Search for a Heavy Photon in Electro-Produced e^+e^- Pairs with the HPS Experiment at JLab

Omar Moreno[1,2*], Nathan Baltzell[3], Mathew Graham[1], John Jaros[1]

[1] *SLAC National Accelerator Laboratory, Menlo Park, CA 94205*
[2] *University of California, Santa Cruz, Santa Cruz, CA 95064*
[3] *Thomas Jefferson National Laboratory, Newport News, VA 23606*

Abstract

The Heavy Photon Search engineering run took place in the spring of 2015 using a 1.056 GeV, 50 nA beam. This note will describe the details of an analysis used to search for a prompt heavy photon with a mass in the range 20 MeV/c^{2}and 90 MeV/c^{2}using the full engineering run dataset which amounts to a luminosity of 1165.7 nb^{-1} (7.28 mC of charge). No significant excess above the e^+e^- invariant mass spectrum that would arise from the decay of the heavy photon ~~at~~ an e^+e^- pair was found. Since no significant signal was observed, 2σ upper limits on the square of the coupling were set at the level of 10^{-5}.

---

* Corresponding author. E-mail: omoreno@slac.stanford.edu

Contents

1 Introduction

This note will present the results of the first search for prompt heavy photon (A') decays by the Heavy Photon Search (HPS) experiment using the data taken in the 2015 HPS Engineering Run at JLab. The search involves scanning the measured e^+e^- invariant mass spectrum for a narrow resonance (the A') above a large continuous background, i.e. conducting a "bump hunt".

The production of a canonical A' is completely described by two parameters, its coupling ϵe^- to the photon, and its mass $m_{A'}$. For the parameters for which this search is sensitive, decays of the A' are prompt so the decay products would emanate from the target. Specifically, the A' production cross-section is related to the production cross-section for trident production via virtual photon bremsstrahlung ("radiative production") as

$$\frac{d\sigma(e^-Z \to e - Z(A' \to l^+l^-))}{d\sigma(e^-Z \to e - Z(\gamma^* \to l^+l^-))} = \frac{3\pi\epsilon^2}{2N_{eff}\alpha}\frac{m_{A'}}{\delta m} \tag{1}$$

Here, N_{eff} is the number of decay channels kinematically accessible (=1 for HPS searches below dimuon threshold) and δ_m is the width of the mass window centered at $m_{A'}$ in which we search. In practice, we take $\delta_m(m) = 2.56\sigma_m$ where σ_m is the experimental invariant mass resolution at mass m. We expect 79.9% of the signal events to appear in this mass window.

Relationship 1 lets us calculate the number of heavy photons of mass $m_{A'}$ and coupling ϵ which would be detected in terms of the number of observed e^+e^- pairs, assuming we know the fraction of all e^+e^- pairs which are "radiative" and assuming we know the mass resolution. HPS can accurately account for the observed rate and kinematics of its e^+e^- sample by attributing it to three sources: trident production, e^+e^- production resulting from converted wide angle bremsstrahlung events (WABs), and a small admixture of accidental coincidences. MadGraph5 physics generators are used to simulate trident and WAB production and to evaluate the radiative fraction in the kinematic region we use to search for heavy photons. The mass resolution is determined in a full simulation of the HPS detector response, scaled by the same factor which matches the mass resolution ~~observance~~ for the e^+e^- Møller mass to that seen in simulation. Møllers have a unique mass at each center of mass energy and serve as a perfect calibration point.

Our search for an A' signal has been conducted as a blind search. We did so by first developing the cuts used to isolate the signal using a 10% unblinded sample of the 2015 events. This sample was also used to optimize the fitting functions for the background shape of the invariant mass distribution, which were then used in the analysis of the 10% sample to establish limits on the expected A' signal and its coupling as a function of mass. After unblinding the entire data sample, these same cuts and fitting procedures were employed to derive the final results below.

This analysis uses all the 2015 Engineering run data satisfying the "golden run" criteria: the SVT was positioned 0.5 mm from the beam and ~~a~~ run had $>$ 5M events. The data satisfying these requirements, corresponding to an integrated luminosity of 1165.7 nb$^{-}1$, were processed by the HPS physics reconstruction pipeline tagged as "tweakpass-6" and the HPS Data Summary Tape version 0.11.

The remainder of this note will review the following: (1) e^+e^- selection criteria and tuning of cuts in the 10% sample; (2) the composition of the e^+e^- sample and determination of the radiative fraction; (3) HPS mass resolution; (4) the e^+e^- invariant mass spectrum and procedures for searching for a resonance; (5) fits and upper limits; (6) systematic errors; (7) physics results.

2 Event Selection

Searching for a heavy photon resonance requires the accurate reconstruction of the e^+e^- invariant mass spectrum and the efficient rejection of background events due to accidental e^+e^- pairs, wide angle bremsstrahlung (WAB) and Bethe-Heitler tridents. With these requirements in mind, a set of cuts was developed and optimized using roughly 10% of the 2015 engineering run dataset processed using the reconstruction pipeline version "tweakpass6". A detailed discussion of the reconstruction can be found in HPS-NOTE xxx. This section will describe the selection used to isolate the final e^+e^- sample along with the procedure used to optimize it.

2.1 Preprocessing

Before kinematic and timing cuts are applied, the data is preprocessed to select events which have signatures that are loosely consistent with a true e^+e^- pair. To begin, events are required to satisfy the "pairs-1" trigger criteria. This selects events with two clusters whose position and energy in the Ecal are consistent with an e^+e^- from either a trident reaction or the decay of an A' [1]. Additionally, only events where the SVT was fully biased (i.e. bias at 180 V) and positioned at 0.5 mm from the beam are considered. Furthermore, only events free of data acquisition errors were included in the final data sample.

Ensuring that true e^+e^- pairs are selected begins by first requiring that an event have a single unique positron in at least one of the volumes of the SVT. Events where more than a single unique positron are found are not considered in this analysis unless their tracks all share a large number of hits, indicating the presence of a single real track. When an SVT volume does have more than a single positron track, and all of the tracks within that volume share at least 4 hits, the track with the best fit χ^2 is chosen.

Currently, the reconstruction takes all possible e^+e^- pairs, vertexes each pair, and creates a collection of vertexed particles (V_0's) without applying any additional requirements. As a result, along with true e^+e^- pairs, there will be "accidental" pairs constructed by combining beam e^-'s or e^-'s that originate in a different bunch with the detected positron. These events can be largely rejected by requiring the V_0 particles satisfy the criteria specified on Table 1.

Cut Description	**Requirement**
The V_0 must have been created from vertexing GBL tracks	particle type > 32
Ecal clusters in opposite volumes	$y_{e^+ \text{ Cluster}} \times y_{e^- \text{ Cluster}} < 0$
Loose track-cluster match	$\chi^2 < 10$
Beam electron cut	$p(e^-) < 0.75 E_{\text{beam}}$

Table 1: Requirements applied to V_0 particles during the preprocessing stage.

The physics trigger used by HPS was tuned to accept time coincident e^+e^- pairs, where the e^+ and e^- reside in opposite detector volumes. Therefore, as an initial requirement, the Ecal clusters associated with the e^+ and e^- are required to be in opposite halves of the detector, i.e. have a y position which satisfies the following relation:

$$y_{e^+ \text{ Cluster}} \times y_{e^- \text{ Cluster}} < 0. \quad (2)$$

Since this analysis only considers pairs formed using tracks that are matched to Ecal clusters, a loose cut is placed on the track-cluster matching χ^2 to guard against the case where a track is grossly mismatched to an Ecal cluster. Specifically, the track-cluster match χ^2 is required to be less than 10. This ensures that the majority of matched tracks and clusters are within 2-3 cm of each other (See Figure 1).

Finally, an e^- associated with a e^+e^- pair may actually be an elastically scattered beam electrons of energy $E_{\text{beam}} = 1.056$ GeV instead of one associated with a true e^+e^- pair. To remove these V_0's from the final sample, the momentum of the electron track is required to be less than $0.75E_{\text{beam}}$. This is illustrated on Figure 2 which shows the beam electron peak (mean at 1.056 GeV) and the cut used to reject it.

The above requirements reject a large fraction of accidental V_0 particles. If a positron continues to have multiple V_0 particles associated with it after all cuts have been applied, the V_0 particle with the best vertex χ^2 is chosen. This is essentially placing a requirement on the χ^2 of the electron track and will remove any additional V_0 particles that result from electron tracks that share hits or mis-reconstructed tracks. The result of applying all of these requirements is a set of unique V_0 particles.

2.2 Radiative Selection

As discussed in Section 1, the kinematics of the radiative trident background are indistinguishable from A' signal events. Although the rate of the Bethe-Heitler process dominates among the two trident processes, its different kinematics can be used to reduce its contributions to the final event sample. Specifically, the A' decay products are highly boosted while the recoiling electron is soft and scatters at large angles. In contrast the Bethe-Heitler process is not enhanced at higher pair energies. Furthermore, only one of the leptons in the Bethe-Heitler pair will be highly boosted, while the other will be much softer. These kinematic differences are illustrated in Figure 3 which plots the momentum of the positron versus that of the electron for both A' (maroon) and Bethe-Heitler (blue) events. As illustrated in the figure, the signal distribution is concentrated in the region where the sum of the energy of the electron and positron ("p-sum") is approximately equal to beam energy. Therefore, p-sum can be used to discriminate between the two processes. Figure 4 shows the distribution of the sum of the momenta of the e^+e^- tracks composing a pair. From studies of simulated events, it was found that requiring the sum of the momentum to be greater than $0.8E_{beam}$ maximizes the rejection of Bethe-Heitler events while minimizing the impact on signal.

2.3 Rejecting Accidentals and Poor Quality Tracks and Vertices

The preprocessing stage results in a set of unique V_0 particles. However, additional cuts were imposed to remove out of time backgrounds, enforce good track and vertex quality and further eliminate contamination by beam electrons. This section will discuss these requirements in detail.

Although the HPS trigger requires time coincident pairs, the trigger window is wide enough that there were instances where a pair is formed using a e^+ and e^- from different beam bunches or even from different interactions within a single bunch. This can be seen from Figure 5 which plots the cluster time of one cluster composing a pair versus that of the other cluster. The central band (yellow) is due to pairs originating in the triggering beam bunch, while the neighboring bands

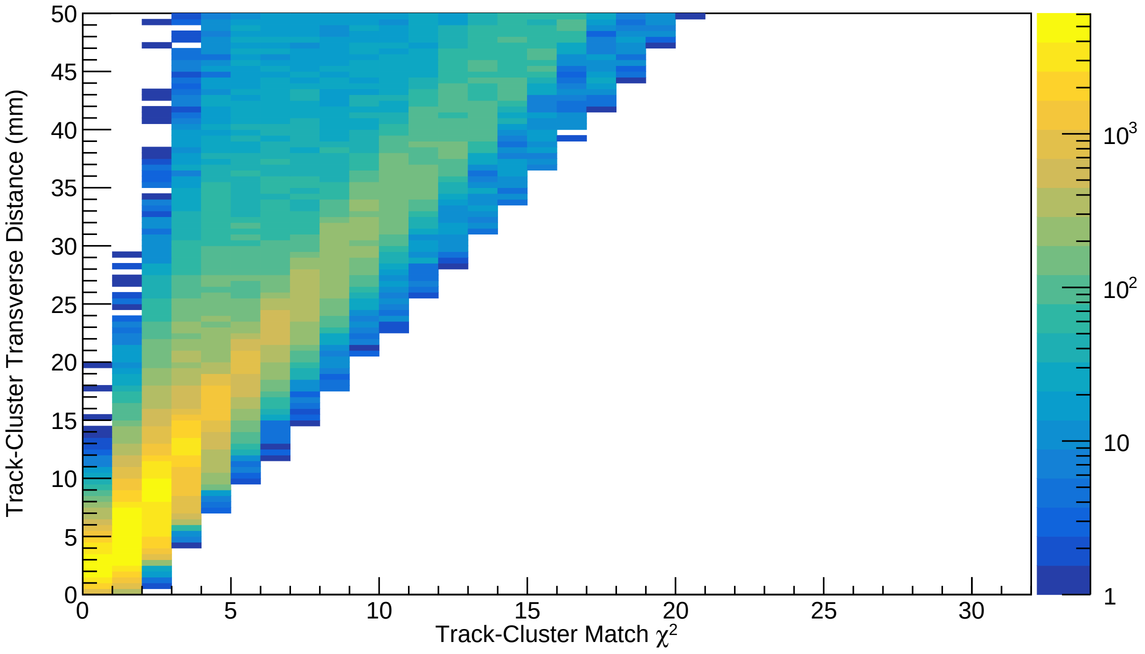

Figure 1: The track-cluster match χ^2 versus the transverse distance of the track from the matched cluster. Requiring the track-cluster match χ^2 to be less than 10 ensures that the transverse distance between the track and cluster is within a couple of mm.

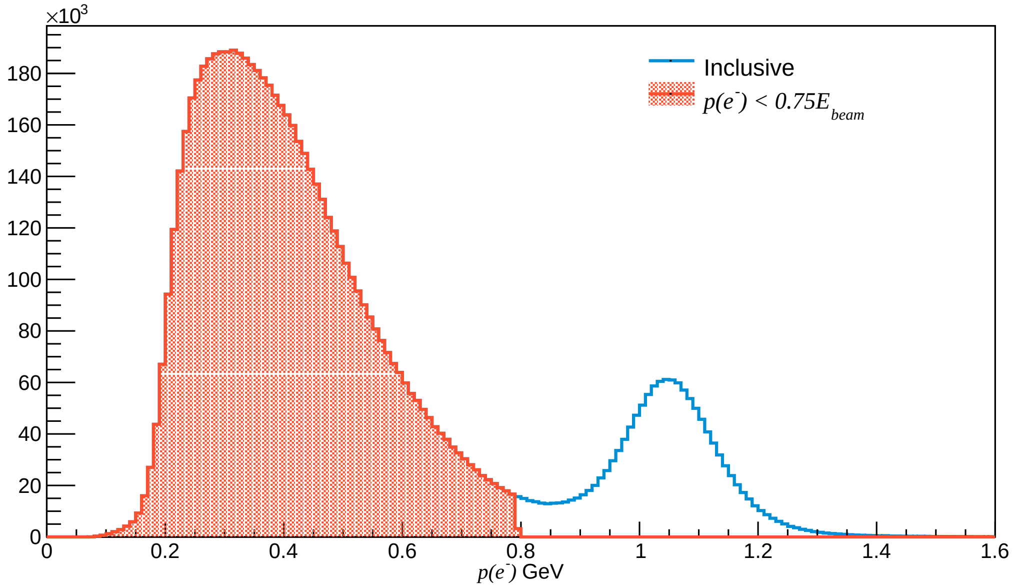

Figure 2: The momentum of electrons associated with a vertexed particle (V_0) is shown in blue. The orange distribution shows the sample that remains after the beam electron cut $(p(e^-) < 0.75E_{\text{beam}})$ is applied.

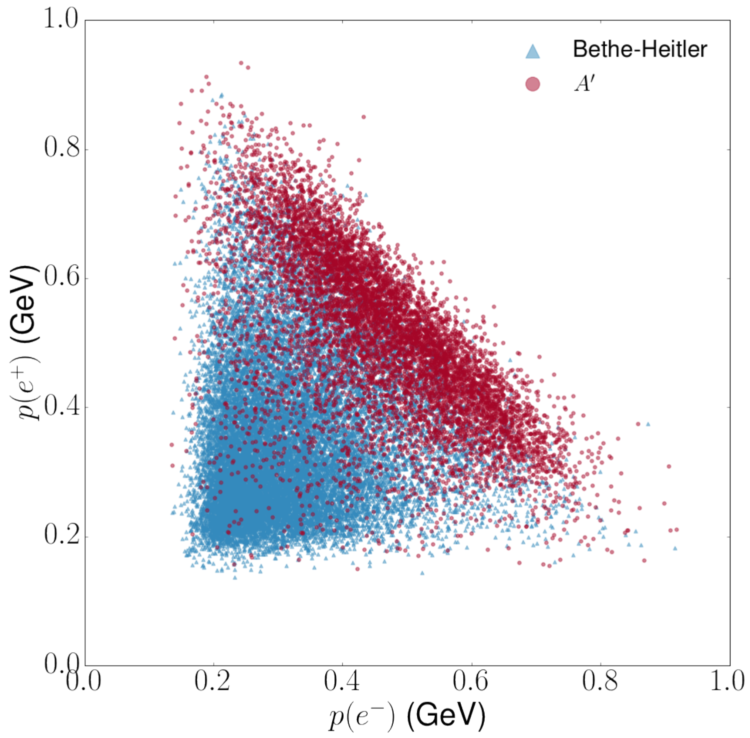

Figure 3: Scatter plot of the positron momentum versus the electron momentum of pairs from Bethe-Heitler background (blue) and 50 MeV A' signal events (red). The kinematics of the irreducible radiative background are indistinguishable from A' signal events within an invariant mass window, δm, centered at $m_{A'}$ and can be used to analyze the rate of A' signal production.

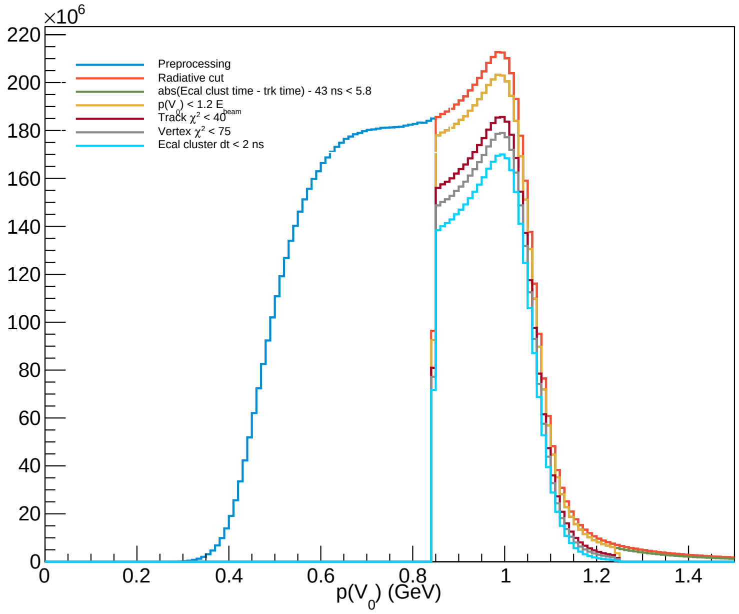

Figure 4: Cutflow of the sum of the momenta of the e^+e^- tracks composing a pair. The sharp cutoff corresponds to the radiative cut.

(green) saw the e^- or e^+ originate from an out-of-time bunch. Therefore, suppression of pairs where the e^- or e^+ originate in an out of time bunch is achieved by requiring that the difference in Ecal cluster times is less than 2 ns. No explicit requirement is placed on the individual cluster times.

As mentioned previously, some track-cluster matches are formed using mis-reconstructed tracks that are mismatched to an Ecal cluster. Given that the track-cluster match χ^2 cut was chosen to be loose, a cut on the difference between the track and cluster times was used to further suppress e^+e^- pairs where one of the tracks is mismatched to the corresponding cluster. Specifically, a track-cluster match is required to satisfy the following condition:

$$\text{abs}(t_{\text{cluster}} - t_{\text{track}} - 43) < 5.8 \text{ ns} \tag{3}$$

Since the track time is referenced to the trigger time and the Ecal cluster time to the start of the Ecal readout window, a constant offset of 43 ns, reflecting the time difference in those references, is subtracted from the track-cluster time difference. A plot showing the effect of all cuts on the track-cluster time difference is shown in Figure 6.

Even after applying the track-cluster match requirements, there are still instances where one or both of the tracks associated with a pair are mis-reconstructed. In order to suppress these types of events, loose requirements on both the quality of the tracks and the vertex they form are enforced. Specifically, the χ^2 of the GBL fit of all tracks is required to be less than 40. Furthermore, the vertex χ^2 must be less than 75. As illustrated in Figures 7 and 8, after applying all other cuts, the quality requirements reject only the worse tracks and vertices which lie on the tails of these distributions.

Finally, the momentum of the V_0 is required to be less than $1.2E_{\text{beam}}$. This further eliminates pairs where a beam electron or mis-reconstructed positron track has been included.

The effect of the above requirements on the Ecal cluster time difference is shown in Figure 9, while the base cut flow is summarized in Table 2. As can be estimated from the figure, the fraction of accidentals contaminating the final event sample is less than 1%. This is estimated from the size of the outer peaks before the cluster time coincidence cut is applied. Additional cuts used to eliminate the background from converted wide-angle bremsstrahlung are discussed below.

Cut Description	**Requirement**
Track-cluster match quality	$\text{abs}(t_{\text{cluster}} - t_{\text{track}} - 43) < 4.5$ ns
Momentum sum cut	$p(e^+e^-) < 1.2E_{\text{beam}}$
Track quality	$\chi^2 < 40$
Vertex quality	$\chi^2 < 75$
Cluster time coincident	< 2 ns

Table 2: Selection used to reject accidental e^+e^- pairs.

2.4 Minimizing Wide-Angle Bremsstrahlung

The most significant non-trident contribution to the pairs sample comes from wide-angle bremsstrahlung when the incident electron radiates a hard photon at a relatively large angle to the beam axis

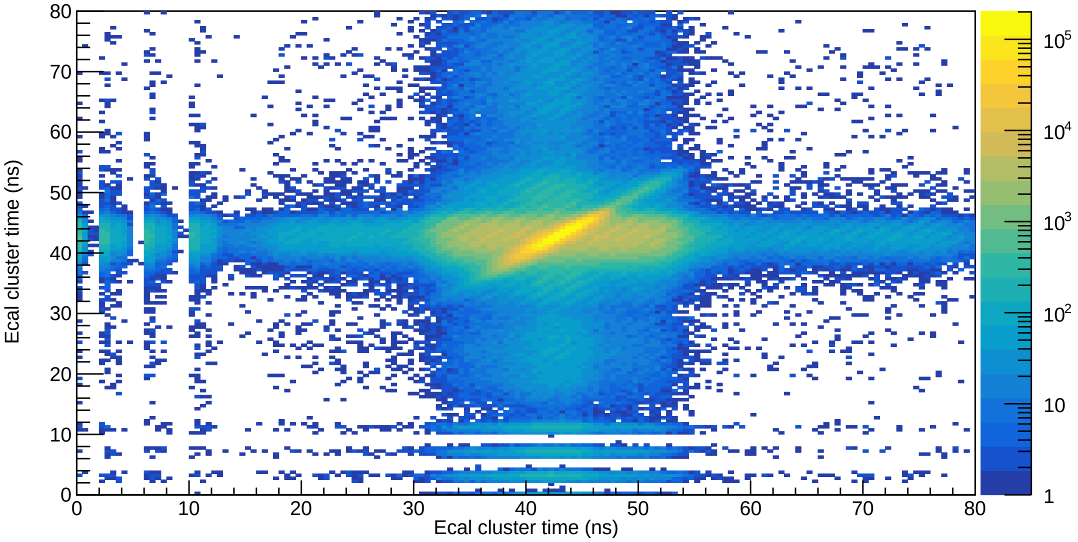

Figure 5: Correlation between Ecal cluster times composing a pair. The central band (yellow) is due to pairs originating in the triggering beam bunch, while the neighboring bands (green) saw the e^- or e^+ originate in an out of time bunch.

