D_s Semileptonic Decays

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Justin Thomas Hietala

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Professor Dan Cronin-Hennessy, Ph.D.

December, 2014

Acknowledgements

Strange as it may sound, I've feared writing this section more than any other part of my dissertation. Partially, this comes from my tendency to avoid "statements of strong personal preference."[1] However, I'm even more concerned by the inevitability that I won't give proper credit to those who have helped me through the years. For the record, then, more people have surely contributed to this work than I will acknowledge here, and those whom I do mention have almost assuredly given me more assistance than I can convey.

In days of yore, Ben Speakman, Chris Stepaniak, and Bernie Becker all helped me learn what progressing beyond an undergraduate degree was like. I may not have learned very quickly, but that's certainly through no fault of theirs. Alex Smith served as a wonderful mentor for a student just cutting his teeth on experimental particle physics with BTeV[2] and then CLEO. Alex was always friendly and willing to help; I regret not having spoken with him more during those early days.

I bothered Mikhail Voloshin far too often with all manner of physics questions, and he was always excessively generous with his time. I still wish that I would have found the η_b through $\chi_{b0}(2P) \to \eta\eta_b$ decays.

Ron Poling was intertwined with everything CLEO at Minnesota, and my work was hardly an exception. He asked good questions, and he didn't criticize when I asked poor ones. Yuichi Kubota contributed many analysis suggestions throughout my time at Minnesota, probably more than any other single person.

Eric Sorell gave me a year's worth of Fridays to take some time out from being a Dad and actually get some work done on this thing.

[1] These are disallowed in Minnesotan speech, according to Garrison Keillor's "Wobegonics."
[2] RIP

Tim Klein, Alexander Scott, Kaiyan Gao, Pete Zweber, Derrick Toth, Nick Howell, and the other denizens of room 253 all contributed various ideas over the years. More importantly, they provided mental and emotional support during my time at Minnesota. Brian Lang was always my first go-to on any conceptual or practical problem, whether it was for BTeV, CLEO, or just physics in general. I've had a hard time wrapping my head around the sum of his various contributions over the years.

It goes without saying that my parents, James Hietala and Becky McCauley, have been significant in all aspects of my education over the years. Plus, there was the whole "being born" thing, without which it seems unlikely that I'd have written this. The financial subsidies over the years have also helped keep me fed and have made life more bearable.

I'd particularly like to thank Dan Cronin-Hennessy for letting me be a full-time Dad while still working on this dissertation. It's something that most people in his shoes wouldn't have done, and since I would certainly have chosen my family over a Ph.D., this would have never been written without him. As if that wasn't enough, he was always friendly, helpful, and interested during a period when that's what I needed most. Dan also made administrative details disappear, which provided me some much needed sanity.

My wife, Jessica, has been with me since I started this work. She's always been my primary emotional support, and she provided financial support once our family grew. She's my partner in all aspects of my life, including this one. I love you, Jess.

Abstract

We measure six exclusive D_s semileptonic branching ratios using CLEO-c data collected at 4170 MeV. We isolate our semileptonic events by reconstructing a tagged D_s to identify a $D_s^* D_s$ event, then we find an electron and the semileptonic hadron. Dropping the D_s^* daughter photon gives us additional events and avoids the need to model soft photon backgrounds, at the expense of a clean neutrino missing mass. We obtain $\mathcal{B}(D_s \to \phi e\nu) = 2.14 \pm 0.17 \pm 0.09\%$ and $\mathcal{B}(D_s \to \eta e\nu) = 2.28 \pm 0.14 \pm 0.20\%$ for the two largest branching ratios.

Contents

List of Tables

List of Figures

Preface

This dissertation uses various conventions that warrant a brief discussion.

I've chosen to use *natural units* throughout the dissertation in derivations or when dealing with particle masses and momenta. In these units, the speed of light and the reduced Planck constant both get set to 1 ($c = \hbar = 1$). This simplifies the equations and the units while also making the relationship between quantities clearer. For instance, using 1968 MeV for the D_s mass makes the maximum energy available for final states obvious; writing 1968 MeV/c^2 provides exactly the same information but has more clutter, while writing the mass in kilograms (3.5×10^{-27} kg) gives almost no useful information.[3]

I retain conventional units when describing most other quantities, like a 20 cm cut on distance through the detector. In some cases, I choose units based on the context, like a 560 nm photon when discussing the visible light in the calorimeter but a 140 MeV photon in the D_s^* decay. Other quantities have their own unit conventions, like 0.916 nb for the $e^+e^- \to D_s^* D_s$ cross-section area (about 9.2×10^{-38} m^2). As a general rule, I've tried to provide the most convenient unit rather than adhere to some arbitrary standard.

Most energies/masses/momenta in this dissertation fall roughly in the 10^7–10^9 eV range, like the η mass at 5.48×10^8 eV. In almost all such cases, I've chosen to write 548 MeV rather than 0.548 GeV. While many particle experimentalists find the GeV scale more natural, I feel that using MeV for particles in this analysis's energy range makes it easier to compare differences with absolute quantities. For instance, I'll use

[3] I suppose that writing the mass in kilograms gives some sense of scale: a person's mass compared to the D_s mass is roughly the same as our sun's mass compared to a person's. This sense of scale really only helps the first time, though; I consider this footnote to provide the "sense of scale" benefit so that I can go ahead with more useful units in the main text.

a 30 MeV range on the 1968 MeV D_s mass, rather than a 0.030 GeV range on the 1.968 GeV mass or (egads!) a 30 MeV range on a 1.968 GeV mass.

I've seen different uses of the terms "branching ratio" and "branching fraction" in the literature. Some people use "branching fraction" to refer to a particle's partial decay width divided by its total width[4] and "branching ratio" to refer to a ratio of two branching fractions. While I see the appeal of this usage, I suspect that the literature more often uses "branching ratio" as the partial width divided by the total width (with "branching fraction" being synonymous). Consequently, I use "branching ratio" in the latter sense throughout and avoid the term "branching fraction" to eliminate any confusion.

Finally, I've chosen to use the plural "we" throughout the rest of this dissertation rather than the singular "I." There's quite a bit of debate over the best first person pronoun in scientific writing, none of which I find terribly conclusive.[5]

When discussing derivations, usage like "we see that A=B" must be plural as it includes the reader.[6] However, this does not extend to most of the text, as the reader (most likely) did not perform this analysis. Additionally, I have been given suggestions on how to proceed in some parts of this analysis. In these cases, "we" seems appropriate. Unfortunately, I have long since forgotten what portions of the analysis this applies to, and I wouldn't want to parse what deserves a "we" versus an "I" in any case.[7]

Ultimately, using "I" or an "I/we" mixture became very jarring and distracted from the writing. This practical effect has led to my choice of "we" throughout the dissertation. In some cases, the "we" can be considered to include the reader; in other cases, it reflects the assistance I've had along the way. The rest of the time, I'm just calling it the majestic plural.

[4] This is just the probability that a decaying particle will wind up in a given final state.

[5] I do think that avoiding the first person entirely tends to go south quickly. It encourages passive voice, which turns already dry, scientific writing into writing that's infrared telescope worthy.

[6] "I see that A=B" or "you see that A=B" don't work nearly as well.

[7] The AIP Style Manual prefers that you determine what deserves a "we" and what deserves an "I." However, their style advice otherwise falls into the category of "We/I sure is a problem; sucks to be us, doesn't it?" In any case, the style guide last saw a serious update in 1990, and I get the sense from the literature before and after that perspectives on this issue have changed in the last couple decades; I won't even go into how they cite a website.

Chapter 1

Introduction

This dissertation describes a measurement of the D_s meson's decay rate through six different semileptonic modes. The D_s meson comprises the lowest energy bound state between a charm quark and a strange quark.[1] Its semileptonic decays involve the charm quark converting into a lighter quark (either strange or down) with the emission of a leptonic doublet, here an electron and a neutrino.[2]

These semileptonic decay rates give insight into the D_s meson's spatial structure. Aside from any intrinsic value, the measured rates play into the accuracy of theoretical techniques used in fundamental measurements (Section 1.2). Further, the decay rates tell us about the structure of the more common particles found in its final states. They also provide a comparison point for other semileptonic measurements that probe fundamental parameters in the Standard Model.

1.1 Standard Model

A desire to know what happens and why has been common across diverse human[3] cultures throughout history. Often, this quest for knowledge has resulted in tangible benefits that improve the quality of human life; other times, it has merely resulted in

[1] Strictly speaking, the D_s^+ consists of a charm and an antistrange quark, while the D_s^- has an anticharm and a strange quark.

[2] Specifically, the D_s^+ emits a positron and neutrino while the D_s^- emits an electron and an antineutrino.

[3] Given recent Neandertal cave art and tool discoveries, humans might not have even been the only curious hominids on the block.

metaphysical satisfaction and a more fulfilled life. Science drives our understanding of the complex world around us by testing and rejecting false ideas until we're only left with theories accurate enough to make sound predictions.

A reoccurring pattern in our successful theories of nature involves finding basic elements whose properties and interactions build up the complex world we observe, from cells in biology to atoms in chemistry. Elementary particle physics takes this pursuit as far as possible, trying to uncover the fundamental components of our universe. The Standard Model represents our current, established understanding of the universe's basic building blocks.[4] Even though physicists established the Standard Model some time ago, it has matched every test so far, with the discovery of the missing top quark (1995), the tau neutrino (2000), and the long-sought-after Higgs boson (2012).

1.1.1 A Quick History

The search for nature's most basic elements has a history that stretches back thousands of years. Even in Ancient Greece, Democritus formed the idea of "atoms" as basic, indivisible units with different geometries that gave matter its various properties. Galileo held the notion of piccolissimi quanti (the "smallest quanta"). In the 1700s, Boscovich speculated about the basic units of matter being geometric points interacting via forces, not terribly unlike our modern concepts [1].

Particle physics really began, however, with the discovery that atoms have constituent particles. In 1897, J.J. Thomson won the race to understand cathode rays when he showed that they consisted of negatively charged particles, which we now call electrons. Ernest Rutherford followed up with a series of experiments that defined the nucleus, culminating in the identification of the proton in 1919. Chadwick rounded out our knowledge of the atom by providing experimental evidence for the neutron in 1932.

Meanwhile, Einstein published his theory of special relativity in 1905, describing the motion of high energy particles that we regularly use in particle physics today.[5] Always the overachiever, Einstein also described the photoelectric effect in the same year, leading to the notion that light consists of quantized units, or photons. Several

[4] Well, the building blocks for the universe if you ignore the 96% made up of dark matter and dark energy, which we don't understand in the slightest.

[5] Indeed, we internally call the subject matter *high energy physics* for good reason, although I find "particle physics" more straight-forward.

physicists contributed to quantum mechanics from the late 1800s/early 1900s on, describing how particles behave and leading Dirac to predict the existence of antimatter with the positron in 1927 (discovered in 1932).

Nature proved more interesting than having just photons, electrons, protons, neutrons, and their antiparticles, however. 1937 led to the discovery of the seemingly unnecessary muon[6] in cosmic rays [2], 1947 saw the discovery of the pion and the kaon [3], and, just for fun, 1950 added the wholly unexpected Λ^0 [4]. The rest of the 1950s and early 1960s followed up with a host of seemingly elementary particles, colloquially termed the "particle zoo."

While the number of basic particles expanded rapidly, our understanding of basic forces cleared up. A theory of gravity had been established in the 1600s by Isaac Newton and improved upon by Einstein with general relativity in 1916.[7] Electromagnetism, governing charges, magnetism, and light, had become fairly well understood by Maxwell's equations in 1873. The early 1960s saw electromagnetism merge with the weak interaction, which had been originally proposed by Fermi [5] in 1933 to explain beta decay (emission of electrons by a nucleus). These combined into the electroweak force [6], adding the $W^{\pm}$ and Z bosons (experimentally confirmed in 1981). The late 1960s featured the incorporation of the Higgs boson [7] into the electroweak model [8, 9].

The beginnings of the quark model came in 1964, independently proposed by Gell-Mann [10] and Zweig [11], with three basic quarks that combine to form composite particles (simplifying the "particle zoo"). Of particular relevance to this analysis, Glashow and Bjorken [12] predicted a fourth quark, the charm, although the idea didn't catch on until 1970 when the GIM mechanism [13] explained the nonexistence of flavor-changing neutral currents (like $K_L \to \mu^+\mu^-$). In 1968, deep inelastic scattering showed substructure to the proton [14, 15], initially called partons. By the early 1970s, the last bits of the Standard Model fell into place with a formal model for the strong force (QCD) [16, 17], the identification of partons with quarks and gluons, the prediction of bottom and top quarks by Kobayashi and Maskawa to explain CP violation [18], and the discovery of the tau [19].

[6] Confused originally (and understandably) with Yukawa's predicted nuclear intermediary, the pion.

[7] General relativity manages to be both fantastically interesting and beyond the scope of this dissertation, as gravity is the only one of the four basic forces that does not play into D_s semileptonic decays.

1.1.2 Model Overview

Each elementary particle has properties distinguishing it from the other particles.[8] A particle's *invariant mass* (normally just called "mass") defines the relationship between the particle's energy and momentum. A particle may have an electric charge that governs how strongly it interacts with electromagnetism, and it may have a color charge that governs its interactions with the strong force. A particle's *spin* refers to its intrinsic angular momentum and determines multiple particle statistics. Spin has natural units of $\hbar$, a common scale in quantum mechanics. We call particles with half-integer spin *fermions* as they follow Fermi-Dirac statistics and obey the exclusion principle (no two particles in the same quantum state). We call particles with integer spin *bosons* as they follow Bose-Einstein statistics (particles prefer the same quantum state, all else equal).

The current formulation of the Standard Model includes three forces[9] and 17 particles,[10] split into 12 spin-$\frac{1}{2}$ fermions, 4 spin-1 gauge bosons, and a spin-0 Higgs boson.[11] The fermions can be further divided into six leptons and six quarks, each in three generations of two particles. Each of these particles has no substructure within the limits of our measurement ability, so we call them elementary particles (or fundamental particles).

1.1.3 Leptons

The six leptons consist of three charged, electron-like particles (the electron, the muon, and the tau) and three neutral, neutrino-like particles (the electron neutrino, the muon neutrino, and the tau neutrino). Leptons have no color charge, so they only undergo weak and electromagnetic interactions (although the neutrinos have no electric charge and so *only* interact through the weak force). We typically group the leptons into pairs,

[8] This section's assertions can be found in most advanced particle physics texts, such as [20]. Particle properties come from [21].

[9] Gravity, with its hypothesized spin-2 graviton, does not get included in the Standard Model. Also, the electromagnetic and weak forces may be integrated together into a single electroweak interaction. However, the Z and W bosons' masses create enough interaction differences from the photon that we tend to talk about electromagnetic and weak interactions separately.

[10] In our 17 particle count, we don't add antiparticles or particles of different color charge separately. For instance, the charm quark counts as one particle instead of 6 (red, blue, or green; charm or anticharm). Counting these separately instead gives 61 particles.

[11] Other models for the Higgs mechanism exist (and have interesting consequences), but the scalar Higgs both fits current data and provides the simplest way to explain massive W and Z bosons.

or *generations*, corresponding to the mass of their electron-like particles:[12]

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix} \qquad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \qquad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}$$

For elementary particles, we base the unit of electric charge on the electron. Consequently, the electron and other electron-like leptons have charge $-1e$ (about -1.6×10^{-19} C), while their antiparticles have charge $+1e$. The tau and the muon both decay quickly (lifetimes of 3×10^{-13} s and 2×10^{-6} s, respectively) as they have masses much larger than the light, stable electron.

This analysis involves both electrons and electron neutrinos created in D_s semileptonic decays. Neutrinos only interact weakly and so pass through our detector with a vanishingly small probability of interaction (we commonly use the rule of thumb that a neutrino has a mean free path of about 1 light-year through lead). The neutrinos' lost energy and momentum creates the primary challenge for this analysis. Electrons, meanwhile, interact electromagnetically with multiple components of the detector and get reconstructed quite cleanly when they have enough momentum to make it through the detector's magnetic field.

1.1.4 Quarks

Like the leptons, the six quarks may also be grouped into pairs, or generations, based on their mass:

$$\begin{pmatrix} u \\ d \end{pmatrix} \qquad \begin{pmatrix} c \\ s \end{pmatrix} \qquad \begin{pmatrix} t \\ b \end{pmatrix}$$

These pairings also correspond to the strongest quark couplings with the W boson. The up-like quarks (up, charm, and top) each carry a $+\frac{2}{3}e$ charge, while the down-like quarks (down, strange, and bottom) carry a $-\frac{1}{3}e$ charge. Consequently, all quarks can interact electromagnetically in addition to their weak interactions.

Quarks also carry a color charge, so they can interact through the strong force. Since free particles must carry no color charge (color confinement, Section 1.1.5), quarks only appear in composite particles named *hadrons*. We call hadrons that consist of only two quarks (one quark and one antiquark) *mesons*, while we term hadrons that consist of

[12] We don't currently know the mass hierarchy for the neutrinos.

three quarks (all three quarks or all three antiquarks) *baryons*. Mesons tend to decay away entirely, with the longest lived mesons having lifetimes shorter than 10^{-7} s. Quarks inside a baryon can't annihilate due to their color charges, so baryons tend to decay down to the proton, the lightest baryon (although protons and neutrons form stable states together when combined, as in a nucleus).

In this analysis, we deal with the D_s meson, which involves both a charm (or anticharm) quark and an antistrange (or strange) quark. The semileptonic decays that we measure all involve the weak decay of the charm quark to a strange or down quark through a W boson.[13] We then see a new meson in the "final" state.

1.1.5 Gauge Bosons and Forces

The photon mediates the electromagnetic force, coupling to particles based on their electric charge. The photon's zero mass allows the electromagnetic interaction's range to extend arbitrarily far and generates the familiar $1/r^2$ force law between charged particles. Additionally, emitted photons don't decay when left to themselves, giving us the practical benefit of being able to see the world around us.

The very massive $W^{\pm}$ and Z^0 bosons mediate the weak force, limiting its range to short distances. While high energy processes can directly generate W and Z bosons (with a mean lifetime of around 3×10^{-25} s), they more often appear in the context of mediating particle decay processes like beta decay or the semileptonic decays considered in this analysis. Of the four fundamental forces known to physics, only the weak force can change lepton or quark *flavor*, break the parity symmetry, or break CP-symmetry.[14]

The massless gluon mediates the strong force, coupling to particles with a color charge. Unlike electric charge, which we can describe with just a positive or negative sign, the color charge has three different states, which we call "red", "green", and "blue". The names for these color charges come from an analogy with the RGB color model, as a composite particle with all three charges (or a charge and its complement)[15] has

[13] Other modes contribute to D_s semileptonic decay, like weak annihilation, but a direct decay through the W should dominate each decay rate.

[14] Changing flavor means the weak force can change the quark or lepton type without an antiparticle annihilation (e.g. muons can decay to electrons and neutrinos). Parity symmetry means that physics seen in a mirror operates the same way as in our own world, while CP-symmetry combines the parity mirror reversal (P) with a swap of the particle into its antiparticle (C).

[15] The complement of red may be called either antired or green-blue.

no color, which we sometimes refer to as "white." Quarks carry a single color charge (or anticharge), while gluons carry both a charge and an anticharge (although never in such a way as to make the gluon colorless).

Since gluons carry a net color, they can have strong interactions with other gluons (self-interactions), unlike the neutral photons in electromagnetism. This leads to novel properties for the strong force, like *asymptotic freedom* in which gluon self-interactions anti-screen a bare color charge. This enhancement causes the strong interaction coupling, α_S, to increase at large distances (or, equivalently, α_S decreases at high energies).[16] The gluon self-interaction also causes the gluon field between two color charges (e.g. quarks) to elongate into a tube instead of spreading out in space as in electromagnetism [22]. This causes the energy between the two quarks to increase linearly (unbounded) as the distance increases, eventually generating new quarks from the vacuum that form colorless hadrons with the original quarks. Ultimately, particles with a color charge can't be isolated, an effect known as *color confinement*.[17]

In this analysis, electromagnetism influences various decays, notably the initial e^+e^- collision and $D_s^* \to D_s\gamma$. Semileptonic decays naturally involve a W boson that couples to the quarks and to the leptonic decay products. The strong force's large coupling provides the motivation for the analysis in the first place, as the mess of gluonic interactions inside the initial and final state hadrons gives rise to their uncertain structure.

1.2 Motivation

Semileptonic decays of mesons with heavy quarks (B, D) have proven very useful in determining the magnitudes of CKM elements. In particular, such decays have led to the estimates for $|V_{cb}|, |V_{ub}|, |V_{cd}|$, and they have contributed to the estimate for $|V_{cs}|$ [21].

While D_s semileptonic decays could yield similar estimates in principle, its lower statistics and more complicated light hadron states make direct estimates of CKM elements less valuable. However, these same decays to light hadrons give insight into

[16] Gluon interactions with quarks provide a color screening effect as well, similar to the screening of charges in electromagnetism that leads to a very slow increase in its coupling (α) with energy. However, the anti-screening effect of gluons dominates for a reality with six quarks and three colors.

[17] This explanation of color confinement has not yet been established definitively; one of the seven Millennium Prize Problems relates to part of the proof.

calculations from lattice QCD, provide a probe into light meson spectroscopy, and comprise much of the inclusive spectrum that can be used for kinematic checks on D and B decays.

1.2.1 Lattice QCD

The lattice QCD formalism allows QCD processes to be computed numerically, using a discrete space-time grid with only the bare quark masses and α_s as inputs. Lattice QCD has proven very useful in a variety of processes, including the extraction of CKM parameters from semileptonic decays to light mesons [21].

However, the required lattice discretization size depends on the quark masses involved, and smaller lattice sizes require more computing power. While a few lattice analyses have recently begun to get down to the level of the physical light (up and down) quark masses [23], lattice calculations still typically set the light quark masses to a higher than physical value, expressed as a fraction of the strange quark mass [24, 25, 26]. They then use different lattice grid sizes and extrapolate to the continuum limit.

These lattice discretization effects tend to dominate lattice calculation errors. D_s semileptonic decays to $s\bar{s}$ states provide an excellent test of lattice QCD procedures as they allow lattice calculations to use the strange quark mass for the valence quarks rather than extrapolating to light quark masses.

1.2.2 Light Meson Spectroscopy

Semileptonic D_s decays most often result in $s\bar{s}$ final states, which affords us an opportunity to probe a sector that can otherwise be difficult to access cleanly [27]. In particular, D_s semileptonic decays can potentially probe the $s\bar{s}$ content of the η and η', and they can shed insight into the $f_0(980)$ quark content and structure.

When decaying semileptonically to pseudoscalars, the D_s couples to the $s\bar{s}$ component of η and η'. Conversely, when charged D semileptonic decays result in an η or η', they couple to the mesons' $d\bar{d}$ component. Since the decays have related kinematics, a comparison of the four decay widths should determine the strange and nonstrange $q\bar{q}$ content of the η/η'.

The $f_0(980)$ has been considered to consist of a $q\bar{q}$ state, a $qq\bar{q}\bar{q}$ state, a $K\bar{K}$ molecule,

or even to have a gluon component [28]. Since the D_s transitions to the f_0 particle's $s\bar{s}$ component, $D_s \to f_0 e\nu$ should provide information on the underlying quark content of the f_0. BaBar may have seen S-wave interference with the $\phi e\nu$, $\phi \to K^+K^-$ final state [29], and this mechanism could also lead to a deeper understanding of the f_0 substructure.

1.2.3 Inclusive D_s

Given CLEO's inclusive D_s measurement [30], the six D_s semileptonic modes considered in this analysis ($\phi e\nu$, $\eta e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K^* e\nu$, $K_s e\nu$) saturate most of the total semileptonic width. Knowing the components of the D_s semileptonic width should improve phenomenological comparisons that use the inclusive D_s spectrum.

In the most prominent example, heavy quark symmetry allows a constraint on the weak annihilation (four-quark, Figure 1.2) component of $B \to X_u l\nu$ semileptonic decays that would otherwise complicate the $|V_{ub}|$ measurement. This constraint comes from comparing the difference of charged and neutral B semileptonic widths to the difference in D_s and D^0 semileptonic widths, which should be related up to factors like m_b^3/m_c^3 [31, 32].

1.3 Theory

In D_s semileptonic decays to light hadrons, the charm quark couples to a down or strange quark through a W boson, which also couples to a $l\nu$ lepton pair (Figure 1.1). The lepton and quark states trivially separate (factorize) in the decay amplitude since the leptons don't feel the strong force. However, we can not simply ignore the spectator strange quark in the $c \to (d, s)$ coupling. The daughter down and strange quark masses both fall well below Λ_{QCD}, so $\lambda_{(d,s)}$ becomes too long to benefit from asymptotic freedom (as might be done with the short heavy quark wavelengths in $b \to c$). The resultant strong interactions with the spectator quark complicate the picture enough that it generally becomes easier to work directly with the meson coupling ($D_s \to X_q$) than with the individual quark coupling.

Separating out the weak decay's lepton pair gives a decay amplitude of[18]

[18]The notation used here comes from [20].

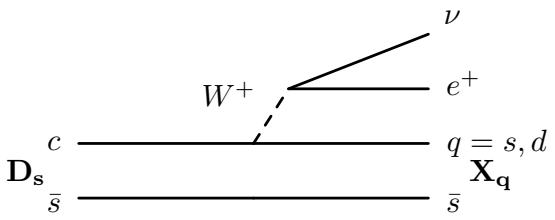

Figure 1.1: Spectator model for $D_s \to X_q e\nu$.

$$\mathcal{A} = \frac{G_F}{\sqrt{2}} V_{cq}^* \bar{u_e} \gamma_\mu (1-\gamma_5) v_\nu \langle X_q | \bar{q} \gamma^\mu (1-\gamma_5) c | D_s \rangle , \tag{1.1}$$

where q corresponds to the d or s final state quark field and X_q represents our final, exclusive meson state ($\phi, \eta, \eta', f_0, K^*$, or K^0). Here, we have implicitly assumed a simple $q\bar{q}$ form for the f_0 and have ignored complications from weak annihilation.

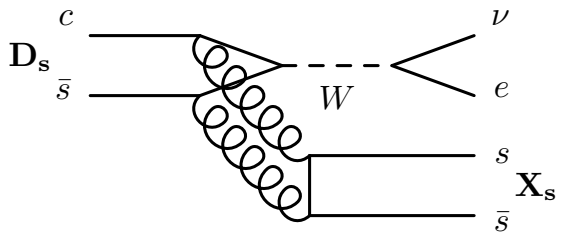

Figure 1.2: An example of the nonfactorizable weak annihilation contribution in D_s semileptonic decays to an $s\bar{s}$ state (e.g. η, η', ϕ). A similar contribution can appear in $B \to X_u l\nu$, which may distort the $|V_{ub}|$ measurement if too large.

After summing over lepton spins, the leptonic part of the squared amplitude

$$\sum_{spins} |\mathcal{A}|^2 = \frac{G_F^2}{2} |V_{cq}|^2 L_{\mu\nu} H^{\mu\nu} \tag{1.2}$$

becomes

$$\begin{aligned} L_{\mu\nu} &= 2\,\mathrm{Tr}[\not{k}\gamma_\mu \not{p}_e \gamma_\nu(1-\gamma_5)] \\ &= 8(q-p_e)^\alpha p_e^\beta \left[g_{\alpha\mu}g_{\beta\nu} + g_{\alpha\nu}g_{\beta\mu} - g_{\alpha\beta}g_{\mu\nu} + i\epsilon_{\alpha\beta\mu\nu}\right], \end{aligned} \tag{1.3}$$

where k represents the neutrino's four momentum, p_e represents the electron's four momentum (which will be taken to be massless in what follows), and $q \equiv k + p_e = p_{D_s} - p_X$.

As mentioned previously, the inability to separate out the spectator quark complicates the hadronic part of the squared amplitude. The standard approach relates the hadronic amplitude to a series of form factors (functions of q^2) that weigh all the possible combinations of physical dependencies. For instance, the hadronic amplitude for a pseudoscalar D_s decay to a pseudoscalar X_q (e.g. $D_s \to \eta e\nu$) can only depend on $p^\mu_{D_s}$ and p^μ_X as no spins or orbital motion are involved (the axial part of the hadronic current can't contribute). We typically combine these into

$$\langle X_q|\,\bar{q}\gamma^\mu c\,|D_s\rangle = f_+(p_{D_s}+p_X)^\mu + f_-(p_{D_s}-p_X)^\mu, \tag{1.4}$$

since the decay rate only depends upon f_+ in the zero lepton mass limit ($q_\mu L^{\mu\nu} = 0$ for massless leptons). Decays to vector particles follow a similar procedure but with additional kinematic dependencies and form factors.

Different models make different assumptions about these form factors' q^2 dependence. We have used the ISGW2 model [33, 34] as a baseline in our Monte Carlo generation, but most results in the literature use some form of a nearest pole dominance model. Although not entirely identical, both models have similar forms (if different parametrizations) in the case of D_s decaying to pseudoscalar or vector particles.

1.3.1 Free Quark

The simplest model of D_s semileptonics comes from ignoring our earlier precaution against separating out the spectator quark and considering only the charm to down/strange coupling with the W in the decay amplitude. While this assumption should be a disaster for the down quark, we can expect it to be merely awful when applied to the strange quark ($m_s \approx 100$ MeV $< \Lambda_{QCD} \approx 200$ MeV). Nonetheless, it gives us a useful baseline to compare against the more sophisticated techniques that include resonance effects.

In the free quark case, the hadronic part of the squared amplitude becomes

$$H^{\mu\nu} = \mathrm{Tr}\left[\not{P}\gamma^{\mu}\not{p}'\gamma^{\nu}(1-\gamma_5)\right], \tag{1.5}$$

where P is the momentum of the charm quark, p' is the momentum of the final state quark, and the factor of 2 seen in $L_{\mu\nu}$ drops when averaging over the two initial charm spin states. After simplifying $L_{\mu\nu}H^{\mu\nu}$ to $8(P\cdot k)(p'\cdot p_e)$ and integrating over the massless electron and neutrino momenta in the charm rest frame, we get

$$\frac{d\Gamma}{dE'} = \frac{G_F^2|V_{cq}|^2}{12\pi^3}\sqrt{E'^2-m^2}\left[m^2(3E'-2M)+ME'(3M-4E')\right], \tag{1.6}$$

where M is the mass of the decaying charm quark, m is the mass of the final state quark, and E' is the energy of the final state quark (which may also be written in terms of $q^2 = M^2+m^2-2ME'$).

The total decay rate after considering all E' between m and $\frac{M}{2}$ has the familiar form

$$\Gamma = \frac{G_F^2|V_{cq}|^2}{192\pi^3} f\left(\frac{m^2}{M^2}\right) \tag{1.7}$$

for $f(x) = 1-8x+8x^3-x^4-12x^2\ln x$.

We could further improve this by including bound state effects (e.g. time dilation of the charm quark due to its momentum inside the D_s) and QCD corrections [35]. However, the simple form above should be sufficient for our purposes since the free quark model only serves to give an order of magnitude estimate on the inclusive decay rate and a rough sense of the dependence on q^2.

1.3.2 ISGW Formalism

The initial ISGW model [33] takes the most general form for the hadronic part of the squared amplitude, then relates it to particular form factors for each exclusive decay (ignoring nonresonant states). It uses a nonrelativistic quark potential model with a Coulomb plus linear potential given by

$$V(r) = -\frac{4\alpha_s}{3r} + c + br, \tag{1.8}$$

which generates wave functions for the mesons and an explicit calculation for the form factors based on those wave functions.

After summing over final state hadron spins, the most general hadronic portion of the squared amplitude has the form

$$\begin{aligned} H^{\mu\nu} = &-\alpha g^{\mu\nu} + \beta_{++}(p_{D_s}+p_X)^\mu(p_{D_s}+p_X)^\nu + \beta_{+-}(p_{D_s}+p_X)^\mu(p_{D_s}-p_X)^\nu \\ &+ \beta_{-+}(p_{D_s}-p_X)^\mu(p_{D_s}+p_X)^\nu + \beta_{--}(p_{D_s}-p_X)^\mu(p_{D_s}-p_X)^\nu \\ &+ i\gamma\epsilon^{\mu\nu\rho\sigma}(p_{D_s}+p_X)_\rho(p_{D_s}-p_X)_\sigma, \end{aligned} \tag{1.9}$$

where α, the various β, and γ may depend on q^2.

Combining with $L_{\mu\nu}$ gives

$$\begin{aligned} L_{\mu\nu}H^{\mu\nu} = 8\{\ &2\alpha(p_e\cdot q) + \beta_{++}\left[2(P\cdot q)(P\cdot p_e) - 2(P\cdot p_e)^2 - (q\cdot p_e)P^2\right] \\ &+ 2\gamma\left[(P\cdot q)(q\cdot p_e) - (P\cdot p_e)q^2\right] \\ &+ \left[q^2 - 2(q\cdot p_e)\right]\left[(\beta_{+-}+\beta_{-+})(P\cdot p_e) + \beta_{--}(q\cdot p_e)\right]\ \}, \end{aligned} \tag{1.10}$$

where $P = p_{D_s} + p_X$.

Using the phase space delta function to integrate over the neutrino momenta and $d(\cos\theta_{qe})$ also conveniently gives $q\cdot p_e = \frac{q^2}{2}$, eliminating all but the α, β_{++}, and γ terms. Defining the dimensionless variables $x \equiv E_e/m_{D_s}$ and $y \equiv q^2/m_{D_s}^2$ leads to the differential decay rate

$$\begin{aligned} \frac{d^2\Gamma}{dxdy} = \frac{G_F^2 m_{D_s}^5 |V_{cq}|^2}{32\pi^3}\Bigg\{ &\frac{\alpha}{m_{D_s}^2}y + 2\beta_{++}\left[-4x^2 + 2x\left(1 - \frac{m_X^2}{m_{D_s}^2} + y\right) - y\right] \\ &+ \gamma y\left(1 - \frac{m_X^2}{m_{D_s}^2} - 4x + y\right)\Bigg\}. \end{aligned} \tag{1.11}$$

The particular α, β_{++}, and γ depend upon the final state meson. For a decay to a pseudoscalar, the hadronic amplitude follows Eq. 1.4 with $\alpha = \gamma = 0$ and $\beta_{++} = f_+^2$.

A decay into a vector particle (e.g. $D_s \to \phi e\nu$) has the vector's polarization (ε) as an allowed kinematic variable in the (non-spin-averaged) amplitude. The axial and vector amplitudes become

$$\langle X|\, A_\mu\, |D_s\rangle \equiv f\varepsilon_\mu + a_+(\varepsilon\cdot p_{D_s})(p_{D_s}+p_X)_\mu + a_-(\varepsilon\cdot p_{D_s})(p_{D_s}-p_X)_\mu \tag{1.12}$$

$$\langle X|\, V_\mu\, |D_s\rangle \equiv ig\epsilon_{\mu\nu\rho\sigma}\varepsilon^\nu(p_{D_s}+p_X)^\rho(p_{D_s}+p_X)^\sigma, \tag{1.13}$$

where f, g, and $a_\pm$ are form factors that depend on q^2. The decay rate in this case requires all three of α, β_{++}, and γ (e.g. $\gamma = 2gf$).

In each case, the meson wave functions used in calculations with the potential are taken from the harmonic oscillator

$$\psi^{1S} = \frac{\beta_S^{3/2}}{\pi^{3/4}} e^{-\beta_S^2 r^2/2}, \tag{1.14}$$

with a free parameter (β_S) that gets fixed by the variational method.

The resulting pseudoscalar and vector form factors (f_+, f, g, and $a_\pm$) depend upon the quark masses, the quark model's meson masses, and the variational parameters β_{D_s} and β_X. Their q^2 dependence comes in the form $e^{-r_1^2(q^2_{\max}-q^2)/6}$, where the constant r_1 carries the mass and β_S dependence. The decay rate parameters α, β_{++}, and γ inherit the exponential q^2 dependence from the form factors and may also contain explicit q^2 coefficients. α and β_{++} contain such terms in the vector decays; for example, α contains a term of the form $g^2(m_{D_s}^2 - q^2)$. The final decay rate does not in general have a simple q^2 dependence with this model.

The updated ISGW2 model [34] that we use in our Monte Carlo generation follows the same approach but makes a series of modifications to the original model. Most notably, it takes advantage of heavy quark symmetry to more accurately match the quark model form factors to the meson form factors, it adds relativistic corrections, and it changes the form factor q^2 dependence from exponential to $[1 + r_2^2(q^2_{\max} - q^2)/6N]^{-N}$, where N=2 for S-wave to S-wave transitions and $r_2^2 = r_1^2 + \mathcal{O}(\frac{1}{m_c m_q})$.

1.3.3 Pole Model

Assuming nearest pole dominance yields an alternative form factor dependence on q^2. Again, the pseudoscalar to pseudoscalar form factors only involve the vector hadronic current $V_\mu = \bar{q}\gamma_\mu c$, now using the parametrization [36]

$$\langle X | V_\mu | D_s \rangle = F_1(q^2) \left[(p_{D_s} + p_X)_\mu - \frac{m_{D_s}^2 - m_X^2}{q^2} q_\mu \right] + F_0(q^2) \left[\frac{m_{D_s}^2 - m_X^2}{q^2} q_\mu \right], \tag{1.15}$$

where $F_1(q^2)$ and $F_0(q^2)$ correspond to the two possible form factors (with the constraint $F_1(0) = F_0(0)$ to avoid an issue at $q^2 = 0$).

This parametrization has the advantage that only the $F_0(q^2)$ term survives when acted on by q^μ, allowing its identification as the scalar ($J_P = 0^+$) component of the current. Similarly, the $F_1(q^2)$ form factor corresponds to the spin-1 component of the current (orthogonal to q^μ).

In the massless electron limit, the pseudoscalar to pseudoscalar decay rate only depends on $F_1(q^2)$ and is given by

$$\frac{d\Gamma}{dq^2} = \frac{G_F^2 |V_{cq}|^2}{24\pi^3} \boldsymbol{p_X}^3 \left|F_1(q^2)\right|^2, \tag{1.16}$$

where $\boldsymbol{p_X}$ is the final state hadron's three momentum in the D_s rest frame.

Since the form factors are analytic in the complex q^2 plane other than singularities when q^2 has an on-shell intermediate particle, we can assume that the nearest meson resonance with proper spin and quark content will dominate the form factor behavior (Fig. 1.3). The dispersion relation gives

$$F_1(q^2) = \frac{1}{\pi} \int \frac{\mathrm{Im}F_1(s) ds}{s - q^2 - i\epsilon}, \tag{1.17}$$

with

$$\mathrm{Im}F_1(s) = \pi C \, \delta(s - M^2) \tag{1.18}$$

for a narrow resonance at the meson with mass M, where C is a constant that contains the coupling strengths between both the resonance with the W and the resonance with the decaying particle. Combining equations gives

$$F_1(q^2) = \frac{F_1(0)}{1 - \frac{q^2}{M^2}}. \tag{1.19}$$

This "simple pole" model serves well as a first order approximation for the form factor. However, D_s (and other) semileptonics have a q^2 range too wide to simply be dominated by a single pole at one extreme point (near $q^2 = q^2_{\mathrm{max}}$). Various modifications to the simple pole have been proposed [37], most of which equate to adding extra effective poles (or multipoles) with the pole strengths and pole masses left as variables for fitting (e.g. $F_1(q^2) = \frac{F_1(0)}{(1-(q/M)^2)(1-\alpha(q/M)^2)}$ [38]). At a certain point, this becomes less about describing a perfect physical model and more about having roughly the right form with enough free parameters to match the data. Nonetheless, having a

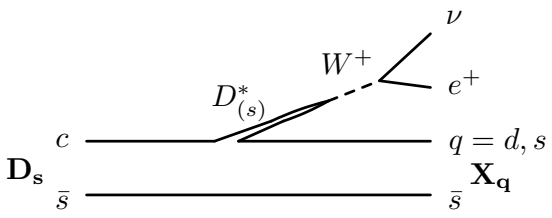

Figure 1.3: Pole dominance in $D_s \to X_q e\nu$. Here we have shown the vector meson pole through the D^* and D_s^* resonances when the charm quark decays to a down or strange quark, respectively.

common parametrization for comparison has value when we have multiple methodologies (Section 1.4).

Pseudoscalar to vector decays in the pole model commonly use the parametrization

$$\langle X|\, V_\mu \,|D_s\rangle = 2i\epsilon_{\mu\nu\alpha\beta}\frac{\varepsilon^\nu p_X^\alpha p_{D_s}^\beta}{m_{D_s}+m_X}V(q^2), \tag{1.20}$$

$$\begin{aligned}\langle X|\, A_\mu \,|D_s\rangle = {} & (m_{D_s}+m_X)\left[\varepsilon_\mu - \frac{(\varepsilon\cdot q)q_\mu}{q^2}\right]A_1(q^2) \\ & - (\varepsilon\cdot q)\left[\frac{(p_{D_s}+p_X)_\mu}{m_{D_s}+m_X} - \frac{(m_{D_s}-m_X)q_\mu}{q^2}\right]A_2(q^2) \\ & + 2m_X\frac{(\varepsilon\cdot q)q_\mu}{q^2}A_0(q^2),\end{aligned} \tag{1.21}$$

where A_μ represents the axial current $(\bar{q}\gamma_\mu\gamma_5 c)$ and ε represents the final vector particle's polarization. Only the $A_0(q^2)$ term survives when acted on by q^μ in this parametrization, thereby representing the pseudoscalar resonance's form factor. $V(q^2)$ corresponds to the vector meson resonance, leaving $A_1(q^2)$ and $A_2(q^2)$ to describe the axial meson resonance exchange. We can also define

$$A_3(q^2) \equiv \frac{m_{D_s}+m_X}{2m_X}A_1(q^2) - \frac{m_{D_s}-m_X}{2m_X}A_2(q^2), \tag{1.22}$$

with the constraint that $A_3(0) = A_0(0)$ so that no pole arises at $q^2 = 0$.

These form factors may also be written as helicity amplitudes [39] by considering the conserved helicity between the vector meson and the W, giving

$$H_0(q^2) \equiv \frac{M_{D_s} + M_X}{2M_X\sqrt{q^2}} \left[(M_{D_s}^2 - M_X^2 - q^2)A_1(q^2) - \frac{4M_{D_s}^2 |\vec{p_X}|^2}{(M_{D_s} + M_X)^2} A_2(q^2) \right], \quad (1.23)$$

$$H_\pm(q^2) \equiv (M_{D_s} + M_X) \left[A_1(q^2) \mp \frac{2M_{D_s}|\vec{p_X}|}{(M_{D_s} + M_X)^2} V(q^2) \right]. \quad (1.24)$$

With these amplitudes, the decay rate takes the compact form

$$\frac{d\Gamma}{dq^2} = \frac{G_F^2 |V_{cq}|^2}{96\pi^3} \frac{q^2 \boldsymbol{p_X}}{M_{D_s}} \sum_{i=0,\pm} \left| H_i(q^2) \right|^2. \quad (1.25)$$

The overall decay rate does not depend upon $A_0(q^2)$ in the limit of a massless electron (the integrated leptonic current $L^{\mu\nu}$ for massless leptons has the form $q^\mu q^\nu - q^2 g^{\mu\nu}$, so $q_\mu L^{\mu\nu} = 0$). From here, the A_1, A_2, A_0 and V form factors each get represented by a simple pole using meson resonances with the proper axial, pseudoscalar, or vector characteristics.

1.4 Decay Rate Predictions

QCD doesn't particularly lend itself to precision calculations since its large α_s coupling prohibits perturbative techniques. This limits an exact calculation of D_s semileptonic decay rates and form factors from first principles. However, versions of constituent quark models, QCD sum rules, and lattice QCD have all been used to avoid the problems intrinsic to QCD processes and predict the various D_s semileptonic decay rates.

1.4.1 $D_s \to (\eta, \eta')e\nu$

D_s semileptonic decays to the pseudoscalar η and η' states hold interest for both the η/η' mixing angle (or η/η'/glue mixing [40]) and the weak annihilation process (Figure 1.2) [41]. These decays have consequently seen the most discussion in recent years, with multiple methods of calculating form factors and decay rates.

Predictions typically require a value for the mixing angle, ϕ, given by

$$|\eta\rangle = \cos\phi\,|\eta_q\rangle - \sin\phi\,|\eta_s\rangle \quad (1.26)$$

$$|\eta'\rangle = \sin\phi\,|\eta_q\rangle + \cos\phi\,|\eta_s\rangle, \quad (1.27)$$

where $|\eta_s\rangle = |s\bar{s}\rangle$ and $|\eta_q\rangle = \frac{1}{\sqrt{2}}\left|u\bar{u} + d\bar{d}\right\rangle$. Methods we discuss take $\phi \approx 40\,^\circ$ unless otherwise noted. The octet-singlet η_8/η_1 mixing angle (θ) can also be used, which differs from ϕ by $\arctan(\sqrt{2})$. The ϕ parametrization leads to the form factor relation

$$\frac{|f_+^{\eta'}(q^2)|}{|f_+^{\eta}(q^2)|} = \cot\phi \tag{1.28}$$

in the limit of no annihilation, where f_+ follows the definition in Equation 1.4.

Methods for predicting $D_s \to (\eta, \eta')e\nu$ decay rates have included constituent quark models constrained by lattice results (CQM) [42], constituent quark models with dynamics defined relativistically along the $x^\pm \equiv t \pm z$ light front (LFQM) [43], light cone QCD sum rules (LCSR) [44], QCD sum rules with a phenomenological adjustment for weak annihilation (QCDS) [45], and kinematics with an intrinsic 1:3 pseudoscalar/vector rate (KIN) [46]. Preliminary lattice calculations have also given some form factor results at $q^2 = 0$ [47]. We have summarized these predictions in Table 1.1.

Table 1.1: $D_s \to \eta(')e\nu$ branching ratio predictions from different analyses. When necessary, we have used $\tau_{D_s} = 500 \times 10^{-15} s$ and $|V_{cs}| = 0.973$ for conversion to a branching ratio. The kinematics-based analysis (KIN) uses ratios with $D_s \to \phi e\nu$, so we've taken $\mathcal{B}(D_s \to \phi e\nu) = 2.0\%$ (this assumption won't affect the $\eta' e\nu/\eta e\nu$ ratio).

Analysis	$\mathcal{B}(D_s \to \eta e\nu)$	$\mathcal{B}(D_s \to \eta' e\nu)$	$\frac{\mathcal{B}(D_s \to \eta' e\nu)}{\mathcal{B}(D_s \to \eta e\nu)}$
CQM [42]	2.5%	0.93%	0.37
LFQM ($f_{D_s} = 270$ MeV) [43]	2.3%	0.91%	0.40
LCSR [44]	3.2%	0.97%	0.31
QCDS [45]	2.3%	1.0%	0.43
KIN ($\phi = 45\,^\circ$) [46]	1.4%	0.43%	0.31
ISGW2 ($\phi = 45\,^\circ$) [34]	2.7%	1.1%	0.43

1.4.2 $D_s \to \phi e\nu$

$D_s \to \phi e\nu$ decay rates and form factors have also seen a variety of predictive models, including a different approach using QCD sum rules (QCDS) [48]. We also state results from a lattice calculation that uses the Highly Improved Staggered Quark action for valence quarks and asqtad sea quarks (LAT) [26]. We've converted these decay rates

into branching ratios and added them to the CQM model's prediction and the ISGW2 model's result in Table 1.2.

Table 1.2: $D_s \to \phi e\nu$ branching ratio predictions from different analyses. When necessary, we have used $\tau_{D_s} = 500 \times 10^{-15} s$ and $|V_{cs}| = 0.973$.

Analysis	$\mathcal{B}(D_s \to \phi e\nu)$
CQM [42]	2.6%
QCDS [48]	1.4%
LAT [26]	2.4%
ISGW2 [34]	2.2%

1.4.3 $D_s \to f_0 e\nu$

$D_s \to f_0 e\nu$ decays provide a clean system for probing the scalar f_0 particle's $s\bar{s}$ component. Consequently, recent years have seen a few predictions for the $f_0 e\nu$ decay rate that depend on the f_0 mixing angle, θ, where

$$|f_0\rangle = \cos\theta\,|s\bar{s}\rangle + \sin\theta\,|n\bar{n}\rangle \tag{1.29}$$

with $|n\bar{n}\rangle = \frac{1}{\sqrt{2}}\left|u\bar{u} + d\bar{d}\right\rangle$.

Table 1.3 gives the results from these analyses, which include a light front quark model (LFQM) [49] and two different QCD sum rule methods (SUM [50] and QCDS [48]).

Table 1.3: $D_s \to f_0 e\nu$ branching ratios from different analyses. Each analysis's branching ratio depends upon the $s\bar{s}$ fraction of the f_0, given by $\cos\theta$.

Analysis	$\mathcal{B}(D_s \to f_0 e\nu)$
LFQM [49]	$4.2 \times 10^{-3}\cos^2(\theta)$
SUM [50]	$4.1 \times 10^{-3}\cos^2(\theta)$
QCDS [48]	$5.5 \times 10^{-3}\cos^2(\theta)$

1.4.4 $D_s \to (K, K^*) e\nu$

The Cabibbo suppressed $D_s \to (K, K^*) e\nu$ decays don't receive the same interest as other D_s semileptonics, presumably due to the lower branching ratios, the less ambiguous

(interesting) final meson states, and the lack of experimental data for comparison. Even so, both the CQM model and ISGW2 extend their results from the charm-strange interactions to the charm-down decays, with results shown in Table 1.4. Additionally, we have used the kinematic model (KIN) to obtain an additional prediction for the ratio of D_s semileptonic decays to K^0 and K^*.

Table 1.4: $D_s \to (K, K^*)e\nu$ branching ratio predictions from different analyses. When necessary, we've used $\tau_{D_s} = 500 \times 10^{-15} s$ and $|V_{cd}| = 0.225$ for conversion to a branching ratio.

Analysis	$\mathcal{B}(D_s \to K^0 e\nu)$	$\mathcal{B}(D_s \to K^* e\nu)$	$\frac{\mathcal{B}(D_s \to K^0 e\nu)}{\mathcal{B}(D_s \to K^* e\nu)}$
KIN [46]	—	—	0.9
CQM [42]	0.32%	0.19%	1.7
ISGW2 ($\phi = 45\,^\circ$) [34]	0.23%	0.11%	2.0

1.5 Past Results

BaBar and CLEO-c have each made recent semileptonic measurements in the D_s system. Table 1.5 contains a summary of their results. Our results include the same data as the prior CLEO-c results and in most modes (all but f_0 and ϕ) include additional data that roughly doubles the integrated luminosity. Further, our results include additional events through improved efficiency, due to a combination of dropping the D_s^* daughter photon and using generally looser particle selections.

1.6 Summary

Measured D_s semileptonic decay rates provide a cross-check on lattice calculations, give insight into light meson quark content, and allow for phenomenological comparisons using the total D_s semileptonic rate. This analysis simplifies comparisons for these purposes by measuring the six dominant D_s semileptonic decay rates with a common procedure. Further, it improves statistics in all six modes through higher efficiency and includes twice the data sample for four modes ($\eta e\nu$, $\eta' e\nu$, $K_s e\nu$, $K^* e\nu$).

Table 1.5: The most recent D_s semileptonic results from BaBar and CLEO-c. CLEO-c has published two $\phi e\nu$ and $f_0 e\nu$ results; here, we include the results with higher luminosity. BaBar saw evidence for $f_0 e\nu$ from a small S-wave component in $D_s \to KKe\nu$.

Mode	BaBar	CLEO-c
$D_s \to \phi e\nu$ [29] [51]	$(2.61 \pm 0.03 \pm 0.17)\%$	$(2.36 \pm 0.23 \pm 0.13)\%$
$D_s \to \eta e\nu$ [52]	—	$(2.48 \pm 0.29 \pm 0.13)\%$
$D_s \to \eta' e\nu$ [52]	—	$(0.91 \pm 0.33 \pm 0.05)\%$
$D_s \to K_s e\nu$ [52]	—	$(0.19 \pm 0.05 \pm 0.01)\%$
$D_s \to K^* e\nu$ [52]	—	$(0.18 \pm 0.07 \pm 0.01)\%$
$D_s \to f_0 e\nu$, $f_0 \to \pi^+\pi^-$ [29] [51]	Seen	$(0.20 \pm 0.03 \pm 0.01)\%$
Inclusive semileptonic [30]	—	$(6.52 \pm 0.39 \pm 0.15)\%$

Chapter 2

Experimental Apparatus

Once upon a time in upstate New York, the CLEO experiment detected the results of electron-positron (e^-e^+) collisions over a variety of energies. This required two major pieces of machinery: CESR, which made the electron and positron beams; and the CLEO-c detector, which measured the collision results. Wilson Synchrotron Laboratory at Cornell University in Ithaca, NY housed both the collider (which still exists) and the detector.

2.1 CESR

The Cornell Electron Storage Ring, or CESR, collided electron and positron beams to generate the particles ultimately detected in CLEO-c [54]. The 768 m circumference storage ring sat about 12 m below an athletic field on the Cornell campus and had one interaction region, located inside the detector at the ring's south end. CESR provided CLEO the highest possible luminosity (interactions per area per time) by balancing limitations on the number of charges per beam and the beams' lifetime.

2.1.1 Collider Layout

To collide two charged particle beams of a given energy, CESR had to perform three general operations: obtain the electrons and positrons used in the beams; ramp the particles up to the desired energy; and smash the beams together as many times as possible to get the most interactions. Matching these three functions, CESR had three

main components: the linear accelerator (linac), the synchrotron, and the eponymous storage ring (Figure 2.1). CESR also had a 150 kV electron gun, where low energy electrons boiled off a filament for injection into the linac, and two transfer lines connecting the synchrotron and the main storage ring.

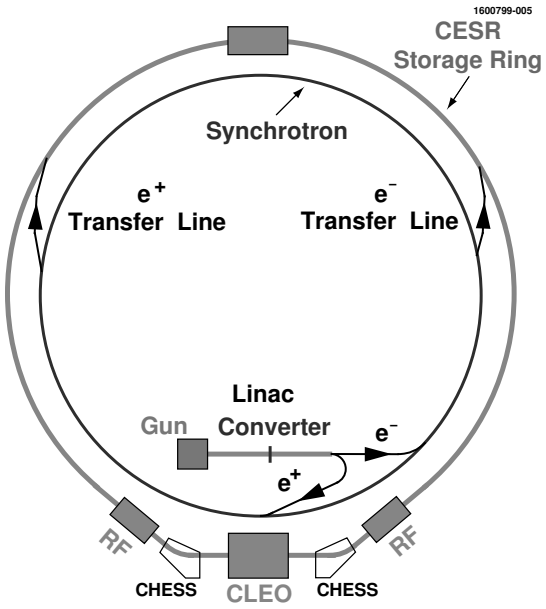

Figure 2.1: CESR schematic showing the three main components (linac, synchrotron, and storage ring) along with the electron gun and transfer lines.

CESR's linac created the electron beam by taking electrons from the electron gun and accelerating them through a series of 8 RF cavities up to an energy of around 300 MeV, then transferring them into the synchrotron. To make the positron beam, CESR inserted a tungsten target halfway down the linac to intercept the electrons. The 150 MeV electron beam hit the high-Z target, creating e^+, e^-, and γ. Magnets selected out the positrons and focused the beam, then the linac accelerated it up to around 200 MeV before injection into the synchrotron.

The synchrotron ring ran just inside the storage ring. It used 4 accelerating cavities, each about 3 m long, that ramped the beam up to CLEO's requested energy (typically around 2 GeV). This process took of order 10 ms, or a few thousand cycles around the

synchrotron. The synchrotron used several dipole magnets to steer the beams around the ring, clockwise for the positrons and counterclockwise for the electrons (as viewed from above). Once the beams had enough energy, they were injected through the appropriate transfer line into the storage ring.

Like the synchrotron, the storage ring used dipole magnets to steer the beams into a circle. Quadrupole magnets provided the primary focusing for the beams, while sextupole magnets countered the divergence that the quadrupoles introduced over the beam's small energy spread. Superconducting RF cavities running at 500 MHz added energy to compensate for the beams' synchrotron radiation losses.

Rather than having a continuous beam, the electrons and positrons were created in bunches, with as many as 5 bunches in a train and 9 bunch trains circling the storage ring at a time. The two counter-rotating beams then had 18 regions in the storage ring with potential crossing points, but only one point inside one region corresponded to the desired interaction point (the detector). CESR avoided 16 of the parasitic crossing regions through the use of 4 electrostatic horizontal separators that gave the two beams a "pretzel" orbit (Figure 2.2). The region immediately across from the detector required an additional pair of vertical electrostatic deflectors to redirect the beams away from the parasitic crossings. Finally, slightly redirecting the beams to have a small crossing angle (2.5 mrad) at the detector ensured that the interaction region had only one crossing point.

2.1.2 Beam Considerations

CESR was initially designed for beam energies near 5 GeV, with most of its running at the Υ(4S) (5.3 GeV beams). When running at energies closer to 2 GeV, the synchrotron radiation decreases significantly (power $\propto E^4$) and doesn't provide the needed damping on the beam's energy spread. CESR compensated for the lower intrinsic synchrotron radiation by adding 12 local bend magnets, or wigglers, composed of a series of alternating magnetic fields (8 poles). These fields created a tight, back-and-forth beam path through the wiggler, greatly increasing the synchrotron radiation and shortening the beam's damping time from 570 ms to 55 ms [55]. The beam size, or emittance, also increased from 30 nm-rad to 220 nm-rad, which improved beam stability by decreasing the beams' charge density through the interaction region.

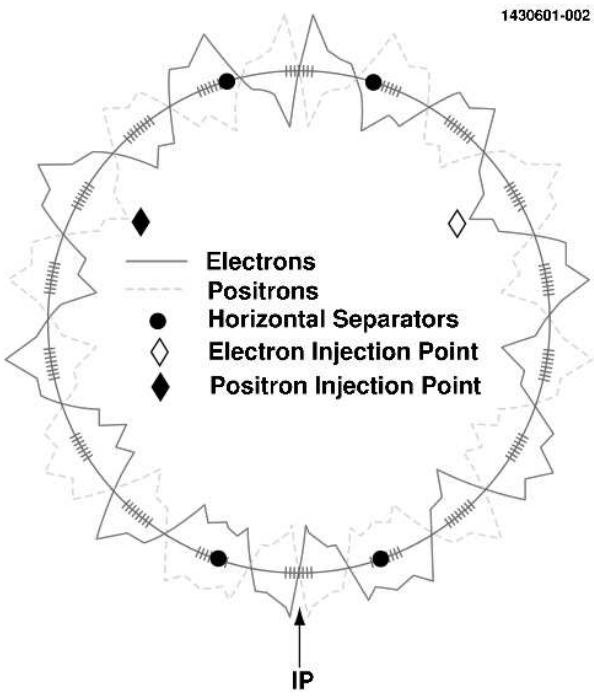

Figure 2.2: Exaggerated depiction of the pretzel orbit. Hash marks indicate the potential crossing regions.

Scattering between particles in a bunch (the Touschek effect) and interactions between the beam and gas particles in the beamline limited the beam lifetime. Intrabeam scattering depends on a variety of factors, including the beam energy, momentum spread and acceptance, and the emittance. Beam-gas scattering depends largely on the vacuum in the beamline, about 10^{-12} atm for CESR. During running, CESR also saw significant beam-beam interactions near the interaction region caused by the CLEO-c solenoid, which were largely corrected by installing two 2 T "anti-solenoids" [55].

The instantaneous luminosity ($\mathcal{L}$) in cgs units ($\mathrm{cm}^{-2}\mathrm{s}^{-1}$) may be calculated from [56]:

$$\mathcal{L} = 2.17 \times 10^{32}(1 + r^*)\frac{I\xi_y E_b}{\beta_y^*}, \tag{2.1}$$

where r^* is the vertical to horizontal beam size ratio at the interaction point, I is the beam current ($\sim$65 mA), E_b is the beam energy in GeV, ξ_y is the vertical beam-beam

parameter ($\sim$0.03), and β_y^* is the vertical focusing function at the interaction point in meters ($\sim$0.012 m). "Horizontal" in this context refers to the direction perpendicular to motion that lies in the plane of the accelerator, while "vertical" extends perpendicular to the accelerator plane. The dimensionless ξ_y depends on the damping rate and carries additional factors of the beam energy. During the D_s running, CESR attained luminosities of around 7×10^{31} cm^{-2}s^{-1}.

For a process with a given cross section, σ, the number of events comes from

$$N = \sigma \int \mathcal{L} dt, \tag{2.2}$$

where we call $\int \mathcal{L} dt$ the integrated luminosity. CESR measured the luminosity via the QED processes $e^+e^- \to e^+e^-, \mu^+\mu^-$, and $\gamma\gamma$, since those modes both have well determined cross sections and involve different detector components.

2.2 CLEO-c Detector

CLEO operated as a general purpose detector, capable of detecting both charged and neutral particles (including hadrons, leptons, and photons) with good resolution. Most of the cylindrically symmetric detector sat inside a solenoid 3.5 m long and 1.5 m in radius, with an active detector region covering 93% of the solid angle. The electron-positron beams from CESR met roughly in the detector's center at a slight crossing angle of 2.5 mrad to ensure a single interaction point for the bunch trains.

The CLEO detector saw a variety of modifications over time [54, 57, 58]. By CLEO-c, the detector held an inner drift chamber (the ZD), an outer drift chamber (the DR), a ring imaging Cherenkov detector (the RICH), a crystal calorimeter (CC), and a muon detector (Figure 2.3). The helium-cooled, superconducting solenoid generated a 1.0 T magnetic field, uniform to $\pm$0.02% over the tracking volume. The solenoid surrounded the drift chambers, the RICH, and the calorimeter. The 3 layered muon detector sat outside the magnetic field to capture muons passing through all inner material, although we will not otherwise discuss the muon chambers as they primarily functioned in a regime outside the scope of both this analysis and most CLEO-c analyses (detector applicable for muon momenta 1 GeV and greater).

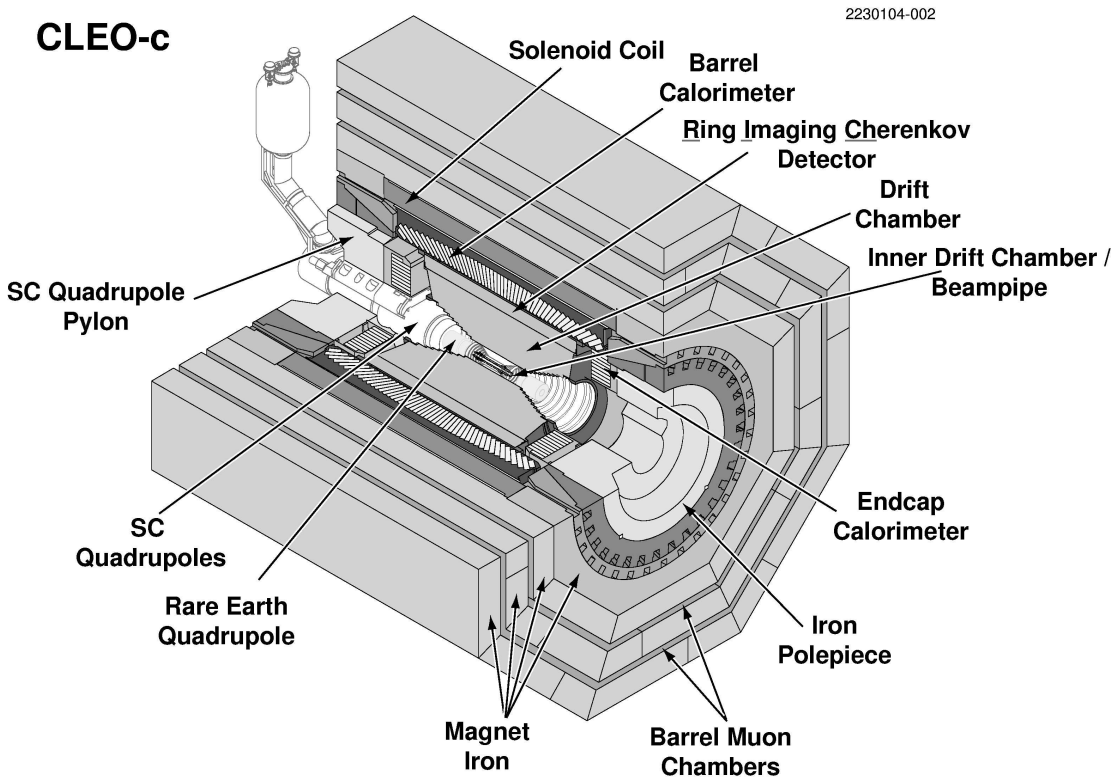

Figure 2.3: The CLEO-c detector.

2.2.1 Inner Drift Chamber

The inner drift chamber, or ZD, sat nearest the beamline inside the CLEO detector. Prior versions of CLEO used a silicon vertex detector as the innermost detector. With the lower energies of CLEO-c, however, only longer lived charged particles ($e^{\pm}, \mu^{\pm}, \pi^{\pm}, K^{\pm}, p/\bar{p}$) traveled a measurable distance, reducing some of the silicon detector's value. Further, the softer tracks required a minimum of material to avoid multiple scattering, prompting a switch from the silicon wafers to a gas and wire system.

The ZD consisted of 300 cells with an aluminum-mylar skin, each 10 mm wide (5 mm half-cell size) with a 20 μm gold-plated tungsten sense wire at +1900 V surrounded by 110 μm gold-plated aluminum field wires held at ground. When a charged particle passed through the gas, the atoms ionized, releasing an electron and leaving an ion.

While the ion didn't move much due to its mass, the electron accelerated in the strong electric field and ionized other atoms in secondary collisions. The secondary electrons then ionized further atoms and created an electron cascade. The pulse height told us the charged particle's specific ionization, or dE/dx (energy lost per length). The gas drift velocity, time of pulse, and bunch crossing time allowed us to determine the particle's position.

While CLEO originally used argon as the ionizing fill gas in its drift chambers (at that time just the main drift chamber), more recent versions switched to helium for better position resolution [59]. Propane (C_3H_8) acted as the quench gas to avoid spurious signals from photon emission in atoms that became excited instead of ionized. CLEO-c settled on a 60:40 helium-propane mixture in both the inner and the main drift chambers, which had a long radiation length of 330 μm (important since multiple scattering dominated momentum resolution at CLEO-c energies).

2.2.2 Main Drift Chamber

As the inner drift chamber had been modeled on the main drift chamber, the two shared many similarities. The main drift chamber had a slightly larger cell size of 14 mm (7 mm maximum drift distance), with the sense wire held at +2100 V. Corresponding to its larger size, the main drift chamber held 47 layers of field and sense wires rather than the inner drift chamber's 6 layers. The first 16 wire layers (1696 wires in total) ran parallel to the beam axis (axial layers), while the remaining layers (8100 wires) alternated in stereo angle.

As a charged particle passed through the main drift chamber, the magnetic field bent it into a helical path. A Kalman (Billoir) fitter reconstructed the wire hits from ionization into this path, incorporating dE/dx information. The path's curvature in the magnetic field then allowed us to determine the particle's momentum.

Aside from being useful in its own right, the momentum measurement combined with the particle's dE/dx to identify the charged particle. The deposited dE/dx depends upon a particle's velocity through the medium. For momenta comparable to or less than the particle's mass (velocities distinguishable from the speed of light), dE/dx and the momentum gave enough combined information to determine the particle's mass.

In the case of pions and kaons (the two most common long-lived hadrons in CLEO-c),

this dE/dx technique gave very good separation up to momenta near the kaon mass, as shown in Figure 2.4. In this analysis, we frequently used the deviation from expected dE/dx under a given particle mass hypothesis, $\sigma_{\pi/K}^{dE/dx}$, to identify (or reject) charged tracks.

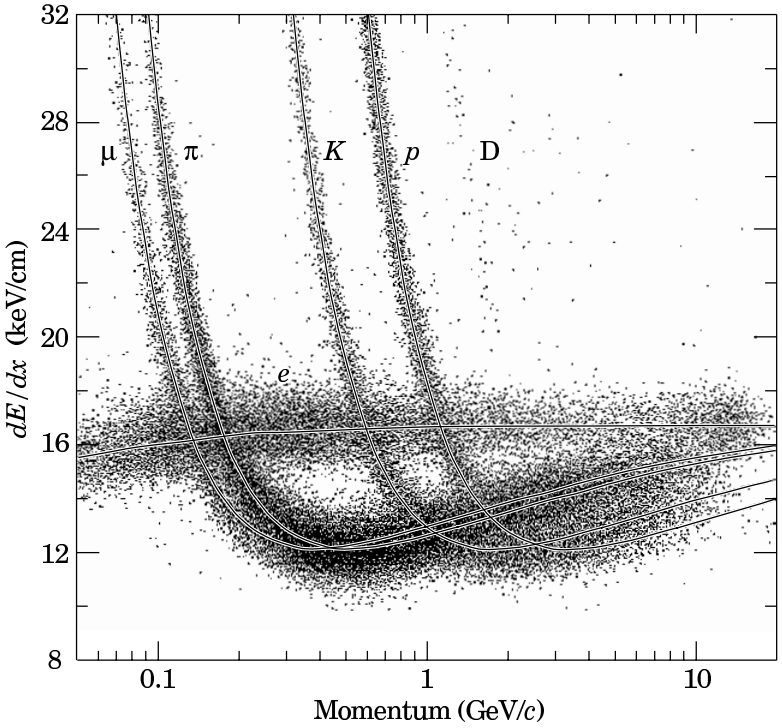

Figure 2.4: dE/dx through the drift chamber as a function of momentum for the long-lived charged particles. Below 600 MeV or so, kaons and pions have good dE/dx separation.

2.2.3 Ring Imaging Cherenkov Detector

While the drift chamber provided good particle identification for momenta below ∼600 MeV, dE/dx did not give sufficient separation to identify charged particles with high momenta. To remedy this flaw, a prior version of CLEO (CLEO III) added a ring imaging Cherenkov detector (Figure 2.5) between the main drift chamber and the crystal calorimeter's barrel, covering about 83% of the total solid angle. The RICH detected

Cherenkov radiation produced when charged particles moved faster than the phase velocity of light in the dielectric radiator ($p > 440$ MeV for the kaon threshold).

The dielectric consisted of LiF crystals about 1.7 cm thick laid out in 14 rows. The outer crystal rows had flat surfaces, but the central four rows had a sawtooth shape to avoid total internal reflection of the 135 nm - 165 nm Cherenkov photons. The emitted photons passed through an expansion gap filled with nitrogen (N_2) to avoid difficulties in working with a vacuum while still being transparent to the VUV photons. Following the gap, the photons passed through CaF_2 windows into a multi-wire proportional chamber filled with methane-TEA (triethylamine), where the photons converted into photoelectrons and created an ionizing cascade much like the drift chambers.

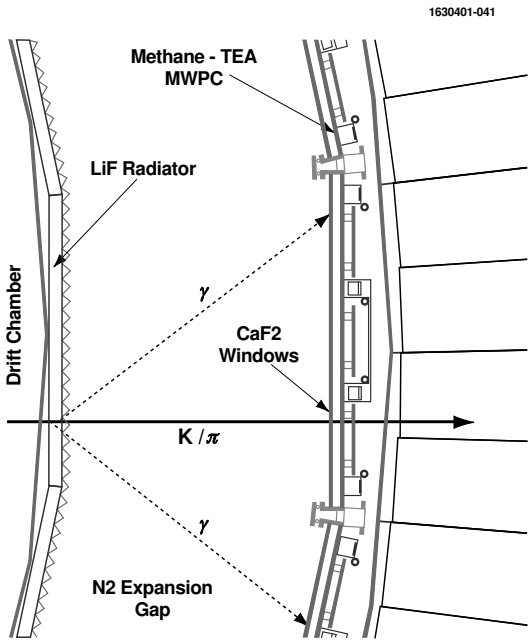

Figure 2.5: CLEO-c RICH detector schematic, including the Cherenkov photon cone from the passage of K/π through the radiator. The radiators here have the sawtooth shape.

The RICH determined a charged particle's mass by measuring its velocity through the Cherenkov effect (Figure 2.6) and comparing it with the momentum from the drift

chamber. When above threshold, the dielectric medium (LiF radiators) emitted light at an angle defined by the charged particle's velocity ($\cos\theta = \frac{1}{\beta n}, n_{\mathrm{LiF}} = 1.4$). Measuring the angle of emitted light then determined the velocity and allowed a likelihood distribution with good separation between different mass hypotheses (Figure 2.7).

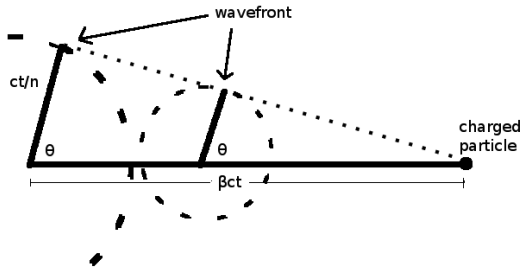

Figure 2.6: When a charged particle moves faster than the speed of light in a material, light emitted at the critical angle $\cos\theta = \frac{ct/n}{\beta ct} = \frac{1}{\beta n}$ sees no interference at the common tangent wavefront (the common envelope). For a thin piece of material, this results in a thin cone of Cherenkov radiation at that angle.

2.2.4 Calorimeter

While the drift chambers and the RICH focused solely on charged particles, the crystal calorimeter detected both charged and neutral particles. The calorimeter energy information proved particularly useful in detecting photons and electrons, both of which we use in this analysis.

The crystal calorimeter sat outside the RICH but still inside CLEO's magnetic field. The full calorimeter used nearly 7,800 thallium-doped cesium iodide crystals, each 30 cm long with a 5 cm × 5 cm face. The crystals' size gave the calorimeter an angular resolution of roughly 10 mrad. About 80% of the crystals formed the calorimeter's barrel region ($|\cos\theta| < 0.80$), with the remainder in one of the two endcaps ($0.85 < |\cos\theta| < 0.93$) or in the barrel-endcap transition region ($0.80 < |\cos\theta| < 0.85$). Material from the RICH and the main drift chamber obstructed the transition region, so most analyses (including this one) rejected data from that region (Figure 2.9).

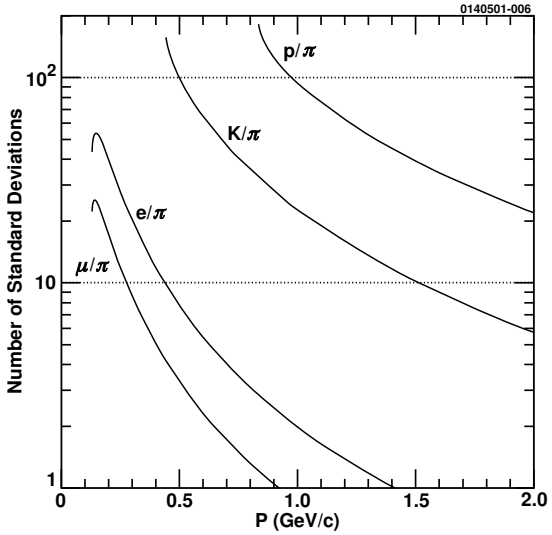

Figure 2.7: RICH separation for different particle hypotheses, by momentum. K and π see good separation at momenta above the RICH threshold (near 500 MeV).

The crystal calorimeter's usefulness in identifying electrons/positrons and photons came from the particles' propensity to give up all their energy to the calorimeter. The small electron/positron mass meant that they deposited an energy roughly equal to their momentum ($E/p \approx 1$), distinguishing $e^{\pm}$ from other charged particles. Photons similarly deposited their energy into the crystals, with their momentum direction determined by the shower location in the calorimeter.

With a CsI radiation length of 1.86 cm [21], photons underwent e^+e^- pair creation shortly after entering the crystal (photon mean free path of 9/7 the radiation length). The newly created electrons and positrons then emitted photons via bremsstrahlung. These new photons also underwent pair creation, with those charged particles then emitting further bremsstrahlung, and so on, resulting in an electromagnetic shower. Incident electrons underwent an identical process, albeit without pair creation as the first step.

The electromagnetic shower continued until the daughter particles' average energy fell low enough that ionization by the charged particles dominated bremsstrahlung. Ionized electrons and their associated holes formed weakly bound Coulomb pairs, or

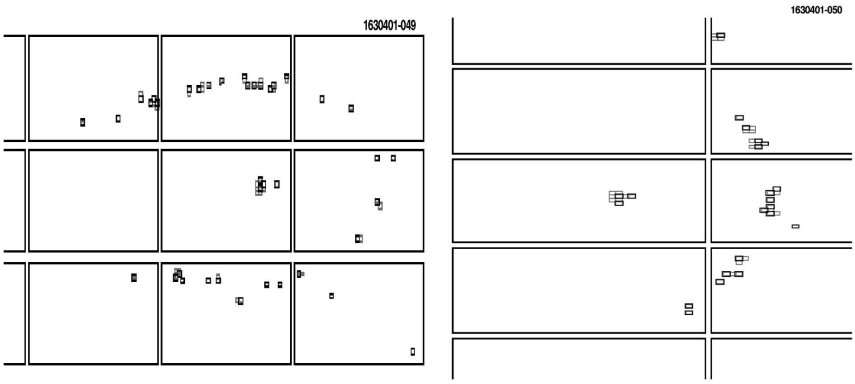

Figure 2.8: Hits in the RICH from tracks passing through the sawtooth (left) and flat (right) LiF radiators. A ring of Cherenkov photon hits surrounds the charged particle's passage. Only half of the ring appears in the right image, as total internal reflection trapped the other half of the Cherenkov photons in the radiator.

excitons, which moved through the crystal until captured by the thallium atoms. The thallium atoms then de-excited from their new metastable states via the emission of visible 560 nm photons, which passed through the transparent crystal.

As the magnetic field inhibited the use of photomultiplier tubes, CLEO used four silicon photodiodes to collect the light at the end of each crystal. Each photodiode had a separate preamplifier, with outputs combined by a mixer/shaper.

With a nuclear interaction length just over 39 cm (compared to a crystal length of 30 cm), hadrons (π, K) passing through the calorimeter sometimes formed hadronic showers. In this case, the hadron underwent a nuclear interaction with the crystal and generated new hadrons, including π^0 that decayed to photons and deposited energy into the calorimeters through the electromagnetic shower process. These showers don't have a particularly well defined energy deposition, so we didn't use them for particle identification (beyond rejecting the electron hypothesis). They did, however, sometimes generate "splitoff" showers in other areas of the calorimeter that otherwise looked like photons (a frequent complication for this analysis given its low energy photons, like the

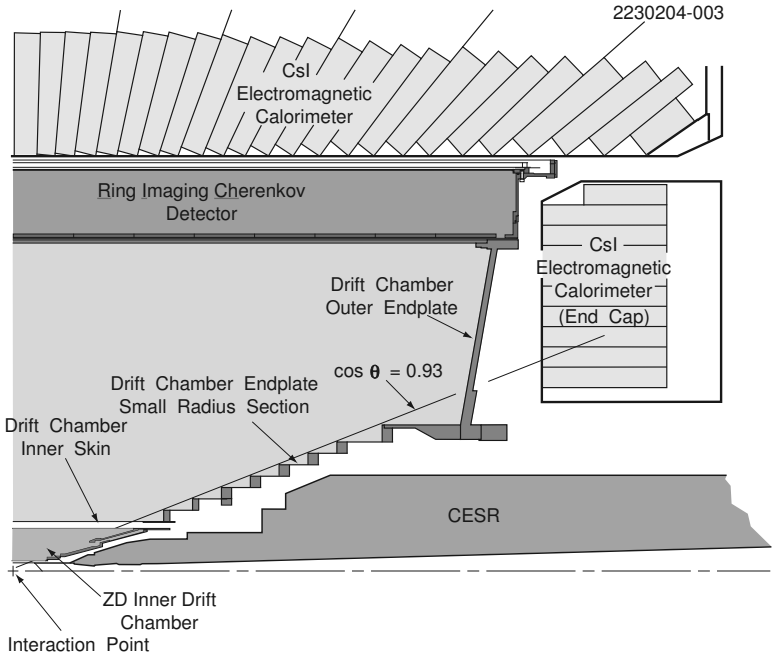

Figure 2.9: A side view of the CLEO-c detector.

