angles (of order 2 degrees) are involved, the particles need penetrate only slightly into the lead, losing a negligible amount of energy. As a result, the probability of these grazing-incidence particles being accepted is relatively energy-independent, as is shown in Figure 24. The exception is in the endpoint region.

On the other hand, particles which penetrate far into the beveled surface and suffer hard, inelastic collisions are usually deflected by a large angle and have little probability of being accepted. This is especially true because of the relatively strong magnetic field which exists in this region: it is not possible for $e^{\pm}$ with less than about half of the tuned momentum to reach $T2$. These arguments also apply to particles hitting the upstream face of $K3$. Despite their large number, they have no significant effect at low energies and only a small effect at high energies.

5.7 Effect of the $K4$ Collimator

The $K4$ collimator is somewhat vulnerable to penetration. This is not because its 5.1 cm of lead is unreasonably thin, but because of the large number of particles incident upon it, compared to those in the spectrometer acceptance—a ratio of around 26. This effect is a major source of spurious triggers at spectrometer tunes of 30 MeV/c and greater.

Many of these events can be discarded: they mostly occur at high energies where the detector energy resolution is good, and these spurious triggers tend to involve mostly small energy deposits. Although complete separation is not possible, the effect on most of the spectrum is small when tight energy-deposition cuts are made. The most noticeable spectrum distortion appears in the line shape measurement at the endpoint: the good event rate falls to zero, while these events (which tend to come from particles below the tune momentum) fall off more slowly.

5.8 Effect of the $K5$ Collimator

The particles incident on $K5$ have momenta similar to the normal events, and most have to penetrate several centimeters of lead to reach the detector. Thus, spurious events at high P_{tune} are eliminated by the energy-deposition cut, while they very seldom occur at low P_{tune}. There is also no large enhancement factor from the spectrum weight of high-energy particles. This collimator does what it should do; the corrections for it are negligible.

5.9 Effect of the $C1$, $C2$ Veto Counters

Due to the proximity of $C1$ and $C2$ to the detector, shower particles from it can hit these counters and veto good events. This reduction of detection efficiency was discussed in Section 5.2.3.

For $C1$, particles striking the beveled surface do so at grazing incidence, so that $e^{\pm}$ can scatter into the acceptance before depositing enough energy to allow rejection. This is more likely for low-momentum particles, which tend to scatter at larger angles for a given distance traversed. The fraction of accepted events from this source has been calculated by Monte Carlo to vary from 1.2% at $P_{tune} = 6.15$ MeV/c down to 0.14% at $P_{tune} = 50$ MeV/c.

Also, $e^{\pm}$ from $T1$ strike the upstream face of $C1$; a small region adjacent to the beveled surface is essentially "dead", in that the particle trajectories pass through so little scintillator that vetoes are unlikely. Thus, the spectrometer line shape broadens slightly toward low momenta. In addition, the effect has a P_{tune} dependence, since low momentum particles are more likely to be scattered out of the acceptance by this edge. By calculation, these events make up about 0.28% of accepted events at high momenta, falling to 0.12% at $P_{tune} = 6.15$ MeV/c—a net distortion of 0.16% in this range. The combined effect of scattering from the $C1$ bevel and face is shown in Figure 25.

In addition, the geometry of $C1$ is non-ideal, meaning azimuthally asymmetric. The scintillator leaves a gap as it rises away from the $K4$ collimator, then overlaps itself for a short distance, crosses the momentum-selection slit and finally joins the light guide. The low-momentum edge of the line shape is affected by both the gap and by the double thickness of $C1$: the former adds a long tail, the latter reduces the FWHM. The crossing of the slit between $C1$ and $C2$ just reduces the total acceptance.

Similar phenomena occur for $C2$. The $e^{\pm}$ striking the $C2$ bevel do so at fairly steep angles and have little chance to scatter out without enough energy deposit to cause a veto. However, $e^{\pm}$ that hit the most downstream part of the bevel will exit from the downstream face, sometimes traversing so little scintillator that no veto results. These tend to hit or not hit $T2$, as they would have done anyway. About 0.45% of the events passing all cuts have gone through this corner, but the principal effect is a slight broadening of the spectrometer line shape toward high momenta. Aside from this, the net spectrum shape distortion is around 0.14%—acceptance being reduced by this much at $P_{tune} = 6.15$ MeV/c, compared to high-momentum tunes, due to the increased scattering. The calculated effect is shown in Figure 26.

The geometry of this counter is also non-ideal. The long strip of scintillator is wrapped circularly, leaving a small gap on the inside radius just before it overlaps itself. This introduces a long tail onto the high-momentum edge of the spectrometer line shape. It also substantially increases the number of particles which trigger $T3$ without hitting $T2$.

5.10 Effect of the Comus Back Plate

The back plate of the Comus vacuum vessel is a 3.8 cm thick plate of aluminum, which corresponds to 10.3 g/cm^2 or 0.43 radiation lengths. Obviously this is not adequate to absorb the showering completely from moderate or high-energy particles that might strike it, although it attenuates the energy of the shower and increases its angular spread. Shielding was never, however, its main purpose.

There are several spectrometer characteristics that make this discussion relevant. First, the magnetic field in Comus deflects particles far above the tune momentum by a few degrees, as one would expect. However, because the field increases in strength with increasing radius, there is a tendency for high-energy particles to be deflected onto paths which are roughly parallel to the spectrometer axis, and at a fairly large radius from it. Many of these paths thread between the thick lead shielding of the $K5$ collimator and the A,B,C,D anticounters, striking the relatively unprotected area of the back plate left for the A,B,C,D counter light guides to exit the spectrometer. Finally, the material separating this portion of the backplate from $T2$ is relatively thin aluminum and provides only minimal additional shielding.

There are, fortunately, other circumstances which mitigate the problem. Many $e^{\pm}$ are vetoed as intended by A,B,C,D; others are vetoed by Čerenkov emission in the light guides of these counters. A few particles will be attenuated by the magnetic shields and support members between the back plate and $T2$. Finally, the distance between the under-shielded part of the back plate and $T2$ is large enough that the probability of false events is reduced due to a solid angle factor. These reductions are very significant.

Nonetheless, a problem remains due to the large number of particles which hit this area and their relatively high energy. The time-of-flight of these particles hardly differs from that of normal events, so TDC cuts do not help. The total-energy cut is extremely helpful at high P_{tune} settings, since the detector energy resolution is good there and events from the back plate tend to be of low energy. In the absence of this cut, the contamination reaches a maximum of 4% near $P_{tune} = 30$ MeV; the cut reduces it to 0.8%. With all cuts applied, the maximum calculated contamination is 1.2% at $P_{tune} = 17$ MeV/c. At low energies, the Monte Carlo calculation needs to be supplemented by the empirical $\overline{T2} \cdot T3 \cdot T4$ data. Figure 27 shows the effect on the spectrum as obtained from the combination of the two.