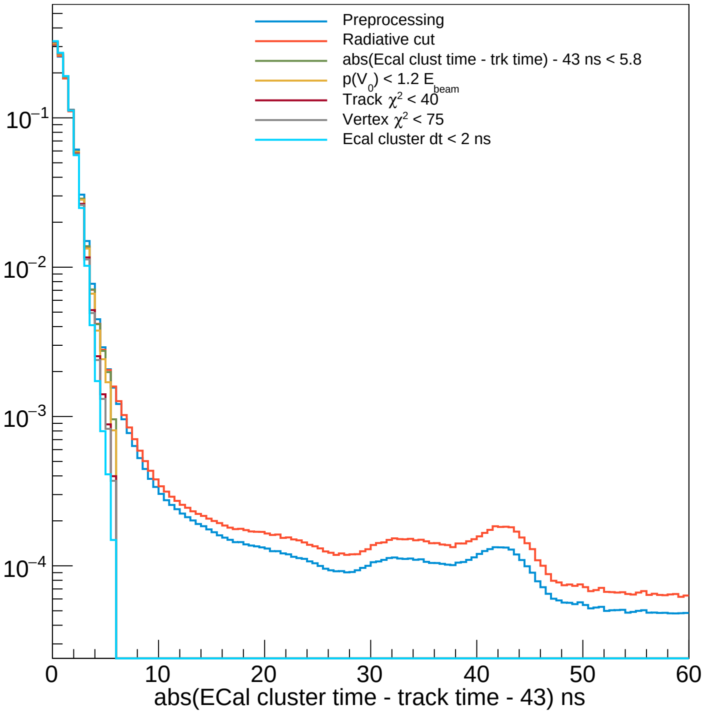

Figure 6: Cutflow of the time difference between a track and cluster. The cut placed at 5.8 ns was used to eliminate e^+e^- pairs where one of the tracks is mismatched to the corresponding cluster.

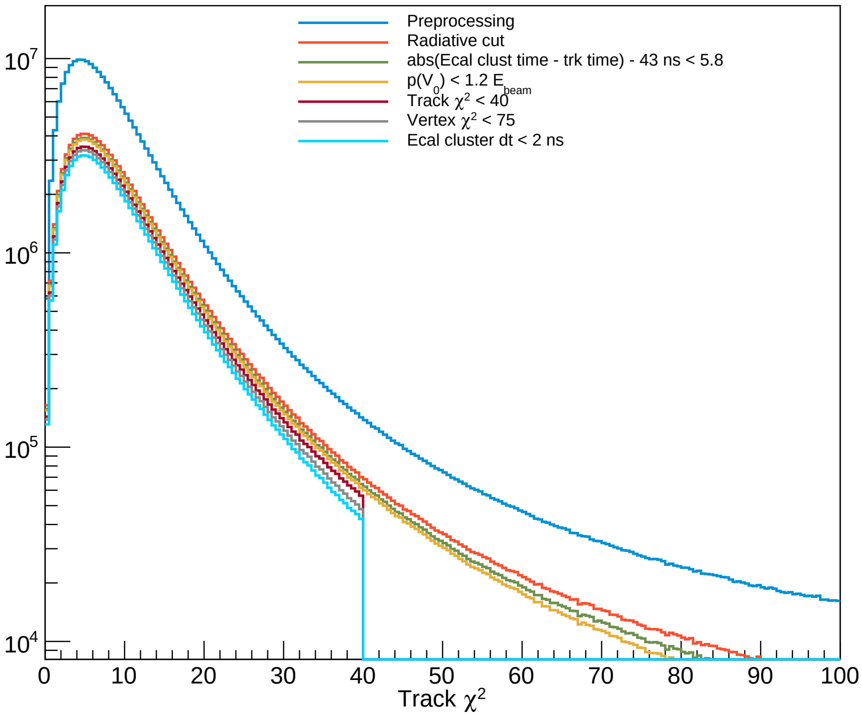

Figure 7: Cutflow of the track χ^2. The sharp cutoff at 40 denotes the requirement used to eliminate grossly misreconstructed tracks.

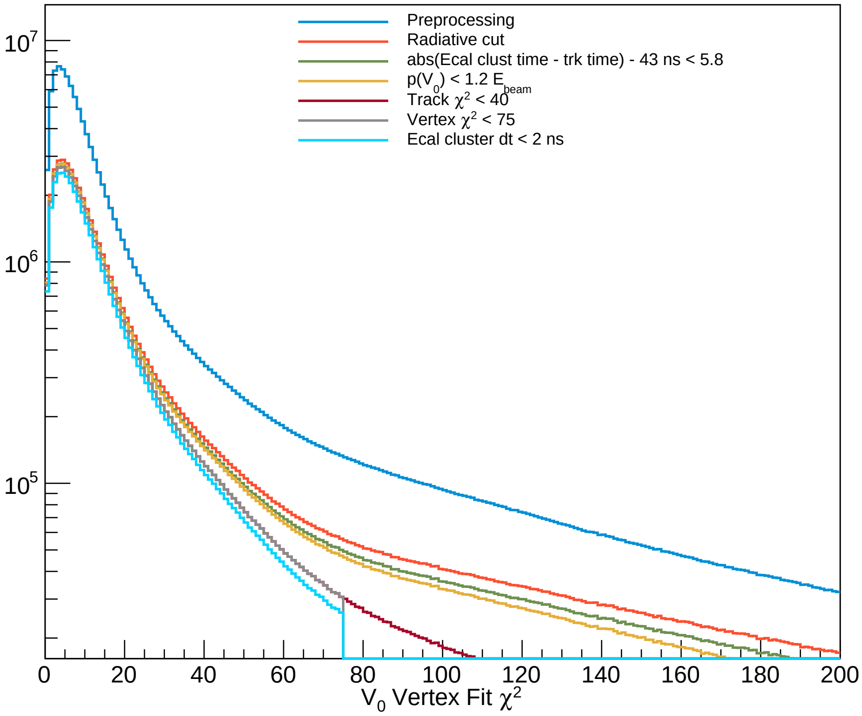

Figure 8: Cutflow of the vertex fit χ^2. In order to eliminate misfit vertices, the vertex fit χ^2 was required to be less than 75.

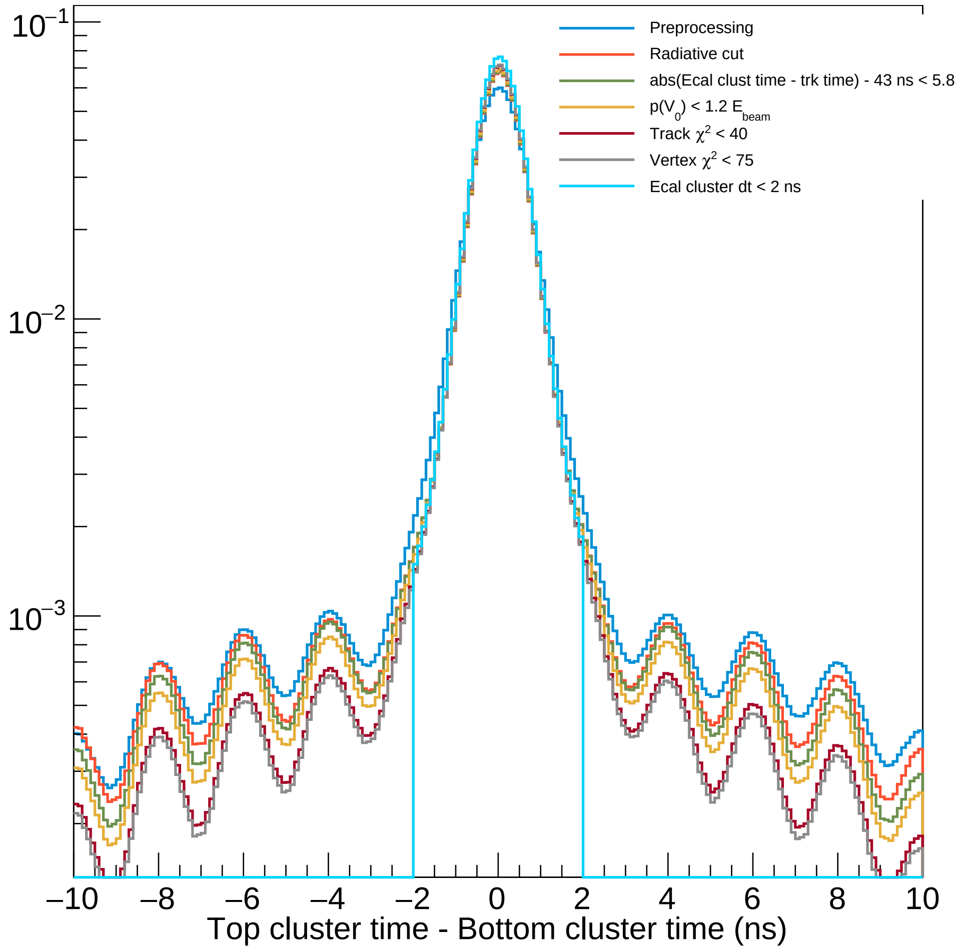

Figure 9: Difference in time between the Ecal clusters composing a pair. After all cuts, the contamination due to accidentals is less than 1%.

which subsequently converts in either the target or first few layers of the SVT. In the V_0 sample, the electron from the pair conversion is rarely reconstructed. It is much more likely that the positron from the conversion along with the recoil electron are within the acceptance, thereby constituting a background to the e^+e^- sample.

Since, for the 5 and 6 layer tracks under consideration, the conversion of the hard photon most frequently occurs in the first layer of the SVT, the positron may not leave enough energy in both layer 1 sensors to record a 3D hit. This can be seen in Figure 10 which plots the rate of MC trident and WAB positron tracks which have a layer 1 hit. It's clear that the positron track associated with a trident event will almost always have a layer 1 hit while the majority of positrons from converted WABs are missing it. As a result, requiring a positron track to have a layer 1 hit will greatly reduce the number of WAB's contaminating the final event sample.

Another variable that can be used to discriminate WAB's is the distance of closet approach. When the WAB conversion occurs in the SVT, the positron track will not extrapolate back to the target because of its finite curvature. Instead, it will extrapolate to one side resulting in a skewed distribution. This can be seen in Figure 11 which shows the distribution of the distance

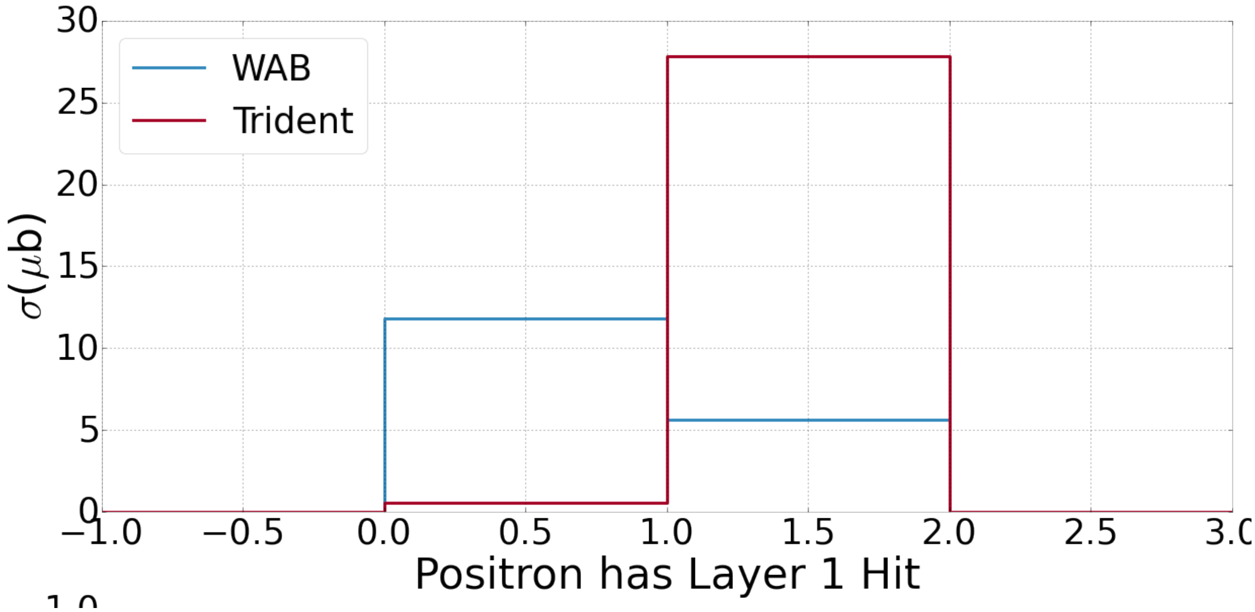

Figure 10: Histogram of the rate of Monte Carlo tridents (maroon) and WAB (blue) positron tracks which have a hit in layer 1. Positron tracks associated with a trident will almost always have a layer 1 hit, while positrons from a WAB are missing it.

of closest approach, d_0, for positron tracks from both trident and wab events. Requiring the d_0 of positrons to be less than 1.1 mm eliminates the WAB tail.

Finally, the p_t asymmetry of the e^+e^- pair can be used to suppress WAB's even further. This can be seen in Figure 12, which shows the p_t asymmetry of WAB's is skewed compared to trident events, since the momentum of the missing e^- from the conversion is unaccounted for. Requiring the p_t asymmetry to be less than 0.47 further reduces the number of converted WAB's in the final event sample.

The cuts used to suppress WAB's are summarized in Table 3. After applying all cuts, $> 80\%$ of WAB's were cut from the final event sample.

Cut Description	**Requirement**
L1 hit and L2 requirement	e^+ is required to have L1 and L2 hit
e^+ d_0	$d_0 < 1.1$ mm
p_t asymmetry	$p_t(e^-) - p_t(e^+)/p_t(e^-) + p_t(e^+) < .47$

Table 3: WAB rejection cuts applied to all e^+e^- pairs used in this analysis.

2.5 Selection Optimization

This section describes how the critical variables for e^+e^- event selection were chosen, and how optimal cuts in those variables were determined. The effectiveness of any given variable at rejecting background events was evaluated using the corresponding Receiver Operating Characteristic

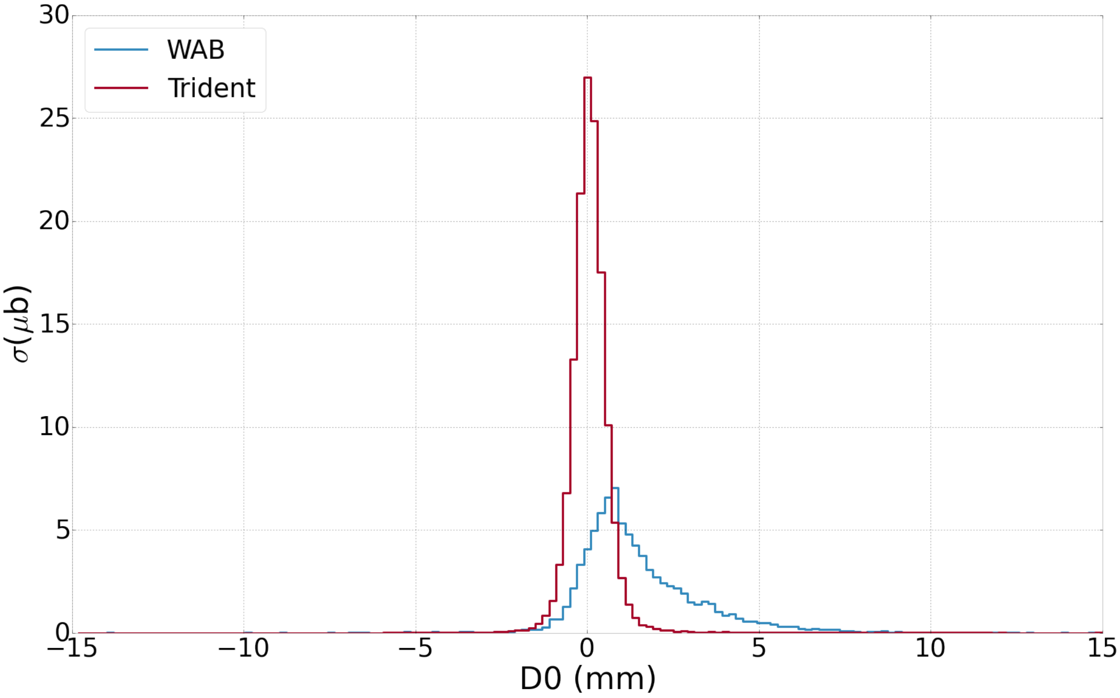

Figure 11: Distance of closest approach of positrons associated with tridents (maroon) and WABs (blue). Requiring the d_0 of positrons to be less than 1.1. mm eliminates the WAB tail.

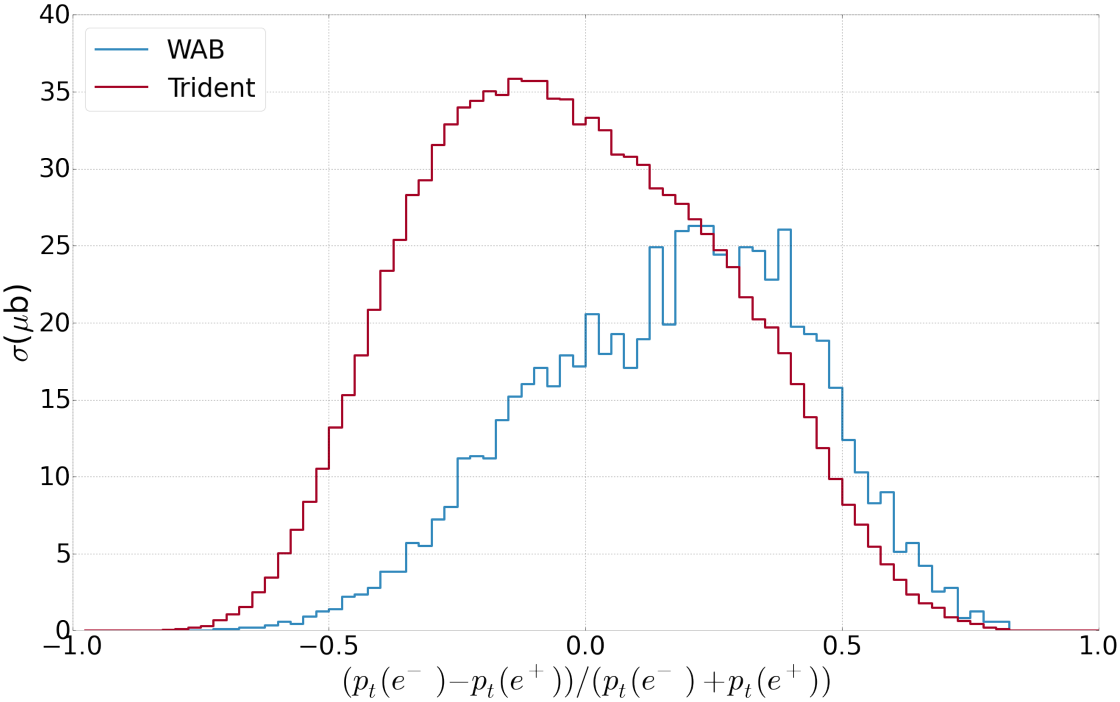

Figure 12: Momentum asymmetry of tridents (maroon) and WABs (blue). Requiring the asymmetry to be less than 0.47 further reduces the number of converted WABs in the final sample.

curve (ROC curve). Such a curve plots the signal fraction vs 1 - the background fraction as a particular variable is varied. The integral of this curve over the full range 0 to 1.0 provides a measure of the effectiveness of a cut on this particular variable, higher integrals corresponding to greater effectiveness. The effectiveness of several potential variables can then be ordered, and the most effective variables chosen. Once an effective variable is chosen, one can select a value for the cut in that variable which optimizes the significance.

As an example, a cut on the electron momentum is important to separate true e^+e^- events from background events arising from in time or out of time beam electrons. Figure 15 shows how the number of preprocessed events (before making a cut on the electron momentum) depends on the cluster time difference in the ECal. It clearly shows a signal region, near 0 time difference, and a background region, where the absolute value of the time difference exceeded 3 ns. The use of these regions is possible because Ecal cluster pairs which satisfy equation 4 are dominated by pairs where at least one of either the e^+ or e^- originates outside the beam bunch of interest while pairs satisfying 5 are dominated by true e^+e^- pairs.

$$\text{Background region: } t_{\text{top}} - t_{\text{bottom}} > 3 \text{ ns} \tag{4}$$

$$\text{Signal region: } t_{\text{top}} - t_{\text{bottom}} < 1 \text{ ns} \tag{5}$$

Figure 14 are plots the electron momentum distributions for all the preprocessed events, those in the signal region, and those in the background region. For each potential value of the cut in electron momentum, the figure determines a signal fraction and a background fraction. The related ROC curve corresponding to a cut in electron momentum is shown in Figure 13. Its integral is relatively large, indicating the effectiveness of this variable. The cut in electron momentum is optimized by plotting the signal divided by the square root of the signal plus background, which is shown in Figure 14. It's clear that a cut at 0.8 GeV/c maximizes the significance of the cut.

This procedure has been used to identify the most effective variables for cuts, and to determine the optimal values of the cuts with those variables. ~~Variables~~ with ROC curve integrals below 0.5 were dropped. Many of these cuts were data driven as in the example above. In some cases, as for example in selecting a sample of WAB events, MC is used. Optimization of the selection used to reject WAB's used a Monte Carlo "tri-trig" sample as the signal and a Monte Carlo WAB sample as the background. Monte Carlo was needed because, unlike accidentals, isolating a pure WAB sample in data is difficult.