D_s^* daughter).

Muons and hadrons that didn't undergo a hadronic shower still deposited energy in the calorimeter through ionization. Such particles over a small energy threshold (a few hundred MeV for μ/π) passed through the calorimeter entirely, depositing just under 200 MeV in the calorimeter as minimum ionizing particles (mips). While useful for calorimeter calibration, we did not need to directly make use of mips for this analysis.

2.2.5 Trigger and Data Acquisition

Under the best running conditions (45 bunches/train, perfect crossings), CESR could deliver bunch crossings near 20 MHz. However, with a luminosity of 10^{32} cm^{-2}s^{-1} and a Bhabha cross section near 500 nb, even vaguely meaningful crossings only occurred at rates near 20 Hz (with "interesting physics," including continuum, less than a tenth that at 4170 MeV). Since no new data could be acquired during an event read-out and CLEO's data acquisition (DAQ) read-out rate peaked at around 80 Hz during data

taking, CLEO's trigger had to efficiently separate meaningful crossings from all possible crossings.

The trigger required a 100 ns time resolution for the interaction, so it binned the data into 42 ns wide slices, three times the 14 ns bunch spacing. During this period, the main drift chamber and the crystal calorimeter both read out limited hit information, time-aligned for the latency of roughly 2 μs on tracking and 2.5 μs on the calorimeter. The number of tracks, number of showers, and event topology determined whether an event triggered a full read-out (Table 2.1). CLEO's selection methodology resulted in the needed background rejection while retaining a 99% single track efficiency.

Table 2.1: CLEO trigger line definitions. "Axial" refers to the 16 drift chamber layers that run along the beam line, while "stereo" refers to the drift chamber's stereo section. Low, medium, and high refer to the shower's calorimeter energy, while "B" and "E" refer to the barrel and endcap calorimeter regions, respectively.

Name	Definition
Hadronic	$(N_{\text{axial}} > 1)$ and $(N_{\text{B low}} > 0)$
μ-pair	two back-to-back stereo tracks
Barrel Bhabha	back-to-back high showers in B
Endcap Bhabha	back-to-back high showers in E
Electron + track	$(N_{\text{axial}} > 0)$ and $(N_{\text{B med}} > 0)$
τ/Radiative	$(N_{\text{stereo}} > 1)$ and $(N_{\text{B low}} > 0)$
Two track	$N_{\text{axial}} > 1$
Random	Random 1 kHz source

Once a trigger passed, CLEO had to digitize some 400,000 detector channels for the event. CLEO managed this with an average read-out time around 30 μs, leaving a small dead-time of around 0.3% while the DAQ wrote to disk. Each event averaged around 10 kB of data, with a throughput rate to disk of order 1 MB/s.

In addition to the data transfer from front end electronics to disk, the DAQ also contained a slow control system. Slow control included alarm handling, run control, calibration constants, and data quality checks. A fraction of the reconstructed events went into "pass1" for the data quality checks. Pass1 included both an online version (Online-pass1) for rapid checks and an offline version (Caliper) that included the ability to isolate interesting physics events by making cuts.

Table 2.2: Resolution and solid angle coverage for various CLEO-c detector components [54].

Component	Coverage and performance summary
Drift chamber	93% of 4π; $\sigma_p/p = 0.35\%$ at $p = 1$ GeV
RICH	80% of 4π; 87% kaon efficiency with 0.2% pion fake at $p = 0.9$ GeV
Calorimeter	93% of 4π; $\sigma_E/E = 4.0\%$ (2.2%) at E = 100 MeV (1 GeV)
Muon chambers	85% of 4π for $p > 1$ GeV

Chapter 3

Analysis Plan

We intend to measure branching ratios for six D_s semileptonic decays ($D_s \to \phi e\nu$, $\eta e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$). These cover all resonant D_s semileptonic decays up to the singly Cabibbo suppressed level. We use CLEO-c's 4170 MeV data, where 95% of the D_s sample comes from $D_s^* D_s$ events [60, 61], and the remainder come from $D_s^+ D_s^-$. The D_s^* decays to $D_s\gamma$ nearly all the time (94%) [21], with $D_s^* \to D_s\pi^0$ making up the difference. Candidate D_s semileptonic events then contain one D_s^+, one D_s^-, and either zero, one, or two photons.

For all six modes, we reconstruct the nonsemileptonic D_s through one of 13 "tag" modes. We also reconstruct the semileptonic side's electron and hadron (ϕ, η, η', f_0, K_s, or K^*). We do not attempt to reconstruct the photon(s) from a possible D_s^* decay, which increases our overall efficiency but costs us a clean neutrino missing mass.

We use the following 13 D_s tag modes to determine candidate events: $K_s K$; $KK\pi$; $K_s K\pi^0$; $K_s K_s \pi$; $KK\pi\pi^0$; $K_s K^+\pi\pi$; $K_s K^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. Of these modes, $K_s K$, $KK\pi$, $KK\pi\pi^0$, $\pi\pi\pi$, $\pi\eta$, $\pi\eta', \eta' \to \pi\pi\eta$ give the most statistical power as they contain over 60% of the tags and have lower relative background than the other 7 modes.

The signal modes $D_s \to \phi e\nu$ and $D_s \to \eta e\nu$ should have the largest branching ratios (around a couple percent) due to their $s\bar{s}$ components, about an order of magnitude higher than the Cabibbo suppressed decays to $K_s e\nu$, $K^* e\nu$, and $f_0 e\nu$. We expect $D_s \to \eta' e\nu$ to have a branching ratio somewhere in between these levels.

Ignoring the D_s^* daughter photons does have the limitation that the branching ratio

results cannot be easily converted into $\frac{d\Gamma}{dq^2}$. Consequently, we do not determine form factors in this analysis.

Chapter 4

Data Samples and Monte Carlo

We use the CLEO-c data collected at a 4170 MeV center-of-mass energy (datasets 39, 40, 41, 47, and 48), with a total integrated luminosity of 586 pb^{-1} (Table 4.1). At this energy, the e^+e^- collisions produce both D_sD_s and $D_s^*D_s$. $D_s^*D_s$ production has the much larger cross section at 4170 MeV, with a $\sigma_{D_s^*D_s}$ of 0.92 nb and a $\sigma_{D_s^+D_s^-}$ of 0.03 nb [60]. With these cross sections, we expect to have about 1.11×10^6 D_s mesons in the data sample.

Table 4.1: Dataset luminosities determined from Bhabha events ($e^+e^- \to e^+e^-$), with statistical and systematic errors.

Dataset	Integrated lum. (pb^{-1})
39	55.12 ± 0.03 ± 0.55
40	123.93 ± 0.05 ± 1.24
41	119.11 ± 0.05 ± 1.19
47	109.78 ± 0.05 ± 1.10
48	178.23 ± 0.06 ± 1.78
Total	586.17 ± 0.11 ± 5.86

We use CLEO's generic `ddmix` Monte Carlo for each of the 4170 MeV datasets, which generates a weighted mixture of all charm production ($DD, D^*D, D^*D^*, D_sD_s, D_s^*D_s$). Each `ddmix` sample has 20× the data luminosity, and in all cases we use the final regenerated sample from CLEO's `20080404_MCGEN_1` release (which includes ISR). The continuum, radiative return, and tau Monte Carlo samples used for backgrounds each

simulate 5× the data luminosity, with a release that differs by dataset. Table 4.2 lists these releases. For all six semileptonic modes, the backgrounds from charm mesons dominate the backgrounds from continuum, radiative return, and tau production.

Table 4.2: CLEO code release used to generate each background MC sample.

Dataset	Release
39	20060426_MCGEN
40	20060426_MCGEN_2
41	20060426_MCGEN_2
47	20080404_MCGEN_1
48	20080404_MCGEN_1

`EVTGEN` [62] generated the charm, continuum, and radiative return samples, while `QQ` generated the tau samples. The continuum generation used the Lund area law generator.

We have also created four signal Monte Carlo samples for each of our six semileptonic modes ($\phi e\nu$, $\eta e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$). The four samples correspond to different D_s production modes at 4170 MeV: $D_s D_s$; $D_s^* D_s$ with $D_s^* \to \gamma(D_s \to he\nu)$; $D_s^* D_s$ with $D_s^* \to \pi^0(D_s \to he\nu)$; and $D_s^* D_s$ with the prompt $D_s \to he\nu$. We generated 250,000 events for each of our $\phi e\nu$ and $\eta e\nu$ signal Monte Carlo samples, with 100,000 events for the other semileptonic mode samples.

We processed both data and Monte Carlo with the `20060224_FULL_A_3` release to maintain consistency with version 2 D skims [63, 64].

Chapter 5

D_S Tagging

CLEO-c produces nearly all of its D_s sample at a 4170 MeV center-of-mass energy. While this energy gives the most D_s mesons, the total D_s cross section at 4170 MeV still falls slightly below 1.0 nb. That compares to a 9 nb total charm cross section (dominated by D^*D^* at 4.7 nb and D^*D at 2.6 nb) [61] and a 12 nb continuum cross section [21, 60] at 4170 MeV. Since the lost neutrino prevents us from reconstructing the entire D_s semileptonic decay, leaving only the hadron (ϕ, η, η', f_0, K_s, or K^*) and electron, we need to find some other way to isolate D_s events lest we be smothered by combinatoric background and other decays containing an electron and target hadron.

As all D_s events contain two D_s mesons, we can entirely reconstruct, or *tag*, one of the mesons as an event requirement for D_s semileptonic decay candidates. When measuring branching ratios, D_s tagging yields the further benefit that the measured tags directly provide the branching ratio denominator rather than needing to independently estimate the data sample's total D_s counts.

We use a total of 13 D_s tag modes, chosen for their relative abundancy and their separation from combinatoric background. We have mostly chosen to identify the tag modes by their final state particles rather than their intermediate particles (e.g. $KK\pi$ instead of $\phi\pi$ or K^*K). This choice maintains consistency with previous CLEO work [65] and avoids the need to worry about overlapping resonances (particularly a concern for $KK\pi$, the most statistically significant mode).

We reconstruct the following 13 tag modes: $D_s^+ \to K_sK^+$, $D_s^+ \to K^+K^-\pi^+$, $D_s^+ \to K_sK^+\pi^0$, $D_s^+ \to K_sK_s\pi^+$, $D_s^+ \to K^+K^-\pi^+\pi^0$, $D_s^+ \to K_sK^+\pi^+\pi^-$, $D_s^+ \to$

$K_s K^- \pi^+ \pi^+$, $D_s^+ \to \pi^+ \pi^+ \pi^-$, $D_s^+ \to \pi^+ \eta$, $D_s^+ \to \pi^+ \pi^0 \eta$, $D_s^+ \to \pi^+ \eta'$ with $\eta' \to \pi^+ \pi^- \eta$, $D_s^+ \to \pi \pi^0 \eta'$ with $\eta' \to \pi^+ \pi^- \eta$, and $D_s^+ \to \pi^+ \eta'$ with $\eta' \to \rho^0 \gamma$. Here, and elsewhere, the charge conjugate tag modes are also implied.

Once we have our D_s tag candidates, we determine each mode's tag counts by fitting their invariant mass. Any event with a D_s tag passing our wide mass window gets treated as a semileptonic decay candidate.

5.1 Basic Selection Criteria

We use a common selection criteria for daughter particles in our 13 exclusive tag modes. We have found little gain in deviating from the standard D-tag cuts, so our selection criteria emulates those selections [64, 66].

5.1.1 Track Selection

Our tag modes include two charged particles that leave tracks: kaons and pions. Our selection for both $K^\pm$ and $\pi^\pm$ have several track quality features in common:[1]

- $|d_b| < 5$ mm
- $|z_0| < 5$ cm
- $\chi^2 < 100,000$
- $|\cot\theta| < 2.53$ (equivalent to $|\cos\theta| < 0.93$)[2]
- Hit Fraction > 0.5

For $K^\pm$, we further require:[3]

- 0.125 GeV $< p_K <$ 2.0 GeV
- $\left|\sigma_K^{dE/dx}\right| < 3.0$

[1] d_b, z_0, dE/dx, the hit fraction, and splitoff are all defined in the glossary (Appendix E). The χ^2 here comes from the track fit, and its requirement only serves to throw out particularly egregious tracks.
[2] Here, θ represents the angle from the beamline.
[3] $p_K^{\min}$ raised from 0.050 GeV to 0.125 GeV for better dE/dx agreement between data and Monte Carlo [65]

- `UsePID true`

We use the standard CLEO parameters for `UsePID`. Specifically, if we have RICH information with both π and K hypotheses analyzed, $p > 0.7$ GeV, and $|cos(\theta)| > 0.8$, then we combine the RICH likelihood and $\sigma^{dE/dx}$ by requiring:

- $\mathcal{L} \equiv (\sigma_\pi^2 - \sigma_K^2) + (L_\pi - L_K) \geq 0$
- At least four RICH photons detected ($N_\gamma^{RICH} > 3$)

Otherwise, we just use dE/dx values by requiring $(\sigma_\pi^2 - \sigma_K^2) \geq 0$ [66].

Similarly, for $\pi^\pm$ we require:

- 0.050 GeV $< p_\pi <$ 2.0 GeV
- $\left|\sigma_\pi^{dE/dx}\right| < 3.0$
- `UsePID true`

The `UsePID true` here matches that for the kaons, although now we require $\mathcal{L} \leq 0$.

We only apply these track cuts to the tag mode daughter particles. The daughters of the semileptonic hadrons have their own similar, but often looser, selection criteria.

5.1.2 K_s Selection

We make a 1.575σ mass cut on our tag modes' K_s mesons. This corresponds to a 6.3 MeV nominal mass cut. Our K_s mesons' $\pi^\pm$ daughters don't have to fulfill the standard π cuts listed in Section 5.1.1 since they don't necessarily originate from the interaction point. Given CLEO-c's lower energies than earlier CLEO analyses, we do not use the `CleanV0` cuts, nor do we add a flight significance or distance cut.

5.1.3 Photon Selection

Several tag modes include particles like π^0 and η that ultimately decay to photons. Also, $D_s^+ \to \pi^+\eta', \eta' \to \rho^0\gamma$ has an explicit photon in the tag mode. These photons share several different selection criteria:

- $E_\gamma > 30$ MeV

- No $\frac{E_9}{E_{25}}$ cut[4]
- No splitoff rejection used
- Showers with a matched track are disallowed
- Showers from hot crystals are disallowed

5.1.4 π^0 and η Selection

In addition to the daughter photon selection, we also consider some additional selection criteria for the π^0 and η mesons used in D_s tags:

- The pull mass for both π^0 and η needs to be within 3.0
- We do not reject π^0 that have both showers in the endcap
- We do reject η that have both showers in the endcap
- Nominal mass less than 1.0 GeV
- Max number of σ from expected mass within 1,000
- $\chi^2 \leq 10,000$
- No additional energy cut on γ in the endcap

5.1.5 η' Selection

The tag modes include η' reconstructed from its $\pi\pi\eta$ decay mode and from its $\rho^0\gamma$ decay mode, where the ρ^0 decays to $\pi^+\pi^-$. Each of these decay modes has additional selection criteria.

The $\eta' \to \pi\pi\eta$ mode involves reconstructing both pions and η mesons. We use the same selection criteria for these as in Section 5.1.1 and Section 5.1.4, respectively. Additionally, we require 947.8 MeV $< M_{\eta'} <$ 967.8 MeV

$\eta' \to \rho^0\gamma$ ultimately involves reconstructing two pions and a photon. Again, pions share the same selection criteria as in Section 5.1.1. The photon inherits our standard photon tagging selection. We further require:

[4]Energy in a shower's $3 \times 3 = 9$ central crystals divided by the energy in a shower's $5 \times 5 = 25$ central crystals.

- 0.5 GeV $< M_{\pi^+\pi^-} <$ 1.0 GeV
- 920 MeV $< M_{\eta'} <$ 995 MeV

5.2 Recoil Mass

The tagged D_s mesons are only created in either $D_s^* D_s$ or $D_s D_s$ events, which constrains their momentum range. The direct D_s momenta depend only on the beam energy (4170 MeV), while the secondary D_s from the D_s^* decay gain a slight boost. Table 5.1 gives the kinematic ranges for D_s momenta at our beam energy.

We restrict the allowed momentum range by cutting on a directly related variable, the recoil mass, which includes the beam momentum and corresponds physically to the other meson's mass in the case of prompt D_s decays. We define the recoil mass by

$$M_{\text{recoil}} = |p_{cm} - p_{D_s}| \equiv \sqrt{\left(E_{cm} - \sqrt{|\vec{p}_{D_s}|^2 + M_{D_s}^2}\right)^2 - |\vec{p}_{cm} - \vec{p}_{D_s}|^2},$$

where p_{cm}, E_{cm}, and $\vec{p}_{cm}$ correspond to the center-of-mass four vector, energy, and momentum; M_{D_s} comes from the PDG [21]; and $\vec{p}_{D_s}$ denotes the reconstructed D_s momentum.

Since we will use both the recoil mass and the D_s invariant mass, we do not use either the beam constrained mass (M_{bc}) or $\Delta E \equiv E_{D_s} - E_{beam}$.

Table 5.1: Kinematically allowed recoil mass and momentum ranges for D_s mesons at 4170 MeV.

D_s origin	Possible Momenta	Possible Recoil Mass
$D_s D_s$	687 MeV	$M_{D_s} \approx 1968$ MeV
Prompt D_s in $D_s^* D_s$	429 MeV	$M_{D_s^*} \approx 2112$ MeV
D_s from $D_s^* \to D_s\gamma$ in $D_s^* D_s$	259 MeV – 542 MeV	2058 MeV – 2169 MeV
D_s from $D_s^* \to D_s\pi^0$ in $D_s^* D_s$	351 MeV – 449 MeV	2104 MeV – 2142 MeV

Our recoil mass cut varies by tag mode and depends upon the shape and combinatoric background for that mode. Table 5.2 gives our cut values by tag mode.

Figure 5.1 and Figures G.1–G.3 in Appendix G show each tag mode's recoil mass spectrum and recoil mass cut in the Monte Carlo and data. The background reduction

Table 5.2: Recoil mass cut, by D_s tag mode.

D_s tag modes	Minimum Recoil Mass	Maximum p_{D_s}
K_sK $KK\pi$ $\pi\eta$ $\pi\eta', \eta' \to \pi\pi\eta$	2051 MeV	555 MeV
$K_sK\pi^0$ $K_sK_s\pi$ $KK\pi\pi^0$ $K_sK^+\pi\pi$ $K_sK^-\pi\pi$ $\pi\pi^0\eta$ $\pi\pi^0\eta', \eta' \to \pi\pi\eta$ $\pi\eta', \eta' \to \rho\gamma$	2099 MeV	459 MeV
$\pi\pi\pi$	2101 MeV	455 MeV

from our recoil mass cut benefits us across the D_s invariant mass spectrum, as the two variables are fairly uncorrelated (shown in Figures G.4–G.9).

5.3 Individual Tag Mode Cuts

Each tag mode has unique backgrounds that we reduce by making a series of additional cuts. These cuts reject D^0 or $D^\pm$ mesons, reject unwanted K_s, or remove excess (and often peaking) combinatoric background arising from soft pions. We chose these cuts to maintain consistency with previous CLEO D_s tagging [65] when applicable. Table 5.3 lists our rejection criteria for each tag mode. We take charges into consideration for our stated invariant mass rejections (e.g. for K_sK, we only apply the $M_{K\pi}$ rejection when the K and π have opposite signs).

After making our previously listed cuts, we choose a best candidate for each tag mode and charge by keeping only the D_s with a recoil mass closest to $M_{D_s^*}$ (2112.3 MeV).

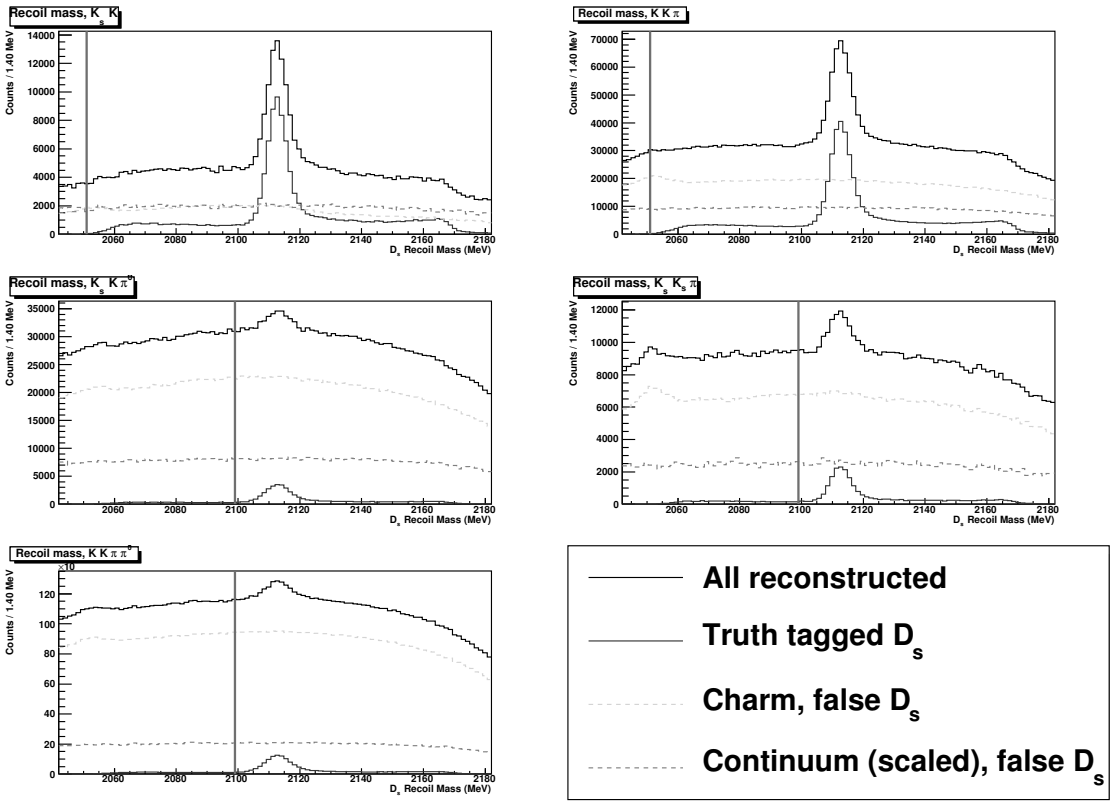

Figure 5.1: Monte Carlo (charm and scaled continuum) simulation of the recoil mass distribution for the D_s tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We reject all D_s with a recoil mass below the cut line.

Table 5.3: Rejections based on individual tag modes' daughter particles.

D_s tag mode	Daughter particle cuts (rejections)
K_sK	$M_{K\pi} \geq 1830$ MeV $p_\pi \leq 100$ MeV
$KK\pi$	1845 MeV $\leq M_{KK} \leq$ 1880 MeV
$KK\pi\pi^0$	1860 MeV $\leq M_{KK\pi} \leq$ 1880 MeV $p_\pi \leq 100$ MeV $p_{\pi^0} \leq 100$ MeV
$\pi\pi\pi$	475 MeV $\leq M_{\pi\pi} \leq$ 520 MeV 1840 MeV $\leq M_{\pi\pi} \leq$ 1885 MeV 1845 MeV $\leq M_{K\pi} \leq$ 1880 MeV, where one of the supposed pions is treated as a kaon. $p_\pi \leq 100$ MeV
$\pi\eta$ $\pi\eta', \eta' \to \pi\pi\eta$	No additional cuts
$K_sK\pi^0$	$M_{K\pi\pi^0} \geq 1830$ MeV 1860 MeV $\leq M_{K\pi\pi} \leq$ 1880 MeV $p_\pi \leq 100$ MeV $p_{\pi^0} \leq 100$ MeV
$K_sK_s\pi$ $K_sK^+\pi\pi$ $K_sK^-\pi\pi$	$p_\pi \leq 100$ MeV, where the D_s is rejected if either the K_s meson's daughter π or the direct π fails the cut.
$\pi\pi^0\eta$	$p_{\pi^0} \leq 100$ MeV
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	480 MeV $\leq M_{\pi\pi} \leq$ 515 MeV $p_{\pi^0} \leq 100$ MeV
$\pi\eta', \eta' \to \rho\gamma$	480 MeV $\leq M_{\pi\pi} \leq$ 515 MeV $p_\pi \leq 100$ MeV

5.4 Fitting Procedure

Once we make the tag cuts and choose a best candidate for each mode/charge combination in the event, we determine the tag counts for a given mode by fitting its D_s mass spectrum (combining D_s^+ and D_s^-). We model our signal mass spectrum by a double gaussian[5] for the tag modes reconstructed with only tracks (and for $\pi\eta', \eta' \to \pi\pi\eta$). We use the sum of a gaussian and a crystal ball function for modes containing photons or nontrivial FSR ($\pi\pi\pi$). We take either a linear function or a quadratic function for our background, depending upon the mode and the shape of its combinatoric background. Table 5.4 lists the particular combination of fit functions for each tag mode.

Table 5.4: D_s invariant mass fit functions, by mode.

D_s tag mode	Function
K_sK $KK\pi$ $K_sK_s\pi$ $K_sK^+\pi\pi$ $K_sK^-\pi\pi$ $\pi\eta', \eta' \to \pi\pi\eta$	Signal: Double Gaussian Background: Linear Polynomial
$K_sK\pi^0$ $\pi\eta$ $\pi\pi^0\eta', \eta' \to \pi\pi\eta$	Signal: Gaussian + Crystal Ball Background: Linear Polynomial
$KK\pi\pi^0$ $\pi\pi\pi$ $\pi\pi^0\eta$ $\pi\eta', \eta' \to \rho\gamma$	Signal: Gaussian + Crystal Ball Background: Quadratic Polynomial

In our signal shape functions, we use a common mean for the two gaussians (or for the gaussian and the gaussian portion of the crystal ball). To reduce the number of free parameters further, we fit our signal shape to a truth-tagged Monte Carlo sample, then we use those results to fix the relative normalization and relative width of the two signal shape component functions. We also fix the two remaining shape parameters in the crystal ball function from the truth-tagged fit, if applicable for the mode. This procedure leaves three free parameters for the signal shape of the reconstructed D_s mass

[5]Sum of two gaussian functions.

spectrum: an overall normalization, an overall width, and the common mean. Combined with the two or three background parameters for the linear or quadratic polynomial, respectively, we end with five or six free parameters for the reconstructed D_s mass fit.

5.5 Results

In the following sections, we present the results from our fits to the reconstructed D_s invariant mass spectrum in both Monte Carlo and data. We only consider statistical errors on the tag counts here. We do consider systematics associated with our tag counting procedure in Section 10.1.1, but we focus on the branching ratio's systematic error from tagging rather than on the error for raw tag counts. We typically get smaller tag-related systematics on the branching ratio than on tag counting alone because our procedure involves the D_s tag shape in both the branching ratio's numerator and denominator. Any comparison with other D_s tagging analyses should keep in mind that our raw tag counts would presumably have a higher systematic error than reflected in just our branching ratio systematics.

As mentioned in Section 5.4, we first fit the truth-tagged Monte Carlo M_{D_s} distribution to fix all but three parameters for our signal shape function. Figures G.10–G.12 show these fits' results, by tag mode. The fit functions closely match the truth-tagged histograms, which gives us the freedom to use our functions rather than less flexible signal histograms when fitting the data.

5.5.1 Monte Carlo

Before we fit the data, we first build confidence in our procedure by ensuring that we get the proper tag counts in the generic Monte Carlo sample (charm plus continuum). Figures G.13–G.15 show our fits to the D_s invariant mass for this sample, resulting in the total tag counts displayed in Table 5.5. While we used a 20× Monte Carlo sample, we have scaled the table's tag counts down to the data's luminosity to make direct comparisons with the data counts easier.

Fits to the modes $\pi\pi^0\eta', \eta' \to \pi\pi\eta$ and $\pi\eta$ gave the most significant deviations from their truth-tagged counts. In both cases, the background shapes predicted by the Monte Carlo bordered on requiring a non-linear function, like the four other crystal ball modes

Table 5.5: Tagging results from the 20× Monte Carlo sample, scaled to data size.

D_s mode	Fit counts	Truth tagged counts	$\frac{N_{fit}-N_{MC}}{\sigma}$
K_sK	5,764.0 ± 100.8	5,693.1	0.70
$KK\pi$	25,242.0 ± 233.9	25,731.6	−2.09
$K_sK\pi^0$	1,670.5 ± 157.7	1,871.2	−1.27
$K_sK_s\pi$	1,141.4 ± 69.3	1,081.5	0.86
$KK\pi\pi^0$	6,693.4 ± 323.6	6,844.5	−0.47
$K_sK^+\pi\pi$	1,744.1 ± 105.5	1,717.3	0.25
$K_sK^-\pi\pi$	3,246.3 ± 92.2	3,200.6	0.50
$\pi\pi\pi$	6,081.6 ± 326.3	6,197.6	−0.36
$\pi\eta$	2,882.3 ± 182.9	3,334.4	−2.47
$\pi\pi^0\eta$	6,825.9 ± 700.7	6,560.0	0.38
$\pi\eta', \eta' \to \pi\pi\eta$	2,132.4 ± 64.3	2,108.4	0.37
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	532.5 ± 84.5	749.3	−2.57
$\pi\eta', \eta' \to \rho\gamma$	3,904.4 ± 245.2	4,079.9	−0.72
Sum	67,860.7 ± 959.8	69,169.5	−1.36

($KK\pi\pi^0$; $\pi\pi\pi$; $\pi\pi^0\eta$; and $\pi\eta', \eta' \to \rho\gamma$). However, we try to avoid such background functions when we have a wide signal shape because the background function can dip inappropriately in the D_s mass region. This lower background leads to an overestimate on the tag counts. The four modes in which we do use a quadratic background have more events than either $\pi\pi^0\eta', \eta' \to \pi\pi\eta$ or $\pi\eta$, and they tend to have narrower signal shapes ($\pi\pi^0\eta$ has a wider shape but more events). These qualities make us less sensitive to the background function when we shift to the data, where we need to be more careful in case the D_s mass reconstruction has a poorer resolution than the Monte Carlo predicts.

In addition to procedure cross-checks, we have used the Monte Carlo to determine our tagging efficiency within semileptonic events. As expected, we see essentially the same tagging efficiency independent of the semileptonic mode (Section 10.1.3). However, our recoil mass cut does create a difference in tagging efficiency based on the D_s production method: D_sD_s, $D_s^*D_s$ with the tagged D_s from the D_s^* ("secondary"), or $D_s^*D_s$ where the tagged D_s does not come from the D_s^* ("prompt"). Table 5.6 gives the efficiencies for each case.

Table 5.6: Overall D_s tagging efficiency from signal Monte Carlo, including all branching ratios. Our D_s recoil mass cut creates the efficiency difference between tags from prompt D_s and tags from secondary D_s.

D_s production mode	$\varepsilon_{\rm tag}$
D_sD_s	0.42% ± 0.01%
$D_s^*D_s$ with prompt $D_s \to tag$	7.21% ± 0.03%
$D_s^*D_s$ with secondary $D_s \to tag$	5.69% ± 0.03%
Weighted MC	6.22% ± 0.02%

5.5.2 Data

Figures 5.2–5.4 show our fits to the combined data from datasets 39, 40, 41, 47, and 48. Table 5.7 summarizes each mode's tag counts resulting from these fits.

Table 5.7: Tagging results from the full data sample (sum of datasets 39, 40, 41, 47, 48).

D_s mode	Fit counts
K_sK	6,226.7 ± 101.2
$KK\pi$	27,373.5 ± 248.4
$K_sK\pi^0$	2,246.8 ± 209.9
$K_sK_s\pi$	1,125.5 ± 76.5
$KK\pi\pi^0$	7,355.5 ± 377.4
$K_sK^+\pi\pi$	1,859.4 ± 120.6
$K_sK^-\pi\pi$	3,377.3 ± 100.0
$\pi\pi\pi$	6,606.3 ± 337.7
$\pi\eta$	3,810.3 ± 190.8
$\pi\pi^0\eta$	9,476.9 ± 529.0
$\pi\eta', \eta' \to \pi\pi\eta$	2,386.6 ± 65.6
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	1,090.5 ± 118.7
$\pi\eta', \eta' \to \rho\gamma$	4,272.3 ± 193.3
Sum	77,207.5 ± 880.2

Although not directly relevant for this analysis, we find it interesting that we see about 14% more tags in the data than we expected from the 20× Monte Carlo sample. This difference persists across each mode and dataset to within errors, as shown in

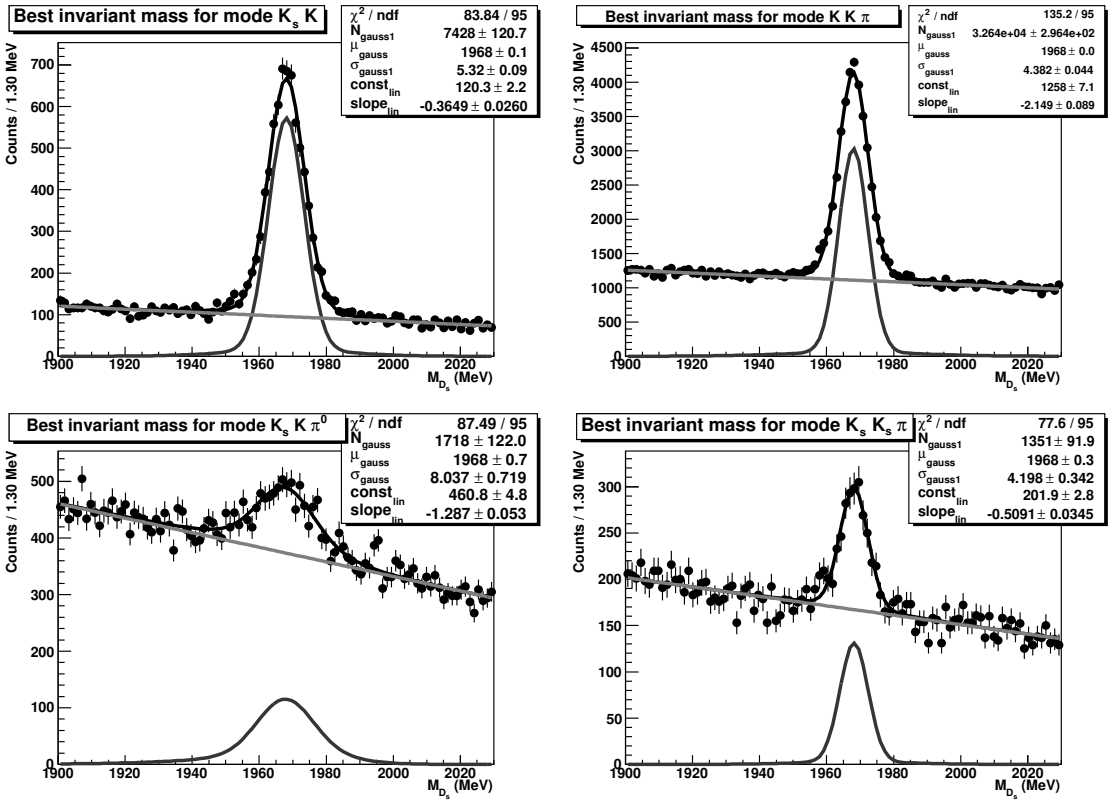

Figure 5.2: D_s invariant mass fits in the data, determining the total number of D_s tags for modes K_sK, $KK\pi$, $K_sK\pi^0$, and $K_sK_s\pi$. The peaking function reflects our fits' signal component, while the linear/quadratic function shows our background.

Tables F.1 and F.2.

5.5.3 Cross-Checks

In addition to our fitting systematics, described in Section 10.1.1, we have performed two cross-checks for our fitting procedure. In the first cross-check, we use the Monte Carlo truth-tagged histogram for our signal shape instead of the double gaussian or gaussian + crystal ball functions. In the second, we ensure that our chosen procedure consistently fits data-sized samples by breaking the 20× Monte Carlo into 20 equal subsets and fitting each individually.

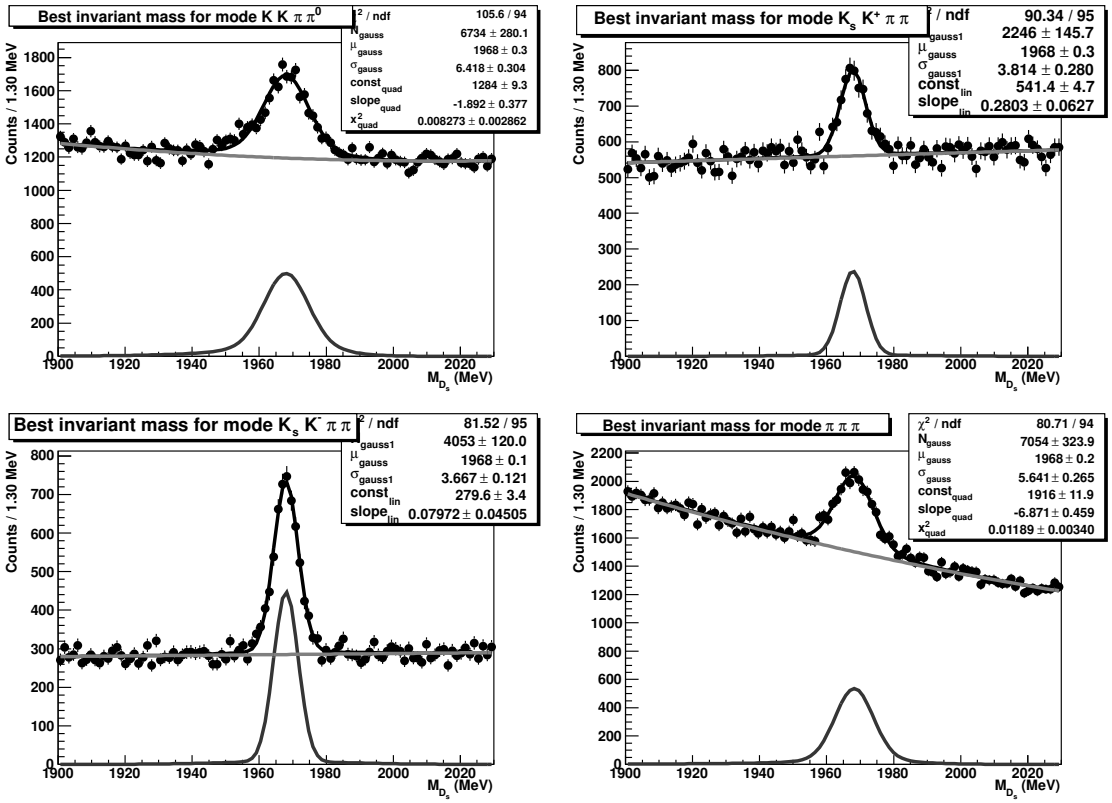

Figure 5.3: D_s invariant mass fits in the data, determining the total number of D_s tags for modes $KK\pi\pi^0$, $K_sK^+\pi\pi$, $K_sK^-\pi\pi$, and $\pi\pi\pi$.

In our signal histogram cross-check, we take the D_s invariant mass spectrum from the truth-tagged Monte Carlo as the signal shape instead of a double gaussian or gaussian + crystal ball function. The overall histogram normalization gives us our only free signal parameter. We then add the same background function as in our standard fit (linear or quadratic, by mode). This leaves either 3 or 4 total free parameters, depending on the D_s mode.

We first fit the full Monte Carlo to ensure that the signal histogram fit returns the truth-tagged counts. This fit does match both our standard fit results and the true number of tags to within error, as shown in Table 5.8.

We then fit the data with the signal histograms and corresponding background

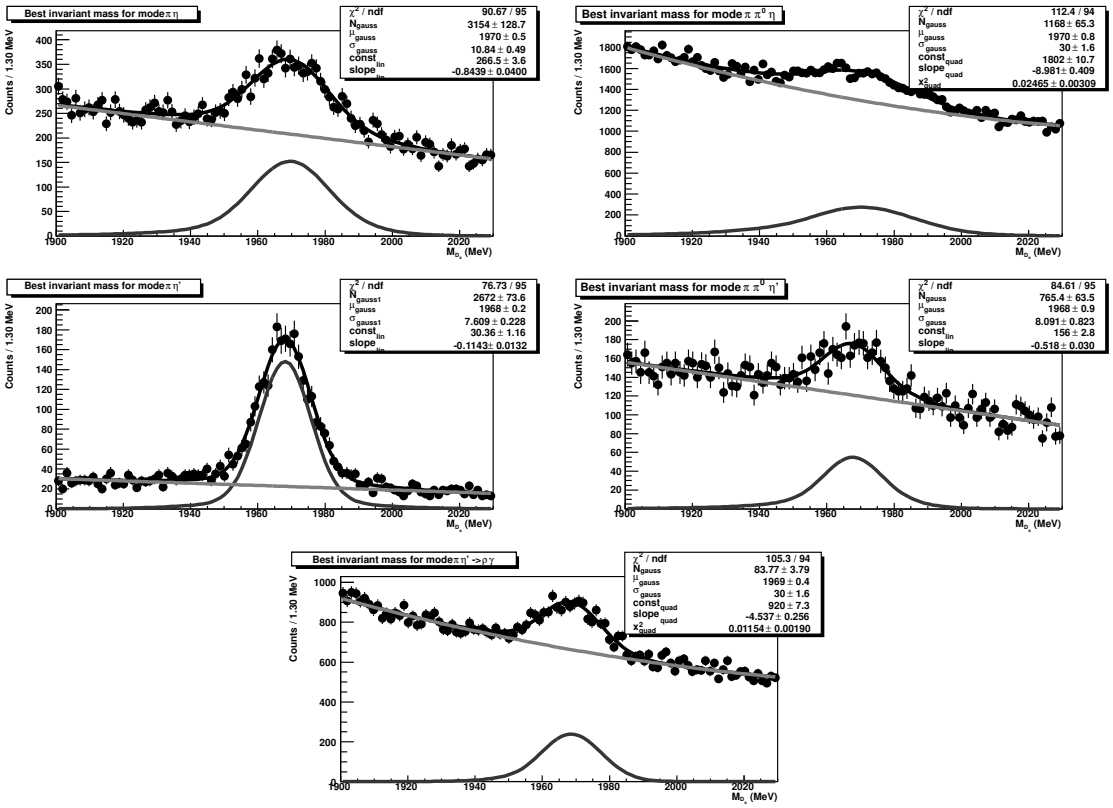

Figure 5.4: D_s invariant mass fits in the data, determining the total number of D_s tags for modes $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

functions. We expect the signal histogram fit to give more or less the same result as our fit function as long as the D_s mass resolution in Monte Carlo accurately represents the true resolution in the data.

The fit results, displayed in Figures G.16–G.21 and summarized in Table 5.9, show consistency between the signal histogram fit and our more flexible double gaussian/gaussian + crystal ball function for all modes, with the exception of $\pi\pi^0\eta$. This discrepancy does not particularly surprise us since $\pi\pi^0\eta$ has the worst signal to background ratio, has a wide signal shape, and has a background shape that requires a quadratic function. The combination of these issues allows the signal shape to trade off with the background's quadratic curvature to some extent. Since the $\pi\pi^0\eta$ signal histogram fit clearly doesn't

Table 5.8: Signal histogram fit results compared to our standard double gaussian/ gaussian+crystal ball fit results in the 20× Monte Carlo sample, scaled to data size.

D_s mode	Signal histogram	Standard fit	Truth-tagged counts
K_sK	5,813.5 ± 94.0	5,764.0 ± 100.8	5,693.1
$KK\pi$	25,834.7 ± 214.9	25,242.0 ± 233.8	25,731.6
$K_sK\pi^0$	1,783.8 ± 123.5	1,670.5 ± 159.2	1,871.2
$K_sK_s\pi$	1,161.4 ± 59.5	1,141.4 ± 69.2	1,081.5
$KK\pi\pi^0$	6,815.9 ± 219.4	6,693.4 ± 323.6	6,844.5
$K_sK^+\pi\pi$	1,766.5 ± 87.9	1,744.0 ± 105.3	1,717.3
$K_sK^-\pi\pi$	3,326.3 ± 83.1	3,246.3 ± 92.2	3,200.6
$\pi\pi\pi$	6,220.3 ± 220.9	6,082.3 ± 309.7	6,197.6
$\pi\eta$	3,043.0 ± 138.6	2,882.3 ± 182.4	3,334.4
$\pi\pi^0\eta$	6,615.8 ± 417.8	6,825.9 ± 700.7	6,560.0
$\pi\eta', \eta' \to \pi\pi\eta$	2,164.2 ± 58.6	2,132.4 ± 64.3	2,108.4
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	624.4 ± 77.1	532.5 ± 84.0	749.3
$\pi\eta', \eta' \to \rho\gamma$	4,074.1 ± 205.0	3,904.4 ± 245.2	4,079.9
Sum	69,244.0 ± 656.2	67,861.2 ± 954.3	69,169.5

fit well (both visually and in terms of χ^2), undershooting the D_s mass distribution's high side and overshooting the low side, we don't feel a need to take an additional systematic based on its results.

For our data-sized cross-check, we split the 20× Monte Carlo sample into 20 separate samples to ensure that our fit function will successfully and reliably converge. Table 5.10 gives our results, where the summation row states the total fit tag counts, the total truth-tagged counts, and the total χ^2 across the 20 samples. Our reduced χ^2 of 1.90 represents a systematic undercounting of the true number of tags.

We believe that the slight, systematic undercounting of about 1% originates from an imperfection in the signal fit shape. The signal function peak falls slightly below the true peak whereas the tails have a small surplus, as seen from our signal Monte Carlo plots in Figures G.10–G.12. When we use this signal shape in a sample with background, the tails of the signal distribution get partially subsumed into a higher background while the slightly low signal shape peak still matches the data peak.

The statistical errors for each tag mode exceed the slight D_s tag undercounting effect. However, since the peak region essentially acts like a double gaussian for all

Table 5.9: Signal histogram fit results compared to our standard double gaussian/ gaussian+crystal ball fit results in the full data sample.

D_s mode	Signal histogram	Standard fit
K_sK	6,236.9 ± 94.9	6,226.7 ± 101.2
$KK\pi$	27,576.2 ± 225.9	27,373.5 ± 248.4
$K_sK\pi^0$	2,039.9 ± 127.7	2,246.8 ± 209.9
$K_sK_s\pi$	1,155.0 ± 63.1	1,125.5 ± 76.5
$KK\pi\pi^0$	7,590.2 ± 256.7	7,355.5 ± 377.4
$K_sK^+\pi\pi$	1,854.6 ± 98.9	1,859.4 ± 120.6
$K_sK^-\pi\pi$	3,434.5 ± 89.4	3,377.3 ± 100.0
$\pi\pi\pi$	6,437.5 ± 233.2	6,606.3 ± 337.7
$\pi\eta$	3,527.5 ± 128.7	3,810.3 ± 190.8
$\pi\pi^0\eta$	5,848.1 ± 485.2	9,476.9 ± 529.0
$\pi\eta', \eta' \to \pi\pi\eta$	2,300.0 ± 58.6	2,386.6 ± 65.6
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	1,069.8 ± 81.6	1,090.5 ± 118.7
$\pi\eta', \eta' \to \rho\gamma$	4,265.9 ± 209.8	4,272.3 ± 193.3
Sum	73,336.1 ± 724.2	77,207.5 ± 880.2

modes (as the crystal ball function acts like a gaussian near the peak), we consider the undercounting rate correlated between modes. We would treat the undercounting as a systematic when determining the branching ratio, but this systematic gets complicated by two factors.

The first issue with a proposed undercounting systematic comes from the fact that we ultimately fit the D_s mass spectrum for the number of semileptonic events as well as the number of tags. The semileptonic event counts also have a linear background function, so the slight undercounting can appear in both our branching ratio's numerator and denominator, significantly shrinking the ultimate systematic.

Even if we dismiss the numerator effect, though, we have a second ameliorating factor. Our conventional systematic (Section 10.1.1) involves changing the D_s mass resolution based on results from the $D^{\pm}$ system. This systematic indicates a wider underlying D_s mass resolution than predicted by the Monte Carlo, which actually makes our signal shape more accurate. Since the conventional systematic involves a wider distribution and gives larger final errors (just due to the systematic's precision), we consider that study to measure essentially the same signal shape concern, and we don't

Table 5.10: Fit results from the Monte Carlo's 20 data-sized samples. The final column in the "Sum" row gives the total χ^2.

Datasize sample	Fit counts	Truth tagged counts	$\frac{N_{fit}-N_{MC}}{\sigma}$
1	67,932.6 ± 1,022.9	70,585	−2.59
2	71,422.1 ± 1,214.3	71,265	0.13
3	71,579.7 ± 1,386.5	71,112	0.34
4	71,409.4 ± 919.6	71,119	0.32
5	69,774.6 ± 1,046.3	71,552	−1.70
6	69,795.4 ± 983.4	71,326	−1.56
7	71,017.8 ± 1,221.3	71,007	0.01
8	72,092.7 ± 1,275.1	71,154	0.74
9	67,852.4 ± 1,133.5	71,084	−2.85
10	70,128.0 ± 809.5	71,044	−1.13
11	69,108.9 ± 926.7	71,265	−2.33
12	70,547.9 ± 1,182.5	71,208	−0.56
13	71,312.2 ± 1,125.7	71,260	0.05
14	71,769.1 ± 1,390.2	71,292	0.34
15	69,466.9 ± 840.6	70,906	−1.71
16	69,036.8 ± 992.4	71,235	−2.21
17	71,385.3 ± 1,250.5	71,528	−0.11
18	71,574.6 ± 1,352.9	70,952	0.46
19	70,312.3 ± 1,266.5	71,439	−0.89
20	69,738.3 ± 1,027.0	70,842	−1.07
Sum	1,407,257.0 ± 5,060.6	1,423,175	38.04

add an additional undercounting systematic.

Chapter 6

Semileptonic Selection Criteria

While we use the D_s mass spectrum in our fits for both the tag counts and the semileptonic counts, we do make cuts on the other particles to ensure that we have a semileptonic event. In particular, we make an electron cut that gives us the best background rejection in our analysis. We have also studied extra track and extra shower event cuts. While splitoff showers make any extra shower cuts problematic, we do find an extra track cut useful in our final selection.

6.1 Electron Selection

We select electrons using the same general track quality requirements applied to D_s tracks in Section 5.1.1. Further, we use the $\mathcal{F}_{w/RICH}$ variable [67], a likelihood ratio that tests the electron hypothesis under a weighted combination of E/p, dE/dx, and RICH information. $\mathcal{F}_{w/RICH}$ runs from 0 to 1, with 1 being electron-like, as shown in Figure 6.1. We require $\mathcal{F}_{w/RICH} \geq 0.8$.

We also add the requirement that the hit fraction falls below 1.2, although this cut has no real impact beyond consistency with previous systematic work. We do not attempt to reconstruct electrons with a momentum below 200 MeV. These soft electrons cost some efficiency, but as Table 6.1 shows, we get substantial background reduction from our combined electron cuts.

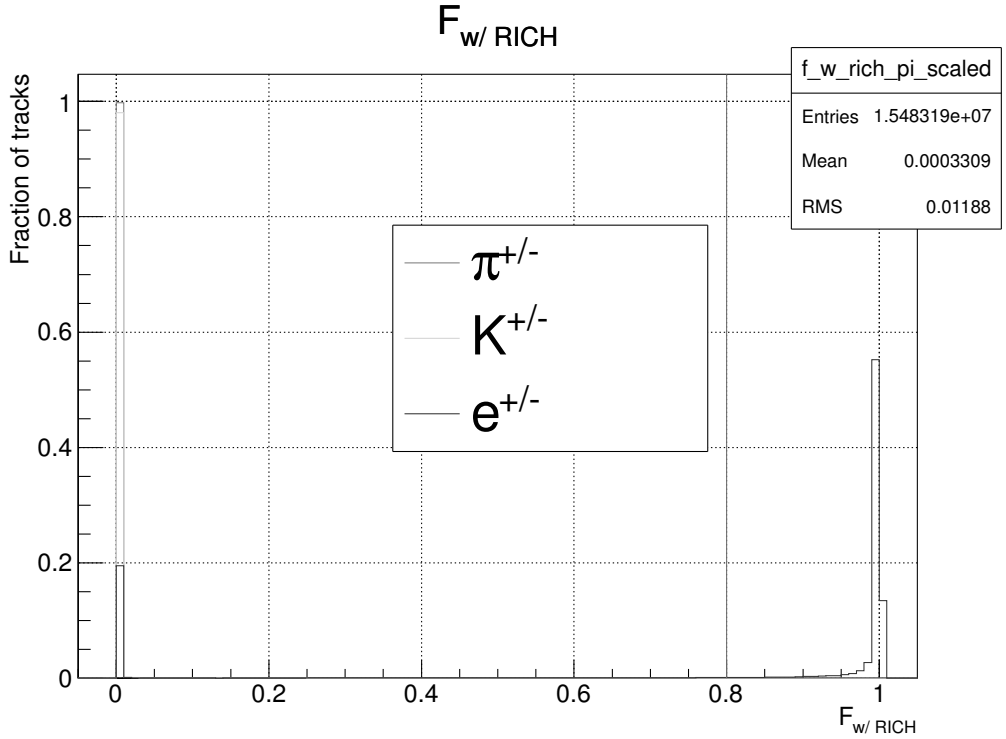

Figure 6.1: $\mathcal{F}_{w/RICH}$ in the 20× MC sample for good tracks that are not used in the tagged D_s. We only plot electrons involved in one of our six semileptonic modes ($\phi e\nu$, $\eta e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$). The electron peak at zero comes primarily from tracks with a momentum below 200 MeV.

6.2 Event Selection

We look for D_s semileptonic decays in $D_s^* D_s$ and $D_s D_s$ events, where the D_s^* decays to a D_s with some photons (either directly to a γ or via a π^0). Consequently, we should not have any tracks other than those from the tagged D_s, the electron, or the semileptonic hadron. We reject any event with an extra track, which cuts out events with an e^+e^- pair that would otherwise pass our electron selection, semileptonic events with the wrong hadron mode (e.g. $D_s \to \eta' e\nu$ faking $D_s \to \eta e\nu$), and some semileptonic events with misreconstructed tags. Not surprisingly, we rarely throw out signal events with our extra track cut, as shown in Table 6.2.

We have considered a rejection on extra showers above various energy thresholds

Table 6.1: Effect of the electron cuts (track and $\mathcal{F}_{w/RICH}$) in the 20× Monte Carlo sample for truth-tagged semileptonic and generic decay modes. These precede any semileptonic hadron cut, but passed D_s tags must fall within the tagging fit window (1900 MeV $< M_{D_s} <$ 2030 MeV).

D_s semileptonic mode	# Passed D_s tags	# Passed electron cuts	Cut efficiency
$\phi e\nu$	55,399	31,864	58%
$\eta e\nu$	50,775	35,772	70%
$\eta' e\nu$	19,282	12,038	62%
$K_s e\nu$	1,022	641	63%
$K^* e\nu$	4,148	2,582	62%
$f_0 e\nu$	7,132	4,411	62%
All other modes	6,741,304	204,944	3%

(25 MeV, 100 MeV, 300 MeV, and 500 MeV), but we did not find them useful (Figures 6.2 and G.22–G.26). Our D_s and electron selections keep the background levels low, so even the handful of signal events improperly rejected by an extra shower cut will cause a reduction in our statistical significance. Several different effects can lead to improper signal rejection under an extra energy cut: modes containing kaons can have decays in flight, leaving a shower without a properly matched track; the D_s^* can decay to a π^0, yielding one extra shower; legitimate single showers sometimes get misreconstructed as two or more showers; and splitoff showers may not get identified as such. Although we can reduce the impact of each effect through various selections, we have not found that such improvements save the extra shower cut.

Table 6.2: Effect of an extra track cut on the signal and the background in the peaking D_s mass region after all other selections (e.g. semileptonic hadron cuts) have been made. We define f as $\frac{S^2}{S+B}$, such that $\frac{f_{\text{post-cut}}}{f_{\text{pre-cut}}}$ measures our statistical improvement.

D_s semileptonic mode	$\varepsilon_{\text{signal}}$	$\varepsilon_{\text{background}}$	$\frac{f_{\text{post-cut}}}{f_{\text{pre-cut}}}$
$\phi e\nu$	99%	61%	1.01
$\eta e\nu$	97%	42%	1.31
$\eta' e\nu$	97%	52%	1.06
$K_s e\nu$	98%	69%	1.38
$K^* e\nu$	98%	64%	1.40
$f_0 e\nu$	99%	59%	1.23

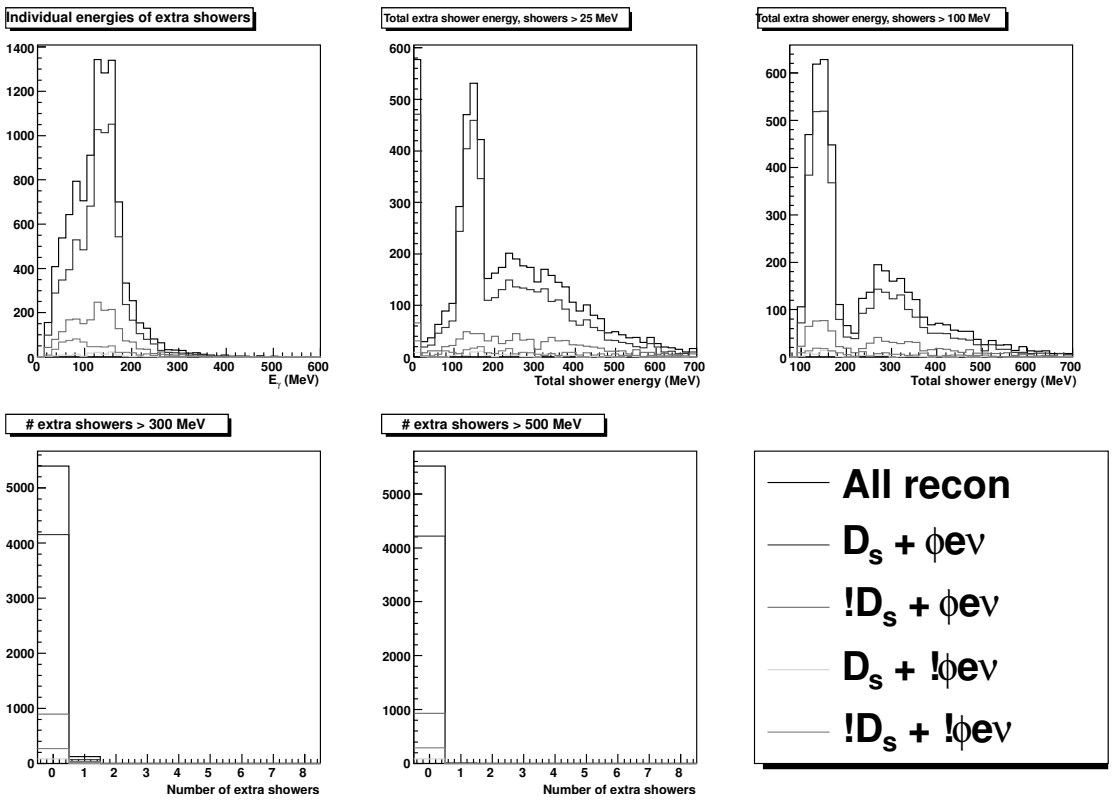

Figure 6.2: Extra showers after finding the tagged D_s, the ϕ, and the electron in $\phi e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays. We don't see an improvement in our results from making a cut on any extra shower variable.

Chapter 7

Measurement of $D_s \to \phi e\nu$

While we ultimately reconstruct six different D_s semileptonic modes, we use a similar reconstruction and fitting procedure for five of them ($\phi e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$). To illustrate the procedure, we first describe how the process works for $D_s \to \phi e\nu$. Then, in section 8, we show how the procedure differs for each of the other four semileptonic modes that have a similar reconstruction but fewer events.

We determine the number of $D_s \to \phi e\nu$ events by fitting the tagged D_s mass spectrum after making the electron selections in Section 6.1, the ϕ cuts in Section 7.1, and the event cuts in Section 6.2. We do not directly use the electron or ϕ kinematic information beyond requiring that they pass our particle cuts, although our ϕ reconstruction gets used indirectly in our $f_0 \to KK$ background subtraction.

7.1 ϕ Selection

We reconstruct the ϕ meson in $D_s \to \phi e\nu$ via the $\phi \to KK$ decay mode, which makes up roughly half of all ϕ decays. $D_s \to \phi e\nu$ presents a challenge in its ϕ reconstruction as the semileptonic ϕ tends to be fairly soft (Figure 7.1a). The soft ϕ and low Q value in the $\phi \to KK$ decay (about 32 MeV) leads to soft daughter kaons, which decreases our detection efficiency significantly relative to higher energy ϕ decays (e.g. $D_s \to \phi\pi$). As an additional challenge, the long tail of the ϕ Breit-Wigner forces us to use a wide ϕ mass window relative to its decay width of $\Gamma_\phi \approx 4.26$ MeV.

The combination of low $D_s \to \phi e\nu$ background and soft kaons encourage loose kaon

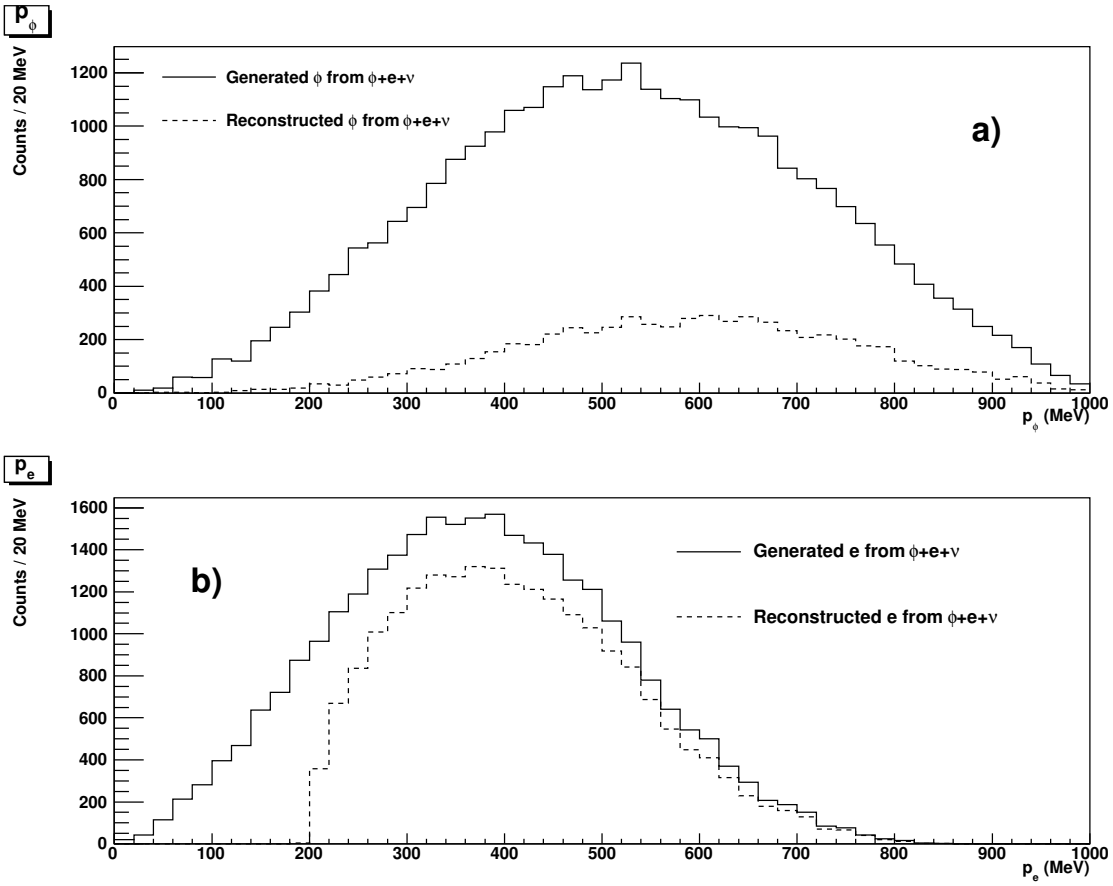

Figure 7.1: Generated and reconstructed momentum spectra from the 20× Monte Carlo. a) ϕ from $D_s \to \phi e\nu$, using our cuts. b) Electron from $D_s \to \phi e\nu$.

selection criteria. In particular, we have essentially dropped any hit fraction cut to accommodate the higher likelihood that a soft kaon will decay in flight (Figure 7.2). We considered loosening other typical track cuts, but we did not find the slight efficiency improvement to be worth deviating from established systematic studies [68]. We ultimately require the following cuts on the ϕ meson's daughter kaons:

- $|d_b| < 5$ mm
- $|z_0| < 5$ cm

- $\chi^2 < 100,000$
- $\left|\sigma_K^{dE/dx}\right| < 3.0$
- Hit Fraction > 0.1

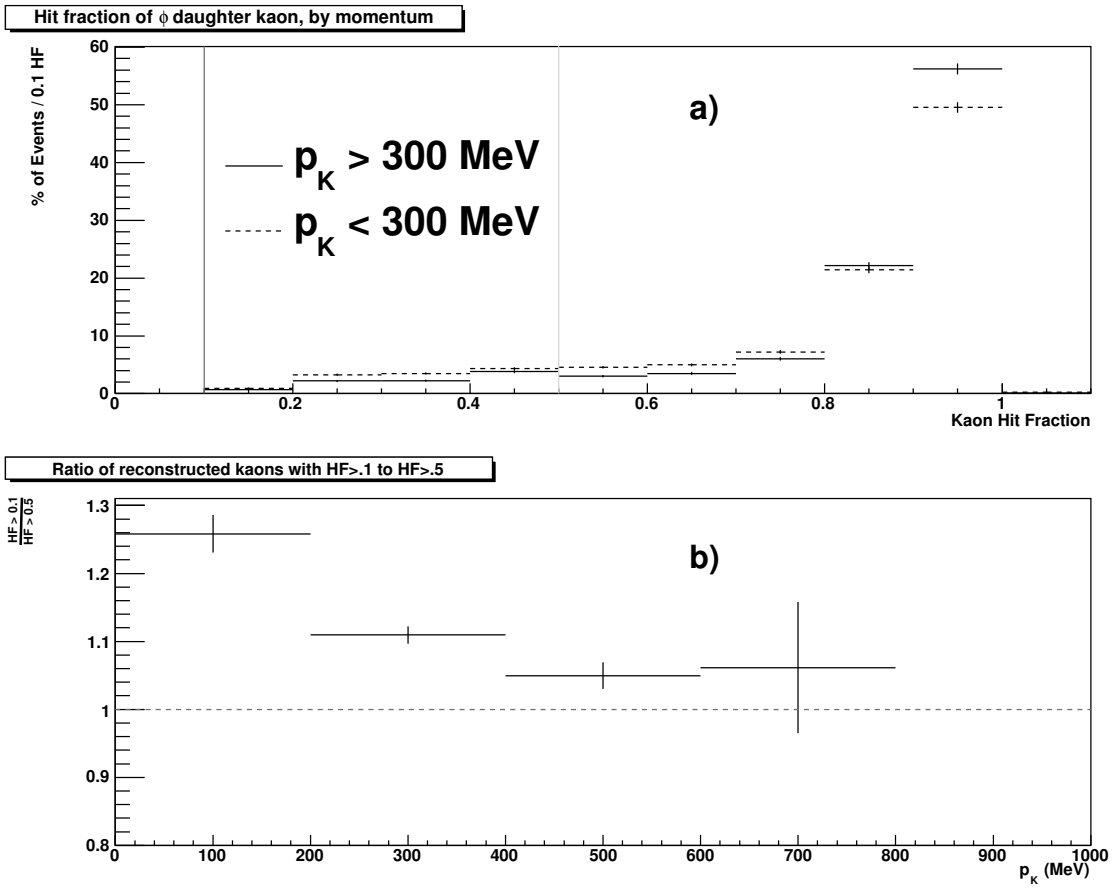

Figure 7.2: Hit fraction distributions for kaons from $\phi \to KK$ in $D_s \to \phi e\nu$. a) Percentage of such kaons with a given hit fraction, above and below 300 MeV (roughly the median momentum). b) Relative improvement in our kaon reconstruction efficiency when loosening the hit fraction cut from .5 to .1, by momentum.

With the low combinatoric background in $D_s \to \phi e\nu$ and the long ϕ Breit-Wigner tail, we extend the high side of our ϕ mass window as far as we reasonably can. On the low side of the ϕ mass window, we run into contamination from $D_s \to f_0 e\nu$, where $f_0 \to KK$. While we do measure $D_s \to f_0 e\nu$ via $f_0 \to \pi\pi$ and use it to guide our background subtraction, the uncertainty in f_0 parameters leads us to restrict the low ϕ mass range somewhat to reduce the impact of these systematic uncertainties on our $\phi e\nu$ result. For our final ϕ mass cut, shown in Figure 7.3, we use:

- -15 MeV $< M_\phi^{\text{reconstructed}} - M_\phi^{\text{PDG}} < 30$ MeV

We do not attempt to reconstruct the ϕ through modes other than KK (e.g. $\pi\pi\pi^0$).

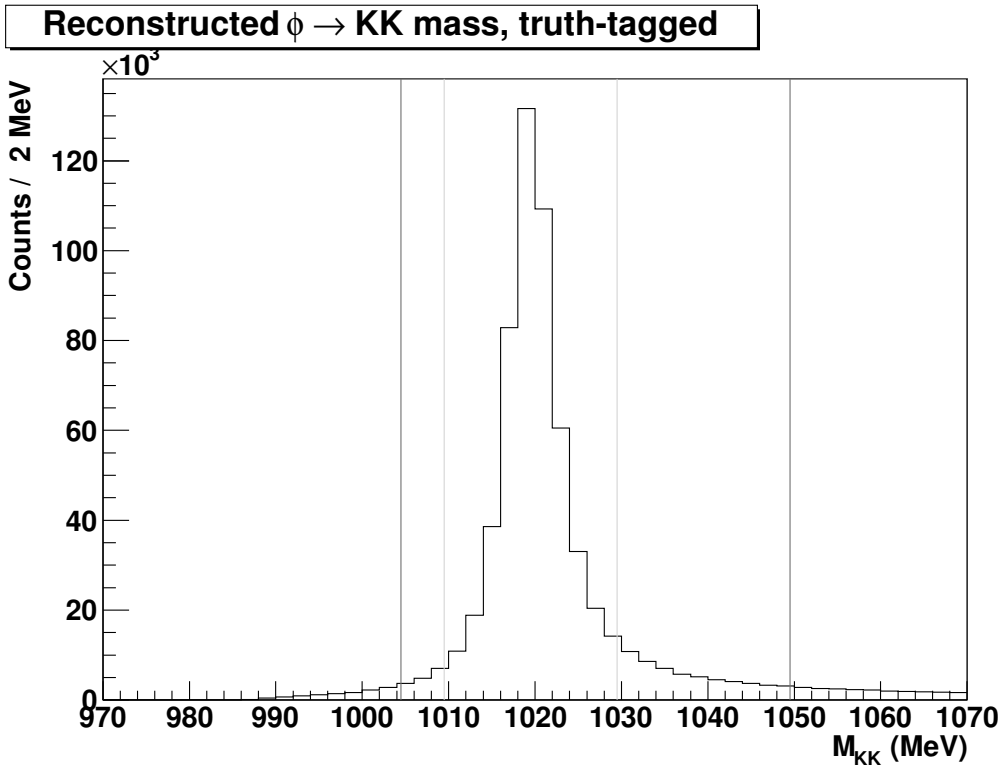

Figure 7.3: ϕ mass distribution when reconstructed from $\phi \to KK$ in the 20× Monte Carlo. The inner vertical lines represent a 10 MeV cut (roughly $2\Gamma_\phi$), which does not capture the high mass tail. We accept ϕ masses within the outer vertical lines.

7.2 $D_s \to \phi e\nu$ Reconstruction

7.2.1 Efficiency

To determine our efficiencies, we generated signal Monte Carlo samples with one D_s decaying to $\phi e\nu$ and the other decaying generically. We made four samples with 250,000 events, one for each different type of semileptonic D_s production at 4170 MeV: one sample for $D_s^+ D_s^-$ and three for $D_s^* D_s$, where the semileptonic D_s can either be "prompt" (not from the D_s^* side) or "secondary" in one of two ways ($D_s^* \to D_s\gamma$ or $D_s^* \to D_s\pi^0$).

We present the efficiencies for $D_s \to \phi e\nu$ in Table 7.1 for typical ϕ cuts and Table 7.2 for this analysis's looser ϕ cuts. We determine all our semileptonic efficiencies after successfully reconstructing a D_s tag within the fit window (1900 MeV $\leq M_{D_s} \leq$ 2030 MeV). The efficiency for both the ϕ and the electron increases with higher momenta, as shown in Figure 7.4. This causes the overall semileptonic efficiency to be slightly lower than the simple product of hadron and electron efficiencies, since high momentum electrons are correlated with low momentum ϕ and vice versa.

Table 7.1: Efficiencies for semileptonic particles in $D_s \to \phi e\nu$, with typical ϕ cuts (HF $>$ 0.5, ϕ mass within 10 MeV). The efficiencies include the $\phi \to KK$ branching ratio.

D_s production mode	ε_e	ε_ϕ	ε_{SL}
$D_s D_s$	70.4% ± 2.6%	12.9% ± 1.1%	9.4% ± 1.0%
$D_s^* D_s$ with $D_s^* \to (D_s \to \phi e\nu)\ \gamma$	70.2% ± 0.6%	15.2% ± 0.3%	10.2% ± 0.2%
$D_s^* D_s$ with $D_s^* \to (D_s \to \phi e\nu)\ \pi^0$	70.3% ± 0.6%	15.1% ± 0.3%	10.1% ± 0.2%
$D_s^* D_s$ with prompt $D_s \to \phi e\nu$	70.7% ± 0.7%	14.1% ± 0.3%	9.5% ± 0.3%
Weighted MC	70.5% ± 0.5%	14.9% ± 0.2%	10.1% ± 0.2%

Table 7.2: Efficiencies for semileptonic particles in $D_s \to \phi e\nu$, with the ϕ cuts used in this analysis. The efficiencies include the $\phi \to KK$ branching ratio.

D_s production mode	ε_e	ε_ϕ	ε_{SL}
D_sD_s	70.4% ± 2.6%	17.3% ± 1.3%	12.1% ± 1.1%
$D_s^*D_s$ with $D_s^* \to (D_s \to \phi e\nu)\ \gamma$	70.2% ± 0.6%	20.9% ± 0.4%	14.2% ± 0.3%
$D_s^*D_s$ with $D_s^* \to (D_s \to \phi e\nu)\ \pi^0$	70.3% ± 0.6%	20.9% ± 0.3%	14.1% ± 0.3%
$D_s^*D_s$ with prompt $D_s \to \phi e\nu$	70.7% ± 0.7%	19.8% ± 0.4%	13.5% ± 0.3%
Weighted MC	70.5% ± 0.5%	20.7% ± 0.3%	14.0% ± 0.2%

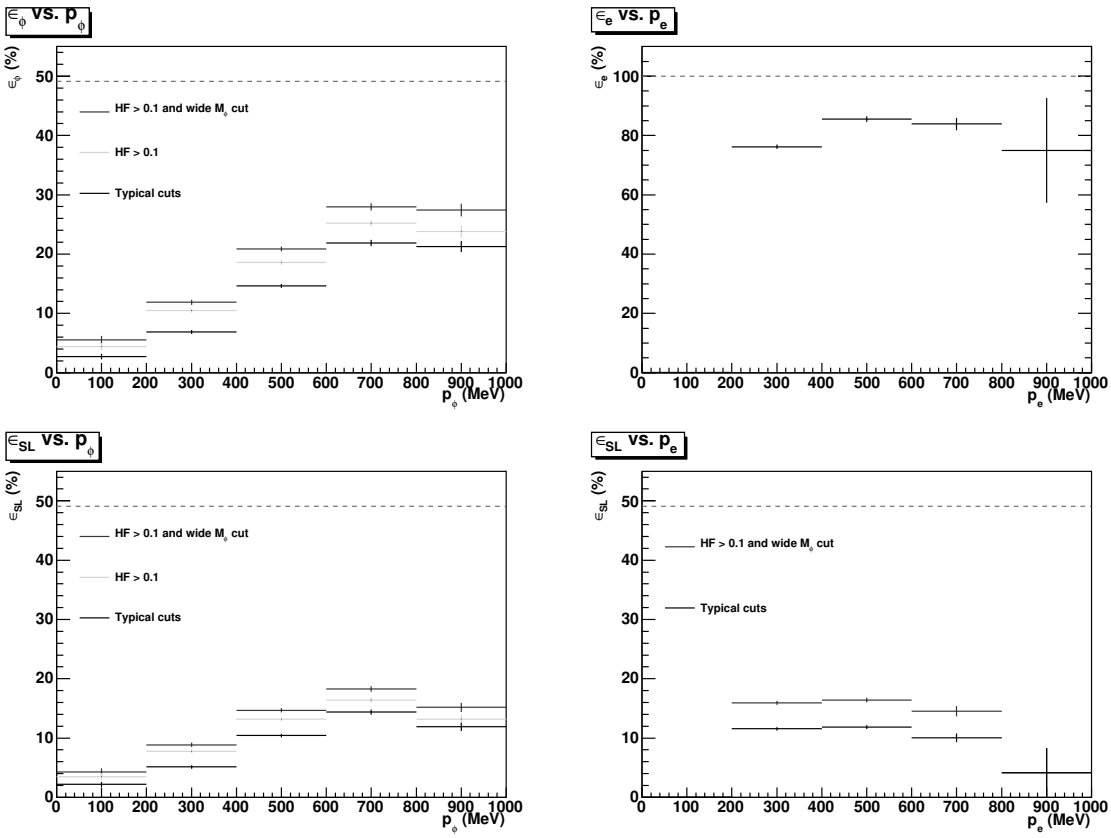

Figure 7.4: Efficiencies for individual semileptonic particles and the overall semileptonic side (ϕ + electron), by momentum. We include the $\phi \to KK$ branching ratio in our efficiencies, so ε_ϕ and ε_{SL} must be less than 49%. The ϕ "typical cuts" have HF $>$ 0.5 and a 10 MeV mass cut. Section 7.1 gives our looser ϕ selection.

7.2.2 Backgrounds

Our predicted $D_s \to \phi e\nu$ background from the 20× Monte Carlo primarily consists of events that don't peak within our D_s mass fit region (Table 7.3). We use a linear background function in our data fit for such events. However, we do have to take special account of events that fake the semileptonic side (electron or $\phi \to KK$) while having a true D_s tag, as these events will look like signal events in our fit to the D_s mass spectrum.