Chapter 6

Data Analysis and Results

6.1 Monte Carlo Calculations

Many separate Monte Carlos, as opposed to one large one, were used to study various aspects of the spectrometer. This simplified the use of variance-reduction schemes and reduced the overall required CPU time to a practical, though still high, level. When the resulting inaccuracy was negligible, products of two or more small effects from separate Monte Carlos were often ignored.

In general, the philosophy of the Monte Carlos was to record experimentally observable quantities (i.e., energy depositions in the counters) event-by-event, along with quantities useful for variance reduction and comprehension. Events were stored in a compact form and, in most cases, could be kept accessible in computer disk files. Counters in which energy was deposited during an event were indicated in a mask word, and the quantities deposited in those counters were stored as 16-bit words.

Experimental considerations, such as counter calibrations, resolutions and cuts were then imposed on these files. This allowed one to vary cuts, study correlated quantities and adjust detector parameters in the analysis with relative ease. Except as noted, the basic interaction physics in each Monte Carlo was that provided by the EGS4 code system, with modifications as discussed in Appendix C.

In order to include magnetic deflection of the charged particles, field maps were calculated by the POISSON[1] code in the approximation of azimuthal symmetry. Measurements at a field 10% above the endpoint found the four flux-return yokes allowed a deviation from symmetry at the ±0.25% level, which, given the 2.8% line width and 90° rotation of transmitted tracks, does not materially affect the results. Then, exploiting the cylindrical symmetry of the problem, an analytic function was fitted to the field on the spectrometer axis; such a function, together with an adequate number of its derivatives, can reproduce the field away from the axis. Particle trajectories were found according to the prescription of Lindgren and Schneider,[2] in which the equations of motion, containing a power series expansion of the field at each point, are numerically integrated. Accepted tracks calculated by this approach typically deviated by only 0.01 cm from those found by the more usual Runge-Kutta integration method, and were obtained in a small fraction of the time. Also, the data could be held in a few parameters, as opposed to a full field map.

6.1.1 Incident Beam Studies

Incident muons pass through a thin beam line window and the Tb counter before stopping in $T1$, whose normal was turned 42° from the beam. Large angle scatters in the window or Tb prevent some particles from striking $T1$—and bias the stopping depths in $T1$ toward slightly larger values, since these particles would usually stop more shallowly than average in a semi-infinite region. Also, the quantity of interest, the muon deposition depth below the $T1$ surface, is strongly coupled with scattering angles and transverse transport by the $T1$ rotation—to a much greater degree than in the usual situation of normal incidence.

For these reasons, the usual values of range and, especially, straggling were modified. A heavily rewritten and extended version of the TRIM85 code (discussed in Appendix D) was used to determine the applicable parameters. The results were checked against, and found to be consistent with, measured range curves obtained by placing varying amounts

---

[1] M. T. Menzel and H. K. Stokes, Report LA-UR-87-115 (1987).

[2] I. Lindgren and W. Schneider, Nucl. Instrum. Methods **22**, 48 (1963).

of degrader between *Tb* and *T*1.

This code was also used to study the lead collimator which was just upstream of the *Tb* counter and which restricted the beam spot width and movement on *T*1. The aperture of this collimator was tapered to allow tighter collimation of the beam than normally possible. Despite the grazing-incidence particles on this bevel, it was found that only a small number of muons were scattered from the collimator onto *T*1 or its light guide; these reduced the average deposition depth by 0.06%. The EGS code was used to check for similar problems for the positrons in the beam, finding none. The explanation is that particles losing a substantial fraction of their energy in the collimator are almost certainly scattered so much that they miss *T*1, while the others present no direct difficulty.

6.1.2 Target Effects Study

The *T*1, *Tb*, *A*1, *A*2, *A*3, *A*, *B*, *C* and *D* counters were included in the target-area Monte Carlo. These counters were studied as a group because of the possibility of vetoes from correlated particles, although the *A*, *B*, *C* and *D* veto counters were also present in other studies. The Monte Carlo used the PRESTA macros as an important addition to EGS4, because of scattering in *T*1.

The number of events required to study the target adequately is very large, partly because of the small, but significant, probability of a high-energy decay positron yielding a low-energy event. Preliminary studies were done to find a distribution of starting events in phase space that substantially reduced the number required (by several orders of magnitude over nature's crude, but effective, approach). The necessary number of high-energy events is finally set by the statistics in the tails of the bremsstrahlung and Bhabha-scattering distributions, while multiple scattering sets the number at low energies. The statistical fluctuations in these distributions could have been lowered only with great difficulty, completely rewriting EGS. For the approach used, 54 million events were started.

In part, the actual variance reduction was accomplished by distributing the event weight very regularly in phase space, eliminating most of the statistical variation due to starting conditions. The separation between occupied points in phase space was small enough that the spectrometer's angular acceptance and momentum bite, along with physical effects, smoothed the discrete distribution adequately. (Multiple scattering, of course, helps to smooth the discrete distribution in angle, but not by as much as one might naively expect. For example, if particles are started at polar angles spaced by 1°, and scatter by a degree or so on the average, the resulting distribution will still be peaked at the starting angles, if they are more than a few degrees from the axis: the problem is two dimensional and many of the scattering angles add in quadrature, rather than linearly, to the starting angles.)

Events are efficiently used when more are started in phase space regions with an acceptance probability near 50%, and fewer are started in regions where acceptance is either very likely or unlikely. A weighting factor for each event allowed the retention of the correct probability density. In general, events were started so as to minimize the standard deviation of the resulting distribution for the number of events used. However, very large weighting factors were avoided so that, in the real world of limited statistics, the tails of the variance distribution were not too large: the probability distribution for the accepted event weight does not necessarily resemble a Gaussian and grave errors can be made by treating it as such. (As an example, if there is a probability of 0.001 that a particle causing a 1% effect will be accepted, the average effect is 0.001% and the standard deviation is 0.03%. If one draws only a single particle from this distribution, the 99.9% chance that the result will be low by only 0.001% does not compensate well for the occasional 1% spikes, and the probability of the 1% spikes is obviously not at the 30σ level of a Gaussian distribution.) Regions of phase space found to have utterly negligible effects were left unpopulated, although these regions were not very large at the lowest energies studied.