The sample used to optimize the accidental cuts was generated by preprocessing 7.5% of the engineering run data as described in Section 2.1. However, the "beam electron cut" was excluded from the requirements as it was also part of the optimization procedure. As shown in Figure 15 , the resulting sample was split into "signal" and "background" regions as defined by Equations 4 5.

2.6 Event Selection Efficiency

The cuts used to isolate the final e^+e^- invariant mass distribution along with their efficiency for data, trident MC, radiative MC, WAB MC and 30 MeV A' events are summarized in Table 4. The effect of each cut on the data e^+e^- invariant mass sample is also shown in Figure 16. In total, after all cuts, the final sample contains 21M events.

Description of Cuts	**Data**		**Trident MC**	**Radiative MC**	**WAB**	**30 MeV A' MC**
	N$_{\text{events}}$	ε (%)	ε (%)	ε (%)	ε (%)	ε (%)
Event Level						
Trigger Count	2367740075	-	-	-	-	-
SVT Quality Cut	2062863479	87.1	-	-	-	-
Events with a track	1554997154	65.7	92.6	92.5	78.0	98.6
Events with e^+ track	242319129	10.2	77.8	75.8	3.0	89.4
e^+e^- **Pairs**						
Preprocessing	113433204	4.8	54.8	56.0	2.4	76.7
Radiative cut	49293122	2.1	19.7	33.2	1.5	56.6
abs($t_{\text{cluster}} - t_{\text{track}} - 43) < 4.5$ ns	46468774	2.0	19.7	33.2	1.5	56.6
$p(e^+e^-) < 1.2E_{\text{beam}}$	45377797	1.9	19.7	33.2	1.5	56.5
Track $\chi^2 < 40$	40171569	1.7	18.8	31.8	1.4	54.5
Vertex Fit $\chi^2 < 75$	38293045	1.6	18.7	31.6	1.4	54.1
Ecal cluster pair $\Delta t < 2$ ns	35810225	1.5	18.4	31.1	1.4	53.3
L1 and L2 hit	24901106	1.1	17.8	29.9	0.4	52.4
$d_0(e^+) < 1.1$ mm	22733909	1.0	17.4	29.5	0.3	51.2
$p_t(e^-) - p_t(e^+)/p_t(e^-) + p_t(e^+) < .47$	21452278	0.9	16.7	28.5	0.3	49.7

Table 4: Table showing the efficiency of each cut for data, a sample of trident MC, pure radiatives and 30 MeV A' events. The trident sample contains both Bethe-Heitler and radiative events.

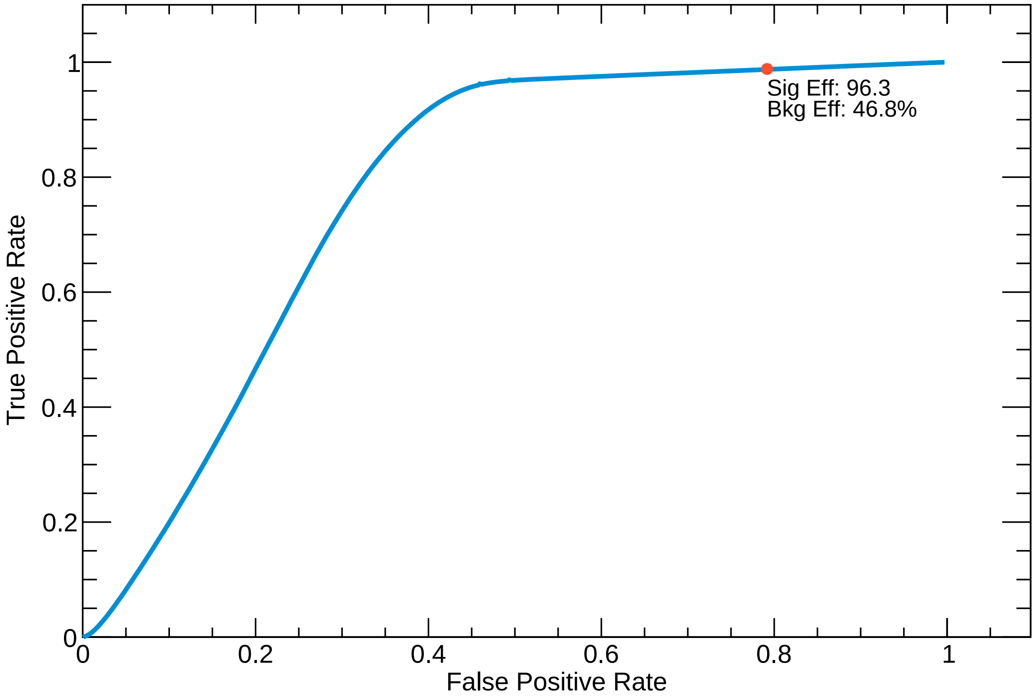

Figure 13: ROC curve showing the true positive versus the false positive rate for a given beam electron cut.

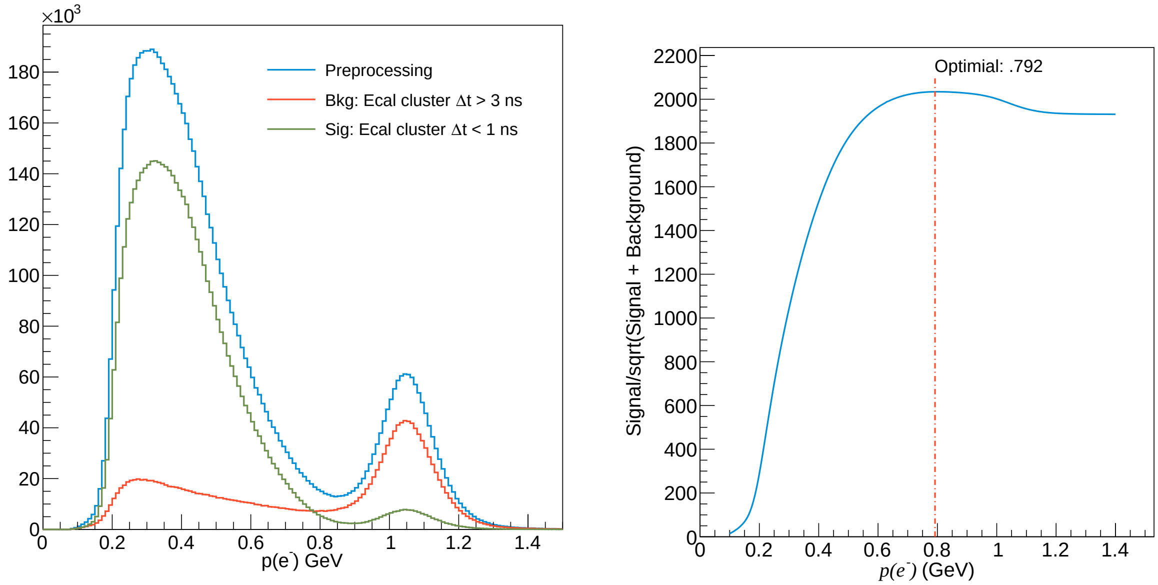

Figure 14: The plot on the left shows the signal (purple) vs background (red) distributions use to optimize the beam electron cut. The plot on the right shows the metric $S/\sqrt(S + B)$ for a given cut.

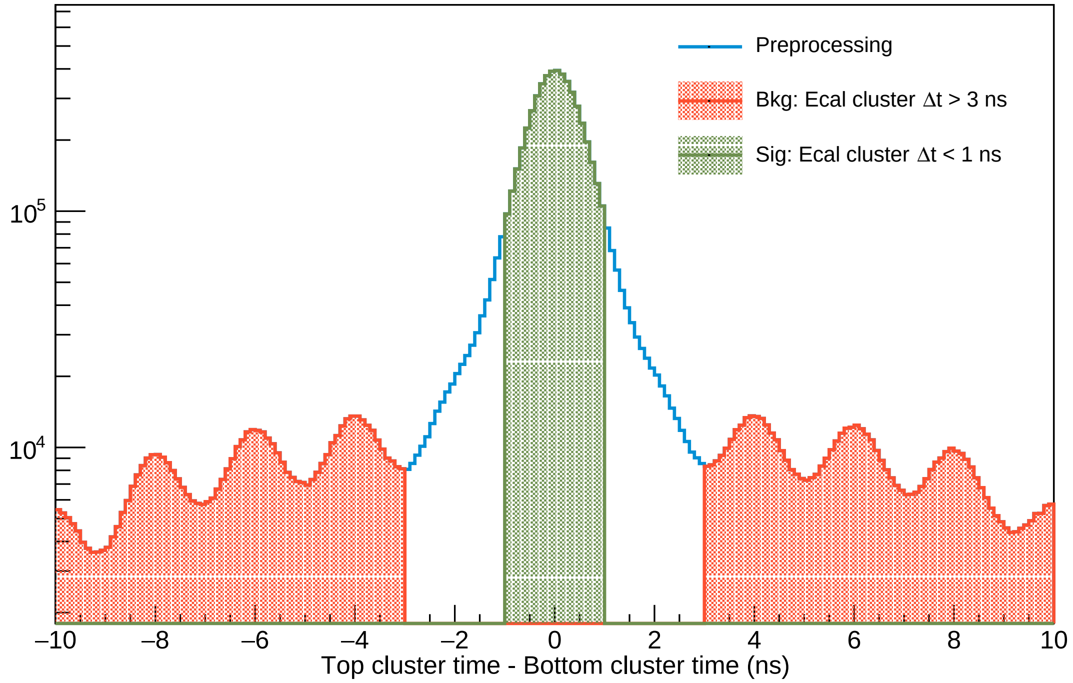

Figure 15: Difference in time between the top and bottom Ecal cluster in an event. The green region defines the signal and the orange region defines the background. These regions were used to optimize the base selection.

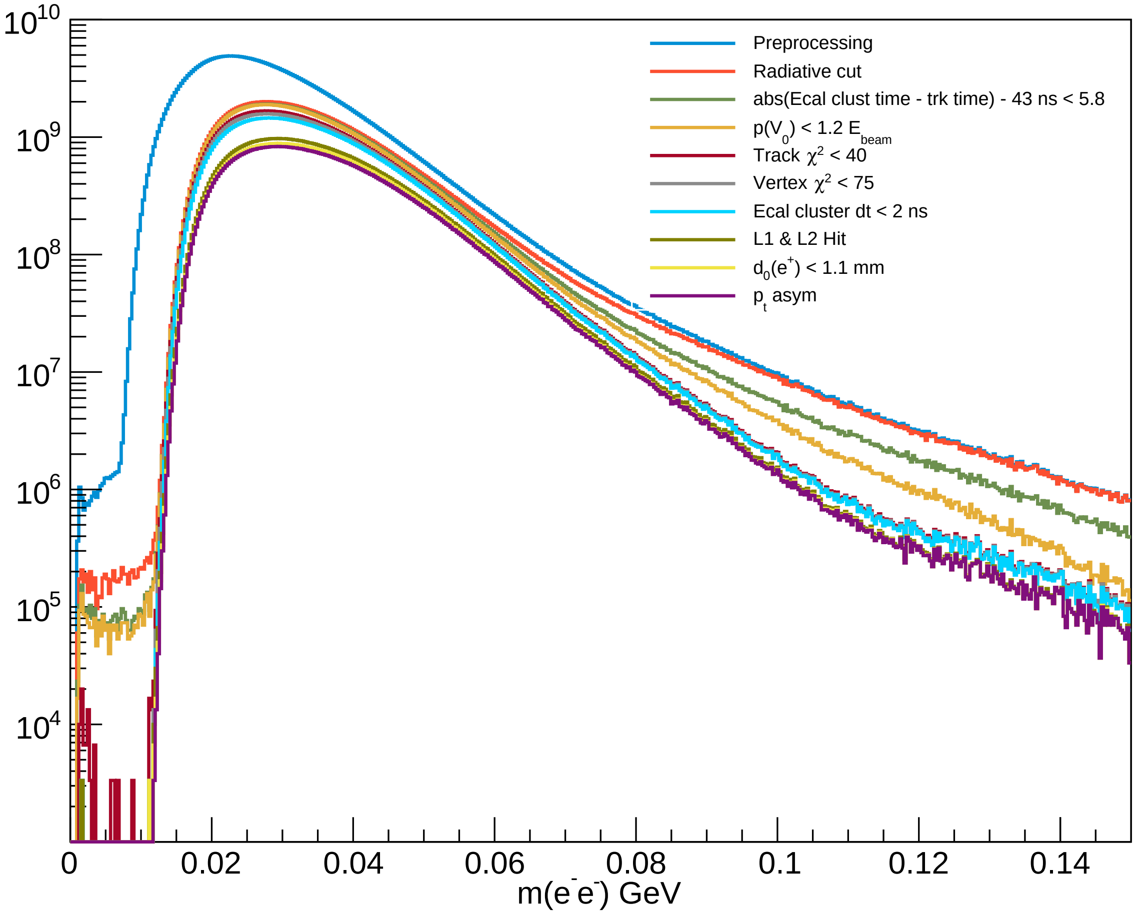

Figure 16: Cut-flow showing the effect of all cuts on the e^+e^- invariant mass distribution. The mass distribution in violet will serve as the starting point for the resonance search.

3 Composition of the e^+e^- Sample

A key component to translating the number of signal events to the coupling epsilon is the fraction of reconstructed events in our sample, after all selection requirements, that come from radiative tridents. This fraction is defined as :

$$f_{rad} = \frac{N_{rad}}{N_{tri} + N_{cWAB}} \tag{6}$$

While the total number of tridents can be taken from data (from the maximum likelihood fit, as a function of mass), the radiative fraction must be computed using Monte Carlo. Therefore it is important that we understand the composition of our data sample, namely, the relative contributions of trident events to converted WAB (cWAB) events .

A dedicated study of the HPS sample composition and comparison of what we see in a data to our Monte Carlo was performed and documented in HPS-NOTE-XXXX [**?**]. The study looked at both γe^- (dominated by WABs) and e^+e^- (both tridents and cWABs) final states. The selection used for the composition study were chosen to be the same as those used in this note, although some of the details of selecting good tracks were somewhat different, particularly when there were multiple tracks using the same hits. There were three primary conclusions from this study:

- it appears that the MC overestimates the overall rates of all MC samples by ~ 0.87
- the SVT hit inefficiency is underestimated in MC
- after correcting for that inefficiency, the data is consistent with MC regarding the ratio of cWABs to tridents as evidenced by a number of distributions which differentiate their contributions

Figure 17 shows the distributions of some kinematic variables for the e^+e^- events from the composition study. The events shown in these distributions have the radiative cut but do not require that the tracks have both L1 and L2 hits. There are a few features in these plots (e.g. skew in track momentum) but generally the agreement between data and MC for these distributions is quite good.

In order to focus on the relative cWAB fraction in this data, we look at some qualities that show differences between cWABs and tridents. First, we remove the positron d0 and pT asymmetry cuts, which are designed to cut out cWABs. In Figure 18, we show the distributions for positron d0 and phi0, and the electron-positron pT asymmetry for events whose positron has an L1 hit (left) and without an L1 hit (right).

The ratio of L1/no L1, another good discriminator of cWABs and tridents, agree well within errors (which is $\sim 3.5\%$). The distributions for the positrons with an L1 hit (dominated by tridents) suggest that the converted WABs are underestimated with respect to the tridents. The difference between this MC/data difference could potentially be from an underestimation of converted WAB events, in some cases by 50% or even more (pT asymmetry). It is not likely that the WAB cross-section is off by so much since we have measured the observed cross-section in data, and the MC actually overestimates that. It is possible, but unlikely, that the conversion probability (which is proportion to the amount of material) is underestimated by a large amount, but it's hard to reconcile that with the rate and shape agreements (which, granted, have a higher

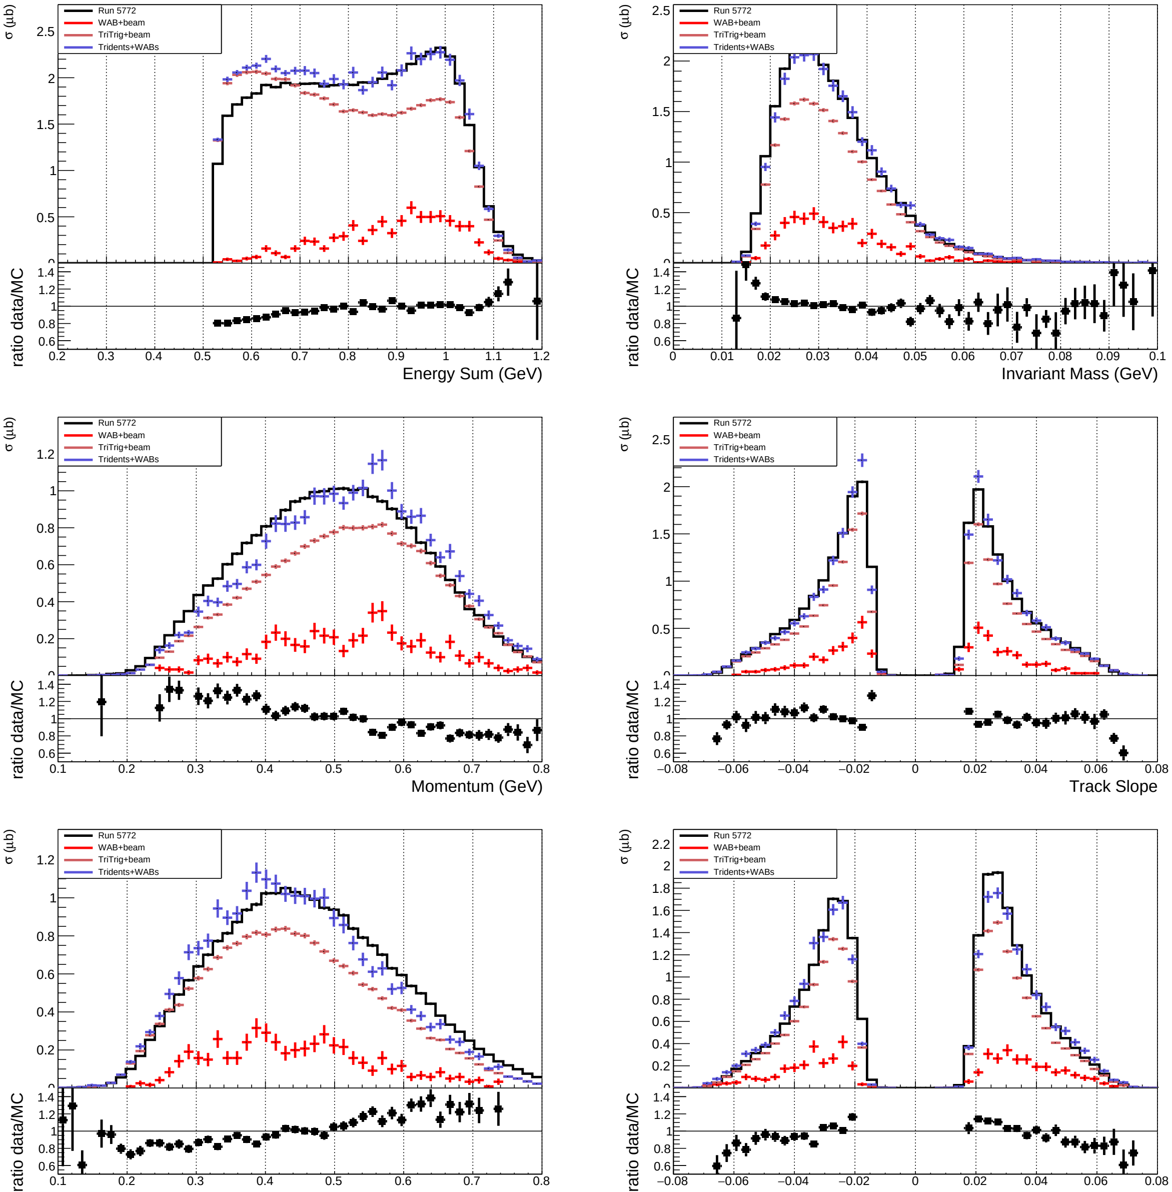

Figure 17: Distributions of e^+e^- events with scaling all cross-sections by 0.87. Clockwise from top left: positron+electron momentum, e^+e^- invariant mass,positron track slope, electron track slope, electron momentum, positron momentum. Each plot shows the raw distribution from data and the various MC samples as well as the ratio of data and the sum of WAB-beam and tritrig-beam.