Table 7.3: Truth-tagged breakdown for $D_s \to \phi e\nu$ candidates passing all cuts in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Passing candidates
True D_s tag, true $D_s \to \phi e\nu$	210.90
True D_s tag, not $D_s \to \phi e\nu$	4.10
False D_s tag, true $D_s \to \phi e\nu$	46.40
False D_s tag, not $D_s \to \phi e\nu$	16.55

Most events with a true D_s tag that are incorrectly reconstructed as $D_s \to \phi e\nu$ come from $D_s \to f_0 e\nu$, where $f_0 \to KK$ (Figure 7.5). The large decay width of the f_0 means that a nontrivial fraction of $f_0 \to KK$ decays have an invariant mass within our $\phi \to KK$ mass window. At the same time, the low overall number of $D_s \to f_0 e\nu$, $f_0 \to KK$ decays relative to other backgrounds (particularly combinatoric background near KK threshold with false D_s tags) prevents us from simply fitting the M_{KK} spectrum to determine the number of $f_0 e\nu$ background events. We instead use our $D_s \to f_0 e\nu$ ($f_0 \to \pi\pi$) measurement from Section 8.5, a range of $\frac{\Gamma_{f_0 \to KK}}{\Gamma_{f_0 \to \pi\pi}}$ values from the Particle Data Book [21], and a model for the $f_0 \to KK$ lineshape to estimate the amount of $D_s \to f_0 e\nu$, $f_0 \to KK$ background that appears within our $D_s \to \phi e\nu$ sample, as described in Appendix A.

We give our correction for the $D_s \to \phi e\nu$ branching ratio from the $D_s \to f_0 e\nu$, $f_0 \to KK$ peaking background and its associated systematic in Table 7.4. We also include a correction for events with a correct D_s tag that fake $D_s \to \phi e\nu$ from sources other than $D_s \to f_0 e\nu$, primarily $\phi\rho^{\pm}$. We use Monte Carlo estimates for the latter correction since the decay kinematics are well understood (relative to their statistical significance).

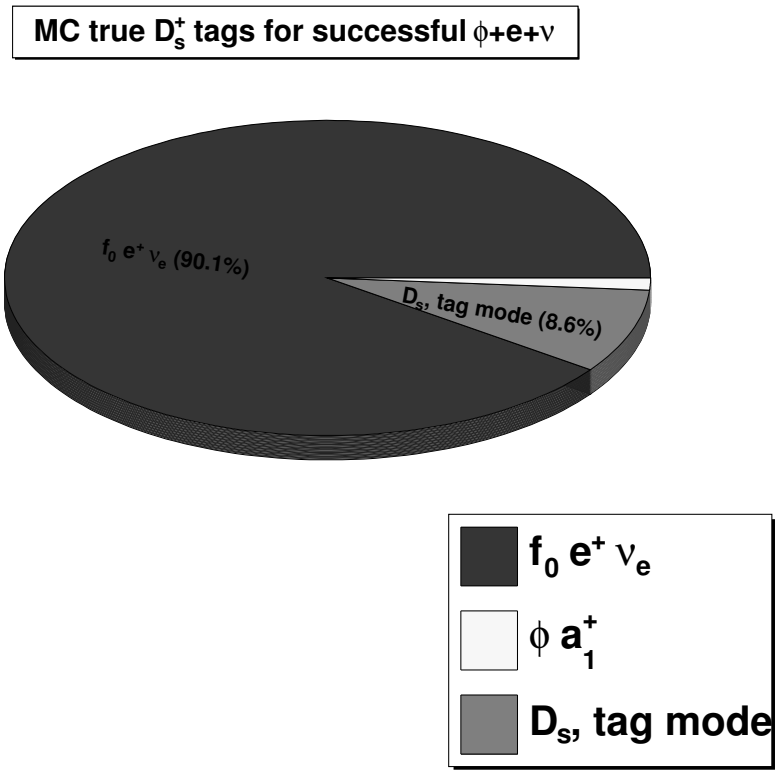

Figure 7.5: $D_s \to \phi e\nu$ backgrounds with a true D_s tag (peaking background) in the 20× Monte Carlo. Only $D_s \to f_0 e\nu$, $f_0 \to KK$ presents a sizable peaking background for $D_s \to \phi e\nu$.

7.2.3 Fit Procedure

In most tagged neutrino analyses (e.g. [69, 70]), the event's missing energy and momentum will be reconstructed and used in some form (e.g. missing mass, ΔE and M_{bc}) to determine the number of signal events. Since we've dropped the D_s^* meson's daughter γ in this analysis, our missing four momentum contains both the neutrino and the photon. The missing variables (mass, energy, momentum) are then distorted and do not give clear separation between signal and all backgrounds (although some background may still be rejected, as described for $K_s e\nu$ in Section 8.2).

When searching for an alternate fit parameter, we have noticed that a major systematic error in our $\phi e\nu$ and $\eta e\nu$ modes comes from the lineshape uncertainty associated

Table 7.4: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to \phi e\nu)$ from peaking background.

Background mode	BR correction
$D_s \to f_0 e\nu$	$(0.0080 \pm 0.0115)\%$
Non-semileptonic	$(0.0041 \pm 0.0014)\%$
Total $D_s \to \phi e\nu$ correction	$(0.0120 \pm 0.0116)\%$

with the D_s tag. To minimize this systematic, we have chosen to fit the D_s mass spectrum for our branching ratio's numerator as well as its denominator, which eliminates most of the impact from our reconstructed D_s lineshape uncertainty. As described previously, we do not directly involve the ϕ or electron kinematic information in our fit, using that particle information only to establish that we have a legitimate $D_s \to \phi e\nu$ event.

Even in our relatively high statistics $D_s \to \phi e\nu$ semileptonic mode, we only expect about 200 events spread across our 13 tag modes, with some tag modes likely to contain only one or two events. Our low statistics semileptonic modes ($\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$) have even fewer events spread across the same number of tag modes. To deal with the errors associated with our low statistics, we use an unbinned, extended maximum likelihood fit to our D_s mass spectrum. Further, since the small number of events per tag mode will distort a weighted sum combining individual results,[1] we instead use a common branching ratio parameter across all 13 tag modes.

We fix our signal D_s mass lineshape from each mode's tag fit results. This leaves only the normalization floating (via the common branching ratio parameter), giving us just one signal parameter in our fit. We determine a linear background for each tag mode from the 20× Monte Carlo, then we allow the overall normalization of this shape to float independently for each tag mode. If we have fewer than 20 background events for a tag mode in the 20× Monte Carlo sample (less than one expected background event in the data), we choose a constant background function instead of a linear function as our background shape for that tag mode.

Our fit function then has 14 floating parameters in our D_s mass fit – one signal

[1] Bienaymé formula for adding errors in quadrature need not apply.

parameter (the branching ratio) and one background normalization for each of our 13 D_s tag modes. We have a log likelihood function to minimize given by:

$$\begin{aligned}\mathcal{F}(\mathcal{B}, \vec{N}_{\rm BG}) &\equiv -\ln \mathcal{L}(\mathcal{B}, \vec{N}_{\rm BG}) \\ &= \mathcal{B} \cdot \varepsilon_{\rm SL} \cdot \left[\sum_{i=1}^{13} N^i_{\rm tag}\right] + \left[\sum_{i=1}^{13} N^i_{\rm BG}\right] \\ &\quad - \sum_{m_j} \ln\left[\mathcal{B} \cdot \varepsilon_{\rm SL} \cdot N^{[j]}_{\rm tag} \cdot f^{[j]}_{\rm sig}(m_j) + N^{[j]}_{\rm BG} \cdot f^{[j]}_{\rm BG}(m_j)\right],\end{aligned}$$

where $\mathcal{B}$ is the branching ratio, $\vec{N}_{\rm BG}$ is the background normalization (one per tag mode), $[j]$ refers to the tag mode associated with reconstructed D_s mass m_j, $f^i_{\rm sig}(m)$ and $f^i_{\rm BG}(m)$ are the normalized mass distributions of the signal and background for the given tag mode, respectively, and $N^i_{\rm tag}$ is the number of D_s tags for mode i. The first two terms in our function just represent the overall signal and background normalizations, while the third term corresponds to the sum of each event's log likelihood given our signal and background shapes.

7.3 Results

The fit results presented in the following subsections only involve the statistical error. We determine our systematic errors in Section 10 and give our full errors with the final, efficiency-corrected result in Section 11.2.

7.3.1 Monte Carlo

To ensure that our procedure properly measures the branching ratio on the couple hundred events expected in data, we have split the 20× Monte Carlo into twenty data-sized subsamples. These subsamples allow us an in-vs.-out test in which we accurately measure the number of truth-tagged semileptonic events, as seen from the lack of fit bias in Table 7.5. We also obtain the proper branching ratio in our Monte Carlo test after correcting for peaking background, as demonstrated in Table 7.6 with our branching ratios' χ^2 of 21.6 over the twenty subsamples.

Table 7.5: Test of potential bias in our fitting procedure for $D_s \to \phi e\nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	234	231.74 ± 16.58	−0.137
2	224	227.52 ± 14.79	0.238
3	199	191.93 ± 15.43	−0.458
4	226	227.01 ± 16.54	0.061
5	208	202.15 ± 15.74	−0.371
6	191	191.44 ± 13.94	0.032
7	240	248.97 ± 17.05	0.526
8	222	229.31 ± 16.71	0.438
9	191	194.93 ± 15.13	0.260
10	208	208.35 ± 14.50	0.024
11	194	201.72 ± 15.31	0.504
12	201	197.51 ± 14.91	−0.234
13	218	217.17 ± 16.13	−0.051
14	219	226.68 ± 16.23	0.473
15	235	242.62 ± 16.59	0.460
16	216	212.87 ± 15.74	−0.199
17	212	216.87 ± 15.85	0.307
18	220	212.31 ± 15.52	−0.496
19	210	212.19 ± 15.63	0.140
20	213	201.87 ± 15.55	−0.716
Sum	4281	4295.17 ± 70.28	0.202

For reference, we present our fits to these twenty data-sized subsamples as Figures G.27–G.31 in Appendix G.

Table 7.6: Monte Carlo comparison of the measured $D_s \to \phi e\nu$ branching ratio to its generating branching ratio (2.170%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios.

Dataset	Raw fit BR	Corrected BR	# σ
1	$(2.439 \pm 0.175)\%$	$(2.405 \pm 0.175)\%$	1.34
2	$(2.306 \pm 0.150)\%$	$(2.271 \pm 0.150)\%$	0.67
3	$(1.917 \pm 0.154)\%$	$(1.883 \pm 0.154)\%$	-1.86
4	$(2.289 \pm 0.167)\%$	$(2.255 \pm 0.167)\%$	0.51
5	$(2.076 \pm 0.162)\%$	$(2.041 \pm 0.162)\%$	-0.80
6	$(1.972 \pm 0.144)\%$	$(1.937 \pm 0.144)\%$	-1.62
7	$(2.520 \pm 0.173)\%$	$(2.485 \pm 0.173)\%$	1.83
8	$(2.300 \pm 0.168)\%$	$(2.266 \pm 0.168)\%$	0.57
9	$(2.014 \pm 0.156)\%$	$(1.979 \pm 0.156)\%$	-1.22
10	$(2.101 \pm 0.146)\%$	$(2.066 \pm 0.146)\%$	-0.71
11	$(2.094 \pm 0.159)\%$	$(2.059 \pm 0.159)\%$	-0.70
12	$(2.013 \pm 0.152)\%$	$(1.978 \pm 0.152)\%$	-1.26
13	$(2.178 \pm 0.162)\%$	$(2.143 \pm 0.162)\%$	-0.17
14	$(2.236 \pm 0.160)\%$	$(2.201 \pm 0.160)\%$	0.19
15	$(2.504 \pm 0.171)\%$	$(2.470 \pm 0.171)\%$	1.75
16	$(2.217 \pm 0.164)\%$	$(2.182 \pm 0.164)\%$	0.07
17	$(2.196 \pm 0.160)\%$	$(2.161 \pm 0.161)\%$	-0.06
18	$(2.138 \pm 0.156)\%$	$(2.104 \pm 0.156)\%$	-0.42
19	$(2.111 \pm 0.156)\%$	$(2.076 \pm 0.156)\%$	-0.60
20	$(2.058 \pm 0.158)\%$	$(2.023 \pm 0.158)\%$	-0.93
Weighted averages/χ^2	$(2.171 \pm 0.036)\%$	$(2.137 \pm 0.036)\%$	21.54

7.3.2 Data

We give our measured branching ratio and number of signal events for $D_s \to \phi e\nu$ in Table 7.7, which includes our correction from background with a peaking D_s mass. The branching ratio here assumes the Monte Carlo's efficiency; our systematics section (Section 10) discusses some corrections to this efficiency (Table F.18) that appear in our final result (Table 11.10).

Table 7.7: $D_s \to \phi e\nu$ measurement in the data, including the peaking background correction from Table 7.4.

Measurement	Branching Ratio	# Events
Raw fit result	$(1.935 \pm 0.152)\%$	208.0 ± 16.6
Peaking BG correction	$(0.012 \pm 0.012)\%$	1.3 ± 1.2
$\mathcal{B}(D_s \to \phi e\nu)$	$(1.923 \pm 0.153)\%$	206.7 ± 16.7

Figure 7.6 shows the result of our likelihood fit on the D_s mass spectrum after our ϕ, electron, and event selections. The plot shows the sum over all D_s masses and fit functions for simplicity, even though the underlying mass distribution varies by tag mode. We give our individual functions and mass plots for each tag mode in Figures 7.7 and 7.8. These individual plots better represent how the likelihood fit operates, although the common branching ratio does connect each of the tag modes' signal normalizations to one another.

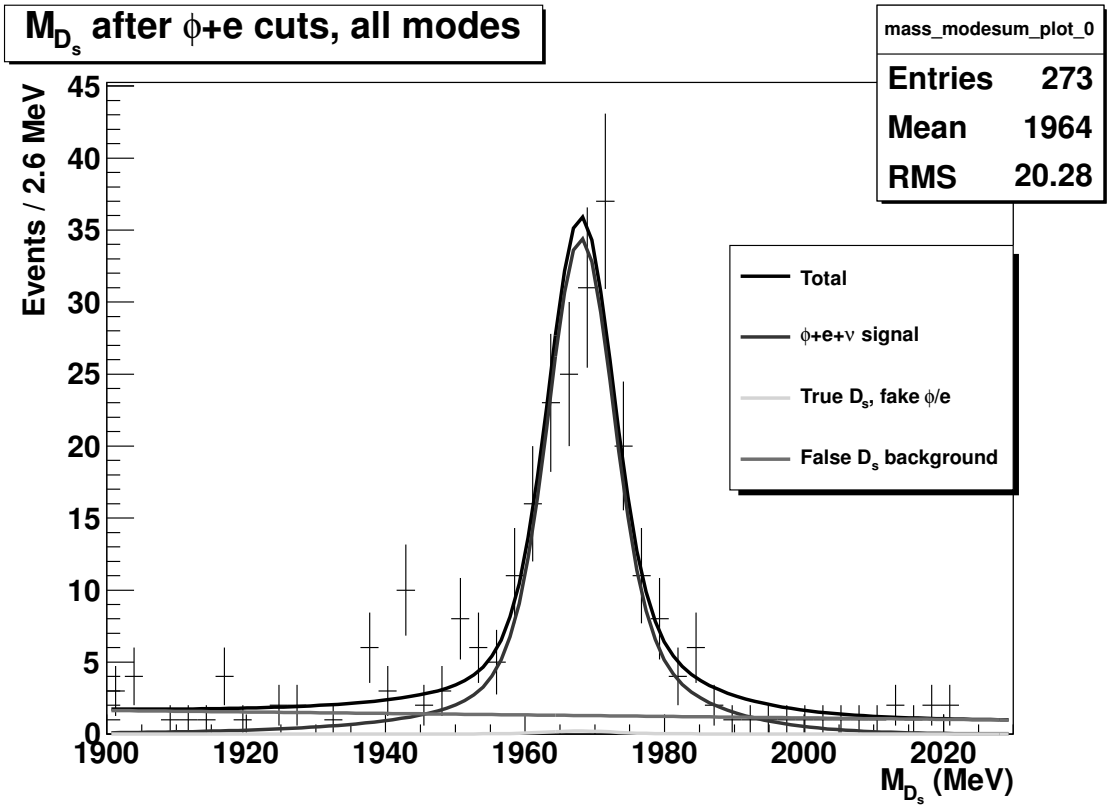

Figure 7.6: $D_s \to \phi e \nu$ data results after our semileptonic selections. We fit the tagged M_{D_s} spectrum with a common branching ratio across all 13 tag modes. The likelihood uses each tag mode's signal shape on its corresponding masses; the above results show a sum over all tag modes.

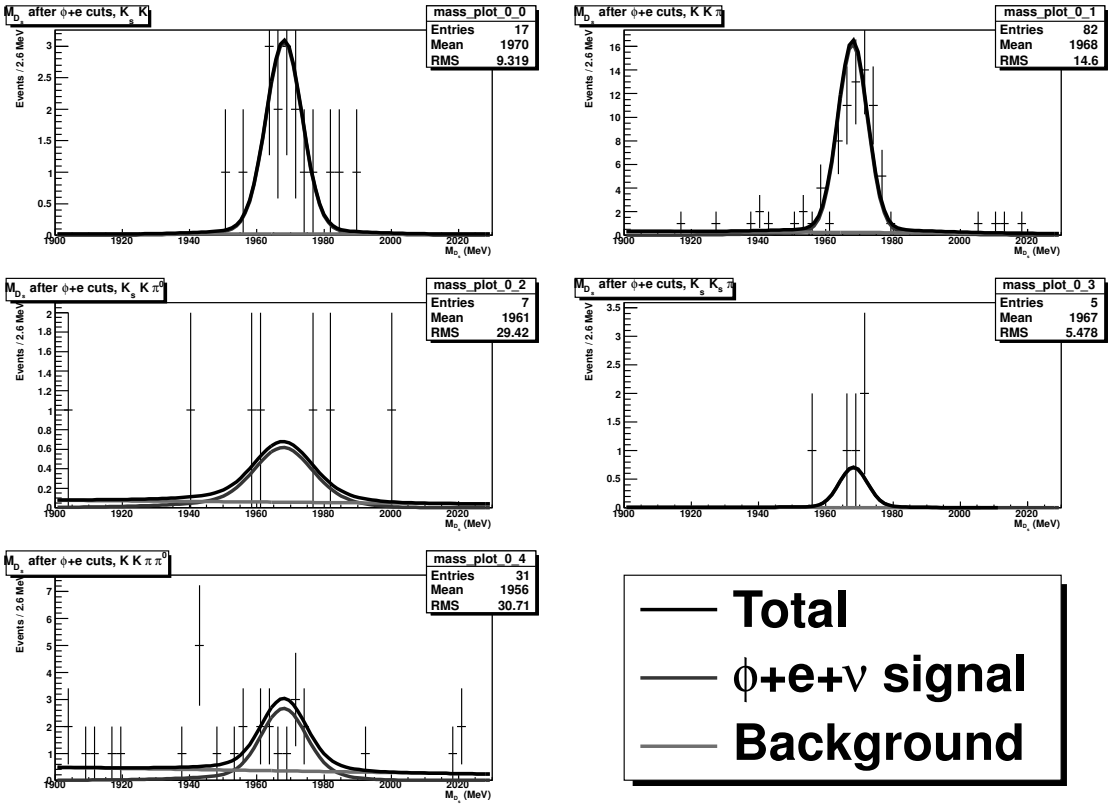

Figure 7.7: Fit results in the data after applying $D_s \to \phi e\nu$ semileptonic cuts for modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

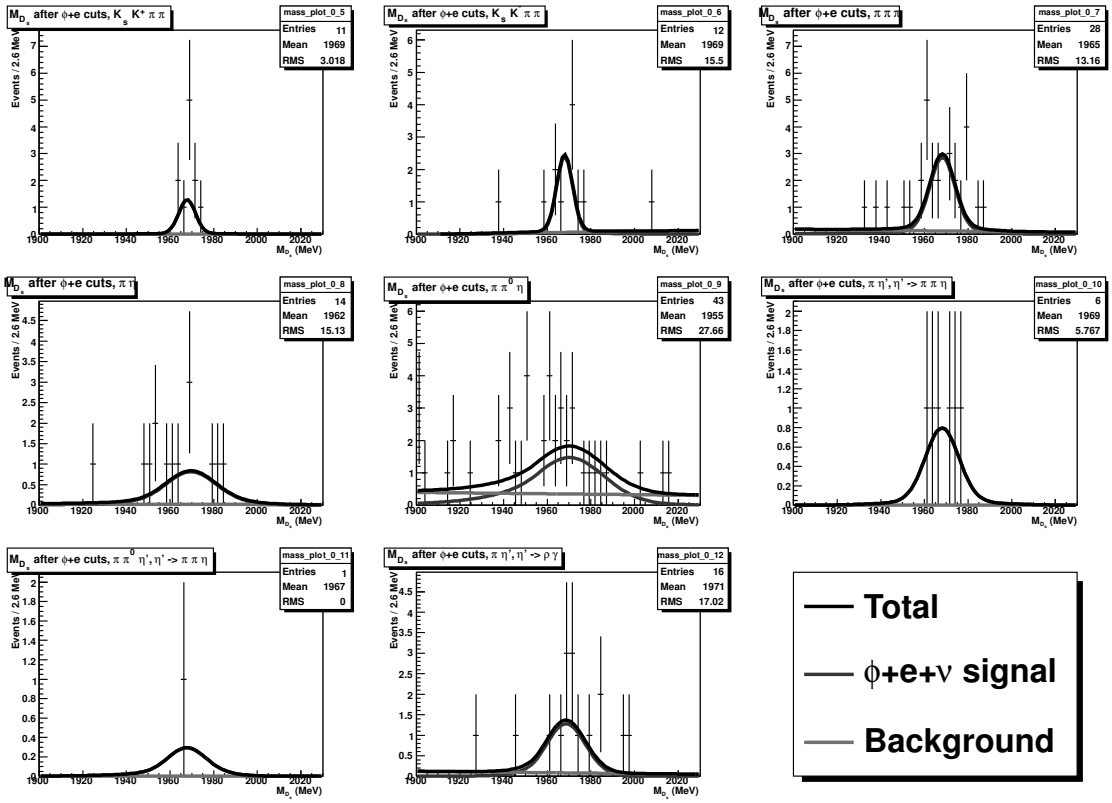

Figure 7.8: Fit results in the data after applying $D_s \to \phi e\nu$ semileptonic cuts for modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode receives an independent background normalization.

Chapter 8

Measurement of $D_s \to (K_s, K^*, f_0, \eta')e\nu$

Our procedure for the D_s semileptonic modes $K_s e\nu$, $K^* e\nu$, $f_0 e\nu$, and $\eta' e\nu$ closely follows that for $D_s \to \phi e\nu$. We use the same D_s tags and electron selections for each of the additional four semileptonic modes, but each mode has its own cuts for the semileptonic-side's hadron. In some cases, we also make additional background rejection cuts and event missing mass cuts, as described in Sections 8.2 and 8.3.

8.1 General Particle Cuts

The semileptonic modes $K^* e\nu$, $f_0 e\nu$, and $\eta' e\nu$ each have pions in their final states. CLEO has made substantial effort to understand pion behavior in the detector over a range of momenta, so we follow the standard pion selections [68] for these modes:

- $\left|\sigma_\pi^{dE/dx}\right| < 3.0$
- $|\cos(\theta)| < 0.93$
- $|d_b| < 5$ mm
- $|z_0| < 5$ cm
- $\chi^2 < 10,000$

- Hit Fraction > 0.5
- If $p_\pi > 550$ MeV and we have info from the RICH, we also require:
  - 2 or more photons in the RICH (pion hypothesis)
  - Particle ID: $(LLH_\pi - LLH_\mathrm{K}) + (\sigma_\pi^2 - \sigma_\mathrm{K}^2) < 0$,
    where LLH refers to the RICH log likelihood and σ comes from the dE/dx measurement.

$K^*e\nu$ also has a kaon in its final state. While $\phi e\nu$ required two kaons to make a ϕ meson, the kaon in a K^* reconstruction gets paired with a pion. This forces us to be more aggressive in our kaon selections to avoid excess combinatoric pairings. Relative to $\phi e\nu$, we have increased the hit fraction requirement (although only to 0.3), we have added a particle ID cut, and we have slightly decreased the χ^2 requirement. For kaons in $K^*e\nu$, we require

- $\left|\sigma_\mathrm{K}^{dE/dx}\right| < 3.0$
- $|\cos(\theta)| < 0.93$
- $|d_b| < 5$ mm
- $|z_0| < 5$ cm
- $\chi^2 < 10,000$
- Hit Fraction > 0.3
- If $p_\mathrm{K} > 550$ MeV and we have info from the RICH, we also require:
  - 1 or more photons in the RICH (kaon hypothesis)
  - Particle ID: $(LLH_\pi - LLH_\mathrm{K}) + (\sigma_\pi^2 - \sigma_\mathrm{K}^2) > 0$,
    where LLH refers to the RICH log likelihood and σ comes from the dE/dx measurement.
- If $p_\mathrm{K} < 550$ MeV and we have info from the RICH, we require 4 or fewer photons (pion hypothesis) in the RICH.

8.2 $D_s \to K_s e\nu$

Unlike $D_s \to \phi e\nu$, where our signal events dominated our background even with fairly loose cuts, $D_s \to K_s e\nu$ has a large background component. While some of this background comes from combinatoric effects, by far the dominant contribution comes from other D_s semileptonic modes that have both a valid D_s tag and a valid electron. $\phi \to K_L K_s$ in $\phi e\nu$ produces the most problematic background as it also has a true K_s, so tightening our K-short cuts will not help our background rejection ($K^* \to K_s \pi^0$ gives a similar but smaller problem).

To deal with the problem of $D_s \to K_s e\nu$ background coming from other semileptonic modes, we have reintroduced the notion of an event missing mass cut. In this case, the missing four vector consists of both the neutrino and the D_s^* meson's daughter photon. Even though the missing mass doesn't peak at zero as it would with only a missing neutrino, we still get good separation between $K_s e\nu$ and $\phi e\nu$ since the soft photon doesn't push the missing mass as far as the extra K_L (or, to a lesser extent, as far as the extra π^0 in $K^* e\nu$). We show this missing mass separation in Figures 8.1 and 8.2. We've optimized our missing mass cut using a simple figure of merit $\left(\frac{S^2}{S+B}\right)$, as seen in Figure 8.3.

In addition to our peaking background, the missing mass cut rejects the majority of our combinatoric background. However, we still retain more total combinatoric events than signal events. As these combinatoric events often lack a true K_s, we have also considered adding a flight significance cut to our K_s selection. We again use a simple figure of merit for events in our signal region to evaluate our potential flight significance cut. We present this figure of merit for various flight significance cuts in Figure 8.4. While the Monte Carlo analysis favors a very large minimum flight significance, we have elected to make the cut at 4.0 to capture most of the benefit while avoiding any potential systematics from data/Monte Carlo deviations at higher flight significance values.

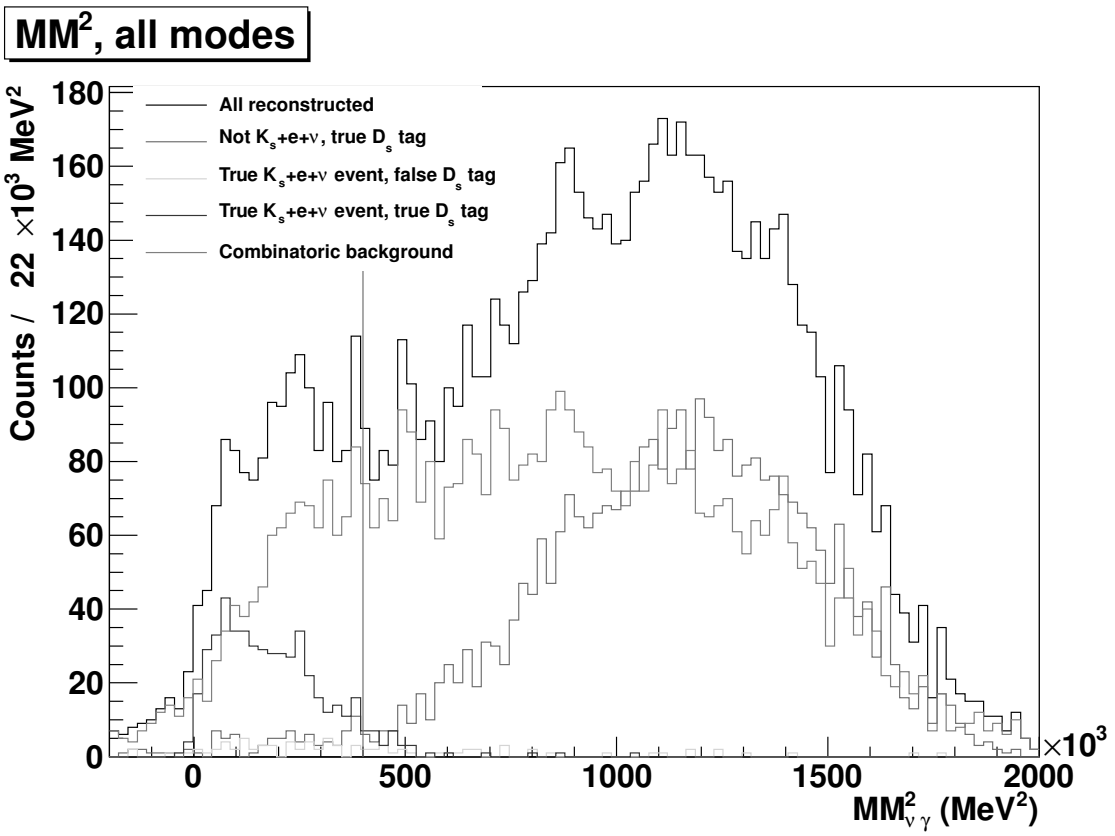

Figure 8.1: Missing mass after finding the D_s tag, K_s, and electron in $D_s \to K_s e\nu$, from the 20× Monte Carlo. We keep all events with a MM^2 below the line at 0.4 GeV^2.

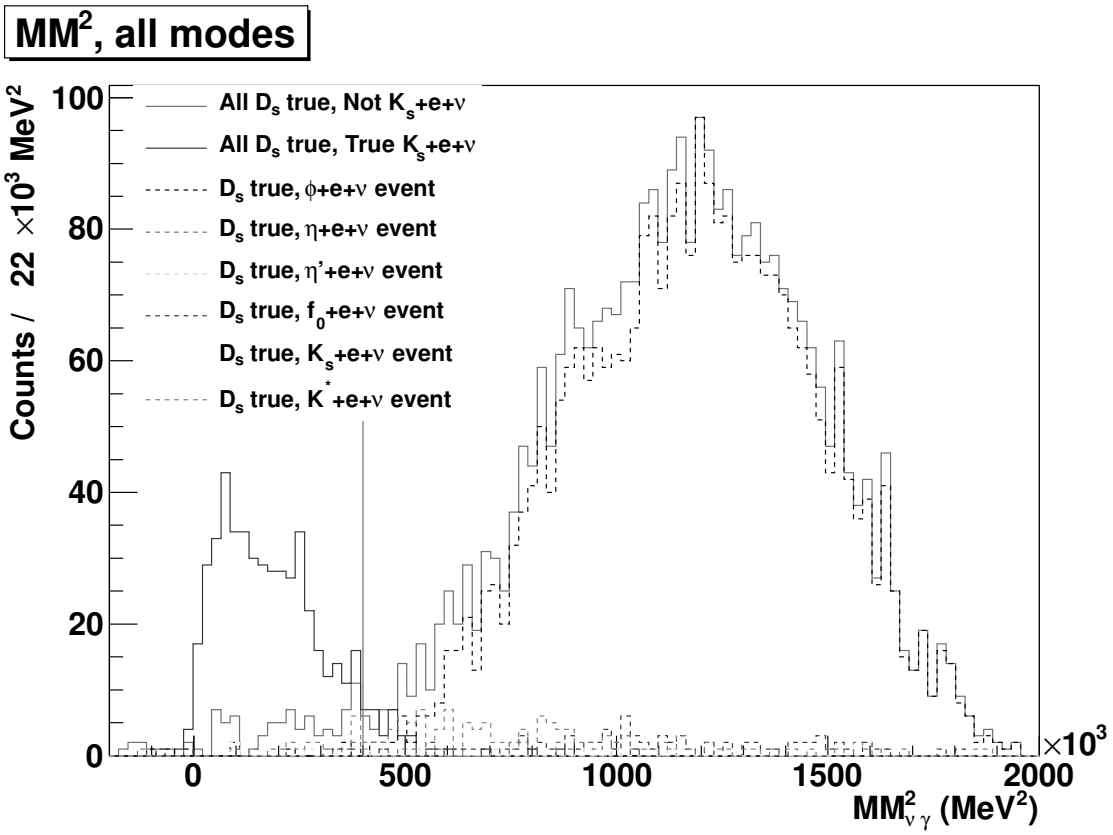

Figure 8.2: $D_s \to K_s e\nu$ missing mass in the 20× Monte Carlo, by background semileptonic mode. The dominant background comes from $D_s \to \phi e\nu$, where $\phi \to K_L K_s$. Our missing mass cut removes most of this background.

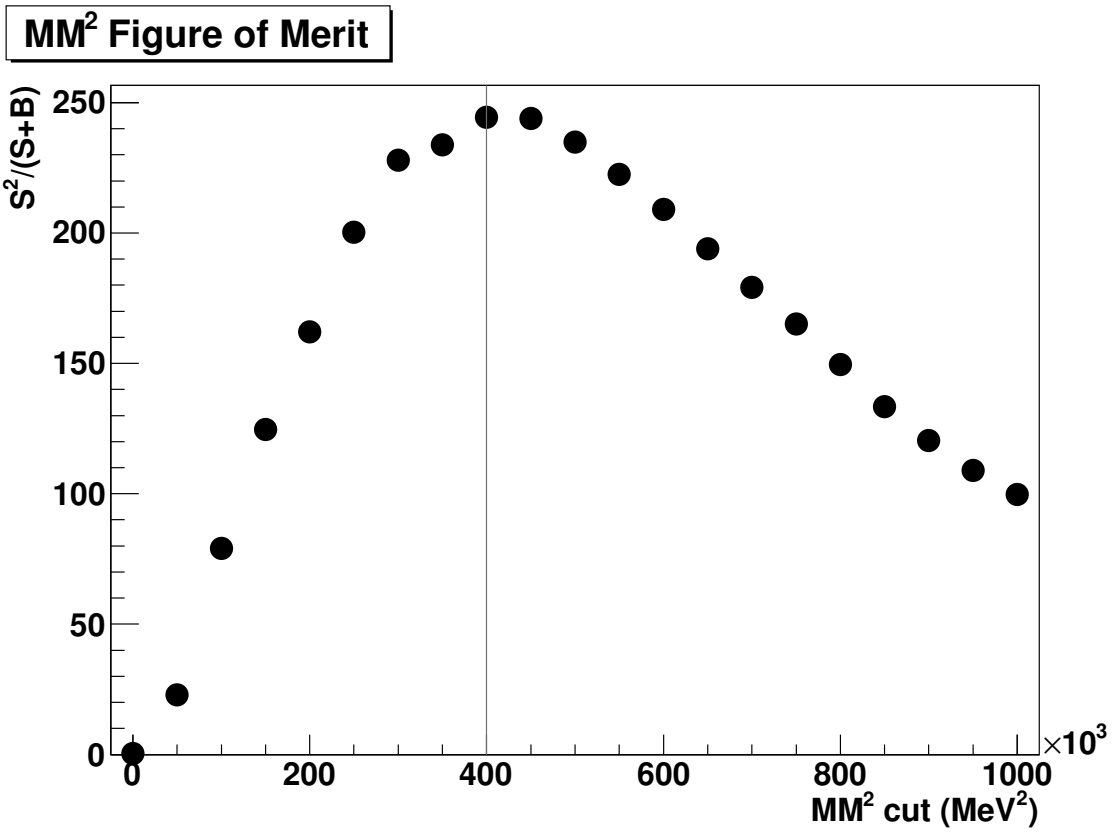

Figure 8.3: Our figure of merit for different missing mass cut values in $D_s \to K_s e\nu$. We only consider signal and background events that have a reasonable D_s mass, between 1955 MeV– 1985 MeV.

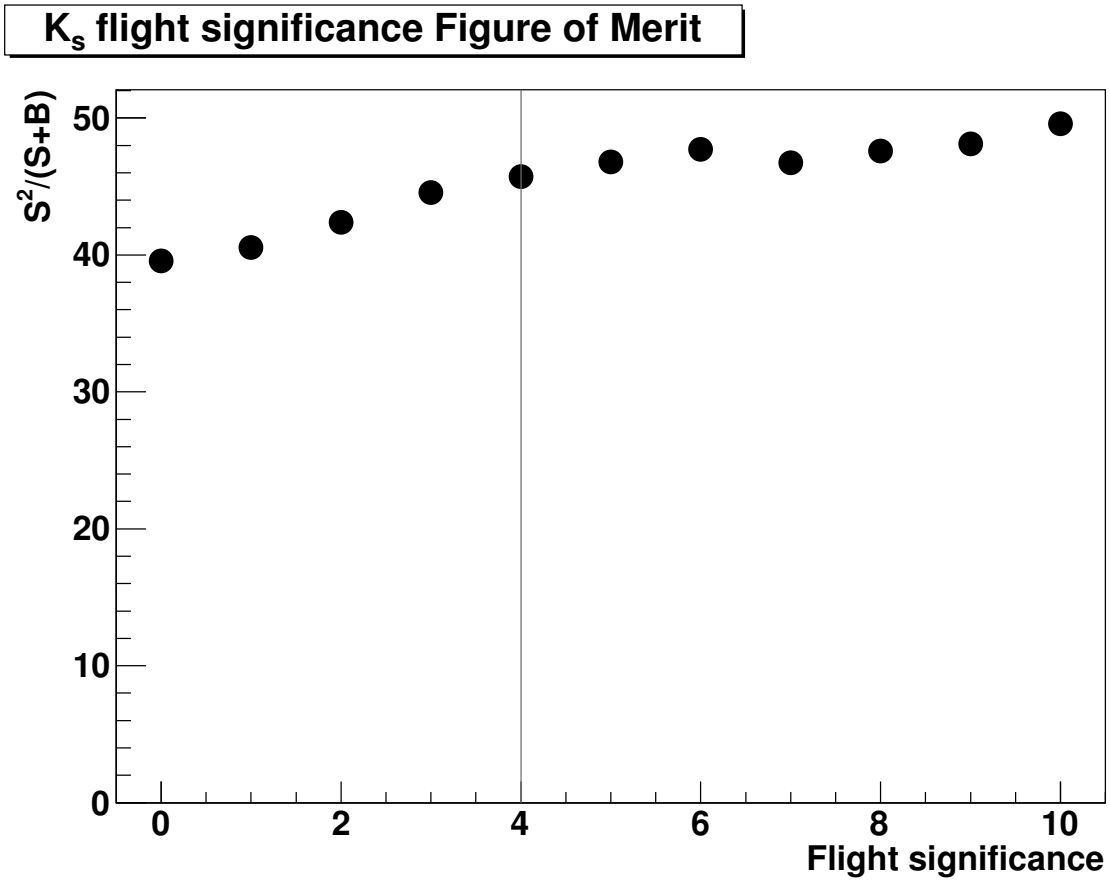

Figure 8.4: Figure of merit for different flight significance cuts in $D_s \to K_s e\nu$. We only consider events with a D_s tag mass within 1955 MeV and 1985 MeV, since events outside that region will be dismissed as background in our final fit.

We make a few additional rejections on potential $K_s e\nu$ events to avoid specific backgrounds. In particular, we have small but identifiable backgrounds from $D_s \to KK_s$ and $D_s \to KK_s\pi^0$ where the kaon fakes an electron. To deal with these, we treat the electron track like a kaon and reject the event if it combines with the K_s (or $K_s + \pi^0$) to have a D_s mass within 10 MeV of the PDG value.

We also see some background with real electrons, notably in $D_s \to \tau\nu$ where a fake K_s gets formed far out in the detector. We eliminate most of these events by requiring that each π track forming the K_s comes from within 20 cm of the origin. This cut also removes a small, similar background from $D_s \to \eta e\nu$.

We give our full list of $D_s \to K_s e\nu$ semileptonic-side requirements below:

- $|M_{K_s}^{\text{recon}} - M_{K_s}^{\text{PDG}}| < 6.3$ MeV
- K_s flight significance > 4.0
- $MM^2_{\gamma\nu} < 400,000$ MeV^2 (.40 GeV^2)
- Reject if $|M_{KK_s}^{\text{recon}} - M_{D_s}^{\text{PDG}}| < 10$ MeV when the electron is treated as a kaon
- If any π^0 are found, reject if $|M_{KK_s\pi^0}^{\text{recon}} - M_{D_s}^{\text{PDG}}| < 10$ MeV when the electron is treated as a kaon
- $\rho_0^\pi \equiv \sqrt{d_b^2 + z_0^2} < 20$ cm for the π from K_s

We obtain a much improved signal relative to background, particularly peaking background, by making these cuts as seen in Table 8.1. In this table, our "before cuts" column only uses the K_s mass cut, while the "after cuts" column contains the events after making the other listed cuts. Figures G.32 (before cuts) and G.33 (after cuts) in Appendix G further breakdown the modes that produce a peaking background.

After all cuts, we get an efficiency for $D_s \to K_s e\nu$ of 30.9% (including the K_s branching ratio). For ease of comparison, we have placed our efficiencies for this and the other non-η branching ratios in Table 8.13, at the end of the section. Our stated efficiencies come from signal Monte Carlo with different D_s production modes; we cover the systematic uncertainty from this weighting for all modes in Section 10.11.

$D_s \to K_s e\nu$ still has some remaining peaking background (true D_s) events from each of the other semileptonic modes we measure, as well as peaking background from events

Table 8.1: Truth-tagged breakdown for $D_s \to K_s e\nu$ candidates in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Candidates before cuts	Candidates after cuts
True D_s tag, true $D_s \to K_s e\nu$	24.95	23.00
True D_s tag, not $D_s \to K_s e\nu$	162.30	3.05
False D_s tag, true $D_s \to K_s e\nu$	4.35	2.55
False D_s tag, not $D_s \to K_s e\nu$	292.30	62.30

with a fake electron. We correct our $D_s \to K_s e\nu$ branching ratio for the semileptonic background sources by using their faking efficiencies and our measured branching ratios. Since the non-semileptonic fakes come from more precisely measured D_s decay modes (primarily through some of the same decay modes that we use for tags), we simply use a Monte Carlo estimate in their branching ratio correction. We give the individual components and the overall sum of these corrections in Table 8.2.

Table 8.2: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to K_s e\nu)$ from peaking background.

Background mode	BR correction
$D_s \to \eta e\nu$	$(0.0010 \pm 0.0005)\%$
$D_s \to f_0 e\nu$	$(0.0001 \pm 0.0001)\%$
$D_s \to \phi e\nu$	$(0.0014 \pm 0.0005)\%$
$D_s \to K^* e\nu$	$(0.0044 \pm 0.0011)\%$
$D_s \to \eta' e\nu$	$(0.0002 \pm 0.0002)\%$
Non-semileptonic	$(0.0052 \pm 0.0011)\%$
Total $D_s \to K_s e\nu$ correction	$(0.0123 \pm 0.0017)\%$

As in $D_s \to \phi e\nu$, we have used our twenty data-sized Monte Carlo samples for an in-vs.-out test to ensure that our fitting procedure accurately measures the input branching ratio. Our χ^2 of 13.3 over the twenty samples and overall fit-to-truth difference of 0.5σ gives us confidence that our underlying procedure works on this sample size and background rate (full results in Tables F.3 and F.4 from Appendix F).

We show our $D_s \to K_s e\nu$ data fit results in Figure 8.5 for the sum across all D_s tag modes, with the component fits from each tag mode in Figures G.34 and G.35

(Appendix G). We present both the raw branching ratio and our branching ratio after correcting for peaking background (but before correcting the Monte Carlo efficiency for systematic biases) in Table 8.3.

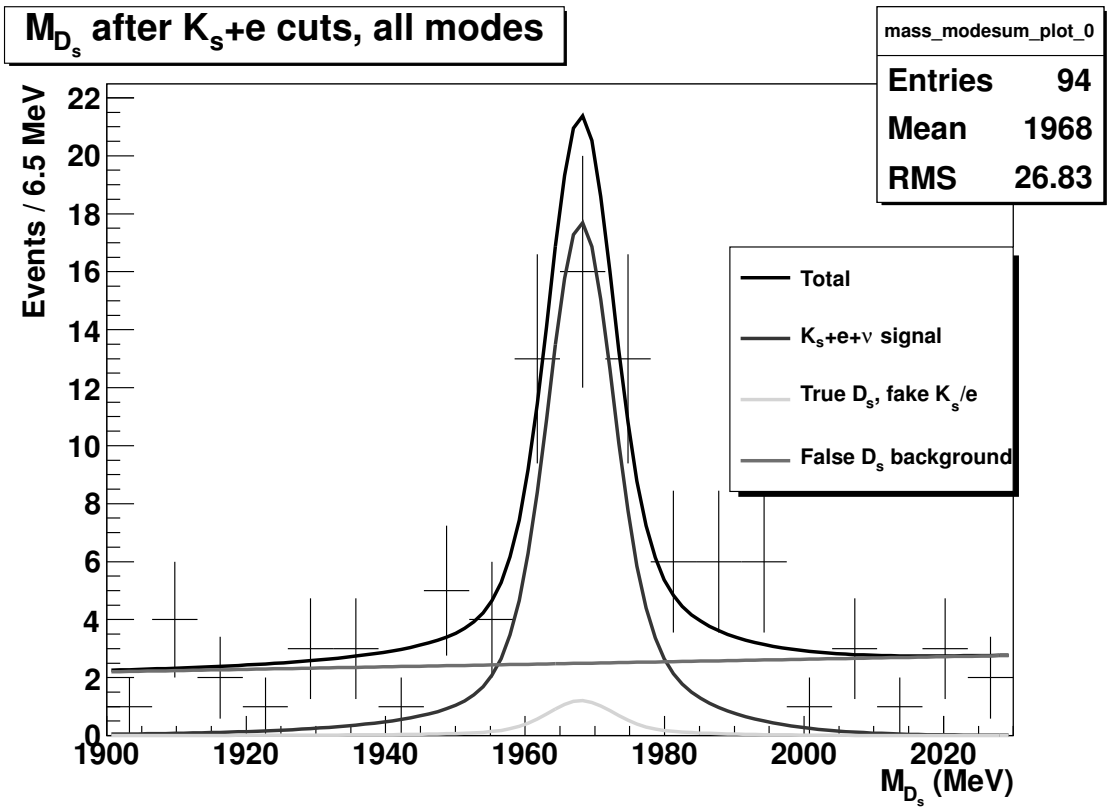

Figure 8.5: Our $D_s \to K_s e\nu$ data fit to the M_{D_s} spectrum, after all semileptonic selections. This fit function represents the sum over the fit functions from each D_s tag mode.

Table 8.3: $D_s \to K_s e\nu$ measurement in the data, including the peaking background correction from Table 8.2.

Measurement	Branching Ratio	# Events
Raw fit result	$(0.186 \pm 0.035)\%$	44.5 ± 8.4
Peaking BG correction	$(0.012 \pm 0.002)\%$	2.9 ± 0.4
$\mathcal{B}(D_s \to K_s e\nu)$	$(0.173 \pm 0.035)\%$	41.5 ± 8.4

8.3 $D_s \to K^* e\nu$

$D_s \to K^* e\nu$ has a small branching ratio relative to the other D_s semileptonic modes considered in this analysis, but it also has a relatively low number of background events. Our only problematic backgrounds come from $D_s \to KK\pi$ (e.g. K^*K) when a kaon fakes an electron and from $D_s \to \phi e\nu$ when one of the kaons fakes the K^* pion.

We can deal with most of the kaon-faking-electron background by simply treating the electron as a kaon and rejecting the event if it combines with the K^* to form a D_s. We similarly deal with the $\phi e\nu$ background by treating the K^* daughter pion as a kaon and rejecting the event if it pairs with the other kaon to form a ϕ. A missing mass cut on the event rejects much of the combinatoric background (Figures 8.6 and 8.7), which would otherwise present the largest remaining challenge to our measurement. As in $D_s \to K_s e\nu$, our missing 4-vector consists of both the neutrino and the unobserved D_s^* daughter photon, shifting the missing mass away from zero.

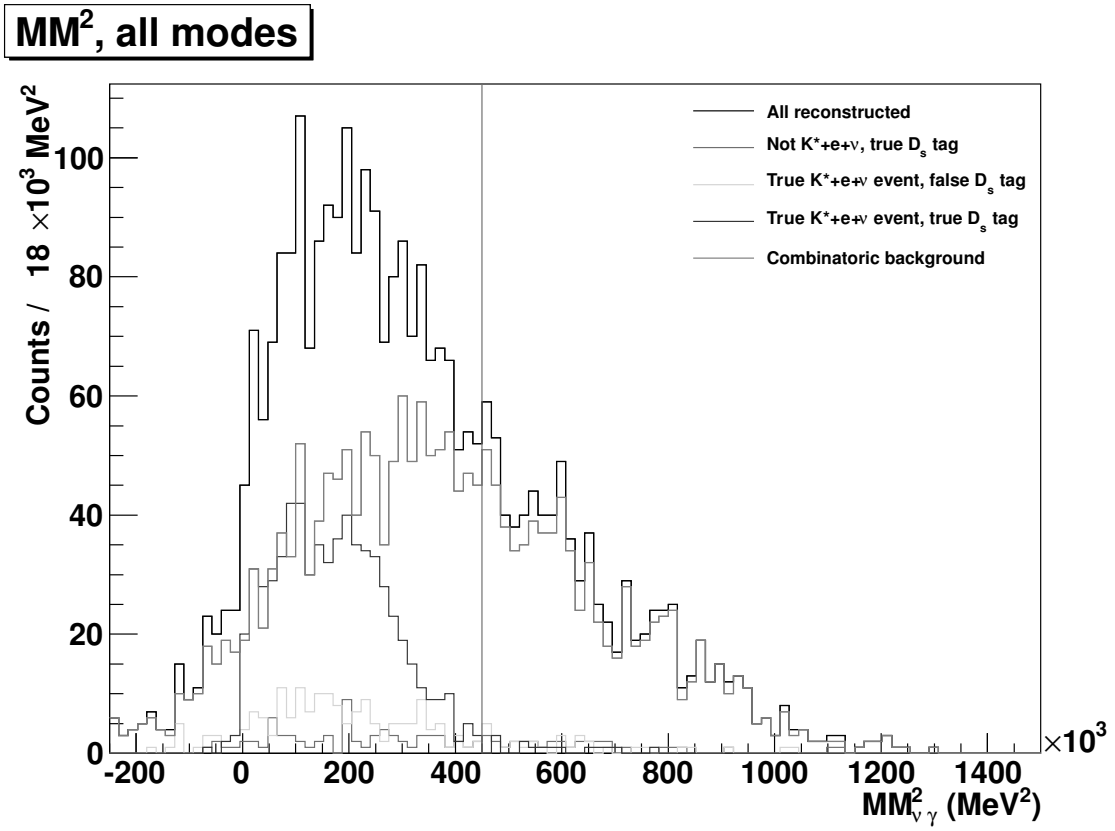

Figure 8.6: $D_s \to K^* e\nu$ missing mass after finding the D_s tag, the electron, and the K^*, from the 20× Monte Carlo. Most of the peaking background has a low $MM^2_{\nu\gamma}$, but we remove almost half of the combinatoric background with our cut.

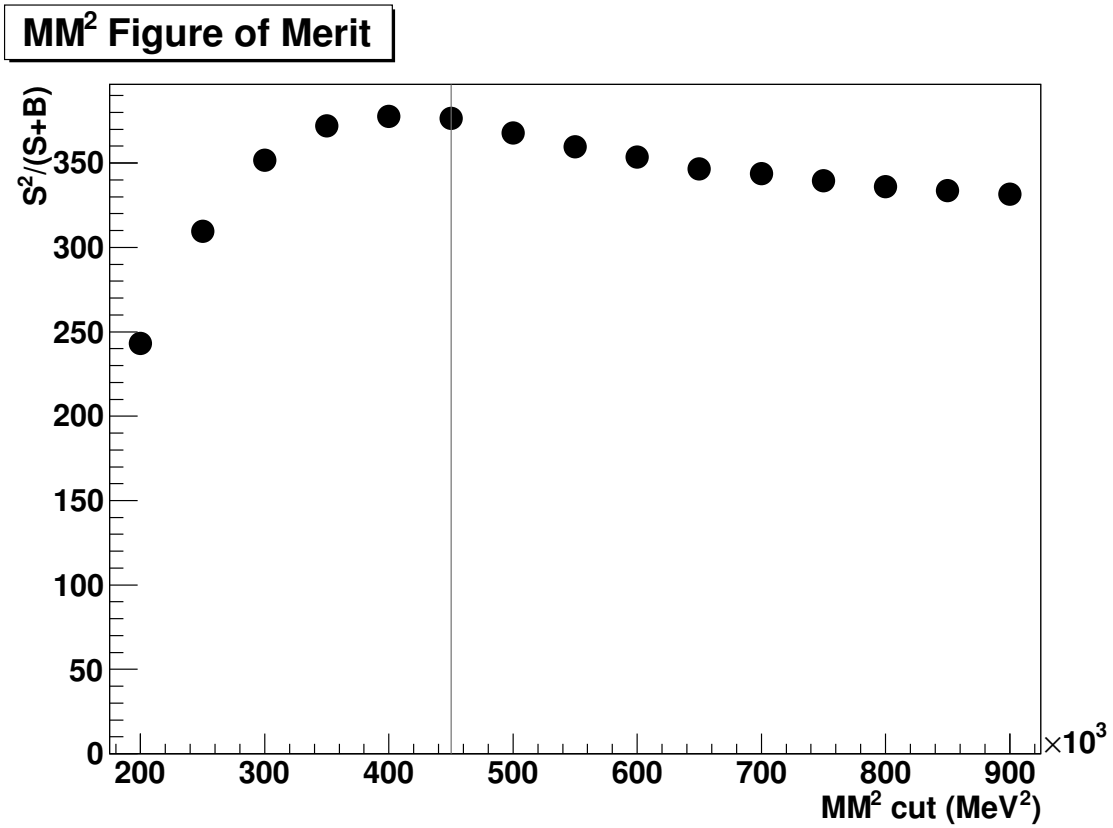

Figure 8.7: Our figure of merit for different missing mass cut values in $D_s \to K^* e\nu$. We only consider signal and background events that have a reasonable D_s mass, between 1955 MeV– 1985 MeV.

Since our prior cuts remove most of the background, we only need a loose selection on the K^* invariant mass, as shown in Figures 8.8 and 8.9. Our figure of merit $\left(\frac{S^2}{S+B}\right)$ for the M_{K^*} cut window plateaus around 100 MeV. We choose to keep all K^* within 106 MeV, which corresponds to a 5σ mass cut (about $2\Gamma_{K^*}$).

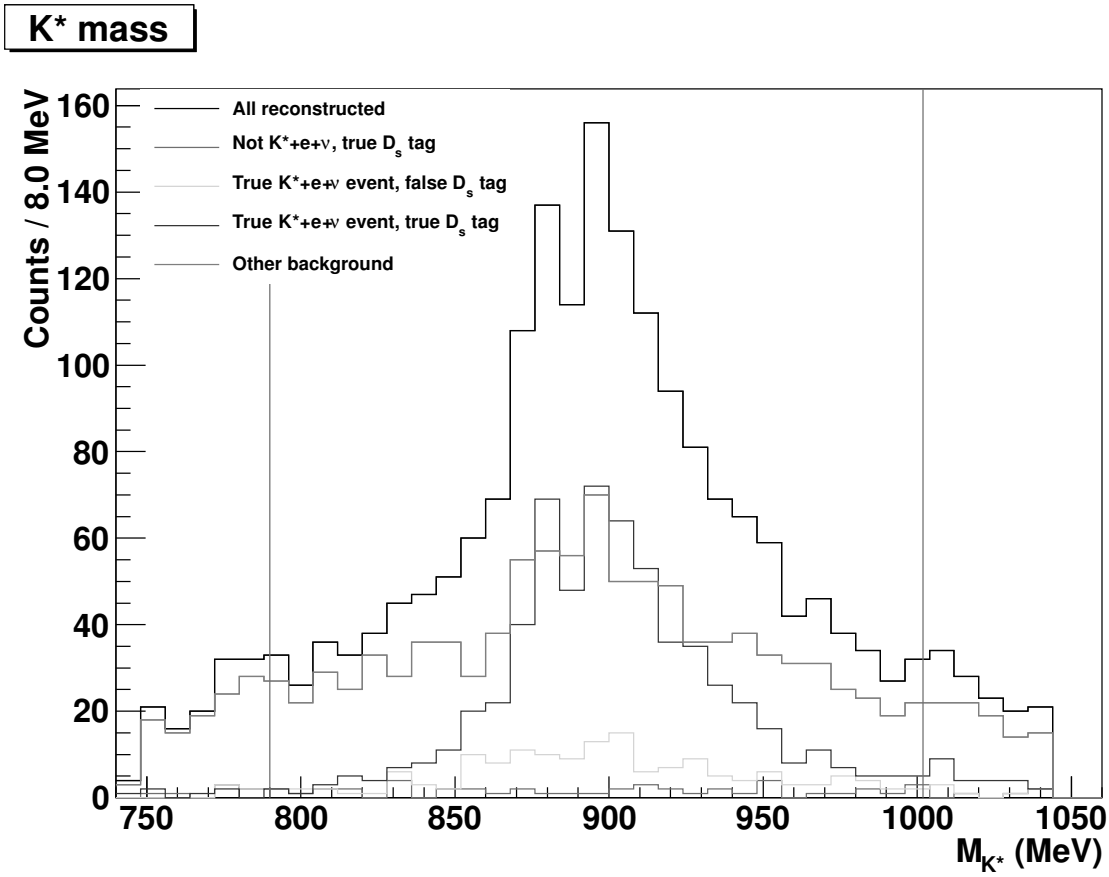

Figure 8.8: K^* mass for events that pass D_s, electron, and our other $D_s \to K^* e\nu$ cuts in the 20× Monte Carlo. Much of our background has a real K^*, so we only obtain moderate background reduction from a tighter mass cut.

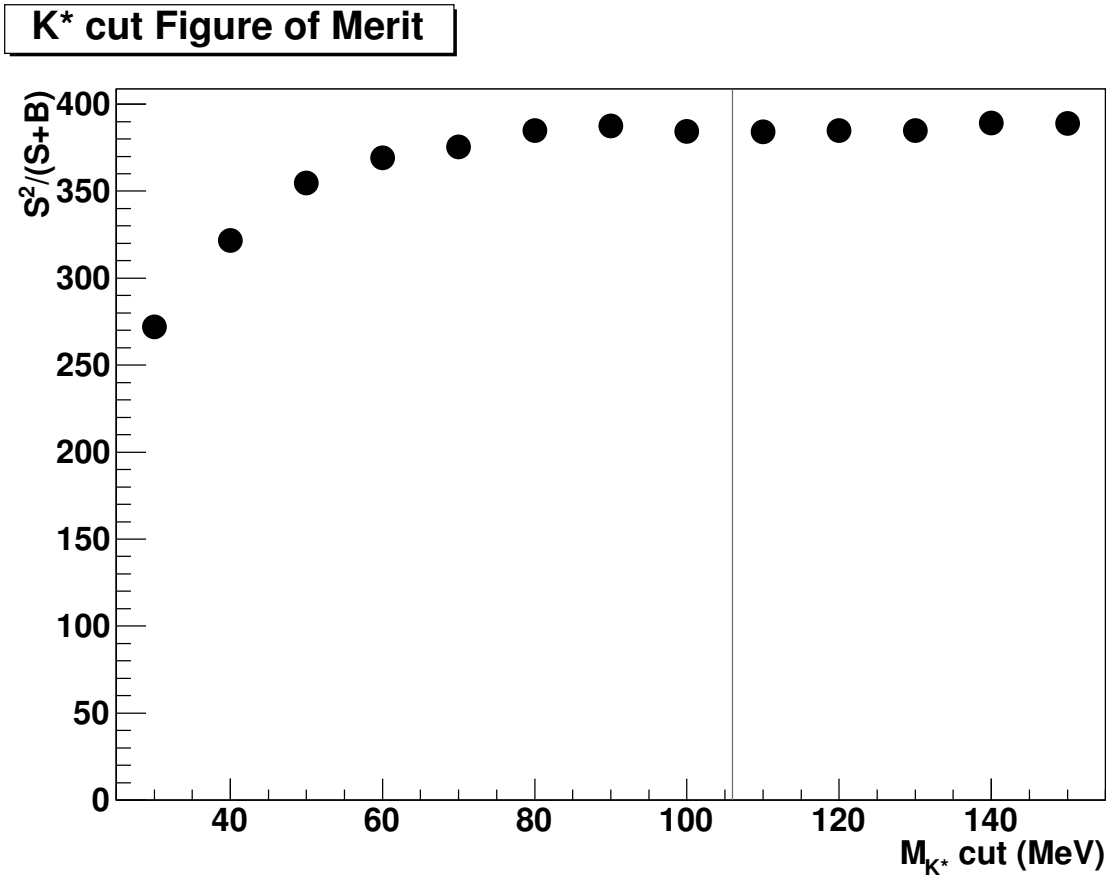

Figure 8.9: Figure of merit for different K^* mass cuts in $D_s \to K^* e\nu$. We only consider events that have a D_s tag mass within 1955 MeV and 1985 MeV, since events outside that region will be dismissed as background in our final fit.

We list our final $D_s \to K^* e\nu$ cuts below:

- The K^* daughter kaon's charge must be opposite to the D_s tag charge
- $|M_{K^*}^{\text{recon}} - M_{K^*}^{\text{PDG}}| < 106$ MeV
- $MM_{\gamma\nu}^2 < 450,000$ MeV^2 (.45 GeV^2)
- Reject if $|M_{KK\pi}^{\text{recon}} - M_{D_s}^{\text{PDG}}| < 20$ MeV when the electron is treated as a kaon
- Reject if $M_{KK} < 1060$ MeV when the K^* daughter pion is treated as a kaon

These cuts remove over half of the peaking background and just under half of the combinatoric background, as seen in Table 8.4. The "before cuts" column only includes the K^* charge requirement and a wide K^* mass cut of 150 MeV. The "after cuts" column includes all our listed cuts. Figure G.36 gives the breakdown of our peaking background before cuts, while Figure G.37 shows the peaking background components after all our cuts. We get a $K^* e\nu$ semileptonic-side efficiency after all cuts of 24.1% (including the K^* branching ratio).

Table 8.4: Truth-tagged breakdown for $D_s \to K^* e\nu$ candidates in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Candidates before cuts	Candidates after cuts
True D_s tag, true $D_s \to K^* e\nu$	33.05	30.15
True D_s tag, not $D_s \to K^* e\nu$	5.35	2.10
False D_s tag, true $D_s \to K^* e\nu$	10.45	7.50
False D_s tag, not $D_s \to K^* e\nu$	126.90	58.85

We correct our raw branching ratio result for peaking background from events with a true D_s. These events come from other semileptonic modes, $\tau\nu$, and D_s decays where a kaon fakes the electron. As in $D_s \to K_s e\nu$, we use our measured branching ratio for the semileptonic correction while using the Monte Carlo rates for the non-semileptonic correction. We show the components of this correction and their sum in Table 8.5.

Table 8.5: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to K^* e\nu)$ from peaking background.

Background mode	BR correction
$D_s \to K_s e\nu$	$(0.0001 \pm 0.0001)\%$
$D_s \to f_0 e\nu$	$(0.0002 \pm 0.0001)\%$
$D_s \to \phi e\nu$	$(0.0055 \pm 0.0012)\%$
$D_s \to \eta' e\nu$	$(0.0017 \pm 0.0007)\%$
Non-semileptonic	$(0.0032 \pm 0.0010)\%$
Total $D_s \to K^* e\nu$ correction	$(0.0107 \pm 0.0017)\%$

Like our other semileptonic measurements, we have used our twenty data-sized Monte Carlo samples for an in-vs.-out test to ensure that our fitting procedure works with the signal and background levels in $D_s \to K^* e\nu$. We placed the results from this test in Tables F.5 and F.6 from Appendix F.

Figure 8.10 contains the result of our data fits, summed over all D_s tag modes. Figures G.38 and G.39 in Appendix G show the fits by individual tag modes. We present our $D_s \to K^* e\nu$ branching ratio measurement before and after peaking background corrections in Table 8.6.

Table 8.6: $D_s \to K^* e\nu$ measurement in the data, including the peaking background correction from Table 8.5.

Measurement	Branching Ratio	# Events
Raw fit result	$(0.180 \pm 0.040)\%$	33.6 ± 7.5
Peaking BG correction	$(0.011 \pm 0.002)\%$	2.0 ± 0.3
$\mathcal{B}(D_s \to K^* e\nu)$	$(0.170 \pm 0.040)\%$	31.6 ± 7.5

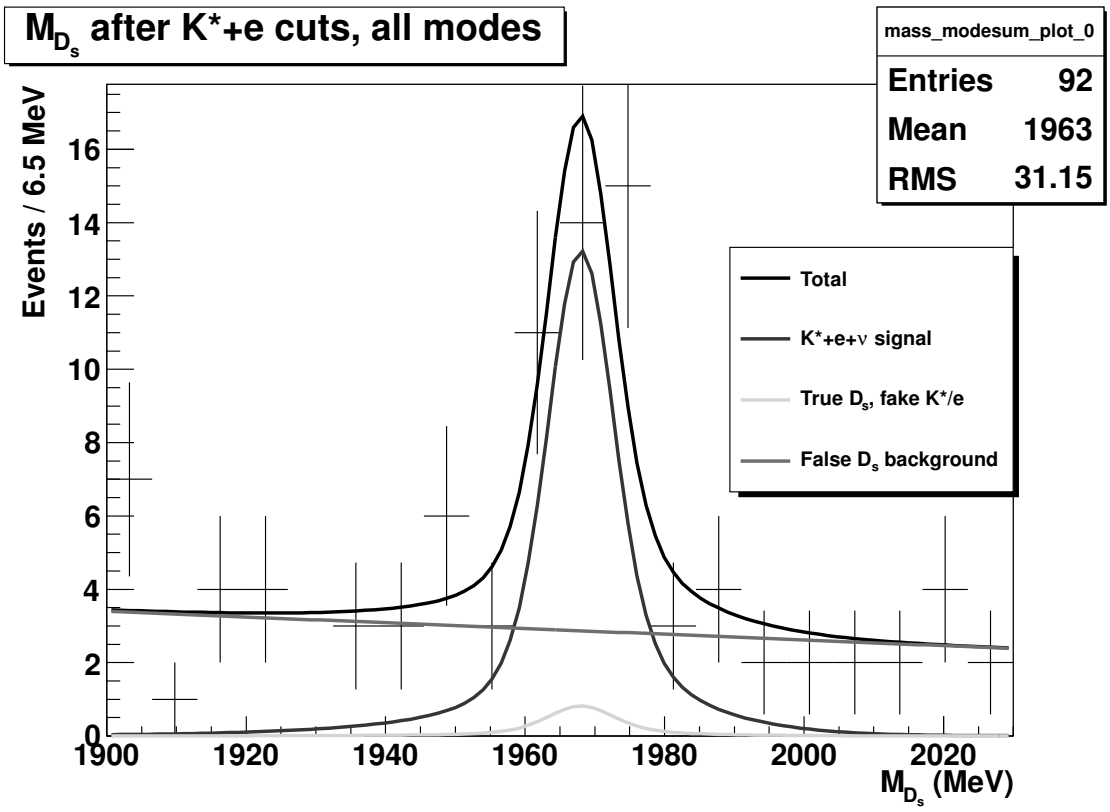

Figure 8.10: Our $D_s \to K^* e\nu$ data fit to the M_{D_s} spectrum, after all semileptonic selections. This fit function represents the sum over the fit functions from each D_s tag mode.

8.4 $D_s \to \eta' e\nu$

Unlike our other semileptonic modes, where we observe the semileptonic side's hadron through its direct decay to two final state particles, we reconstruct the η' in $D_s \to \eta' e\nu$ through its decay to $\pi\pi\eta$ with $\eta \to \gamma\gamma$. This gives us two mass constraints, reducing both our combinatoric and peaking background to very low levels (with signal-to-background comparable to the much higher statistics $\phi e\nu$ mode). The extra constraint allows us to use minimal rejections on the semileptonic side and maintain a high efficiency.

Although we could use the invariant mass of the η as a constraint, we instead choose to use the η pull mass, which takes the different errors on each daughter photon measurement into account. Adding on a comfortable η' mass cut of 10 MeV, we then just have the $D_s \to \eta' e\nu$ cuts

- $|\sigma_\eta| < 3.0$
- $|M_{\eta'}^{\text{recon}} - M_{\eta'}^{\text{PDG}}| < 10$ MeV

As seen in Table 8.7, these cuts give a marginal improvement over our extremely loose "before cuts," with a 5.0 pull mass on the η and a 30 MeV mass cut on the η'. We see a semileptonic-side efficiency for $\eta' e\nu$ of 4.0% after cuts, including all branching ratios.

Table 8.7: Truth-tagged breakdown for $D_s \to \eta' e\nu$ candidates in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Candidates before cuts	Candidates after cuts
True D_s tag, true $D_s \to \eta' e\nu$	24.60	22.75
True D_s tag, not $D_s \to \eta' e\nu$	0.60	0.30
False D_s tag, true $D_s \to \eta' e\nu$	3.75	3.10
False D_s tag, not $D_s \to \eta' e\nu$	12.40	6.25

What little peaking background we do see (Figure G.40) tends to come from $D_s \to \eta e\nu$, where the $\eta \to \pi\pi\pi^0$. A photon from the π^0 then combines with another shower (like the D_s^* daughter photon) to make a fake η. This $\eta e\nu$ peaking background contains few enough events relative to signal that we reject too many true events when we try a direct π^0 reconstruction. We instead just do a peaking background subtraction based on

our $\eta e\nu$ measurement. We've given this correction and our smaller peaking background corrections in Table 8.8.

Table 8.8: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to \eta' e\nu)$ from peaking background.

Background mode	BR correction
$D_s \to \eta e\nu$	$(0.0065 \pm 0.0033)\%$
$D_s \to \phi e\nu$	$(0.0002 \pm 0.0002)\%$
Non-semileptonic	$(0.0035 \pm 0.0024)\%$
Total $D_s \to \eta' e\nu$ correction	$(0.0102 \pm 0.0041)\%$

Our in-vs.-out Monte Carlo test results for $D_s \to \eta' e\nu$ can be seen in Section F, Tables F.7 and F.8.

Figure 8.11 shows our data fits for the sum over all D_s tag modes. We have placed our individual tag mode fits in Appendix G, Figures G.41 and G.42. Table 8.9 gives our raw $D_s \to \eta' e\nu$ measurement and the result after correcting for peaking background.

Table 8.9: $D_s \to \eta' e\nu$ measurement in the data, including the peaking background correction from Table 8.8.

Measurement	Branching Ratio	# Events
Raw fit result	$(0.646 \pm 0.140)\%$	20.4 ± 4.5
Peaking BG correction	$(0.010 \pm 0.004)\%$	0.3 ± 0.1
$\mathcal{B}(D_s \to \eta' e\nu)$	$(0.636 \pm 0.140)\%$	20.1 ± 4.5

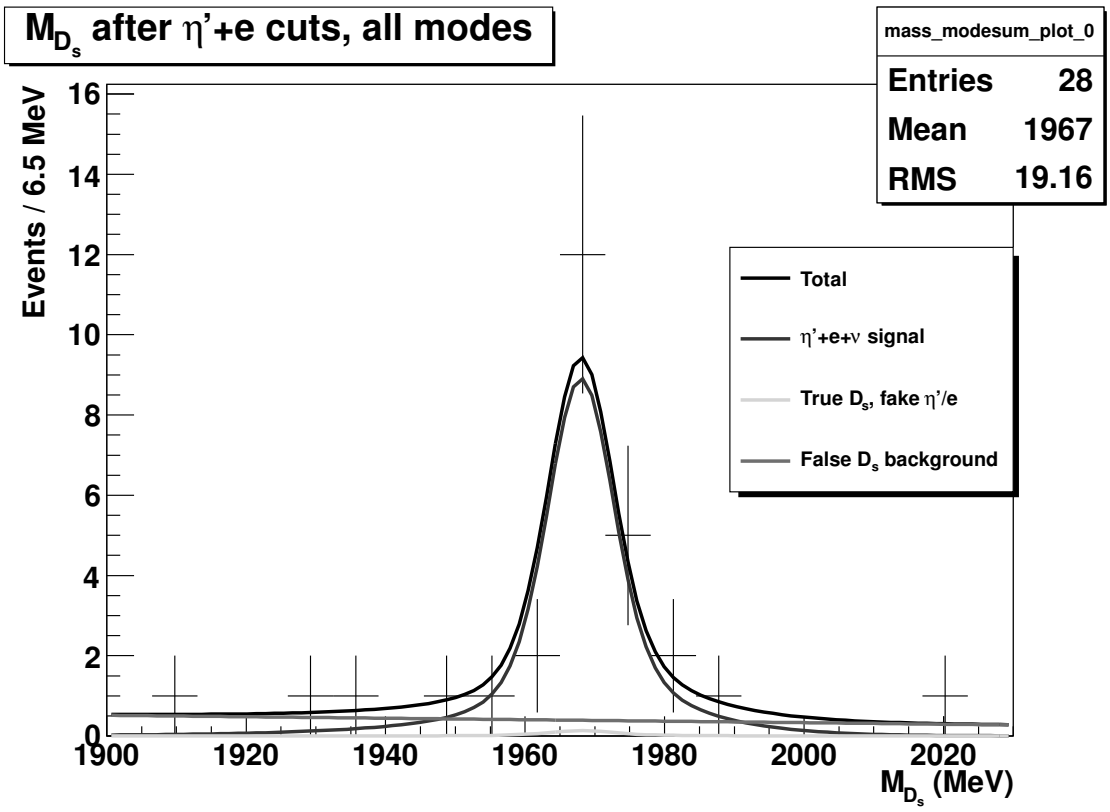

Figure 8.11: Our $D_s \to \eta' e\nu$ data fit to the M_{D_s} spectrum, after all semileptonic selections. This fit function represents the sum over the fit functions from each D_s tag mode.

8.5 $D_s \to f_0 e\nu$

D_s decays often contain kaon final states, which makes the $f_0 \to \pi\pi$ reconstruction mode for $D_s \to f_0 e\nu$ fairly clean. Although $f_0 e\nu$ has some peaking background from other semileptonic modes (like $D_s \to \eta' e\nu$, $\eta' \to \pi\pi\gamma$), most of our peaking background comes from events with a kaon faking the electron. These events also generally involve another kaon faking a pion, as the large V_{cs} ensures that kaons in D_s decays tend to come in pairs. Due to this double-fake rarity, none of the individual modes with electron faking occur with significant frequency. Since we can't simply reconstruct all such modes to reject the event without hitting our signal through the combinatorics, we instead just apply a correction using the Monte Carlo expected rates for such fakes.

As shown in Figures 8.12 and 8.13, we have a fairly low background in $D_s \to f_0 e\nu$, so we use a relatively broad mass cut. Other than our standard pion and electron cuts, our $f_0 e\nu$ reconstruction only contains the f_0 mass cut:

- $|M_{f_0}^{\text{recon}} - M_{f_0}^{\text{PDG}}| < 60$ MeV

We take 980 MeV to be the f_0 PDG mass.

Table 8.10 shows our peaking and combinatoric background levels compared to our signal events. Figure G.43 in Appendix G shows the breakdown of our true D_s tag background, while Table 8.11 has our branching ratio corrections for this peaking background. Given the uncertainty in $\mathcal{B}(f_0 \to \pi\pi)$, we've chosen to quote a result for $D_s \to f_0 e\nu, f_0 \to \pi\pi$ rather than assuming any particular branching ratio.

Table 8.10: Truth-tagged breakdown for $D_s \to f_0 e\nu$ candidates passing all cuts in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Passing candidates
True D_s tag, true $D_s \to f_0 e\nu$	55.10
True D_s tag, not $D_s \to f_0 e\nu$	3.75
False D_s tag, true $D_s \to f_0 e\nu$	6.70
False D_s tag, not $D_s \to f_0 e\nu$	171.80

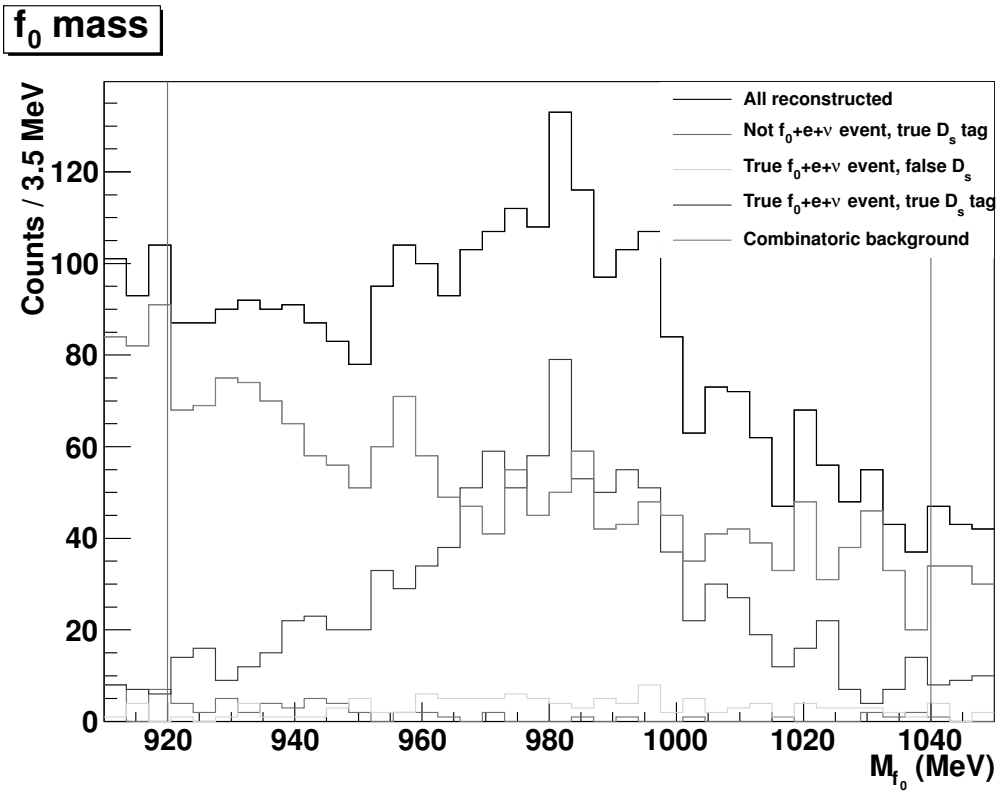

Figure 8.12: f_0 mass for events passing D_s and electron cuts in the 20× Monte Carlo.

Table 8.11: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to f_0 e\nu, f_0 \to \pi\pi)$ from peaking background.