The magnetic field was ignored at this stage. Except for extremely low-momentum particles (less than a few hundred KeV/c), particles upstream of $A1$ move through the vacuum in straight lines. The extremely low-momentum particles are not important because they are far below the range over which the spectrometer was tuned for spectrum measurements, and the cuts on the target-area counters were set high enough that these particles do not cause vetoes. It is therefore necessary to run only a single set of target-area Monte Carlos, as opposed to doing one set for each of many P_{tune} settings.

The distribution of particle starting positions within the target was that predicted by the incident beam Monte Carlos, including the effects of collimation and straggling. Events yielding one or more particles downstream of $A1$ with a polar angle in the accepted range were written to an intermediate file on magnetic tape. These trajectories were traced to their intersection with a plane transverse to the spectrometer axis at the center of $T1$. The event record then consisted of the direction cosines, intersection point and particle energy, as well as the energy deposited in the counters. Any accompanying charged particles downstream of $A1$ were recorded at the $A1$ position, while photons directed at A, B, C or D were allowed to strike them, with any resulting energy deposition recorded.

These files were then inspected by a second program containing maps of two phase-space regions: the one for which a charged particle would strike the A, B, C or D anti-counters, and the spectrometer acceptance. Events in which a charged particle hit A, B, C or D were assumed to be vetoed efficiently (in order to reach these anti-counters, the particles necessarily have high momentum). The program also applied the detection efficiency, as determined by a separate set of Monte Carlos, for the $T2 \cdot T3$ and $\overline{T2} \cdot T3 \cdot T4$ event classes. To use the event statistics most efficiently, the efficiencies were applied as factors to the event weights; the probability that each counter ADC would pass experimental cuts, when the resolution was folded with the energy deposition, was determined; the product gave the overall probability. (The alternative, inefficient approach would be to select from the ADC distributions randomly and compare with the cuts for these specific values.) Four histograms in P_{tune}, separate for the two triggers and for e^+ and e^-, were then incremented with the appropriate weight for each field setting at which the event could be detected.

In this approach, the histograms have been automatically smoothed by the spectrometer line width. However, because this is very narrow at low momenta, and because single events with large weight are found here (from Bhabha scatters or bremsstrahlung), further smoothing was needed. A third-level program thus combined histograms from separate Monte Carlo runs, divided these by a function representing the approximate result, smoothed them and then multiplied by the same function. Different functions were used for electron and positron histograms, because of their very different shapes. This avoids most of the error from blindly smoothing a nonlinear function. The statistical variation in the final histogram is better than 0.2%, and was calculated for decay energies down to 6 MeV.

There is one further complication: the spectrum shape is the thing which we are trying to measure; it is not available as pure input to a Monte Carlo. The solution is to parameterize the spectrum as

$$\frac{d\Gamma(x,\eta)}{dx} = \frac{d\Gamma_1(x)}{dx} + \eta\frac{d\Gamma_2(x)}{dx} ,$$

which is possible using Eq. 2.4. Grotch has already set $\rho = 3/4$, so it does not appear as a second parameter. (Deviation of ρ from 3/4 is studied later in an approximate way, neglecting products of this deviation with the radiative corrections and experimental spectrum distortion.) It is not necessary to run separate Monte Carlos to calculate the experimental modifications to Γ_1 and Γ_2; one simply applies different initial event weights to the same events.

6.1.3 Internal Bremsstrahlung Studies

The principal target-area Monte Carlo involved only the decay positrons, not the internal bremsstrahlung photons. This was natural because Grotch has integrated over the photon phase space, assuming that they are not detected, and an acceptable approximation because most were, in fact, not detected. Their effects were studied separately in two specialized programs.

The first program found the probability that correlated photons would veto events and the variation of this with positron energy. Rather than using EGS to make a detailed simulation of each event, the photon-interaction physics were applied in a much faster, simpler way. Tables of the probability that a photon of given energy would cause a veto for a given path length in a veto counter were calculated for specific cuts; these tables were interpolated to find the probability that the photons accompanying an accepted positron would veto the event. Fortunately, the veto probability was found to be small and insensitive to the anti-counter cuts in the spectrum region used to determine η, as discussed in Section 5.1.6. The correction was finally applied to the calculated spectra.

The radiative-decay photons also produce contaminating $e^{\pm}$ on the $K1$ collimator and the aluminum vacuum chamber walls in its vicinity. This is studied with EGS4, finding the probability that photons will produce accepted $e^{\pm}$, without the correlated positron vetoing the event.

6.1.4 Collimator Studies

The collimator surfaces downstream of $A1$ were all studied with the same basic program—but separately, allowing events to be distributed in an efficient manner and making the relevant physics more transparent. It also made interpolation much simpler, as will be discussed below. The surfaces studied were the

- the upstream face and bevel of $K2$,
- the upstream face and outer radius of $K3$,
- the beveled edge of $K3$,
- the upstream face of $C1$,
- the beveled surface of $C1$,
- the upstream face of $C2$,
- the beveled surface of $C2$,
- the upstream face of $K4$,
- the upstream face of $K5$, and

- the spectrometer back plate.

The required CPU time was reduced by factoring effects into several parts. Typically, three functions were found—

- The physical processes giving events from particles incident on a surface usually vary slowly with energy, but are expensive to simulate. (Typically one second of VAX-780 CPU time is needed for each event, this slowness being due to the extensive showers which particles can produce, as well as the cost of transporting charged particles in a magnetic field.) These Monte Carlos were therefore only run at a moderate number of specific spectrometer tunes, spaced in roughly 1 MeV/c steps at low momenta, up to 10 MeV/c steps at high momenta. Events were run and saved, the experimental counter resolutions and cuts being applied later. After the cuts were applied, giving the probability that a particle incident on a surface would yield an event for a given trigger, the Monte Carlo results were (usually) smoothed by hand. These energy-dependent probabilities were then interpolated by three-point, Lagrange interpolation.

- The number of particles incident on a surface can be rapidly varying with P_{tune}, but is relatively cheap to calculate when only direct tracks from $T1$ are considered. Normalized to the number of accepted particles when the central incident momentum is tuned, this value can be cheaply calculated at closely-spaced points as needed; it changes rapidly when the spectrum endpoint excludes impacts as P_{tune} increases, or when the acceptance at the normalization tune approaches the endpoint.

- Finally, the number of particles entering the acceptance at the normalization tune was available from the target-area Monte Carlo. By using these results, some higher-order corrections were included in, at least, an approximate way.

The approximations involved in this sort of separation are usually small; the principal errors occur when particles from the endpoint are incident on a surface, so that the particle *distribution*, as well as number, may vary rapidly. In some cases, Monte Carlos were closely spaced in P_{tune} to reduce this error.