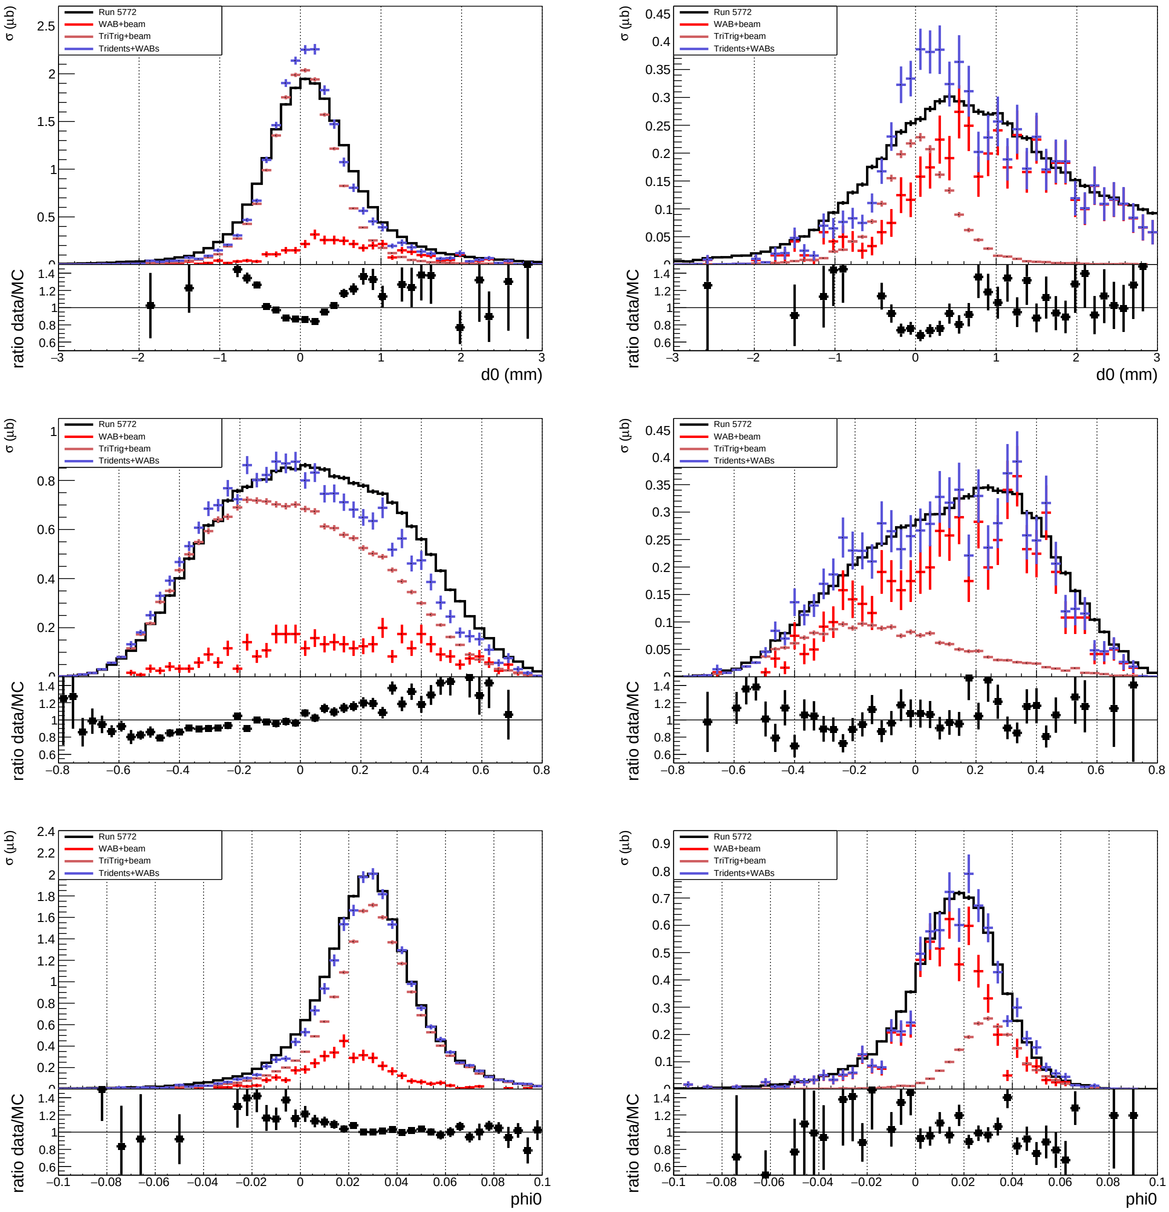

Figure 18: Distributions for e^+e^- events requiring both electron and positron have L1 and L2 hits (left column) and where the positron does not have an L1 hit (right column). Top to bottom: positron d0, pT asymmetry, and positron phi0.

statistical uncertainty). It is most likely that the MC generation and simulation gives slightly skewed shapes compared to data and that is the effect we are seeing here.

The conclusion of this study is that the Monte Carlo does a reasonably good job at estimating the relative fractions of tridents, WABs, and converted WABs. In a later section, we will estimate a systematic uncertainty on the radiative fraction due to the estimated uncertainty in the cWAB-to-trident ratio.

4 Mass Resolution

The heavy photon signal is expected to appear as a Gaussian resonance above the e^+e^- invariant mass spectrum with its width reflecting the mass resolution of the experiment. Determining that resolution experimentally is thus a critical component of the resonance search. In the section that follows, the procedure for determining the mass resolution is described.

4.1 Møller Event Selection

Determination of the mass resolution from data was accomplished by using electron-electron elastic scattering (Møller scattering) which will have a well-defined invariant mass at any given beam energy. For this particular study, only events which satisfy the "singles1" trigger are considered. Additionally, only events passing the SVT quality cuts described in Section 2.1 are included in the final sample.

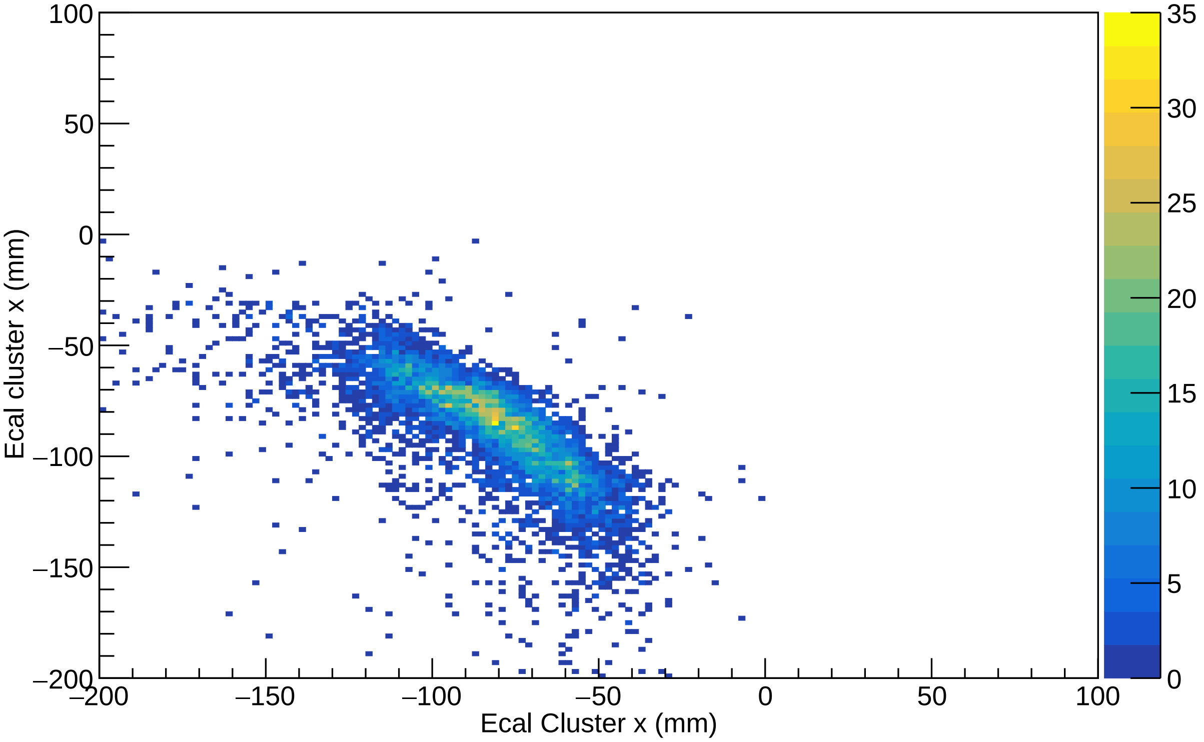

Figure 19: Plot showing the correlation between the x position of Ecal clusters associated with a Møller e^-e^- pair. A Møller pair has a very distinctive signature in x-x space.

Selection of Møller e^-e^- pairs begins by requiring that an event have a single e^- track in each of the SVT volumes. Both tracks are required to loosely match clusters in the Ecal in order to guard against mismatches due to mis-reconstructed tracks. Specifically, a track-cluster match χ^2 of less than 10 has been found to ensure that the distance between the extrapolated track position to the Ecal and the Ecal cluster is reasonable. Furthermore, a track fit χ^2 of less than 40

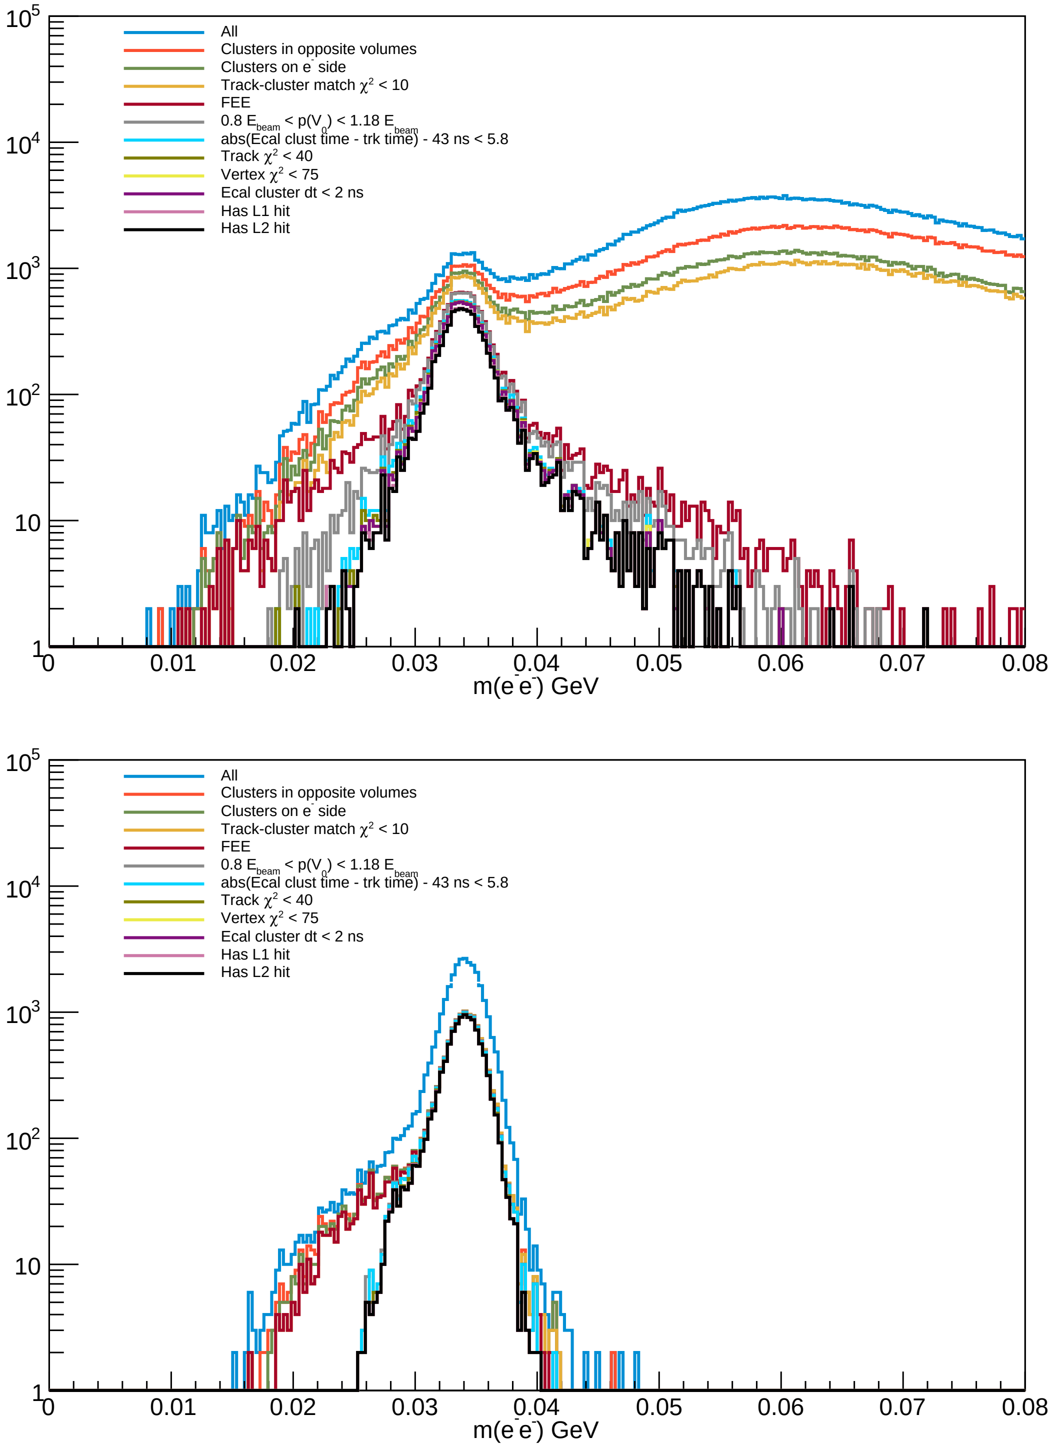

Figure 20: Cut-flow showing the effect of all cuts on the e^-e^- invariant mass distribution in both data (top) and MC (bottom). These cuts are used to isolate the Møller peak (black) which is used to measure the mass resolution.

is enforced to ensure high quality tracks are used. In addition, a track-cluster match is required to satisfy the following condition:

$$\text{abs}(t_{\text{cluster}} - t_{\text{track}} - 43) < 4.5 \text{ ns} \tag{7}$$

The presence of a magnetic field means that the Ecal clusters associated with the Møller electrons will both appear at negative x, i.e. to beam's right (the electron side) in the calorimeter (see Figure 19). With this in mind, the clusters are required to satisfy the following relation:

$$x_{e-\text{ Cluster}} \times x_{e-\text{ Cluster}} >= 0 \tag{8}$$

Because of the track-cluster match requirement, requiring the clusters to be on the electron side also ensures that both tracks are on the electron side.

There are two main sources of background affecting this analysis:

- Accidental pairs where one or both e^- are actually elastically scattered beam electrons.
- e^-e^- pairs where one of the electrons comes from a conversion in or downstream of the target.

To suppress these backgrounds, first, the momentum of each electron track is required to be less than 0.75E_{beam}. Second, the sum of the track momentum of the pair is required to satisfy the following:

$$0.8E_{\text{beam}} < p(e^-e^-) < 1.18E_{\text{beam}} \tag{9}$$

To suppress pairs which include an electron originating before or after the triggering bunch, the clusters associated with the tracks are required to be coincident to within 2 ns. Requiring that both electron tracks have a hit in both layer 1 and layer 2 will eliminate most pairs where an electron comes from a conversion. Finally, if the vertex χ^2 is greater than 75, the e^-e^- pair is disregarded.

The above cuts are summarized in Table 5. The effect of these cuts on both the data and MC e^-e^- invariant mass distribution is shown in Figure 20.

4.2 Mass Resolution

Determination of the mass resolution was accomplished by fitting the core of the final e^-e^- invariant mass distribution (black in Figure 20) to a Crystal Ball function [4]. The resulting fit to both the data and MC Møller peaks is shown in Figure 21. In data, the mass peak is found to be at 33.9 MeV and the mass resolution was measured to be 1.609 MeV. When compared to the expected values, the mass peak is within $\sim 3\%$ of the expected value, however, the mass resolution differs by $\sim 23\%$. The discrepancy in the mass value between the data and MC is accounted for by the fact that the target position assumed in the tweakpass6 reconstruction, $z = 0$, did not agree with that observed in the data. When accounting for this effect, the measured Møller mass was within 1% of the expected value. The discrepancy between the mass resolution seen in data and MC can be attributed to the difference in the data and MC momentum resolution. As shown in Figure 22, the momentum resolutions measured using data and MC differ by 1.13/5.9

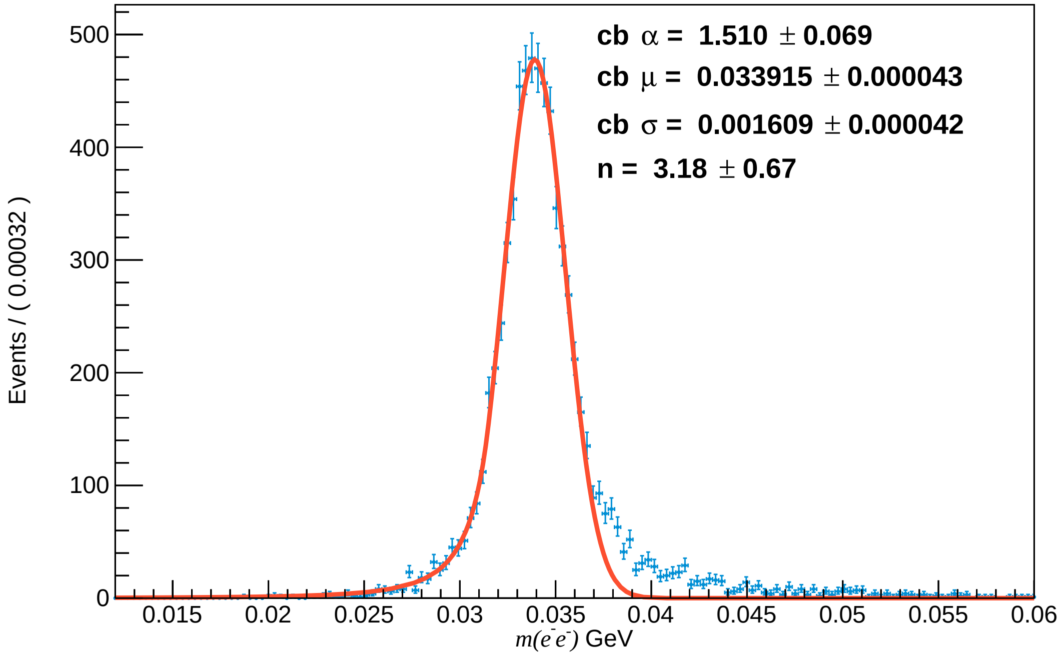

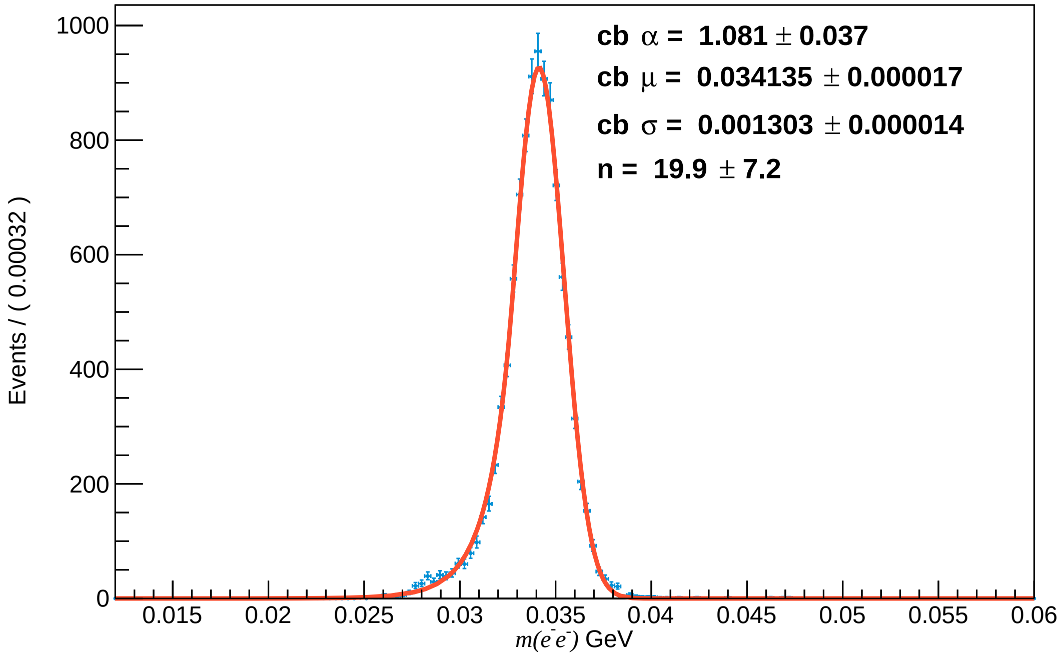

Figure 21: The data (top) and Monte Carlo Møller mass peaks used to measure the mass resolution. Both peaks were fit with a Crystal Ball function and the σ was taken as the mass resolution.

$\sim 20\%$. Now, using the small-angle approximation for the opening angle, the invariant mass and mass resolution is given by

$$m \sim \frac{1}{\sqrt(2)}\theta\sqrt(p_{e^+}p_{e^-}) \tag{10}$$

$$\sigma_m = \frac{1}{\sqrt(2)}(\theta\frac{\sqrt(p_{e^+}p_{e^-})}{2}(\frac{\sigma_{p_{e+}}}{p^{e^+}} + \frac{\sigma_{p_{e-}}}{p^{e^-}} + \sigma_\theta\sqrt(p_{e^+}p_{e^-})) \tag{11}$$

From equation 11, it can be seen that a 20 % increase in momentum resolution results in an increase in the mass resolution of $\sim 20\%$, which would largely account for the mass resolution discrepancy.

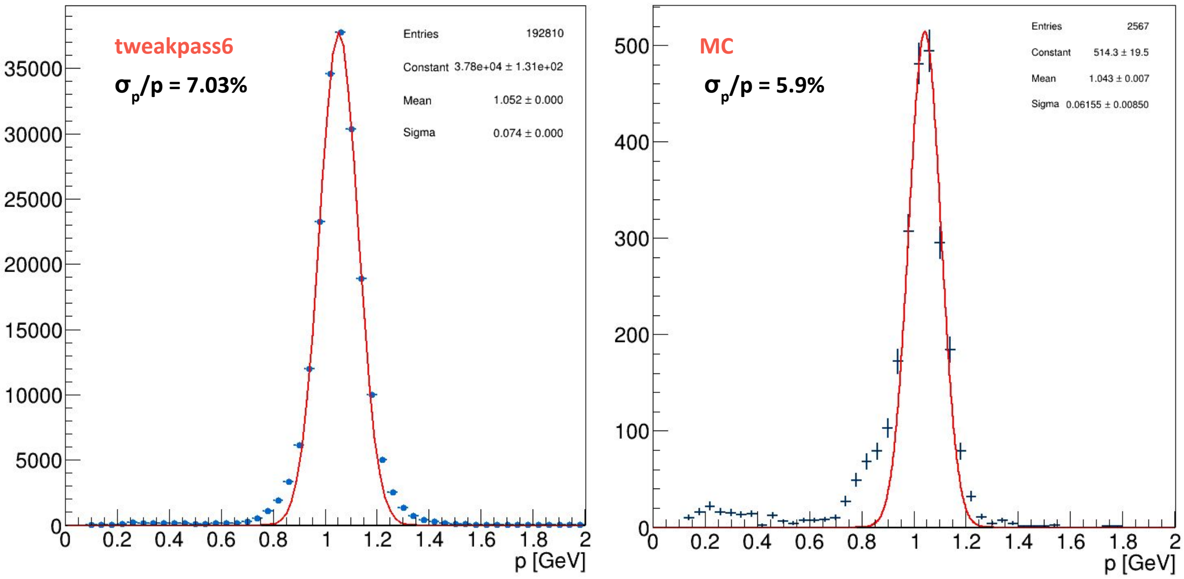

Figure 22: Beam electron momentum peaks for both data (left) and Monte Carlo (right). The core of each of the peaks is fit to a Gaussian and in order to extract the momentum resolution.