Background mode	BR correction
$D_s \to \phi e\nu$	$(0.0002 \pm 0.0001)\%$
$D_s \to K^* e\nu$	$(0.0001 \pm 0.0001)\%$
$D_s \to \eta' e\nu$	$(0.0014 \pm 0.0005)\%$
Non-semileptonic	$(0.0037 \pm 0.0008)\%$
Total $D_s \to f_0 e\nu, f_0 \to \pi\pi$ correction	$(0.0054 \pm 0.0010)\%$

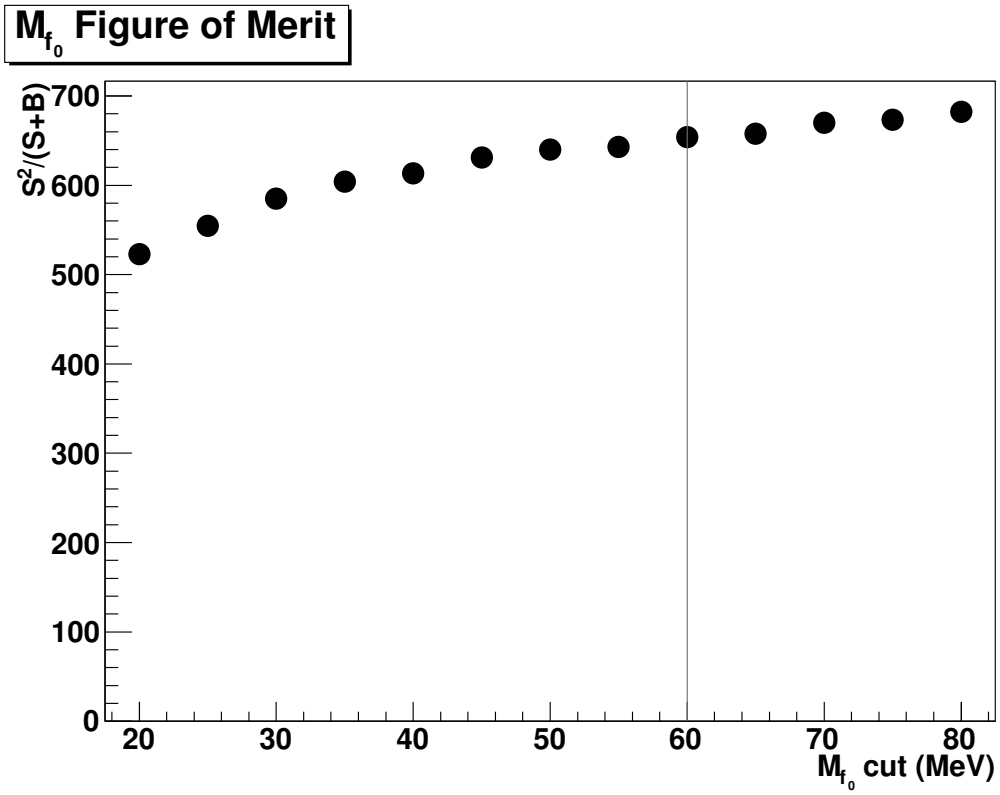

Figure 8.13: Figure of merit for different f_0 mass cuts in $D_s \to f_0 e\nu$, considering only events with a D_s tag mass within 1955 MeV and 1985 MeV. Since the f_0 width has some uncertainty, a 60 MeV mass cut gives us a good balance between retaining most of the signal while not allowing too much excess background.

As with our other modes, we have done an in-vs.-out test for $D_s \to f_0 e\nu$ using the twenty Monte Carlo data-sized samples. Tables F.9 and F.10 in Appendix F contain the results of this comparison.

Figure 8.14 has the summed results of our data fits across all D_s tag modes. Figures G.44 and G.45 in Appendix G contain the individual tag mode fits. We give our raw measurement and background corrected branching ratio for $D_s \to f_0 e\nu, f_0 \to \pi\pi$ in Table 8.12.

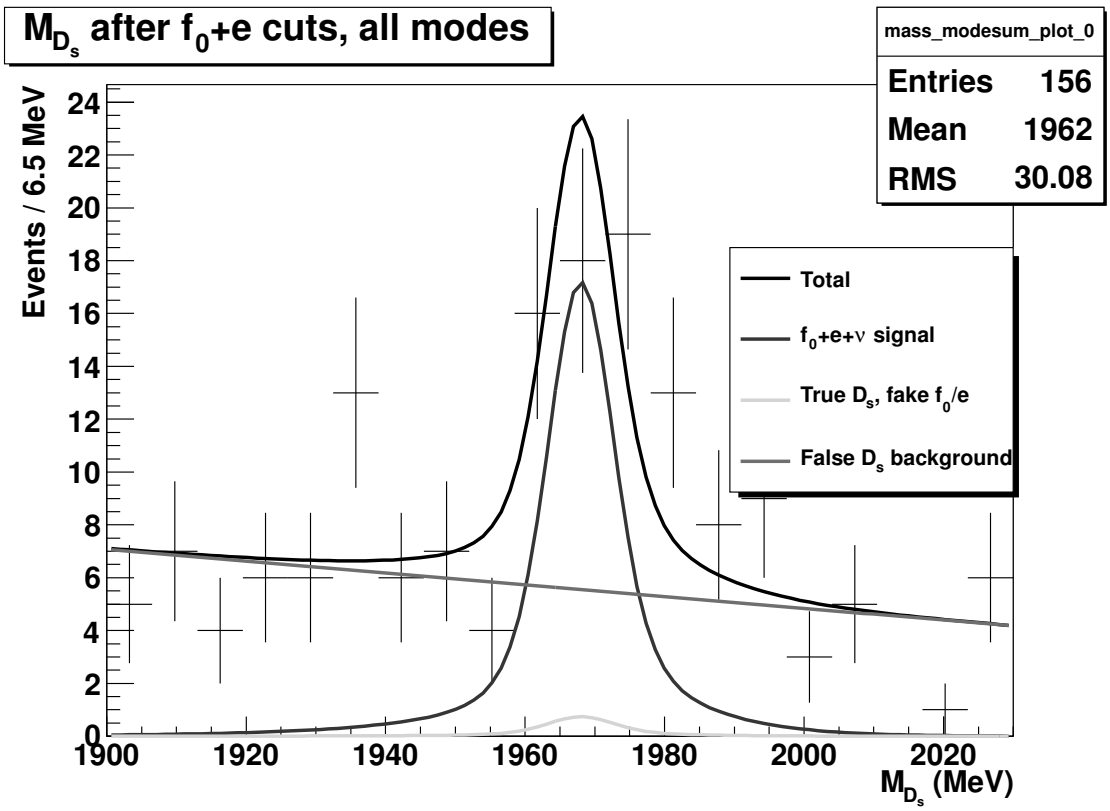

Figure 8.14: Our $D_s \to f_0 e\nu$ data fit to the M_{D_s} spectrum, after all semileptonic selections. This fit function represents the sum over the fit functions from each D_s tag mode.

Table 8.12: $D_s \to f_0 e\nu, f_0 \to \pi\pi$ measurement in the data, including the peaking background correction from Table 8.11.

Measurement	Branching Ratio	# Events
Raw fit result	$(0.135 \pm 0.024)\%$	43.6 ± 7.8
Peaking BG correction	$(0.005 \pm 0.001)\%$	1.7 ± 0.3
$\mathcal{B}(D_s \to f_0 e\nu, f_0 \to \pi\pi)$	$(0.130 \pm 0.024)\%$	41.9 ± 7.8

For comparison, we summarize all our Monte Carlo efficiencies, by semileptonic mode, in Table 8.13. We will later correct these efficiencies through the various effects in Section 10 (Tables F.20–F.23), leading to the final branching ratios in Table 11.10.

Table 8.13: Semileptonic-side efficiencies for $\phi e\nu$, $\eta' e\nu$, $f_0 e\nu$, $K_s e\nu$, and $K^* e\nu$, after all cuts. The first column includes the hadron branching ratios into the efficiency, while the second column gives the efficiency considering only hadron decays to the reconstruced decay mode.

Semileptonic Mode	$\varepsilon_{\rm SL}$	$\varepsilon_{\rm SL}^{\rm excl\ BR}$
$D_s \to \phi e\nu$	$(13.91 \pm 0.18)\%$	$(28.33 \pm 0.38)\%$
$D_s \to K_s e\nu$	$(30.93 \pm 0.53)\%$	$(45.09 \pm 0.77)\%$
$D_s \to K^* e\nu$	$(24.13 \pm 0.35)\%$	$(36.25 \pm 0.53)\%$
$D_s \to \eta' e\nu$	$(4.02 \pm 0.16)\%$	$(23.46 \pm 0.96)\%$
$D_s \to f_0 e\nu$	$(21.69 \pm 0.34)\%$	$(41.72 \pm 0.66)\%$

Chapter 9

Measurement of $D_s \to \eta e \nu$

We find $D_s \to \eta e \nu$ events in much the same way as $D_s \to \phi e \nu$ events: we reconstruct the D_s tag, the electron, and the semileptonic-side hadron, while ignoring the $D_s^* \to D_s \gamma$ transition photon or other extra showers in the event. We originally made the choice to ignore the transition photon in $D_s \to \phi e \nu$ because of the photon's low efficiency, generally low backgrounds in D_s semileptonics, and complications from splitoff showers. We still benefit from leaving the transition photon out of our reconstruction for $D_s \to \eta e \nu$, but the $\eta \to \gamma\gamma$ reconstruction has higher backgrounds from both combinatoric background and splitoff showers that contribute to fake η. The rise in combinatoric background doesn't present a huge problem, as we still have a potential sideband subtraction and fewer background than signal events in the signal region. Splitoff showers used in fake η do cause difficulties, however, since true $\eta e \nu$ events can have improperly reconstructed η with a rate not always well modeled by the Monte Carlo.

With $D_s \to \phi e \nu$, we reconstructed the ϕ but only used it as a consistency check rather than directly involving it in the fit since false ϕ didn't present much of an issue. For $D_s \to \eta e \nu$, however, false η in background modes become a problem, so we need to use the η kinematic information directly in our fit. Specifically, we do a two-dimensional fit to the η pull mass and the D_s mass to determine the number of $D_s \to \eta e \nu$ events. We retain the electron selections and event cuts from Sections 6.1 and 6.2, respectively.

9.1 η Selection

We considered a cut on $\frac{E_9}{E_{25}}$ for the η daughter photons, but we found that the reduction in background did not compensate for the relative 6% efficiency loss (Figure 9.1). In general, we have found that our 2D fit to the η pull mass and D_s invariant mass separates signal from background well enough that we can use a fairly loose selection on the η.

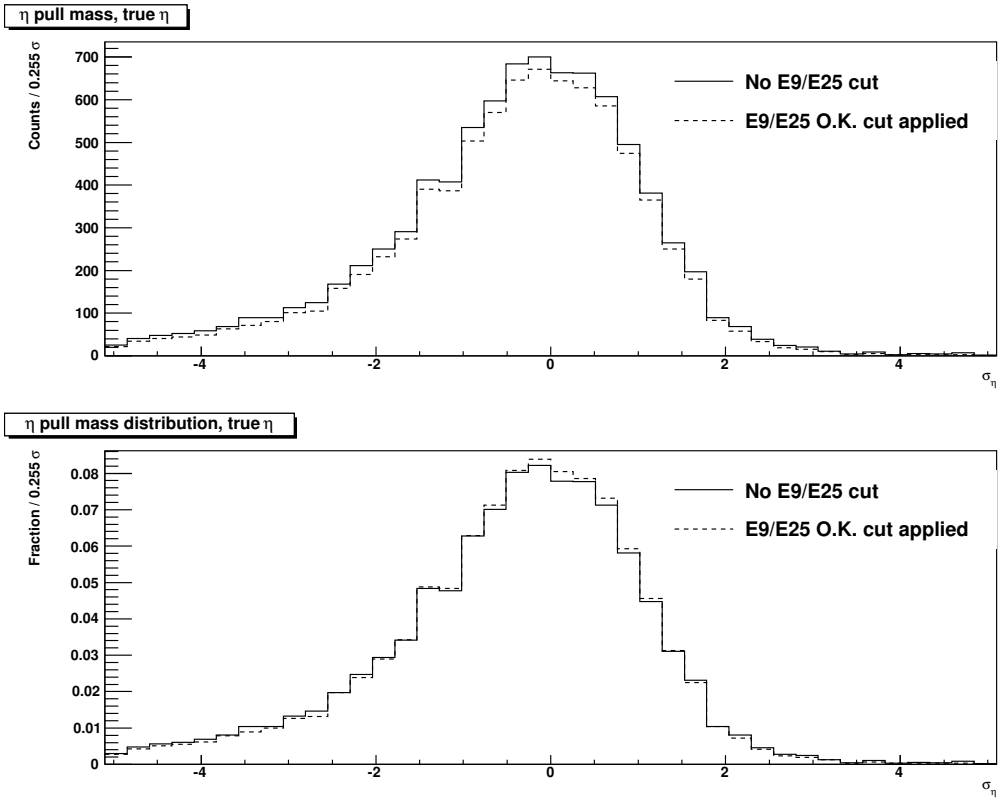

Figure 9.1: The $\frac{E_9}{E_{25}}$ O.K. cut's effect on the η pull mass distribution. Top: Reconstructed η spectrum with and without the $\frac{E_9}{E_{25}}$ O.K. cut. Bottom: Normalized η spectrum with and without the $\frac{E_9}{E_{25}}$ O.K. cut, showing that the cut doesn't disproportionally change the pull mass distribution (slightly lower efficiency for large pull masses).

We make the following selections on the η daughter photons:

- No track matches the shower location

- No showers may come from hot crystals
- The shower must come from the barrel or the endcap of the calorimeter, not the transition region in between
- $E_\gamma > 30$ MeV

Beyond the individual photon cuts, we also make a simple η selection. We use the η pull mass instead of the nominal mass to take advantage of the two showers' uncertainty information. Our loose pull mass cut of 5.0σ ensures that we have a sideband region for false η in our eventual 2D fit.

As in Section 8.2 with $D_s \to K_s e\nu$, we have made a cut on the mass of the missing four vector. This missing mass includes both the neutrino and the $D_s^* \to D_s\gamma$ transition photon. The soft transition photon's low energy ensures that the missing mass for properly reconstructed events stays closer to the neutrino's missing mass (zero) than incorrectly reconstructed events. This allows us to cut out most of the combinatoric background, the majority of fakes from other D_s semileptonic modes (e.g. $\eta' e\nu$, $\eta' \to \pi^0\pi^0\eta$), and a significant portion of misreconstructed η from true events, as seen in Figures 9.2 and 9.3.

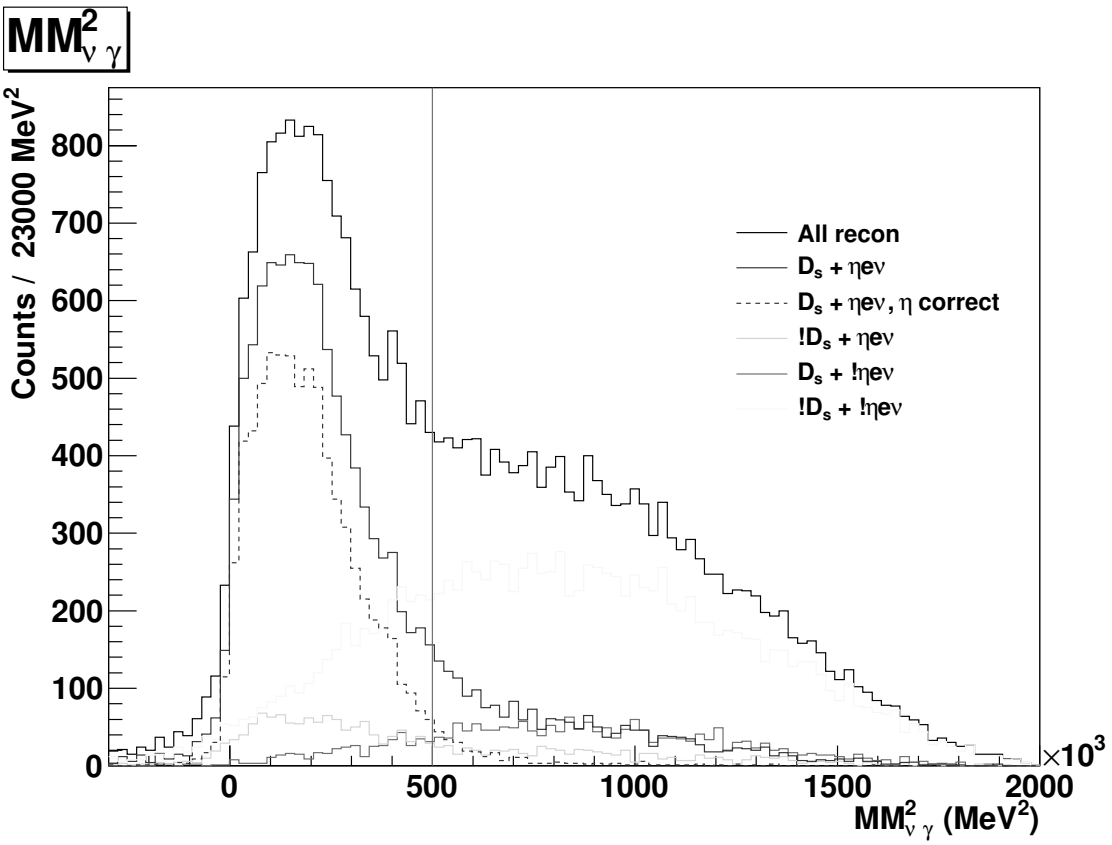

Figure 9.2: $D_s \to \eta e\nu$ missing mass after reconstructing a D_s tag, an η, and the electron in the 20× Monte Carlo. The solid $D_s + \eta e\nu$ line represents all generated $D_s \to \eta e\nu$ events that have a correct D_s tag, while the dotted $D_s + \eta e\nu$ line has the additional requirement that the η gets properly reconstructed from its daughter photons (no splitoff or transition photon fakes).

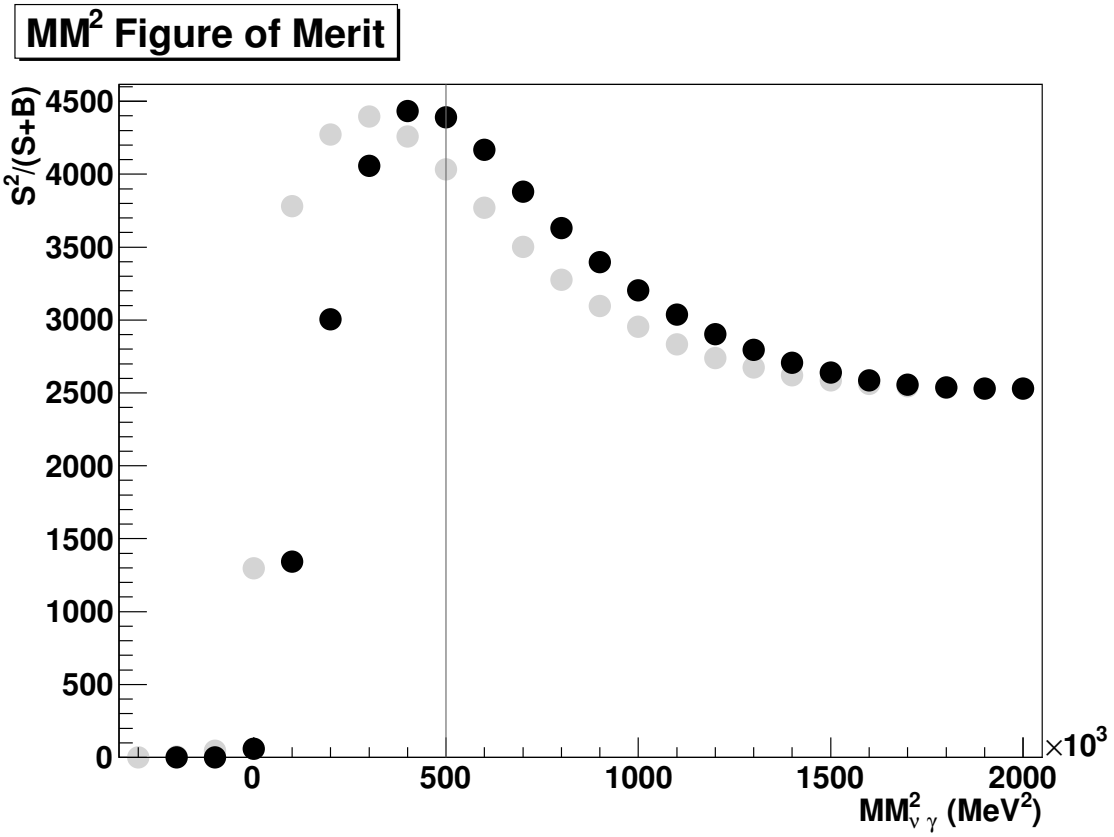

Figure 9.3: Our figure of merit for different missing mass cuts in the $D_s \to \eta e \nu$ signal region (within 3σ of a correct D_s mass and η pull mass). The dark dots correspond to the cuts from this analysis, in which we choose a maximum missing mass cut of 500×10^3 MeV^2 to err on the side of high efficiency. We have also tried reconstructing the best $D_s^* \to D_s \gamma$ transition photon when available and incorporating it into the missing four vector (light dots). However, we don't see an improvement in our figure of merit within the Monte Carlo by including the transition photon, and using it would expose us to potential problems from the modeling of splitoff showers.

Overall, we require the following cuts on the η beyond those for its daughter photons:

- $|\sigma_\eta| < 5.0$
- $MM^2_{\gamma\nu} < 500{,}000$ MeV^2 (.50 GeV^2)

9.2 $D_s \to \eta e\nu$ Reconstruction

9.2.1 Efficiency

As with $D_s \to \phi e\nu$ (Section 7.2.1), we generated four signal Monte Carlo samples of 250,000 events for each D_s production mode ($D_s^+ D_s^-$, $D_s^* D_s$ with prompt $D_s \to \eta e\nu$, $D_s^* D_s$ with $D_s^* \to D_s\gamma$ or $D_s^* \to D_s\pi^0$). We give our efficiencies for our wide two-dimensional fit region (1900 MeV $\leq M_{D_s} \leq$ 2030 MeV and $\left|\sigma_\eta\right| < 5.0$) in Table 9.1. For comparison, we also provide the efficiency with a tighter η selection ($\left|\sigma_\eta\right| < 3.0$) in Table 9.2. The semileptonic efficiencies include our event missing mass cut, so they're slightly smaller than the simple product of the electron and η efficiencies.

Table 9.1: Efficiencies for semileptonic particles in $D_s \to \eta e\nu$, with the η cuts used in this analysis. The η and semileptonic efficiencies include the $\eta \to \gamma\gamma$ branching ratio.

D_s production mode	ϵ_e	ϵ_η	ϵ_{SL}
$D_s D_s$	81.7% ± 2.7%	26.8% ± 1.6%	20.7% ± 1.4%
$D_s^* D_s$ with $D_s^* \to (D_s \to \eta e\nu)\ \gamma$	79.9% ± 0.7%	26.4% ± 0.4%	20.7% ± 0.4%
$D_s^* D_s$ with $D_s^* \to (D_s \to \eta e\nu)\ \pi^0$	80.7% ± 0.7%	26.5% ± 0.4%	20.7% ± 0.4%
$D_s^* D_s$ with prompt $D_s \to \phi e\nu$	80.4% ± 0.7%	26.7% ± 0.5%	20.5% ± 0.4%
Weighted signal MC	80.2% ± 0.5%	26.6% ± 0.3%	20.6% ± 0.3%
Generic MC	80.3% ± 0.5%	26.5% ± 0.3%	20.4% ± 0.2%

We have a higher efficiency for electron detection in $\eta e\nu$ than we do in $\phi e\nu$ or most of our other semileptonic modes. This improvement comes from the lower mass of the η, leading to a higher Q value in $D_s \to \eta e\nu$ and fewer of the low efficiency, slow electrons (Figure 9.4). We also find that the $\eta \to \gamma\gamma$ efficiency does not have a strong dependence on the η momentum (Figure 9.5), in contrast to our other semileptonic modes where slower hadrons have difficult-to-reconstruct charged tracks.

Table 9.2: Efficiencies for semileptonic particles in $D_s \to \eta e\nu$, with $|\sigma_\eta| < 3.0$. The efficiencies include the $\eta \to \gamma\gamma$ branching ratio.

D_s production mode	ϵ_e	ϵ_η	ϵ_{SL}
D_sD_s	81.7% ± 2.7%	25.4% ± 1.5%	19.8% ± 1.4%
$D_s^*D_s$ with $D_s^* \to (D_s \to \eta e\nu)\ \gamma$	79.9% ± 0.7%	24.5% ± 0.4%	19.3% ± 0.3%
$D_s^*D_s$ with $D_s^* \to (D_s \to \eta e\nu)\ \pi^0$	80.7% ± 0.7%	24.6% ± 0.4%	19.3% ± 0.3%
$D_s^*D_s$ with prompt $D_s \to \phi e\nu$	80.4% ± 0.7%	24.9% ± 0.4%	19.2% ± 0.4%
Weighted signal MC	80.2% ± 0.5%	24.7% ± 0.3%	19.3% ± 0.3%
Generic MC	80.3% ± 0.5%	24.7% ± 0.3%	19.1% ± 0.2%

Previous CLEO η studies [71] suggest a correction to the relative η efficiency of -5.6% with a relative systematic of 5.9%. We have done our own systematic (discussed in Section 10.6.2 and Appendix C) that does not show a needed correction, although we get a large, relative systematic error of 7.9%. We use this systematic in our final results, as it comes from a run environment that more closely matches our own. For completeness, however, we also include the $D_s \to \eta e\nu$ branching ratio when using the corrected efficiency and smaller η efficiency systematic (final results in Table 9.11).

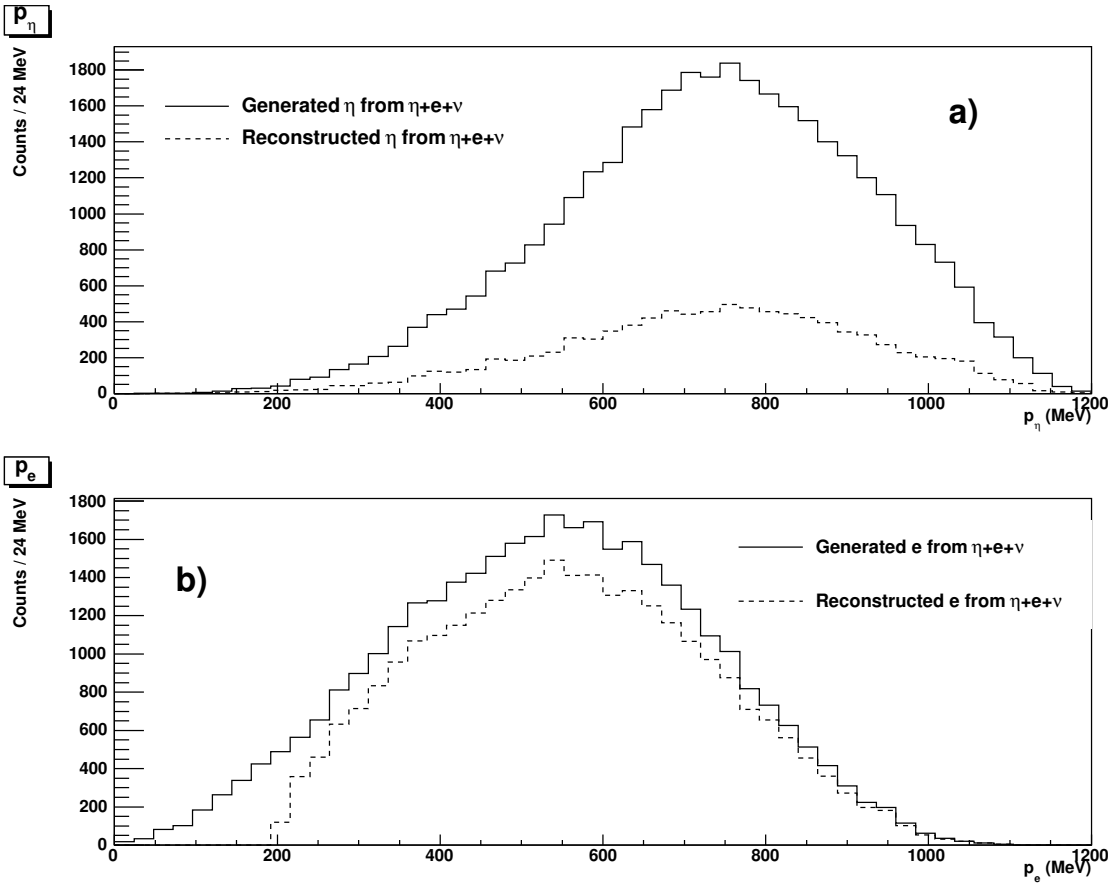

Figure 9.4: Generated and reconstructed momentum spectra from the 20× Monte Carlo. a) Lab frame η momentum from $D_s \to \eta e\nu$. b) Electron momentum in $D_s \to \eta e\nu$.

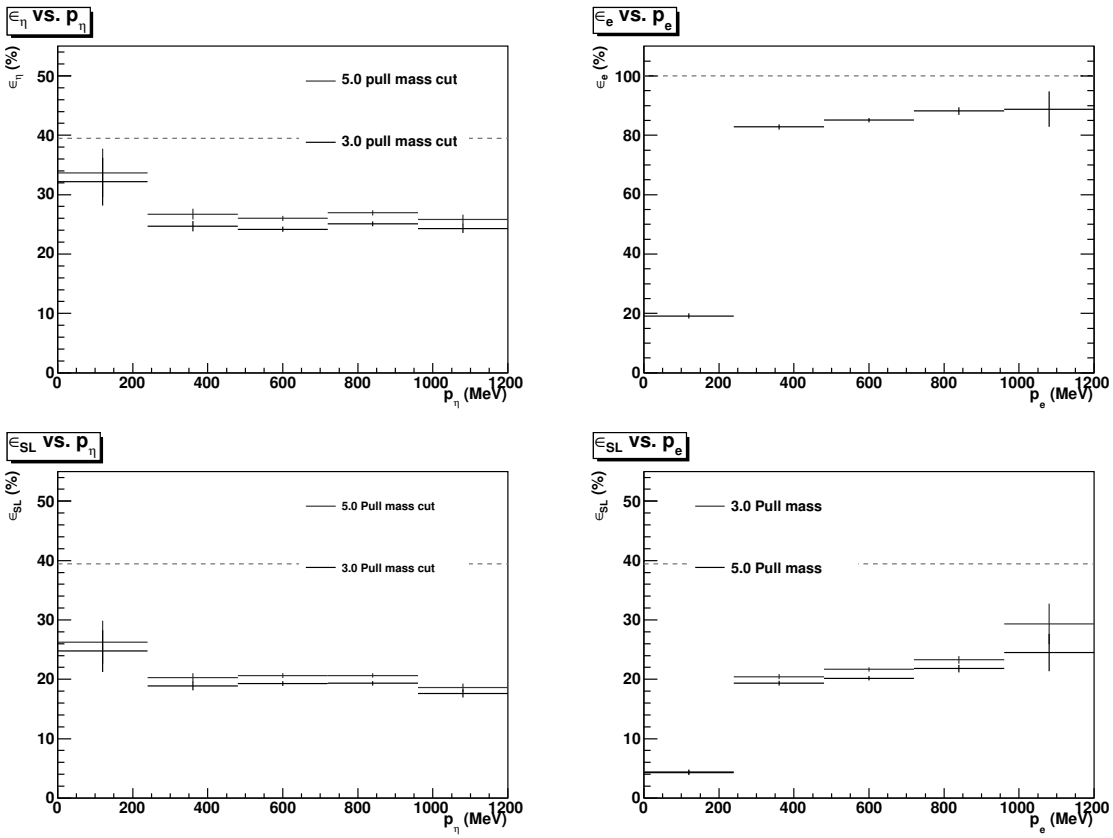

Figure 9.5: Efficiencies for η, electron, and the overall semileptonic side (η+electron, with MM^2 cut), by momentum. Our η and semileptonic efficiencies include the $\eta \to \gamma\gamma$ branching ratio of 39.5%.

9.2.2 Backgrounds

Our event missing mass cut eliminates most of our background combinations from events other than $D_s \to \eta e\nu$ or with incorrect D_s tags (Table 9.3). We can fit the remaining background combinations fairly well with either a linear background function in the η pull mass, a linear background in the D_s mass, or both, as described in Section 9.2.3.

Table 9.3: Truth-tagged breakdown for $D_s \to \eta e\nu$ candidates in the 20× `ddmix` and 5× `continuum` Monte Carlo, scaled to data size.

Event truth	Candidates before cuts	Candidates after cuts
True D_s tag, true $D_s \to \eta e\nu$	604.55	499.30
True D_s tag, not $D_s \to \eta e\nu$	109.60	22.20
False D_s tag, true $D_s \to \eta e\nu$	90.80	59.05
False D_s tag, not $D_s \to \eta e\nu$	787.75	198.70

We still have a problem with combinations from true $D_s \to \eta e\nu$ events where the tagged D_s and electron get properly constructed, but where we have an improper η reconstruction ("volunteer" events). The η can be formed by using a splitoff shower from one of the tagged D_s tracks or by using the D_s^* daughter transition photon (Figure 9.6). These misreconstructed η combinations make up 27% of all true D_s tag, true $D_s \to \eta e\nu$ combinations in the Monte Carlo, which can be seen in the difference between the solid and dotted $D_s + \eta e\nu$ lines in Figure 9.2.

We attempted a best candidate selection on the η, such that each true event only gets counted once whether it has a volunteer combination or not. However, this shapes our other false η backgrounds away from a simple linear distribution, and it still requires us to estimate how many volunteer-only events we have for our efficiency systematic. While we expect the Monte Carlo to model the volunteer η combinations from a real η photon paired with the transition photon fairly well, we have found instances (e.g. $D_s \to \phi e\nu$) where the Monte Carlo underestimated the number of splitoff showers. Simply rejecting all splitoff showers costs us too many true signal events, so we ultimately ran a separate systematic correction to account for false η pairings with splitoff events.

For our splitoff systematic, we take advantage of CLEO's large sample of D^0 and $D^\pm$ events from $\psi(2S) \to DD$ events at 3770 MeV. In particular, we use $D^0 \to K^*\eta$

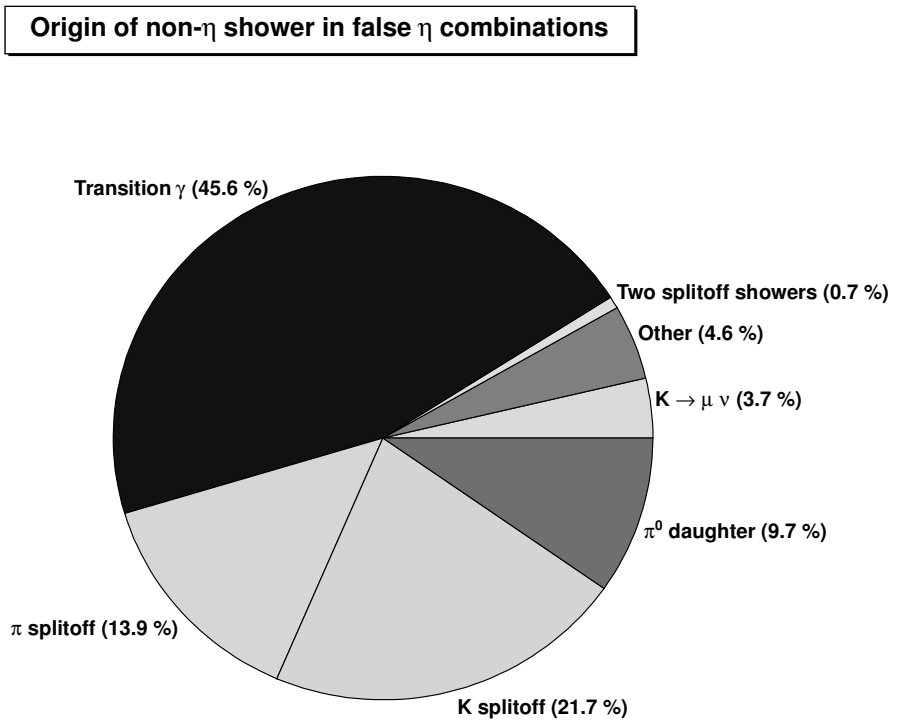

Figure 9.6: Cause of the shower that leads to a false η combination when we have a correct D_s in an $D_s \to \eta e\nu$ event, from our 20$\times$ Monte Carlo. These false combinations account for 27% of our counts in $D_s \to \eta e\nu$ events with a valid D_s tag. Our systematic addresses possible modeling flaws with the data for the three large, light slices (π splitoff, K splitoff, and $K \to \mu\nu$).

decays, which gives us a fairly pure η sample after we do the full event reconstruction. From this sample, we see how often we get an extra η combination in events with a correct reconstruction, which tells us the false η rate from splitoff showers (or similar causes, like $K \to \mu\nu$ decays from the tag side). We then scale the number of anticipated splitoff showers from each D_s mode by the rate of extra splitoffs that we observed in the data from similar D^0 modes.

We use four different D^0 tags ($K\pi$, $K\pi\pi^0$, $K_s\pi\pi$, $K\pi\pi\pi$), which we link with our 13 D_s tag modes (Table 9.4) to model the possible splitoff opportunities in $D_s \to \eta e\nu$ events. We cut fairly harshly on the D^0 to get as clean a sample as possible, requiring the beam constrained mass to be within 5 MeV of the D^0 mass and the ΔE within 20 MeV

of zero. We then reconstruct $K^* \to K\pi$ from the other tracks, requiring consistent charges with the D^0 tag, good tracks (within 5 cm of z_0 and 5 mm of d_b), and that particle ID matches the K or π track ($\mathcal{L} = (\sigma_\pi^2 - \sigma_K^2) + (L_\pi - L_K)$ less than zero for π, greater than zero for K). We also require that the reconstructed K^* mass be within 35 MeV of its PDG mass.

Table 9.4: D^0 tag modes used to estimate splitoff systematic for D_s modes.

D^0 mode	Corresponding D_s modes
$K\pi$	K_sK
	$KK\pi$
$K\pi\pi^0$	$K_sK\pi^0$
	$KK\pi\pi^0$
	$\pi\eta$
	$\pi\pi^0\eta$
	$\pi\eta', \eta' \to \pi\pi\eta$
	$\pi\pi^0\eta', \eta' \to \pi\pi\eta$
	$\pi\eta', \eta' \to \rho\gamma$
$K_s\pi\pi$	Cross-check only
$K\pi\pi\pi$	$K_sK_s\pi$
	$K_sK^+\pi\pi$
	$K_sK^-\pi\pi$
	$\pi\pi\pi$

Once we have a reconstructed D^0 tag and a K^*, we ensure that we have a $D^0 \to K^*\eta$ event by requiring that the recoil of the event fall near the η mass (525 MeV to 600 MeV), as seen in Figures 9.7 and 9.8. We then reconstruct $\eta \to \gamma\gamma$ with a 5σ pull mass cut and ensure that we have only properly reconstructed η by cutting tightly ($\pm$10 MeV) on the event's missing mass with the η included (Figures 9.9 and 9.10).

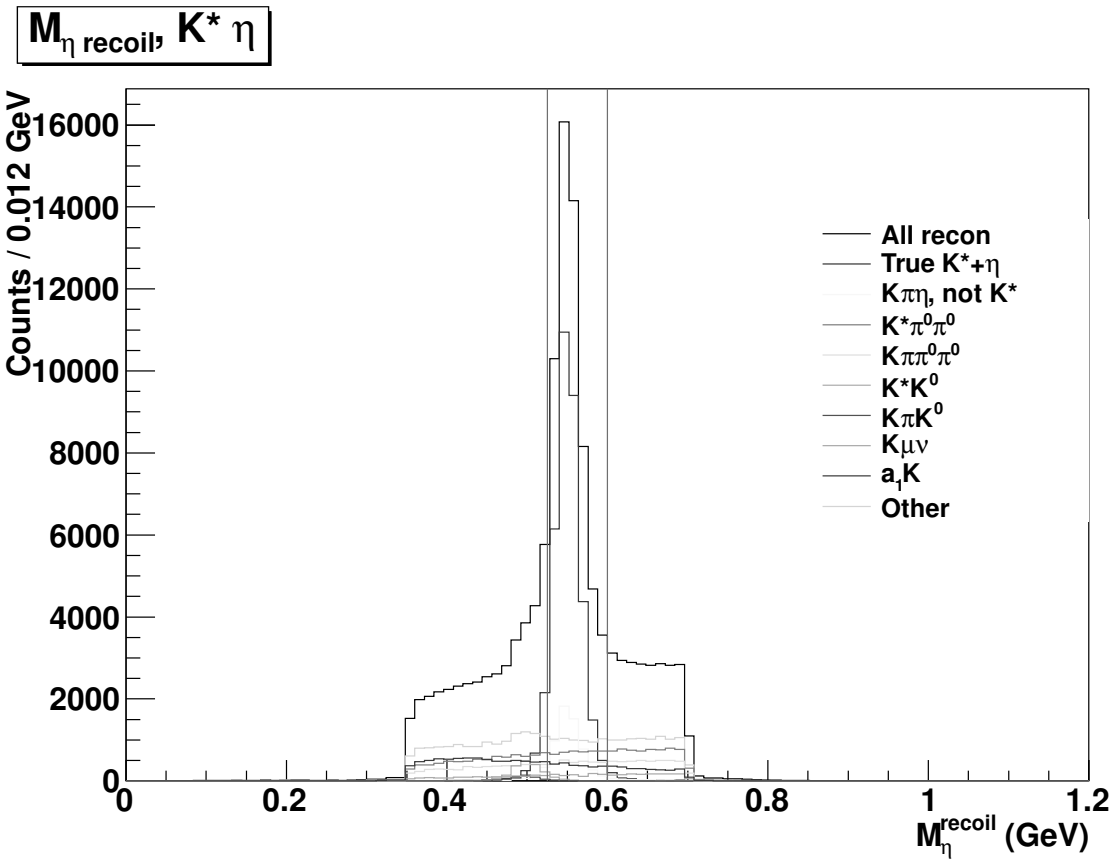

Figure 9.7: Recoil mass against the $D^0 + K^*$ in a 20× Monte Carlo sample. We keep events between the vertical lines.

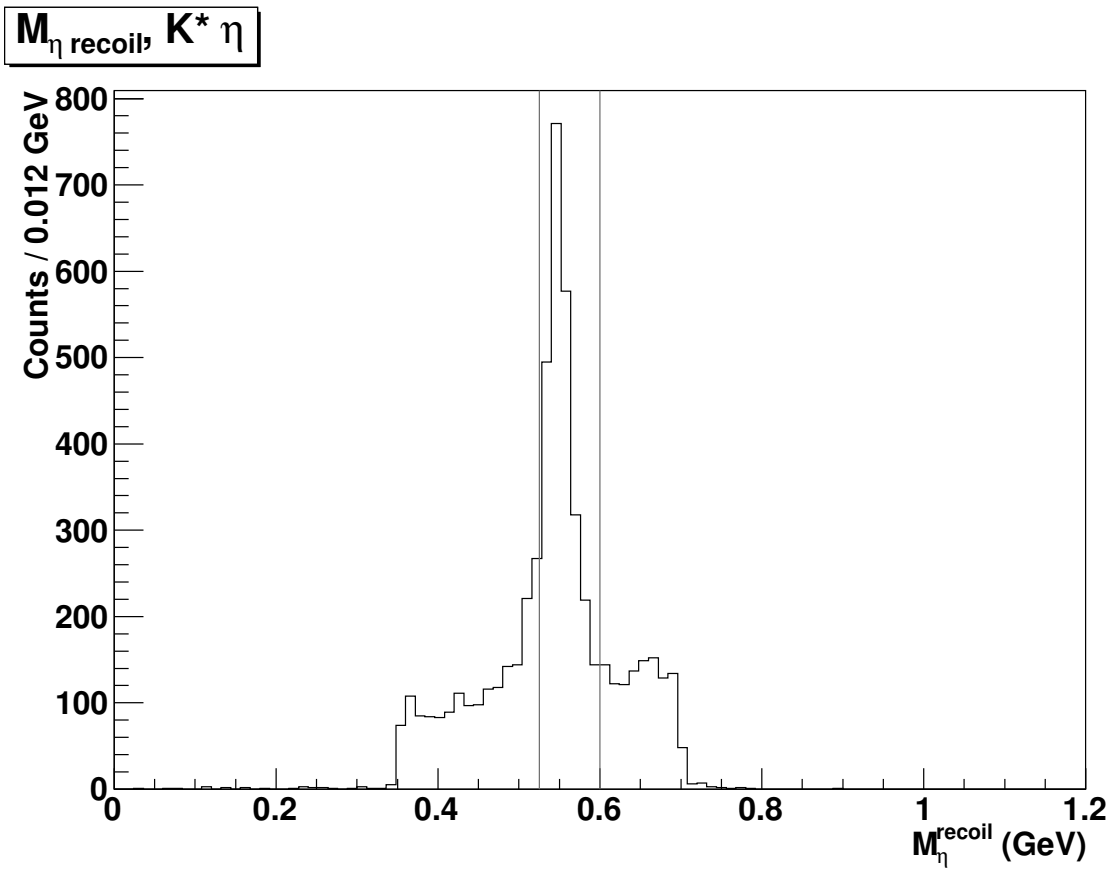

Figure 9.8: Recoil mass against the $D^0 + K^*$ from the 3770 MeV data.

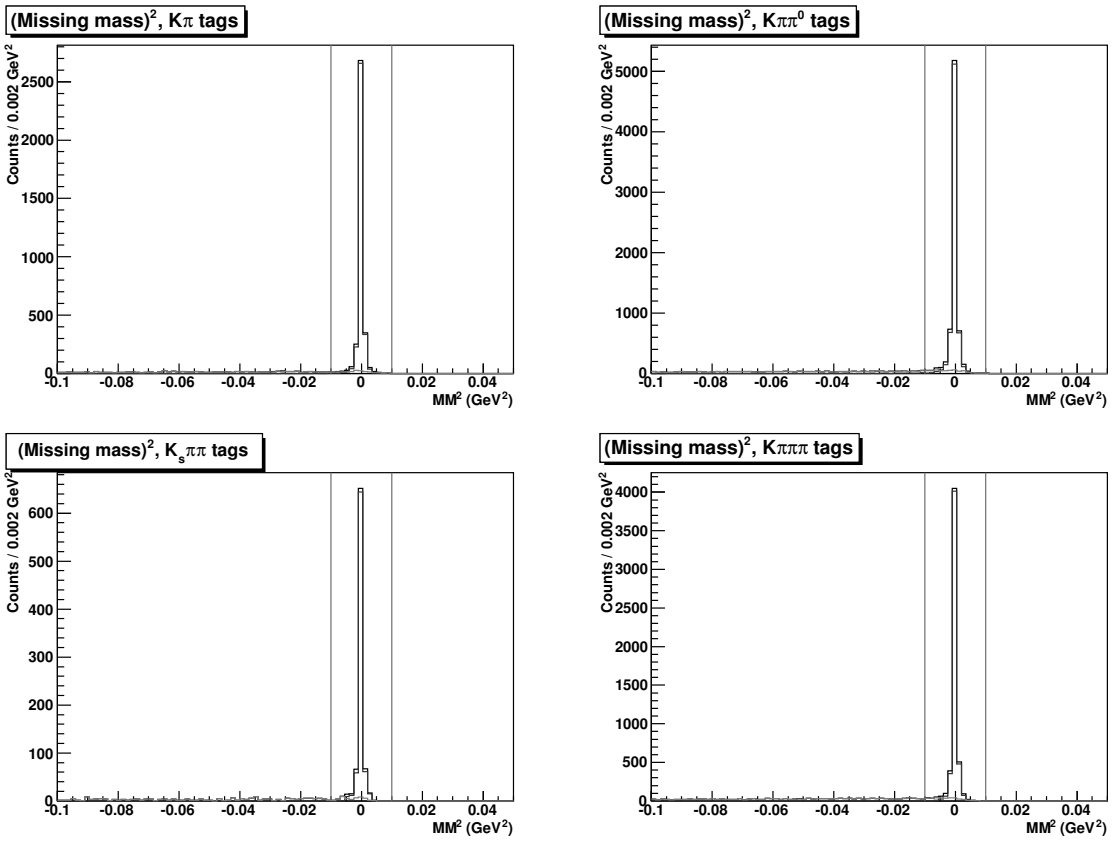

Figure 9.9: Full event missing mass, by D^0 tag mode for $D^0 + K^*\eta$ in a 20× Monte Carlo sample. We select η combinations between the vertical lines so that we know we have only true η, then we see if there are any other splitoff η combinations in the event.

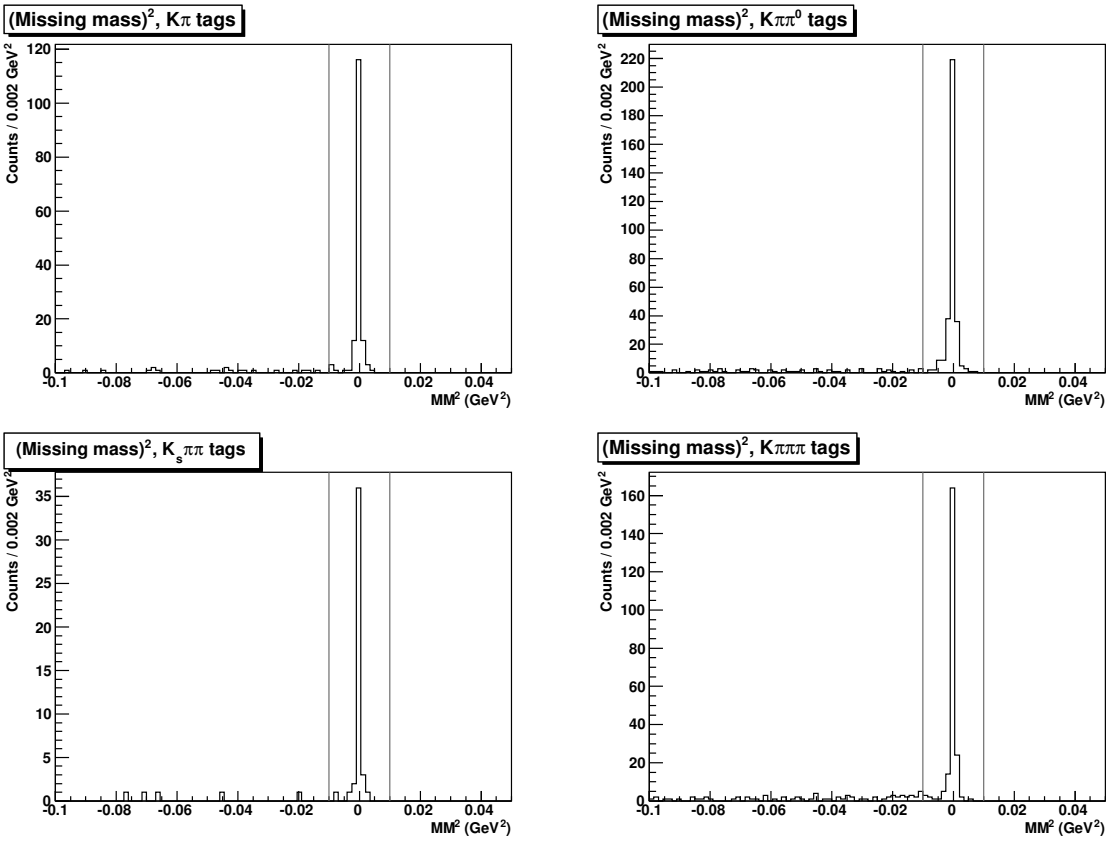

Figure 9.10: Full event missing mass, by D^0 tag mode for $D^0 + K^*\eta$ in the 3770 MeV data. We select η combinations between the vertical lines to determine true η, then we see if there are any other splitoff η combinations in the event.

When we have identified events with a correct η reconstruction, we see how many proper events have additional η from splitoff showers after dropping the missing mass requirement. By comparing the extra η from splitoff in Monte Carlo and in data, we can determine the difference in splitoff rate (Table 9.5). We don't see a large needed splitoff correction for η reconstruction with our cuts, mostly because the photons involved have larger energies where the Monte Carlo models splitoff better. Since our $D_s \to \eta e\nu$ fits still depend somewhat on the Monte Carlo's splitoff rate, we incorporate these results in a systematic, as described in Section 10.8.

Table 9.5: Rate of additional η formed using splitoff showers, by D^0 mode. The data/MC splitoff correction error (extra splitoff factor) includes a small systematic from combinatoric background.

Mode	Recon Data %	Recon MC %	Truth-tagged MC %	Extra splitoff factor
$K\pi$	(25.9 ± 3.6)%	(27.0 ± 0.8)%	27.1%	0.959 ± 0.137
$K\pi\pi^0$	(27.4 ± 2.5)%	(29.9 ± 0.5)%	30.0%	0.915 ± 0.085
$K_s\pi\pi$	(20.9 ± 6.2)%	(34.7 ± 1.6)%	34.6%	0.603 ± 0.181
$K\pi\pi\pi$	(39.2 ± 3.4)%	(38.4 ± 0.7)%	38.7%	1.020 ± 0.089
Integrated	(30.1 ± 1.7)%	(32.2 ± 0.4)%	32.3%	0.936 ± 0.054

Foregoing any cuts on $\frac{E_9}{E_{25}}$ increases our signal at the expense of some extra splitoff background. However, the Monte Carlo models the extra rate of background about as well with no $\frac{E_9}{E_{25}}$ cut as when we do apply an $\frac{E_9}{E_{25}}$ O.K. cut, shown in Table 9.6. For our full analysis, we simply drop $\frac{E_9}{E_{25}}$ to maximize our signal and use our splitoff systematic results to correct for the difference in splitoffs between the data and the Monte Carlo.

We also have a small, peaking background from events with a true D_s and a correct η that aren't $D_s \to \eta e\nu$ events. These events come from $D_s \to \phi e\nu$ where $\phi \to \eta\gamma$, $D_s \to \eta' e\nu$ where the η' decays to a state with no tracks and an η, and events where a kaon fakes the electron. Most such peaking background get rejected by our event missing mass cut, but we use our measured $D_s \to \phi e\nu$ and $D_s \to \eta' e\nu$ branching factions to correct the remainder. We give the resultant correction to our $D_s \to \eta e\nu$ branching ratio from these corrections in Table 9.7.

Table 9.6: Rate of additional η formed with splitoff showers after applying an $\frac{E_9}{E_{25}}$ cut, by D^0 mode. The data/MC splitoff correction error (extra splitoff factor) includes a small systematic from combinatoric background.

Mode	Recon Data %	Recon MC %	Truth-tagged MC %	Extra splitoff factor
$K\pi$	(10.9 ± 2.8)%	(10.6 ± 0.6)%	10.7%	1.033 ± 0.266
$K\pi\pi^0$	(13.3 ± 2.1)%	(12.5 ± 0.4)%	12.5%	1.064 ± 0.169
$K_s\pi\pi$	(16.7 ± 6.2)%	(13.5 ± 1.3)%	13.8%	1.239 ± 0.476
$K\pi\pi\pi$	(20.2 ± 3.0)%	(16.4 ± 0.5)%	16.5%	1.236 ± 0.186
Integrated	(15.1 ± 1.4)%	(13.4 ± 0.3)%	13.4%	1.127 ± 0.110

Table 9.7: Absolute branching ratio correction and systematic error for $\mathcal{B}(D_s \to \eta e\nu)$ from peaking background.

Background mode	BR correction
$D_s \to \phi e\nu$	(0.0037 ± 0.0011)%
$D_s \to \eta' e\nu$	(0.0104 ± 0.0029)%
Non-semileptonic	(0.0017 ± 0.0008)%
Total $D_s \to \eta e\nu$ correction	(0.0158 ± 0.0032)%

9.2.3 Fit Procedure

In our other semileptonic modes, we had low background on the semileptonic side. This let us fit the D_s invariant mass for both our tags and semileptonic events without worrying about the specific reconstruction of the electron or hadron. $D_s \to \eta e\nu$ has somewhat more background on the semileptonic side due to the relative ease of making an η. We solve this problem by directly incorporating the η into our fit, doing a 2D fit to the η pull mass and the tagged D_s mass after getting a good electron.

We find no correlation between the η pull mass and the D_s mass (as expected), so we can use a simple product of the two distributions for our fit functions. This allows us to reapply the D_s mass signal lineshape that we determined previously from our tag fits for each D_s mode. We then generate a truth tagged η lineshape with our cuts in the Monte Carlo, which we take to be the η signal distribution.

We use a linear background for combinations with a false η or a false D_s. We

determine a normalized slope for the η background and for the background in each of the D_s modes by fitting their 1D background projections in the Monte Carlo.

As in $D_s \to \phi e\nu$, we do an unbinned, extended maximum likelihood fit over each D_s tag mode, with the fits linked by a common branching ratio parameter. However, instead of one background normalization for each tag mode due to false D_s, we now have the possibility of a false η, a false D_s, or both.

We don't require an extra parameter for combinations with a false η and true D_s as we determine their rate by one of two methods. The larger component of false η, true D_s combinations comes from volunteer (true) events, where the reconstructed η contains a splitoff shower or the D_s^* transition photon. We use our splitoff study to estimate the splitoff combination rate, and we use the Monte Carlo for the transition photon combination rate (which only involves kinematics). We then tie both to the rate for correctly reconstructed, true events (the branching ratio).

The smaller component of false η, true D_s combinations come from events that do not have a $D_s \to \eta e\nu$. The majority of these combinations arise from other D_s semileptonic modes, where either splitoff showers combine with a real η shower or where the semileptonic hadron decays to multiple π^0. We determine the rate of such combinations from the Monte Carlo and correct that rate using our measured branching ratio for each of those D_s semileptonic modes. We have a small remaining component (half an event) expected from all other sources; we use the Monte Carlo rate for such combinations.

This leaves 13 parameters (one per tag mode) for false D_s combinations with a true η. Similarly, we have another 13 parameters for combinatoric background (false D_s, false η), one per tag mode. We have a total of 1 signal parameter (the common branching ratio) and 26 background parameters in our fit.

Our log likelihood function to be minimized is given by

$$\begin{aligned}\mathcal{F}(\mathcal{B},\vec{N}_{\rm BG},\vec{K}_{\rm BG}) &\equiv -\ln\mathcal{L}(\mathcal{B},\vec{N}_{\rm BG},\vec{K}_{\rm BG}) \\ &= \mathcal{B}\cdot\varepsilon_{\rm SL}\cdot\left[\sum_{i=1}^{13}N^i_{\rm tag}\cdot(1+r^i)\right] \\ &\quad+\epsilon_{!\eta e\nu}\cdot\left[\sum_{i=1}^{13}N^{[j]}_{\rm tag}\right] \\ &\quad+\left[\sum_{i=1}^{13}N^i_{\rm BG}\right]+\left[\sum_{i=1}^{13}K^i_{\rm BG}\right] \\ &\quad-\sum_{m_j}\ln F(m_j,\sigma_\eta^{[j]}),\end{aligned}$$

where

$$\begin{aligned}F(m_j,\sigma_\eta^{[j]}) \equiv\ & \mathcal{B}\cdot\varepsilon_{\rm SL}\cdot N^{[j]}_{\rm tag}\cdot f^{[j]}_{\rm sig}(m_j)\cdot g_{\rm sig}(\sigma_\eta^{[j]}) \\ &+\mathcal{B}\cdot\varepsilon_{\rm SL}\cdot N^{[j]}_{\rm tag}\cdot f^{[j]}_{\rm sig}(m_j)\cdot g_{\rm BG}(\sigma_\eta^{[j]})\cdot r^{[j]} \\ &+\epsilon_{!\eta e\nu}\cdot N^{[j]}_{\rm tag}\cdot f^{[j]}_{\rm sig}(m_j)\cdot g_{\rm BG}(\sigma_\eta^{[j]}) \\ &+N^{[j]}_{BG}\cdot f^{[j]}_{\rm BG}(m_j)\cdot g_{\rm sig}(\sigma_\eta^{[j]}) \\ &+K^{[j]}_{BG}\cdot f^{[j]}_{\rm BG}(m_j)\cdot g_{\rm BG}(\sigma_\eta^{[j]}).\end{aligned}$$

Here, $\mathcal{B}$ is the $D_s\to\eta e\nu$ branching ratio, $\varepsilon_{\rm SL}$ is our $\eta e\nu$ efficiency, r^i is the total rate of false η from splitoff and transition photon combinations for D_s tag mode i, $\epsilon_{!\eta e\nu}$ is the total rate for false η combinations from all non-$\eta e\nu$ sources, $\vec{N}_{\rm BG}$ is the background normalization for false D_s/true η events (one per tag mode), and $\vec{K}_{\rm BG}$ is the combinatoric background. $[j]$ refers to the tag mode associated with the D_s mass, m_j. Our distributions are given by $f^i_{\rm sig}$ for the normalized signal mass distribution of true D_s with tag mode i, $f^i_{\rm BG}$ for the normalized, linear background function from false D_s of tag mode i, $g_{\rm sig}$ for the normalized pull mass distribution from true η, and $g_{\rm BG}$ for the normalized, linear background function from false η.

In $\mathcal{F}$, the first term corresponds to the number of signal $D_s\to\eta e\nu$ events. The second term (with r^i) gives the extra false η combinations formed from true events, while the third term gives false η combinations from other sources. The fourth and fifth

terms combine to form the total background from false D_s tags. The last, log term gives the sum over the distributions for each of the different signal and background sources, explicitly stated in the definition of $F(m_j, \sigma_\eta^{[j]})$.

9.3 Results

Unlike the other semileptonic modes where our statistical errors dominate, our $D_s \to \eta e\nu$ measurement has comparable levels of statistical and systematic error. The difficulty of obtaining a clean and comparable sample for the η efficiency drives the systematic error, so we have included that systematic (described with more detail in Section 10.6.2) in the following results. Section 11.2 contains the final result with our additional, less dominant systematic errors from Section 10.

9.3.1 Monte Carlo

We first break our 20× Monte Carlo into twenty data-sized samples to test our analysis technique with a limited statistics data set. Our comparison across these data-sized samples, given in Tables 9.8 and 9.9, show that our analysis reproduces the generating $D_s \to \eta e\nu$ branching ratio and number of signal events to within statistical error (χ^2 of 26.9 over 20 samples).

Table 9.8: Monte Carlo comparison of the measured $D_s \to \eta e\nu$ branching ratio to its generating branching ratio (2.480%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios.

Dataset	Raw fit BR	Corrected BR	# σ
1	$(2.712 \pm 0.132)\%$	$(2.692 \pm 0.132)\%$	1.60
2	$(2.423 \pm 0.124)\%$	$(2.404 \pm 0.124)\%$	−0.62
3	$(2.251 \pm 0.119)\%$	$(2.231 \pm 0.119)\%$	−2.09
4	$(2.602 \pm 0.128)\%$	$(2.583 \pm 0.128)\%$	0.80
5	$(2.313 \pm 0.122)\%$	$(2.293 \pm 0.122)\%$	−1.53
6	$(2.588 \pm 0.125)\%$	$(2.568 \pm 0.125)\%$	0.70
7	$(2.485 \pm 0.122)\%$	$(2.465 \pm 0.122)\%$	−0.12
8	$(2.492 \pm 0.124)\%$	$(2.473 \pm 0.124)\%$	−0.06
9	$(2.321 \pm 0.113)\%$	$(2.301 \pm 0.113)\%$	−1.58
10	$(2.263 \pm 0.121)\%$	$(2.243 \pm 0.121)\%$	−1.97
11	$(2.530 \pm 0.125)\%$	$(2.511 \pm 0.126)\%$	0.25
12	$(2.384 \pm 0.123)\%$	$(2.364 \pm 0.123)\%$	−0.94
13	$(2.511 \pm 0.123)\%$	$(2.491 \pm 0.123)\%$	0.09
14	$(2.480 \pm 0.124)\%$	$(2.460 \pm 0.124)\%$	−0.16
15	$(2.365 \pm 0.122)\%$	$(2.346 \pm 0.122)\%$	−1.10
16	$(2.758 \pm 0.131)\%$	$(2.738 \pm 0.131)\%$	1.98
17	$(2.539 \pm 0.123)\%$	$(2.519 \pm 0.123)\%$	0.32
18	$(2.354 \pm 0.121)\%$	$(2.335 \pm 0.121)\%$	−1.20
19	$(2.343 \pm 0.120)\%$	$(2.323 \pm 0.120)\%$	−1.31
20	$(2.569 \pm 0.125)\%$	$(2.549 \pm 0.125)\%$	0.55
Weighted averages/χ^2	$(2.457 \pm 0.028)\%$	$(2.437 \pm 0.028)\%$	26.87

Table 9.9: Test of potential bias in our fitting procedure for $D_s \to \eta e \nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	369	382.00 ± 22.30	0.583
2	369	354.60 ± 21.61	−0.666
3	350	334.04 ± 21.13	−0.755
4	397	382.58 ± 22.38	−0.644
5	343	334.00 ± 21.03	−0.428
6	373	372.55 ± 21.62	−0.021
7	367	364.06 ± 21.34	−0.138
8	348	368.40 ± 21.82	0.935
9	324	333.11 ± 19.36	0.471
10	369	332.66 ± 21.24	−1.711
11	369	361.51 ± 21.47	−0.349
12	348	346.87 ± 21.35	−0.053
13	374	371.25 ± 21.68	−0.127
14	380	372.84 ± 22.19	−0.323
15	361	339.79 ± 20.99	−1.011
16	370	392.65 ± 22.23	1.019
17	390	371.83 ± 21.54	−0.844
18	350	346.62 ± 21.28	−0.159
19	348	349.24 ± 21.32	0.058
20	377	373.66 ± 21.68	−0.154
Sum	7276	7184.27 ± 96.09	−0.955

Figure 9.11 shows the 1D projections of our 2D fit to the D_s invariant mass and η pull mass in the Monte Carlo, after summing over all twenty datasets and each tag mode. We have also added the 1D projections of our Monte Carlo results for the four highest statistics tag modes (summed over all twenty datasets) and for the first four datasets (summed over all tag modes) in Figures G.48–G.52 from Appendix G.

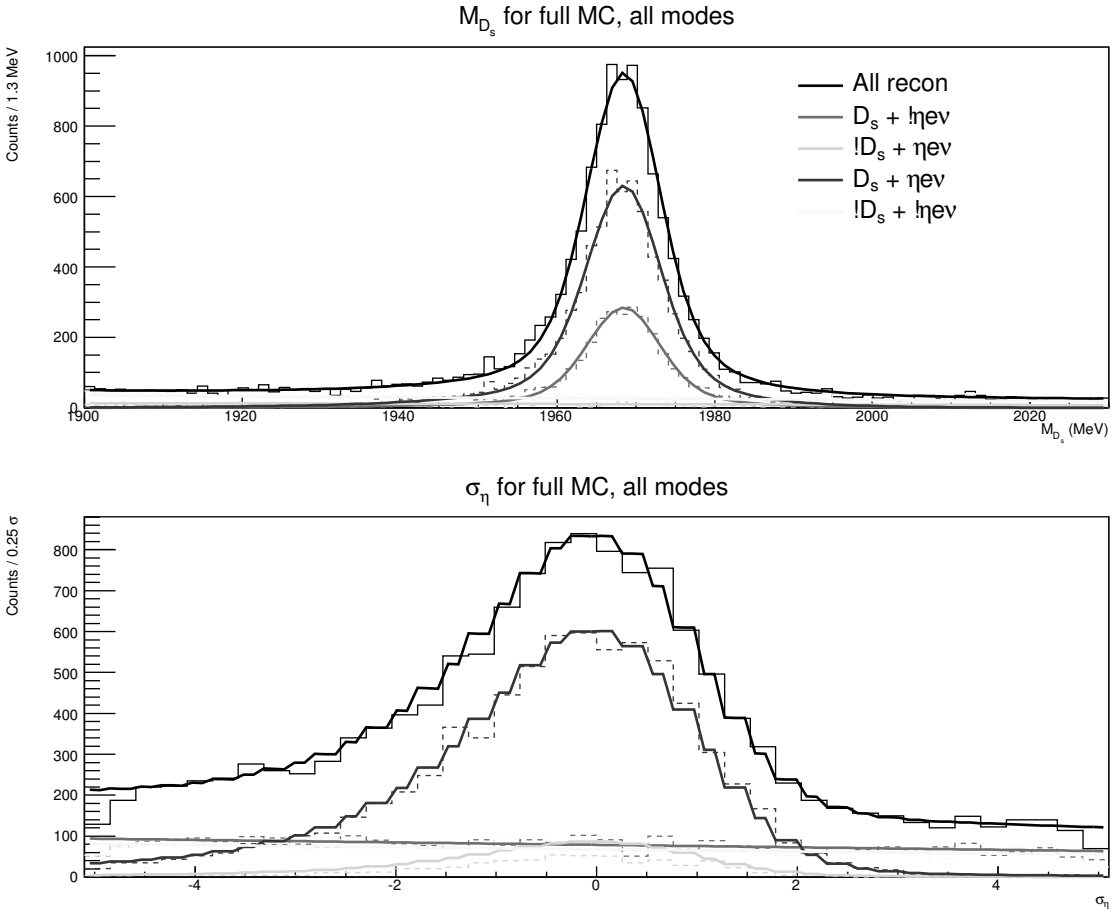

Figure 9.11: $D_s \to \eta e\nu$ 2D fit projections for the reconstructed D_s mass (top) and η pull mass (bottom) in the 20× Monte Carlo, summing over all tag modes.

9.3.2 Data

Table 9.10 contains our $D_s \to \eta e\nu$ branching ratio with statistical error after we correct for true D_s, true η feedthrough from other semileptonic modes. Figure 9.12 shows the projections of our reconstruction and fit along the D_s invariant mass and the η pull mass.

Table 9.10: $D_s \to \eta e\nu$ measurement in the data, including the peaking background correction from Table 9.7.

Measurement	Branching Ratio	# Events
Raw fit result	$(2.265 \pm 0.136)\%$	360.7 ± 21.9
Peaking BG correction	$(0.016 \pm 0.003)\%$	2.5 ± 0.5
$\mathcal{B}(D_s \to \eta e\nu)$	$(2.249 \pm 0.136)\%$	358.2 ± 21.9

As the η efficiency systematic (Section 10.6.2) dominates our error for $D_s \to \eta e\nu$ (a relative 7.9% systematic versus a relative 6.0% statistical error), we show the branching ratio with just that systematic error added in Table 9.11. For comparison, we've also included the $D_s \to \eta e\nu$ result when using the standard CLEO η efficiency systematic [71]. That analysis uses different η selections and has a cleaner environment ($\psi' \to \eta J/\psi$), extrapolating the systematic on their monoenergetic η from a π^0 study. They saw a relative systematic error of 5.9% with a relative efficiency correction of -5.6%. Although we feel that the standard systematic provides a viable alternative, we have chosen to use our own systematic for the final result, believing that we gain improved accuracy at the expense of finer precision.

Table 9.11: $D_s \to \eta e\nu$ branching ratio and errors under both η efficiency systematic scenarios.

η efficiency systematic	$D_s \to \eta e\nu$ branching ratio
This analysis	$(2.249 \pm 0.136 \pm 0.179)\%$
Alternate systematic	$(2.375 \pm 0.143 \pm 0.134)\%$

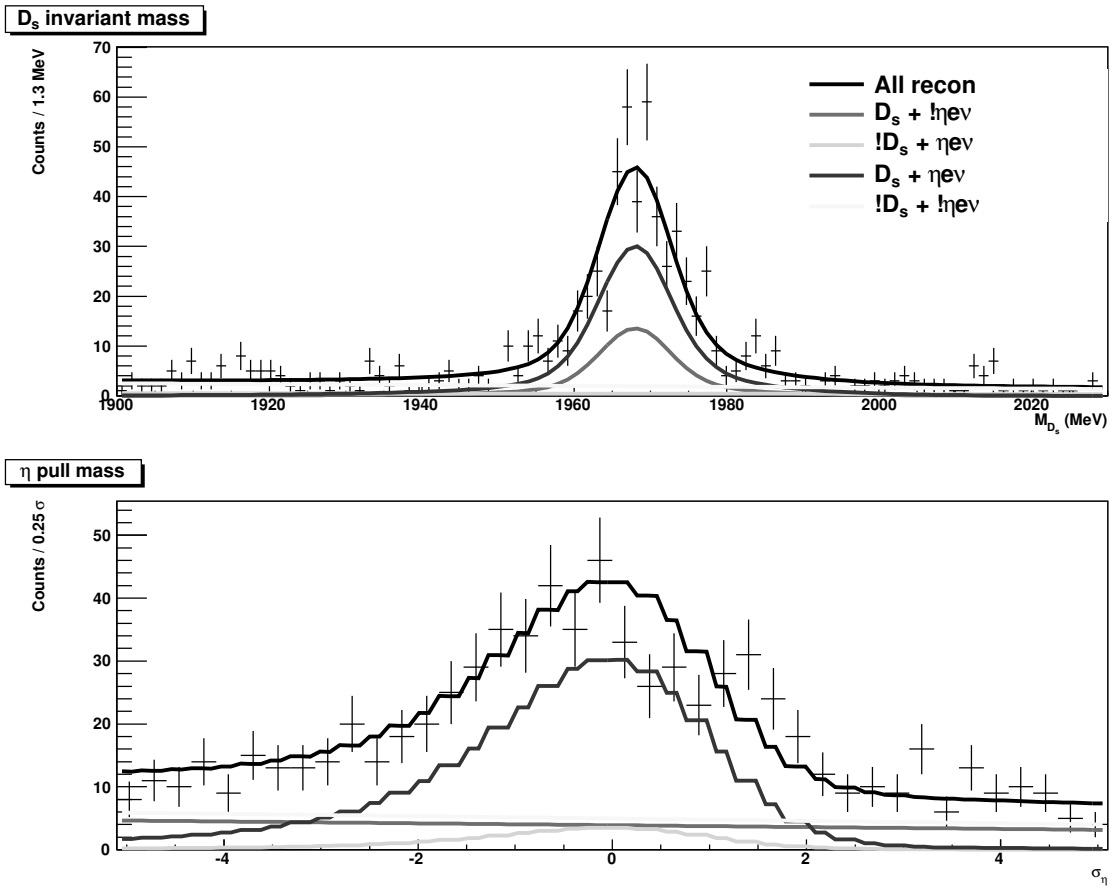

Figure 9.12: $D_s \to \eta e\nu$ 2D fit projections for the reconstructed D_s mass (top) and η pull mass (bottom) in the data, summed over all tag modes.

Chapter 10

Systematic Uncertainties

Our main limit on measurement precision in D_s semileptonic decays with our sample comes from the amount of available data. However, all six of our semileptonic modes have non-trivial errors of the "how well do we know what we know" variety, which we have determined and enumerated to the best of our abilities in this chapter.

10.1 D_s Tagging

We have a variety of possible systematic effects that could cause a mismatch between the Monte Carlo efficiency and the true efficiency for reconstructing D_s tags, like our recoil mass resolution or resonant substructure (e.g. $KK\pi$ can come from $\phi\pi$, K^*K, or be nonresonant). However, by using the D_s invariant mass for both the numerator and denominator of our branching ratios, the effect of any such errors in our tagging efficiency cancels out. We only have to worry about systematic effects that create different biases in our tag counting than in our semileptonic fits.

10.1.1 Signal Shape Variation

Our D_s tag fits have either a double gaussian or a gaussian plus crystal ball for their signal components, as mentioned in Section 5.4 (Table 5.4). We use this same shape for both the number of tags (our branching ratio denominator) and the number of semileptonic counts (our numerator), so we don't expect our result to depend strongly on minor errors with our lineshape. However, we do see far more background relative

to signal for our tag counts than for our semileptonic counts, so we could have some bias due to a wider than expected signal looking like background in our tag fits.

We allow our D_s mass signal shape's overall normalization, mean, and overall width to vary when we do the tag counts, so we should expect no bias between the data and Monte Carlo from those parameters. However, we use the sum of two shapes for our signal functions, and we fixed the relative normalization and relative width between them using our predicted (Monte Carlo) histograms. While leaving the overall width to float deals with most M_{D_s} resolution issues, we can tell a story about the poorer quality tracks having a worse than expected resolution while the high quality tracks match well, or one where we have more poor quality tracks than the Monte Carlo expects. The relative width or normalization, respectively, would then need to adjust to properly match the true signal shape in data.

Our tagged D_s backgrounds make it impractical to simply allow the relative width or normalization to float, so we need to look elsewhere to study any potential biases between the Monte Carlo and data. Since the $D^{\pm}$ has similar decay modes to the D_s (often with just a K to π conversion), we look at that system to study our tag signal shapes. While probably overkill, we wanted to keep the procedure as close to our D_s tagging as possible, so we use DD^* events at 4170 MeV instead of moving to the high data running at 3770 MeV. This costs us some precision (and generates much more work), but it allows us to use a similar choice of a best recoil mass to protect us from the (unlikely) possibility that a best choice somehow biases the track quality in a way not predicted from the Monte Carlo.

We use seven different $D^{\pm}$ tag modes, each of which corresponds to one or more D_s tag modes (Table 10.1). We reconstruct the $D^{\pm}$ tags with the same daughter particle cuts as listed in Section 5.1 for the D_s. Our best recoil mass selection for each charge now takes the $D^{\pm}$ with a recoil mass closest to the D^{*+} mass instead of the recoil mass closest to a D_s^*.

We fit each $M_{D^{\pm}}$ spectrum with our analogous D_s signal function (double gaussian or gaussian plus crystal ball), fixing the relative normalization and relative width to the Monte Carlo signal. Then, we fit the $M_{D^{\pm}}$ spectrum in data with the relative normalization or relative width allowed to float. We treat the ratio between the data's relative normalization (or width) and the Monte Carlo's relative normalization (or width) as

Table 10.1: $D^{\pm}$ tag mode used for each D_s mode's relative normalization and relative width systematics.

$D^{\pm}$ mode	Corresponding D_s modes
$K\pi\pi$	$\pi\pi\pi$
$K_s\pi\pi^0$	$K_sK\pi^0$ $\pi\eta', \eta' \to \pi\pi\eta$
$K\pi\pi\pi^0$	$KK\pi\pi^0$ $\pi\eta$
$K_s\pi$	$K_sK_s\pi$
K_sK	K_sK $K_sK^+\pi\pi$ $K_sK^-\pi\pi$
$KK\pi$	$KK\pi$
$\pi\pi^0\pi^0$	$\pi\pi^0\eta$ $\pi\pi^0\eta', \eta' \to \pi\pi\eta$ $\pi\eta', \eta' \to \rho\gamma$

our 1σ systematic variation.

Once we have the systematic variation on the relative normalizations or widths, we redo the D_s semileptonic analyses with our rescaled values. We show the systematic on our semileptonic branching ratios for each mode in Table 10.2.

Table 10.2: Systematic errors from our D_s tag fits.