The input file of trajectories encountering each surface of interest was created by tracking particles from the predicted beam spot distribution on $T1$. They were uniformly started within some region of phase space whose boundaries had been determined to contain all trajectories which could reach the surface from $T1$. The intersection with, and momentum vector at, the surface of interest was recorded for those events which did not first strike another surface. The weights and momenta for these events were later scaled at each chosen value of P_{tune}.

6.1.5 Upstream Support-Cable Studies

The three upstream sets of support cables were also Monte Carloed at only a limited number of spectrometer tunes, for the same reasons as above. To a reasonable approximation, the three cable sets lie on conic surfaces. Simulated particle trajectories were traced through the spectrometer, and their intersections with these surfaces were recorded. For each cone, a file was made of the trajectories that would be successful in the absence of the cable set; another was made of those that would fail after crossing the surface.

Next, an azimuthal angle was chosen from the range in which the trajectory could intersect the given cable set, the event was weighted appropriately and the particle was tracked back to its original impact on the cable. The motivation for this approach was to reduce the size of the input file and the cost of producing it. Particles were transported through the cable (approximated as a cylinder of iron with the measured global density) by EGS4; particles exiting the cable were tracked through the magnetic field until they again encountered material. If at least one particle from the shower reached the vacuum window at the spectrometer exit, the starting conditions of the event, including the random number seeds, were recorded. These were only a miniscule fraction of the events started.

The events could then be reproduced, and the same code used to Monte Carlo the spectrometer downstream of $A1$ was used to simulate the remainder of the event. Energy depositions in counters and other event-specific information were recorded for later use, as before.

6.1.6 Detection Efficiency Studies

Detection efficiency was studied in a way similar to that used for the collimator impacts; the same basic code was also used. Files of positrons striking the vacuum window and last set of collimator-support cables were created at several spectrometer tunes by tracking particles, saved by the target-area Monte Carlo, through the spectrometer. Separate efficiencies were found for e^+ and e^-, as these differ considerably, but the number of e^- events produced by Monte Carlo was too small; thus, the e^+ input file was used again with opposite charge. The approximation is insignificant.

The result variance was reduced by using the known, average fraction (0.71%) of events hitting the last set of collimator-support cables; this was useful because these events represent a major fraction of the inefficiency at some spectrometer tunes. Energy depositions in the counters were recorded as in the previously-discussed studies, again allowing later application of counter resolutions and cuts.

It was necessary that these studies include the azimuthal irregularities in the $C1$ and $C2$ counters, as these affect the particle distribution on the trigger counters, and, hence, the detection efficiency. Runs were spaced rather closely for tunes in the endpoint region, as the particle distribution changes rapidly.

6.2 Applying Cuts

In general, the measured spectrum shape will depend on the counter pulse-height cuts chosen. In this experiment, the sensitivity is large for some counters: the dead areas of $C1$ and $C2$, for example, are proportional to the applied cuts; and the loss of detection efficiency at a given energy due to annihilation is sensitive to the cut on $T3$. Uncorrected, both of these effects roughly mimic a shift in η. Therefore, it is necessary to apply the experimental counter resolutions and cuts carefully to the Monte Carlo results. The required information comes from several sources, depending on the counter.

Most TDC spectra were corrected for walk, due to pulse-height variation, using a simple, two-parameter, exponential function of the ADC—with excellent results. $A1$ and $A2$ are more complicated because the ADC and TDC are also correlated due to the attenuation of light as it reflects (and is delayed) around the circumference of the counter rings. The sum of two exponentials, with four parameters, substantially improves the quality of the correction although this form has no particularly strong basis. $C1$ and $C2$ have a similar problem in principle, but the light attenuation is much less and the trigger rate in these counters is so low that there is no motivation to make small improvements in the TDC resolution. To improve resolution further and simplify their TDC spectra, the high-rate counters, $T1$, $A1$ and $A2$, are calibrated in picoseconds and referenced to the calibrated TDC of a low-rate trigger counter, $T2$ or $T4$.

*T*1, *Tb*

The $T1$ ADC energy calibration was determined by passing beam positrons. The resolution for energy deposition is dominated by photoelectron statistics; the calibration of photoelectrons per deposited MeV was the value for which photoelectron statistics, combined with the Monte Carlo energy-deposition distributions of beam and decay positrons, reproduced the measured distributions.

The ADC was corrected for rate, by comparing the ADC distributions at the same P_{tune} for different beam rates, and for magnetic field effects when $P_{tune} > 37$ MeV/c. It was also corrected in a time-dependent way for the tail and reflections of the muon pulse, which had effects up to 2 μs afterwards, by using the nanosecond clock reading.

The ADC cuts on this counter were moderately tight, rejecting about 3% of the events, especially from the high energy-loss tail. This was possible because positron energy loss is nearly constant over the momentum range studied, and desirable because most events scattering out of the $T1$ plane were thereby rejected, along with some Bhabha scatters and decays of muons that stopped in the edges of $T1$. Even tighter cuts would have rejected more undesirable events, but would also have increased the vulnerability to inaccuracies in the ADC corrections and calibrations.

The lower ADC cut was made at the point where noise and good events made roughly equal contributions to $P_{tune} = 6.15$ MeV/c runs. Thus, the accepted event sample contained a certain amount of contamination, which was handled by using TDC information. A simple cut was not possible because of the high rates in $T1$; the TDC was not infrequently stopped before the desired pulse by beam positrons, muon pulse reflections, noise, other decay positrons, etc. Its statistical use is, however, extremely important and will be discussed in Section 6.3.

The Tb counter was calibrated by oscilloscope measurements of passing beam positron pulse heights. It was part of the hardware trigger, and no off-line cuts were applied, but the calibration for the discriminator level was needed in the target-area Monte Carlo.

$A1$,$A2$,$A3$

The $A1$,$A2$ counters were calibrated partly by the band of ~80-KeV lead X rays and partly by passing $e^{\pm}$. The X rays, produced by particle impacts on $K1$ and $K2$, were visible in high-P_{tune} measurements—due to the high statistics and relatively few hits in the target-area counters in these runs. The $A3$ calibration relied upon the X rays, since $A3$ is struck by very few high-energy $e^{\pm}$ in triggered events. Position dependence of the light-collection efficiency in all three was determined with passing electrons from a ^{106}Ru source.