Since data only provides a single mass resolution point, characterizing the parameterization of the mass resolution as a function of mass required the use of A' MC. A' samples ranging in mass from 20 MeV to 90 MeV were subject to the selection described in Section 2. This resulted in the isolation of each of the A' peaks. The core of each of the peaks was then fit to a Crystal Ball function (see Figure 23) in order to extract the mass resolution at each mass point. As shown in Figure 24, the resulting mass resolutions as a function of mass is then fit to a line and scaled to match the observed data mass resolution. Since the difference in the mass resolution is due to the momentum resolution difference, scaling the mass resolution parameterization to match data is equivalent to smearing the momentum resolution to match data. The resulting parameterization given by the equation

$$\sigma(m_{e^+e^-}) = 0.0516690134619m_{e^+e^-} - 0.0000895420419565 \tag{12}$$

is used as an input to the resonance search.

Cut Description	**Requirement**
The e^-e^- pair must have been created from vertexing GBL tracks	particle type > 32
Clusters on e^- side	$x_{e-Cluster} \times x_{e-Cluster} >= 0$
Loose track-cluster match	$\chi^2 < 10$
Beam electron cut	$p(e^-) < 0.75E_{\text{beam}}$
Track-cluster match quality	abs$(t_{\text{cluster}} - t_{\text{track}} - 43) < 4.5$ ns
Momentum sum cut	$0.8E_{\text{beam}} < p(e^-e^-) < 1.2E_{\text{beam}}$
Track quality	$\chi^2 < 40$
Vertex quality	$\chi^2 < 75$
Cluster time coincident	< 2 ns
L1 hit and L2 requirement	e^+ is required to have L1 and L2 hit

Table 5: Summary of selection used to isolate Møller pairs.

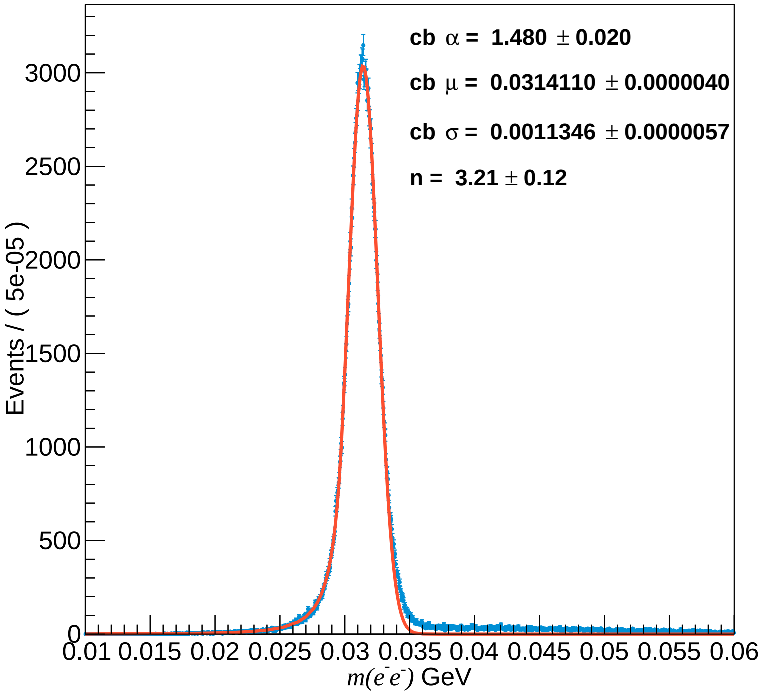

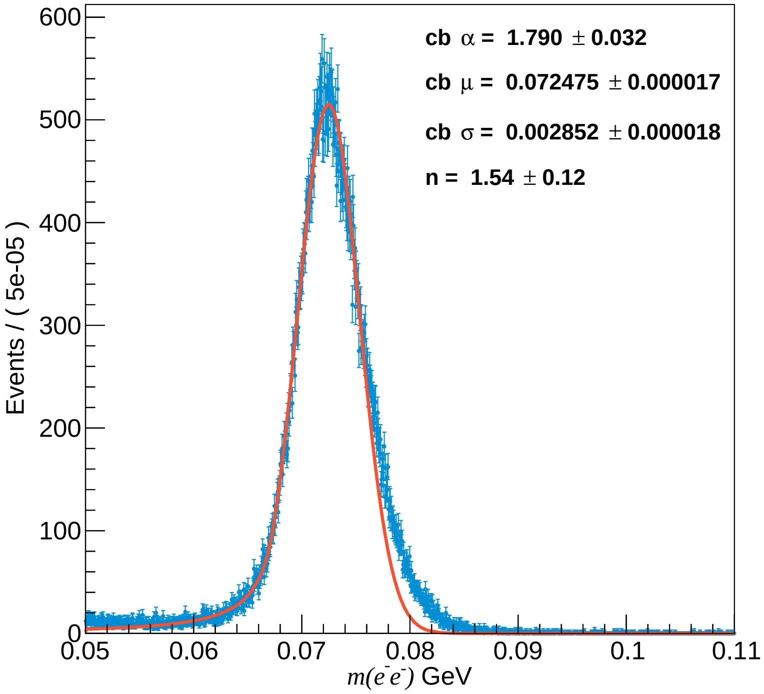

Figure 23: The invariant mass distributions for 30 MeV and 70 MeV A' along with the resulting Crystal Ball fits.

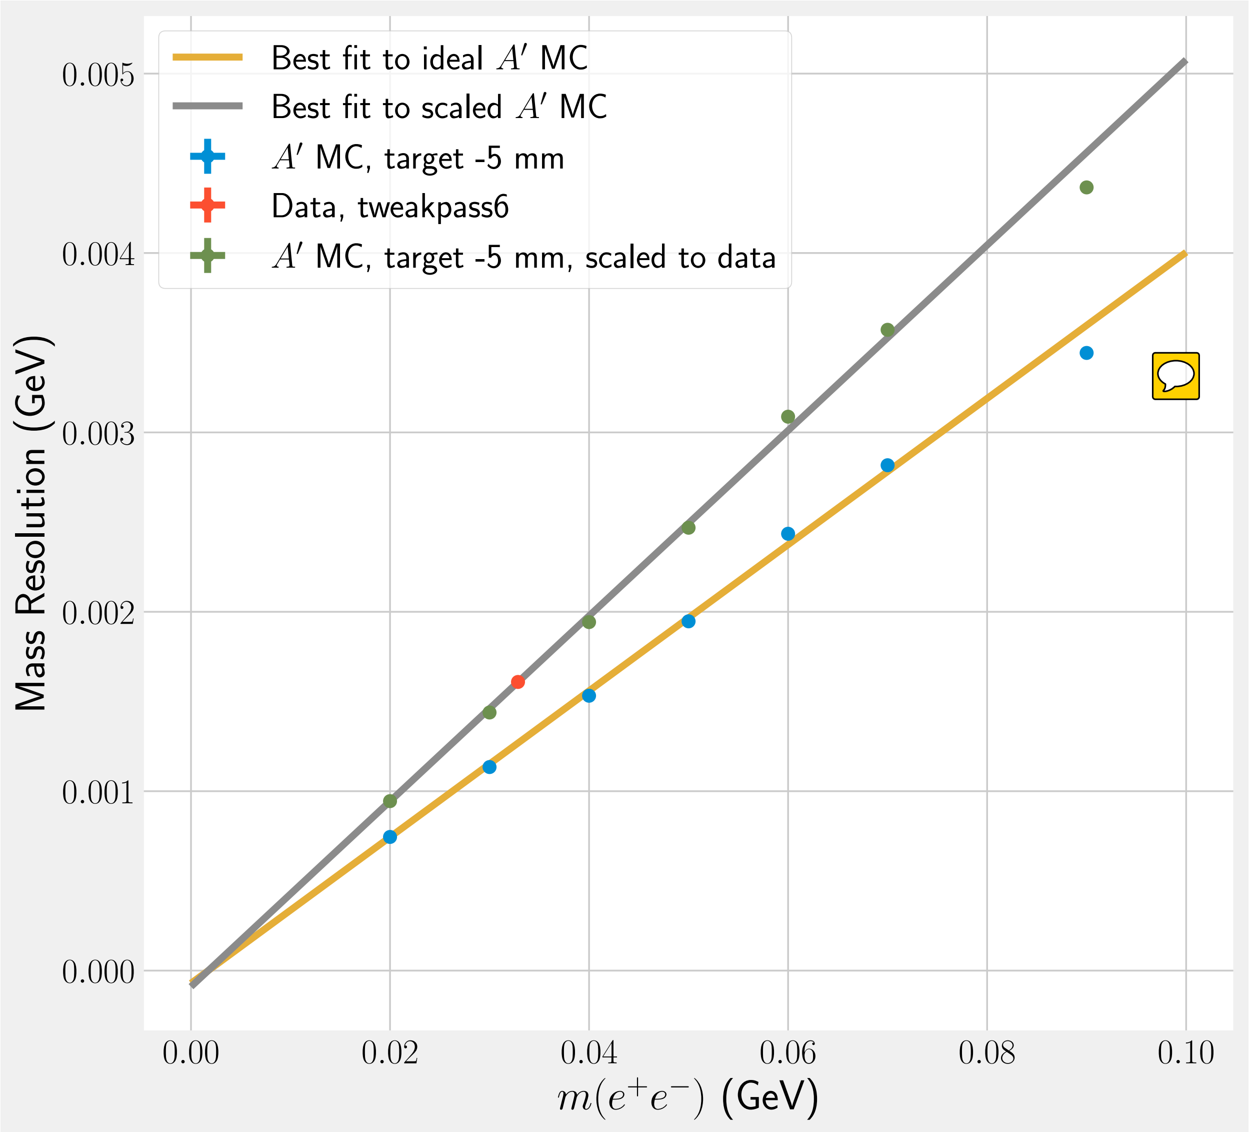

Figure 24: The mass resolution as a function of mass calculated using the invariant mass distribution of A' (blue) and Møller data (orange). The ideal mass resolution calculated using Monte Carlo is scaled to match data (green) and fit to a line (grey). The resulting parameterization is used as an input to the resonance search.

5 Searching for a Resonance

A heavy photon signal is expected to appear as resonance above the e^+e^- invariant mass spectrum. Such a signal is expected to be Gaussian in nature, with a mean equal to the mass $m_{A'}$ of the A' and with a mass dependent width, $\sigma_{m_{A'}}$, given by the mass resolution parameterization defined in 12. With this in mind, the invariant mass distribution measured by HPS (purple distribution in Figure 16) will serve as the starting point for this analysis.

Since the mass of the A' is unknown a priori, the e^+e^- invariant mass spectrum needs to be scanned for any significant peaks. Customarily, a search for a resonance is performed within a window constructed around the mass hypothesis of interest. Within the window, the distribution of A' signal events is modeled using the probability distribution function

$$P(m_{e^+e^-}) = \mu \cdot \phi(m_{e^+e^-}|m_{A'}, \sigma_{m_{A'}}) + B \cdot p(m_{e^+e^-}|\mathbf{t}) \tag{13}$$

where $m_{e^+e^-}$ is the e^+e^- invariant mass, μ is the signal yield, B is the number of background events within the window, $\phi(m_{e^+e^-}|m_{A'}, \sigma_{m_{A'}})$ is a Gaussian probability distribution describing the signal and $p(m_{e^+e^-}|\mathbf{t})$ is a Chebyshev polynomial of the first kind with coefficients $\mathbf{t} = (t_1, ...t_j)$ that is used to describe the background shape. In this instance, a 7th order Chebyshev polynomial was found to best describe the background (see Section 6.3). Furthermore, $m_{A'}$ and $\sigma_{m_{A'}}$ are constant and set to the A' mass hypothesis and expected experimental mass resolution, respectively. Estimating the signal yield as well the background normalization and shape within a window can be done by the method of maximum likelihood. The theoretical formalism used to do this will be outlined here but a detailed discussion can be found in [2].

5.1 Maximum Likelihood Fit

Assume the events within the window are binned as $\mathbf{n} = (n_1, ...n_i)$. Furthermore, assume the center of the ith bin is given by b_i and has a width equal to ϵ. The expected number of events of the ith bin is given by

$$E[n_i] = S_i + B_i \tag{14}$$

where

$$S_i = \mu \int_{b_i-\epsilon/2}^{b_i+\epsilon/2} \phi(m_{e^+e^-}|m_{A'}, \sigma_{m_{A'}}) d(m_{e^+e^-}) \tag{15}$$

$$B_i = B \int_{b_i-\epsilon/2}^{b_i+\epsilon/2} p(m_{e^+e^-}|t_j) d(m_{e^+e^-}). \tag{16}$$

Denoting the parameters that are not of immediate interest, i.e. the nuisance parameters, by $\theta = (B, \mathbf{t})$, an estimate of μ and θ can be obtained by finding the parameters $\hat{\mu}$ and $\hat{\theta}$ that maximize the Poisson likelihood function, $\mathcal{L}$

$$\mathcal{L}(\mu, \theta) = \prod_{k=1}^{n_{\text{bins}}} \frac{(S_k + B_k)^{n_k}}{n_k!} e^{-(S_k+B_k)} \tag{17}$$

where the sum is over all bins within the window, n_{bins}. In the case where the invariant mass is scanned for a resonance, the Poisson likelihood function is maximized within the window constructed around each A' mass hypothesis. This yields estimators for the signal yield and nuisance parameters at each A' mass hypothesis which are used to determine if a significant resonance was found.

5.2 Likelihood Ratio

When searching for a resonance above a background distribution, it is necessary to discriminate between two scenarios:

- The background only or null hypothesis, $H_0 : \mu = 0$.
- The signal+background hypothesis or alternative, $H_1 : \mu > 0$.

Establishing whether the signal+background model is significantly different from the background only model is typically done using the profile likelihood ratio

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}})}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \tag{18}$$

where $\hat{\hat{\theta}}$ is the conditional estimator for the nuisance parameters obtained by maximizing the Poisson likelihood assuming that the null or background only hypothesis is true i.e. $\mu = 0$. The unconditional estimators $\hat{\mu}$ and $\hat{\theta}$ are obtained by maximizing the Poisson likelihood without any constraints on μ. As can be seen from 18, if the estimator of the signal yield, $\hat{\mu}$, is compatible (incompatible) with the hypothesized μ, the likelihood ratio will tend to 1 (0).

A more convenient test statistic is the log likelihood ratio defined as

$$q_0 = \begin{cases} -2\ln\frac{\mathcal{L}(0,\hat{\hat{\theta}})}{\mathcal{L}(\hat{\mu},\hat{\theta})} & \hat{\mu} > 0 \\ 0 & \hat{\mu} < 0. \end{cases} \tag{19}$$

In the large sample limit, the test statistic q_0 can be shown to follow a $1/2\chi^2$ distribution defined in [2] as

$$f(q_0|0) = \frac{1}{2}\left(\delta(q_0) + \frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{q_0}}e^{-q_0/2}\right) \tag{20}$$

where the first term on the right side of the equation is a delta function at 0 and the second term is a χ^2 distribution with one degree of freedom.

Quantifying how extreme the observation is can be done by calculating a p-value as

$$p = \int_{q_{0,obs}}^{\infty} f(q_0|0)dq_0. \tag{21}$$

This is shown graphically in Figure 25. Typically, the observed p-value is compared against a significance level α. The significance level denotes the probability of incorrectly rejecting the null hypothesis in favor of the alternative (type-I error). In other words, it denotes the probability of there being a statistical fluctuation in the background large enough to mimic a signal. If a p-value is found to be less than α, the measurement is claimed to be significant. Typically, in particle physics, an α on the order of 3×10^{-7} (5σ) is required to claim discovery of new phenomena. This means that there is a 1 in about 3.5 million chance that the observation is due to a fluctuation in the background.

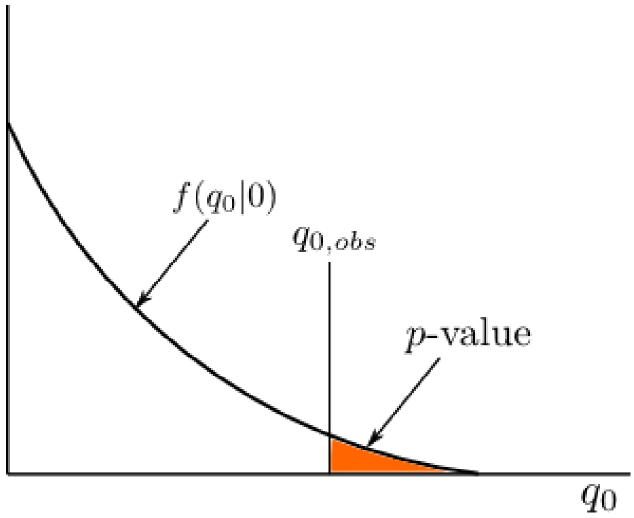

Figure 25: Graphical representation of a p-value.

5.3 The Look-Elsewhere Effect

As discussed previously in Section 5.2, a result is determined significant if the p-value is smaller than some pre-determined threshold, α. However, when performing multiple tests, as is the case when scanning a mass distribution for a resonance, an observation with a p-value that is as extreme as α is bound to occur at a rate of $n \times \alpha$ where n is the number of measurements. This phenomenon is known as the "Look-Elsewhere Effect" (LEE) and needs to be taken into account through a correction to the "local" p-value observed at each mass hypothesis.

Assuming that only a single heavy photon can be observed within the HPS invariant mass distribution, the correction can be estimated using a large number of pseudo-data sets and generating the distribution $f(q_{0,max}|0)$ composed of the largest q_0 (i.e. smallest p-value) from each of the invariant mass scans. However, generating a distribution of $f(q_{0,max}|0)$ that would allow an estimation of a "global" p-value (i.e. local p-value after correction) down to the level of 5σ with any accuracy would require running $> 10^6$ pseudo experiments. Generating so many pseudo-data sets is often not feasible within a reasonable amount of time.

Instead, the smallest p-values obtained from a series of resonance searches on 10,000 pseudo data sets were ranked and the corresponding quantile was calculated [5]. A mapping from a local p-value to a global p-value is then created. The mapping created for this analysis is shown in Figure 26. As can be seen from the figure, a local p-value equal to 0.05 corresponds to a global p-value of $\sim$ 0.5.

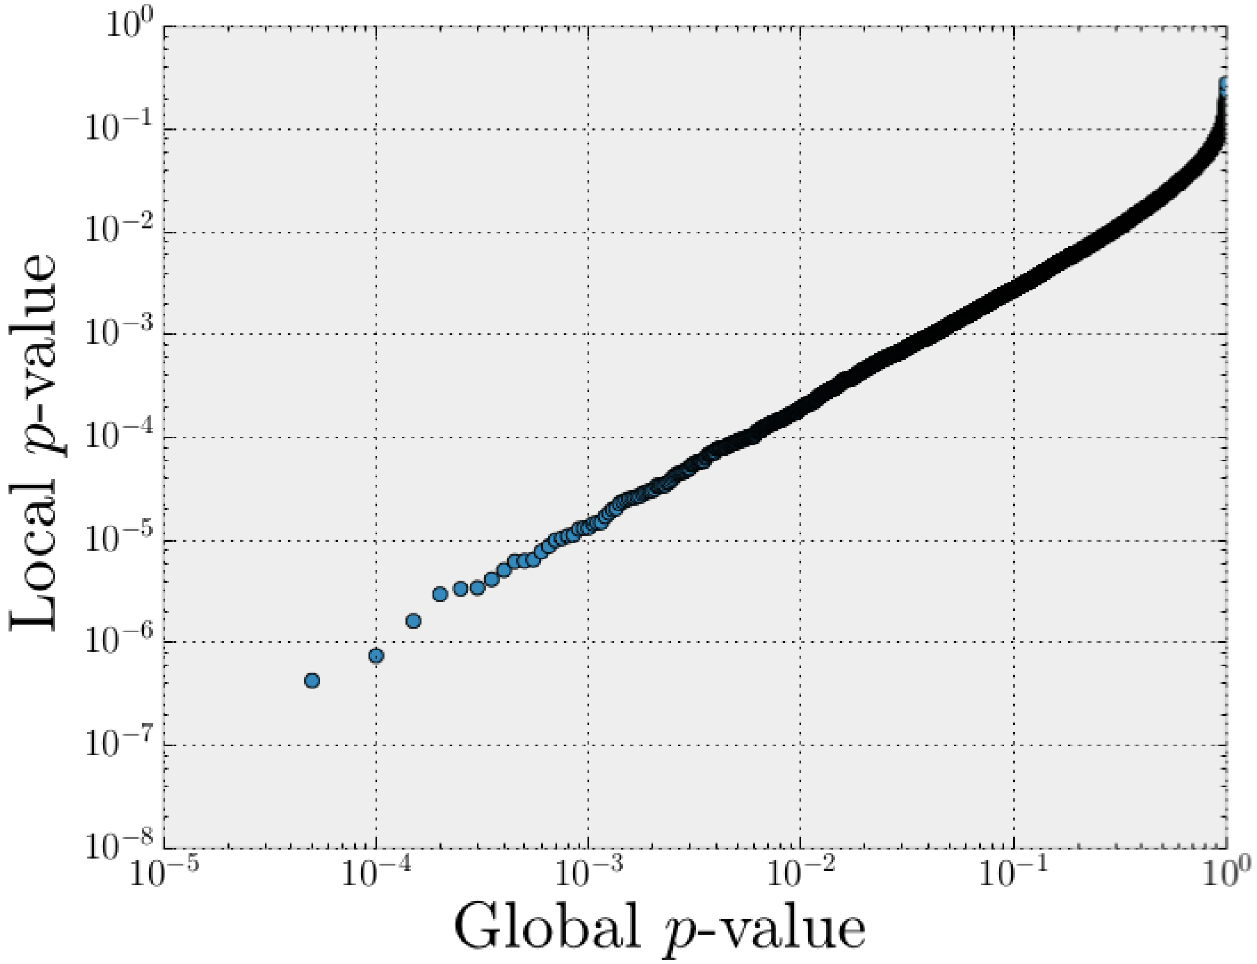

Figure 26: Mapping between local and global p-values.

6 Fit Optimization

Given that this is a blind analysis, optimization of fit parameters was done using the unblinded portion (∼10%) of the 2015 engineering run dataset. After applying the selection described in Section 2 to the unblinded dataset, the resulting e^+e^- spectrum was fit, and pseudo data sets generated using the parameters of the fit. A resonance search was then performed on each of the pseudo datasets and several key parameters were varied. These parameters included the size of the fit window, the binning of the e^+e^- mass spectrum, and the order of the polynomial used to model the background. The optimal fit parameters were chosen such that both the signal pull (bias), defined as

$$\text{pull} = \frac{\mu_{\text{fit}} - \mu_{\text{inserted}}}{\mu_{\text{fit error}}}, \tag{22}$$

and the signal upper limit were minimized. The following section will discuss this procedure in detail.