Semileptonic mode	Tag statistics	Fit shape, $\frac{N_2}{N_1}$	Fit shape, $\frac{\sigma_2}{\sigma_1}$	Total systematic
$\phi e\nu$	1.0%	1.5%	0.3%	1.8%
$\eta e\nu$	1.0%	1.9%	0.7%	2.2%
$\eta' e\nu$	1.3%	1.6%	0.3%	2.1%
$f_0 e\nu$	1.0%	1.2%	0.3%	1.6%
$K_s e\nu$	0.9%	1.3%	1.5%	2.2%
$K^* e\nu$	1.2%	0.5%	2.7%	3.0%

Table 10.2 also gives the branching ratio systematic from the statistical uncertainty in our D_s tag count measurements. Since the tag counts for each mode provide an explicit weighting in our linked semileptonic fits, we can't simply add their errors in

quadrature with the numerator errors. Instead, we follow each statistical variation through the entire process and add the final results from each tag mode's fluctuations in quadrature. Whether this error should be considered statistical in nature because it arises from underlying statistics on the tagging or systematic in nature because it creates a bias in our numerator fits' weights is a matter of philosophy; we include it here as it has more in common with our tagging systematics than with our semileptonic measurements.

10.1.2 Background Functions

We approximate the combinatoric backgrounds on our D_s tag fits with either a linear or quadratic background function, depending on the tag mode (Table 5.4). This gives us some flexibility in case the data has slightly different combinatorics than predicted by the Monte Carlo, but it may give the background too much freedom to add or steal counts from our signal (particularly the quadratic backgrounds).

To estimate our systematic error from our choice of background function, we also fit using a one parameter background histogram with our signal function. Our background histogram includes both a charm and a continuum component, which we simply fix to the Monte Carlo expectation so that we only have one free background parameter.

The histogram background fits give us roughly the same tagging results as our normal procedure, with a relative difference in total D_s counts of less than 1%. However, since each mode's tag counts also act as a weighing function for our semileptonic fits, we follow the changes through the entire procedure to our branching ratio. After adding the branching ratio variation from each tag mode's background histogram fit in quadrature, we obtain the systematic errors shown in Table 10.3.

10.1.3 Multiple Candidate Choice

In our D_s tag selection, we make a best candidate choice based on the D_s recoil mass. When multiple candidates exist in the event, this selection can occasionally cause us to throw out the proper tag and instead choose the extra candidate (with a non-peaking D_s mass). The multiple candidate rate when the other side D_s decays semileptonically differs slightly from when it decays generically, creating a slight bias between tags

Table 10.3: Systematic errors from our D_s tag background shape.

Semileptonic mode	Relative systematic error
$\phi e\nu$	1.3%
$\eta e\nu$	0.9%
$\eta' e\nu$	1.3%
$f_0 e\nu$	0.8%
$K_s e\nu$	0.9%
$K^* e\nu$	2.1%

in semileptonic events and tags without a semileptonic event. The D_s tags' multiple candidate efficiency comes primarily from kinematics and differences in charged/neutral daughter hadron decay rates, both of which should be well modeled by the Monte Carlo. We determine the systematic shown in Table 10.4 by combining the tag multiple candidate efficiency, each semileptonic mode's branching ratio, and the small difference in multiple candidate efficiency for the semileptonic mode compared to the overall multiple candidate efficiency.

Table 10.4: Relative systematic error from the multiple candidate efficiency difference between semileptonic and all other D_s decay modes.

Semileptonic mode	Relative systematic
$D_s \to \phi e\nu$	0.11%
$D_s \to \eta e\nu$	1.67%
$D_s \to \eta' e\nu$	0.21%
$D_s \to f_0 e\nu$	2.20%
$D_s \to K_s e\nu$	3.05%
$D_s \to K^* e\nu$	0.28%

10.2 Tracking

10.2.1 Reconstruction

Our semileptonic hadron reconstruction efficiencies in Section 10.6 depend in part upon the reconstruction efficiency for their daughter pion and kaon tracks. CLEO has performed tracking systematics for π and K reconstruction using the baseline track selections that we also adopt for our D_s tagging [72]. They see no correction needed between the Monte Carlo and data track reconstruction efficiencies, with a 0.3% systematic on π tracks and a 0.6% systematic on kaon tracks.

The standard CLEO systematic analysis works for our typical π reconstruction, used in $f_0 e\nu$, $K^* e\nu$, and $\eta' e\nu$. However, the K systematic does not carry over to our analysis for a couple reasons. The 0.6% kaon reconstruction systematic depends on a momentum distribution peaking near 500 MeV, while we have a much softer momentum distribution with a peak closer to 250 MeV in $D_s \to \phi e\nu$. This alone might be dealt with using CLEO's kaon systematic broken into momentum bins, but we also loosen the kaon cuts for both ϕ and K^* reconstruction to increase the efficiency.

To this end, we have performed our own kaon systematic with the track selections used in this analysis and a momentum binning that better follows the kaon momentum distribution in our events. Since the standard CLEO particle ID cuts also differ from our own, we have combined both the kaon track reconstruction and kaon particle ID into one study, given in Appendix D. We find that the Monte Carlo does a good job of modeling higher momentum kaon tracks but that it overestimates the efficiency for low momentum tracks. CLEO's particle ID study [73] also found a large correction required at low kaon momenta, with our work matching those results.

We incorporate the π tracking systematic and the kaon systematics from Appendix D into our overall hadron efficiency systematic in Section 10.6.

10.2.2 π and K Particle ID

CLEO has performed a systematic study for the data/Monte Carlo efficiency difference in its standard π and K particle ID selections [73] using D^0 and D^+ decay modes in the 3770 MeV data. While we have a customized kaon systematic for our looser cuts (Appendix D), we do follow the standard cuts for pions in K^*, η', and f_0.

The standard study sees a systematic error of 0.02% for pions with an efficiency correction that has a strong momentum dependence. We correct each mode's particle ID efficiency based on that mode's pion momenta. Our f_0 and K^* pions have an efficiency correction close to the average from the original study (-0.49% per pion), but our slower η' pions give us a slightly larger correction. We've summarized the integrated corrections and systematics for these three modes from all constituent particles in Section 10.6.4.

10.3 Photon reconstruction

We only reconstruct photons for use in our D_s tags and in our η modes ($\eta e\nu$, $\eta' e\nu$). In the case of the D_s tags, our reconstruction efficiency doesn't impact our branching ratio measurements, as mentioned in Section 10.1. We also don't need to make an explicit photon reconstruction systematic for the η since we roll both daughter photon efficiencies into our overall η systematic (described later in Section 10.6.2 and in Appendix C).

10.4 Electron ID

Electron systematic errors have a strong momentum dependence due both to direct particle identification and to changes in the electron identification efficiency when in the presence of other tracks and showers [74]. While the systematic from direct electron identification dominates, we have also included the non-trivial environmental effects (following the $D^+ \to Xe^+\nu$ procedure). Each semileptonic mode has a distinct electron momentum distribution, causing each mode to have its own electron systematic.

Table 10.5 gives our final electron identification systematics and efficiency corrections after integrating over momentum and combining each systematic effect. We've included the different components for these systematics and corrections in Tables F.12 and F.13 (Appendix F).

10.4.1 Wrong Sign Electron

Any peaking background in our six semileptonic modes requires that we have a true D_s and that some track passes the electron cut. We explicitly correct each semileptonic mode for peaking background due to other semileptonic modes, which leaves a small

Table 10.5: Electron particle identification systematic and efficiency correction, by semileptonic mode.

Semileptonic mode	Electron PID systematic	Electron PID correction
$\phi e\nu$	0.68%	-1.91%
$\eta e\nu$	0.37%	-1.24%
$\eta' e\nu$	0.59%	-1.71%
$f_0 e\nu$	0.59%	-1.69%
$K_s e\nu$	0.38%	-1.24%
$K^* e\nu$	0.60%	-1.69%

peaking component from events with no direct electrons. Some of our modes ($K_s e\nu$, $K^* e\nu$) have problems with kaon tracks faking electron tracks, which we deal with by applying a missing mass cut. After correcting for events with a direct electron and those where another track fakes an electron, we only have to worry about real electrons that get produced indirectly, like through photon interactions in the detector.

Since these indirect electrons don't have a preferred sign, we can test the impact of such events by reconstructing the entire event with the required electron sign flipped (matching the D_s tag charge instead of the semileptonic D_s charge). We compare our expected numbers of wrong sign electron events from the Monte Carlo to our reconstructed wrong sign events in the data to see if we have an inconsistency.

We find very few wrong sign electron events in the data, with fit errors higher than the number of reconstructed events in each case. As Table 10.6 shows, this consistency with zero events matches our prediction from the Monte Carlo. Only $D_s \to f_0 e\nu$ even had a single measured wrong sign event in the data, and this mode also (not coincidentally) had the highest background from false D_s.

The predicted indirect electron events already have a statistical error from our Monte Carlo determined correction. We take no additional systematic for the Monte Carlo's modeling of such events.

Table 10.6: Passing events with a good D_s when reconstructing each semileptonic mode using an electron of the wrong charge. Our errors for the reconstructed events in each mode slightly exceed that mode's counts (all six modes are consistent with zero).

Semileptonic mode	Expected data events (from MC)	Actual data events
$\phi e\nu$	0.1	0.0
$\eta e\nu$	0.2	0.7
$\eta' e\nu$	0.0	0.0
$f_0 e\nu$	0.2	1.1
$K_s e\nu$	0.1	0.1
$K^* e\nu$	0.1	0.1

10.5 Monte Carlo Consistency

CLEO collected the 4170 MeV data over five data sets, which correspond to roughly two calendar time periods. The Monte Carlo generation reflects this separation in time, as datasets 39-41 use a different code release from datasets 47 and 48. While we expect no noticeable impact from the different generating time periods on our efficiencies, CLEO regularly tweaked its code to improve the accuracy between data and Monte Carlo. We have thus checked the generic Monte Carlo efficiencies for our six semileptonic modes across each dataset, as well as checking our signal Monte Carlo against the generic Monte Carlo (Figure 10.1).

The fluctuation across datasets and signal Monte Carlo for $D_s \to \eta e\nu$ pushed the limits of random variation. This, combined with a different run environment and η momentum range than in the standard CLEO η efficiency systematic, led to our performing a custom η systematic that we discuss in Appendix C.

We also find a discrepancy in the $D_s \to K_s e\nu$ efficiency between our signal Monte Carlo and the generic Monte Carlo. We traced this to a random 3σ fluctuation between the two in $\mathcal{B}(K_s \to \pi^+\pi^-)$. When we change our denominator to only consider events where the K_s decays to charged pions, this efficiency discrepancy disappears.

We do not take any systematic to our efficiency from internal Monte Carlo inconsistency for the other semileptonic modes, as the efficiencies in the signal Monte Carlo and across the five datasets falls within a reasonable distribution around the average.

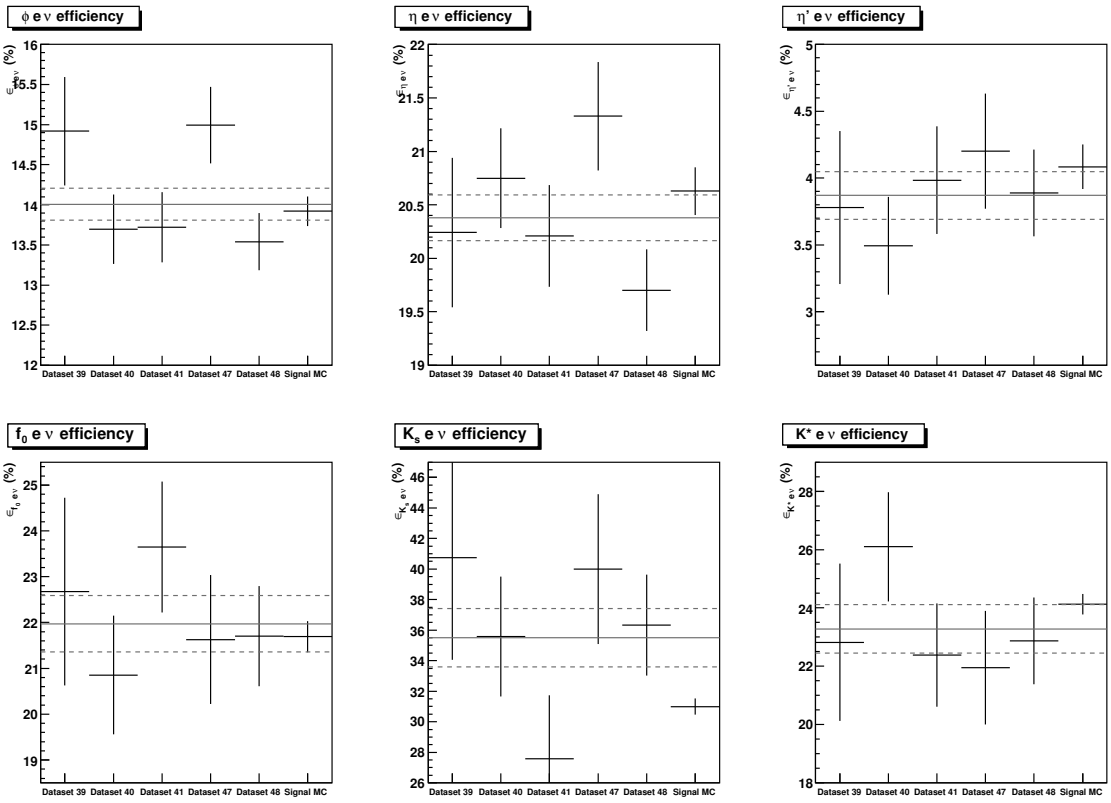

Figure 10.1: Efficiency (including hadron branching ratios) for each semileptonic mode, by dataset. The solid lines give the average across the full generic Monte Carlo sample, while the dotted lines show the 1σ range on this average.

10.6 Hadron Efficiencies

10.6.1 ϕ

We have two sources of possible systematic error to consider for our Monte Carlo determined ϕ efficiency. The largest concern comes from the kaon reconstruction and particle ID efficiencies, discussed in Appendix D, which requires both a systematic error and an efficiency correction. The other possible systematic arises from the possibility that the combined kaon tracks in the data could have a broader invariant mass distribution (poorer resolution) than predicted by the Monte Carlo, extending part of the ϕ distribution beyond our mass window.

Our kaon reconstruction and identification efficiency measurement shows that we tend to properly reconstruct soft kaons in the Monte Carlo more often than we actually do in the data. Since the two daughter kaons from the ϕ have correlated momenta, we need to correct our ϕ efficiencies based on each kaon momentum pair. Figure 10.2 shows the result of our ϕ efficiency correction, by momentum. For the predicted ϕ momentum distribution from $D_s \to \phi e\nu$ in the Monte Carlo (ISGW2 model), this results in a $\phi e\nu$ semileptonic efficiency change of -8.2% (relative).

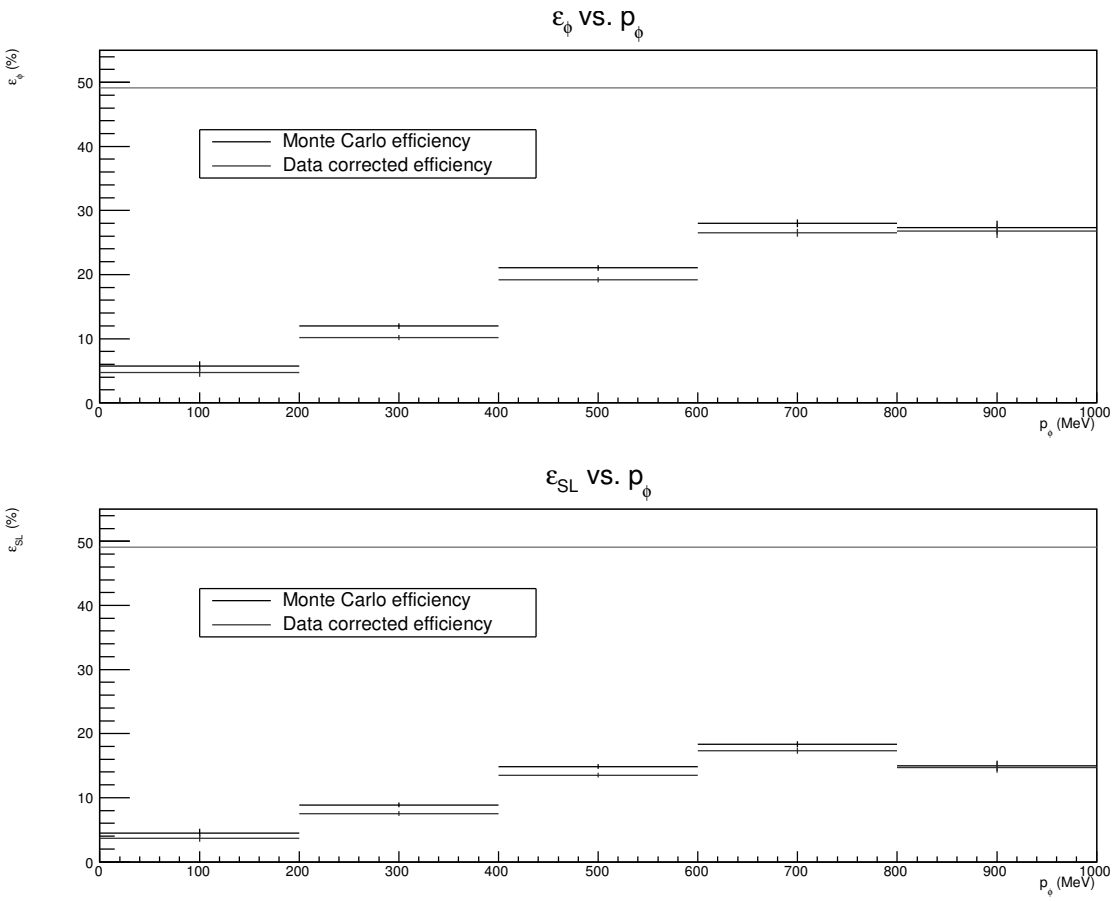

Figure 10.2: Top: ϕ efficiency in the Monte Carlo, by momentum, before and after correcting the efficiency based on the kaon systematic study in Appendix D. Bottom: $D_s \to \phi e\nu$ semileptonic efficiency, by ϕ momentum, before and after correction.

We additionally obtain two systematic errors from our kaon reconstruction study in

Appendix D. The first comes about directly from our measurement limitations on the kaon efficiency in each momentum bin. We treat the individual kaon systematics as correlated and obtain a relative systematic error of 1.7% for $\phi e\nu$ given our ϕ momentum distribution. The second systematic results from the process of splitting the kaon momenta into bins in the first place, given that each bin may not have a constant efficiency. Considering different efficiency distributions across the bins gives us a relative systematic error of 0.5%.

In addition to finding and correctly identifying the kaons, they also have to combine to form a ϕ mass that falls within our cut window (-15 MeV $< M_\phi^{\text{recon}} - M_\phi^{\text{PDG}} <$ 30 MeV). Since our mass window already captures most of the ϕ spectrum, we don't expect any resolution difference to significantly affect our efficiency. We explicitly test this by assuming that the data ϕ could peak in a slightly different location (a shift) and by taking the data resolution to have a gaussian smear convoluted with the Monte Carlo resolution.

We don't have enough data to test the ϕ resolution explicitly in our tagged $D_s \to \phi e\nu$ analysis, so we instead use an inclusive approach by plotting the KK spectrum when we find an electron in the event (no D_s tagging). We use our standard $D_s \to \phi e\nu$ kaon cuts but to avoid electron-only events (e.g. $ee \to eeee$, where two electrons fake kaons), we also require that the kaons not pass electron cuts and that the tracks not be too close to the beamline ($|\cos(\theta)| < 0.8$). We have redone the systematic relaxing these additional kaon restrictions with a similar result, but we get less precision due to the extra background.

We use the Monte Carlo signal and background functions to fit our data spectrum, allowing the signal to shift or have a poorer resolution from a gaussian smear. Figure 10.3 shows our best fit to the data (a peak shift of -0.05 MeV and smear with a σ of 0.1 MeV) alongside the fit with no shift or smear allowed. Not only do we obtain a small relative systematic of 0.025% from the ϕ resolution over our large window, but we find that the Monte Carlo matches the data well enough that we'd see a small systematic even with a tighter mass window. We also show more detailed fit results for our different shifts and smears with Figures G.46 and G.47 in Appendix G.

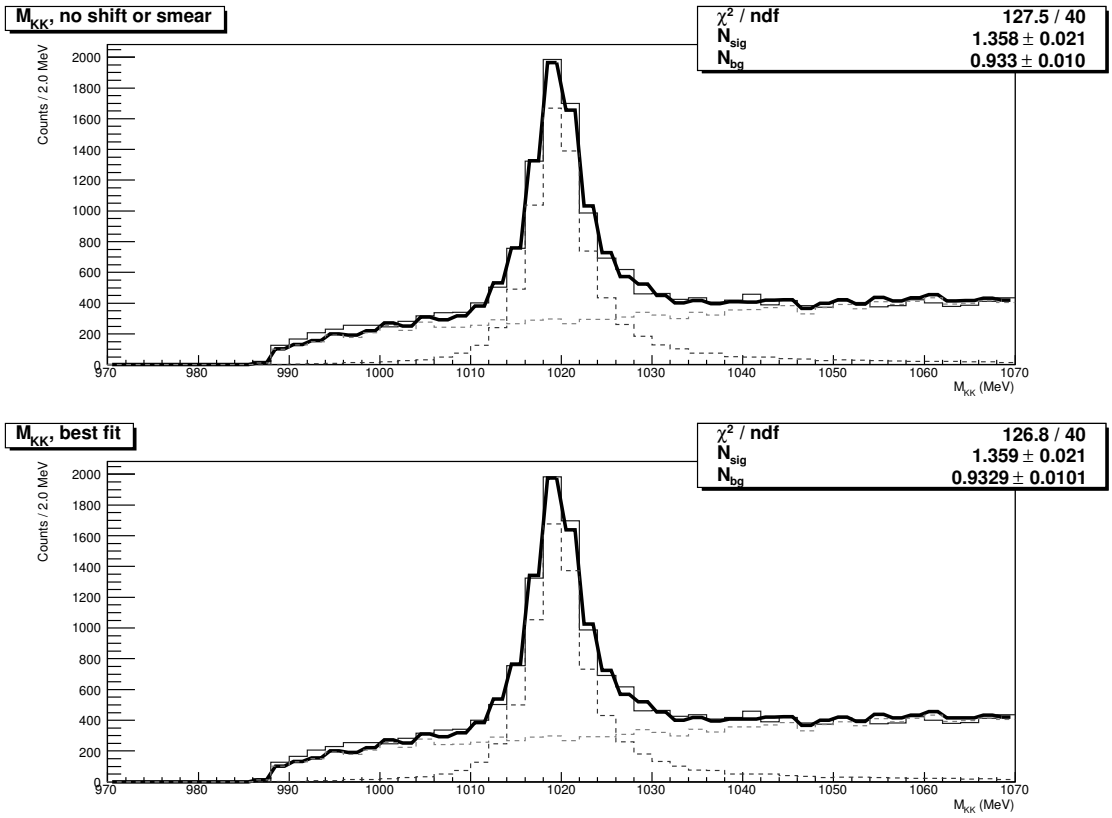

Figure 10.3: Top: ϕ mass fit in the data, using the signal and background produced in the Monte Carlo. Bottom: Best ϕ mass fit in the data after allowing the signal Monte Carlo histogram to shift its peak and convoluting it with variable width gaussians.

10.6.2 η

The best existing study of η efficiency systematics at CLEO [71] uses a 3σ pull mass cut in its η selection, compared to the 5σ cut we use in this analysis. That study also uses monoenergetic η from $\psi' \to \eta J/\psi$ with a relatively soft η momentum of 199 MeV. $D_s \to \eta e\nu$ involves a much wider η momentum spectrum, as seen in Figure 9.4, which gives us a much wider range of resultant photon energies. Furthermore, $\psi' \to \eta J/\psi$ with $J/\psi \to \mu^+\mu^-$ produces a cleaner environment (fewer photons and poorly tracked hadrons) than the environment we see in $D_s^* D_s$ events.

For these reasons, we have decided to do our own η efficiency systematic study. This results in worse precision on our systematic than the high statistics ψ' study, but we feel that it more accurately represents the η selections and environment in our analysis.

Our η efficiency systematic uses the same D_s tag modes as our standard analysis, but it takes advantage of the relatively large $D_s^\pm \to \rho^\pm \eta$ branching ratio to get a clean η sample. We have described our full efficiency systematic technique and results in Appendix C.

Unlike the previous study of η efficiency systematics at CLEO, we do not find that we need a Monte Carlo correction to match the data. That study suggested a relative correction of -5.6% to the η efficiency for η without an $\frac{E_9}{E_{25}}$ O.K. cut (-6.5% for those with an $\frac{E_9}{E_{25}}$ O.K. cut). In our study, with a broader pull mass cut and wider η momentum range, we find that the data and Monte Carlo agree to within 1.2%. This falls well within the precision of our study, so we take no efficiency correction.

Our systematic procedure has the downside of requiring a large systematic error given the statistical error in our $\eta e\nu$ measurement. While the previous η study had a relative 5.9% efficiency systematic after extrapolating to a wider momentum region, our η systematic procedure yields a 7.9% relative systematic. This makes it a limiting error in our overall $D_s \to \eta e\nu$ measurement.

10.6.3 K_s

The standard CLEO systematic study on K_s reconstruction [75] shows no efficiency difference between Monte Carlo and data up to 0.8%, as long as both daughter π tracks have been found. However, our backgrounds in $D_s \to K_s e\nu$ lead us to make tighter

selections on the K_s than the prior study. While they used a 12 MeV mass cut with no flight significance selection for the K_s, we use a 6.3 MeV mass cut and a flight significance greater than 4.0. Furthermore, we require that the K_s daughter tracks fall within 20 cm of the origin to avoid $\tau\nu$ backgrounds, which the generic K_s study did not have to concern itself with.

Since our selections lead to a significantly different efficiency than the loose cuts from the standard study (a relative difference of about 30%), we have run our own systematic for K_s reconstruction. The low statistics in our $D_s \to K_s e\nu$ measurement mean that we can have a fairly forgiving precision from our K_s systematic without impacting our overall error. We consequently try to keep our systematic study's environment as close to $D_s \to K_s e\nu$ as possible by using tagged $D_s^* D_s$ decays (4170 MeV data) in our systematic measurement rather than the higher statistics 3770 MeV data.

We compared the K_s momentum spectrum in $D_s \to K_s e\nu$ to that from several other K_s modes, and we ultimately chose $D_s \to K_s K$ for our systematic above 650 MeV and $D_s \to K^* K^*$ ($D_s^\pm \to K_s K^\mp \pi^\pm \pi^\pm$) for our lower momentum systematic (Figure 10.4). We use a procedure similar to that for our kaon systematic (Appendix D) by reconstructing all particles other than the K_s, then fitting the recoil mass both when we successfully reconstruct a K_s ("found" events) and when we don't find a K_s ("not found" events). Rather than trust that the general π tracking systematic applies for K_s daughter tracks, we combine our track and K_s reconstruction into one systematic (i.e. we don't require two extra tracks before looking for candidate K_s events).

In reconstructing all other particles in the event, we need to find a D_s tag, the D_s^* daughter photon, and either a single kaon (for $K_s K$) or a kaon and two pions (for $K^* K^*$). Our D_s tag involves the same 13 tag modes as our full analysis, with selections given in Section 5. For a cleaner sample, we also add a mass cut to the D_s, with each tag mode's cut listed in Table F.11 (Appendix F).

We use the selections from Section 9.1 for our D_s^* daughter photons, excepting the minimum energy cut (irrelevant here). After ensuring that the $D_s + \gamma$ recoil mass matches a D_s (M_{recoil} between 1950 MeV and 1990 MeV), we improve the later resolution slightly by rescaling our photon energy to match the $D_s^* D_s$ kinematic requirement. Our kaons and non-K_s pions must have the proper charges relative to the tagged D_s and pass the selections from Section 8.1. We also require $p_\pi > 100$ MeV to eliminate soft

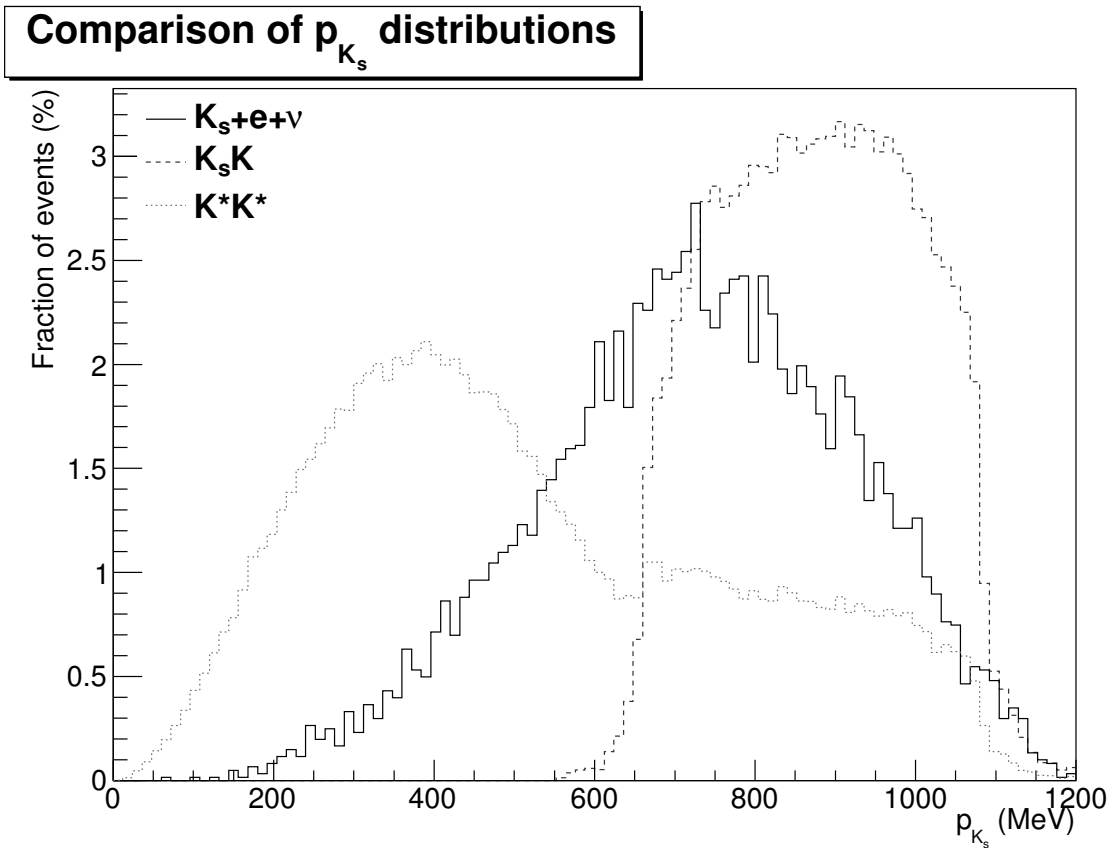

Figure 10.4: Normalized K_s momentum distributions in D_s decays to $K_s e\nu$, $K_s K$, and K^*K^* ($K_s K^{\mp}\pi^{\pm}\pi^{\pm}$). We use $K_s K$ to study K_s reconstruction above 650 MeV and K^*K^* to study the systematic below 650 MeV.

pion swaps.

After reconstructing all other particles, we separate candidate events into three systematic regions based on their recoil momentum: a high momentum K_s region (650+ MeV), a medium p_{K_s} region (400 MeV– 650 MeV), and a low p_{K_s} region (200 MeV–400 MeV). We then attempt to reconstruct a K_s using the selections from our $D_s \to K_s e\nu$ analysis in Section 8.2 ($|M_{K_s}^{\rm recon} - M_{K_s}^{\rm PDG}| < 6.3$ MeV, K_s flight significance > 4.0, and $\rho_0^{\pi} = \sqrt{d_b^2 + z_0^2} < 20$ cm for the π from K_s). We make both a "found" and "not found" plot for the recoil mass against the K_s candidate in each momentum region.

We fit each momentum region's recoil mass plots with a double gaussian for the signal shape and either a linear background function, a scaled histogram background, or both depending on the characteristics of each mode (e.g. the K_sK "not found" recoil mass fit requires a $K\eta$ background shape; the K^*K^* requires an extra shape for softly peaking false D_s^* daughter photons). Figure 10.5 contains the "found" and "not found" fits for our K_sK data, while Figure 10.6 has the "found" and "not found" data fit results for the two K^*K^* momentum regions. For completeness, we have also included the fits from Monte Carlo in Figures G.53 and G.54 as part of our extra figures section (Appendix G).

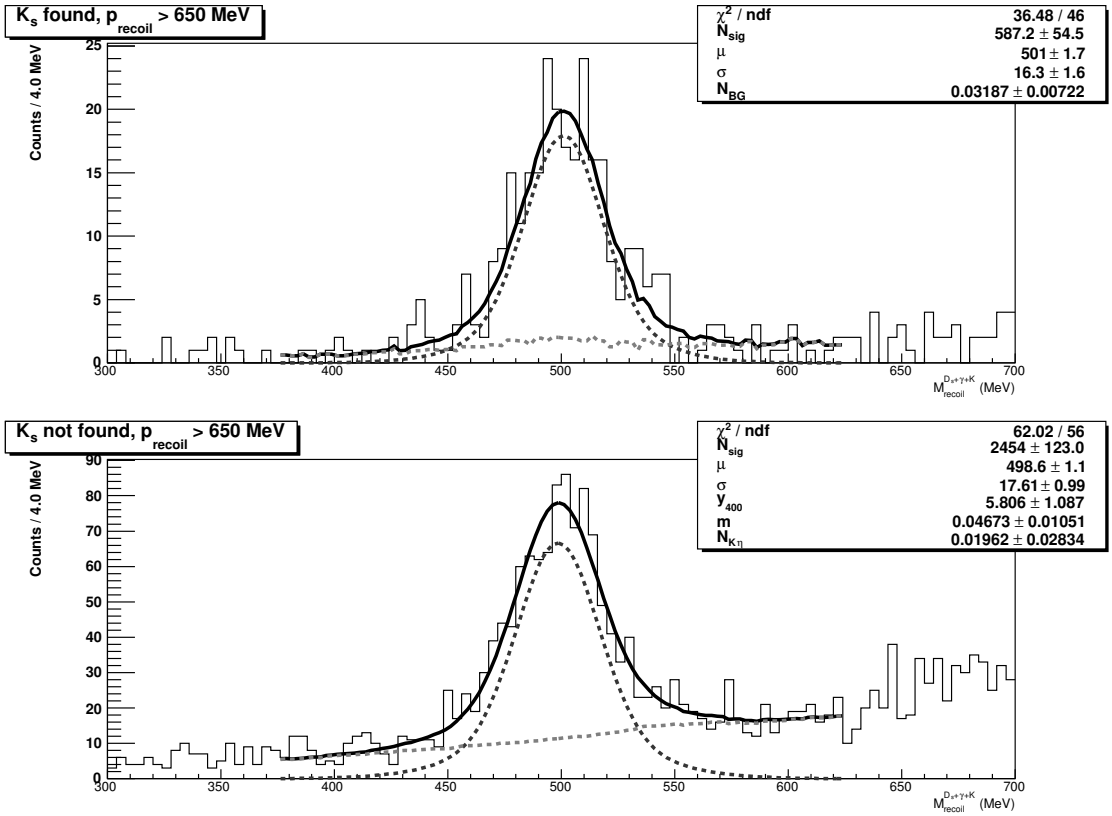

Figure 10.5: $D_s + \gamma + K$ recoil mass in data K_sK events for "found" and "not found" K_s.

Table 10.7 gives the efficiency results from our various momentum region fits. We

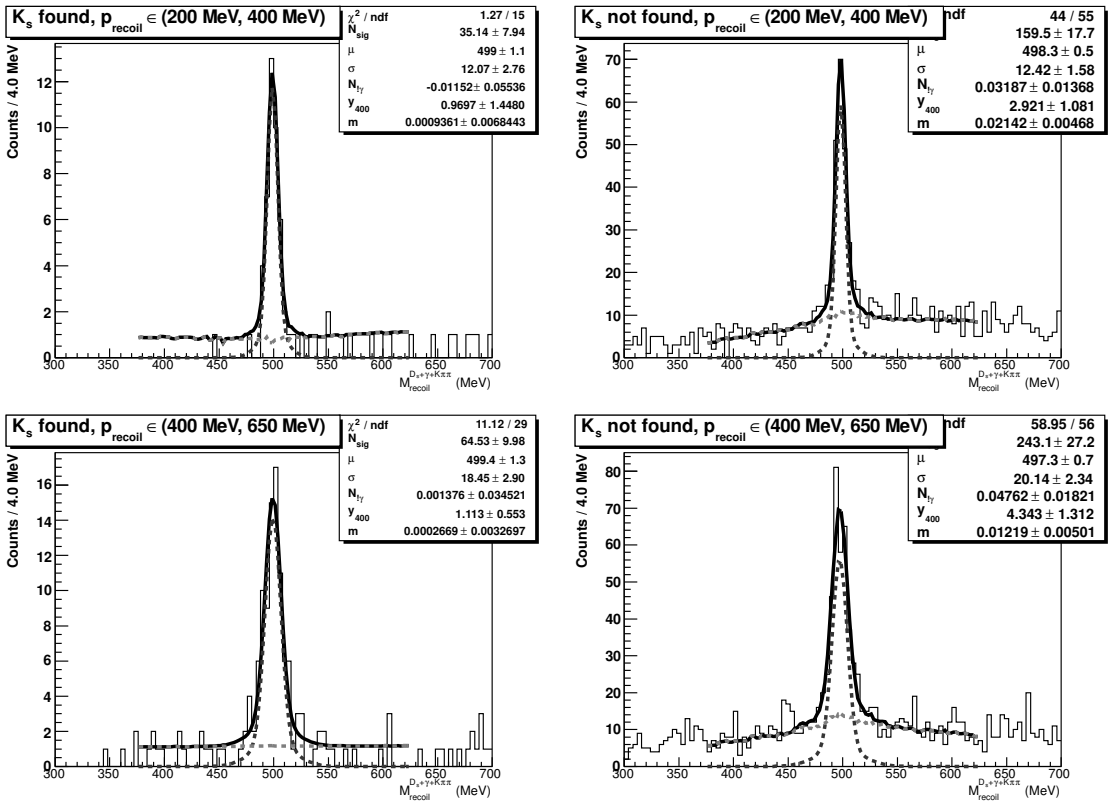

Figure 10.6: $D_s+\gamma+K\pi\pi$ recoil mass in data K^*K^* events for "found" and "not found" K_s. The top row shows only low momentum K_s while the bottom row gives results in our medium K_s momentum region, with p_{K_s} determined by the recoil momentum.

had decided before the study to take a correction to our K_s efficiency if we found a 2σ or larger difference between data and Monte Carlo, and the high momentum region just reached this threshold. The lower K_s efficiency in this region comes from the daughter tracks themselves not being reconstructed properly. When we repeat the analysis requiring that the event has two candidate K_s tracks (with invariant mass between 300 MeV–700 MeV and a combined momentum within 60 MeV of the K_s, following [75]), the difference between data and Monte Carlo disappears (with both about 80% efficient). We get a final correction to our overall K_s efficiency of -11.1% after weighting by the K_s momentum distribution in $D_s \to K_s e\nu$.

We base our systematic on the combination of our data precision and the Monte

Table 10.7: K_s efficiency systematic and correction from our found/not found recoil mass fits in each momentum region.

p_{K_s} region	Syst. mode	$\varepsilon_{\mathrm{MC}}^{\mathrm{true}}$	$\varepsilon_{\mathrm{data}}$	Correction	Systematic
200 MeV–400 MeV	K^*K^*	18.5%	(17.7 ± 3.7)%	—	21.1%
400 MeV–650 MeV	K^*K^*	19.2%	(20.6 ± 3.1)%	—	15.5%
650+ MeV	K_sK	23.7%	(20.0 ± 1.7)%	-15.2%	8.6%
Integrated	Combined	N/A	N/A	-11.1%	7.3%

Carlo in-vs.-out precision. Since each K_s momentum region has the same efficiency to within error, we do not apply an additional systematic to account for using finite-sized K_s momentum bins. Like our K_s efficiency correction, we have weighted each momentum region's systematic error based on the $K_se\nu$ momentum spectrum to get an integrated $D_s \to K_se\nu$ systematic of 7.3%.

10.6.4 K^*, η', and f_0

Our $K^*e\nu$, $\eta'e\nu$, and $f_0e\nu$ modes all have hadrons with relatively broad mass distributions that decay to a final state involving pions (with one kaon for K^* and one η for η'). Since we do make a (broad) cut on each resonant particle's mass distribution to avoid allowing in too much background, we need to ensure that the mass resolution in data matches the Monte Carlo. Additionally, each mode has a distinct momentum spectrum for its final particles, which we need to incorporate into our momentum dependent tracking and particle identification systematics.

In all three semileptonic modes, we determine the mass resolution by reconstructing the candidate hadron (K^*, η', or f_0) in a fairly clean D_s mode ($D_s \to K^*K$, $D_s \to \pi\eta', \eta' \to \pi\pi\eta$, or $D_s \to f_0\pi$, respectively). We fit the candidate hadron's mass in the data by using the Monte Carlo signal shape. However, we allow the signal mass distribution to shift either direction, and we convolute the signal shape with a gaussian to model potentially poorer data resolution. We take our systematic to be the relative change in events passing our mass cut window for the smeared and unsmeared distributions. In all three cases, the systematic fell well within the precision of our measurement, so we find no need for a correction to our efficiency.

To get a clean sample for our mass resolution, we reconstruct a D_s tag, a D_s^* daughter photon, and a kaon (K^*) or pion (η', f_0), following the same procedures and particle selections as for our K_s systematic (Section 10.6.3). We reject any events with extra tracks after the candidate hadron reconstruction, and we cut on the momentum or missing mass of the event to further clean up our sample, as follows:

- 650 MeV $< p_K^{\text{in } D_s \text{ rest}} <$ 750 MeV for the kaon in K^*K
- 650 MeV $< p_\pi^{\text{in } D_s \text{ rest}} <$ 800 MeV for the pion in $f_0\pi$
- 900 MeV $< M_{\text{recoil}}^{D_s+\gamma+\pi} <$ 1000 MeV in $\pi\eta', \eta' \to \pi\pi\eta$
- $\left| M_{\text{recoil}}^{D_s+\gamma+\pi+\eta'} \right| <$ 75 MeV in $\pi\eta', \eta' \to \pi\pi\eta$

We have placed the resulting K^*, η', and f_0 mass resolution plots in Appendix G (Figures G.55 – G.57). We have put our final mass resolution systematic in Table 10.8 (along with our tracking and particle identification systematics, described below).

We've taken a correlated 0.3% track reconstruction systematic for each π and electron track [72]. Since our failures in kaon track reconstruction come primarily from decays in flight, we consider this uncorrelated to the other tracking systematics and have instead incorporated it into our momentum-dependent kaon particle ID systematic. For simplicity, we've also absorbed our rather large η reconstruction efficiency systematic into the η' particle identification systematic. Only pion and kaon particle identification require a correction to our efficiency, which we have summarized in Table 10.9.

Table 10.8: Summary of semileptonic hadron systematic errors. Our kaon and η systematics have been included into the K^* and η' PID columns, respectively.

Semileptonic hadron	Track reconstruction	PID	Mass resolution
η'	0.90%	7.90%	3.15%
f_0	0.90%	0.04%	2.63%
K^*	0.60%	1.21%	2.59%

Table 10.9: Summary of hadron efficiency corrections from particle identification.

Semileptonic hadron	Relative ε correction
η'	-2.94%
f_0	-0.50%
K^*	-2.88%

10.7 Decays in Flight

For all modes other than $D_s \to \eta e\nu$, we consider any passing event containing the semileptonic mode in question as a true event, regardless of whether or not we correctly reconstructed the semileptonic side. Normally, we don't have any ambiguity in this procedure because we actually did reconstruct the semileptonic side correctly according to the Monte Carlo. However, we do see a few cases in each mode where a true semileptonic event passes our cuts with an incorrectly reconstructed semileptonic-side hadron. The vast majority of these cases come from either a kaon or pion on the semileptonic side decaying in flight, nearly always to a muon.

The Monte Carlo should model the decay-in-flight kinematics without difficulty. However, the "kinked" track could cause problems with tracking reconstruction and the resultant track momentum, used indirectly (and sometimes directly) in our event reconstruction. Given the small effect, we simply take 50% of the efficiency for true but incorrectly reconstructed events as our systematic error to model any possible data/Monte Carlo differences.

Table 10.10 gives the efficiency for all such true but incorrectly reconstructed events and our ultimate systematic (expressed as a relative error). This systematic includes both decay-in-flight events and all other events, although the efficiency from other events always falls well below the threshold at which we include systematic errors (a relative 0.3% error).

10.8 Splitoff Rate

Our $D_s \to \eta e\nu$ fits need to correct for "volunteer" combinations, which predominantly come from either the D_s^* daughter photon or a splitoff shower combining with a true η

Table 10.10: Systematic for true semileptonic events that pass with incorrect particle identification, mostly due to π or K decays in flight to μ. We take 50% of the effect's size in Monte Carlo as our systematic.

Semileptonic mode	$\varepsilon_{\text{SL pass, wrong MC tag}}$	Relative systematic
$D_s \to \phi e\nu$	0.023%	0.08%
$D_s \to \eta' e\nu$	0.040%	0.49%
$D_s \to f_0 e\nu$	0.226%	0.52%
$D_s \to K_s e\nu$	0.390%	0.63%
$D_s \to K^* e\nu$	0.343%	0.71%

daughter photon to make an extra η candidate. Section 9.2.2 discusses this effect with a procedure using $D^0 \to K^*\eta$ that lets us correct the splitoff rate from the Monte Carlo. We obtained a splitoff correction consistent with the Monte Carlo rate (1σ difference), so we don't take a bias correction from splitoff. However, our splitoff rate procedure has an associated uncertainty, which we take as the systematic error for the Monte Carlo's splitoff model. Varying the splitoff rate across its 1σ range gives us a $\mathcal{B}(D_s \to \eta e\nu)$ relative systematic of 1.16%.

10.9 Hadronic Branching Ratios

We measure each semileptonic hadron decay through a particular decay mode (e.g. $\eta \to \gamma\gamma$). The semileptonic modes' efficiencies depend upon the branching ratio for these hadronic decays. However, updates to the known hadronic branching ratios from more recent measurements require an efficiency correction, while uncertainties in those branching ratios contribute a systematic error. Table 10.11 gives these corrections and systematics using the most recent branching ratios from the Particle Data Group [21]. The η' decay includes both the direct uncertainty in $\mathcal{B}(\eta' \to \pi\pi\eta)$ and the uncertainty from $\mathcal{B}(\eta \to \gamma\gamma)$.

Table 10.11: Systematic errors and efficiency corrections from uncertain or changed branching ratios in semileptonic daughter hadron decays.

Hadron decay	$\mathcal{B}_{\rm MC}$	$\mathcal{B}_{\rm PDG}$	Systematic	Correction
$\phi \to KK$	49.1%	$(48.9 \pm 0.5)\%$	1.0%	—
$\eta \to \gamma\gamma$	39.5%	$(39.4 \pm 0.2)\%$	0.5%	—
$\eta' \to \pi\pi\eta$	43.7%	$(42.9 \pm 0.7)\%$	1.7%	-1.8%
$K_s \to \pi\pi$	68.6%	$(69.2 \pm 0.1)\%$	0.1%	0.9%

10.10 Semileptonic Fit Functions

After making our semileptonic cuts, we fit each mode's D_s mass spectrum using a linear background function and a signal shape from the tagging fits. The D_s mass signal shape shouldn't generate an additional systematic beyond that discussed in Section 10.1.1 because we use the same shape for our branching ratio's numerator and denominator. However, we made the choice of a linear background function empirically, with parameters from a fit to the Monte Carlo's predicted background.

To investigate a potential systematic from our choice of background function, we have replaced our linear background function with a constant function and compared the branching ratio results. A constant function generally goes beyond the 1σ variation on our linear fit to background, but we take this as a worst case scenario on the Monte Carlo's effective background model. The results from Table 10.12 show that we get a negligible systematic even for this worst case.

The $D_s \to \eta e\nu$ mode also includes an explicit fit to the η pull mass spectrum, using a signal histogram shape and a linear background. For this mode, we independently take a constant background on the η pull mass and each D_s mode. Table 10.12 contains the results after combining all $\eta e\nu$ background systematics in quadrature, which still yields a negligible systematic.

We took the η signal shape for our $D_s \to \eta e\nu$ fit from the Monte Carlo. We've used two techniques to obtain a systematic on this signal shape. In the first technique, we extract an η pull mass spectrum from a clean $D_s \to \pi\eta$ sample, we fit the data's pull mass spectrum with a crystal ball shape, and then we use the resultant crystal ball function as our η signal shape in the branching ratio fit. For our second technique, we

Table 10.12: Branching ratio change from a different semileptonic background function. The $D_s \to \eta e\nu$ line combines changes to both the pull mass and D_s mass backgrounds. In all cases, the systematic from choosing a different background shape falls well below the statistical or systematic error.

Semileptonic mode	$\delta\mathcal{B}/\mathcal{B}$	$\sigma_{\text{BG syst}}/\sigma_{\text{stat}}$
$\phi e\nu$	0.50%	0.063
$\eta e\nu$	0.39%	0.075
$\eta' e\nu$	1.04%	0.048
$f_0 e\nu$	−0.86%	0.048
$K_s e\nu$	0.63%	0.034
$K^* e\nu$	−0.35%	0.016

convolute the Monte Carlo's η pull mass spectrum with a gaussian of varying widths and compare the best fit branching ratio to our standard branching ratio (without a gaussian smear). Both techniques come to similar relative systematic estimations (1.0% vs. 0.8%), but we have chosen the $D_s \to \pi\eta$ method as most closely representing the uncertainty in signal shape.

To get our clean η pull mass spectrum from $D_s \to \pi\eta$, we reconstruct a tagged $D_s + \gamma$, do a basic kinematic fit on the D_s^* daughter photon, then find the other side π and η. We use the same D_s tag modes and cuts as in our normal analysis (Section 5), dropping the four tag modes with η daughters to avoid any possible complications. We cut on the D_s mass based on its tag mode (Table F.11 in Appendix F), on the $D_s + \gamma$ recoil mass (1950 MeV–1990 MeV), and on the π momentum in the D_s rest frame (within 20 MeV of the ideal 902 MeV). Once we have a reconstructed η, we also require that the event's missing mass fall within 100 MeV of zero.

Events passing all our $D_s \to \pi\eta$ cuts have nearly no background, giving us a very pure η sample. We fit the data's pull mass spectrum to a crystal ball function and use that (slightly wider than Monte Carlo) shape in our branching ratio fit. Our branching ratio changes by a relative 1.0%, so we take that to be our η signal shape's systematic.

10.11 D_s Production Efficiencies

Tagged and semileptonic D_s get created through multiple modes at 4170 MeV. The e^+e^- collision can directly produce $D_s^+D_s^-$ ($\sigma_{D_sD_s} = 0.034$ nb), or it can produce $D_s^*D_s$ ($\sigma_{D_s^*D_s} = 0.916$ nb) [61]. Further, the D_s^* may decay to either $D_s\gamma$ (94%) or to $D_s\pi^0$ (6%) [21]. Each of these D_s production mechanisms have associated uncertainties, while the Monte Carlo that we use to determine our semileptonic efficiency simply takes each production mode's most likely value.

To incorporate the D_s production mode uncertainty into our overall error, we have determined each process's tag and semileptonic efficiencies (Table 5.6 for tags; Tables 7.2, 9.1, and F.14–F.17 for semileptonics). We then vary each D_s production uncertainty by 1σ and take the change in our average efficiency as a systematic. In practice, $D_sD_s^*, D_s^* \to D_s\gamma$ production dominates our efficiency. This dominance and the fact that each production mode has similar semileptonic efficiencies makes the D_s production systematic negligible (Table 10.13).

Table 10.13: Relative systematic for various D_s production rate uncertainties. This combines the uncertainties from the D_sD_s and $D_s^*D_s$ cross sections at 4170 MeV with the uncertainty from the D_s^* branching ratio (the fraction going to $D_s\gamma$ vs. $D_s\pi^0$). These combined effects still contribute a negligible systematic.

Semileptonic mode	$\sigma^{\text{syst}}_{D_sD_s}$	$\sigma^{\text{syst}}_{D_s^*D_s}$	$\mathcal{B}(D_s^* \to D_s\pi^0)^{\text{syst}}$	Total relative systematic
$D_s \to \phi e\nu$	0.0033%	-0.0018%	-0.0028%	0.0047%
$D_s \to \eta e\nu$	-0.0001%	0.0001%	-0.0000%	0.0001%
$D_s \to \eta' e\nu$	-0.0091%	0.0050%	-0.0192%	0.0218%
$D_s \to f_0 e\nu$	-0.0035%	0.0019%	0.0072%	0.0082%
$D_s \to K_s e\nu$	-0.0019%	0.0011%	0.0316%	0.0317%
$D_s \to K^* e\nu$	0.0029%	-0.0016%	0.0113%	0.0118%

10.12 Final State Radiation

When the D_s decays to charged particles, the decay can also include photons emitted via an electromagnetic interaction with the final state charged particle. This final state

radiation (FSR) doesn't cause us a problem in tagged D_s, as any tag efficiency drop will be reflected proportionally in our branching ratio's numerator and denominator. However, the quarks that make up the semileptonic hadron and particularly the electron produced in the semileptonic decay may have FSR that distorts the semileptonic efficiency. We use the `PHOTOS 2.0` package to estimate FSR in our Monte Carlo.

Since FSR emission from charged particles mostly results in soft photons and our particle efficiencies stay fairly flat outside the extreme regions, our decays' efficiencies don't change much with the inclusion of FSR. Only about 2% of semileptonic decays (varying slightly by mode) have FSR that alters daughter particle momenta enough to push the combined $he\nu$ momenta outside its allowed kinematic range. Of those decays, 90% still have relative efficiencies within 5% of the non-FSR efficiency. Roughly 0.2% of semileptonic decays see a significant efficiency drop, mostly due to the electron momentum falling below threshold.

Table 10.14 gives the efficiency difference with and without FSR for each semileptonic mode. Past work [65] has taken 30% of this difference as a systematic, but none of our efficiency variations affect the overall systematic error even if we take the entire drop as our systematic. We include the results here for reference, but we otherwise dismiss FSR as a systematic effect.

Table 10.14: Efficiency difference due to final state radiation, by D_s semileptonic mode.

Semileptonic mode	Relative systematic
$D_s \to \phi e\nu$	0.19%
$D_s \to \eta e\nu$	0.06%
$D_s \to \eta' e\nu$	0.06%
$D_s \to f_0 e\nu$	0.28%
$D_s \to K_s e\nu$	0.19%
$D_s \to K^* e\nu$	0.28%

10.13 Initial State Radiation

In the initial e^+e^- collision, one of the two charged particles may emit a soft photon. This initial state radiation (ISR) lowers the collision's center-of-mass energy. Since

CLEO-c ran at 4170 MeV, just above the $D_s^* D_s$ threshold of 4081 MeV, the D_s momenta for events with even moderate ISR can vary significantly.

Fortunately, the Monte Carlo provides a good model for ISR, with D_s single tags at 4170 MeV matching the Monte Carlo's ISR prediction to within 0.6% [65]. Nonetheless, the Monte Carlo predicts that just over 10% of events in our sample will have a center of mass below 4160 MeV, so we have checked the semileptonic efficiency difference for events produced at lower center-of-mass energies.

Table 10.15 gives the efficiency difference between events produced without ISR and events that include ISR. Not surprisingly, we find very little difference between the two given the fairly flat efficiency across D_s momenta and the fact that most events don't have significant ISR. Even if the Monte Carlo had too little ISR by 30% (well above the precision extrapolated from the single tag study mentioned previously), we could ignore this systematic. We thus take no additional systematic from ISR effects.

Table 10.15: Efficiency difference due to initial state radiation, by D_s semileptonic mode.

Semileptonic mode	$\Delta\varepsilon/\varepsilon$
$D_s \to \phi e\nu$	0.12%
$D_s \to \eta e\nu$	0.79%
$D_s \to \eta' e\nu$	0.09%
$D_s \to f_0 e\nu$	0.09%
$D_s \to K_s e\nu$	0.77%
$D_s \to K^* e\nu$	0.27%

10.14 Generating Models

Our semileptonic efficiency primarily depends on the electron and hadron momenta in the detector (the lab frame). These momentum distributions and their correlations are determined in part by the form factors' q^2 dependence, which isn't easily calculable from first principles. Our Monte Carlo uses the ISGW2 [34] quark model when generating the semileptonic D_s decay, but various pole dominance models [36] offer alternate form factor dependencies and consequent momentum distributions. While we believe the

ISGW2 model best represents the underlying physics given the relatively heavy c quark and the wide q^2 range relative to the number of close resonances in the charm system, the pole model has been used most often in the literature and gives us a simple alternative to estimate our efficiencies' form factor model dependence.

We have used the pole model's simplest form as our point of comparison, in which a single resonance dominates the form factors. In this case, each form factor has a $(1 - \frac{q^2}{M^2})^{-1}$ dependence on q^2, where M is the mass of the nearest meson resonance with appropriate quantum numbers. We use a D_s^* pole mass for our vector form factors and a $D_{s1}(2460)$ pole mass for our axial form factors, matching prior work [53]. For D_s decays to vector hadrons (ϕ, K^*), we have three form factors and also need the relative normalizations between them; we use $r_v = 1.81$ and $r_2 = 0.82$, where $r_v = \frac{V(0)}{A_1(0)}$ and $r_2 = \frac{A_2(0)}{A_1(0)}$ are the relative normalizations at $q^2 = 0$ for the vector/axial and axial/axial form factors, respectively.

We generate our D_s decays using both our baseline model (ISGW2) and the simple pole model, then we treat the difference between the two as a 1σ systematic arising from the generating model. The default CLEO Monte Carlo had some minor coding errors in the masses for its ISGW2 implementation; we corrected those in our own implementation but found that the final systematic didn't change. For a further comparison point between models, we also include the original, less sophisticated ISGW model (no relativistic corrections, exponential form factor dependence). We did not use the ISGW model in our systematic estimate, although it would have had only a minor effect in any case.

Table 10.16 contains our systematic for each semileptonic decay mode. In all cases, the pole model created events with a higher efficiency than the ISGW2 model. The slightly more energetic electron spectrum in pole model events dominated this efficiency increase with fewer electrons below our 200 MeV minimum p_e cut (Figure 10.7). While decays to pseudoscalar and scalar hadrons have similar q^2/E_e correlations between models and thus don't have a significant efficiency change beyond the higher electron efficiency, decays to vector hadrons see a further efficiency increase in the pole model from a more beneficial q^2/E_e correlation, pairing lower q^2 (higher E_{hadron}) with more energetic electrons (Figure 10.8).

Table 10.16: Relative systematic from different generating models' reconstruction efficiency.

D_s mode	$\frac{\varepsilon_{\rm Pole}-\varepsilon_{\rm MC}}{\varepsilon_{\rm Pole}}$
$D_s \to \phi e\nu$	2.9%
$D_s \to \eta e\nu$	0.7%
$D_s \to \eta' e\nu$	1.6%
$D_s \to f_0 e\nu$	2.3%
$D_s \to K_s e\nu$	1.4%
$D_s \to K^* e\nu$	5.1%

In our extra figures section, we have included the lab frame hadron and electron momenta for each of the other four semileptonic modes (Figures G.58 and Figures G.59). We've also included the q^2 and q^2 vs. E_e distributions for the different models in $D_s \to \phi e\nu$ and $D_s \to \eta e\nu$ decays (Figures G.60–G.62).

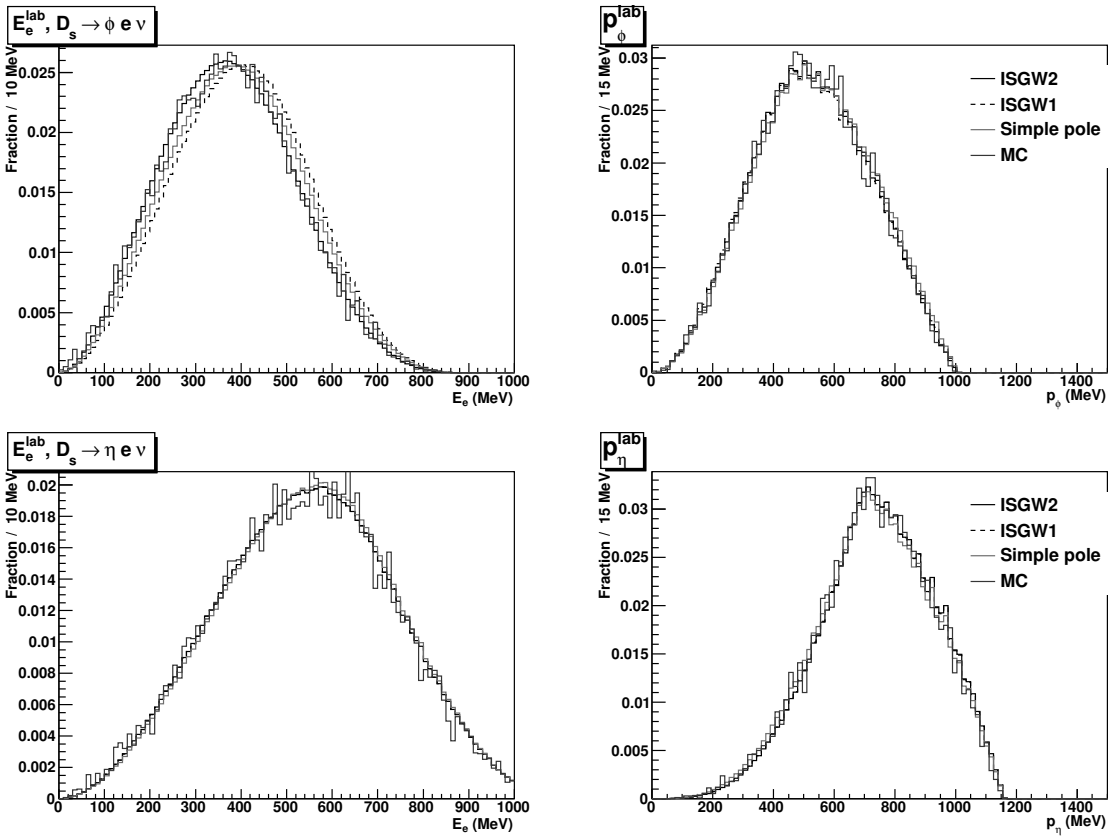

Figure 10.7: Top: Lab frame electron energy (left) and ϕ momentum (right) in $D_s \to \phi e\nu$ for the ISGW2 and pole models. The electron energy has a noticeable increase from ISGW2 to the pole model. Bottom: Lab frame electron energy and η momentum in $D_s \to \eta e\nu$. The decay to a pseudoscalar has a smaller but still positive electron energy shift.

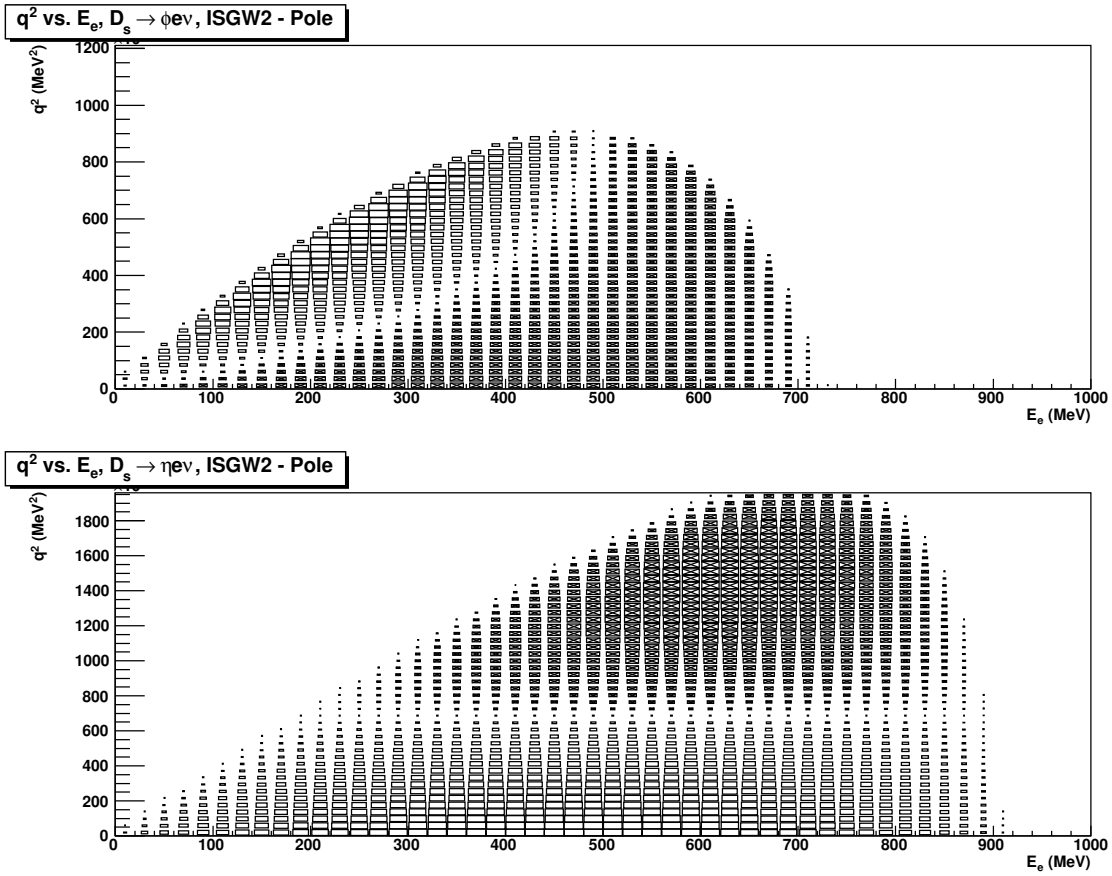

Figure 10.8: Difference between the ISGW2 and pole models in the q^2 and electron energy distributions for $D_s \to \phi e\nu$ (top) and $D_s \to \eta e\nu$ (bottom). Empty boxes indicate a surplus in the ISGW2 model, while darkened boxes with an "x" indicate a surplus in the simple pole model. The pole model has higher electron energies in both cases, although it has higher q^2 values for the pseudoscalar η decay and lower q^2 values for the vector ϕ decay.

10.15 Sum of Systematic Errors

Each D_s semileptonic decay mode has its own collection of systematic biases and errors, as described earlier in this section. Our extra tables section contains a full systematic error and bias listing for each mode. The systematic bias corrections can be found in Tables F.18-F.23, while the relative systematic errors can be found in Tables F.24-F.29. We've included a summary of our total systematics here, with Table 10.17 giving the efficiency corrections from biasing effects and Table 10.18 giving each mode's relative systematic errors.

Table 10.17: Efficiency for each D_s semileptonic mode before and after corrections from systematic biases. These efficiencies include the hadronic branching ratio (taking $\mathcal{B}(f_0 \to \pi\pi) = 52\%$ for $f_0 e\nu$).

Semileptonic mode	$\varepsilon_{\mathrm{MC}}$	$\varepsilon_{\mathrm{corrected}}$
$\phi e\nu$	13.9%	12.5%
$\eta e\nu$	20.6%	20.4%
$\eta' e\nu$	4.0%	3.8%
$f_0 e\nu$	21.7%	21.2%
$K_s e\nu$	30.9%	27.4%
$K^* e\nu$	24.1%	23.0%

Table 10.18: Total systematic errors (relative) for each D_s semileptonic decay mode.

Semileptonic mode	Relative systematic error
$\phi e\nu$	4.46%
$\eta e\nu$	8.70%
$\eta' e\nu$	10.11%
$f_0 e\nu$	4.91%
$K_s e\nu$	8.56%
$K^* e\nu$	7.13%

In five of our six modes, the statistics available limits our branching ratio measurement rather than our systematic errors. Even the one exception, $D_s \to \eta e\nu$, effectively gets restricted by statistics because that's the limiting factor in the dominant $D_s \to \rho + \eta$

systematic. Our largest required efficiency bias corrections come in $D_s \to \phi e\nu$, driven by soft kaon track reconstruction (true either with our custom kaon selections or with the default CLEO cuts), and in $D_s \to K_s e\nu$, driven by high momentum K_s reconstruction.

Incorporating our efficiency corrections and systematic errors into our measurements gives us the final branching ratio results shown in Table 10.19.

Table 10.19: Branching ratios for each D_s semileptonic mode before and after our systematic biases and errors.

Decay mode	Raw $\mathcal{B}$	Corrected $\mathcal{B}$
$D_s \to \phi e\nu$	$(1.92 \pm 0.15)\%$	$(2.14 \pm 0.17 \pm 0.09)\%$
$D_s \to \eta e\nu$	$(2.25 \pm 0.14)\%$	$(2.28 \pm 0.14 \pm 0.20)\%$
$D_s \to \eta' e\nu$	$(0.64 \pm 0.14)\%$	$(0.68 \pm 0.15 \pm 0.06)\%$
$D_s \to f_0 e\nu, f_0 \to \pi\pi$	$(0.13 \pm 0.02)\%$	$(0.13 \pm 0.02 \pm 0.01)\%$
$D_s \to K_s e\nu$	$(0.17 \pm 0.03)\%$	$(0.20 \pm 0.04 \pm 0.01)\%$
$D_s \to K^* e\nu$	$(0.17 \pm 0.04)\%$	$(0.18 \pm 0.04 \pm 0.01)\%$

Chapter 11

Conclusion

D_s semileptonic decays have seen relatively few measurements over the years despite their use in lattice calculations, light meson spectroscopy, and comparisons to other mesons' semileptonic decays. CLEO-c's 4170 MeV run has given us an opportunity to improve the exclusive D_s semileptonic measurements for the six dominant modes in a fairly clean environment. Low backgrounds made maximizing the signal a priority, so Table 11.1 summarizes the number of signal events we obtained over CLEO-c's integrated luminosity of 586 pb^{-1}.

Table 11.1: Number of observed signal events for each of our six semileptonic modes. We include the branching ratios with their statistical errors for reference.

Semileptonic mode	$\mathcal{B}$	Signal Events
$D_s \to \phi e\nu$	$(2.14 \pm 0.17)\%$	206.7 ± 16.4
$D_s \to \eta e\nu$	$(2.28 \pm 0.14)\%$	358.2 ± 21.6
$D_s \to \eta' e\nu$	$(0.68 \pm 0.15)\%$	20.1 ± 4.4
$D_s \to f_0 e\nu, f_0 \to \pi\pi$	$(0.13 \pm 0.03)\%$	41.9 ± 7.8
$D_s \to K_s e\nu$	$(0.20 \pm 0.04)\%$	41.5 ± 8.3
$D_s \to K^* e\nu$	$(0.18 \pm 0.04)\%$	31.6 ± 7.5

Table 11.10 in Section 11.2 contains our full results, including all systematics. Section 11.2 also contains a synopsis of this D_s semileptonic analysis.

11.1 Comparisons

Section 1.4 and Section 1.5 in our introduction contained theoretical predictions from the literature and a summary of prior D_s semileptonic measurements, respectively. Here, we compare our results to those predictions and measurements with a brief discussion.

11.1.1 Theory

Our six D_s semileptonic modes each provide tests for different theoretical models and parameters. In the following section, we've grouped together modes that explore similar terrain and give a short discussion of how our measurements fit into the various predictions.

In the spectator model, D_s semileptonic decays to the η and η' pseudoscalar states probe the η/η' mixing angle by coupling to the mesons' $s\bar{s}$ component. The possibility of weak annihilation (Figure 1.2) or η/η' glueball mixing complicates this interpretation, but an explicit measurement should still give a first order estimate for the mixing angle. Table 11.2 compares the ranges from six different predictions to our branching ratios for η, η', and their ratio.

Table 11.2: Theoretical prediction range for $\mathcal{B}(D_s \to \eta e\nu)$ and $\mathcal{B}(D_s \to \eta' e\nu)$ compared to our measurements. Our η measurement matches four of the six predictions, while our measured η' falls below the expected values.

	$\mathcal{B}(D_s \to \eta' e\nu)$	$\mathcal{B}(D_s \to \eta e\nu)$	$\frac{\mathcal{B}(D_s \to \eta' e\nu)}{\mathcal{B}(D_s \to \eta e\nu)}$
Predictions [44, 46, 42, 43, 45, 34]	(0.43 – 1.1)%	(1.4 – 3.2)%	0.31 – 0.43
This analysis	$(0.68 \pm 0.16)\%$	$(2.28 \pm 0.24)\%$	0.30 ± 0.07

Our $\mathcal{B}(D_s \to \eta e\nu)$ matches all but the two extreme predictions (1.4% and 3.2%), while our $\mathcal{B}(D_s \to \eta' e\nu)$ measurement falls about 1.5σ below all but the smallest prediction. Consequently, we see a lower η'/η ratio than any of our predictions, although our large error on the ratio covers most of the predicted range. Even at 1.5σ, our η'/η ratio comes within the range of reasonable prediction and measurement error. More exotically but less likely, a low η'/η ratio could indicate some glue mixing or a mixing

angle smaller than $\phi \approx 40\,^\circ$.

Since the ϕ meson contains only an $s\bar{s}$ component, $\mathcal{B}(D_s \to \phi e\nu)$ creates a good testing ground for a variety of theoretical approaches. Our measurement matches the ISGW2 prediction [34] almost exactly, and it matches a basic lattice calculation [26] to within combined error. Our measurement lies about 2.5σ and 4σ from the other two models' predictions (a constituent quark model constrained by lattice results [42] and an approach using QCD sum rules [48]). Table 11.3 shows how our measurement compares to the full range of predictions.

Table 11.3: Theoretical prediction range for $\mathcal{B}(D_s \to \phi e\nu)$ compared to our measurement. Our result matches the ISGW2 and lattice predictions from Section 1.4.2.

	$\mathcal{B}(D_s \to \phi e\nu)$
Predictions [48, 34, 26, 42]	(1.4 – 2.6)%
This analysis	$(2.14 \pm 0.19)\%$

Like η and η', the $f_0(980)$ composition has an $s\bar{s}$ component along with a component from lighter quarks (or even a possible gluonic contribution [28]). By comparing to the predictions from various models, we can estimate a reasonable range for the $s\bar{s}$ mixing angle. If we take $f_0 = \cos\theta\,|s\bar{s}\rangle + \sin\theta\,|\text{other}\rangle$ and use the predicted branching ratio range in Table 11.4, we get an $s\bar{s}$ mixing angle, θ, of 30°–50° for the f_0.

Table 11.4: Theoretical prediction range for $\mathcal{B}(D_s \to f_0 e\nu)$ compared to our measurement. We have assumed $\mathcal{B}(f_0 \to \pi\pi) = 0.52$ for the comparison. Our result with these predictions implies an $s\bar{s}$ mixing angle (θ) in f_0 of 30°–50°.

	$\mathcal{B}(D_s \to f_0 e\nu)$
Predictions [50, 49, 48]	$(0.41 - 0.55)\% \times \cos^2\theta$
This analysis	$(0.26 \pm 0.05)\%$

Our $\mathcal{B}(D_s \to K_s e\nu)$ and $\mathcal{B}(D_s \to K^* e\nu)$ results match one of the two predictions we consider to within 1σ (a constituent quark model with lattice input [42]). The absolute $K_s e\nu$ and $K^* e\nu$ branching ratio measurements each fall 2σ below the ISGW2 prediction [34], but the ratio between our two results comes within 0.3σ of ISGW2's predicted ratio. Table 11.5 gives our measurements compared to the range of predictions,

including an extension of a straight-forward kinematic model [46] that only yields a ratio prediction (about 2σ different from our measurement and a similar distance from the other predictions).

Table 11.5: Theoretical prediction range for $\mathcal{B}(D_s \to K_s e\nu)$ and $\mathcal{B}(D_s \to K^* e\nu)$ compared to our measurements. Our $K_s e\nu / K^* e\nu$ ratio matches two of three predictions closely, but it differs from a simple model by 2σ.

	$\mathcal{B}(D_s \to K_s e\nu)$	$\mathcal{B}(D_s \to K^* e\nu)$	$\frac{\mathcal{B}(D_s \to K_s e\nu)}{\mathcal{B}(D_s \to K^* e\nu)}$
Predictions [46, 42, 34]	(0.115 – 0.16)%	(0.11 – 0.19)%	0.45 – 1.0
This analysis	$(0.20 \pm 0.04)\%$	$(0.18 \pm 0.04)\%$	1.10 ± 0.35

11.1.2 Previous Measurements

The Particle Data Group [21] currently uses three different sets of measurements for their exclusive D_s semileptonic branching ratios. One set comes from BaBar [29], which used charm continuum production when running at the $\Upsilon(4S)$ resonance. They reconstructed a $D_s \to \phi\pi$ tag and used the two jets formed during $c\bar{c}$ hadronization to isolate $D_s \to KKe\nu$ on the other side. Aside from detector differences, the backgrounds involved in the BaBar analysis differ significantly from ours. We thus have a statistically and systematically independent measurement from BaBar.

However, the other two sets of measurements both come from CLEO-c [51, 52]. One set measures $D_s \to \phi e\nu$ and $D_s \to f_0 e\nu$ over the same sample that we use, while the other set measures all modes but over only half the sample. The other two CLEO measurements use roughly the same technique, which differs significantly from our own.

The main difference between the other CLEO-c analyses and our analysis lies in the D_s^* daughter photon reconstruction. Spurious soft photons in the event create a peaking background on the event's missing mass. In our analysis, we avoid this problem by dropping the photon entirely, sacrificing the event's zero missing mass in the process (but gaining statistics from the photon reconstruction efficiency).

The other two analyses retain the photon reconstruction but only keep the best candidate photon. The branching ratio denominator then becomes a $D_s+\gamma$ tag (determined by a 2D $M_{D_s}/M^{\text{recoil}}_{D_s+\gamma}$ fit). This combination requires those analyses to either treat false

γ tags as valid for the branching ratio denominator, or it requires the analyses to reject the soft, peaking background from false γ events in the branching ratio numerator (the event missing mass). The analyses take the latter approach, using the Monte Carlo to estimate the false γ rate and deducting that from the signal.

In addition to the major analysis difference in D_s^* daughter photon reconstruction, the other two CLEO-c analyses also have other, minor differences from our own. They use 9 D_s tagging modes, while we use 13 modes. Eight of the modes overlap closely, but we have looser constraints on $K_sK^-\pi\pi$, add three more K_s modes ($K_sK\pi^0$, $K_sK_s\pi$, $K_sK^+\pi\pi$), and add $\pi\pi^0\eta', \eta' \to \pi\pi\eta$. We also try to take advantage of the low background by using looser semileptonic hadron cuts.

Table 11.6 shows our $\mathcal{B}(D_s \to \phi e\nu)$ and $\mathcal{B}(D_s \to f_0 e\nu)$ measurements compared to both the BaBar and prior CLEO measurements. BaBar obtains a much larger $\mathcal{B}(D_s \to \phi e\nu)$ than we see, with their systematic error limiting their measurement. We obtain a lower branching ratio (statistically significant) for both $\phi e\nu$ and $f_0 e\nu$ than the alternate CLEO measurement [51], with improved error in $\phi e\nu$ due to our higher statistics. Our $\phi e\nu$ and $f_0 e\nu$ measurements do match the previous CLEO measurements [52] that used the same general technique as the newest incarnation but involved half the data sample.[1]

Table 11.6: We find lower $\mathcal{B}(D_s \to \phi e\nu)$ and $\mathcal{B}(D_s \to f_0 e\nu)$ than prior measurements. BaBar observed $D_s \to f_0 e\nu$ using $f_0 \to KK$.