The ADC cuts in $A1$, $A2$ and $A3$ were chosen to be above the level of the lead X rays to simplify the calculation of the veto effects; Monte Carlos need not then propagate particles, including photons from internal bremsstrahlung, into the lead collimators to find the resulting fluorescence photons. Further, vetoes from the very low-energy $e^{\pm}$ which can be deflected by the P_{tune}-dependent magnetic fringe field, are avoided when an energy deposition of at least 90 KeV is required. The cut is still low enough to reject efficiently any energetic $e^{\pm}$ striking $A1$, $A2$, $A3$.

The TDC spectra of $A1$, $A2$ and $A3$ are also affected by low-energy $e^{\pm}$. These particles can pass through the apertures of $A1$ and $A2$ initially, only to be reflected from the spectrometer's magnetic field and strike an anti-counter on the return pass. Such an effect is obviously a function of P_{tune}. Delayed vetoes can also result from particle impacts on the collimators. One might, in principle, prefer to reject events with a delayed veto, since many are caused by Bhabha scatters from the target. However, accurate calculation is difficult, and they cannot be separated from uncorrelated impacts on these counters.

Correction was made for $A1$ or $A2$ vetoes, uncorrelated with an event, in a statistical way, subtracting the TDC background in them for each run and adding vetoes that would have occurred if an early firing had not already stopped the TDC. (A small increase in the uncertainty of the data results, of course, and is included in the statistical error.) Correlations between $A1$ and $A2$ were necessarily included in this TDC correction. Also, because of the existence of delayed, correlated vetoes, only early vetoes were used to determine the background level.

A,B,C,D

The A,B,C,D counters were calibrated using positrons from the target with the spectrometer magnet turned off, since, under usual conditions, charged particles pass through these counters at a wide variety of angles and do not give a clear peak. This was checked against the 341 KeV Compton edge (from 511 KeV annihilation gamma rays) which is visible in the ADC spectra for high-P_{tune} measurements. The ADC cuts were made as close to the pedestals as possible—below this Compton edge and far below the level for passing charged particles. This helps to dispose of events in which a collimator has been struck, as well as those with direct hits. While small energy depositions (mostly from annihilation quanta) veto about 0.12% of the good events, the fraction depends only weakly on energy.

C_1, C_2

The $C1$,$C2$ counters were calibrated by the positrons that traversed them at high-P_{tune} settings. These results were fitted to Monte Carlo calculations incorporating resolution measurements done with a ^{106}Ru source.

The $C1$,$C2$ ADC's were cut as close to the pedestals as possible, eliminating most particles which grazed the bevel, or struck an edge, of a counter. This reduced the correction for, and therefore the uncertainty of, rejecting flawed events. Also, a small upper limit on the path length through scintillator reduces the probable scattering angles and the resulting energy-dependence of detection efficiency.

About 0.5% of events deposit energy in $C1$, $C2$ from showers in the detector, and it would be desirable to retain these events, identifying them by their TDC signature. The expected delay is 4 ns in $C1$ and 3 ns in $C2$. Unfortunately, they cannot be cleanly separated from other events because the optical length of these counters—bent circularly, but observed from one end—is nearly 6 ns. There is a further complication from the range of pulse heights which are important. One must simply correct for the loss of these events; the fraction is only weakly dependent on P_{tune}.

T_2, T_3, T_4, NaI

The $T2$, $T3$ and $T4$ counters were calibrated by fitting the measured ADC distributions to Monte Carlo results. The light-collection efficiency varies only slightly over the surfaces of these counters, as determined with a ^{106}Ru source; photoelectron statistics dominate the resolution. There are small gain losses at high spectrometer settings; the correction at the endpoint is 0.7% for $T2$ and 2.5% for $T3$.

The ADC cuts on these counters were made low enough to avoid rejecting any significant number of events from the main distributions although, especially for $T3$, some annihilation events were inevitably lost and rely on the calculated detection efficiency for compensation. There are also indirect ADC cuts on these counters via the total-energy deposition cuts.

The NaI calibration was somewhat complicated. A variable attenuator before the ADC was set to five different values over the range of spectrum measurements, to reduce the dynamic range problem. The attenuator settings were cross-calibrated with 1% corrections, typically, to the nominal values. This done, the energy loss in the NaI window was checked, since this was not known exactly: the manufacturer, Harshaw Chemical Company, only specifies (and knows) the unpacked geometrical thickness of the window components, which is vague for a material like MgO powder. (Naturally, when used as a γ-ray detector, this is not very important.) Thus, the most probable energy losses were compared to Monte Carlo prediction for several low-momentum tunes, and a thickness for the MgO powder was assigned to bring these into agreement. Finally, the magnetic field degrades the phototube gain for $P_{tune} > 40$ MeV/c; shielding of these tubes is difficult because of the large NaI diameter, and the gain loss reaches 7% at the endpoint.

With calibrations for all counters in the detector, an estimate of the total particle energy was possible. As discussed earlier, this estimate is not highly precise because of backscattering, inactive materials, photoelectron statistics, etc. There is also some effect from the finite range of the ADC's: the Monte Carlos must include the ADC upper limit for $T2$, $T3$ and $T4$, since some backscattered events lead to overflows.

The upper limit on the total energy was chosen to reject very few good events ($\leq 0.05\%$, based on the Monte Carlo results), but is useful because many background events were at much higher energies, especially at low spectrometer tunes. The lower cut was more aggressive, cutting about 0.5% of the good events from each run, since there is no clear separation between contaminating and good events. It succeeds in rejecting as many as half of the undesirable events. This level of cut rejects events fairly far out in the tail of the distribution for good events, and is not very sensitive to small calibration errors. An even tighter cut was not made because the energy-deposition distribution for normal events is beginning to rise rapidly by this point, so that additional gains would have had a much smaller rejection ratio. More importantly, a cut on the total energy can very easily bias the spectrum, through calibration errors or physics approximations.

6.3 Background Subtraction

The $T1$ TDC spectrum is made up of several components: normal "good" events, which occur in a narrow distribution; background events for which the same signal caused both the trigger and TDC stop, and which occur in an almost flat distribution over the combined gate widths of the $T1$ and $T2$ counters; and both good and background events in which the TDC was stopped early by an uncorrelated signal. The uncorrelated TDC stops decrease proportionally to the probability of a TDC stop at a given time, falling to zero by the end of the $T1 \cdot T2$ overlap range.