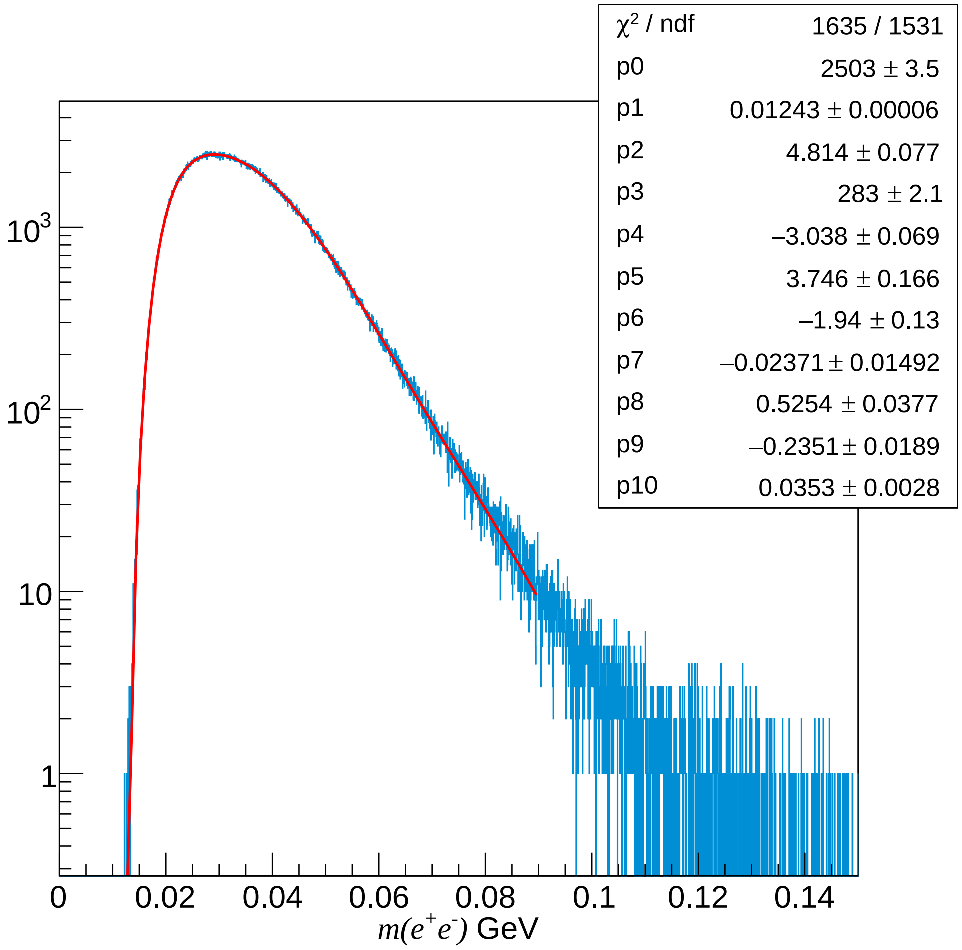

Figure 27: Global fit to the e^+e^- invariant mass distribution generated using 10% of the engineering run data. The fit is used to generate the toys used to evaluate the fitting procedure.

6.1 Pseudo Data Sets

Optimizing the fit parameters and understanding the systematics associated with the fit procedure required the use of pseudo data sets. The pseudo data sets were generated from a fit to the e^+e^- invariant mass distribution obtained from the unblinded 10% of the engineering run data. Specifically, the observed mass distribution was fit between 0.01-0.09 GeV using the function

$$p_0 \left(\frac{m - p_1}{0.03 - p_1}\right)^{p_2} e^{-p_3(m+0.03)} \times \left[\left(1 + \sum_{i=1}^{7} p_{i+3} \left(\frac{m}{0.03}\right)^{i}\right) / \left(1 + \sum_{i=1}^{7} p_{i+3}\right)\right]. \tag{23}$$

The resulting fit is shown in Figure 27. This fit function was then sampled between 0.014-0.115 GeV and used to generate 2000 pseudo experiments each with a number of events (1,297,890) equal to that in the data. The optimizations discussed in this section used these pseudo data distributions.

6.2 Mass Binning

Ideally, an unbinned maximum likelihood fit should be used to estimate the mass and significance of any potential signal(s) in the mass distribution, as described in Section 5. However, given the large statistics of the final e^+e^- sample, it's not possible to do an unbinned maximum likelihood fit in a reasonable amount of time. Instead, a binned likelihood fit is performed. This made it necessary to understand how the bin size of the e^+e^- mass distribution impacts the fit and its systematics.

Understanding the fit systematic associated with the bin size was studied using 2000 pseudo data sets by trying three different bin sizes, 0.2 MeV, 0.1 MeV and 0.05 MeV. At the time of this study, the final fit parameters had not been fully optimized, so a 7th order polynomial and a window size of $13\sigma_{mass}$ were chosen. They produced reasonable results during initial tests of the fitter. Using these parameters, a resonance search was performed using all the pseudo datasets and the signal pull spectra were generated. Examples of these distributions resulting from a search for a resonance at a mass of 33.1 MeV are shown in Figure 28. Ideally, the mean and standard deviation of these distributions should be 0 and 1 respectively. However, as can be seen from Figure 29, there is always a signal bias which get worse at low mass when using larger bins. As a result, a bin size of 0.01 MeV was chosen since it minimizes the signal bias across the whole mass spectrum.

6.3 Fit Window and Polynomial Order

The pseudo data sets were also used to determine the optimal fit window and polynomial order that best models the background. First, 2000 pseudo data sets were generated and binned using a bin size of 0.01 MeV. Next, a resonance search between 20-90 MeV was repeated on all these data sets using polynomials of 3rd, 5th and 7th order and varying window sizes, where a window is defined as

$$n \times \sigma_{m_{A'}(m_{e^+e^-})} \; n = 9, 11, 13, 15, 17. \tag{24}$$

Only odd order polynomials were considered to avoid having a turning point in the middle of the window that may artificially enhance a signal. For each of the combinations of polynomial and

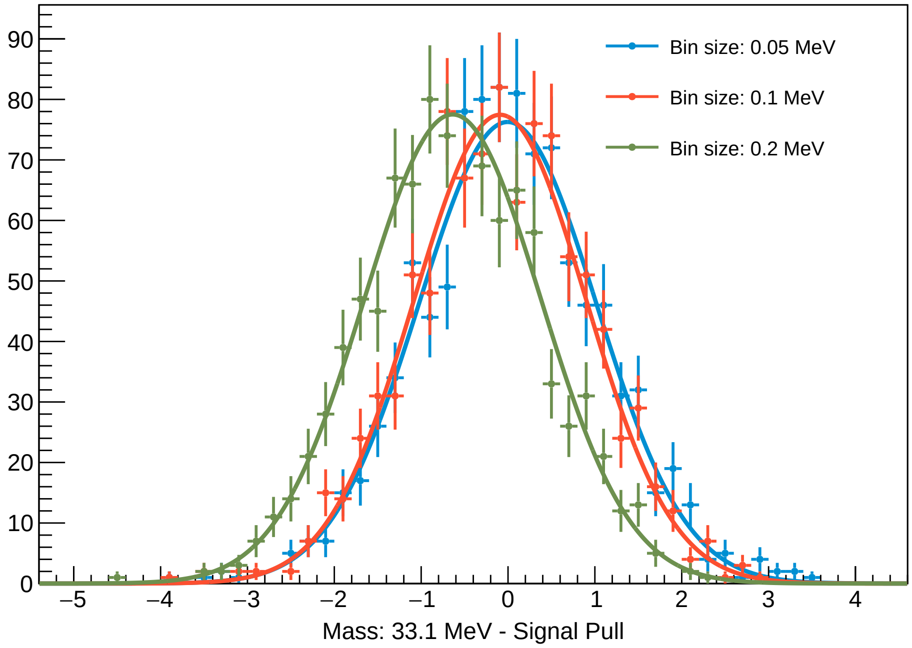

Figure 28: Distribution of signal pulls resulting from fits to toys with no signal injected. Ideally, the mean and standard deviation of these distributions should be 0 and 1 respectively.

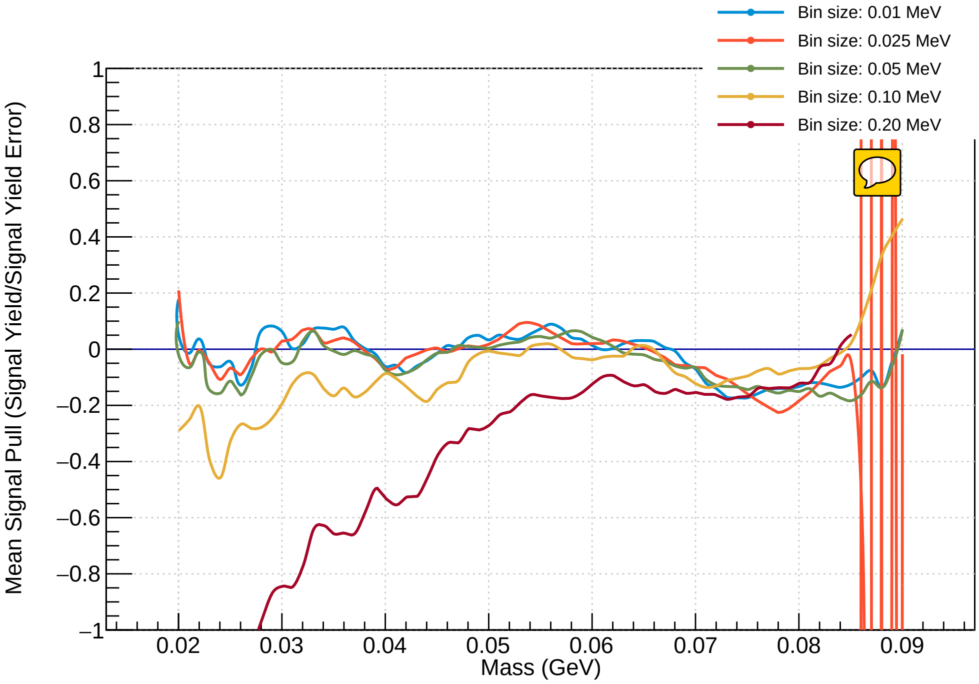

Figure 29: Signal pulls at each mass hypothesis considered during a scan for a resonance. Each point represents the mean of the distribution of pulls at each mass point.

window size, the pulls and upper limits were calculated and used as metrics in determining the optimal combination. Specifically, a combination that minimizes the signal bias while maintaining a reasonable upper limits is desired. As can be seen from Figures 30 and 31, it was found that using a 7th order polynomial with a window size of $11\sigma_{m_{A'}}$ below 70 MeV and a window size of $9\sigma_{m_{A'}}$ from 70-90 MeV achieves the desired results.

6.4 Injection Tests

To study the ability of the fitting methods to find a real signal, pseudo data sets with a variety of known signals were generated and fit to extract the signal. The signals were modeled using Gaussian distributions with centroids at a given mass and σ equal the measured mass resolution (see Section 4). The resulting distributions were then fit with the same polynomial background and fit range described in Section 6.3, and the mass hypothesis was scanned across the relevant mass range in search of the maximum signal yield.

For each combination of injected signal mass and significance, 50 toy experiments were generated from the same mother distributions but with different random number seeds. The ratio of the fitted and generated signal yield as a function of the generated signal significance is shown in Figure 32 for a series of masses. At high significance, the fit returns a signal compatible with but marginally higher than the generated signal, except at the very highest mass, 65 MeV, where it is 20 $\%$ high. For signals near the limit of sensitivity, $< 5\sigma$, the fit overestimates the signal by upwards of 40 $\%$.

This procedure was also tested with toys generated from a simpler background function without the polynomial terms in the square bracket in Equation 23. While not describing the shape of the real data as well, this avoids the possibility of any small wiggles in the parent distributions influencing the result. The results were very similar to the full 11-parameter fit.

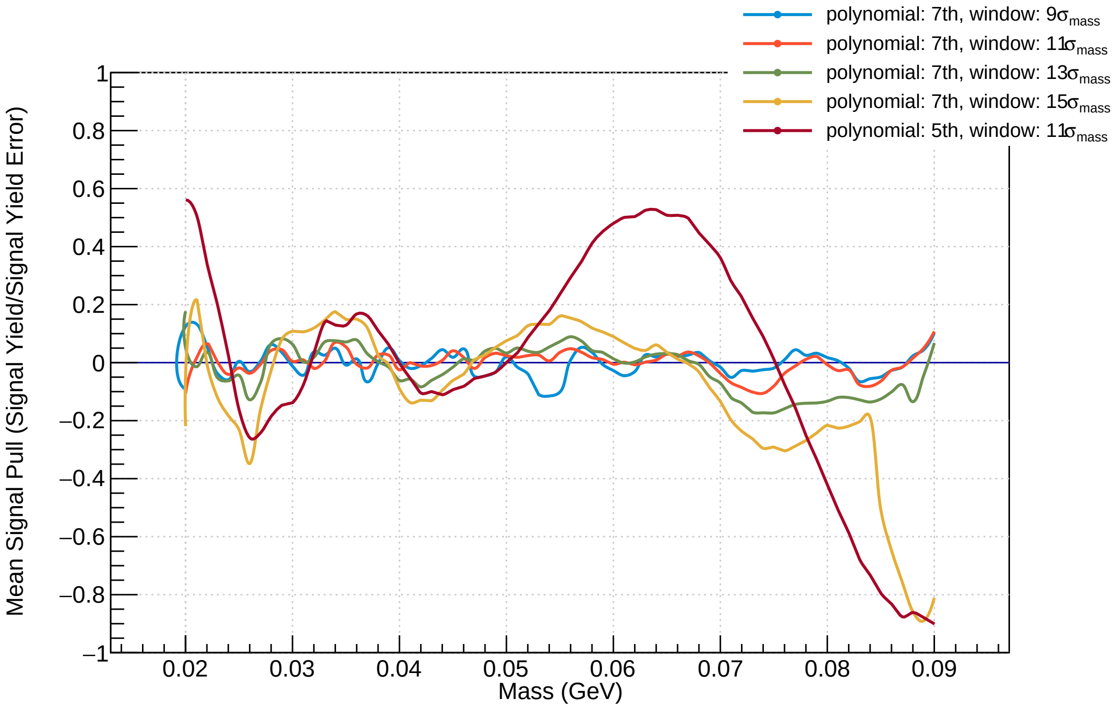

Figure 30: Signal pulls obtained from resonance scans over pseudo data sets using different window sizes and polynomials.

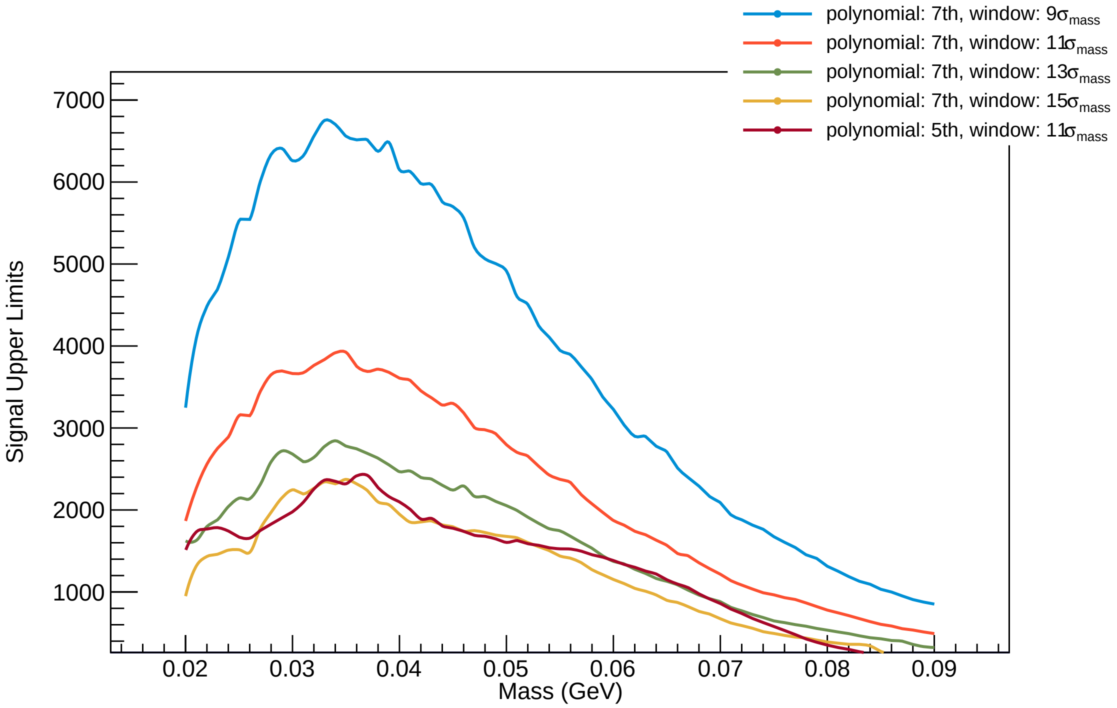

Figure 31: Median signal upper limit at each mass obtained from resonance scans over pseudo data sets using different window sizes and polynomials.

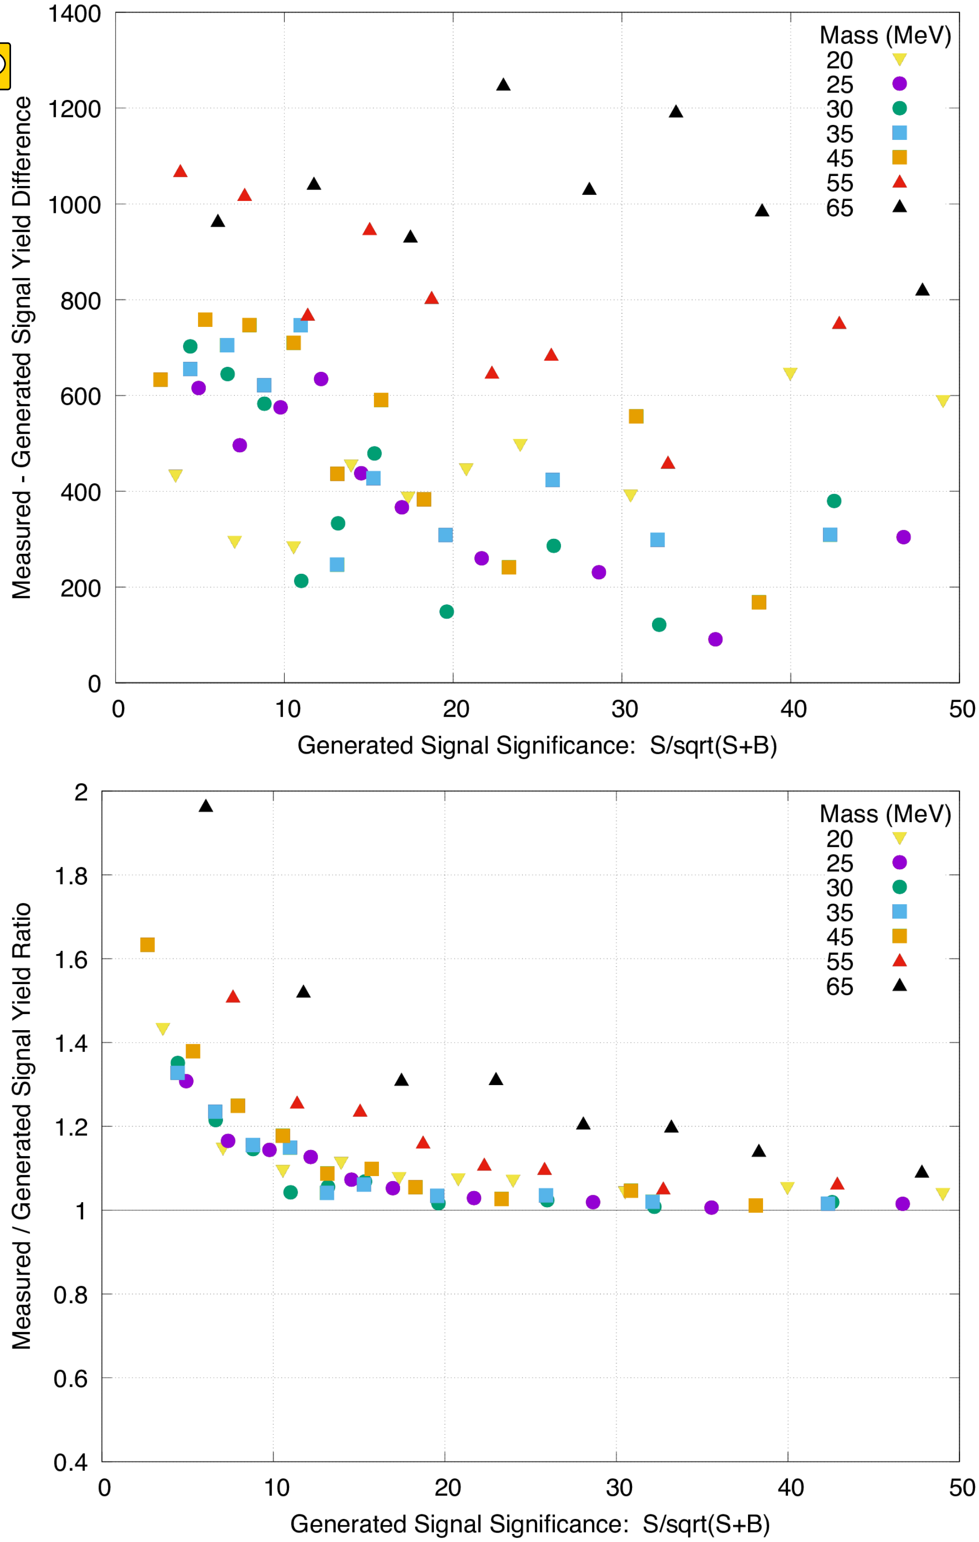

Figure 32: Difference and ratio between the injected and extracted signal yields for different signal masses and significances.

7 Systematic Errors

There are two types of systematic errors that need to be taken into account in the final result: the systematics that impact the signal upper limit and those that effect the scaling of the upper limit to a limit on the coupling. Systematic errors that effect the signal yields impact the sensitivity to a given signal and need to be accounted for in the discovery criteria. These errors include:

- The uncertainty on the mass resolution
- Intrinsic bias of the fits (i.e. pulls)

Furthermore, systematic errors that effect the scaling of the signal to a value of epsilon include:

- The uncertainty in the radiative/trident fraction after all cuts
- The uncertainty in the wab/trident fraction after all cuts

The section that follows will discuss the details of how these uncertainties were evaluated and their impact on the final result.