Experiment	$\mathcal{B}(D_s \to \phi e\nu)$	$\mathcal{B}(D_s \to f_0 e\nu, f_0 \to \pi\pi)$
BaBar [29]	$(2.61 \pm 0.03 \pm 0.17)\%$	Seen
CLEO [51]	$(2.36 \pm 0.23 \pm 0.13)\%$	$(0.20 \pm 0.03 \pm 0.01)\%$
This analysis	$(2.14 \pm 0.17 \pm 0.09)\%$	$(0.13 \pm 0.03 \pm 0.01)\%$

We see lower $\mathcal{B}(D_s \to \eta e\nu)$ and $\mathcal{B}(D_s \to \eta' e\nu)$ than the previous measurements by CLEO over half the data sample, as shown in Table 11.7. The different systematics for $\eta e\nu$ and the low statistics in $\eta' e\nu$ do make the difference in measurements reasonable. We obtain a lower statistical error in both cases but a larger systematic error. The applied

[1] We do not include the prior CLEO measurements in the table as the Particle Data Group does not use them. For completeness, they saw $\mathcal{B}(D_s \to \phi e\nu) = 2.29 \pm 0.37 \pm 0.11$ and $\mathcal{B}(D_s \to f_0 e\nu) = 0.13 \pm 0.04 \pm 0.01$.

η systematic (Section 10.6.2) dominates the difference between systematic errors in the two analyses.

The prior CLEO measurement adopts a 2% relative systematic for η reconstruction, based simply on a 1% relative systematic for each photon reconstruction. After considering CLEO's standard 5.9% relative systematic on η reconstruction with its -5.6% relative efficiency correction [71][2] we decided to do our own η systematic as described in Appendix C. This gave us our 7.9% relative systematic for η reconstruction, leading to a much larger final systematic error on $\mathcal{B}(D_s \to \eta e\nu)$ and $\mathcal{B}(D_s \to \eta' e\nu)$ than the other analysis.

Table 11.7: Our $\mathcal{B}(D_s \to \eta e\nu)$ and $\mathcal{B}(D_s \to \eta' e\nu)$ both fall below the prior measurement using half the CLEO-c data. We also differ with the previous experiment on the reconstruction systematic for the final state η in both modes.

Experiment	$\mathcal{B}(D_s \to \eta e\nu)$	$\mathcal{B}(D_s \to \eta' e\nu)$
CLEO [52]	$(2.48 \pm 0.29 \pm 0.13)\%$	$(0.91 \pm 0.33 \pm 0.05)\%$
This analysis	$(2.28 \pm 0.14 \pm 0.19)\%$	$(0.68 \pm 0.15 \pm 0.06)\%$

As Table 11.8 indicates, our $\mathcal{B}(D_s \to K_s e\nu)$ and $\mathcal{B}(D_s \to K^* e\nu)$ measurements almost exactly match CLEO's prior result. By virtue of having more data available, this analysis's measurement has lower statistical error while having essentially the same systematic error.

Table 11.8: Our $\mathcal{B}(D_s \to K_s e\nu)$ and $\mathcal{B}(D_s \to K^* e\nu)$ almost exactly match the previous CLEO-c result on half the data sample.

Experiment	$\mathcal{B}(D_s \to K_s e\nu)$	$\mathcal{B}(D_s \to K^* e\nu)$
CLEO [52]	$(0.19 \pm 0.05 \pm 0.01)\%$	$(0.18 \pm 0.07 \pm 0.01)\%$
This analysis	$(0.20 \pm 0.04 \pm 0.01)\%$	$(0.18 \pm 0.04 \pm 0.01)\%$

In addition to the exclusive measurements, CLEO-c has measured inclusive semileptonic decays [30]. Table 11.9 shows that our six exclusive modes cover most of the D_s semileptonic spectrum. Roughly 14% of D_s semileptonic decays may go to nonresonant

[2]Determined from $\psi' \to \eta J/\psi$ with a momentum extrapolation based on similar π^0

or other modes, although the difference between the inclusive rate and the exclusive sum has a large associated error. Only 1.6σ separates the inclusive/exclusive difference.

Table 11.9: Comparison of CLEO-c's inclusive D_s semileptonic branching ratio with the sum of this analysis's six branching ratios. The difference shows some room for other semileptonic modes, with relatively large error.

	$\mathcal{B}(D_s \to he\nu)$
Inclusive [30]	$(6.52 \pm 0.39 \pm 0.15)\%$
Sum of exclusive	$(5.60 \pm 0.27 \pm 0.28)\%$
Difference	$(0.92 \pm 0.48 \pm 0.32)\%$

11.2 Summary

We've measured D_s semileptonic branching ratios for six modes: $D_s \to \phi e\nu$, $D_s \to \eta e\nu$, $D_s \to \eta' e\nu$, $D_s \to f_0 e\nu$, $D_s \to K_s e\nu$, and $D_s \to K^* e\nu$. We primarily used CLEO-c's $D_s^* D_s$ sample, where the D_s^* generally decays to a D_s via a soft photon emission. One D_s gave us a tag to identify the event, while the other became a candidate for our target semileptonic decays.

As the soft D_s^* photon had a low reconstruction efficiency and dubious background predictions from our Monte Carlo, we sacrificed the ability to reconstruct a zero neutrino missing mass in exchange for additional events by dropping the photon reconstruction. Since each mode showed fairly low background even without the photon, we could safely loosen our other particle cuts in the key $\phi e\nu$ and $\eta e\nu$ modes to gain further signal events. These looser cuts required new analysis for these particles' reconstruction efficiencies, but the atypically slow kaons and relatively unexplored η momentum range warranted such study in any case.

Table 11.10 states our final results for all six D_s semileptonic modes, including all statistical and systematic errors.

Table 11.10: This analysis's measured branching ratios for each D_s semileptonic mode.

Decay Mode	Branching ratio
$D_s \to \phi e\nu$	$(2.139 \pm 0.170 \pm 0.086)\%$
$D_s \to \eta e\nu$	$(2.277 \pm 0.137 \pm 0.196)\%$
$D_s \to \eta' e\nu$	$(0.680 \pm 0.150 \pm 0.064)\%$
$D_s \to f_0 e\nu, f_0 \to \pi\pi$	$(0.133 \pm 0.025 \pm 0.006)\%$
$D_s \to K_s e\nu$	$(0.196 \pm 0.039 \pm 0.015)\%$
$D_s \to K^* e\nu$	$(0.178 \pm 0.042 \pm 0.012)\%$

References

[1] L. Lederman and D. Teresi, "The God particle: If the universe is the answer, what is the question?," Boston, USA: Houghton Mifflin (1993) 434 p

[2] J. C. Street and E. C. Stevenson, Phys. Rev. **52**, 1003 (1937).

[3] G. D. Rochester and C. C. Butler, Nature **160**, 855 (1947).

[4] V. D. Hopper and S. Biswas, Phys. Rev. **80**, 1099 (1950).

[5] E. Fermi, Nuovo Cim. **11**, 1 (1934).

[6] S. L. Glashow, Nucl. Phys. **22**, 579 (1961).

[7] P. W. Higgs, Phys. Rev. Lett. **13**, 508 (1964).

[8] S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967).

[9] A. Salam, Conf. Proc. C **680519**, 367 (1968).

[10] M. Gell-Mann, Phys. Lett. **8**, 214 (1964).

[11] G. Zweig, CERN-TH-401.

[12] J. D. Bjorken and S. L. Glashow, Phys. Lett. **11**, 255 (1964).

[13] S. L. Glashow, J. Iliopoulos and L. Maiani, Phys. Rev. D **2**, 1285 (1970).

[14] E. D. Bloom, D. H. Coward, H. C. DeStaebler, J. Drees, G. Miller, L. W. Mo, R. E. Taylor and M. Breidenbach *et al.*, Phys. Rev. Lett. **23**, 930 (1969).

[15] J. D. Bjorken and E. A. Paschos, Phys. Rev. **185**, 1975 (1969).

[16] D. J. Gross and F. Wilczek, Phys. Rev. Lett. **30**, 1343 (1973).

[17] H. D. Politzer, Phys. Rev. Lett. **30**, 1346 (1973).

[18] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).

[19] M. L. Perl, G. S. Abrams, A. Boyarski, M. Breidenbach, D. Briggs, F. Bulos, W. Chinowsky and J. T. Dakin *et al.*, Phys. Rev. Lett. **35**, 1489 (1975).

[20] H. -K. Quang and X. -Y. Pham, "Elementary particles and their interactions: concepts and phenomena," Berlin, Germany: Springer-Verlag, 1998. 661 p.

[21] J. Beringer *et al.* [Particle Data Group Collaboration], Phys. Rev. D **86**, 010001 (2012).

[22] K. G. Wilson, Phys. Rev. D **10**, 2445 (1974).

[23] J. Koponen, arXiv:1311.6931 [hep-lat].

[24] R. Gupta, hep-lat/9807028.

[25] A. Bazavov, D. Toussaint, C. Bernard, J. Laiho, C. DeTar, L. Levkova, M. B. Oktay and S. Gottlieb *et al.*, Rev. Mod. Phys. **82**, 1349 (2010) [hep-lat/0903.3598].

[26] G. C. Donald, C. T. H. Davies, J. Koponen and G. P. Lepage, arXiv:1311.6669 [hep-lat].

[27] S. Bianco, F. L. Fabbri, D. Benson and I. Bigi, Riv. Nuovo Cim. **26N7**, 1 (2003) [hep-ex/0309021].

[28] W. Ochs, J. Phys. G **40**, 043001 (2013) [arXiv:1301.5183 [hep-ph]].

[29] B. Aubert *et al.* [BaBar Collaboration], Phys. Rev. D **78**, 051101 (2008) [arXiv:0807.1599 [hep-ex]].

[30] D. M. Asner *et al.* [CLEO Collaboration], Phys. Rev. D **81**, 052007 (2010) [arXiv:0912.4232 [hep-ex]].

[31] M. B. Voloshin, Phys. Lett. B **515**, 74 (2001) [hep-ph/0106040].

[32] Z. Ligeti, M. Luke and A. V. Manohar, Phys. Rev. D **82**, 033003 (2010) [arXiv:1003.1351 [hep-ph]].

[33] N. Isgur, D. Scora, B. Grinstein and M. B. Wise, Phys. Rev. D **39**, 799 (1989).

[34] D. Scora and N. Isgur, Phys. Rev. D **52**, 2783 (1995) [hep-ph/9503486].

[35] Y. Nir, Phys. Lett. B **221**, 184 (1989).

[36] M. Wirbel, B. Stech and M. Bauer, Z. Phys. C **29**, 637 (1985).

[37] D. Becirevic and A. B. Kaidalov, Phys. Lett. B **478**, 417 (2000) [hep-ph/9904490].

[38] M. Artuso, B. Meadows and A. A. Petrov, Ann. Rev. Nucl. Part. Sci. **58**, 249 (2008) [arXiv:0802.2934 [hep-ph]].

[39] J. G. Korner and G. A. Schuler, Z. Phys. C **46**, 93 (1990).

[40] V. V. Anisovich, D. V. Bugg, D. I. Melikhov and V. A. Nikonov, Phys. Lett. B **404**, 166 (1997) [hep-ph/9702383].

[41] C. Di Donato, G. Ricciardi and I. Bigi, Phys. Rev. D **85**, 013016 (2012) [arXiv:1105.3557 [hep-ph]].

[42] D. Melikhov and B. Stech, Phys. Rev. D **62**, 014006 (2000) [hep-ph/0001113].

[43] Z. -T. Wei, H. -W. Ke and X. -F. Yang, Phys. Rev. D **80**, 015022 (2009) [arXiv:0905.3069 [hep-ph]].

[44] K. Azizi, R. Khosravi and F. Falahati, J. Phys. G **38**, 095001 (2011) [arXiv:1011.6046 [hep-ph]].

[45] P. Colangelo and F. De Fazio, Phys. Lett. B **520**, 78 (2001) [hep-ph/0107137].

[46] M. Gronau and J. L. Rosner, Phys. Rev. D **83**, 034025 (2011) [arXiv:1012.5098 [hep-ph]].

[47] S. Collins, I. Kanamori and J. Najjar, arXiv:1311.7393 [hep-lat].

[48] I. Bediaga and M. Nielsen, Phys. Rev. D **68**, 036001 (2003) [hep-ph/0304193].

[49] H. -W. Ke, X. -Q. Li and Z. -T. Wei, Phys. Rev. D **80**, 074030 (2009) [arXiv:0907.5465 [hep-ph]].

[50] T. M. Aliev and M. Savci, Europhys. Lett. **90**, 61001 (2010) [hep-ph/0701108].

[51] K. M. Ecklund *et al.* [CLEO Collaboration], Phys. Rev. D **80**, 052009 (2009) [arXiv:0907.3201 [hep-ex]].

[52] J. Yelton *et al.* [CLEO Collaboration], Phys. Rev. D **80**, 052007 (2009) [arXiv:0903.0601 [hep-ex]].

[53] B. Aubert *et al.* [BaBar Collaboration], hep-ex/0607085.

[54] R. A. Briere *et al.* [CLEO Collaboration], CLNS-01-1742.

[55] J. A. Crittenden, Conf. Proc. C **060626**, 992 (2006).

[56] K. Berkelman, River Edge, USA: World Scientific (2004) 151 p

[57] D. Peterson, K. Berkelman, R. A. Briere, G. Chen, D. Cronin-Hennessy, S. Csorna, M. Dickson and S. von Dombrowski *et al.*, Nucl. Instrum. Meth. A **478**, 142 (2002).

[58] Y. Kubota *et al.* [CLEO Collaboration], Nucl. Instrum. Meth. A **320**, 66 (1992).

[59] K. K. Gan, C. M. Daubenmier, H. Kagan and R. Kass, In *Brussels 1995, High energy physics* 607-608

[60] CLEO Collaboration internal document, B. Lang, Y. Kubota and R. Poling CBX 07-24 (2007).

[61] D. Cronin-Hennessy *et al.* [CLEO Collaboration], Phys. Rev. D **80**, 072001 (2009) [arXiv:0801.3418 [hep-ex]].

[62] A. Ryd, D. Lange, N. Kuznetsova, S. Versille, M. Rotondo, D. P. Kirkby, F. K. Wuerthwein and A. Ishikawa, EVTGEN-V00-11-07.

[63] DSkim web page:
`https://www.lepp.cornell.edu/ dskim/private/dskim.html`

[64] CLEO Collaboration internal document, P. Onyisi and W. Sun, CBX 06-11 (2006).

[65] CLEO Collaboration internal document, P. Onyisi and A. Ryd, CBX 07-14 (2007).

[66] CLEO Collaboration internal document, R. Briere and G. Chen, CBX 07-30 (2005).

[67] EID web page:
`https://www.lepp.cornell.edu/restricted/CLEO/CLEO3/soft/hints/EID.html`

[68] CLEO Collaboration internal document, I. Brock and R. Briere, CBX 05-43 (2005).

[69] CLEO Collaboration internal document, K. Randrianarivony and M. Artuso, CBX 2008-050 (2008).

[70] H. Ha *et al.* [BELLE Collaboration], Phys. Rev. D **83**, 071101 (2011) [arXiv:1012.0090 [hep-ex]].

[71] CLEO Collaboration internal document, P. Onyisi, CBX 07-15 (2007).

[72] CLEO Collaboration internal document, S. Stroiney, A. Ryd, W. Sun, P. Onyisi and D. Cassel,, CBX 2008-040 (2008).

[73] CLEO Collaboration internal document, J. Ge and I. Shipsey, CBX 2009-009 (2009).

[74] CLEO Collaboration internal document, C. Park and E. Thorndike, CBX 2008-049 (2008).

[75] CLEO Collaboration internal document, S. Stroiney, A. Ryd, W. Sun, P. Onyisi and D. Cassel, CBX 2008-041 (2008).

[76] J. D. Richman and P. R. Burchat, Rev. Mod. Phys. **67**, 893 (1995) [hep-ph/9508250].

[77] J. P. Alexander *et al.* [CLEO Collaboration], hep-ex/0007052.

[78] K. M. Ecklund, PoS hf **8**, 024 (1999) [hep-ex/9912034].

[79] S. M. Flatte, Phys. Lett. B **63**, 224 (1976).

[80] N. N. Achasov, V. V. Gubin and V. I. Shevchenko, Phys. Rev. D **56**, 203 (1997) [hep-ph/9605245].

[81] S. M. Flatte, M. Alston-Garnjost, A. Barbaro-Galtieri, J. H. Friedman, G. R. Lynch, S. D. Protopopescu, M. S. Rabin and F. T. Solmitz, Phys. Lett. B **38**, 232 (1972).

[82] D. Barberis *et al.* [WA102 Collaboration], Phys. Lett. B **462**, 462 (1999) [hep-ex/9907055].

[83] K. Ackerstaff *et al.* [OPAL Collaboration], Eur. Phys. J. C **4**, 19 (1998) [hep-ex/9802013].

[84] E. M. Aitala *et al.* [E791 Collaboration], Phys. Rev. Lett. **86**, 765 (2001) [hep-ex/0007027].

[85] M. N. Achasov, S. E. Baru, K. I. Beloborodov, A. V. Berdyugin, A. V. Bozhenok, A. D. Bukin, D. A. Bukin and S. V. Burdin *et al.*, Phys. Lett. B **485**, 349 (2000) [hep-ex/0005017].

[86] R. R. Akhmetshin *et al.* [CMD-2 Collaboration], Phys. Lett. B **462**, 380 (1999) [hep-ex/9907006].

[87] A. Aloisio *et al.* [KLOE Collaboration], Phys. Lett. B **537**, 21 (2002) [hep-ex/0204013].

[88] CLEO Collaboration internal document, S. Stone and L. Zhang, CBX 2009-011

[89] P. Naik *et al.* [CLEO Collaboration], Phys. Rev. D **80**, 112004 (2009) [arXiv:0910.3602 [hep-ex]].

Appendix A

$f_0 \to KK$ Models

$f_0 \to KK$ decays present a unique problem for the $D_s \to \phi e\nu$ analysis as the f_0 resonance lies very near (and likely below) the KK threshold, yet the f_0's wide decay width extends its mass spectrum well into the ϕ mass region. Thus, $D_s \to f_0 e\nu$ events where the $f_0 \to KK$ invariant mass falls within the ϕ mass region become effectively indistinguishable from $D_s \to \phi e\nu$ events and need to be deducted from our $\phi e\nu$ signal. To remove these $f_0 e\nu$ events, we use our $D_s \to f_0 e\nu$ measurement from the $f_0 \to \pi\pi$ mode via

$$\mathcal{B}(D_s \to \phi e\nu)^{\text{correction}} = \mathcal{B}(D_s \to f_0 e\nu, f_0 \to \pi\pi) * \frac{\mathcal{B}(f_0 \to K^+K^-)}{\mathcal{B}(f_0 \to \pi^+\pi^-)} * \frac{\varepsilon_{f_0 \to KK}}{\varepsilon_{\phi \to KK}} * \mathrm{f}_{\text{window}}, \tag{A.1}$$

where $\mathrm{f}_{\text{window}}$ represents the fraction of $f_0 \to KK$ decays with an invariant mass in the ϕ mass cut window (-15 MeV $< M_\phi^{\text{recon}} - M_\phi^{\text{PDG}} < 30$ MeV) and $\varepsilon_{f_0 \to KK}$ is the reconstruction efficiency for f_0 within the mass window.

The fraction of $f_0 \to KK$ decays that fall within our ϕ mass window ($\mathrm{f}_{\text{window}}$) depends upon the $f_0 \to KK$ mass lineshape. Regardless of the underlying model, this lineshape will necessarily depend upon parameters such as the f_0 resonant mass and width, which have large uncertainties from previous measurements [21]. Unfortunately, `EvtGen` does not describe this lineshape in a consistent manner for resonant masses across the KK mass threshold, as described below, so we have used our own model based on a Flatté parametrization [79, 80].

The Particle Data Book estimates the f_0 physical mass at 980 MeV $\pm$ 10 MeV [21], which extends both above and below the K^+K^- threshold near 987.4 MeV. We chose to use the central mass (980 MeV) as the input value for our $\mathcal{B}(D_s \to \phi e\nu)$ correction from $D_s \to f_0 e\nu$, $f_0 \to KK$, then we estimate one systematic error on the correction by varying the physical mass throughout the 1σ mass range (10 MeV). We use a similar method to vary the total and partial f_0 widths. However, since the Particle Data Book does not give a central value for either, we chose 50 MeV as a reasonable value for the total width (with 40 MeV to 100 MeV as our 1σ systematic variation) and 0.80 as the central value for $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$ (with 0.52 to 0.82 for our 1σ systematic variation). We vary each of these values independently in our systematic. This may not be entirely appropriate since the three different f_0 parameter values contributed by each experiment are correlated, but we find it prohibitively time-consuming and of marginal benefit to disentangle each experiment's correlations (if even possible without delving into unpublished results).

In section A.1, we discuss the models available in `EvtGen` that we have chosen not to use, as the information may prove useful to others using `EvtGen` or similar software. We discuss the Flatté model that we instead use, along with its results, in section A.2.

A.1 `EvtGen` Models

A.1.1 Default Model (Breit-Wigner)

In the CLEO Monte Carlo, the mass and width of a particle can be altered by changing its values from the `evt.pdl` file and passing the modified `pdl` file to `EvtGen`. However, `EvtGen` switches the generating model used when the f_0 resonant mass sits below the KK mass threshold from its model for f_0 above the KK mass threshold. Specifically, `EvtGen` uses a non-relativistic Breit-Wigner (Equation A.2) for the $f_0 \to KK$ mass lineshape when the f_0 resonant mass lies below threshold, and it uses a relativistic Breit-Wigner (Equation A.3) for the lineshape when the resonant mass lies above KK threshold:

$$\frac{d\Gamma}{dm}^{\text{NR B-W}} \propto \left| \frac{1}{m - m_0 + i\frac{\Gamma_0}{2}} \right|^2 , \tag{A.2}$$

$$\frac{d\Gamma}{dm}^{\text{Rel B-W}} \propto \left| \frac{\Gamma_0 \left(\frac{p}{p_0}\right)}{(m^2 - m_0^2) + i m_0 \Gamma_0 \left(\frac{m_0}{m}\right) \left(\frac{p}{p_0}\right)} \right|^2 , \tag{A.3}$$

where Γ_0 is the f_0 width, m_0 is the resonant mass, m is the invariant KK mass, p is the daughter kaon momentum in the rest frame of m, and p_0 is the daughter kaon momentum in the rest frame of m_0.

The different lineshapes and their dependence on different f_0 masses can be seen in Figure A.1.

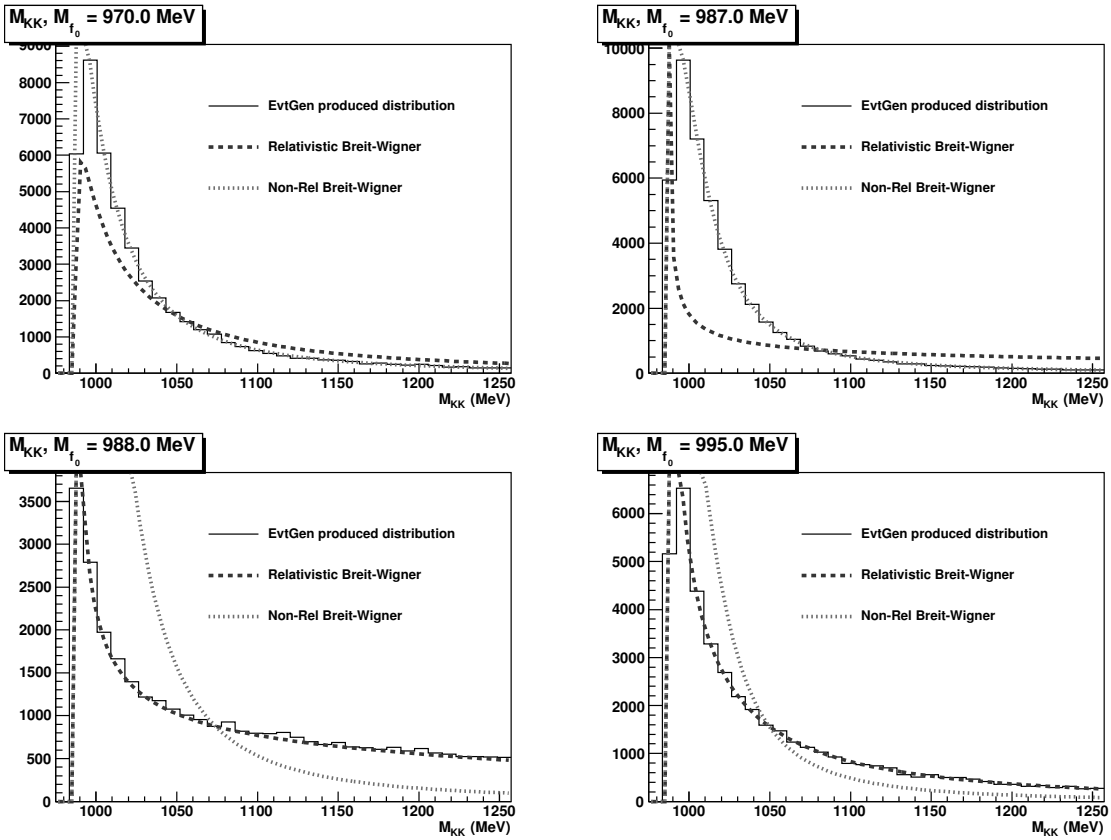

Figure A.1: `EvtGen` produced lineshape for f_0 masses above and below the KK threshold at 987.4 MeV. `EvtGen` changes its behavior from a non-relativistic Breit-Wigner to a relativistic Breit-Wigner discontinuously as the mass crosses threshold.

The discontinuous change in lineshape models as the f_0 resonant mass crosses the

KK threshold produces the dominant effect in our systematic when we try these models, with f_{window} between the relativistic and non-relativistic Breit-Wigners differing by a factor of two or more. Ultimately, we do not believe this systematic represents true variation across the threshold, nor do we think that a Breit-Wigner properly models the f_0 near threshold in any case, so we instead use a Flatté model to describe the lineshape.

A.1.2 Flatté Model

The f_0 mass lies very near (likely just below) $2m_{K^+}$, which leads to substantial threshold effects in $f_0 \to K^+K^-$ decays that a simple Breit-Wigner does not model well. One can fix the biggest issue by changing the constant width, Γ, to a momentum-dependent width. However, the f_0 requires still more work, as the opening of the KK decay mode also alters the $\pi\pi$ mass lineshape below threshold due to analyticity, with non-trivial effects for both modes. The Flatté model gives a form for the lineshape that preserves unitarity and analyticity in the threshold region [81], making it appropriate for analysis of the f_0.

`EvtGen` does have a Flatté model available, and while CLEO doesn't use it for all f_0 decays (as shown in the previous section), it does use the model for one of six resonances in the $D_s \to KK\pi$ Dalitz decay. `EvtGen`'s Flatté model for this mode uses the formula:

$$\frac{d\Gamma_{KK}}{dm}^{\text{CLEO Flatté}} \propto \left| \frac{1}{m^2 - m_0^2 + i\left(g_{KK}^2\rho_K + g_{\pi\pi}^2\rho_\pi\right)} \right|^2 \rho_3, \tag{A.4}$$

$$\begin{aligned} \rho_\pi(m) &= \sqrt{1 - \left(\frac{2m_\pi}{m}\right)^2}, \\ \rho_K(m) &= \begin{cases} \sqrt{1 - \left(\frac{2m_K}{m}\right)^2} & \text{above KK threshold,} \\ i\sqrt{\left(\frac{2m_K}{m}\right)^2 - 1} & \text{below KK threshold,} \end{cases} \end{aligned}$$

where m_0 here is the bare mass of the f_0 and $\rho_3(m)$ is the three-body phase space factor (relevant in $D_s \to KK\pi$ but not our semileptonic decay):

$$\rho_3(m) = m\rho_K\sqrt{\lambda\left(1,\left[\frac{m}{m_{D_s}}\right]^2,\left[\frac{m_\pi}{m_{D_s}}\right]^2\right)},$$
$$\lambda(a,b,c) = a^2+b^2+c^2-2ab-2bc-2ac.$$

Figure A.2 shows the $f_0 \to KK$ mass lineshape generated from this model with the default parameters $m_0 = 965$ MeV, $g_{KK} = 800$ MeV, and $g_{\pi\pi} = 406$ MeV.

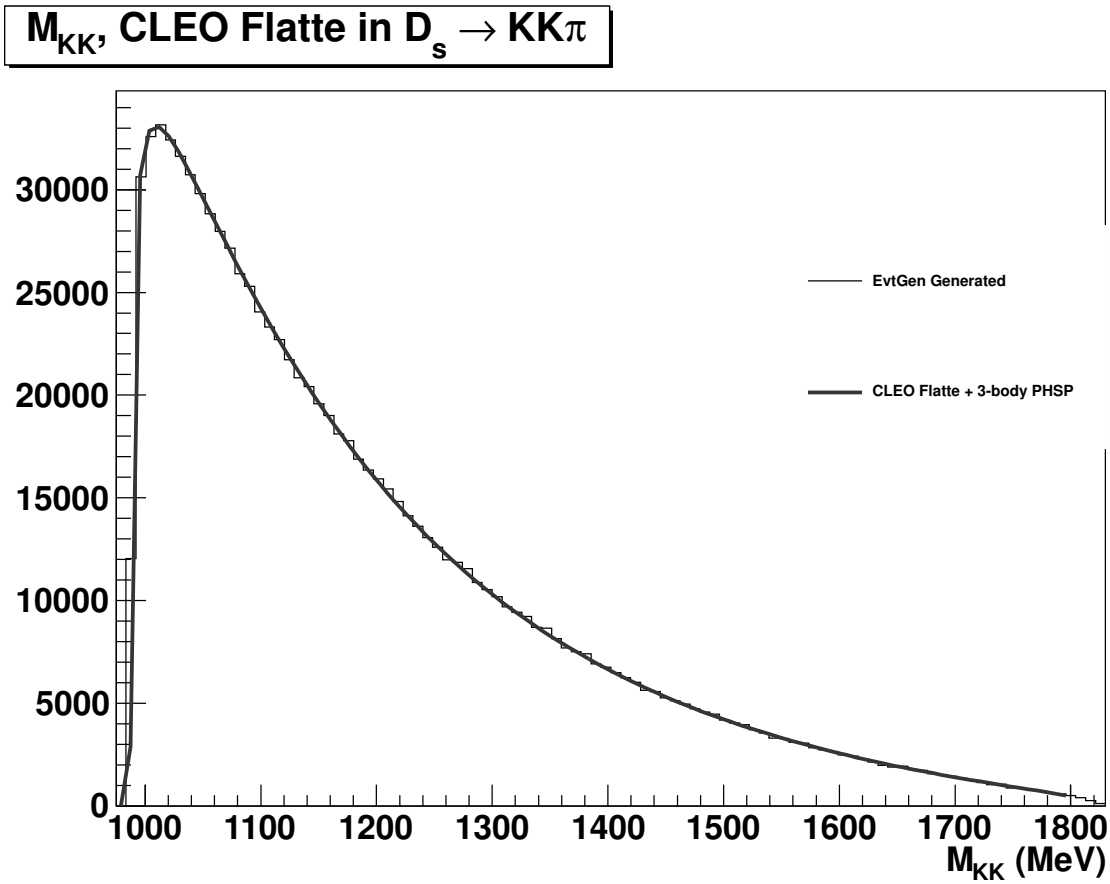

Figure A.2: CLEO Flatté mass lineshape for $f_0 \to KK$ in the decay $D_s \to KK\pi$ using the default parameters.

While this Flatté model uses reasonable parameters and gives a more sensible f_0 mass lineshape than the basic Breit-Wigner, we have chosen to use our own model for

a couple reasons. The `EvtGen` parameters are essentially hard-wired, such that we have to recompile `EvtGen` each time we want to try a new parameter set. Further, the Flatté model only incorporates the K^+K^- and $\pi^+\pi^-$ decay modes of the f_0, while the f_0 can also decay to $K^0\overline{K}^0$ (and $\pi^0\pi^0$). The $K^0\overline{K}^0$ mode becomes relevant when we do our systematic variation across the f_0 mass range listed in the Particle Data Book, as the range extends to 990 MeV— higher than $2m_{K^+}$ but below $2m_{K^0}$, leading the K-coupling to split into both real and imaginary parts.

A.2 Flatté Parametrization

Our Flatté model follows the notation of the original paper [79] with the dimensionless coupling constants g_K and g_π:

$$\frac{d\Gamma_{KK}}{dm}^{\text{Flatté}} \propto \left| \frac{\sqrt{\Gamma_0 \Gamma_K}}{m^2 - m_r^2 + i m_r \left(\Gamma_K + \Gamma_{K^0} + \Gamma_\pi + \Gamma_{\pi^0}\right)} \right|^2, \tag{A.5}$$

where

$$\begin{aligned} \Gamma_\pi(m) &= g_\pi p_\pi, \\ \Gamma_K(m) &= \begin{cases} g_K \sqrt{\left(\frac{m}{2}\right)^2 - m_K^2} & \text{above KK threshold,} \\ i g_K \sqrt{m_K^2 - \left(\frac{m}{2}\right)^2} & \text{below KK threshold,} \end{cases} \end{aligned}$$

with m_r as the resonant (bare) mass of the f_0 and Γ_0 as the f_0 width at the physical mass. Γ_{π^0} and Γ_{K^0} follow the same equations as Γ_π and Γ_K, respectively, with the appropriate mass differences and $\Gamma_{\pi^0} = \frac{1}{2} g_\pi p_{\pi^0}$ due to isospin.[1]

To get the $D_s \to f_0 e\nu$ correction on our $D_s \to \phi e\nu$ measurement from equation A.1, we need to determine both the relative amounts of $\pi^+\pi^-$ to K^+K^- and the fraction of K^+K^- that falls within our ϕ mass window. Specifically, we need to use our Flatté model to get the product $\frac{\mathcal{B}(f_0 \to K^+K^-)}{\mathcal{B}(f_0 \to \pi^+\pi^-)} \times \mathrm{f}_{\text{window}}$ for our parameters' central values and for their 1σ variations.

In our formula, we have three parameters we can vary: the bare mass; Γ_0; and the ratio of couplings, $\frac{g_K}{g_\pi}$. As stated previously, the Particle Data Book gives experimental

[1] Or due to the fundamental behavior of states with identical bosons in quantum mechanics. Whatever explanation strikes your fancy.

ranges for three related f_0 parameters: the physical mass, Γ (which we take as the width at the physical mass, Γ_0), and $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$. We can convert the physical mass, M_0, to the bare mass, m_r, using the quadratic given by

$$M_0^2 = \begin{cases} m_r^2 + m_r g_K \left[\sqrt{m_{K^+}^2 - \left(\frac{M_0}{2}\right)^2} + \sqrt{m_{K^0}^2 - \left(\frac{M_0}{2}\right)^2} \right] & \text{below } K^+K^- \text{ threshold,} \\ m_r^2 + m_r g_K \sqrt{m_{K^0}^2 - \left(\frac{M_0}{2}\right)^2} & \text{above } K^+K^- \text{ threshold.} \end{cases}$$

We vary $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$ by running across a range of different $\frac{g_K}{g_\pi}$ values, then integrating the resulting lineshapes for $\Gamma_{\pi\pi}$ and Γ_{KK}. Our default value of 0.80 stated previously for $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$ corresponds to a default $\frac{g_K}{g_\pi}$ coupling ratio of about 2, with a 1σ range from about 1.8 to 9.2. This range also covers most of the results given from experiments in the literature [82, 83], although not quite all [84].

We find that $\frac{g_K}{g_\pi}$, Γ_0, and M_0 are weakly correlated in their effect on $\frac{\mathcal{B}(f_0\to K^+K^-)}{\mathcal{B}(f_0\to\pi^+\pi^-)} \times \mathrm{f_{window}}$, so we get our final systematic by adding the results of each variation in quadrature. We give the default value of $\frac{\mathcal{B}(f_0\to K^+K^-)}{\mathcal{B}(f_0\to\pi^+\pi^-)} \times \mathrm{f_{window}}$ and the extreme values for each parameter variation in Table A.1. The mass lineshapes for each f_0 mode are shown with the same parameter values in Figure A.3 and Figure A.4.

Table A.1: f_0 parameter variations used to determine our $f_0 \to KK$ correction in $D_s \to \phi e\nu$. Our variations correspond to the PDG ranges for the physical mass, Γ_0, and $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$. In practice, we vary $\frac{g_K}{g_\pi}$ instead of directly varying $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$ since $\frac{g_K}{g_\pi}$ has less correlation with the mass and Γ_0. We use $f \times \frac{\mathcal{B}_{KK}}{\mathcal{B}_{\pi\pi}}$ as shorthand for $f_{\mathrm{window}} \times \frac{\mathcal{B}(f_0\to K^+K^-)}{\mathcal{B}(f_0\to\pi^+\pi^-)}$.

Physical mass (MeV)	Γ_0 (MeV)	$\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$	Bare mass (MeV)	g_K	$\frac{g_K}{g_\pi}$	f_{window}	$f \times \frac{\mathcal{B}_{KK}}{\mathcal{B}_{\pi\pi}}$
980	50	0.796	969.6	0.142	2	0.281	0.0587
990	50	0.733	986.4	0.140	2	0.314	0.0939
970	50	0.834	955.5	0.143	2	0.257	0.0411
980	79	0.749	963.6	0.224	2	0.233	0.0622
980	40	0.818	971.7	0.113	2	0.299	0.0547
980	50	0.523	933.2	0.652	9.196	0.187	0.1355
980	50	0.811	970.6	0.128	1.804	0.286	0.0544

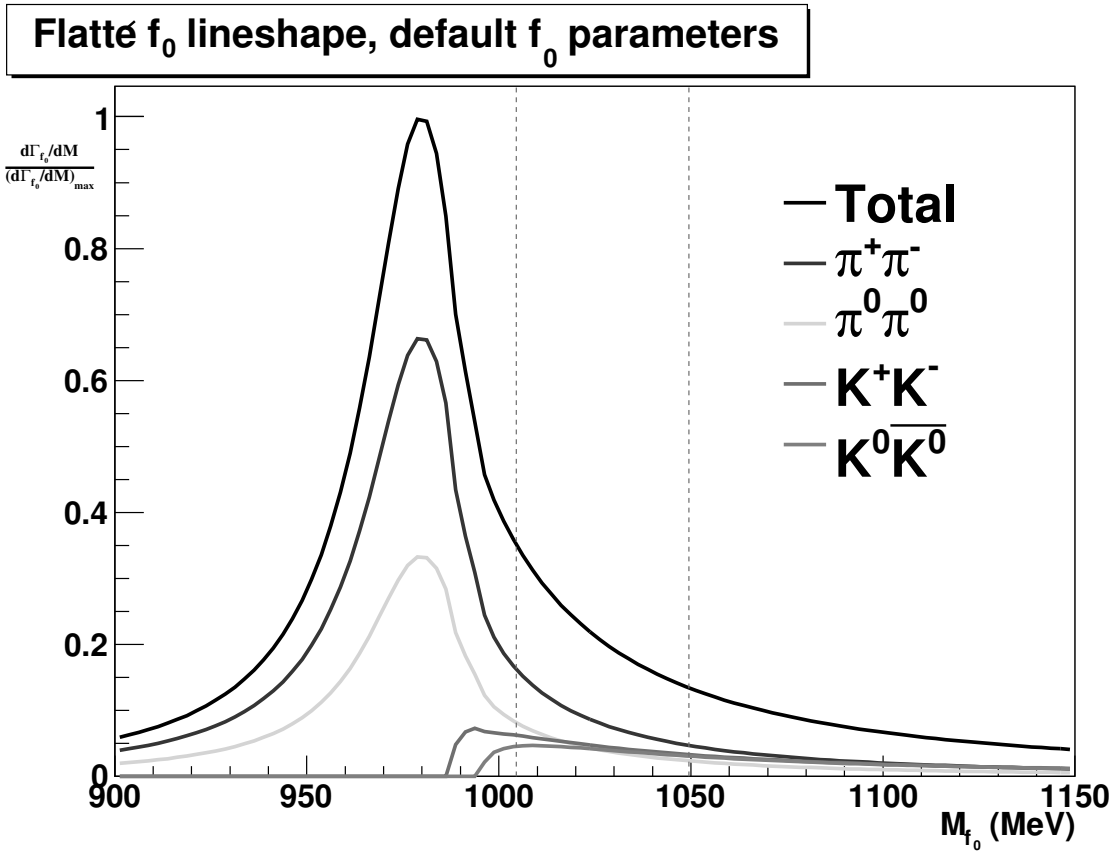

Figure A.3: f_0 mass lineshapes for $M_0 = 980$ MeV, $\frac{g_K}{g_\pi} = 2$, and $\Gamma_0 = 50$ MeV. The dotted lines indicate our ϕ mass window.

The extreme values of $\frac{\mathcal{B}(f_0 \to K^+K^-)}{\mathcal{B}(f_0 \to \pi^+\pi^-)} \times \mathrm{f_{window}}$ correspond to the extreme values of our systematic range for the physical mass and for $\frac{\Gamma_{\pi\pi}}{\Gamma_{\pi\pi}+\Gamma_{KK}}$. However, the maximum $\frac{\mathcal{B}(f_0 \to K^+K^-)}{\mathcal{B}(f_0 \to \pi^+\pi^-)} \times \mathrm{f_{window}}$ occurs in the middle of our Γ_0 range (79 MeV). This maxima remains even if we extend our possible Γ_0 up to 200 MeV, to match some values found in the literature [85, 86, 87].

We present our final correction to the $D_s \to \phi e\nu$ branching ratio after combining the variations in M_0, $\frac{g_K}{g_\pi}$, and Γ_0 with the uncertainty in the $D_s \to f_0 e\nu$ branching ratio in Table A.2.

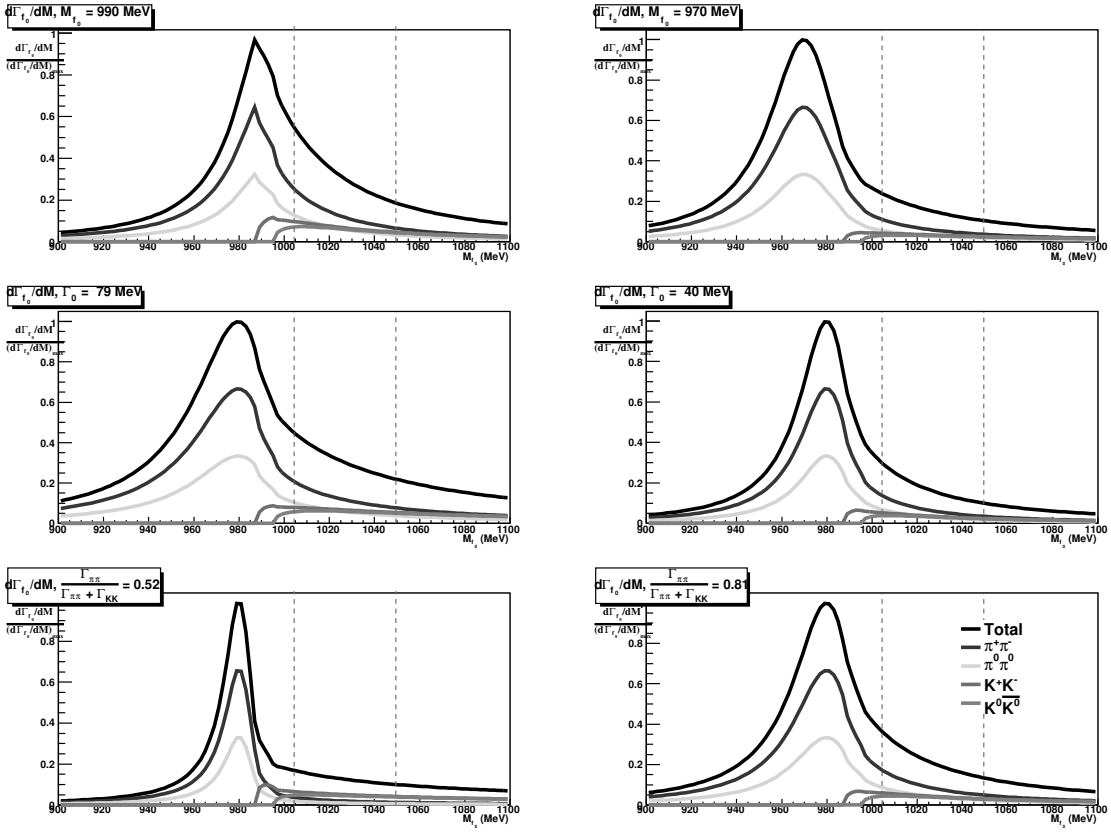

Figure A.4: f_0 mass lineshapes for the M_0, $\frac{g_K}{g_\pi}$, and Γ_0 variations. The sharp peak occurs when $M_0 > 2m_{K^+}$.

Table A.2: Correction and systematic for $\mathcal{B}(D_s \to \phi e\nu)$ from $D_s \to f_0 e\nu$, $f_0 \to KK$ background.

$\mathcal{B}(D_s \to f_0 e\nu, f_0 \to \pi\pi)$	$\frac{\varepsilon_{f_0 \to KK}}{\varepsilon_{\phi \to KK}}$	$\frac{\mathcal{B}(f_0 \to K^+K^-)}{\mathcal{B}(f_0 \to \pi^+\pi^-)} \times \mathrm{f_{window}}$	$\mathcal{B}(D_s \to \phi e\nu)^{\mathrm{cor}}$	$N^{\mathrm{cor}}_{\mathrm{events}}$
$(0.135 \pm 0.025)\%$	1.0	$0.059\,^{0.143}_{0.040}$	$(0.008^{+0.012}_{-0.003})\%$	$0.85\,^{2.09}_{0.54}$

Appendix B

Alternate $D_s \to \phi e \nu$ Reconstruction Methods

In our various D_s semileptonic measurements, we elected not to reconstruct the photon from $D_s^* \to D_s \gamma$ decays. We made this choice for multiple reasons, notably the efficiency hit for the soft photon, potential systematic issues arising from extra candidate showers in the event (from splitoffs or decay-in-flight tracking failures), and the otherwise low background in D_s semileptonic measurements thanks to the electron and D_s tag.

However, we are aware that analyses of this type (e.g. [77]) typically use full event reconstruction to identify the neutrino by its missing mass. Further, other analyses on the CLEO data sample [69, 88] have arrived at a somewhat different $\mathcal{B}(D_s \to \phi e \nu)$ using the missing mass and the D_s^* meson's daughter photon in their tagging procedure. To accommodate any misgivings over the approach that we have taken in our primary analysis, we have also measured $\mathcal{B}(D_s \to \phi e \nu)$ using six variant procedures that approximate the CLEO analyses and other potential approaches.

In four of our six alternate approaches, we reconstruct the D_s^* meson's daughter photon and estimate the number of tags from the $D_s + \gamma$ recoil mass spectrum. We measure the number of semileptonic decays by using the event's missing mass; we get four slightly different combinations by toggling the missing mass fit range and whether or not we make a best candidate choice for the photon. In another alternate approach, we reconstruct the D_s^* daughter photon and require that the $D_s + \gamma$ recoil mass falls

in a reasonable range, but we fit the D_s invariant mass spectrum without the photon for the number of tags. This allows true semileptonic events containing false photon combinations to pass our selection, and it gives us a data estimate for the rate of such false photon combinations when combined with our $D_s + \gamma$ recoil mass tags. Finally, we do an intermediate approach where we use the same D_s tag modes as the other approaches, but we do not reconstruct the D_s^* daughter photon. In this case, we do a simple fit to the D_s invariant mass after all cuts to determine the number of semileptonic events, much like our standard analysis.

We have also tested an approach that uses a two dimensional tagging fit to the D_s invariant mass and the $D_s + \gamma$ recoil mass. We found that this method gave us essentially the same results as when we cut on the D_s mass and fit the $D_s + \gamma$ recoil spectrum. However, our fits to the two dimensional spectrum become sensitive to initial parameters, possibly due to a small remaining correlation between the two variables. Since our fit shape systematic dominates the tagging errors, we have chosen to drop this method and focus on the one dimensional fits instead.

B.1 Particle Selections for Alternate Methods

Given the significant differences in $\mathcal{B}(D_s \to \phi e\nu)$ between our primary analysis and another analysis on the CLEO data sample [88], we have attempted to eliminate any comparison complications by using that analysis's particle selections for all six of our alternate measurement approaches.

We restrict ourselves to 9 of our 13 tag modes: K_sK, $KK\pi$, $KK\pi\pi^0$, $K_sK^-\pi\pi$, $\pi\pi\pi$, $\pi\eta$, $\pi\pi^0\eta$, $\pi\eta', \eta' \to \pi\pi\eta$, and $\pi\eta', \eta' \to \rho\gamma$. We drop our normal D_s momentum cut (in the form of a recoil mass cut that varies by mode) since we will instead be using the $D_s + \gamma$ recoil mass for our fits and selections. Also, we add a 150 MeV ρ cut in $D_s \to \pi\pi^0\eta$. We otherwise retain the individual tag mode cuts listed in Table 5.3.[1]

We use the same electron cuts as in our primary analysis, with the sole exception that we adopt a slightly more conservative $|\cos\theta_e| < 0.90$ angle cut instead of $|\cos(\theta_e)| < 0.93$.

[1]For our intermediate method (Method 6), we do not find the D_s^* meson's daughter photon. While we still restrict ourselves to the 9 tag modes, we otherwise maintain the D_s tag cuts from our primary analysis.

Although we loosen the overall ϕ mass requirement to 60 MeV for all of our alternate methods, we otherwise adopt tighter kaon cuts. Specifically, we institute a minimum hit fraction cut of 0.5 rather than the token hit fraction cut of 0.1 from our primary analysis. In addition to our standard dE/dx consistency cut of $\left|\sigma_K^{dE/dx}\right| < 3.0$, we also add an additional rejection from a basic particle ID. If we have $|p_K > 700$ MeV$|$, then we use both $\sigma^{dE/dx}$ and the RICH likelihood by requiring $(\sigma_\pi^2 - \sigma_K^2) + (L_\pi - L_K) \geq 0$; otherwise, we drop the RICH likelihood and simply require $(\sigma_\pi^2 - \sigma_K^2) \geq 0$.

We require that candidate showers for the D_s^* meson's daughter photon do not come from hot channels in the calorimeter, they can't have an associated track, and they need to pass CLEO's $\frac{E_9}{E_{25}}$ O.K. cut.[2] Showers must have an energy above 50 MeV if in the endcap or 30 MeV in the barrel, although the kinematic range of the D_s^* photon limits these extremes in any case. We also reject any event with an unused shower that meets the above criteria but has an energy above 300 MeV.

B.2 Methods 1-4: Cut on D_s Invariant Mass, Fit $D_s + \gamma$ Recoil Mass

Our first four alternate methods all use the $D_s + \gamma$ recoil mass to tag candidate events. We first restrict the D_s invariant mass to the range $|M_{D_s}^{\rm recon} - M_{D_s}^{\rm PDG}| < 17.5$ MeV. We then allow each passing D_s to pair with any valid shower to form a $D_s + \gamma$ tag candidate. We determine the number of $D_s + \gamma$ tags by fitting the recoil mass spectrum, where the recoil four momentum is given by $p_{\rm recoil} = p_{\rm beam} - (p_{D_s}^{\text{mass constrained}} + p_\gamma)$. By using the mass constrained D_s four vector for our $D_s + \gamma$ recoil, we make the recoil mass fairly independent of the reconstructed invariant mass (Fig. B.1). Aside from conceptual simplicity, this also reduces the remaining recoil background after our basic invariant mass cut.

We fit the recoil mass from each D_s tag mode separately, using a crystal ball function and 4th degree polynomial background function. This gives us a total of five signal parameters and five background parameters. We use the Monte Carlo to fix the crystal

[2] The energy dependent "$\frac{E_9}{E_{25}}$ O.K." cut requires a minimum $\frac{E_9}{E_{25}}$ value for the shower. For low energy photons, the central nine crystals must contain around 80% of the total shower energy. This minimum smoothly scales to requiring roughly 90% of the energy in the central crystals for higher energy photons.

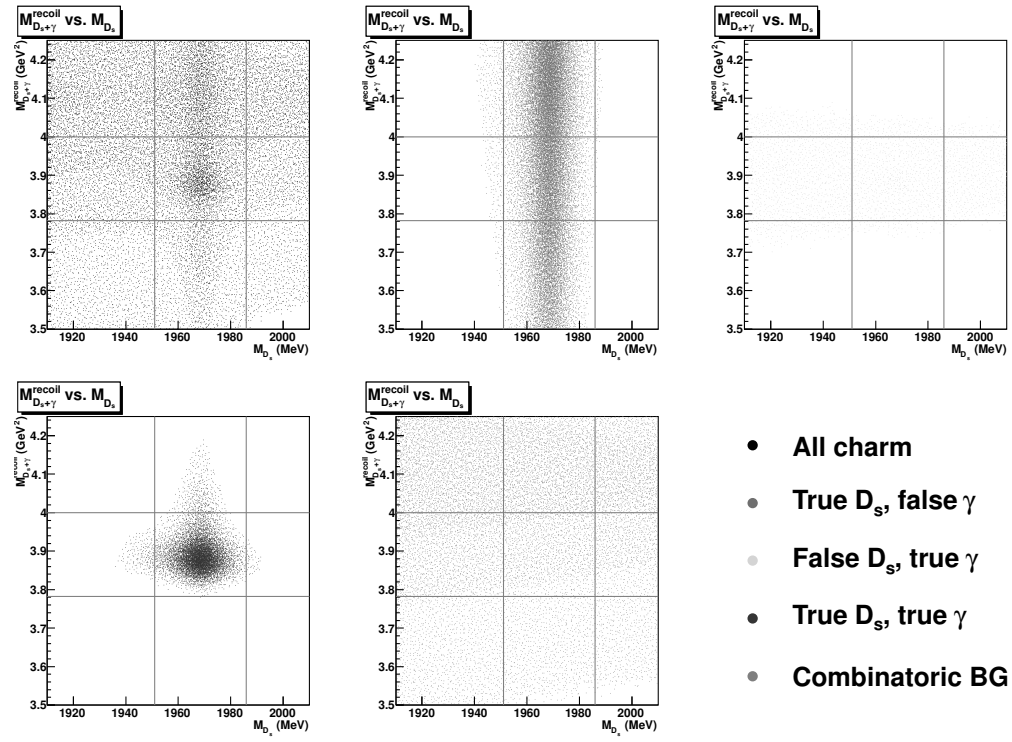

Figure B.1: $D_s + \gamma$ recoil mass vs. D_s invariant mass in the charm Monte Carlo. The lower left plot shows that the two are fairly independent after mass constraining the D_s four vector. The lines indicate our tag fit's invariant mass and recoil mass cuts. We scale the total and combinatoric plots by a factor of 1/10th relative to the others to keep those plots less visually congested.

ball function's mean, its α (the number of σ at which the gaussian turns into a polynomial), and its n (polynomial power). We allow the signal normalization, the signal width (σ), and all of the background parameters to float. We only count $D_s + \gamma$ tags that have a recoil mass within 3.782 GeV^2 $< M^2_{\rm recoil} <$ 4.0 GeV^2, in accordance with [88]; not only does this match our recoil mass cut for the full event reconstruction, but it prevents our tagging counts from being unduly influenced by the long crystal ball tail. We give our $D_s + \gamma$ tag fit results in Table B.1, with Figure B.2 (Monte Carlo) and Figure B.3 (data) showing the actual fits.

For two of our four methods, we use all $D_s + \gamma$ combinations with a valid recoil mass

Table B.1: $D_s + \gamma$ recoil mass tags in the data and Monte Carlo. The crystal ball function tends to undercount the number of tags across all modes, so we adjust the final branching ratio for this systematic effect.

D_s mode	Data tags	MC, fit tags	MC, truth tags
K_sK	3,772 ± 90	72,405 ± 370	74,138
$KK\pi$	16,069 ± 218	326,072 ± 887	341,134
$KK\pi\pi^0$	5,254 ± 251	97,446 ± 934	118,900
$K_sK^-\pi\pi$	2,273 ± 109	50,753 ± 436	53,110
$\pi\pi\pi$	4,618 ± 189	89,019 ± 511	96,594
$\pi\eta$	1,863 ± 78	36,095 ± 265	37,198
$\pi\pi^0\eta$	3,168 ± 152	74,488 ± 516	74,635
$\pi\eta', \eta' \to \pi\pi\eta$	1,419 ± 56	26,888 ± 226	26,681
$\pi\eta', \eta' \to \rho\gamma$	3,351 ± 144	62,202 ± 444	63,918
Sum	41,788 ± 468	835,368 ± 1,683	886,308

when reconstructing a semileptonic event. In the other two methods, we choose a best γ candidate from among those that pass the recoil mass window. We choose this best γ by determining what its lab energy should have been given the $D_s + \gamma$ recoil mass and the shower's position in the calorimeter, then choosing the photon whose energy lies closest to its ideal energy. This procedure provides a slight improvement toward choosing the correct photon in Monte Carlo over simply taking the candidate with the best recoil mass, and it produces less shaping of the final missing mass spectrum. We call the former two methods our "multiple candidate" methods and the latter two our "best candidate" methods, in both cases referring to candidate photons allowed to pair with our D_s.

Once we have our $D_s + \gamma$ tags, we look for a passing ϕ meson and a passing electron from the semileptonic D_s decay. We then calculate the event four vector. The ϕ width prevents us from improving its resolution with a kinematic fit, but we are able to improve the D_s^* daughter photon resolution. For all four of our methods, we adjust the photon energy to its ideal energy, essentially doing a one variable kinematic fit by using the fact that we know the photon location much better than its energy. We then calculate the missing mass from the sum of all four vectors, the beam energy, and the beam momentum: $p_{\text{missing mass}} = p_{\text{beam}} - (p_{D_s}^{\text{mass constrained}} + p_\gamma^{\text{corrected}} + p_\phi + p_e)$.

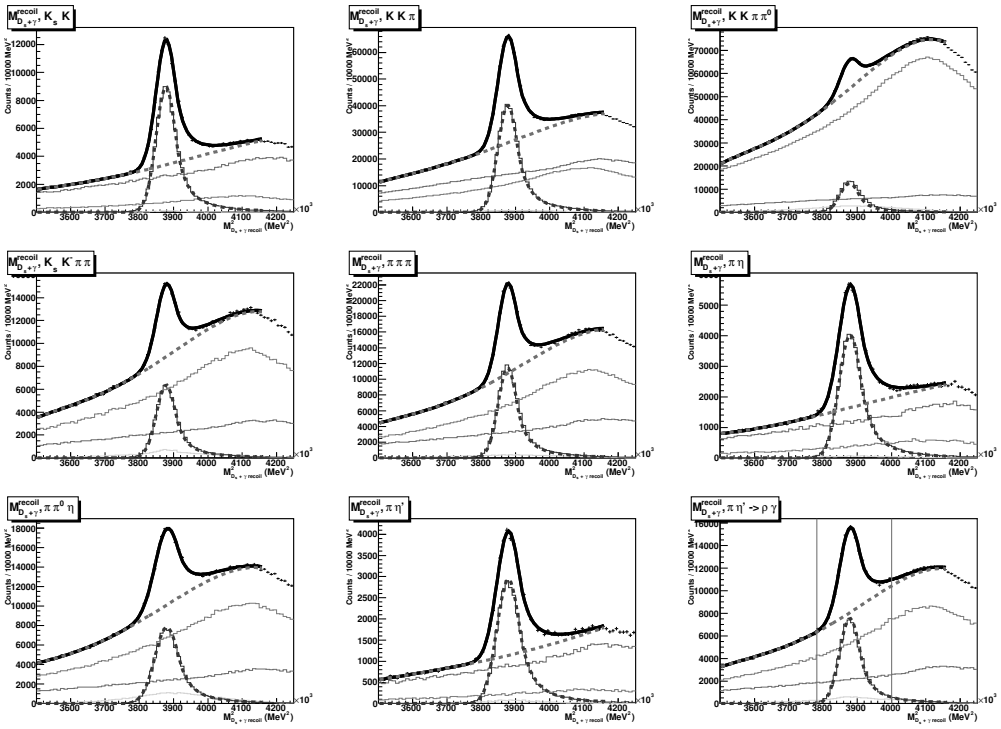

Figure B.2: $D_s + \gamma$ recoil mass distribution for the 9 tag modes in Monte Carlo. The dotted lines give our fit results' signal and background components. The solid lines give the truth-tagged information: signal, combinatoric background, true D_s pairing with a non-D_s^* daughter γ (false γ), and false D_s paired with the D_s^* daughter γ.

We determine the number of signal counts from the missing mass plot, with two signal region possibilities. For our "tight" missing mass range, we consider a signal region with $|MM^2| < 40,000$ MeV^2. Our "wide" missing mass range extends over $|MM^2| < 400,000$ MeV^2.

The two MM^2 ranges deal differently with events that have a true $D_s \to \phi e \nu$ decay but whose D_s tag has been paired with an incorrect D_s^* daughter photon. Such false γ events do not peak in our $D_s + \gamma$ recoil mass spectrum and consequently do not get counted as tags. We then need to cut such events out of our branching ratio's numerator with a background subtraction (the "tight" methods), or we need to trust the Monte Carlo to correct our efficiency properly for such surplus events (the "wide" methods).

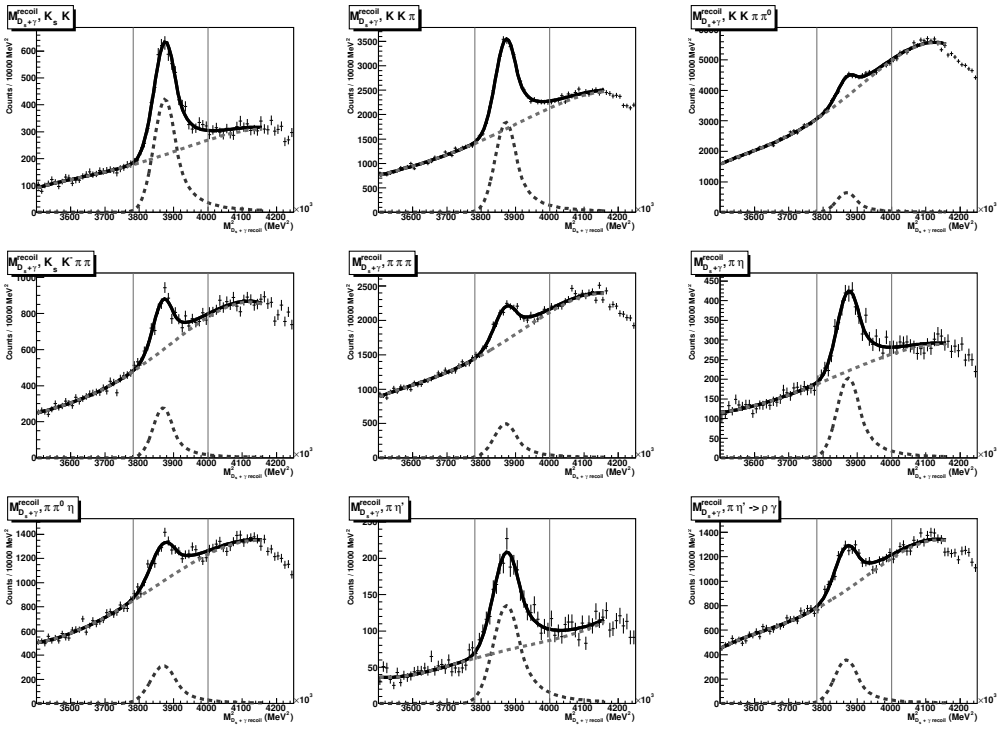

Figure B.3: $D_s + \gamma$ recoil mass distribution for the 9 tag modes in the data. The dotted lines give our fit results' signal and background components.

Both the "tight" and "wide" cut regions have their own difficulties. The "wide" missing mass range covers nearly all semileptonic events (true D_s^* daughter photon or otherwise). We have doubts about using the Monte Carlo to get the false γ rate correct for this method, as we discuss in Section B.4.[3] Combinatoric background also enters into the signal region, which we estimate with the Monte Carlo.[4]

The "tight" background subtraction gets complicated because false γ events peak softly in the $|MM^2|$ distribution. We need to rely on the Monte Carlo to determine the

[3]We can also try estimating the false γ rate from a $D_s + \gamma$ recoil sideband. However, this gets conflated with combinatoric background in the sideband, and it requires us to trust the Monte Carlo to correctly extrapolate the false γ distribution from the sideband to signal region in any case.

[4]We could estimate the background with a D_s mass sideband, but we'd still be relying on the Monte Carlo to estimate how that sideband propagates through the recoil mass distribution with its false γ combinations. We find it best to make Monte Carlo dependence explicit, especially with a small effect like this.

soft peak's shape, as a flat background non-trivially overestimates the branching ratio.

Every method also has the same peaking background from $D_s \to f_0 e\nu$, $f_0 \to KK$ that we see in our standard analysis. For this comparison, we just use the Monte Carlo to correct the $f_0 e\nu$ rate.

Overall, we get four slightly different methods of determining a branching ratio, by taking either a tight/wide missing mass range and by allowing either the best candidate/multiple candidates for the D_s^* daughter photon. We give our missing mass plots from the Monte Carlo and the data for each type of candidate selection in Figures B.4–B.7.

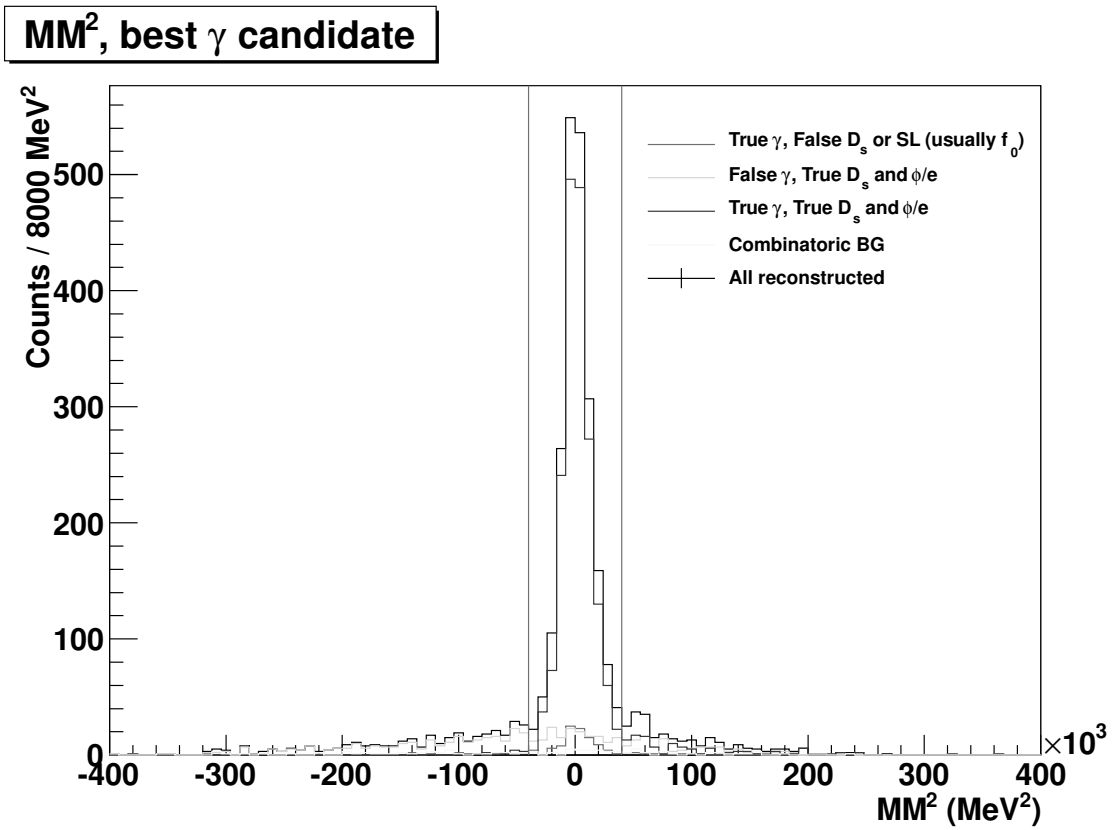

Figure B.4: $D_s \to \phi e\nu$ event's missing mass distribution (ν missing mass) given a best photon candidate selection in the Monte Carlo. The vertical lines indicate our "tight" cut. The wide histogram shows true $D_s \to \phi e\nu$ events that get reconstructed with a false D_s^* daughter photon.

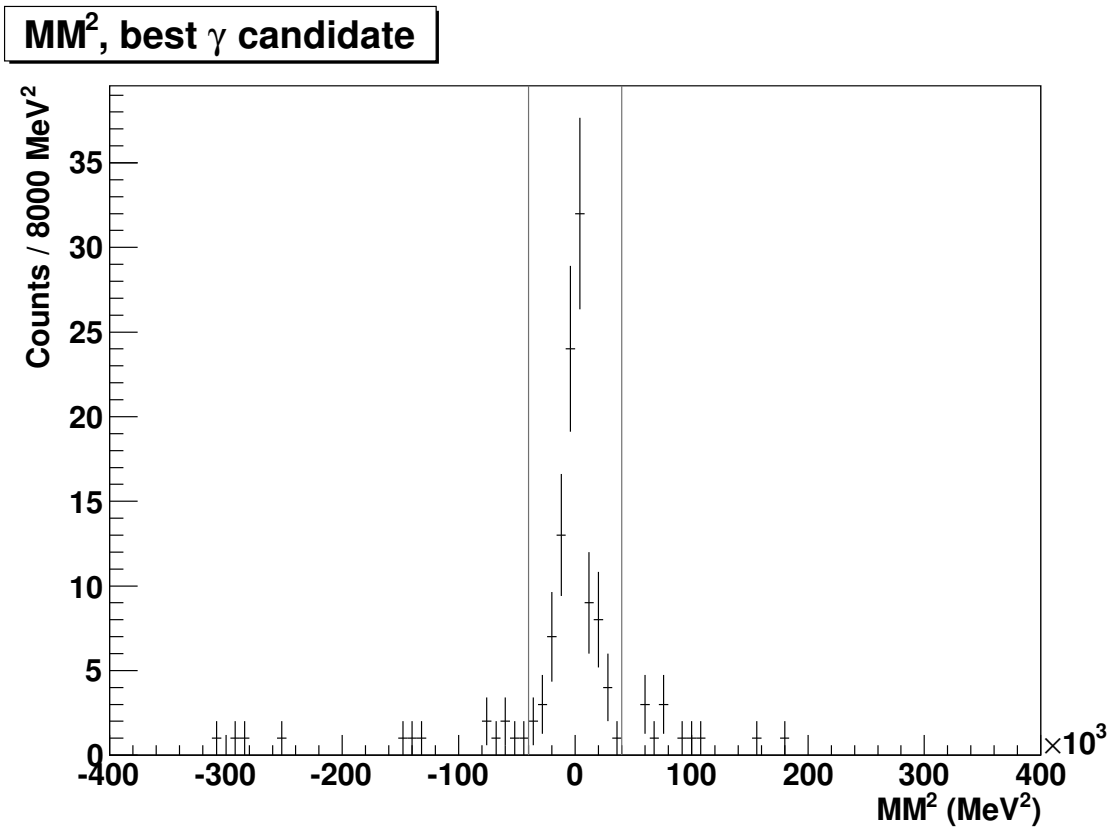

Figure B.5: $D_s \to \phi e \nu$ event's missing mass distribution (ν missing mass) given a best photon candidate selection in the data.

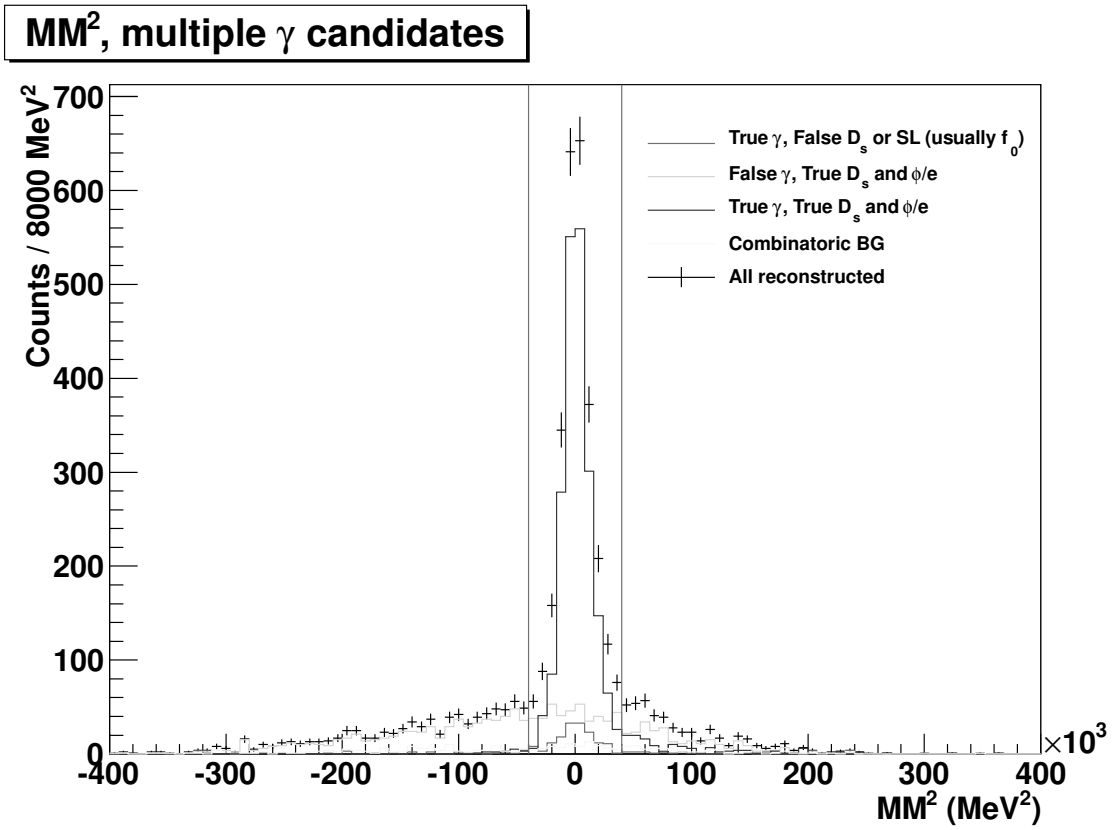

Figure B.6: $D_s \to \phi e \nu$ event's missing mass distribution in the Monte Carlo when we allow multiple photon candidates.

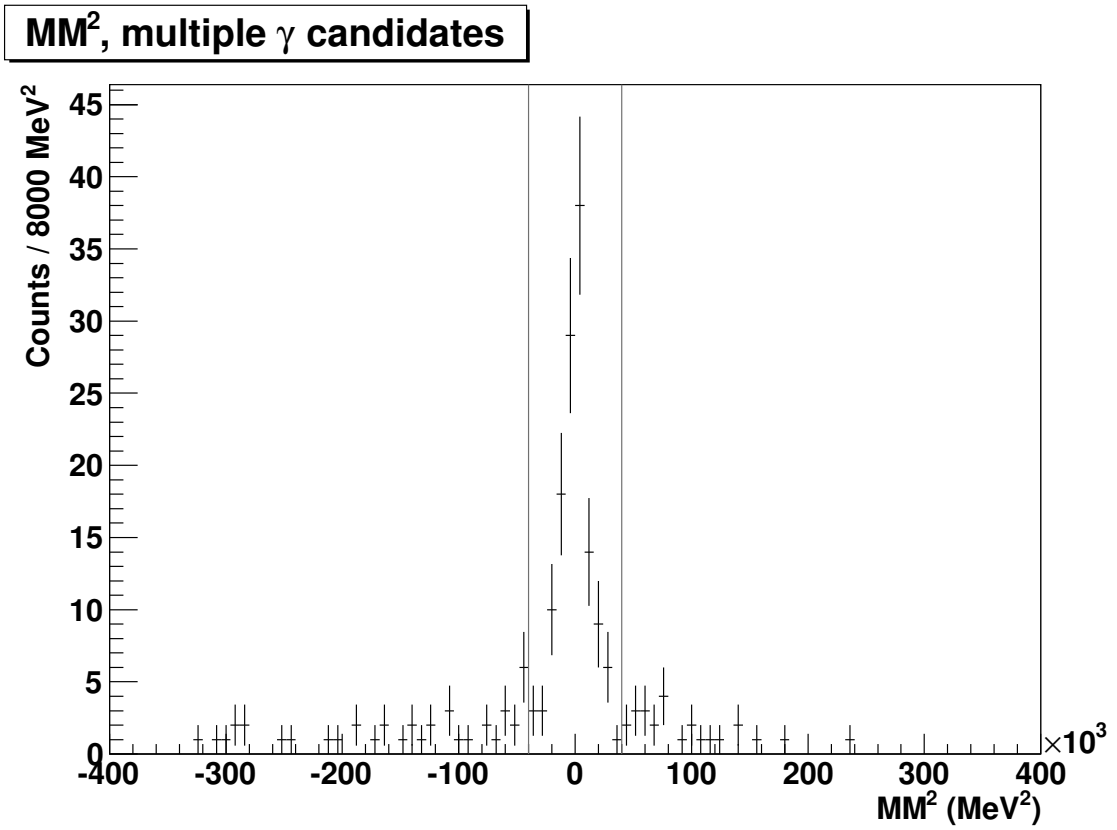

Figure B.7: $D_s \to \phi e \nu$ event's missing mass distribution in the data when we allow multiple photon candidates.

B.2.1 Comparison of Methods 1-4

Each method's photon candidate multiplicity and $|MM^2|$ cut range gives it particular advantages and drawbacks, which we briefly discuss below. As in the previous section, we refer to $D_s + \gamma$ tags with a γ that did not come from a D_s^* daughter photon as "false γ" combinations. In particular, we address the case with more false γ in the data than predicted by the Monte Carlo (although the same problems would occur in reverse if the data has fewer false γ candidates than the Monte Carlo predicts).

Our *best tight* method (best photon candidate, tight $|MM^2|$ range) has the potential flaw that when we get more false γ tag combinations than expected from the Monte Carlo, the false γ contribution to the MM^2 spectrum becomes more peaked than predicted (a consequence of choosing the best γ). This effect causes us to slightly under subtract false γ background from the signal region, leading to an overestimate of the branching ratio. On the other hand, we choose the correct γ less often than predicted with our efficiency (since more false photons exist to potentially make a best candidate), giving us a slight underestimate of the branching ratio. The former effect dominates, as false γ events get pushed into the signal region with the extra combinations more frequently than events with a correct γ get thrown out.

The *best wide* method allows both false and true γ combinations into our $|MM^2|$ range, so we don't have to worry about whether we chose the correct γ combination or not. However, since the $D_s + \gamma$ tag spectrum only peaks with true D_s^* daughter photons, the efficiency determined by the Monte Carlo implicitly assumes a certain rate of extra events from false γ. When we have more false γ than expected, we get more events in the $|MM^2|$ range than we took credit for given our tags and efficiency, causing us to overestimate the branching ratio. We also need to trust the Monte Carlo to determine the number of combinatoric and false D_s background, since our $|MM^2|$ range extends far enough that a sideband subtraction isn't reasonable; this estimate could be systematically high or low, sending the branching ratio either way.