The lower and upper bounds on this overlap range will be called t_1,t_2. The background is estimated by first finding the maximum likelihood values of the parameters b and u in the function

$$Q(t) = b + u \int_{t_1}^{t} S(t')\, dt' \ ,$$

where $S(t')$ is the measured TDC spectrum at time t'. The function is fitted to the TDC spectrum for $t \epsilon (t_1, t_2)$, exclusive of the spike corresponding to normal events. $Q(t)$ represents events in which the $T1$ TDC is not stopped by the departure of a positron that hits $T2$. The background in the overlap region is then

$$B(t) = b + u \int_{t_1}^{t} B(t')\, dt' \ ,$$

with $B(t)$ being found from a numerical fit to the data, using the previously determined values for b and u. The fraction of good events in the overlap region is

$$f_{good} = \frac{\int_{t_1}^{t_2} [S(t) - B(t)]\, dt}{\int_{t_1}^{t_2} S(t)\, dt} \ ,$$

which is also the fraction of good events for $t < t_1$, since these early TDC stops are followed, in reality, by the overlap region which triggered the event.

This argument is, however, somewhat flawed: events with early TDC stops do not necessarily have the identical characteristics as those without them, i.e. it may not be purely statistical which events have early TDC stops. Several $T1_e$ pulses are often triggered by an arriving muon, and events which closely follow the muon arrival are both more likely to have early TDC stops and to have been triggered, in part, by a $T1_e$ pulse not related to the $T2$ pulse. Because of this correlation, events for each run were separated into three classes. The first class was composed of events in which only one muon had stopped in $T1$ in the 4-8 μs preceding the event, and in which the event occurred at least 1 μs after the muon arrival; the other two classes were composed of events which were more likely to have involved extraneous $T1_e$ pulses, the one class more so than the other. The background subtraction was done separately for each class, giving somewhat different results from the naive approach. This was especially true for the $\overline{T2} \cdot T3 \cdot T4$ events, which contain a larger fraction of background. The background subtraction for these events used the same approach as for $T2 \cdot T3$, although the discussion above only refers to the latter.

6.4 Magnetic-Field Corrections

The magnetic field shape in the spectrometer changes somewhat as the spectrum endpoint is approached, due to saturation effects in the iron. There is some change in the volume through which particles travel, affecting the tuned momentum by 0.05% at the endpoint. Although not of much significance, this was corrected in the analysis.

However, the field shape is more affected at larger distances from the magnet center, and the largest effect is the shift of the field at the NMR and Hall probes, relative to the fields deflecting accepted particles. This affects the measurement of the tuned momentum by 0.30%. Small shifts also occur in the probe cross calibrations, since they are not in exactly equivalent positions in the magnet. Corrections are made with a combination of measurements and calculation.

6.5 Scaler Corrections

The $T1_\mu$ and Tb signals create hardware event vetoes, hence deadtime. Also, run-to-run calibration is based on knowing the number of muons stopped in $T1$, given by the $T1_\mu \cdot \overline{Tb_\mu}$ scaler. The involved scalers must be known accurately, and corrections are necessary because conditions were not identical for all runs. When needed, rates were extracted through scalers dedicated to counting oscillator pulses.

While deadtime due to event readout was incorporated into the gates for most scalers, other corrections had to be applied in the analysis. An imperfection in the data-acquisition system enabled the scalers slightly before events could occur that would be recorded (for most later runs, data acquisition was manually enabled to eliminate this problem); the computer system also occasionally lost data buffers (depending on event rate and, hence, P_{tune}). Because the number of triggered events was recorded as a scaler, and all scalers were recorded with each data buffer, the scalers could be trivially adjusted to correct for these errors.

Pulse pile-up effects were not negligible in some scalers and needed careful correction, especially $T1_\mu$ (because of its relatively long veto into the trigger of 154 ns) and Tb (because of its high rate). The long $T1_\mu$ veto also caused the loss of 6.8% of potential events from each muon due to the time correlation between muon arrivals and events, though this is not of great import in a spectrum shape measurement, except in understanding the spectrometer acceptance.

Each event was assigned a weight of

$$W_{event} = e^{R_{T1_\mu}(\Delta_{T1_\mu}+\delta_1)} e^{R_{Tb}(\Delta_{Tb}+\delta_1)} e^{\Delta_{T1_\mu}/\tau_\mu} ,$$

where R_{T1_μ} and R_{Tb} are the average, deadtime-corrected $T1_\mu$ and Tb rates for the buffer holding the event; Δ_{T1_μ} and Δ_{Tb} are the logic pulse widths; δ_1 is the scaler deadtime after each pulse; τ_μ is the muon lifetime. The deadtime correction was done by solving a trivial transcendental equation; as an example, for a measured $T1_\mu$ rate of R''_{T1_μ},

$$R''_{T1_\mu} = R_{T1_\mu} e^{-R_{T1_\mu}(\Delta_{T1_\mu}+\delta_1)} .$$

Newton's method converges very quickly to the solution for R_{T1_μ}. The uncertainty in the correction as a whole introduces about 0.08% error in any given run.

Another correction must be applied to the $T1_\mu \cdot Tb_\mu$ scaler: the coincidence was between the 154 ns $T1_\mu$ and the 20 ns Tb_μ pulses, leaving a window between the end of Tb_μ and the end of $T1_\mu$ during which a muon could strike Tb, miss $T1$ and still be counted on the $T1_\mu \cdot Tb_\mu$ scaler. For a typical run, where $R_{T1_\mu} = 50$ KHz and 39% of the muons that hit Tb went on to miss $T1$, the correction is 0.65%. The correction formula is

$$\frac{R_{T1_\mu \cdot Tb_\mu}}{R'_{T1_\mu \cdot Tb_\mu}} = 1 - \frac{R_{Tb_\mu}}{R'_{T1_\mu \cdot Tb_\mu}} + e^{\left(R'_{T1_\mu \cdot Tb_\mu} - R_{Tb_\mu}\right)\Delta_{Tb_\mu}} \left[1 - e^{-R_{T1_\mu}\left(\Delta_{T1_\mu} - \Delta_{Tb_\mu} - \delta_2 - \delta_3\right)}\right] ,$$

where $R_{T1_\mu \cdot Tb_\mu}$ is the real, and $R'_{T1_\mu \cdot Tb_\mu}$ is the measured rate of $T1_\mu \cdot Tb_\mu$ coincidences corrected only for deadtime; R_{Tb_μ} is the muon rate in Tb (including those missing $T1$); Δ_{Tb_μ} is the width of the Tb_μ pulse; δ_2 is the minimum time, between the end of one pulse and the beginning of the next, for the coincidence module to produce two distinct output pulses; and δ_3 is the pulse overlap requirement of the coincidence module. As before, the rates are the average for the buffer holding the event, corrected for pulse pile-up. A complication is that Tb_μ was not recorded on a regular basis. Some information is available—Tb was recorded, but the positron:muon ratio varies enough to leave a 0.19% uncertainty in the number of stopped muons in a run, relative to other runs. This does not cause a major problem, because the statistical uncertainty in the number of events in a given run is much larger than this.