7.1 Radiative Fraction

The uncertainty on the radiative fraction impacts the translation from the signal upper limit to the associated coupling strength, ϵ. As seen in Equation 6, the radiative fraction denominator is the sum of the full-diagram trident and converted WAB events that pass our final selection. This sum is determined directly from the final data sample as a function of mass.

Rewriting Equation 6 to a more useful form from the perspective of systematic errors specification:

$$f_{rad} = \frac{\frac{N_{rad}}{N_{tri}}}{1 + \frac{N_{cWAB}}{N_{tri}}} \tag{25}$$

we see that the important quantities are the radiative-to-full trident ratio and the cWAB-to-full trident ratio. Both of these quantities must be derived from Monte Carlo. The ratio of cwabs/tridents accordingly has a systematic error associated with it.

The cWAB-to-full trident ratio in MC was checked against the data as described in Section 3. In that section, we concluded that the data and MC are in reasonable agreement for the cWAB/trident ratio however with fairly large uncertainties. Fortunately, the magnitude of the cWAB/trident is quite small so even being very conservative, here taking a 50$\%$ uncertainty on the cWAB rate, the corresponding systematic uncertainty on the radiative fraction, f_{rad} is only $\sim 7\%$. This is what is taken as the systematic associated with the sample composition. As shown in Figure 33, taking into account the sample composition systematic yields a radiative fraction of 8.5%.

7.2 Mass Resolution

Given that the mass resolution is a crucial component of the resonance search, it was necessary to understand any sources of uncertainty associated with it and their impact on the final measurement. The major sources of uncertainty include

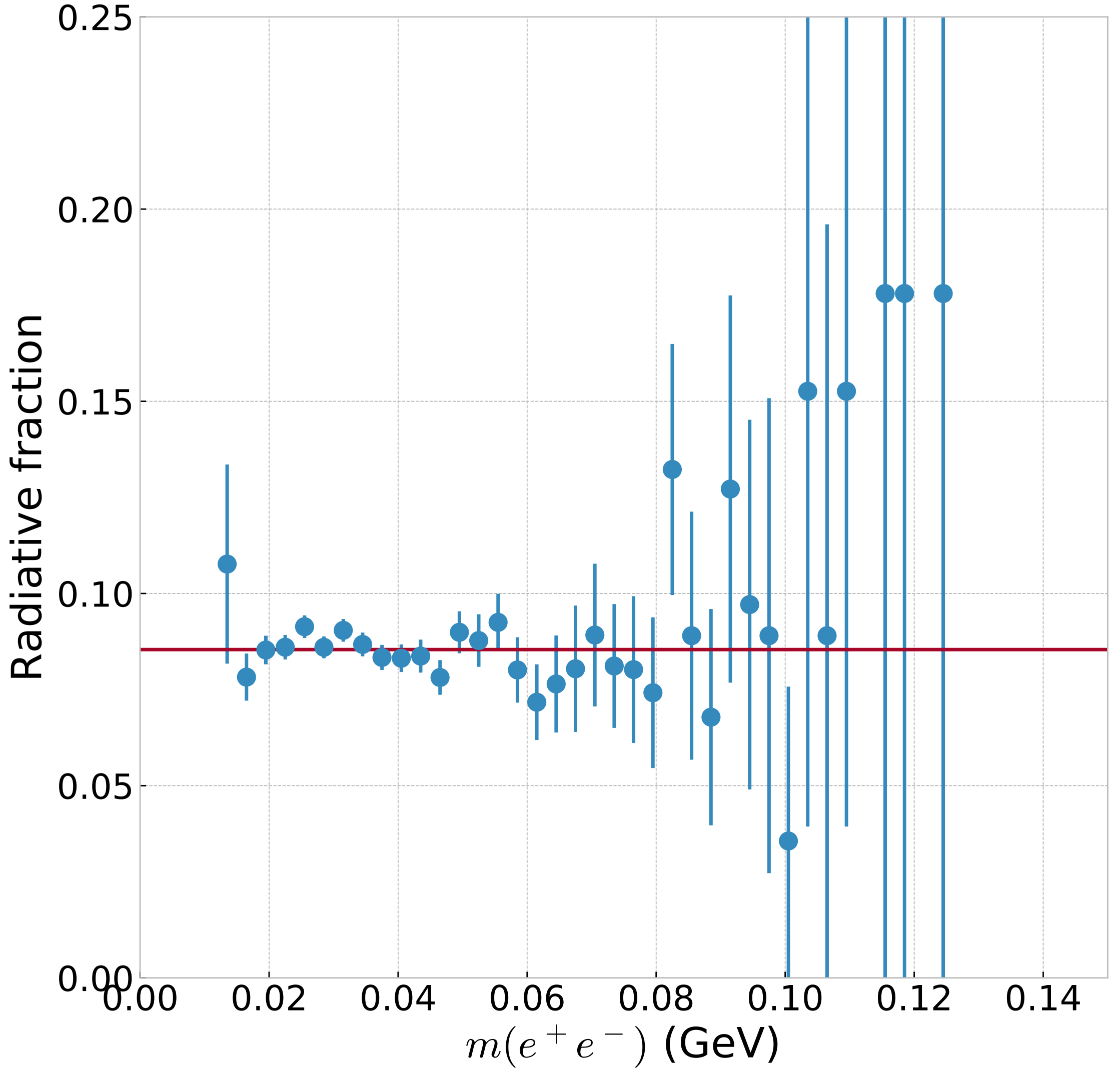

Figure 33: The ratio of the pure radiative cross-section to the full trident plus wab cross-section as a function of mass.

- Error from the fit to the Møller invariant mass peak
- Position of the target

Those systematics that directly impact the mass resolution are combined in quadrature and used to smear the mass resolution parameterization given by Equation 12. The systematic impacting the mass scale is used to shift the mass hypothesis that serves as an input to calculating the mass resolution at a given mass. Furthermore, once the resonance search is complete, the mass scale itself is completely shifted to account for the systematic. This procedure allows the systematics to be taken into account directly in the likelihood used to set an upper limit, thereby minimizing any residual systematic uncertainty.

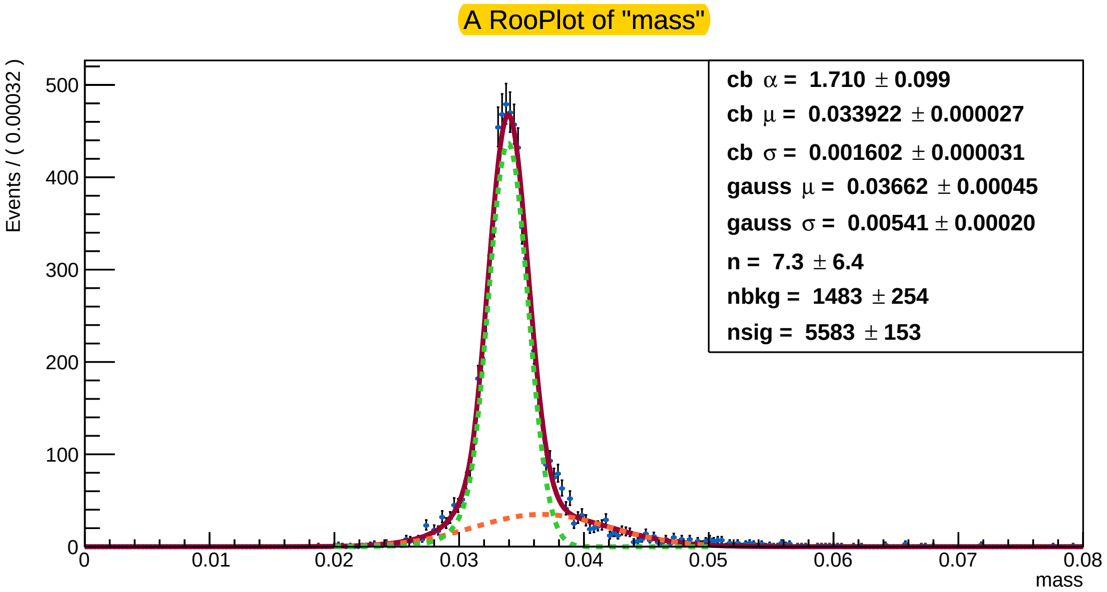

Figure 34: Fit to the Møller invariant mass spectrum that takes into account the background. The spectrum is fit with a Crystal Ball plus a Gaussian on the high side. Taking into account in the background impacts that mass resolution by less than 0.5%.

There are a couple sources of uncertainty in the mass resolution related to the fit of the Møller invariant mass spectrum. As described in Section 4, the core of the Møller mass distribution is fit using a Crystal Ball. One source of uncertainty comes from not taking into account the high side tail background. However, as shown in Figure 34, even after taking into account the background, the mass resolution improves by less than 0.5%, so this won't have much of an impact on the upper limits. In fact, the largest uncertainty come from error on the fit itself, which can be seen from Figure 21 to be 2.6%.

The invariant mass used to measure the mass resolution is calculated using the momentum constrained to the target position. However, knowledge of the target position is not exact and, in fact, is only accurate to within 500μm. As a result, the impact on the mass resolution of constraining the momentum to a position that is $\pm$500 μm off from the true vertex position needs to be evaluated. To do this, first A' MC with the A' vertex positions at $\pm$500 μm was generated and the invariant mass was calculated using the momentum constrained to 0 mm. All events were then subject to the selection described in Section 2. Those that pass are used to extract the mass resolution at each mass point via a Crystal Ball fit to the core of each distribution. The resulting mass resolutions of heavy photons originating at 0 mm and $\pm$500 μm along with best fit line to each of the datasets are shown in Figure 35. As can be seen from the figure, the largest difference in the mass resolution between events constrained to the correct position and those off by 0.5 mm is 1.5%.

As discussed above, the momentum used to calculate the invariant mass is constrained to the position of the target which was assumed to be 0 mm. However, in data, the unconstrained position along z of e^+e^- pairs was found to be at $\sim$-5 mm. As a result, the momentum that is being used to calculate the invariant mass is being constrained to the wrong position, resulting in a shift in the mass scale. The size of the shift needs to be calculated and used to find the correct mass resolution at a given point. In order to estimate the shift, A' MC with masses in the range of 20-90 MeV and vertex positions at -5 mm were generated. The invariant mass at each mass point was calculated using the momentum constrained to 0 mm and the resulting distribution was fit to a Crystal Ball function. The means from the fit to each distribution were then compared to the true mass at each point. Figure 36 shows the percent change at each of the points considered. As can be seen from the figure, constraining to the wrong position results in an upward mass shift as high as 5% at low mass and 1% at high mass. This relationship can be parameterized using a 3rd order polynomial as follows

$$\text{offset} = -11989.232m_{e^+e^-} + 1501.96798m_{e^+e^-} - 83.89m_{e^+e^-} + 6.232 \tag{26}$$

Equation 26 is used to shift the mass used to calculate the mass resolution downward during the resonance search. Furthermore, the shift is also taken into account in the final result.

Adding the two systematics impacting the mass resolution in quadrature yields $\sigma_{\text{target}} = 3\%$. As shown in Figure 37, the systematic is taken into account by degrading the mass resolution at each point by the size of σ_{target}. The parameterization given by equation

$$\sigma(m_{e^+e^-}) = 0.0532190838657m_{e^+e^-} - 0.0000922283032152 \tag{27}$$

is then used as an input into the final resonance search.

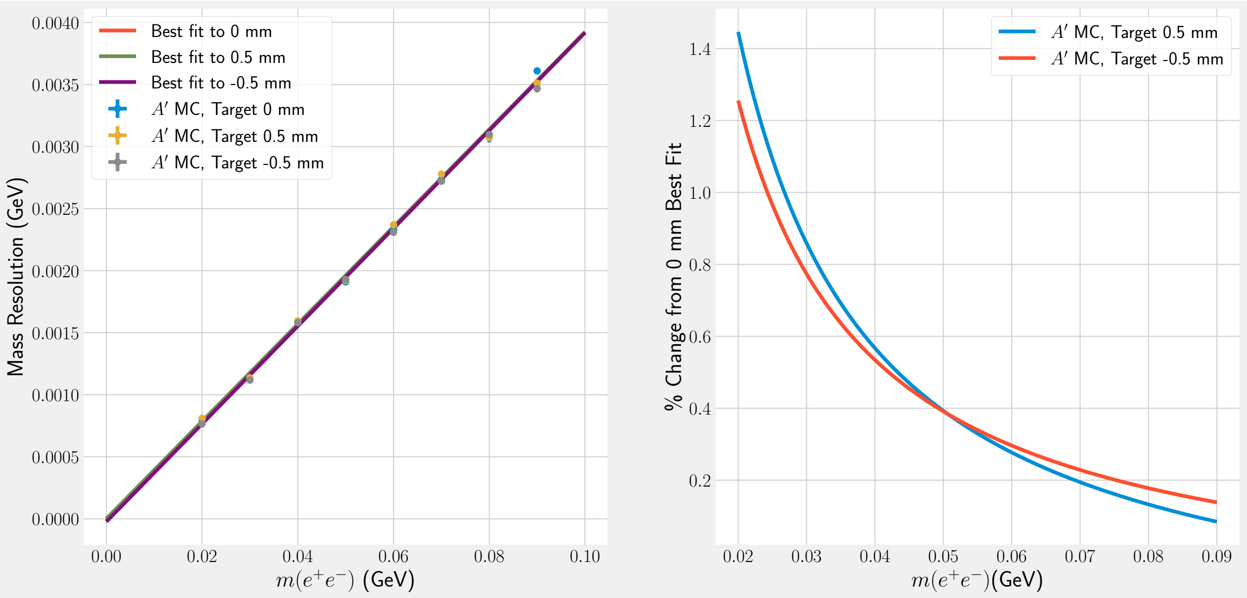

Figure 35: The mass resolutions of heavy photons originating at 0 and ±500 μm along with the best fit line to each of the datasets (left). The largest difference in the mass resolution between events constrained to the correct position and those off by 0.5 mm is 1.5% (right).

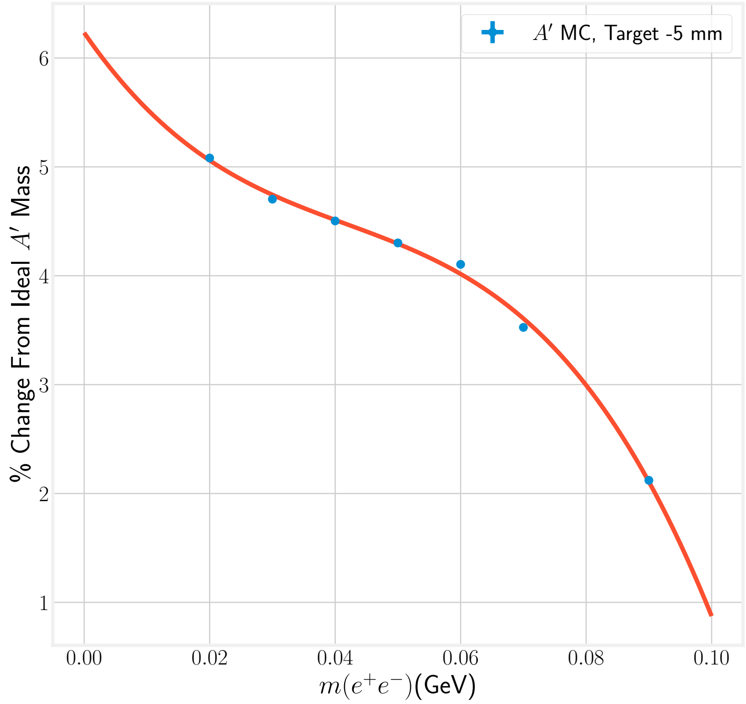

Figure 36: The percent change from the true A' mass when using the momentum constrained to 0 mm instead of -5 mm in the calculation of the invariant mass. The relationship is parameterized and used to adjust the mass scale in the final result.

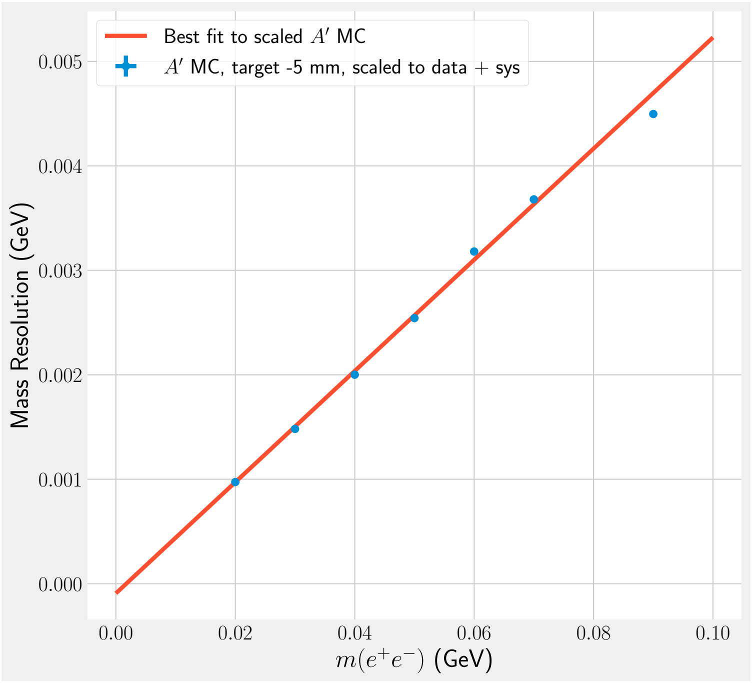

Figure 37: Mass resolution and parameterization after taking into the 3% systematic.

7.3 Fit Systematics

As discussed in Section 6, the fit parameters (polynomial, window size, bin size) were chosen such that both the signal bias and upper limit were minimized. However, as shown in Figure 30, the chosen fit parameters still had a signal bias associated with them. Given that these biases impact the upper limit, the signal pulls needs to be estimated using the full statistics. In order to do this, 2000 pseudo data sets were generated with a number of events equal to that in the full distributions. The signal pulls were then calculated as described in Section 6 for the two different window size regions ($<$ 70 MeV and $\geq$ 70 MeV) that were used in the final search (see Figure 38). The fit systematic was then estimated as 1σ of the distribution of mean pulls in each of the regions (see Figure 39). The systematic is found to be 2% in the region $<$ 70 MeV and 5% in the remaining region. The upper limits are slightly degraded to account for this uncertainty.

7.4 Summary of Systematics

A summary of all systematics used in this analysis is listed in Table 6.

Systematic Description	**Value**
Radiative Fraction	
e^+e^- Composition	7%
Mass Resolution	
Fit to Møller mass spectrum	2.6%
Target position	1.5%
Fits	
Fit systematic $<$ 70 MeV	9%
Fit systematic $>=$ 70 MeV	6%

Table 6: Summary of systematics found to impact the resonance search.

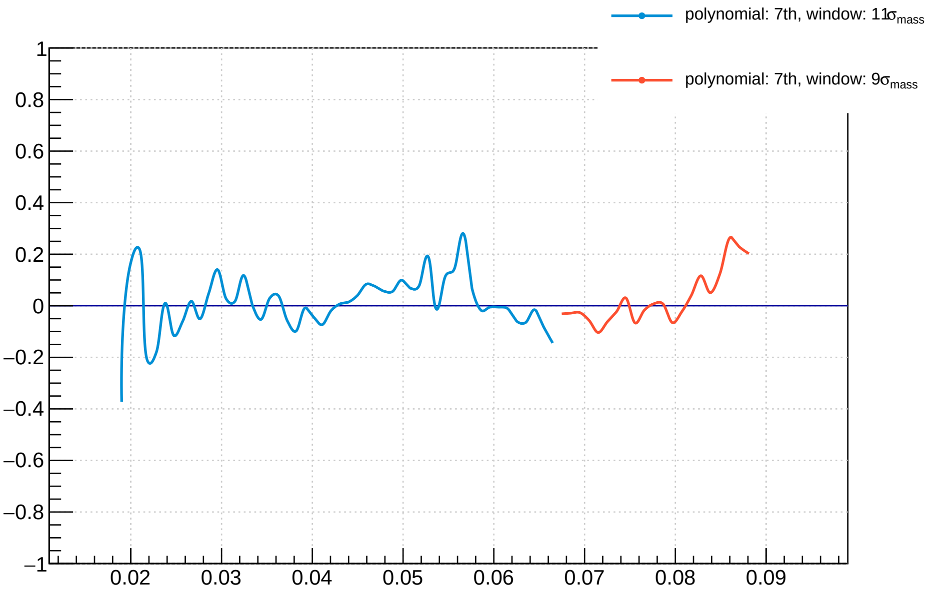

Figure 38: Signal pulls obtained from resonance scans over pseudo data sets with an equal number as the final e^+e^- invariant mass distribution.

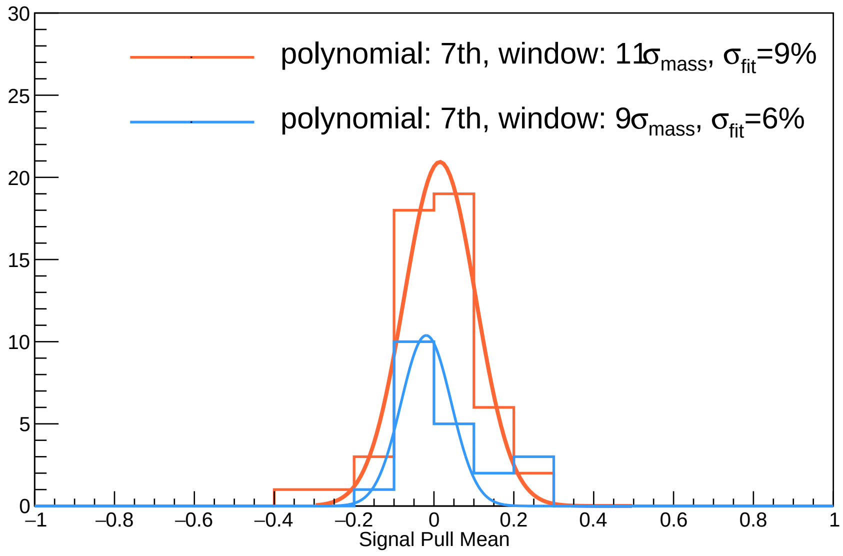

Figure 39: The distribution of pulls for the two window size regions ($<$ 70 MeV and $\geq$ 70 MeV) used in the final search. The fit systematic in each of these regions is estimated by taking 1σ of each distribution.