Our *multiple tight* method solves both flaws of the *best tight* method since we don't have a best candidate choice to shape the false γ background, nor do we have an efficiency issue when choosing false γ in place of true γ. We do have the statistical drawback that we aren't making use of true $D_s \to \phi e \nu$ events that happen to have only false D_s^* daughter photons, but that's also true of the *best tight* method and doesn't

drive our error in any case. Although different from the best candidate methods used in prior analyses [88], we consider this the most accurate of our four methods that use $D_s + \gamma$ recoil mass tags.

We include a *multiple wide* method for completeness, although this suffers from a larger efficiency sensitivity to false γ than the other methods. This occurs because we can frequently get false γ combinations even when the correct γ was found and had the best combination, with no sideband available to estimate such combinations. We expect this method to overestimate the branching ratio due to the extra false γ events (or due to the low Monte Carlo efficiency, depending on your perspective).

Keeping these potential systematic biases in mind, we compare the results from each of our four methods in Table B.2 for the Monte Carlo and in Table B.3 for the data sample. We used a Monte Carlo input $D_s \to \phi e\nu$ branching ratio of 2.170%.

Table B.2: Branching ratios in Monte Carlo for each of the four methods that use $D_s+\gamma$ tags. Errors are statistical only.

Allowed candidates	MM^2 range	ε_{SL}	$N_{SL+D_s+\gamma}$	$\mathcal{B}(D_s \to \phi e\nu)$
Best candidate	Narrow	9.63%	1836 ± 51	(2.150 ± 0.059)%
Best candidate	Wide	13.24%	2505 ± 52	(2.134 ± 0.045)%
Multiple candidate	Narrow	10.89%	2080 ± 59	(2.156 ± 0.061)%
Multiple candidate	Wide	18.85%	3561 ± 64	(2.132 ± 0.038)%

Table B.3: Branching ratios in the data for each of the four methods that use $D_s + \gamma$ tags. Errors are statistical only.

Allowed candidates	MM^2 range	ε_{SL}	$N_{SL+D_s+\gamma}$	$\mathcal{B}(D_s \to \phi e\nu)$
Best candidate	Narrow	9.63%	90.1 ± 11.1	(2.110 ± 0.260)%
Best candidate	Wide	13.24%	118.5 ± 11.4	(2.019 ± 0.195)%
Multiple candidate	Narrow	10.89%	98.3 ± 12.9	(2.036 ± 0.269)%
Multiple candidate	Wide	18.85%	173.7 ± 14.0	(2.079 ± 0.170)%

We get largely similar branching ratios from each of our four methods that use $D_s + \gamma$ tags. At worst, our relative systematic error based on the $|MM^2|$ cut window and choice of best candidate comes to around 4%, which can be ignored given our 13%

relative statistical error and our 10% relative systematic error from the tag fits.

B.3 Method 5: Cut on $D_s+\gamma$ Recoil Mass, Fit D_s Invariant Mass

Fitting the $D_s + \gamma$ recoil mass spectrum suffers from the problem of non-linear background, which gets exacerbated by the signal's long power law tail from photon reconstruction. In contrast, the D_s invariant mass spectrum has a fairly flat background, and a smaller power law tail even in photon tag modes since we have one fewer photon to reconstruct. To take advantage of the cleaner fitting while retaining the D_s^* daughter photon reconstruction, we have tried one method using D_s invariant mass tags after cutting on the $D_s + \gamma$ recoil mass.

We allow a D_s tag to enter our invariant mass plot once if it pairs with one or more photons to create a $D_s + \gamma$ recoil mass within $3.782\ \text{GeV}^2 < M^2_{\text{recoil}} < 4.0\ \text{GeV}^2$. We then fit the D_s invariant mass spectrum for each tag mode with a linear background function and the sum of a gaussian and crystal ball to represent the signal. We fix the relative normalization and relative width of the gaussian and crystal ball in the Monte Carlo, as well as the crystal ball function's α and n parameters. We allow all other parameters to float in our fit, including both linear background parameters and the common mean for the gaussian and crystal ball.

We count tags within $|M^{\text{recon}}_{D_s} - M^{\text{PDG}}_{D_s}| < 17.5$ MeV, since we only allow D_s within this mass range to later combine with an electron and ϕ meson for our full semileptonic event. We give our D_s invariant mass tagging results from both the Monte Carlo and the data in Table B.4, with the fits shown in Figure G.63 and Figure G.64 (Appendix G).

Since we use the same D_s invariant mass and $D_s+\gamma$ recoil mass ranges as our previous four methods, we have exactly the same missing mass reconstruction (Figures B.4–B.7). Unlike the $D_s+\gamma$ recoil mass tagging, our D_s invariant mass tags peak whether the D_s^* daughter photon is true or false. Rather than trust the Monte Carlo to estimate how many tags come from true photon combinations, we use a best candidate selection over the "wide" range for the $|MM^2|$, giving one potential event for each tag and keeping both true and false photons. This procedure means that we only have to use our Monte Carlo to provide the efficiency for the ϕ and electron, independent of the false photon

Table B.4: D_s invariant mass tags in the data and Monte Carlo after cutting on the $D_s+\gamma$ recoil mass. We only allow each D_s mass to enter once, regardless of the number of $D_s+\gamma$ combinations.

D_s mode	Data tags	MC, fit tags	MC, truth tags
K_sK	4,726 ± 76	87,061 ± 317	86,050
$KK\pi$	21,603 ± 181	402,906 ± 742	405,485
$KK\pi\pi^0$	8,498 ± 241	149,721 ± 920	145,063
$K_sK^-\pi\pi$	3,887 ± 96	66,833 ± 348	63,501
$\pi\pi\pi$	6,490 ± 186	109,144 ± 460	111,756
$\pi\eta$	2,341 ± 80	42,248 ± 229	42,876
$\pi\pi^0\eta$	3,329 ± 143	90,113 ± 505	86,663
$\pi\eta', \eta' \to \pi\pi\eta$	1,754 ± 47	31,320 ± 191	30,930
$\pi\eta', \eta' \to \rho\gamma$	3,388 ± 135	71,189 ± 425	74,086
Sum	56,017 ± 434	1,050,530 ± 1,534	1,046,410

rate.

Aside from the possibility that the Monte Carlo underestimates the combinatoric background (less than a 4% effect), we do have a slight complication with our D_s tags in this method. Generic D_s decays may have photons as final decay products, while real $D_s \to \phi e\nu$ events do not (at least not for our reconstructed mode of $\phi \to KK$). This means that generic D_s decays have a slightly higher rate of photon candidates available to create a passing $D_s+\gamma$ recoil mass than $D_s \to \phi e\nu$ decays, allowing relatively more D_s tags with fake D_s^* daughter photons in generic decays than we see with real $\phi e\nu$ events (about a 3% correction). If the Monte Carlo underestimates the number of generic D_s decay photons, we will also slightly underestimate our branching ratio.

Overall, we expect a relative systematic from this method of 6%, which is partially independent of the systematic from our recoil mass tag methods. We state our branching ratio results for this method in our summary section, with Tables B.8 and B.9 giving the results in Monte Carlo and data, respectively.

B.4 D_s^* Daughter Photon Efficiencies

Some of our alternate $D_s \to \phi e\nu$ methods depend upon the Monte Carlo predicted rates for reconstructing the D_s^* daughter photon. Fortunately, with the various tagging methods that we've performed (and an additional fit for the number of D_s before combining with photons), we can get an estimate for the actual photon efficiency in data. We can also get an estimate for the rate at which events without a true D_s^* daughter photon still pass the $D_s+\gamma$ recoil mass cut by pairing with another shower (the "false γ" rate). This cross-check indicates both that the D_s^* daughter photon gets successfully reconstructed at a lower rate than expected from the Monte Carlo and that false showers allow the D_s to pass our $D_s + \gamma$ cut more often than expected.

To determine our true efficiency for the D_s^* daughter photon reconstruction, we perform an additional fit to the D_s invariant mass before it combines with a photon. This gives us our efficiency denominator, which we can combine with our number of $D_s + \gamma$ tags to get the efficiency. We use the same fitting methodology for this M_{D_s} fit that we used when determining invariant mass tags (method 5), although here we don't cut on the $D_s + \gamma$ recoil mass before getting our invariant mass distribution. As seen in Table B.5, we obtain a lower photon efficiency in data by about 8% (relative) than we get in the Monte Carlo. We don't list the errors on the efficiency, as the systematics from tagging dominate the small (sub-percent) level statistical errors (we haven't done a thorough systematic evaluation, but a relative 3%-4% seems likely given our previous work).

Table B.5: D_s^* daughter photon efficiency in data and Monte Carlo.

Sample type	$N_{D_s+\gamma}^{\text{recoil}}$	$N_{D_s}^{\text{pre-recoil}}$	ε_γ
MC, truth-tagged	886,308	1,267,860	69.9%
MC, fit	886,309 ± 1,785	1,268,470 ± 1,821	69.9%
Data, fit	44,336 ± 497	68,999 ± 561	64.3%

The D_s tag can combine with showers from sources other than the D_s^* daughter photon (false γ) to form a passing $D_s + \gamma$ recoil mass. These showers can come from real photons on the other side (untagged) D_s decay, or they can arise when the D_s

decays to soft kaons. The kaons have a high rate of unmatched showers due to splitoff and decays in flight that may not be well determined in the Monte Carlo. Although we don't use this result directly, we can get a feel for the data/Monte Carlo difference by looking at how often a tag passes the $D_s + \gamma$ recoil by pairing with a false γ when the real photon was not found.

We use the number of tags from our recoil mass fit (methods 1-4), our invariant mass fit tags after a $D_s + \gamma$ recoil cut (method 5), and the previously determined photon efficiency to determine the rate of fake γ. As seen in Table B.6, the data has a 10% higher likelihood of finding such a fake shower than the Monte Carlo predicts. This gives us the sense that either generic D_s decays have more photons than the Monte Carlo, or we have more non-photon extra showers from D_s decay products than we'd expect. The former effect will distort our efficiency estimate for the invariant mass tag procedure (method 5), while the latter particularly affects $D_s \to \phi e\nu$ and can distort all of our recoil mass results other than the *multiple tight* method.

Table B.6: Rate at which valid D_s without a correctly reconstructed D_s^* daughter photon will still pass all tagging cuts (including the $D_s + \gamma$ recoil mass).

Sample type	$N_{D_s+\gamma}^{\text{recoil}}$	$N_{D_s}^{\text{post-recoil}}$	ε_γ	$f_{\text{fake }\gamma}$
MC, truth-tagged	886,308	1,046,410	69.9%	42.0%
MC, fit	886,309 ± 1,785	1,050,530 ± 1,534	69.9%	43.0%
Data, fit	44,336 ± 497	56,017 ± 434	64.3%	47.4%

We can also get a more direct estimate for the rate of false D_s^* daughter photons in $D_s \to \phi e\nu$ decays by comparing our *multiple tight* method's background estimates in data and Monte Carlo. Background makes up about 45.4% of events in the Monte Carlo compared to 47.3% for the data. Since combinatoric background only makes up some 5% of the total events (according to the Monte Carlo), it's likely that these surplus data background primarily come from false photon combinations. If we interpret the extra background in data as entirely false photon combinations, we estimate a relative 7% higher rate of false γ combinations for $D_s \to \phi e\nu$ events over what the Monte Carlo predicts.

B.5 Method 6: No D_s^* Photon Reconstruction

For ease of comparison, we also include a simplified version of our primary analysis that uses cuts and fits similar to the other five alternate methods. As in our primary analysis, we do not attempt to reconstruct the D_s^* daughter photon, instead using the D_s invariant mass spectrum for both tags and the number of semileptonic events.

We reconstruct the same 9 D_s tag modes listed in B.1, leaving out four of the higher background modes from our primary analysis ($K_sK\pi^0$; $K_sK_s\pi$; $K_sK^+\pi\pi$; and $\pi\pi^0\eta', \eta' \rightarrow \pi\pi\eta$). Since we do not reconstruct a potential D_s^* daughter photon, we don't have a $D_s + \gamma$ recoil mass available for our selection. Instead, we use our primary analysis's D_s momentum cut in the form of the D_s recoil mass range from Table 5.2.

Once we've selected our D_s candidates, we determine the number of D_s tags by fitting the invariant mass spectrum. As in our previous $D_s \rightarrow \phi e\nu$ reconstruction method using the D_s invariant mass for tags (method 5), we fit each tag mode to the sum of a gaussian and crystal ball for the signal with a linear background function (using fixed signal shape parameters from the Monte Carlo). This fit choice differs from our primary analysis, where we sometimes use a double gaussian or a quadratic background, based on the tag mode. While our choice of fit function for this alternate method may not be as accurate as in our primary analysis, it does reduce the likelihood that any difference in results between the alternate methods came from a fit systematic on the tags. Overall, we see an 8% difference in tags from our primary analysis over these modes, although $\pi\pi^0\eta$ drives nearly the entire difference with its non-linear background shape. Our tag results for this method are given in Table B.7, with our data fits shown in Figure G.65.

After making the same electron and ϕ cuts as for our other alternate methods (rather than the ϕ/e cuts from our primary analysis), we again plot the D_s invariant mass. We fit each tag mode with the signal shape determined from our tag fits (only the overall normalization floats), plus a linear background function. Unlike our primary analysis, we refrain from using a common branching ratio to ensure that this method remains both simple and as similar as possible to the other alternate methods. We thus fit each tag mode independently.

Figures B.8 and B.9 show our Monte Carlo and data plots for the D_s mass after making our semileptonic selections. We include the results in our summary section with

Table B.7: D_s invariant mass tags in the data and Monte Carlo after a D_s momentum cut. We do not require a pairing with a photon.

D_s mode	Data tags	MC, fit tags	MC, truth tags
K_sK	6,277 ± 95	115,958 ± 380	113,885
$KK\pi$	27,676 ± 225	510,195 ± 888	514,865
$KK\pi\pi^0$	6,731 ± 203	119,558 ± 750	136,976
$K_sK^-\pi\pi$	3,411 ± 88	65,749 ± 351	64,034
$\pi\pi\pi$	5,992 ± 201	118,575 ± 521	124,024
$\pi\eta$	3,809 ± 135	64,179 ± 317	66,758
$\pi\pi^0\eta$	5,718 ± 331	108,063 ± 800	131,334
$\pi\eta', \eta' \to \pi\pi\eta$	2,412 ± 60	42,550 ± 232	42,181
$\pi\eta', \eta' \to \rho\gamma$	3,487 ± 158	72,327 ± 453	81,611
Sum	65,513 ± 553	1,217,150 ± 1,700	1,275,670

Tables B.8 and B.9.

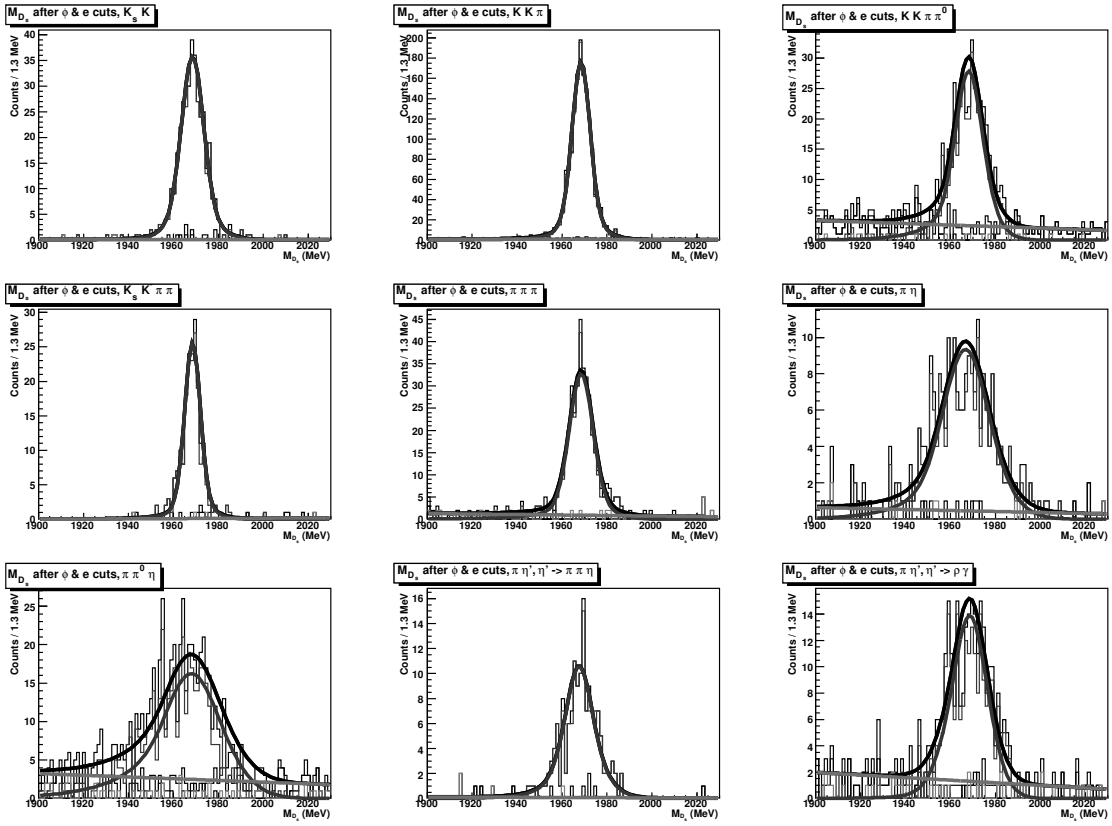

Figure B.8: D_s invariant mass for events passing the semileptonic ϕ and electron cuts in the Monte Carlo.

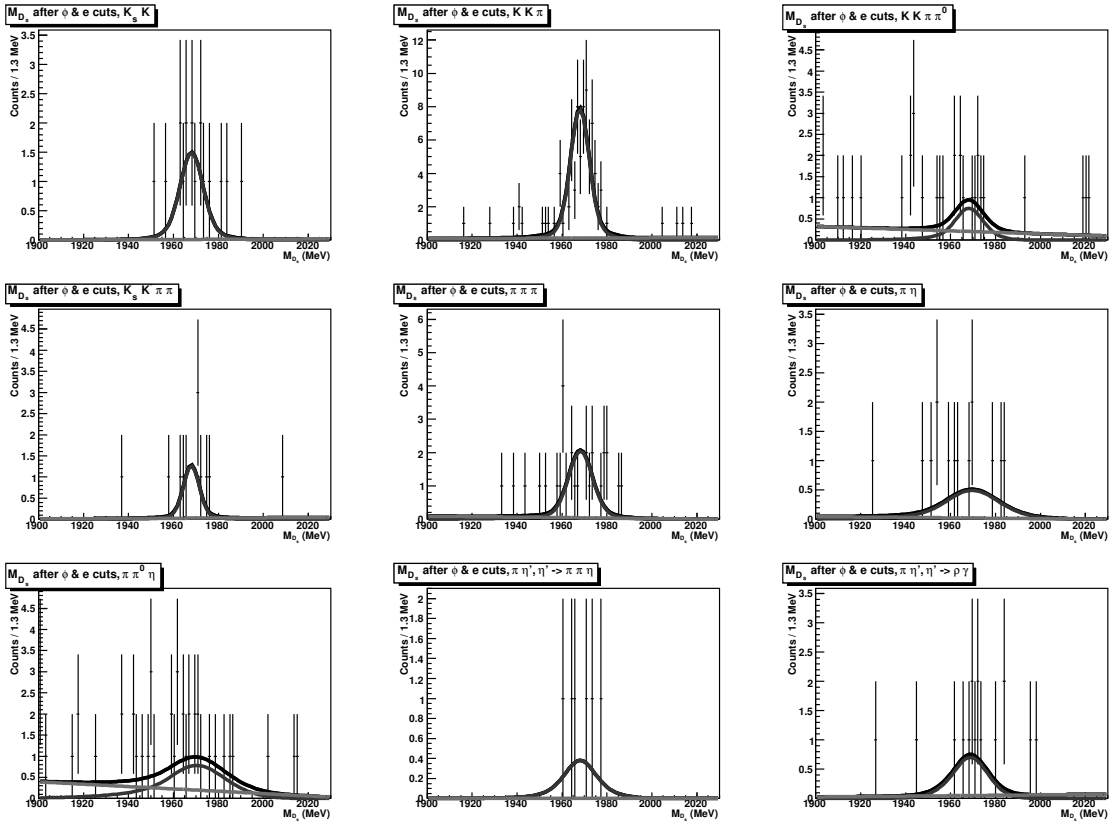

Figure B.9: D_s invariant mass for events passing the semileptonic ϕ and electron cuts in the data.

B.6 Comparison of Alternate Methods

Our best $D_s \to \phi e\nu$ measurement using $D_s + \gamma$ recoil mass tags comes from the *multiple tight* method, so we compare it, the D_s invariant mass tag method, and our simplified no D_s^* photon analysis in Tables B.8 and B.9 for the Monte Carlo and data, respectively. The *multiple tight* method and the simplified no photon method give us the biggest branching ratio range, with a relative difference of about 10%. We expected our non-correlated systematic errors between methods to also be at about the 10% relative level, so this relative difference seems reasonable. Note that while the statistical errors also cover this range, they don't completely explain the difference since they have a correlation with each other (all methods run over the same data sample).

Table B.8: Branching ratios in Monte Carlo for our different $D_s \to \phi e\nu$ alternate methodologies. Errors are statistical only.

$D_s \to \phi e\nu$ method	$N_{\rm tags}^{\rm raw}$	ε_{SL}	$N_{SL+D_s+\gamma}$	$\mathcal{B}(D_s \to \phi e\nu)$
$M_{D_s+\gamma}^{\rm recoil}$ tags, *multiple tight*	835,368 ± 1,683	10.89%	2080 ± 59	(2.156 ± 0.061)%
M_{D_s} tags, *best wide*	1,050,530 ± 1,534	11.48%	2505 ± 52	(2.150 ± 0.045)%
M_{D_s} tags & signal, no γ	1,216,450 ± 1,698	13.97%	3160 ± 61	(2.150 ± 0.042)%

Table B.9: Branching ratios in the data for our different $D_s \to \phi e\nu$ alternate methodologies. Errors are statistical only.

$D_s \to \phi e\nu$ method	$N_{\rm tags}^{\rm raw}$	ε_{SL}	$N_{SL+D_s+\gamma}$	$\mathcal{B}(D_s \to \phi e\nu)$
$M_{D_s+\gamma}^{\rm recoil}$ tags, *multiple tight*	41,787.7 ± 468.1	10.89%	98.3 ± 12.9	(2.036 ± 0.269)%
M_{D_s} tags, *best wide*	56,017.1 ± 433.6	11.48%	118.5 ± 11.4	(1.908 ± 0.183)%
M_{D_s} tags & signal, no γ	65,476.6 ± 551.6	13.97%	144.8 ± 13.3	(1.831 ± 0.169)%

Our primary analysis yields a branching ratio roughly in the middle range of our alternate methods, with the full range covered by its statistical error. We take this to mean that our primary $D_s \to \phi e\nu$ result is fairly robust to the changes in fit function, particle cuts, and D_s^* photon reconstruction considered in this appendix.

Our *multiple tight* method and our D_s invariant mass tag method both involve an

$f_0 \to KK$ correction based on our Monte Carlo. Given the discussion in Appendix A, the Monte Carlo likely has more $f_0 \to KK$ than the data, so this could lead to an underestimate of the branching ratio by a relative 2%-3%. We do not attempt an $f_0 \to KK$ correction in our simplified no D_s^* photon analysis; this could lead to an overestimate of our branching ratio by up to a relative 2%. However, the systematic error from our (different) f_0 correction methodologies comes out to roughly the size of the correction itself, so this doesn't alter the consistency between our different methods.

Appendix C

$D_s \to \eta e \nu$ Efficiency Systematic

As described in Section 10.6.2, we have elected to create our own η efficiency systematic rather than adopting a preexisting one. This gives us an η reconstruction efficiency with our exact η selections, a comparable η lab momentum range (Figure C.1), and a run environment similar to that in the $D_s \to \eta e \nu$ analysis.

In addition to these benefits, our $D_s \to \eta e \nu$ Monte Carlo datasets show an $\eta e \nu$ efficiency variation that borders on the edge of allowed random variation (Figure C.2). Of course, with six semileptonic modes, it shouldn't be shocking if one mode's efficiency variations have a one in six chance of being consistent with a random distribution. Even so, creating our own systematic for the efficiency from the same datasets gives us more confidence that our analysis has a sound foundation.

We obtain an η sample by taking advantage of the large (8.9%) $D_s^{\pm} \to \rho^{\pm}\eta$ branching ratio [89]. We determine the presence of an η by finding the recoil mass after reconstructing a D_s tag, the D_s^* daughter photon, and a ρ. Then, we explicitly reconstruct the $\eta \to \gamma\gamma$ with our analysis's η selection to get an η efficiency. As this can be done in both data and Monte Carlo, we can compare the efficiencies for our overall systematic.

For our D_s tag, we use the same 13 tag modes as our full analysis, described in Section 5. Since we don't need to fit the D_s mass spectrum, we make the additional mode-dependent D_s mass cuts shown in Table F.11 (Appendix F). We follow the selections from Section 9.1 for our D_s^* daughter photon candidates, with the exception of the irrelevant minimum energy cut.

For the $\rho^{\pm}$, we reconstruct a $\pi^{\pm}$ and a π^0. Our charged π follow the same selection

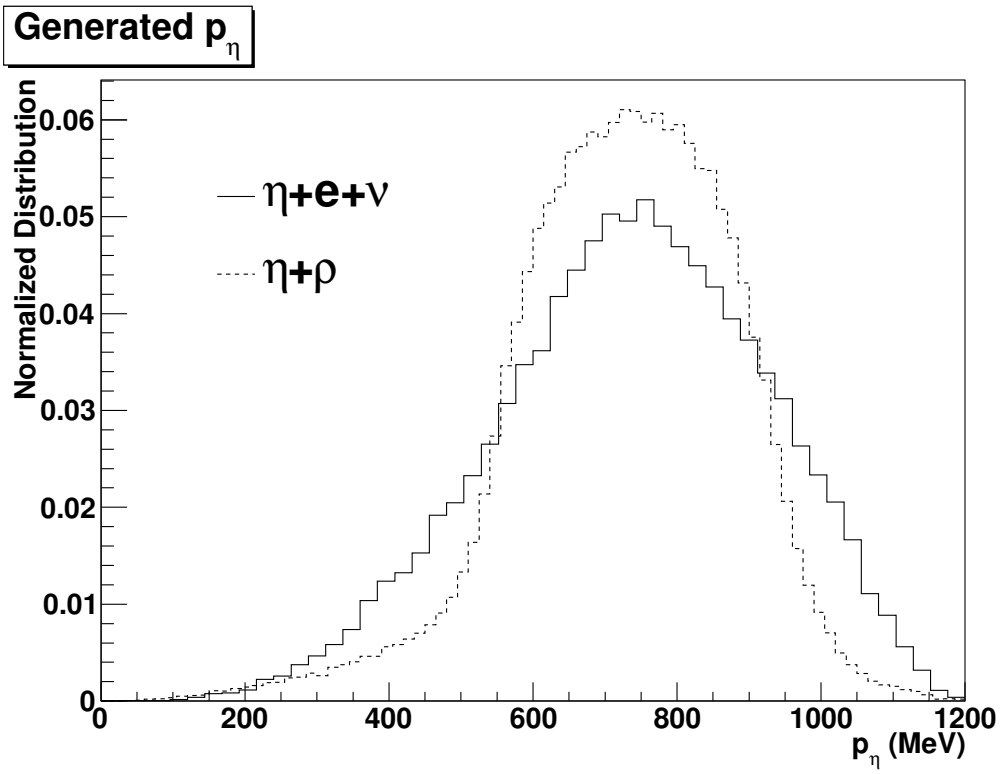

Figure C.1: The η momentum distributions in $D_s \to \eta e \nu$ and $D_s \to \eta \rho$ both peak near 750 MeV and have comparable widths. This contrasts with the alternate source for a clean η sample, $\psi' \to \eta J/\psi$, which creates monoenergetic η with a momentum of 199 MeV.

criteria as in Section 8.1, although we make minor adjustments by lowering the χ^2 threshold to 1,000, we don't allow hit fractions above 1.2, and we drop the $\left|\sigma_\pi^{dE/dx}\right| < 3.0$ cut when we have both RICH information and $p_\pi > 550$ MeV. The π and D_s tag must have opposite charges, and the event can have no other charged tracks. We require a 3.0 pull mass cut on the π^0, and its showers must meet the requirements from Section 9.1 (although we drop the distinction between barrel and endcap showers).

The background to $D_s \to \rho\eta$ mostly consists of events with soft pions. We can eliminate much of this background by adding cuts on the ρ. Specifically, we require 600 MeV $< M_\rho <$ 960 MeV and 500 MeV $< p_\rho <$ 1000 MeV. We also eliminate particular backgrounds by rejecting $\pi^\pm$ that have a rest frame momentum within 5 MeV

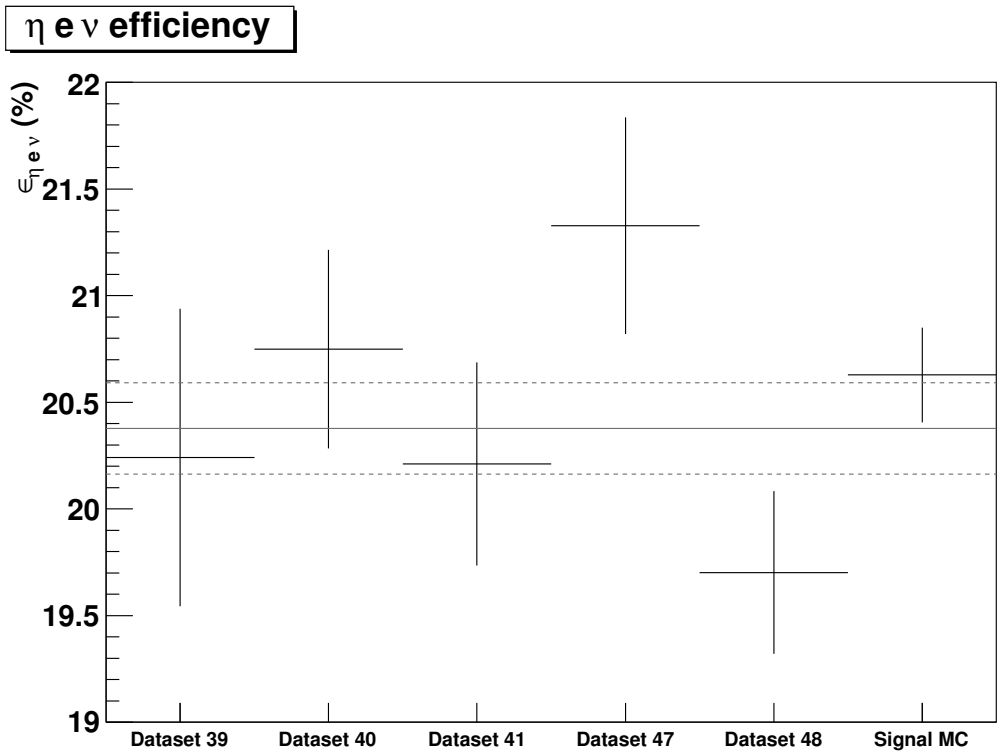

Figure C.2: $D_s \to \eta e \nu$ efficiency (including the $\eta \to \gamma\gamma$ branching ratio) by dataset. The solid line gives the average across the full generic Monte Carlo sample, while the dotted lines show the 1σ range on this average.

of 712 MeV, 743 MeV, or 902 MeV for D_s to $\pi\phi$, $\pi\eta'$, and $\pi\eta$, respectively.

To determine our number of $D_s \to \rho\eta$ events, we perform a 2D fit to the $D_s + \gamma$ recoil mass and the $D_s + \gamma + \rho$ recoil mass. We use the Monte Carlo to get four lineshapes for the fit corresponding to events with true or false η and true or false $D_s + \gamma + \rho$. The Monte Carlo accurately reproduces the widths of these lineshapes in data, but the peak locations have a slight shift. We allow the distributions to shift in each dimension and take the best χ^2. In the data, this shifts our $D_s + \gamma$ recoil mass fit function by 1.3 MeV and our $D_s + \gamma + \rho$ recoil fit function by 4.5 MeV. We give our projections for each fit dimension in Figures C.3 (Monte Carlo) and C.4 (data).

After identifying events with an η, we explicitly reconstruct the $\eta \to \gamma\gamma$ using the same shower selections as for the π^0 from the ρ. We only consider events with a $D_s + \gamma$

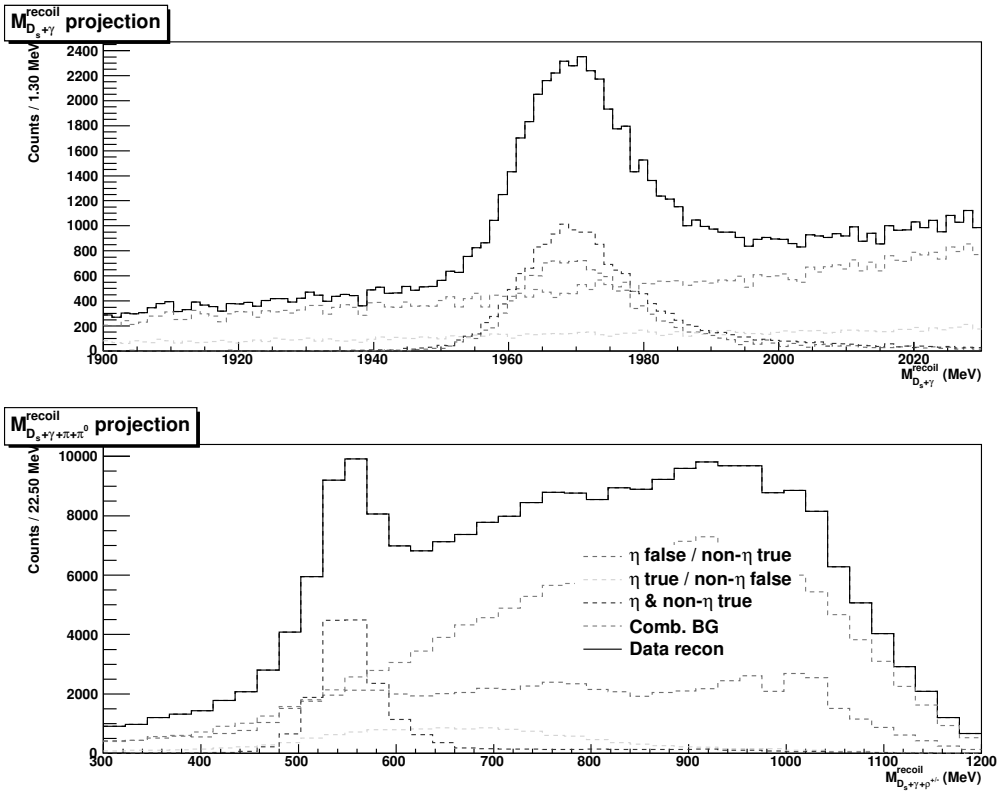

Figure C.3: Top: $D_s + \gamma$ recoil mass in the Monte Carlo for events with a $D_s + \gamma + \rho$ recoil mass between 500 MeV and 600 MeV. Bottom: $D_s + \gamma + \rho$ recoil mass for Monte Carlo events that have a $D_s + \gamma$ recoil between 1955 MeV and 1990 MeV.

recoil mass between 1950 MeV and 1990 MeV.

Once we have our reconstructed η, we get our efficiency's numerator by doing a 2D fit to the η pull mass and the $D_s + \gamma + \rho$ recoil mass. We again take lineshapes from the Monte Carlo, keeping the same 4.5 MeV shift to the $D_s + \gamma + \rho$ recoil mass that we obtained from our previous fit. We show our distributions from this fit for the Monte Carlo in Figure C.5 and for the data in Figure C.6.

We expected to see a lower η efficiency from the data than the Monte Carlo, in accordance with a previous CLEO analysis using $\psi' \to \eta J/\psi$ that saw a relative 5.6% correction to the η efficiency [71]. However, with our η environment and selections, we only see a relative 1.2% lower efficiency in the data compared to the Monte Carlo

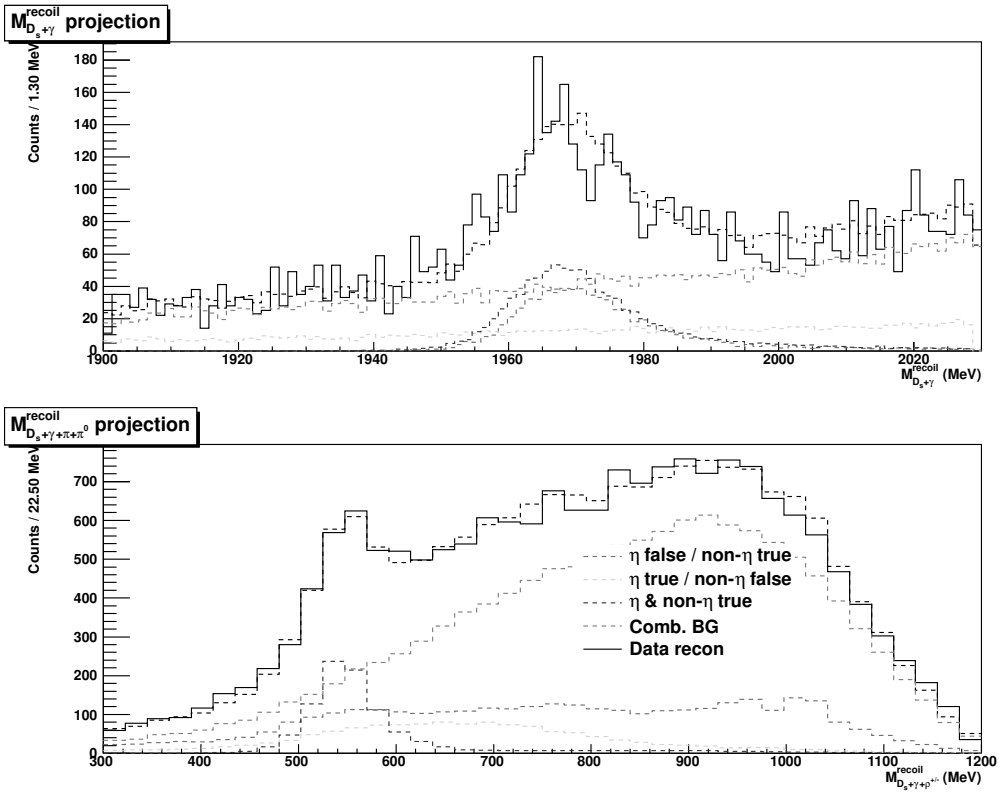

Figure C.4: Top: 2D fit projections for the $D_s + \gamma$ recoil mass in the data from events with a $D_s + \gamma + \rho$ recoil mass between 500 MeV and 600 MeV. Bottom: Fit projection for the $D_s + \gamma + \rho$ recoil mass from data events that have a $D_s + \gamma$ recoil between 1955 MeV and 1990 MeV.

(32.7% to 33.1%), well within our error. Consequently, we do not take a correction to our η efficiency.

Our systematic error on the efficiency comes almost equally from our error on the numerator (2D pull mass/$D_s + \gamma + \rho$ recoil fit) and our error on the denominator (2D $D_s + \gamma + \rho$ recoil/$D_s + \gamma$ recoil fit). We use a binomial error for the our efficiency, although we have to adjust it upward by a factor of 1.2 to account for fit backgrounds. Ultimately, we obtain an η efficiency in the data of (32.7$\pm$ 2.6)%, giving us a relative η efficiency systematic of 7.9%.

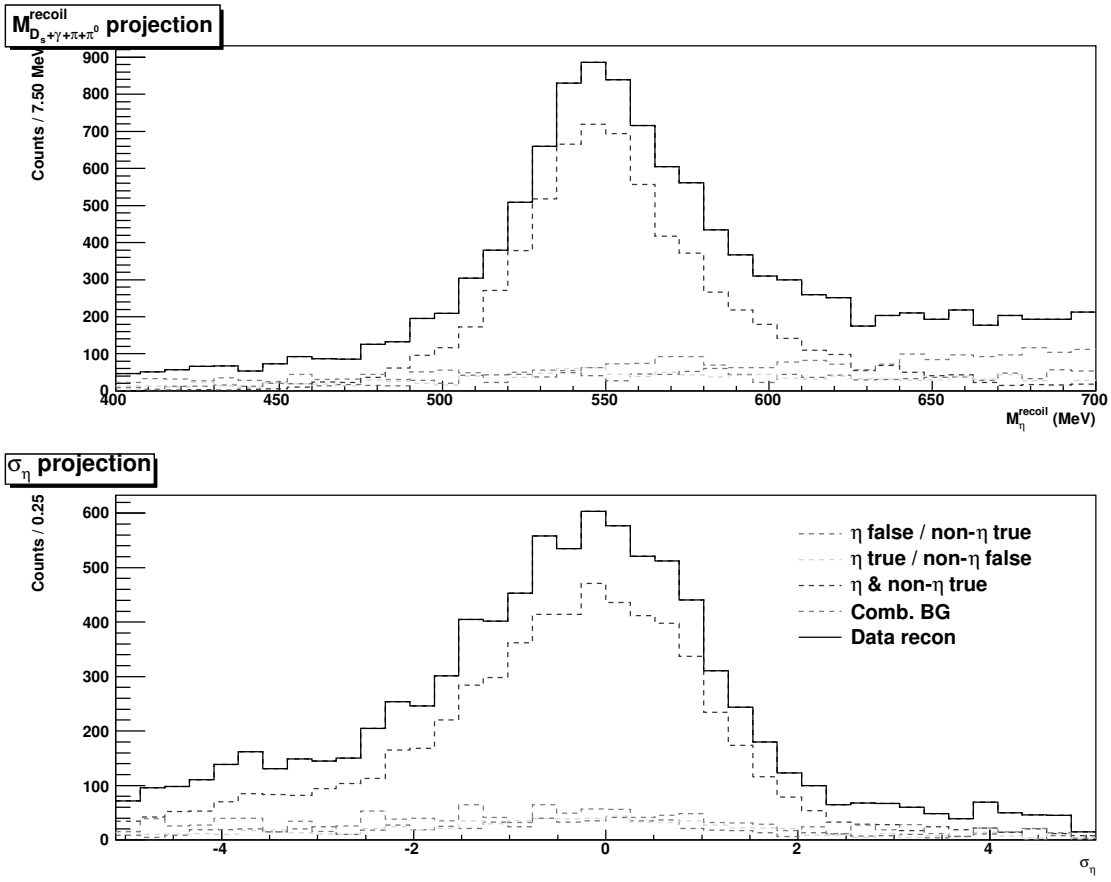

Figure C.5: Top: $D_s + \gamma + \rho$ recoil mass in the Monte Carlo for events with $-3.5 < \sigma_\eta < 2.5$. Bottom: η pull mass for Monte Carlo events that have a $D_s + \gamma + \rho$ recoil between 500 MeV and 600 MeV.

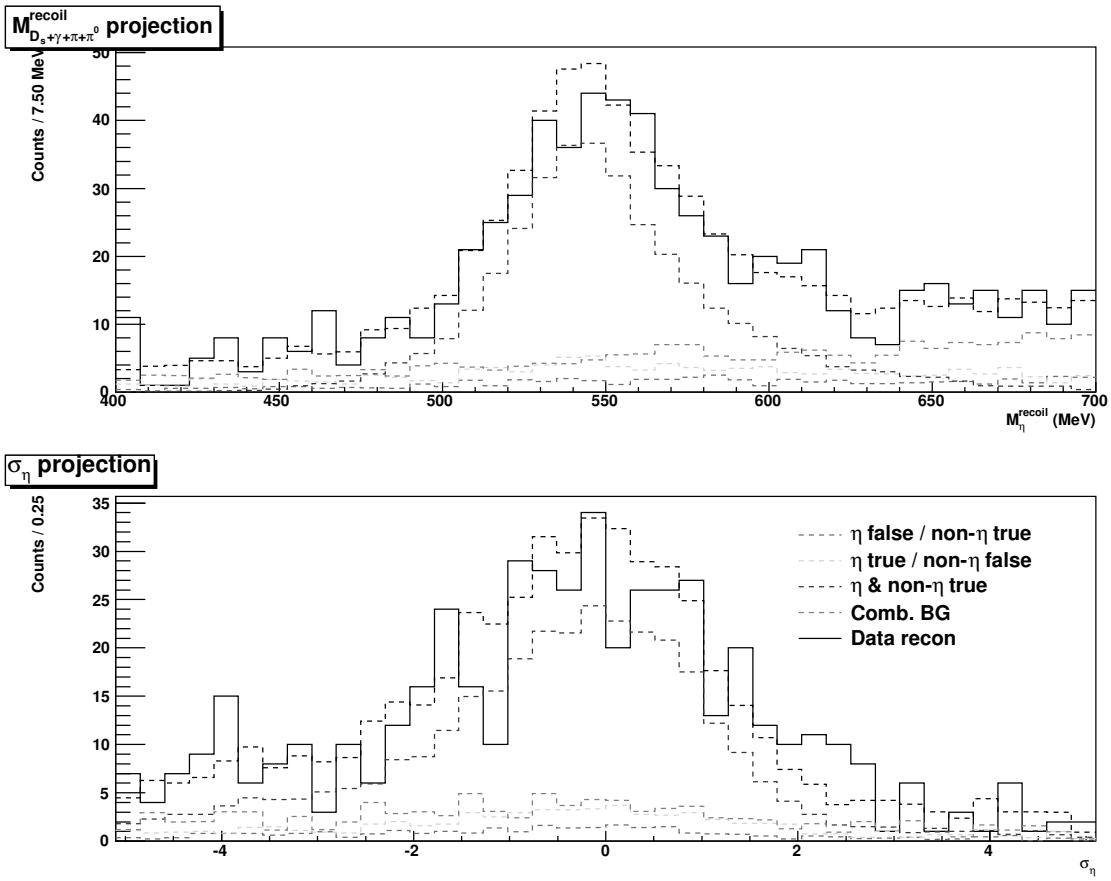

Figure C.6: Top: 2D fit projections of the $D_s + \gamma + \rho$ recoil mass for data events with $-3.5 < \sigma_\eta < 2.5$. Bottom: Fit projection for the η pull mass from data events that have a $D_s + \gamma + \rho$ recoil between 500 MeV and 600 MeV.

Appendix D

Kaon Tracking and Particle ID Systematic

While all track reconstructions depend on the particle's momentum to some extent, kaon reconstruction and identification have a particularly strong dependence on momentum given the possibility of the kaon decaying within the detector. Kaons in D_s semileptonic decays tend to have lower momenta than kaons in D^0/D^+ decays, so we use loose cuts to gain efficiency for those otherwise low yield kaons. However, that means that our cuts don't match the prior CLEO systematics, requiring us to perform our own systematic study for kaons. We choose an approach that combines the reconstruction and particle ID efficiencies, giving a single, momentum-dependent systematic error for kaons.

We initially looked in the D_s system for a kaon systematic, given the relatively large numbers of kaons produced thanks to the D_s meson's strange component. However, our systematic errors from these studies typically came out to about 3% per kaon, or about 6% for $D_s \to \phi e\nu$ as a whole given the correlation between the systematics on the two kaons. The higher statistics from $D^+ \to K^-\pi^+\pi^+$ (and its charge conjugate) during CLEO's 3770 MeV running allow us to achieve a lower systematic on the kaons, at the expense of not having an exactly identical run environment. While we don't use it directly, our original $D_s \to KK\pi$ study yielded roughly the same needed kaon efficiency correction (with higher errors) that we will see from our D^+ study below.

To isolate $D \to K\pi\pi$ decays, we first reconstruct a $D^\pm$ tag. We use five modes in

our $D^{\pm}$ tag reconstruction: $K\pi\pi$, $KK\pi$, $K_s\pi$, $K_s\pi\pi^0$, and $K\pi\pi\pi^0$. We ensure that we have a D^+D^- event by requiring that the tag's beam constrained mass falls within 5 MeV of the D^+ mass and that the ΔE falls within 20 MeV of zero. We then choose each charge's best tag by $D^{\pm}$ invariant mass. Once we have a D^+ tag candidate, we rescale its total momentum to match D^+D^- production, improving the resolution of our later recoil masses.

After finding a tag, we look for two additional tracks with proper charges passing pion particle ID. We reject any event with a total extra energy above 250 MeV or with an extra track passing simple electron cuts, avoiding backgrounds from π^0 modes and semileptonic modes, respectively. We can then identify $D \to K\pi\pi$ events by checking that the recoil mass against the D and two π matches a kaon.

From here, we have two ways of calculating the kaon efficiency. We can take all events with a recoil mass near the kaon mass as the denominator, then obtain our numerator by explicitly reconstructing the kaon and finding a zero missing mass for the event. Alternately, we can try to find the kaon, plotting the "found" recoil mass when we reconstruct it and the "not found" recoil mass when we don't, with $\varepsilon_K = \frac{N_{\text{found}}}{N_{\text{found}}+N_{\text{not found}}}$ giving the total kaon efficiency. As it turns out, the found/not found approach makes it slightly easier in practice to get good precision from our fitting because nearly all the background comes from events without a found kaon, allowing us to focus on those events as the source of any nonstatistical error.

To get separate efficiencies for each kaon momentum region, we split our sample into three bins based on the $D^+ + \pi\pi$ recoil momentum: one for kaon momenta below 250 MeV, one for kaons between 250 MeV and 500 MeV, and one for kaon momenta above 500 MeV. We've chosen these momentum regions so that $D_s \to \phi e\nu$ kaons split roughly evenly between the lowest and middle bins. We fit to the "found" and "not found" plots in each of these regions using histogram shapes from the Monte Carlo, with a fit systematic error determined by doing a simple cut and count to the same plots.

When fitting, we discovered that the recoil kaon mass in the data tends to fall slightly below the recoil mass in the Monte Carlo. To account for this, we allow our signal shape to shift by small amounts to the left, and we take the best χ^2 from all such shifts (which results in a recoil shift of about 0.8 MeV). We show the final Monte Carlo and data plots from each momentum region for our ϕ meson's kaon selections in Figures D.1 and

D.2, respectively.

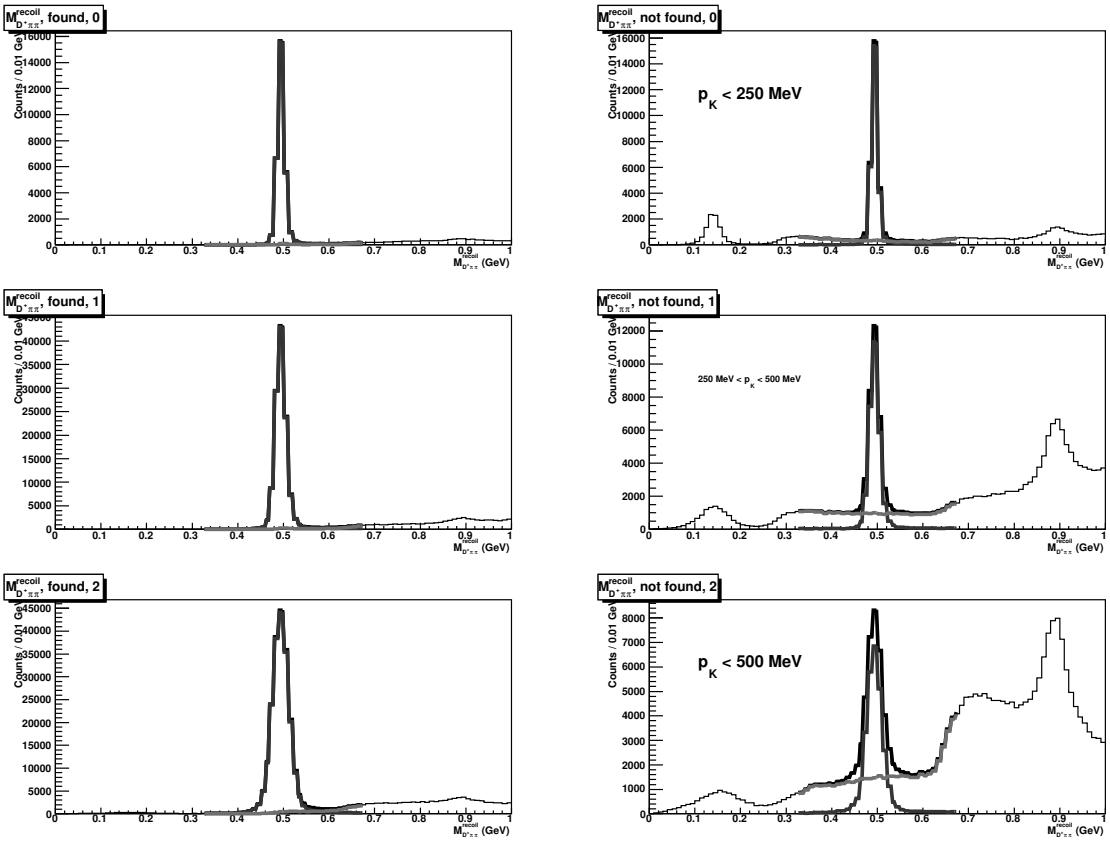

Figure D.1: $\pi^{\pm}\pi^{\pm}D^{\mp}$ recoil mass fits for our ϕ kaon selections in the Monte Carlo. The left plots show the recoil mass when we find a kaon, while the right plots show the recoil mass when we don't find the kaon. The top plots contain recoil momenta below 250 MeV, the middle plots have recoil momenta between 250 MeV and 500 MeV, and the bottom plots have recoil momenta above 500 MeV.

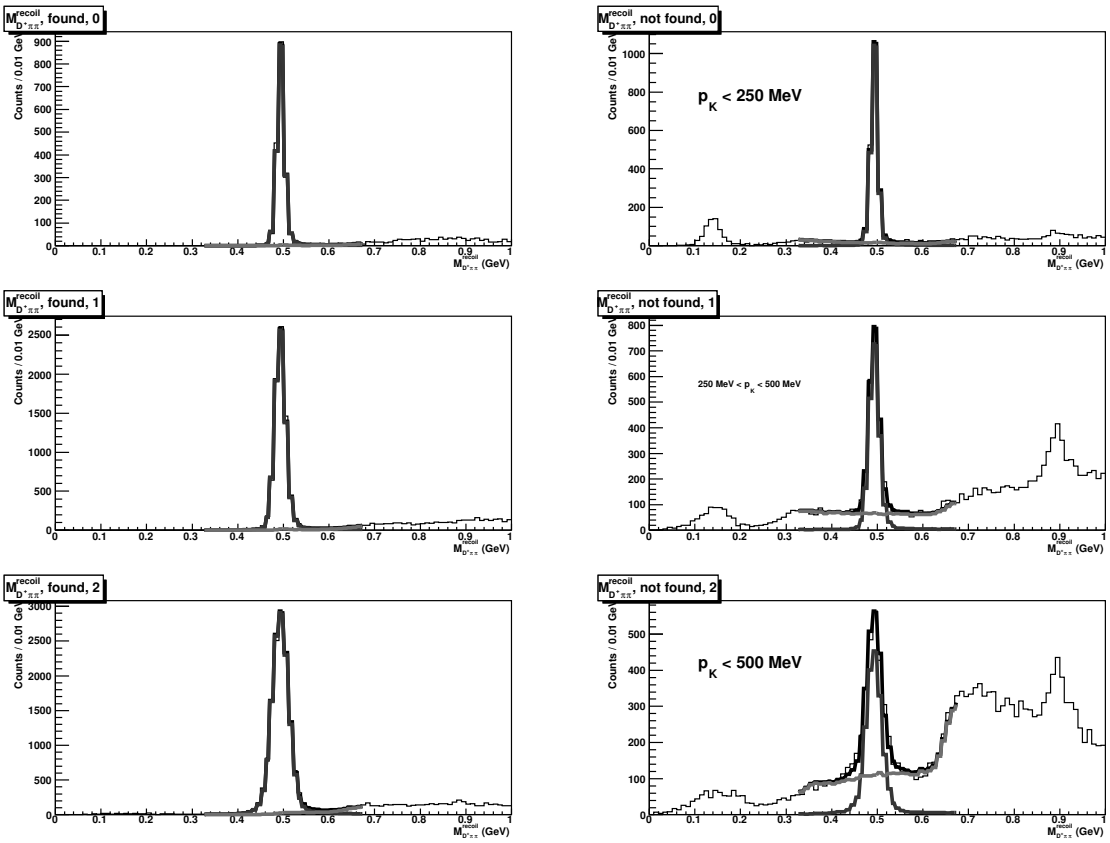

Figure D.2: $\pi\pi D$ recoil mass fits for our ϕ kaon selections in the data. The left plots show the recoil mass when we find a kaon, while the right plots show the recoil mass when we don't find the kaon. The top plots contain recoil momenta below 250 MeV, the middle plots have recoil momenta between 250 MeV and 500 MeV, and the bottom plots have recoil momenta above 500 MeV.

In our analysis, we use three different sets of selections for kaons: one for ϕ (Section 7.1); one for K^* (Section 8.1); and one for D_s tags (Section 5.1.1), where our kaon selections for the D_s tags follow the standard CLEO kaon cuts. We have repeated our systematic for each of these kaon cuts, with the final efficiencies across our kaon selections and momentum regions shown in Figure D.3 for the Monte Carlo and Figure D.4 for the data reconstruction. In all cases, we find that the data efficiency deviates from the Monte Carlo efficiency for soft kaons and requires an efficiency correction, as shown in Figure D.5 and from our final results in Tables D.1–D.3.

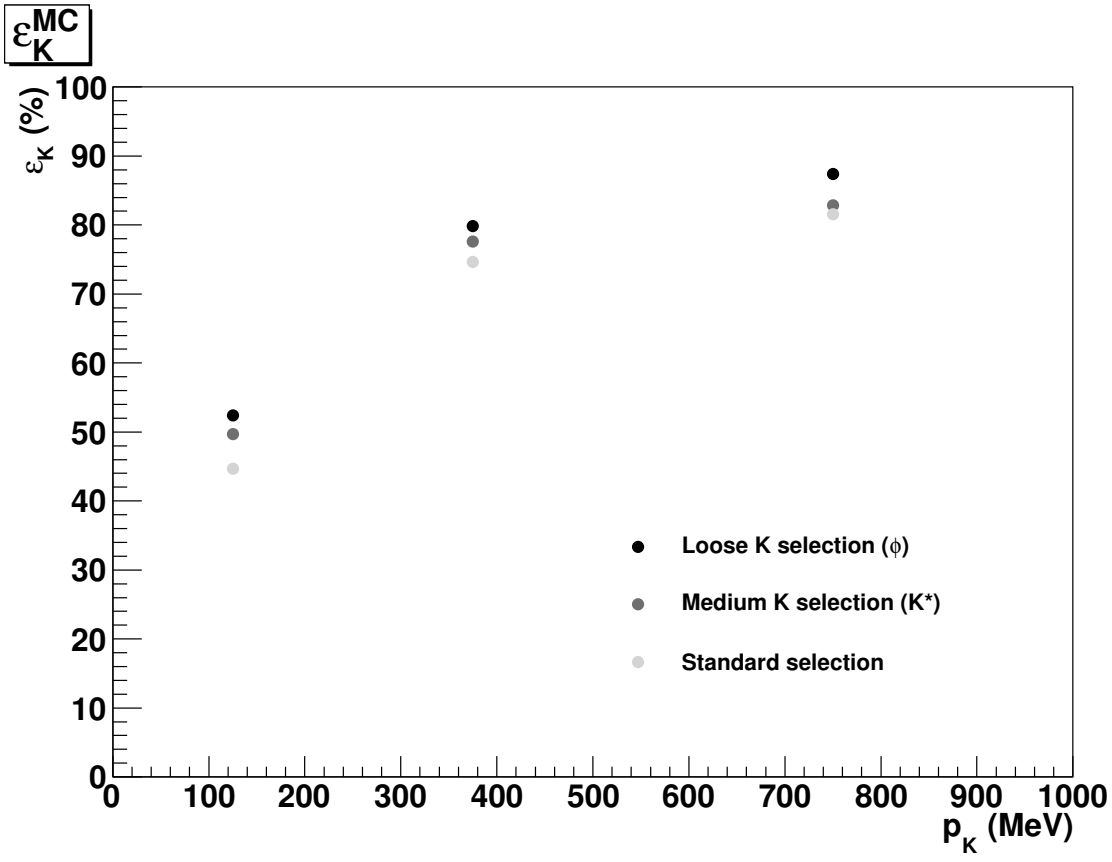

Figure D.3: Monte Carlo kaon efficiency for each set of kaon selections, by momentum. Our ϕ kaon cuts (hit fraction dropped) show a higher efficiency in each momentum range, with a particular relative advantage in the important low momentum region.

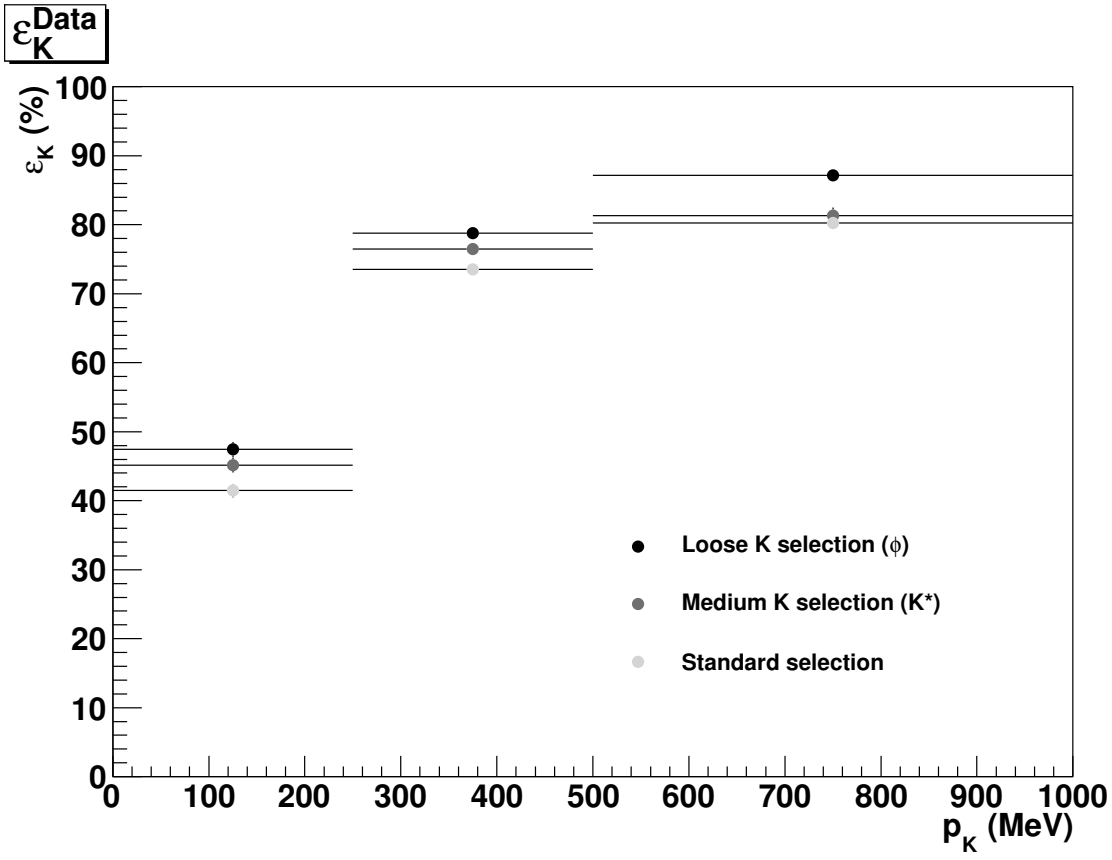

Figure D.4: Kaon efficiencies in the data for each set of kaon selections, by momentum. The error bars on the efficiencies (barely visible) include both a statistical error and the systematic error from fitting. The relative difference between the selection efficiencies roughly matches the Monte Carlo, although the absolute efficiencies for soft kaons all fall below their corresponding Monte Carlo efficiencies.

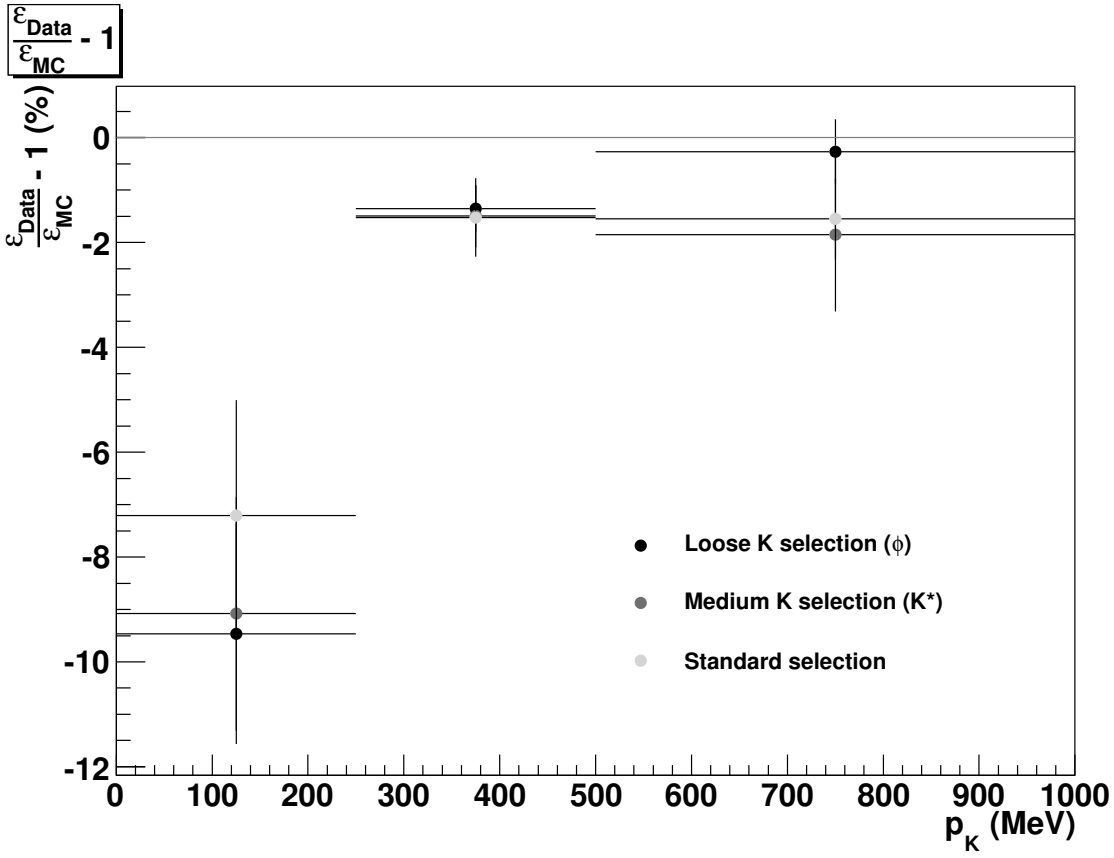

Figure D.5: The relative difference between the kaon efficiency in data and the kaon efficiency in Monte Carlo, by selections and momentum range. The high momentum region requires no correction, the middle momentum region requires a slight efficiency correction, and softest kaons require a sizable efficiency correction. The error bars include both statistical and systematic errors from our tracking/PID reconstruction.

Table D.1: Relative kaon systematic and efficiency correction for our loose ($\phi e\nu$) kaon selection.

Momentum region	Statistical error	Fit systematic	Total systematic	ε_K correction
$p_K < 250$ MeV	1.59%	1.69%	2.32%	-9.47%
250 MeV $< p_K <$ 500 MeV	0.43%	0.08%	0.44%	-1.35%
500 MeV $< p_K$	0.18%	0.59%	0.62%	N/A

Table D.2: Relative kaon systematic and efficiency correction for our medium ($K^*e\nu$) kaon selection.

Momentum region	Statistical error	Fit systematic	Total systematic	ε_K correction
$p_K < 250$ MeV	1.78%	1.69%	2.45%	-9.08%
250 MeV $< p_K <$ 500 MeV	0.60%	0.08%	0.60%	-1.50%
500 MeV $< p_K$	1.37%	0.59%	1.49%	N/A

Table D.3: Relative kaon systematic and efficiency correction for our standard kaon selection.

Momentum region	Statistical error	Fit systematic	Total systematic	ε_K correction
$p_K < 250$ MeV	1.67%	1.69%	2.38%	-7.21%
250 MeV $< p_K <$ 500 MeV	0.76%	0.08%	0.76%	N/A
500 MeV $< p_K$	0.51%	0.59%	0.78%	N/A

Our large momentum bins mean that the kaon momentum distribution within each bin (e.g. 250 MeV $< p_K <$ 500 MeV) can differ between $D \to K\pi\pi$ and the semileptonic mode, shown for $D_s \to \phi e\nu$ in Figure D.6. We add an additional systematic to the efficiency error for this effect by splitting the momentum bin into two halves and allowing each half of the bin to have a different efficiency correction, constrained by the bin's total kaon efficiency and the adjacent bins' efficiency corrections. This procedure results in a relative 0.5% systematic error for $D_s \to \phi e\nu$.

Since the ϕ meson's daughter kaons tend to have strongly correlated momenta, we correct our $\phi e\nu$ efficiency based on the daughter kaon momentum pairs rather than on the individual $\phi e\nu$ kaon momentum distribution.

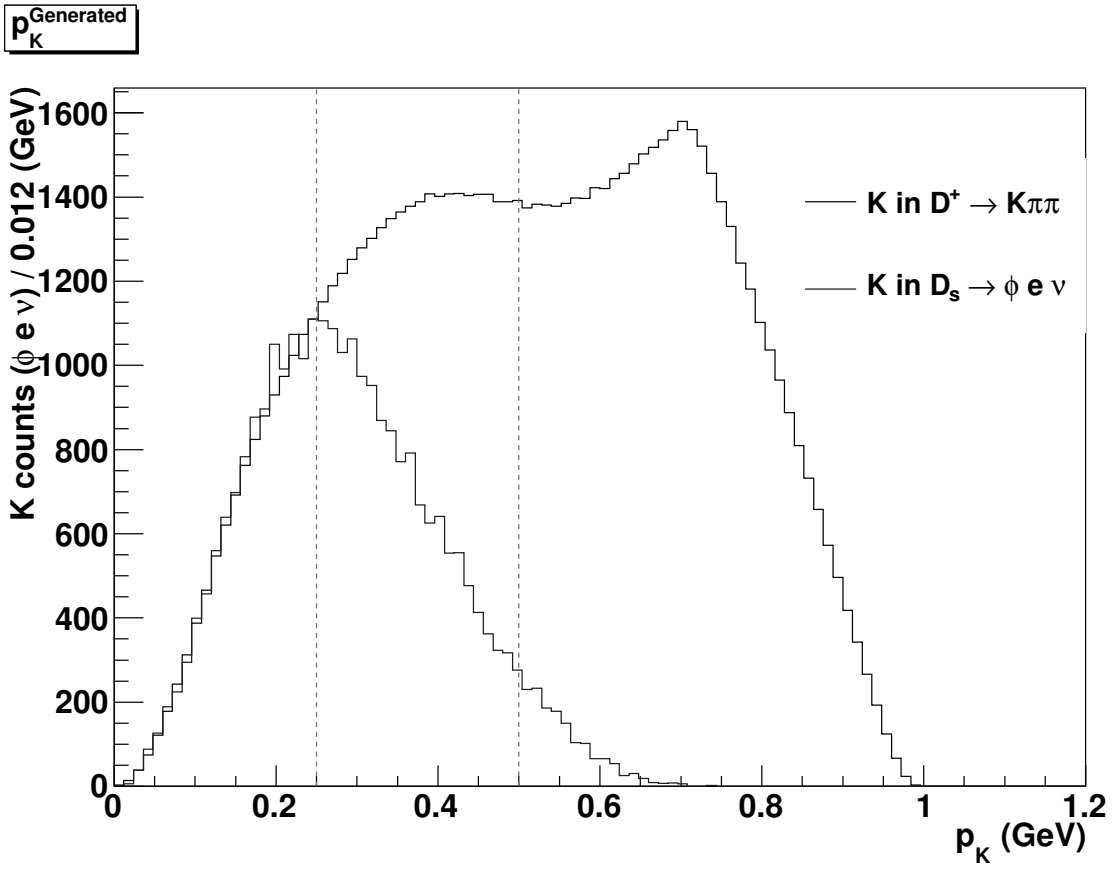

Figure D.6: $D \to K\pi\pi$ and $D_s \to \phi e\nu$ daughter kaon momentum distributions, scaled to line up the left momentum bin. Because the kaon momentum distribution for $D_s \to \phi e\nu$ falls off so sharply in the middle bin relative to $K\pi\pi$, we perform an additional systematic by splitting the bin into two halves and doing a separate efficiency correction for each.

Appendix E

Glossary

As a discipline becomes more and more specialized, it gains its own terminology. While I can't speak to the blechyuckiness of business speak, scientific disciplines really do require some arcane terminology because the English language frankly lacks words with the proper precision.[1] With this glossary, I intend to clarify some of the terms such that someone with a general science background doesn't feel totally lost (just partially).[2] I've also added a quick note on the particles relevant to this analysis.

E.1 General Terminology

- **Breit-Wigner Function** – A probability density function that describes the mass distribution for a particular resonance/particle. The relativistic version has the form:
$$\text{Prob}(M) = \frac{C}{(M^2 - M_0^2)^2 + M^2\Gamma^2},$$
where C is a normalization constant, M is the particle's mass, M_0 is the resonance's mass, and Γ is the resonance width.

- **Branching Ratio** – The fraction of a particle's decays that end in a given final

[1] Let alone the proper compactness. Without detailed terminology and implicit understanding, this dissertation's three word title could easily become "Measuring the rate at which an atom-like object made up of a charm and a strange quark transitions into other particles while giving off an electron and a neutrino, compared to how often that same atom-like object transitions into all other particles."
[2] I'm using my Dad, an electrical engineer, to calibrate the target audience for this section.

state, often expressed as a percentage. This is the same as the decay's partial width divided by the particle's total width.

- **Bremsstrahlung** – "Braking radiation," emitted when a charged particle slows down (decelerates) or deflects.

- **Cherenkov (Čerenkov) Radiation** – Light emitted in a cone when a charged particle moves faster than the speed of light (phase velocity) in a material.

- **CKM Matrix** – The Cabibbo-Kobayahsi-Maskawa matrix that connects the free quark eigenstates to states that take part in the weak interaction. Matrix elements along the diagonal (V_{ud}, V_{cs}, V_{tb}) have the largest magnitudes, and interactions that involve those elements are called *Cabibbo favored*, while interactions involving the off-diagonal elements are *Cabibbo suppressed*.

- **Combinatoric Background** – The e^+e^- collisions create all kinds of charm events that we don't care about (background). Sometimes through sheer numbers, these background events will happen to have the same particles we're looking for with just the right combinations of energy and momenta to fake a signal event. Every extra particle requirement tends to reduce this background, and it can usually be estimated from sideband regions.

- **Continuum** – e^+e^- collision events that don't generate the targeted particle states (nonresonant). While e^+e^- collisions can make a charm-anticharm pair, some of which split into the D_s^* we want, they can also make up-antiup, down-antidown, or strange-antistrange quark pairs that we don't care about in the slightest (continuum).

- **Cross Section** – The quantum mechanical analogue of the classical area exposed by an object to a third dimension. In particle physics, the cross section relates to the interaction probability but carries units of area, typically measured in barns (10^{-28} m^2).

- **Crystal Ball Function** – A function named after the Crystal Ball experiment that splices a Gaussian together with a power law tail, keeping the function and

first derivative continuous. We usually use these in the context of photons detected in a calorimeter, as their measured energies tend to exhibit a power law tail.

- $\mathbf{d_b}$ – Radial distance from the beamline. We use this as a constraint on our drift chamber tracks to ensure that the particle started from the e^+e^- collision point.

- **dE/dx** – Energy lost due to ionization per length by a charged particle moving through material. Particles with different masses have different amounts of dE/dx at a given momentum, allowing us to identify particles by comparing our measured dE/dx to the particle's ideal dE/dx ($\sigma^{dE/dx}$).

- $\frac{\mathbf{E_9}}{\mathbf{E_{25}}}$ – Energy in a shower's $3 \times 3 = 9$ central crystals divided by the energy in a shower's $5 \times 5 = 25$ central crystals. We often discuss an $\frac{E_9}{E_{25}}$ O.K. cut, which scales smoothly from requiring the center crystals to hold about 80% of the shower energy for soft photons to around 90% of the shower energy for higher energy photons. Photon showers show this pattern, while hadronic showers tend to be more spread out.

- **Flight Significance** – A K_s variable that ensures the K_s travels a minimum distance from the interaction point before decaying. This helps distinguish K_s from nonresonant $\pi\pi$ or random combinations.

- **Form Factor** – The Fourier transform (in momentum space) of a hadron's spatial structure. This relates to both the physical size of the hadron and its decays to particles with various spin/parity characteristics.

- **FSR** – Final state radiation, where a final state charged particle emits a soft photon, decreasing its energy and momentum.

- **Hit Fraction** – The number of actual drift chamber wire "hits" in a charged particle's track relative to the expected number of hits. Good tracks tend to have a hit fraction close to 1.0, but particles that decay in flight don't make it all the way through the drift chamber and so leave fewer hits than expected.

- **Hot Crystal** – A crystal/photodiode in our calorimeter known to register signals when we have none. We typically throw out signals from such crystals in our analysis.

- **ISR** – Initial state radiation, where one of the colliding e^+ or e^- emits a soft photon, lowering the collision's energy.

- **Lattice QCD** – A formulation of quantum chromodynamics (QCD), the theory of strong interactions, on a discrete space-time lattice. QCD calculations tend to be intractable analytically, but the lattice formalism brings computers into play by making the calculation precision dependent on the lattice size and spacing. This allows the result precision to depend on the computing power available, although lattice calculations still frequently need additional simplifying assumptions.

- **Lineshape** – Shorthand for the mass lineshape, a particle's mass distribution (e.g. a *Breit-Wigner* function).

- **Luminosity** – Interactions per area per time. We often integrate the luminosity over the run time for an experiment, giving the *integrated luminosity*, the interactions per area. We can then just multiply by an interaction's cross section (an area) to get the number of such events.

- **Missing Mass/Recoil Mass** – The mass created by missing energy and missing momentum via $E^2 = p^2 + m^2$. The missing energy and momentum comes from deducting all the reconstructed energy and momenta from the initial energy and momenta. If the reconstruction left exactly one missing particle, the missing/recoil mass should match up with that particle's mass. Informally, we tend to use "missing mass" when we don't intend to look for the particle and "recoil mass" when we expect the particle to have left some trace that may factor in the later analysis.

- **Mixing Angle** – A way of describing the connection between a particle's state and the linear combination of states that make it up. The component states' coefficients add in quadrature to one, much like cosine and sine, making the notion of an angle a useful way to describe the coefficients. For example, the π^0 is a combination of up-antiup and down-antidown states in the form $|\pi^0\rangle = \cos\theta\,|u\bar{u}\rangle + \sin\theta\,|d\bar{d}\rangle$, where $\theta = -45\,^\circ$.

- **Monte Carlo** – A computer generating technique where you set up the general rules (the physics) but use random numbers to construct the particular events.