6.6 $\overline{T2} \cdot T3 \cdot T4$ Events

A challenging indicator of the accuracy of the spectrum distortion modeling for Comus is the event class ratio $\overline{T2} \cdot T3 \cdot T4 / T2 \cdot T3$ (where both classes are understood to have had all other trigger criteria, cuts and time-uncorrelated background subtractions applied as well). In an ideal spectrometer, almost no $\overline{T2} \cdot T3 \cdot T4$ events would have occurred; impacts on $T4$ are very unlikely unless particles scatter, annihilate or pass through "imperfections" in the $C1$, $C2$ anti-counters. These events thus provide an independent check of spectrum contamination calculations—but are in no way a measurement of the contamination, since the number of events in the $T2 \cdot T3$ and $\overline{T2} \cdot T3 \cdot T4$ classes have no simple relationship.

The calculated ratio, together with experimental data, is plotted against x_p in Figure 28, where $x_p \equiv P_e/P_e(max)$. No free parameters are used in this figure. The error bars exhibited for the experimental data are statistical only. The relative flatness of the plot for $x_p \epsilon (0.2, 0.6)$ is not an intrinsic characteristic, but results from the particular cuts chosen. For example, removing the total detected-energy cuts would result in a variation in $\overline{T2} \cdot T3 \cdot T4/T2 \cdot T3$ of several percent in this range.

In general, the prediction and data appear to be in good agreement, although a few points near the endpoint do not coincide. Regardless, the value of η is obtained from the lower parts of the spectrum, and the endpoint calibration is not sensitive to small inaccuracies. The discrepancy at low momenta is of concern, however. The source cannot be conclusively established, but it probably derives from scattering of $e^{\pm}$ on the spectrometer backplate and vacuum can. Veto inefficiency in the A, B, C, D anti-counters would permit this. While a high efficiency is expected, the 100% assumed in the Monte Carlo was certainly overly optimistic. Calculation of the inefficiency is not possible because many of the vetoes arise from Čerenkov light in the twisting, adiabatic light guides. The geometry, both for the incident particles which deflect in the magnetic field to intersect these guides and for the directional light passing through the guides, is complex.

Thus, one must subtract additional background from the data, based on the discrepancy and the calculated ratio of $T2 \cdot T3/\overline{T2} \cdot T3 \cdot T4$ events from scattering in this region. The background is reduced at higher momentum tunes because fewer e^{+} reach the vulnerable part of the vacuum vessel and total energy cuts reject more background; the fractional effect is even more strongly reduced. The fitted value of η is changed by -0.082, a 1σ shift. The ratio $\overline{T2} \cdot T3 \cdot T4/T2 \cdot T3$, with the additional background added to the prediction, is plotted against x_p in Figure 29. This single-parameter fit, using the calculated energy dependence of the contaminating events, has a χ^2/(degree of freedom) of 7.2/7 for the lower eight points in the graph. (It is also encouraging to note that the inclusion of this correction improves the χ^2/(degree of freedom) for the fit of the data to the $T2 \cdot T3$ spectrum from 11.9/11 to 10.0/11).

6.7 Extracted Value of η

Having finally obtained a theoretical prediction for $d\Gamma(x_p, \eta)/dx_p$ including experimental effects, this was fitted to the data with three parameters: the spectrometer momentum calibration, the acceptance calibration and η. The fit was done iteratively in two parts: the minimum χ^2 fit for acceptance calibration and η were found (with an approximate momentum calibration) for $x_p \epsilon (0.117, 0.756)$, followed by a fit of the momentum calibration to the endpoint data. The theoretical spectrum depends weakly on the acceptance calibration and the fitted values were thus used to correct it. This procedure was repeated a few times until the parameters converged, giving a χ^2/(degree of freedom) of 10.0/11 and the value

$$\eta = -0.080 \pm 0.088 \text{ (for } \rho \text{ constrained to } 3/4).$$

The experimental data, with statistical error bars only, are plotted on the Monte Carlo-calculated theory curve in Figure 30, with $\eta = -0.080$ and ρ constrained to 3/4. Nearby data points have been combined in the figure for clarity. Calculations did not extend below 6 MeV, and data at lower momenta were not included in the fit. The Monte Carlo curve is not completely smooth; this stems partly from statistical fluctuations and partly from real spectrum distortions. Figure 31 shows the data divided by the theory curve, along with a dashed line showing where data for $\eta = 0$ would be centered.

Unless the statistics in the spectrum mid-range are extremely high, compared to those in the lower part of the spectrum, a significant correlation between the acceptance calibration and η will appear in the fit; for these data, the correlation coefficient is -0.884. The uncertainty in the fitted acceptance normalization is ±0.42%.

It was not possible to determine the acceptance more accurately than this from first principles; the actual acceptance was about 4% lower than calculated. (Neither, of course, was it ever intended to use the calculated acceptance in fitting the spectrum. The required accuracy of geometrical measurements would be around 0.01 mm, and several other parameters would need to be known to high accuracy.)

The statistical error on η would decrease with the inclusion of higher momentum points because these help to provide amplitude normalization. On the other hand, the systematic errors increase near the endpoint due to several rapidly varying effects there, and the sensitivity to the ρ parameter becomes excessive. Thus, data above 40 MeV were excluded from the fit.

Instead of fixing ρ to $3/4$, the measurement can be specified using the published experimental value of $\rho = 0.7518 \pm 0.0026$:

$$\eta = -0.058 \pm 0.088 \ \ (\text{for } \rho \text{ constrained to } 0.7518).$$

For $\Delta\rho = 0.0026$ and the data in the range $x_p \epsilon (0.117, 0.756)$, the corresponding uncertainty in η is 0.032.

6.8 Systematic Error Estimate

Systematic error estimates for η, determined using data in the range $x_p \epsilon (0.117, 0.756)$, are given in Table 6.1. Correlated effects have been directly added in the table entries, while those which should be largely uncorrelated have been added in quadrature.

The column labeled "Shift in η_{fit}" is the shift an *uncorrected* systematic effect would cause in the determined value for η. The errors assigned to "theory" in the next column are due to neglected higher-order QED terms, uncertainty in atomic effects, etc. in the cross sections for particle interactions. The errors given are those that remain after approximate corrections were applied, if known. Inaccuracies in large-angle scattering, and hence in the detection efficiency, were approximately corrected in the results, rather than inside the EGS4 code. This was also done for Bhabha scattering. Bremsstrahlung in the electron field in $T1$ was largely corrected by modifying EGS4, but an additional correction was applied to the results. Simple, after-the-fact corrections were not possible in thick regions.