8 Fits and Signal Upper Limits

8.1 Fit Results

Using the method described in section 5 and the optimized fit parameters discussed in section 6, a search for a resonance in the e^+e^- spectrum was conducted between 19 MeV and 85 MeV. The resulting p-values from the search are shown in Figure 40. The most significant signals were found at 38.1 MeV and 73.5 MeV with local p-values of 0.0978 and 0.0032 respectively. The signal plus background fit (maroon) resulting in these p-values are show in Figure 41.

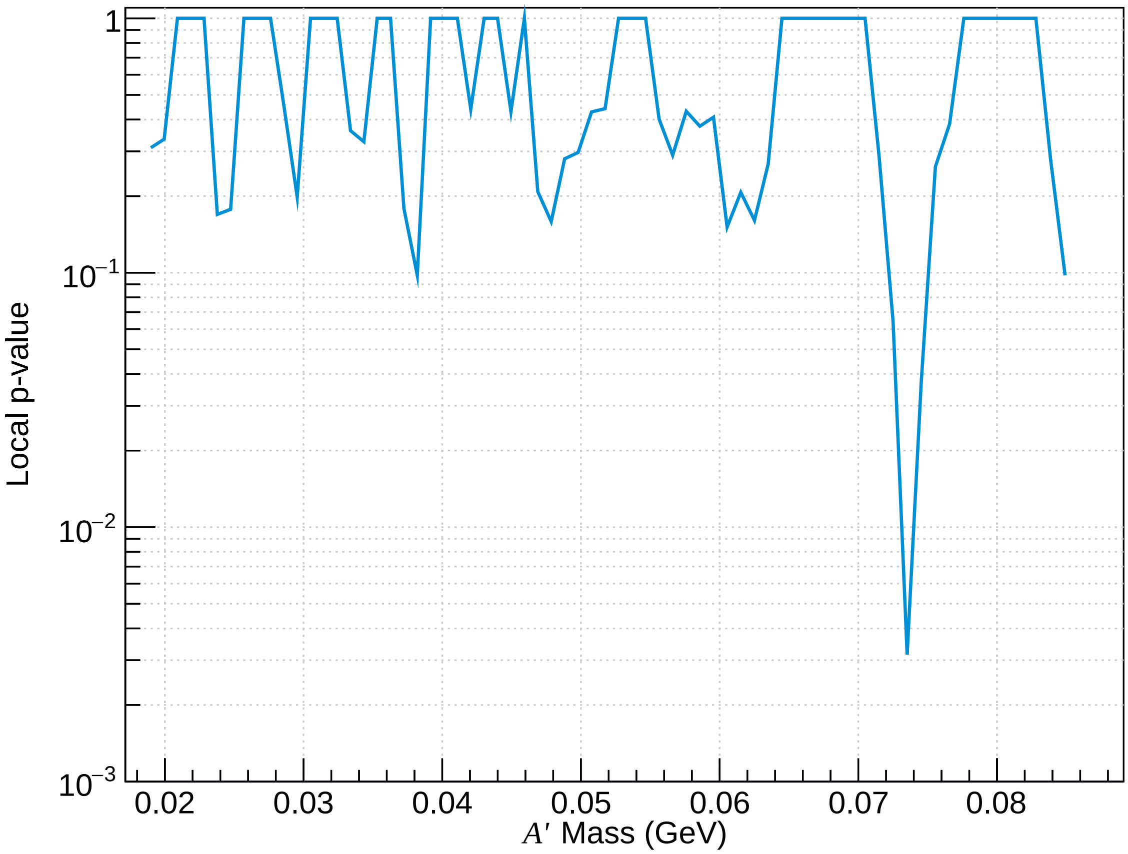

Figure 40: Resulting p-values from a resonance search for an A' across the e^+e^- invariant mass between 19-85 MeV.

As discussed in Section 5, the act of performing multiple searches across the e^+e^- mass spectrum may lead to the observation of significant local p-values simply as statistical fluctuations. Accordingly, a correction to account for the LEE effect needs to be applied to the local p-value before determining its global significance. As can be seen from Figure 26, after the application of the LEE correction, the most significant p-values are found to correspond to global p-values of .9 and .1, both of which are highly probable, and less than 2σ in significance. There are no significant signs of a signal.

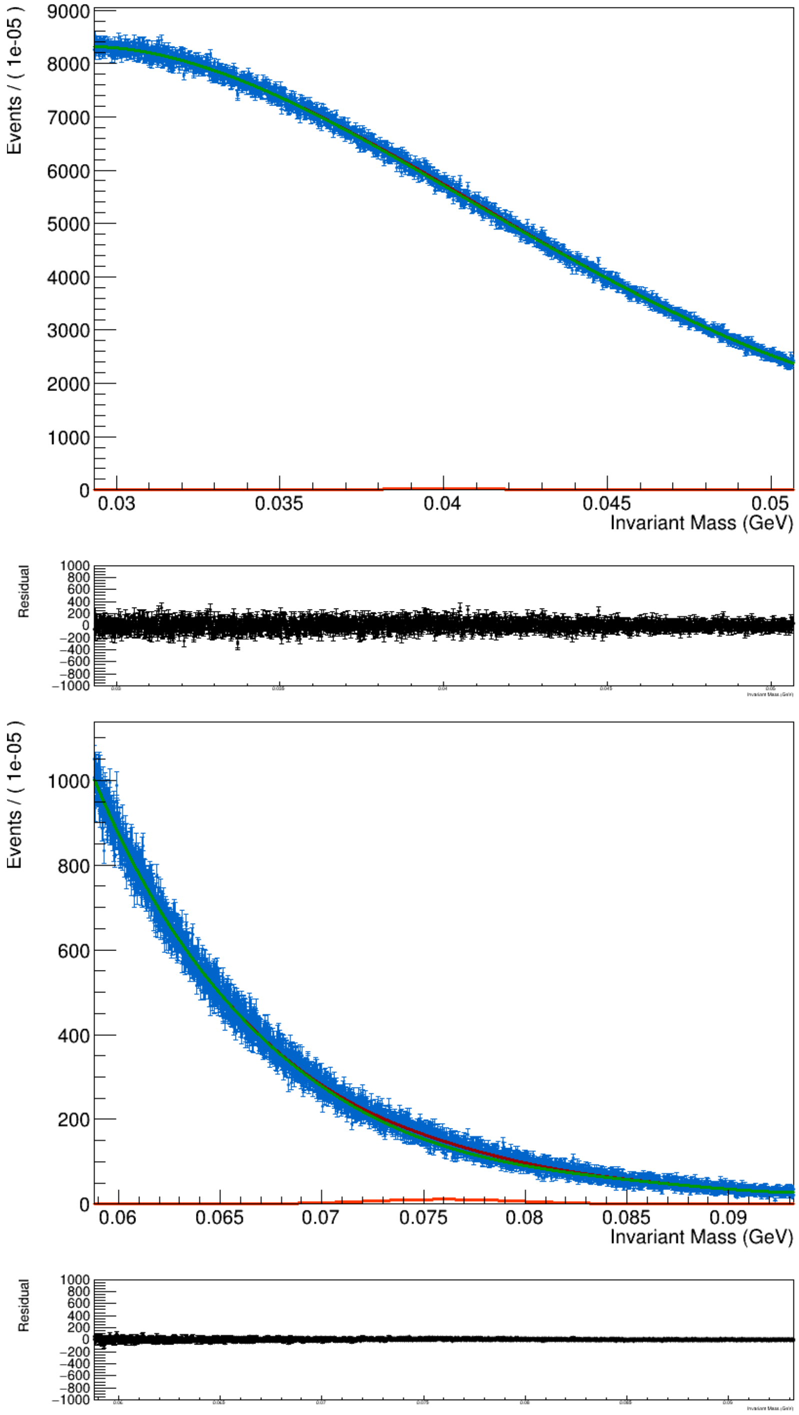

Figure 41: Resulting signal ~~plug~~ background fit (maroon) assuming an A' mass hypothesis of 38.1 MeV (top) and 73.5 MeV (bottom) of the most significant signals. The signal component is shown in red while the background component is shown green.

8.2 Signal Upper Limits

Since no significant resonances were found, a 2σ confidence upper limit on the number of signal events at each mass hypothesis was set. For the purpose of setting an upper limit, the likelihood ratio is inverted. The statistic used to set an upper limit is then

$$q_\mu = \begin{cases} -2\ln\frac{\mathcal{L}(\mu,\hat{\hat{\theta}})}{\mathcal{L}(0,\hat{\hat{\theta}})} & \hat{\mu} < 0 \\ -2\ln\frac{\mathcal{L}(\mu,\hat{\hat{\theta}})}{\mathcal{L}(\hat{\mu},\hat{\theta})} & 0 \leq \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \tag{28}$$

with the corresponding p-value being given by

$$p = \int_{q_{\mu,obs}}^{\infty} f(q_\mu|\mu)dq_\mu \tag{29}$$

where $f(q_\mu|\mu)$ is the probability distribution of q_μ given the hypothesized value of μ. In order to find the upper limit, μ_{up}, the test above is carried out over a range of signal yields until a p-value of 0.045 (2σ) is found. The signal yield value that corresponds to a p-value of 0.045 is μ_{up} and is often referred to as the unconstrained limit. The resulting unconstrained upper limits are shown in orange in Figure 42.

As shown in blue in Figure 42, it is often the case that the estimator for the signal yield, at a given mass hypothesis, is zero or even negative. In such cases, the probability distribution function of the test statistic q_μ assuming μ_{up} will nearly coincide with the distribution of q_μ assuming $\mu = 0$, i.e. the background only hypothesis. As a result, there is a lack of sensitivity to a signal measurement at those mass hypotheses.

In such cases, a 50% power-constrained upper limit on the signal is set [3]. At each mass hypothesis, a distribution of signal upper limits is generated from background only pseudo-data sets and the median (50% quantile) upper limit is calculated, μ_{median}. The upper limit in that region is then set to the larger of either the unconstrained limit or the median limit

$$\mu_{pc} = \text{max}(\mu_{\text{up}}, \mu_{\text{median}}). \tag{30}$$

The power constrained limits are shown in green in Figure 42.

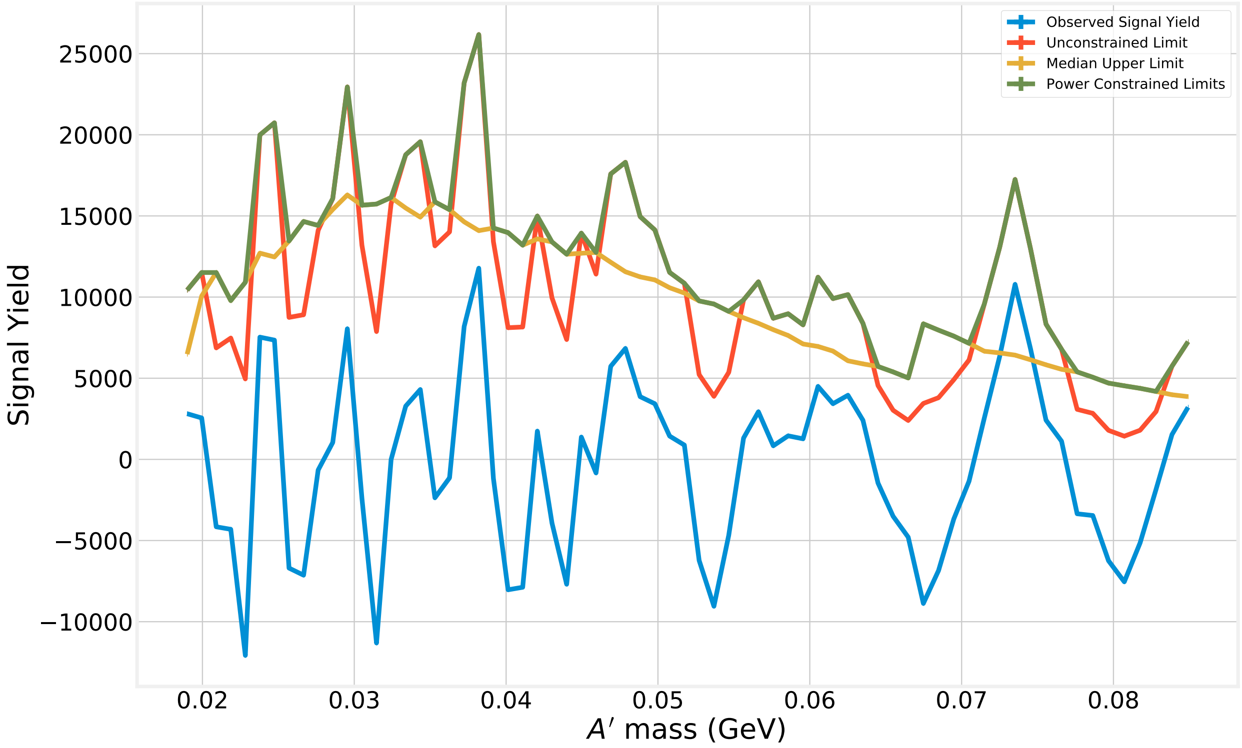

Figure 42: Unconstrained, median and power constrained upper limits on the signal yield at each mass hypothesis. The power constrained limits are what is used to set a final limit on the A' coupling strength.

9 Physics Results

As previously discussed, the kinematic similarities between heavy photons and radiative trident production allows their cross sections to be related within a mass window, δm, near $m_{A'}$ as

$$\frac{d\sigma(e^- Z \to e^- A' Z(A' \to e^+ e^-))}{d\sigma(e^- Z \to e^- \gamma^* Z(\gamma^* \to e^+ e^-))} = \left(\frac{3\pi\epsilon^2}{2N_{eff}\alpha}\right)\left(\frac{m_{A'}}{\delta m_{A'}}\right) \tag{31}$$

where N_{eff} is the number of available decay channels. For the A' masses considered in this analysis, $N_{eff} = 1$. Using Equation 31, the upper limit on the signal, S_{up}, can be related to an upper limit on the A' coupling strength as

$$\epsilon^2 = \left(\frac{S_{\text{up}}/m_{A'}}{f\Delta B/\Delta m}\right)\left(\frac{2N_{eff}\alpha}{3\pi}\right) \tag{32}$$

where $\Delta B/\Delta m$ is the number of background events per MeV and $f = 8.5\%$ is the radiative fraction. Estimating the number of background events per MeV is done by integrating the e^+e^- invariant mass spectrum within a 2.56σ_m window constructed around the A' mass hypothesis. The resulting number of events are then divided by the size of the window. The resulting number of background events per MeV at each mass hypothesis are shown in Figure 43.

The limits on the coupling derived using Equation 32 and using the full 2015 HPS engineering run dataset are shown Figure 44. Using the full 0.5 mm dataset, HPS is able to set a limit at the level of 10^{-5}.

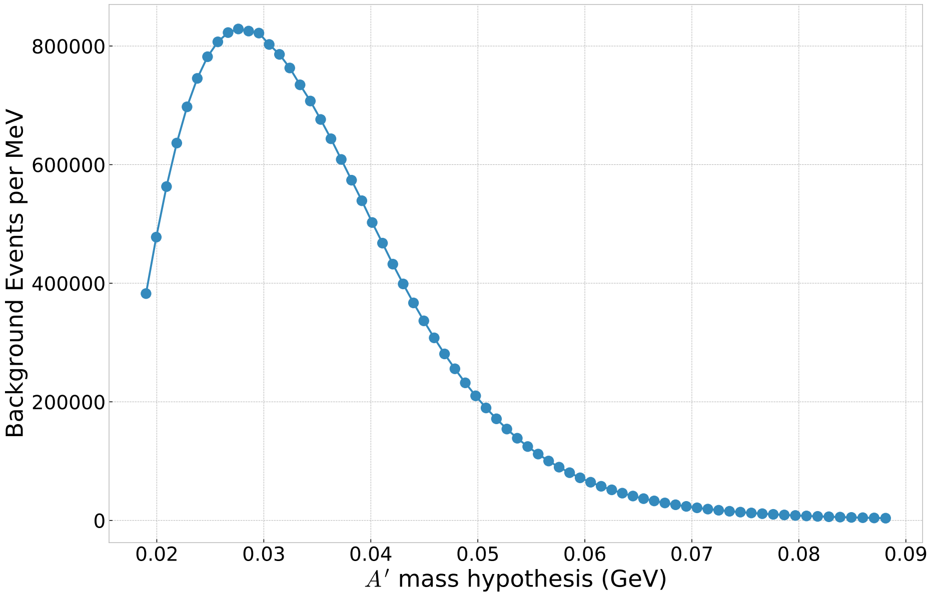

Figure 43: The number of background events per MeV at each mass hypothesis. The value at each mass is used as an input in calculating the coupling associated with a signal upper limit.

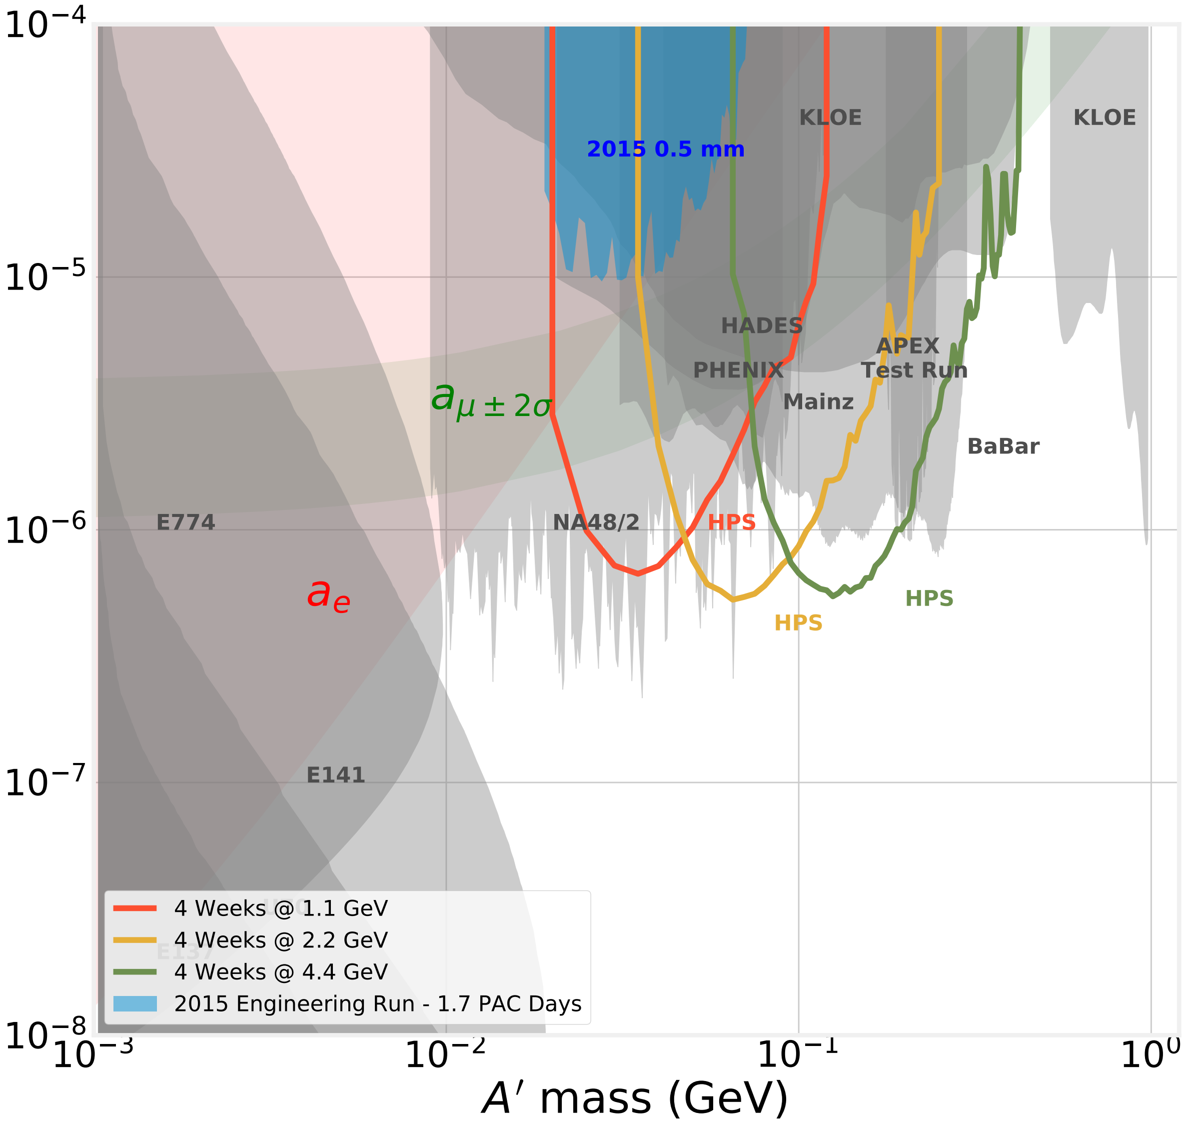

Figure 44: The 2σ upper limits obtained using the full 0.5 mm engineering run dataset (1.7 PAC days). A limit at the level of 10^{-5} is set. The estimated reach assuming 4 PAC weeks of running at 1.1 (orange), 2.2 (yellow) and 4.4 (green) GeV are also shown.

10 Conclusion

A resonance search for a heavy photon with a mass ranging between 19 MeV and 85 MeV and decaying to an e^+e^- pair was performed using the 0.5 mm 2015 HPS engineering run dataset (1165.7 nb^{-1}, 7.28 mC). A search for a bump in the e^+e^- invariant mass spectrum that would arise from the decay of the heavy photon to an e^+e^- pair was found to yield no significant excess. Since no significant signal was observed, 2σ upper limits on the square of the coupling were set at the level of 10^{-5}. This result is a factor of 1.5 worse than was projected at the time of the proposal. The difference is essentially all due to not properly including the electron hole of the ECal in the acceptance calculations. The future projections shown in Figure 44 account for this effect and are consistent with this result. The experiment is in agreement with limits established earlier by other experiments. This result, based on just 1.7 days of running, will be improved significantly with future running, as indicated in Figure 44.

References

[1] I. Balossino et al. The HPS electromagnetic calorimeter. *Nucl. Instrum. Meth.*, A854:89–99, 2017.

[2] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J.*, C71:1554, 2011. [Erratum: Eur. Phys. J.C73,2501(2013)].

[3] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Power-Constrained Limits. 2011.

[4] J. Gaiser. *Charmonium Spectroscopy From Radiative Decays of the J/ψ and ψ'*. PhD thesis, SLAC, 1982.

[5] Eilam Gross and Ofer Vitells. Trial factors or the look elsewhere effect in high energy physics. *Eur. Phys. J.*, C70:525–530, 2010.