The errors assigned to "EGS4" are the uncompensated errors which exist in the modified version of EGS4 that was used, in excess of the errors in the best available theory. In light materials these errors arise mostly from the multiple scattering treatment; in lead, the inaccuracies in the treatment of annihilation and bremsstrahlung also contribute.

The next column—calculational and experimental uncertainties—includes all other effects. Geometrical approximations, counter calibration errors, Monte Carlo statistics, etc. are included.

The total systematic error in the fitted value of η is estimated to be ± 0.056, not including the uncertainty in ρ. This is broken down somewhat in Table 6.1; the error estimates are discussed in more detail in Appendix E, and other references are given in Table 6.2. The "miscellaneous" entry includes events caused by conversion of γ's which start from the target or $K2$, events from shower penetration through $K2/K4$ and errors from NMR probe cross-calibration.

6.9 Conclusions

In the accepted, existing theoretical treatments of the muon decay spectrum and its radiative corrections, the W^+ vector boson has been taken to be infinitely heavy. As discussed in Section 2.4, this yields an excellent approximation to the exact muon decay spectrum, provided that there are no tensor, scalar or pseudoscalar couplings. Since there is no evidence of these couplings in this or other measurements, one may rely upon the radiative corrections that were used—primarily those calculated by Grotch to first order. The higher-order corrections had little effect on the lower and middle portions of the spectrum used in fitting a value of η.

The distortions to the theoretical spectrum, which are inevitable when measuring the low-energy end of the spectrum (especially from stopped muons), have been studied and are understood. The calculation of distortions through variance-reduced and interpolated Monte Carlos has given predictions in close agreement with experimental results, although experimental input was required at low energy. Systematic errors due to other phenomena (such as errors in scaler values and magnetic-field calibrations) made only minor contributions to the overall error.

Thus, it has been possible to largely correct the instrumental effects in the spectrum measurement and to make the experiment self-normalizing. For ρ constrained to 3/4, the best-fit value was $\eta = -0.080 \pm 0.088(\text{statistical}) \pm 0.056(\text{systematic})$. This value is in agreement with the V-A prediction of zero, and there is no indication of exotic physics, such as massive mixed neutrinos, in the $P_e \epsilon (6.15, 40)$ MeV/c range of the spectrum. The determination of η is somewhat sensitive to ρ, and values of ρ different from 3/4 would shift the fitted value by $\Delta\eta = 12.2\,(\rho - 3/4)$.

Source	Shift in η_{fit}	$\Delta\eta$ (theory)	$\Delta\eta$ (EGS4)	$\Delta\eta$ (exp., calc.)	$\Delta\eta$
μ^+ depth	+0.408			±0.029	±0.029
Back plate	+0.133	±0.006	±0.004	±0.027	±0.028
$T1$	+0.408	±0.009	±0.014	±0.017	±0.025
Detector ineff.	−0.218		±0.005	±0.014	±0.015
μ^+ spin angle				±0.014	±0.014
Cables	−0.057	±0.003	±0.005	±0.009	±0.011
$K2$	+0.046	±0.005	±0.005	±0.009	±0.011
$C1$	+0.069		±0.002	±0.009	±0.009
P_{max} calib.				±0.008	±0.008
Beam centering				±0.007	±0.007
Muon stops				±0.006	±0.006
$K3$	+0.014	±0.002	±0.002	±0.004	±0.005
Line shape				±0.005	±0.005
$C2$	−0.008		±0.002	±0.002	±0.003
Asymmetry				±0.002	±0.002
Miscellaneous	−0.020	±0.002		±0.004	±0.005
Total		±0.013	±0.018	±0.056	±0.056

Table 6.1: Error estimates for systematic effects on the fitted value of η (after corrections).

Source	Geometry	References for Effect
μ^+ depth	Fig. 15, Sec. 4.1.5	Sec. 6.1.1, E.2.1
Back plate	Fig. 8	Fig. 27, Sec. 5.10, 6.6, E.2.2
$T1$	Fig. 15, Sec. 4.1.5	Fig. 17-20, Sec. 5.1, 6.1.2, E.2.3
Detector ineff.	Fig. 8, 11, Sec. 4.1.8	Fig. 21, Sec. 5.2, 6.1.6, E.2.4
μ^+ spin angle		Sec. A.2, E.2.5
Cables	Fig. 8	Fig. 22, Sec. 5.3, 6.1.5, E.2.6
$K2$	Fig. 8, Sec. 4.1.6	Fig. 23, Sec. 5.5, 6.1.4, E.2.7
$C1$	Fig. 8, Sec. 4.1.6	Fig. 25, Sec. 5.9, 6.1.4, E.2.8
P_{max} calib.	Sec. 4.1.1	Sec. 6.4, B.4, E.2.9, 4.1.2
Beam centering		Sec. A.1, E.2.10
Muon stops	Fig. 15, 16	Sec. 6.5, E.2.11
$K3$	Fig. 8, Sec. 4.1.6	Fig. 24, Sec. 5.6, 6.1.4, E.2.12
Line shape	Fig. 13, 14, 30	Sec. E.2.13
$C2$	Fig. 8, Sec. 4.1.6	Fig. 26, Sec. 5.9, 6.1.4, E.2.14
Asymmetry		Sec. A.3, E.2.15
Miscellaneous	Fig. 8	Sec. 5.4, 5.7, 6.1.4, E.2.16

Table 6.2: References to information on the systematic errors.

Appendix A

Muon-Polarization Effects

Since a stopping muon retains about 25% of its initial polarization in plastic scintillator, one must ensure that one's spectrometer cancels the asymmetries in the muon decay. In other words, one must verify that it is the unpolarized decay spectrum being measured in an experiment, and not something else. While an azimuthally symmetric, axial-focusing spectrometer whose axis and magnetic field at the muon position is perpendicular to the muon spin will do this in principle, one must consider the effects of deviations from the idealized situation.

A.1 Misaligned $\vec{B}$ at Target

In the Comus spectrometer, there is a magnetic field at the target which points along the symmetry axis of the spectrometer, defined as $\hat{z}$. The muons are polarized in a direction perpendicular to this, defined here as $\hat{x}$, so that their spins precess around $\hat{z}$. As long as the spectrometer is azimuthally symmetric around $\hat{z}$ (the possibility that it is not will be discussed later), there is no difficulty. However, suppose that the magnetic field has non-zero B_x or B_y components: the muon spins will precess around around $\vec{B}$, not $\hat{z}$. These radial field components could arise from several sources—

- The spectrometer magnet might be slightly asymmetric.