Typically, we run many more simulated events than we have real events to investigate any biases or flaws in methodology.

- **Nonresonant** – Decays to final states that don't go through intermediate particles on the way. For instance, this analysis considers the resonant $D_s \to \phi e\nu$, where $\phi \to KK$ and we observe the two kaons. In principle, we could also look for nonresonant $D_s \to KKe\nu$, although these decays happen to be less favored.

- **Pole Dominance** – Tendency of a particular interaction to favor an intermediate resonance, or pole. The pole has the associated particle's mass and other characteristics (propagator).

- **Pull Mass** – Reconstructed mass precision, given the error in constituent particles. We usually use this in the context of $\pi^0 \to \gamma\gamma$ or $\eta \to \gamma\gamma$, where each photon may have a very different associated error. In this case, we want to incorporate those errors into our mass reconstruction precision rather than just taking the energies without error and calculating a nominal mass.

- **Q Value** – The energy difference between the decaying particle's mass and the rest masses of the final state particles, generally in the decaying particle's rest frame. This corresponds to the total kinetic energy released in the decay.

- **QCD** – Quantum chromodynamics, the theory that describes strong force interactions.

- **Recoil Mass** – See *missing mass*.

- **RF Cavity** – A radio frequency cavity with an oscillating electric field timed to accelerate charged particles when they pass through (here, e^+ and e^-). Somewhat self-correcting: a particle with too much energy has a larger radius and arrives at the cavity late, gaining less energy since it's not timed as precisely with the oscillating electric field.

- **Semileptonic Decay** – A decay that includes both leptons (e.g. an electron and an antineutrino) and hadrons (e.g. a ϕ meson) in its ending state. We call decays with only leptons *leptonic*, and we call decays with only hadrons *hadronic*.

- **Sideband** – Reconstructed events that miss some cut window and fall just outside the signal region. These kind of events give an idea of how much random background extends into the signal region (combinatoric background, background with one misreconstructed particle, etc.).

- **Spin** – Intrinsic angular momentum of a particle. When we're discussing a composite particle, we use *spin* to mean the particle's total spin (J) rather than its constituent particles' angular momentum (L) or spin (S).

- **Splitoff Showers** – When particles, particularly kaons, hit the crystals in our calorimeter, they induce reactions that sometimes involve "backscatter," where some of the resultant particles get shunted off and interact with nearby crystals. This looks like a new hit in our calorimeter without an associated drift chamber path, which we usually (and here, incorrectly) interpret as a photon.

- **Synchrotron Radiation** – Photon emission from a charged particle (here, e^+ and e^-) when accelerated into a circle (radially), since accelerating charged particles can radiate.

- **Systematic Errors** – Errors (imprecision or biases) due to the measurement procedure/apparatus rather than statistical fluctuations.

- **Tag (D_s tag)** – A reconstruction that identifies the right kind of event. We want only semileptonic D_s decays from the whole mess produced by e^+e^- collisions at 4170 MeV. Since D_s get produced in pairs ($D_s^* D_s$), we find one D_s in the event through a complete reconstruction; this lets us know that there's another D_s in the event and that we should look for a semileptonic decay.

- **$\mathbf{z_0}$** – Longitudinal (*axial*) distance from the e^+e^- collision point. We use this as a constraint on our drift chamber tracks to ensure that the particle started from the interaction region.

E.2 Particle Zoo

Particles and related terms used in this dissertation.

- **Baryon** – Particle consisting of three quarks (all three quarks or all three antiquarks).
- **Boson** – Particle with integer spin. Multiple bosons can exist in the same quantum state. The force carriers (e.g. photon, gluon) are all bosons.
- **Fermion** – Particle with half-integer spin (typically spin-$\frac{1}{2}$). No two fermions can exist in the same quantum state (the exclusion principle).
- **Hadron** – General term for a composite particle consisting of multiple quarks.
- **Meson** – Hadron consisting of two quarks (a quark and an antiquark).
- **Pseudoscalar (Scalar) Particle** – A spin-0 meson whose quantum state flips sign (does not flip sign) under a parity inversion ($\vec{x} \to -\vec{x}$).
- **Quark** – Spin-$\frac{1}{2}$ elementary particles that have color and so interact under the strong force. We sometimes separate these into light quarks (up, down, and strange) with masses less than the characteristic strong force scale, Λ_{QCD}, and heavy quarks (charm, bottom, and top) with masses above Λ_{QCD}.
- **Vector (Axial) Particle** – A spin-1 meson whose quantum state flips sign (does not flip sign) under a parity inversion ($\vec{x} \to -\vec{x}$). Axial vector mesons are also sometimes called pseudovector mesons.
- D_s – The lowest energy bound state formed by a charm quark and an antistrange quark (or an anticharm quark and a strange quark). It has a mean lifetime of a whopping 5.00×10^{-13} seconds and travels a good 3 micrometers in our detector before breaking up, both of which are actually quite long compared to some other particles we deal with.
- D_s^* – The next-to-lowest energy bound state formed by a charm and an antistrange quark (or anticharm/strange). Unlike the spin-0 D_s, it has spin-1, which can *roughly* be thought of as having the charm and strange spins aligned instead of antialigned. It quickly transitions to a D_s in a *spin-flip* decay via a photon emission.

- ϕ – Vector meson consisting of a $s\bar{s}$ state. The ϕ usually decays to KK, but its mass sits just above the KK threshold. It has a relatively wide mass distribution ($\Gamma = 4.26$ MeV) with a long, high-side tail.

- η – Pseudoscalar mixture of $u\bar{u}$, $d\bar{d}$, and $s\bar{s}$, with a plurality decay to $\gamma\gamma$. The η mostly consists of the uds meson octet's η_8.

- η' – Pseudoscalar mixture of $u\bar{u}$, $d\bar{d}$, and $s\bar{s}$. Heavier than the η, the η' mostly consists of the $\eta_1 = \frac{u\bar{u}+d\bar{d}+s\bar{s}}{\sqrt{3}}$ singlet. The η' meson has a plurality decay to $\pi\pi\eta$.

- K_s – The shorter-lived neutral kaon (lifetime of 9×10^{-11} s) that decays to $\pi^+\pi^-$ just over 2/3 of the time. The pseudoscalar K_s consists of a down and strange quark, roughly in the form $\frac{d\bar{s}-s\bar{d}}{\sqrt{2}}$.

- K^* – The neutral vector K^* meson consists of a $d\bar{s}/\bar{d}s$ state that decays to Kπ. Charged K^* mesons also exist ($u\bar{s}/\bar{u}s$), with the K^* charge usually clear from context. Our semileptonic decays involve the neutral K^*.

- f_0 – A scalar state with unknown composition, probably with an $s\bar{s}$ component. It has also been thought to have a $K\bar{K}$ molecule, four quark, or gluon component.

- ρ^0 – A wide, vector meson resonance that appears frequently and quickly decays to $\pi^+\pi^-$. The related $\rho^\pm$ decays to $\pi^\pm\pi^0$.

Appendix F

Extra Tables

Table F.1: Number of D_s tags in data and Monte Carlo, by dataset. We fit each dataset independently for this comparison and scale the Monte Carlo to data size.

Dataset	Data fit counts	MC fit counts	$\frac{N_{fit}-N_{MC}}{N_{MC}}$
39	7,246.8 ± 255.3	6,482.8 ± 308.6	(11.8 ± 6.6)%
40	15,609.7 ± 414.2	14,278.5 ± 422.0	(9.3 ± 4.3)%
41	16,308.6 ± 443.9	13,886.4 ± 438.2	(17.4 ± 4.9)%
47	14,686.3 ± 408.9	12,940.5 ± 396.6	(13.5 ± 4.7)%
48	23,823.5 ± 752.0	20,283.1 ± 580.2	(17.5 ± 5.0)%
Sum	77,674.9 ± 2,274.3	67,871.3 ± 2,145.6	(14.4 ± 4.9)%

Table F.2: Number of D_s tags in data and Monte Carlo, by mode. We scale the Monte Carlo to the data luminosity.

D_s mode	Data fit counts	MC fit counts	$\frac{N_{fit}-N_{MC}}{N_{MC}}$
K_sK	6,226.7 ± 101.2	5,764.0 ± 100.8	(8.0 ± 2.6)%
$KK\pi$	27,373.5 ± 248.4	25,242.0 ± 233.9	(8.4 ± 1.4)%
$K_sK\pi^0$	2,246.8 ± 209.9	1,670.5 ± 157.7	(34.5 ± 17.9)%
$K_sK_s\pi$	1,125.5 ± 76.5	1,141.4 ± 69.3	(-1.4 ± 9.0)%
$KK\pi\pi^0$	7,355.5 ± 377.4	6,693.4 ± 323.6	(9.9 ± 7.7)%
$K_sK^+\pi\pi$	1,859.4 ± 120.6	1,744.1 ± 105.5	(6.6 ± 9.5)%
$K_sK^-\pi\pi$	3,377.3 ± 100.0	3,246.3 ± 92.2	(4.0 ± 4.3)%
$\pi\pi\pi$	6,606.3 ± 337.7	6,081.6 ± 326.3	(8.6 ± 8.0)%
$\pi\eta$	3,810.3 ± 190.8	2,882.3 ± 182.9	(32.2 ± 10.7)%
$\pi\pi^0\eta$	9,476.9 ± 529.0	6,825.9 ± 700.7	(38.8 ± 16.2)%
$\pi\eta', \eta' \to \pi\pi\eta$	2,386.6 ± 65.6	2,132.4 ± 64.3	(11.9 ± 4.6)%
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	1,090.5 ± 118.7	532.5 ± 84.5	(104.8 ± 39.4)%
$\pi\eta', \eta' \to \rho\gamma$	4,272.3 ± 193.3	3,904.4 ± 245.2	(9.4 ± 8.5)%
Sum	77,207.5 ± 880.2	67,860.7 ± 959.8	(13.8 ± 2.1)%

Table F.3: Test of potential bias in our fitting procedure for $D_s \to K_s e\nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	27	23.95 ± 5.72	−0.533
2	25	26.86 ± 6.08	0.306
3	37	39.67 ± 7.79	0.343
4	35	28.34 ± 6.15	−1.083
5	23	25.55 ± 5.94	0.429
6	23	19.51 ± 5.79	−0.603
7	18	20.18 ± 5.91	0.368
8	29	23.20 ± 5.90	−0.982
9	26	21.46 ± 6.59	−0.689
10	29	26.82 ± 6.19	−0.353
11	20	22.74 ± 5.78	0.474
12	22	23.99 ± 5.88	0.339
13	22	27.44 ± 6.13	0.888
14	22	18.36 ± 5.09	−0.715
15	37	31.85 ± 6.52	−0.790
16	28	32.67 ± 6.59	0.710
17	25	17.67 ± 5.99	−1.224
18	26	26.27 ± 6.26	0.043
19	20	23.32 ± 5.61	0.592
20	27	23.80 ± 5.83	−0.549
Sum	521	503.66 ± 27.31	−0.635

Table F.4: Monte Carlo comparison of the measured $D_s \to K_s e\nu$ branching ratio to its generating branching ratio (0.090%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios.

Dataset	Raw fit BR	Corrected BR	# σ
1	$(0.113 \pm 0.027)\%$	$(0.100 \pm 0.027)\%$	0.36
2	$(0.122 \pm 0.028)\%$	$(0.109 \pm 0.028)\%$	0.68
3	$(0.178 \pm 0.035)\%$	$(0.165 \pm 0.035)\%$	2.13
4	$(0.129 \pm 0.028)\%$	$(0.115 \pm 0.028)\%$	0.89
5	$(0.118 \pm 0.027)\%$	$(0.104 \pm 0.027)\%$	0.52
6	$(0.090 \pm 0.027)\%$	$(0.077 \pm 0.027)\%$	−0.49
7	$(0.092 \pm 0.027)\%$	$(0.078 \pm 0.027)\%$	−0.44
8	$(0.105 \pm 0.027)\%$	$(0.091 \pm 0.027)\%$	0.04
9	$(0.100 \pm 0.031)\%$	$(0.086 \pm 0.031)\%$	−0.13
10	$(0.122 \pm 0.028)\%$	$(0.108 \pm 0.028)\%$	0.64
11	$(0.106 \pm 0.027)\%$	$(0.093 \pm 0.027)\%$	0.09
12	$(0.110 \pm 0.027)\%$	$(0.096 \pm 0.027)\%$	0.24
13	$(0.124 \pm 0.028)\%$	$(0.110 \pm 0.028)\%$	0.73
14	$(0.081 \pm 0.023)\%$	$(0.068 \pm 0.023)\%$	−0.98
15	$(0.148 \pm 0.030)\%$	$(0.134 \pm 0.030)\%$	1.46
16	$(0.153 \pm 0.031)\%$	$(0.139 \pm 0.031)\%$	1.60
17	$(0.080 \pm 0.027)\%$	$(0.067 \pm 0.027)\%$	−0.85
18	$(0.119 \pm 0.028)\%$	$(0.105 \pm 0.028)\%$	0.54
19	$(0.104 \pm 0.025)\%$	$(0.091 \pm 0.025)\%$	0.03
20	$(0.109 \pm 0.027)\%$	$(0.096 \pm 0.027)\%$	0.21
Weighted averages/χ^2	$(0.112 \pm 0.006)\%$	$(0.099 \pm 0.006)\%$	14.38

Table F.5: Test of potential bias in our fitting procedure for $D_s \to K^* e\nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	28	27.80 ± 6.33	−0.032
2	37	33.39 ± 5.24	−0.689
3	35	35.63 ± 7.18	0.087
4	28	29.57 ± 6.48	0.242
5	33	34.40 ± 6.73	0.208
6	34	29.57 ± 6.72	−0.660
7	33	26.81 ± 6.75	−0.916
8	35	35.68 ± 6.96	0.098
9	27	26.01 ± 6.83	−0.144
10	38	40.91 ± 7.10	0.411
11	36	30.72 ± 6.98	−0.756
12	33	38.05 ± 7.58	0.667
13	37	42.45 ± 7.28	0.749
14	28	30.47 ± 6.74	0.367
15	38	31.65 ± 6.73	−0.945
16	29	24.58 ± 6.64	−0.665
17	31	34.54 ± 6.63	0.534
18	24	25.21 ± 5.65	0.214
19	29	29.34 ± 6.52	0.052
20	30	24.11 ± 5.97	−0.987
Sum	643	630.89 ± 29.84	−0.406

Table F.6: Monte Carlo comparison of the measured $D_s \to K^* e\nu$ branching ratio to its generating branching ratio (0.190%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios.

Dataset	Raw fit BR	Corrected BR	# σ
1	$(0.169 \pm 0.038)\%$	$(0.157 \pm 0.038)\%$	-0.86
2	$(0.195 \pm 0.031)\%$	$(0.183 \pm 0.031)\%$	-0.22
3	$(0.205 \pm 0.041)\%$	$(0.193 \pm 0.041)\%$	0.08
4	$(0.172 \pm 0.038)\%$	$(0.160 \pm 0.038)\%$	-0.79
5	$(0.204 \pm 0.040)\%$	$(0.192 \pm 0.040)\%$	0.04
6	$(0.176 \pm 0.040)\%$	$(0.164 \pm 0.040)\%$	-0.66
7	$(0.156 \pm 0.039)\%$	$(0.145 \pm 0.039)\%$	-1.15
8	$(0.206 \pm 0.040)\%$	$(0.194 \pm 0.040)\%$	0.11
9	$(0.155 \pm 0.041)\%$	$(0.143 \pm 0.041)\%$	-1.15
10	$(0.238 \pm 0.041)\%$	$(0.226 \pm 0.041)\%$	0.87
11	$(0.184 \pm 0.042)\%$	$(0.172 \pm 0.042)\%$	-0.43
12	$(0.224 \pm 0.045)\%$	$(0.212 \pm 0.045)\%$	0.49
13	$(0.245 \pm 0.042)\%$	$(0.234 \pm 0.042)\%$	1.03
14	$(0.173 \pm 0.038)\%$	$(0.161 \pm 0.038)\%$	-0.75
15	$(0.188 \pm 0.040)\%$	$(0.176 \pm 0.040)\%$	-0.34
16	$(0.148 \pm 0.040)\%$	$(0.136 \pm 0.040)\%$	-1.36
17	$(0.202 \pm 0.039)\%$	$(0.190 \pm 0.039)\%$	-0.01
18	$(0.146 \pm 0.033)\%$	$(0.134 \pm 0.033)\%$	-1.69
19	$(0.168 \pm 0.037)\%$	$(0.156 \pm 0.037)\%$	-0.90
20	$(0.142 \pm 0.035)\%$	$(0.130 \pm 0.035)\%$	-1.72
Weighted averages/χ^2	$(0.183 \pm 0.009)\%$	$(0.171 \pm 0.009)\%$	15.92

Table F.7: Test of potential bias in our fitting procedure for $D_s \to \eta' e\nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	32	29.34 ± 5.77	−0.461
2	10	8.33 ± 4.01	−0.417
3	27	21.86 ± 5.02	−1.023
4	24	25.28 ± 5.52	0.231
5	15	13.87 ± 4.29	−0.264
6	19	17.85 ± 4.34	−0.265
7	32	33.64 ± 5.46	0.300
8	23	21.90 ± 4.88	−0.226
9	23	22.66 ± 5.42	−0.062
10	16	11.48 ± 4.08	−1.109
11	27	27.47 ± 5.37	0.087
12	16	15.18 ± 4.70	−0.176
13	22	21.65 ± 4.84	−0.073
14	26	24.09 ± 4.96	−0.386
15	18	15.75 ± 4.38	−0.513
16	17	20.49 ± 4.74	0.736
17	19	20.03 ± 4.15	0.249
18	23	20.65 ± 4.94	−0.476
19	27	28.88 ± 5.35	0.351
20	31	22.83 ± 5.41	−1.509
Sum	447	423.20 ± 21.96	−1.084

Table F.8: Monte Carlo comparison of the measured $D_s \to \eta' e\nu$ branching ratio to its generating branching ratio (0.860%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios. The abnormally high χ^2 just reflects low $D_s \to \eta' e\nu$ statistics that distort gaussian error sums (Table F.7 gives a more meaningful comparison for this mode).

Dataset	Raw fit BR	Corrected BR	# σ
1	$(1.068 \pm 0.210)\%$	$(1.057 \pm 0.210)\%$	0.94
2	$(0.292 \pm 0.141)\%$	$(0.281 \pm 0.141)\%$	−4.12
3	$(0.755 \pm 0.173)\%$	$(0.744 \pm 0.173)\%$	−0.67
4	$(0.881 \pm 0.193)\%$	$(0.871 \pm 0.193)\%$	0.05
5	$(0.492 \pm 0.152)\%$	$(0.481 \pm 0.152)\%$	−2.48
6	$(0.636 \pm 0.154)\%$	$(0.625 \pm 0.155)\%$	−1.52
7	$(1.177 \pm 0.191)\%$	$(1.166 \pm 0.191)\%$	1.60
8	$(0.760 \pm 0.169)\%$	$(0.749 \pm 0.169)\%$	−0.66
9	$(0.809 \pm 0.194)\%$	$(0.799 \pm 0.194)\%$	−0.32
10	$(0.400 \pm 0.142)\%$	$(0.389 \pm 0.142)\%$	−3.31
11	$(0.986 \pm 0.193)\%$	$(0.975 \pm 0.193)\%$	0.60
12	$(0.535 \pm 0.165)\%$	$(0.524 \pm 0.165)\%$	−2.03
13	$(0.751 \pm 0.168)\%$	$(0.740 \pm 0.168)\%$	−0.72
14	$(0.821 \pm 0.169)\%$	$(0.810 \pm 0.169)\%$	−0.29
15	$(0.562 \pm 0.156)\%$	$(0.551 \pm 0.157)\%$	−1.97
16	$(0.738 \pm 0.171)\%$	$(0.727 \pm 0.171)\%$	−0.78
17	$(0.701 \pm 0.145)\%$	$(0.690 \pm 0.145)\%$	−1.17
18	$(0.719 \pm 0.172)\%$	$(0.708 \pm 0.172)\%$	−0.88
19	$(0.993 \pm 0.184)\%$	$(0.982 \pm 0.184)\%$	0.66
20	$(0.805 \pm 0.191)\%$	$(0.794 \pm 0.191)\%$	−0.35
Weighted averages/χ^2	$(0.702 \pm 0.038)\%$	$(0.691 \pm 0.038)\%$	53.13

Table F.9: Test of potential bias in our fitting procedure for $D_s \to f_0 e\nu$ by comparing the number of truth-tagged semileptonic events to the fit result. We allow cross-feed from other semileptonic modes for this fitting comparison, as those events produce real peaking background that we deal with outside the fitting apparatus.

Dataset	$N_{SL+tag}^{\text{truth-tagged}}$	N_{SL+tag}^{fit}	Difference (# σ)
1	56	62.93 ± 9.28	0.747
2	52	45.11 ± 8.09	−0.851
3	49	45.96 ± 7.14	−0.425
4	51	42.88 ± 8.17	−0.994
5	57	48.99 ± 7.61	−1.053
6	51	53.61 ± 8.27	0.316
7	52	48.57 ± 8.07	−0.425
8	49	44.73 ± 8.02	−0.533
9	56	57.19 ± 8.96	0.133
10	46	40.18 ± 7.58	−0.769
11	70	77.54 ± 9.43	0.799
12	42	43.89 ± 7.42	0.255
13	45	43.39 ± 8.09	−0.199
14	59	65.11 ± 9.08	0.674
15	53	52.42 ± 7.20	−0.080
16	53	63.73 ± 8.71	1.232
17	63	61.51 ± 9.16	−0.163
18	56	53.97 ± 8.40	−0.242
19	48	53.70 ± 8.33	0.684
20	46	41.28 ± 7.31	−0.647
Sum	1054	1046.68 ± 36.87	−0.198

Table F.10: Monte Carlo comparison of the measured $D_s \to f_0 e\nu$ branching ratio to its generating branching ratio (0.310%), in data-sized samples. The weighted average line contains the χ^2 across the 20 samples rather than the number of σ between the measured/generated branching ratios.

Dataset	Raw fit BR	Corrected BR	# σ
1	$(0.425 \pm 0.063)\%$	$(0.413 \pm 0.063)\%$	1.65
2	$(0.293 \pm 0.053)\%$	$(0.282 \pm 0.053)\%$	−0.54
3	$(0.294 \pm 0.046)\%$	$(0.283 \pm 0.046)\%$	−0.59
4	$(0.277 \pm 0.053)\%$	$(0.266 \pm 0.053)\%$	−0.83
5	$(0.323 \pm 0.050)\%$	$(0.311 \pm 0.050)\%$	0.03
6	$(0.354 \pm 0.055)\%$	$(0.343 \pm 0.055)\%$	0.60
7	$(0.315 \pm 0.052)\%$	$(0.304 \pm 0.052)\%$	−0.12
8	$(0.288 \pm 0.052)\%$	$(0.276 \pm 0.052)\%$	−0.65
9	$(0.379 \pm 0.059)\%$	$(0.368 \pm 0.059)\%$	0.97
10	$(0.260 \pm 0.049)\%$	$(0.248 \pm 0.049)\%$	−1.25
11	$(0.516 \pm 0.063)\%$	$(0.505 \pm 0.063)\%$	3.10
12	$(0.287 \pm 0.048)\%$	$(0.276 \pm 0.049)\%$	−0.71
13	$(0.279 \pm 0.052)\%$	$(0.268 \pm 0.052)\%$	−0.81
14	$(0.412 \pm 0.057)\%$	$(0.401 \pm 0.057)\%$	1.58
15	$(0.347 \pm 0.048)\%$	$(0.336 \pm 0.048)\%$	0.54
16	$(0.426 \pm 0.058)\%$	$(0.414 \pm 0.058)\%$	1.79
17	$(0.399 \pm 0.059)\%$	$(0.388 \pm 0.059)\%$	1.31
18	$(0.349 \pm 0.054)\%$	$(0.337 \pm 0.054)\%$	0.50
19	$(0.343 \pm 0.053)\%$	$(0.331 \pm 0.053)\%$	0.40
20	$(0.270 \pm 0.048)\%$	$(0.258 \pm 0.048)\%$	−1.08
Weighted averages/χ^2	$(0.333 \pm 0.012)\%$	$(0.321 \pm 0.012)\%$	27.42

Table F.11: Allowed D_s mass range at 3σ, from a gaussian fit. We allow a broader range of masses for the full analysis, but we use this restricted range for systematic checks.

D_s tag mode	Minimum D_s mass (MeV)	Maximum D_s mass (MeV)
K_sK	1,949.69	1,987.31
$KK\pi$	1,952.93	1,984.07
$K_sK\pi^0$	1,941.32	1,995.68
$K_sK_s\pi$	1,951.94	1,985.06
$KK\pi\pi^0$	1,944.48	1,992.52
$K_sK^+\pi\pi$	1,953.76	1,983.24
$K_sK^-\pi\pi$	1,953.60	1,983.40
$\pi\pi\pi$	1,948.80	1,988.20
$\pi\eta$	1,934.89	2,002.11
$\pi\pi^0\eta$	1,930.60	2,006.40
$\pi\eta', \eta' \to \pi\pi\eta$	1,945.81	1,991.19
$\pi\pi^0\eta', \eta' \to \pi\pi\eta$	1,939.39	1,997.61
$\pi\eta', \eta' \to \rho\gamma$	1,938.02	1,998.98

Table F.12: Summary of various systematic errors for our electron identification.

SL mode	Base eID systematic	Base eID correction uncert.	Base eID correction adjust.	Event environ. correct.	Event environ. uncert.	Total systematic
$\phi e\nu$	0.23%	0.03%	0.39%	0.48%	0.16%	0.68%
$\eta e\nu$	0.10%	0.03%	0.28%	0.20%	0.10%	0.37%
$\eta' e\nu$	0.16%	0.03%	0.35%	0.42%	0.14%	0.59%
$f_0 e\nu$	0.17%	0.03%	0.34%	0.42%	0.14%	0.59%
$K_s e\nu$	0.13%	0.03%	0.28%	0.19%	0.11%	0.38%
$K^* e\nu$	0.18%	0.03%	0.34%	0.44%	0.14%	0.60%

Table F.13: Relative corrections to the electron identification efficiency for each of our six semileptonic modes.

Semileptonic mode	Base electron ID	Event environment	Total correction
$\phi e\nu$	-1.55%	-0.36%	-1.91%
$\eta e\nu$	-1.10%	-0.14%	-1.24%
$\eta' e\nu$	-1.38%	-0.32%	-1.71%
$f_0 e\nu$	-1.37%	-0.32%	-1.69%
$K_s e\nu$	-1.12%	-0.13%	-1.24%
$K^* e\nu$	-1.36%	-0.33%	-1.69%

Table F.14: Semileptonic-side efficiencies in $D_s \to \eta' e\nu$, including the $\eta' \to \pi\pi\eta$ and $\eta \to \gamma\gamma$ branching ratios.

D_s production mode	ε_e	$\varepsilon_{\eta'}$	ε_{SL}
D_sD_s	75.5% ± 4.2%	7.1% ± 1.3%	5.5% ± 1.1%
$D_s^*D_s$ with $D_s^* \to (D_s \to \eta' e\nu)\ \gamma$	74.9% ± 1.0%	5.5% ± 0.3%	3.8% ± 0.2%
$D_s^*D_s$ with $D_s^* \to (D_s \to \eta' e\nu)\ \pi^0$	75.5% ± 1.1%	5.0% ± 0.3%	3.6% ± 0.2%
$D_s^*D_s$ with prompt $D_s \to \eta' e\nu$	74.0% ± 1.1%	6.2% ± 0.3%	4.4% ± 0.3%

Table F.15: Semileptonic-side efficiencies in $D_s \to f_0 e\nu$, including the $f_0 \to \pi\pi$ branching ratio.

D_s production mode	ε_e	ε_{f_0}	ε_{SL}
D_sD_s	78.7% ± 4.3%	32.7% ± 2.8%	24.6% ± 2.4%
$D_s^*D_s$ with $D_s^* \to (D_s \to f_0 e\nu)\ \gamma$	72.6% ± 1.0%	30.1% ± 0.6%	21.7% ± 0.6%
$D_s^*D_s$ with $D_s^* \to (D_s \to f_0 e\nu)\ \pi^0$	73.2% ± 1.0%	30.5% ± 0.7%	22.1% ± 0.6%
$D_s^*D_s$ with prompt $D_s \to f_0 e\nu$	72.7% ± 1.1%	29.7% ± 0.7%	21.6% ± 0.6%

Table F.16: Semileptonic-side efficiencies in $D_s \to K_s e\nu$, including the $K_s \to \pi\pi$ branching ratio.

D_s production mode	ε_e	ε_{K_s}	ε_{SL}
D_sD_s	81.7% ± 6.3%	44.2% ± 4.6%	33.2% ± 4.0%
$D_s^*D_s$ with $D_s^* \to (D_s \to K_s e\nu)\ \gamma$	80.5% ± 1.5%	41.8% ± 1.1%	30.6% ± 0.9%
$D_s^*D_s$ with $D_s^* \to (D_s \to K_s e\nu)\ \pi^0$	81.8% ± 1.5%	44.0% ± 1.1%	33.1% ± 1.0%
$D_s^*D_s$ with prompt $D_s \to K_s e\nu$	81.0% ± 1.6%	43.3% ± 1.2%	31.0% ± 1.0%

Table F.17: Semileptonic-side efficiencies in $D_s \to K^* e\nu$, including the $K^* \to K\pi$ branching ratio.

D_s production mode	ε_e	ε_{K^*}	ε_{SL}
D_sD_s	65.1% ± 4.1%	32.9% ± 2.9%	21.3% ± 2.3%
$D_s^*D_s$ with $D_s^* \to (D_s \to K^* e\nu)\ \gamma$	71.8% ± 1.0%	34.8% ± 0.7%	24.1% ± 0.6%
$D_s^*D_s$ with $D_s^* \to (D_s \to K^* e\nu)\ \pi^0$	71.6% ± 1.0%	35.7% ± 0.7%	24.8% ± 0.6%
$D_s^*D_s$ with prompt $D_s \to K^* e\nu$	71.6% ± 1.1%	35.1% ± 0.8%	24.0% ± 0.7%

Table F.18: All systematic efficiency corrections (relative) for $D_s \to \phi e\nu$.

Systematic	Relative ε correction
Kaon track efficiency	-8.17%
Electron ID	-1.91%
Total	-10.08%

Table F.19: All systematic efficiency corrections (relative) for $D_s \to \eta e\nu$.

Systematic	Relative ε correction
Electron ID	-1.24%
Total	-1.24%

Table F.20: All systematic efficiency corrections (relative) for $D_s \to \eta' e\nu$.

Systematic	Relative ε correction
Electron ID	-1.71%
π (and K) ID	-2.94%
Semileptonic hadron $\mathcal{B}$	-1.83%
Total	-6.48%

Table F.21: All systematic efficiency corrections (relative) for $D_s \to f_0 e\nu$.

Systematic	Relative ε correction
Electron ID	-1.69%
π (and K) ID	-0.50%
Total	-2.19%

Table F.22: All systematic efficiency corrections (relative) for $D_s \to K_s e\nu$.

Systematic	Relative ε correction
K_s efficiency	-11.08%
Electron ID	-1.24%
Semileptonic hadron $\mathcal{B}$	0.86%
Total	-11.46%

Table F.23: All systematic efficiency corrections (relative) for $D_s \to K^* e\nu$.

Systematic	Relative ε correction
Electron ID	-1.69%
π (and K) ID	-2.88%
Total	-4.57%

Table F.24: All systematic errors (relative) for $D_s \to \phi e\nu$.

Systematic	Relative systematic error
Distribution within p_K bin	0.54%
Kaon track reconstruction	1.71%
D_s tag signal fitting	1.80%
D_s tag BG shape	1.27%
Electron ID	0.68%
Form factor model	2.91%
Multiple candidate choice	0.11%
Semileptonic hadron $\mathcal{B}$	1.02%
Track reconstruction	0.30%
Efficiency statistics	1.33%
Total	4.46%

Table F.25: All systematic errors (relative) for $D_s \to \eta e\nu$.

Systematic	Relative systematic error
D_s tag signal fitting	2.23%
D_s tag BG shape	0.92%
Electron ID	0.37%
Form factor model	0.73%
Signal shape	1.04%
Multiple candidate choice	1.67%
Semileptonic hadron $\mathcal{B}$	0.51%
Particle ID	7.90%
Track reconstruction	0.30%
Efficiency statistics	1.07%
Splitoff rate	1.16%
Total	8.70%

Table F.26: All systematic errors (relative) for $D_s \to \eta' e\nu$.

Systematic	Relative systematic error
D_s tag signal fitting	2.07%
D_s tag BG shape	1.31%
Particle ID	7.90%
Track reconstruction	0.90%
Mass resolution	3.15%
Electron ID	0.59%
Form factor model	1.64%
Multiple candidate choice	0.21%
Decay in flight	0.49%
Semileptonic hadron $\mathcal{B}$	1.71%
Efficiency statistics	4.09%
Total	10.11%

Table F.27: All systematic errors (relative) for $D_s \to f_0 e\nu$.

Systematic	Relative systematic error
D_s tag signal fitting	1.60%
D_s tag BG shape	0.80%
Particle ID	0.04%
Track reconstruction	0.90%
Mass resolution	2.63%
Electron ID	0.59%
Form factor model	2.29%
Multiple candidate choice	2.20%
Decay in flight	0.52%
Efficiency statistics	1.57%
Total	4.91%

Table F.28: All systematic errors (relative) for $D_s \to K_s e\nu$.

Systematic	Relative systematic error
D_s tag signal fitting	2.20%
D_s tag BG shape	0.86%
K_s efficiency	7.28%
Electron ID	0.38%
Form factor model	1.35%
Multiple candidate choice	3.05%
Decay in flight	0.63%
Semileptonic hadron $\mathcal{B}$	0.07%
Track reconstruction	0.30%
Efficiency statistics	1.72%
Total	8.56%

Table F.29: All systematic errors (relative) for $D_s \to K^* e\nu$.

Systematic	Relative systematic error
D_s tag signal fitting	2.97%
D_s tag BG shape	2.08%
Particle ID	1.21%
Track reconstruction	0.60%
Mass resolution	2.59%
Electron ID	0.60%
Form factor model	5.10%
Multiple candidate choice	0.28%
Decay in flight	0.71%
Efficiency statistics	1.47%
Total	7.13%

Appendix G

Extra Figures

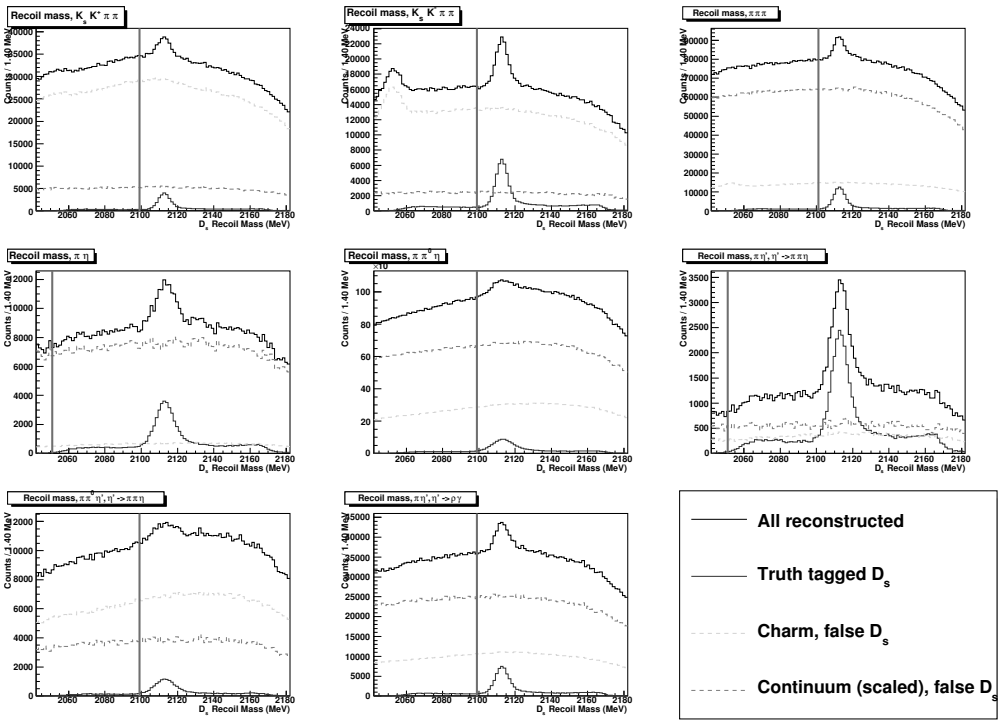

Figure G.1: Monte Carlo (charm and scaled continuum) simulation of the recoil mass distribution for the D_s tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. Again, we reject all D_s with a recoil mass below the cut line.

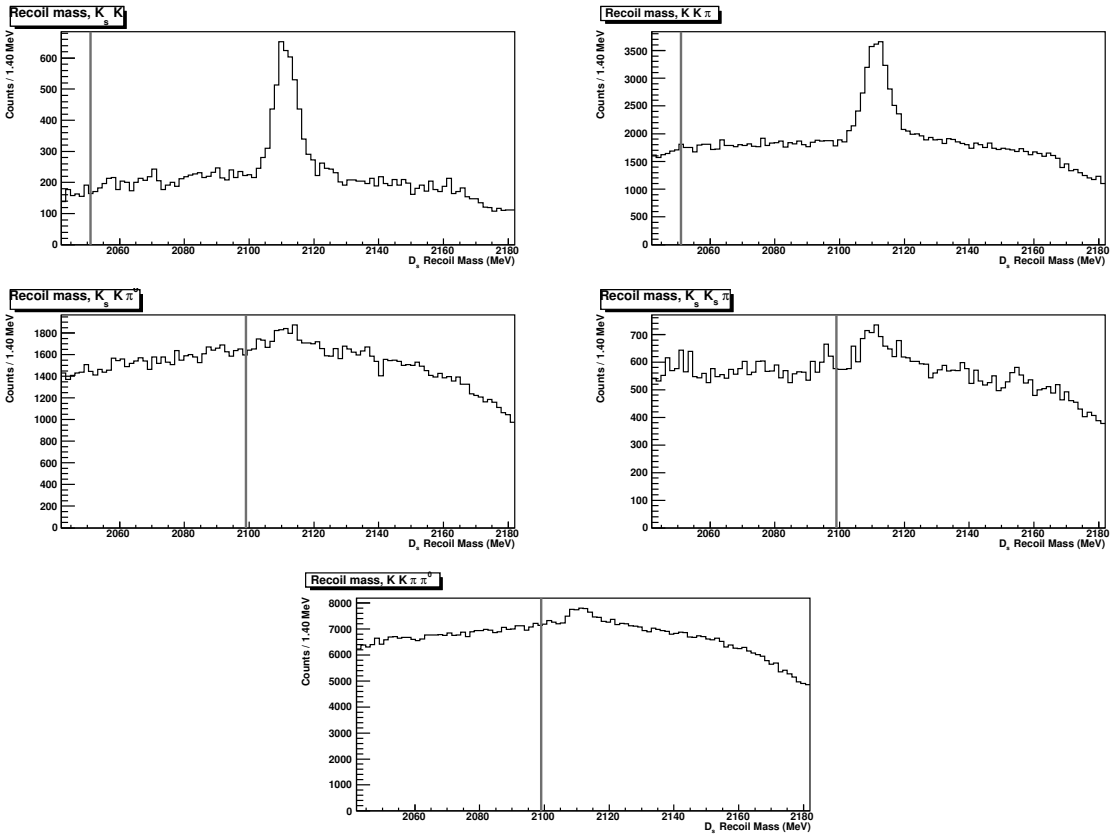

Figure G.2: Recoil mass distribution for the D_s tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$ in the data. We reject D_s with a recoil mass below the cut line.

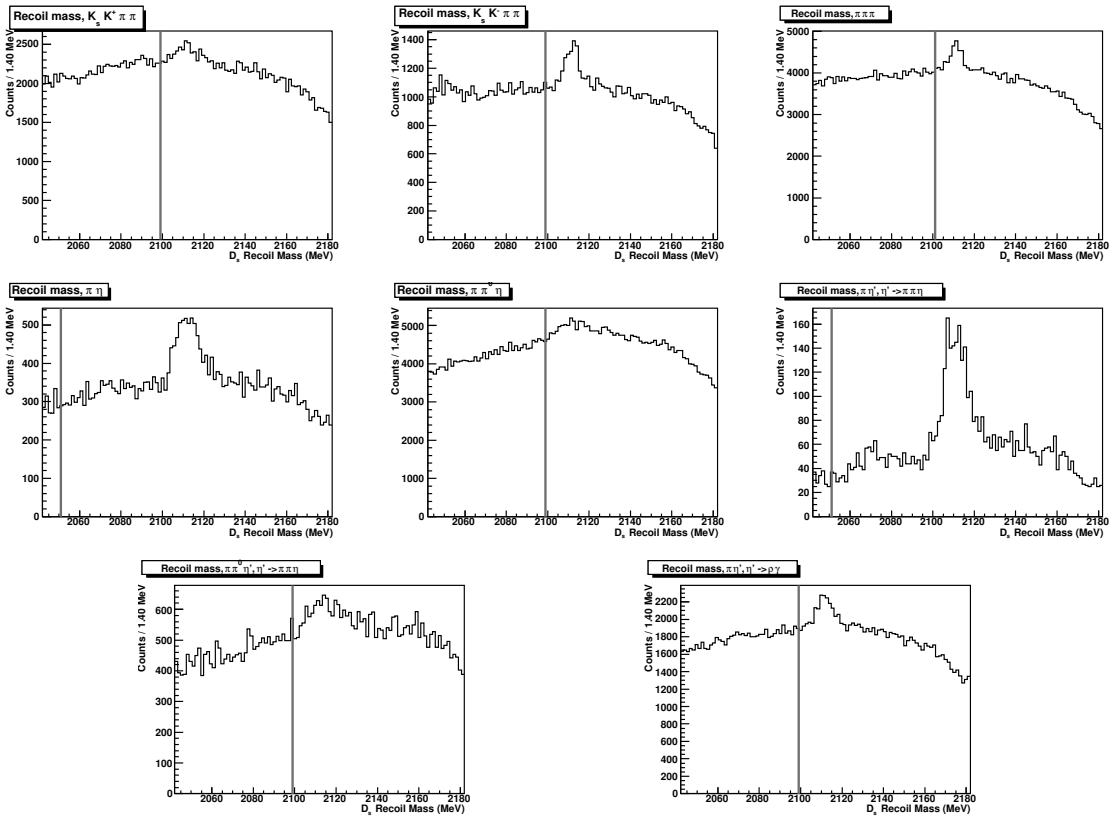

Figure G.3: Recoil mass distribution for the D_s tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$ in the data. We reject D_s with a recoil mass below the cut line.

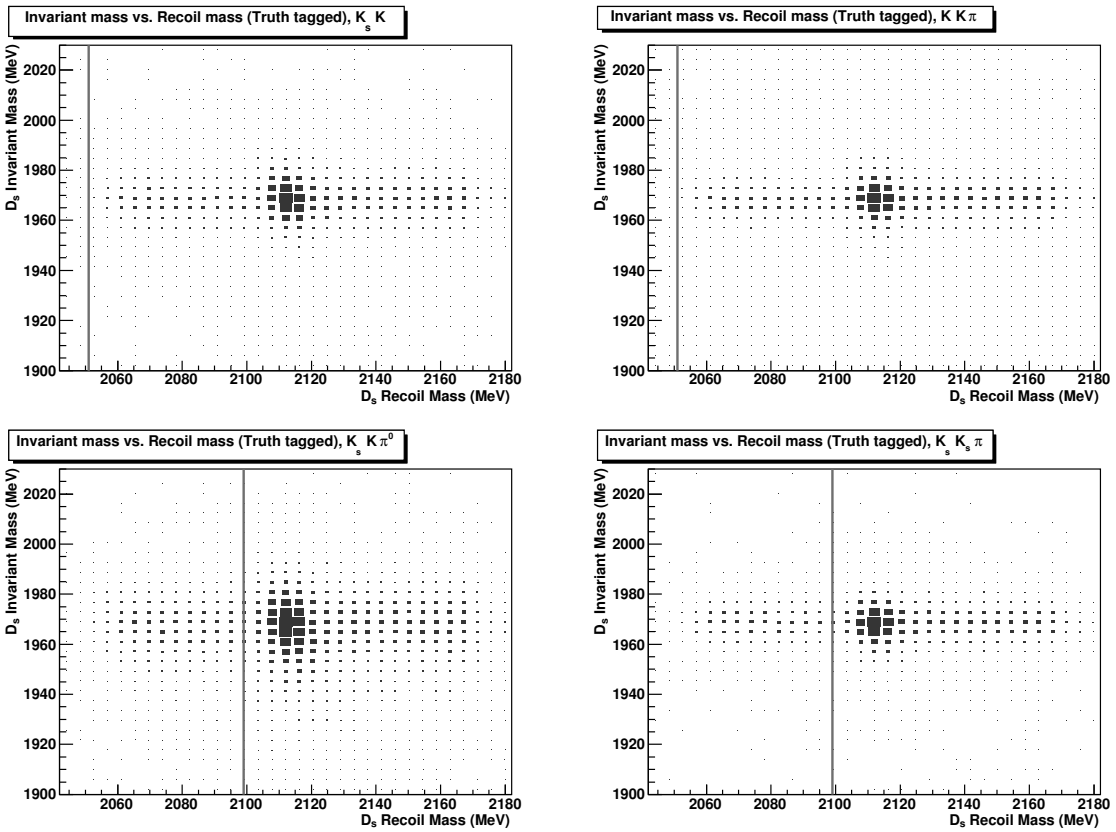

Figure G.4: Monte Carlo truth tagged plots of the invariant mass vs. recoil mass distribution for D_s tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. Note that the recoil mass and the invariant mass are almost entirely uncorrelated.

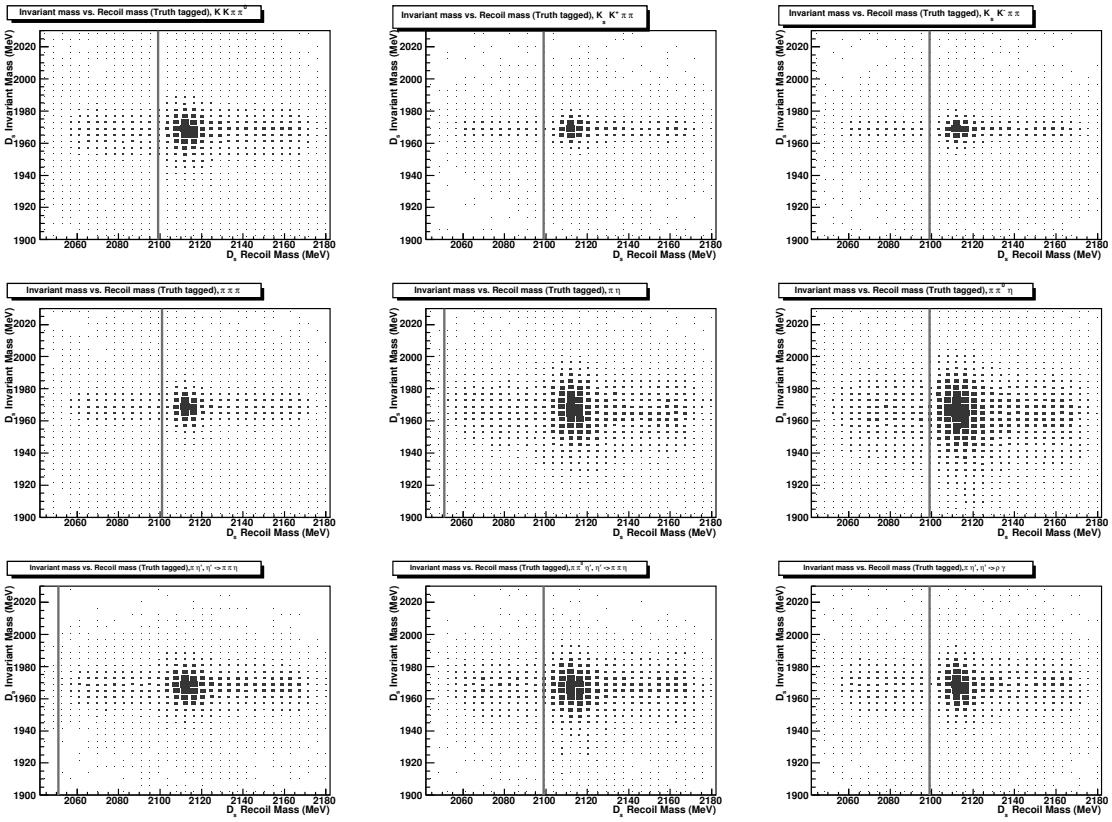

Figure G.5: Monte Carlo truth tagged plots of the invariant mass vs. recoil mass distribution for D_s tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. Again, the recoil mass and the invariant mass show little correlation.

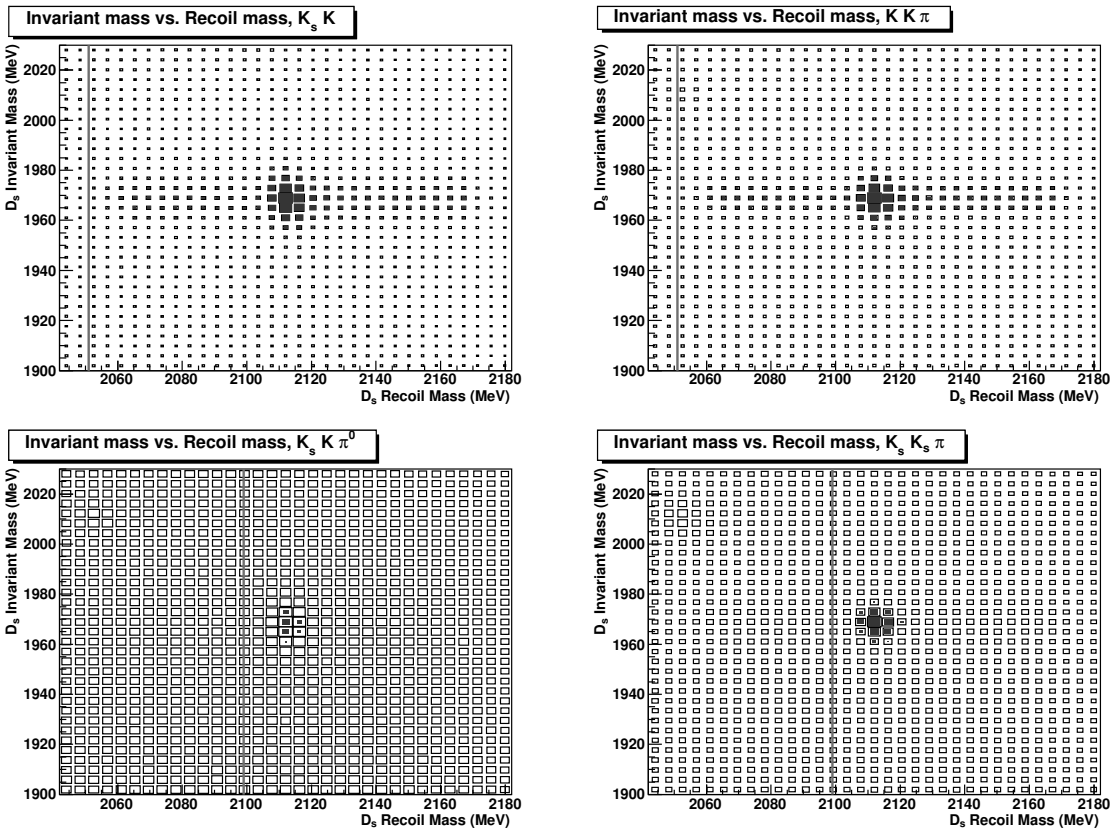

Figure G.6: Monte Carlo plots of the invariant mass vs. recoil mass distribution, including properly weighted charm and continuum background, for D_s tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$.

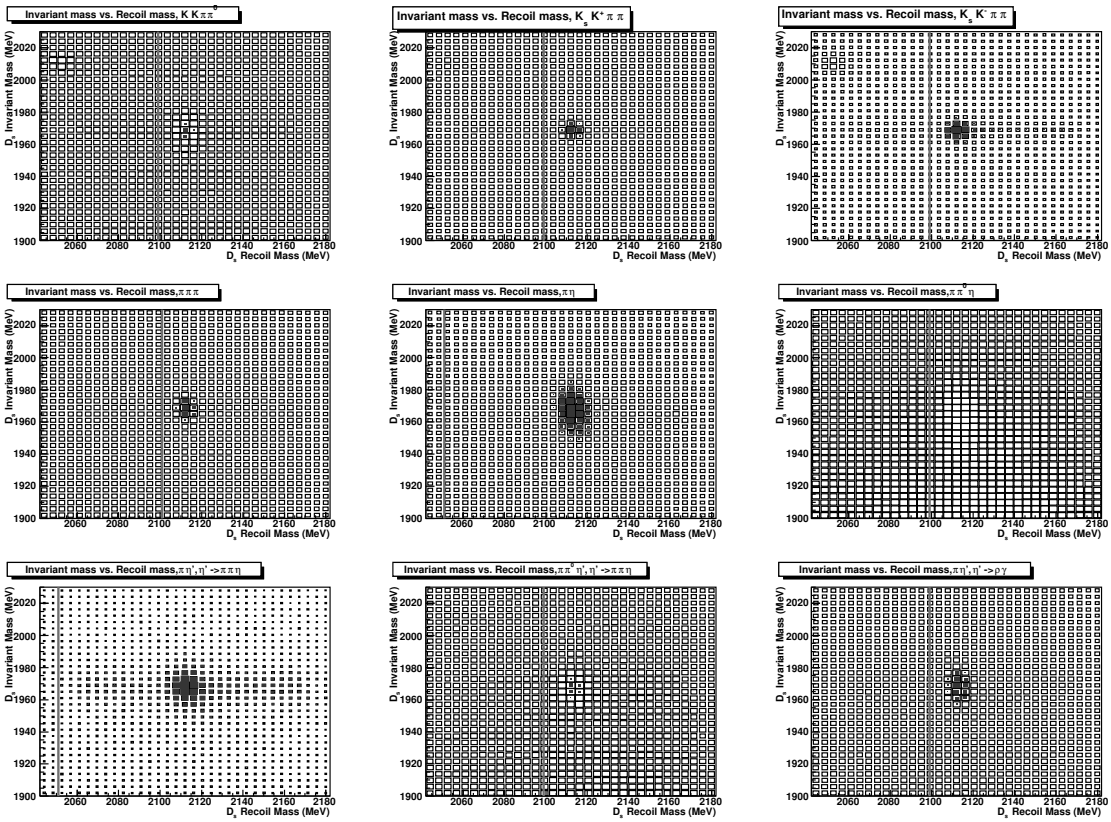

Figure G.7: Monte Carlo plots of the invariant mass vs. recoil mass distribution, including properly weighted charm and continuum background, for D_s tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

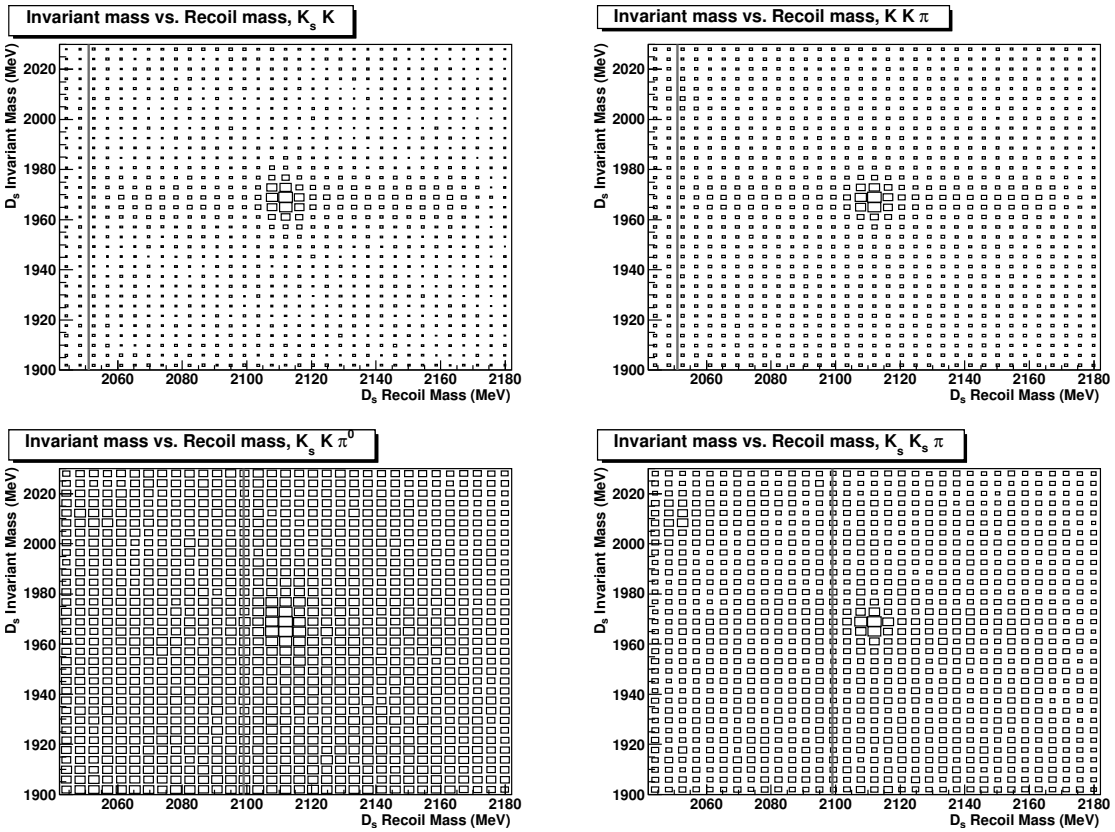

Figure G.8: The invariant mass vs. recoil mass distribution in data for D_s tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. The data distribution doesn't show any unexpected behavior relative to the Monte Carlo expectation.

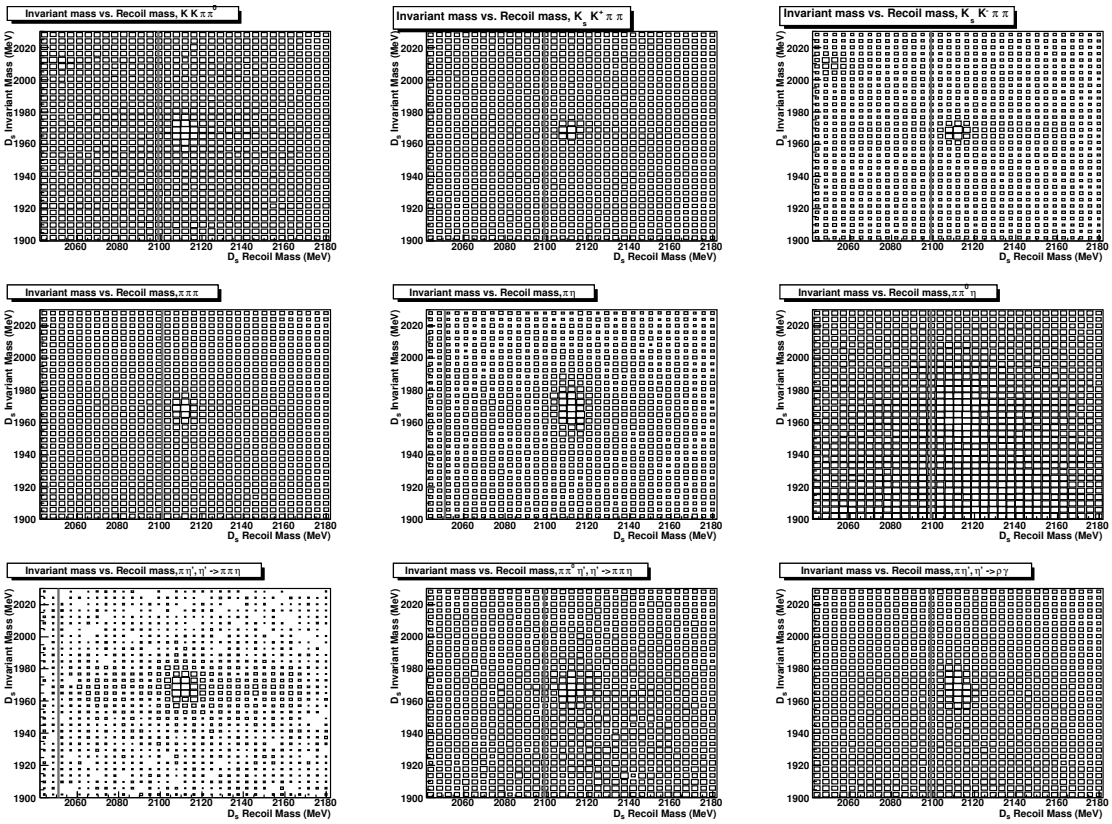

Figure G.9: The invariant mass vs. recoil mass distribution in data for D_s tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

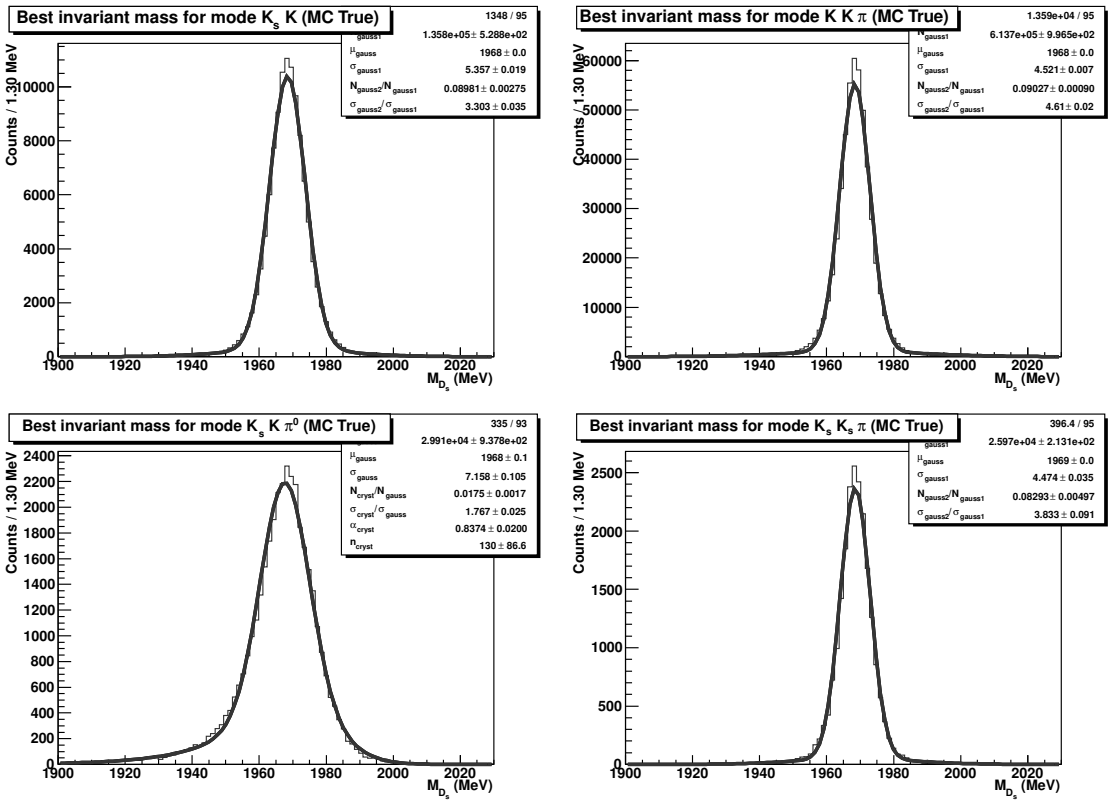

Figure G.10: Fits to the truth-tagged D_s invariant mass from the Monte Carlo. We fix the fit function's shape parameters (relative normalization, relative width, and crystal ball power law tail) from these results. These plots show the fit results for D_s to K_sK, $KK\pi$, $K_sK\pi^0$, and $K_sK_s\pi$.

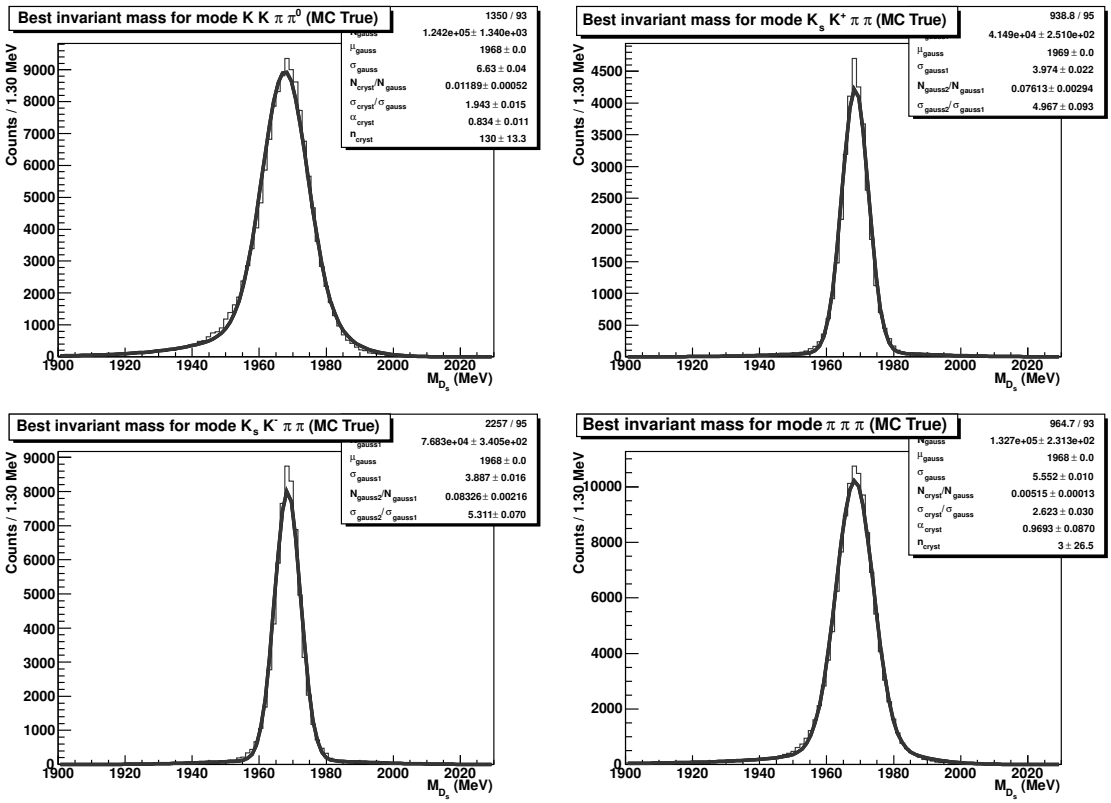

Figure G.11: Fits to the truth-tagged D_s invariant mass from the Monte Carlo. We fix the fit function's shape parameters (relative normalization, relative width, and crystal ball power law tail) from these results. These plots show the fit results for D_s to $KK\pi\pi^0$, $K_sK^+\pi\pi$, $K_sK^-\pi\pi$, and $\pi\pi\pi$.

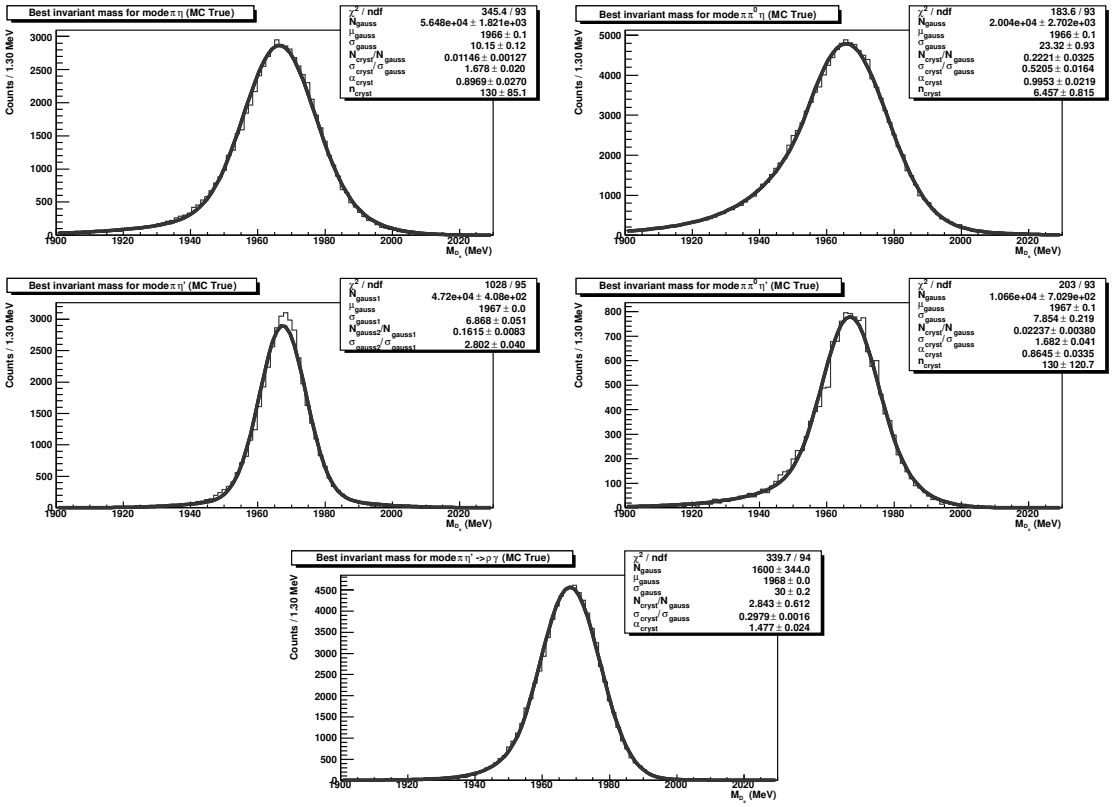

Figure G.12: Fits to the truth-tagged D_s invariant mass from the Monte Carlo. We fix the fit function's shape parameters (relative normalization, relative width, and crystal ball power law tail) from these results. These plots show the fit results for D_s to $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

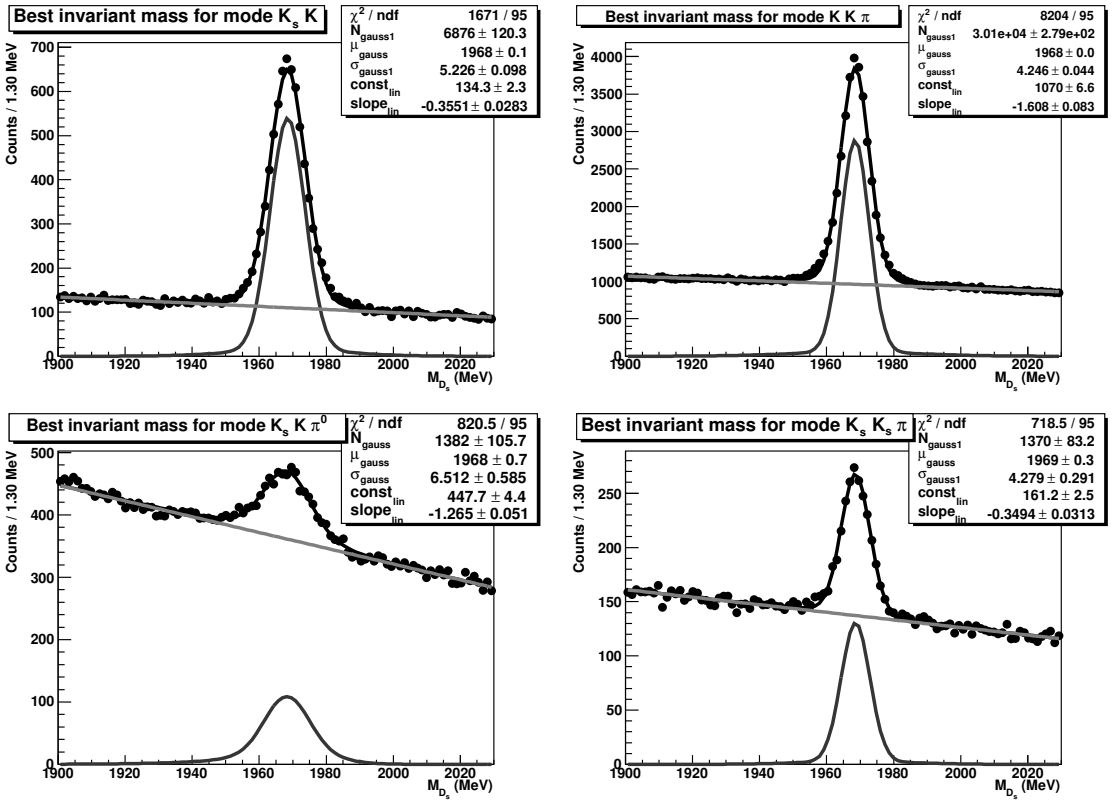

Figure G.13: D_s invariant mass fits in the weighted 20× Monte Carlo sample (charm + continuum), determining the total number of D_s tags for modes K_sK, $KK\pi$, $K_sK\pi^0$, and $K_sK_s\pi$.

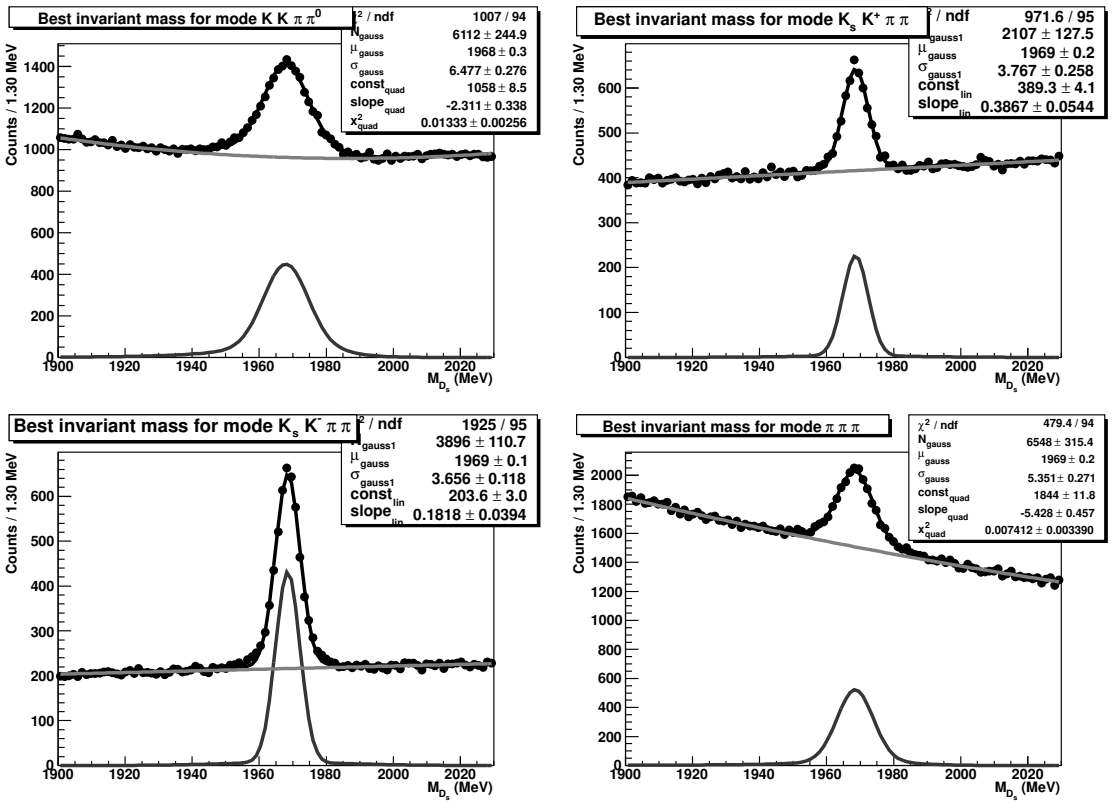

Figure G.14: D_s invariant mass fits in the weighted 20× Monte Carlo sample (charm + continuum), determining the total number of D_s tags for modes $KK\pi\pi^0$, $K_sK^+\pi\pi$, $K_sK^-\pi\pi$, and $\pi\pi\pi$.

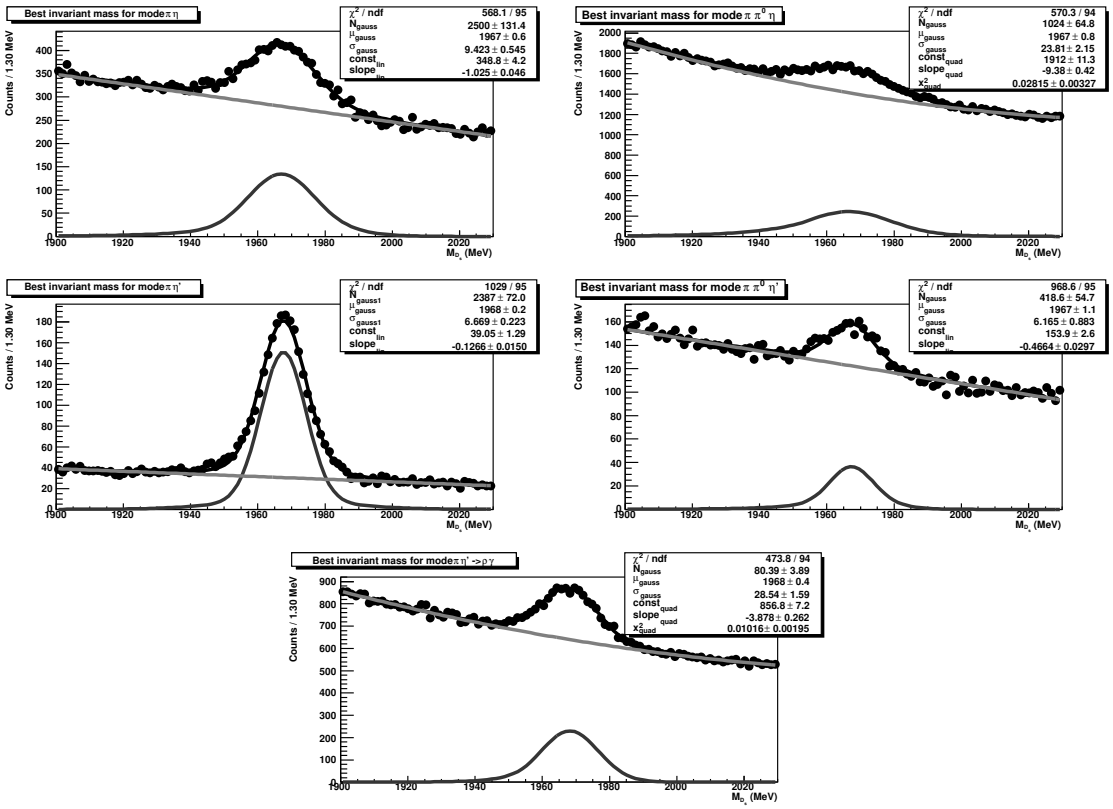

Figure G.15: D_s invariant mass fits in the weighted 20× Monte Carlo sample (charm + continuum), determining the total number of D_s tags for modes $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

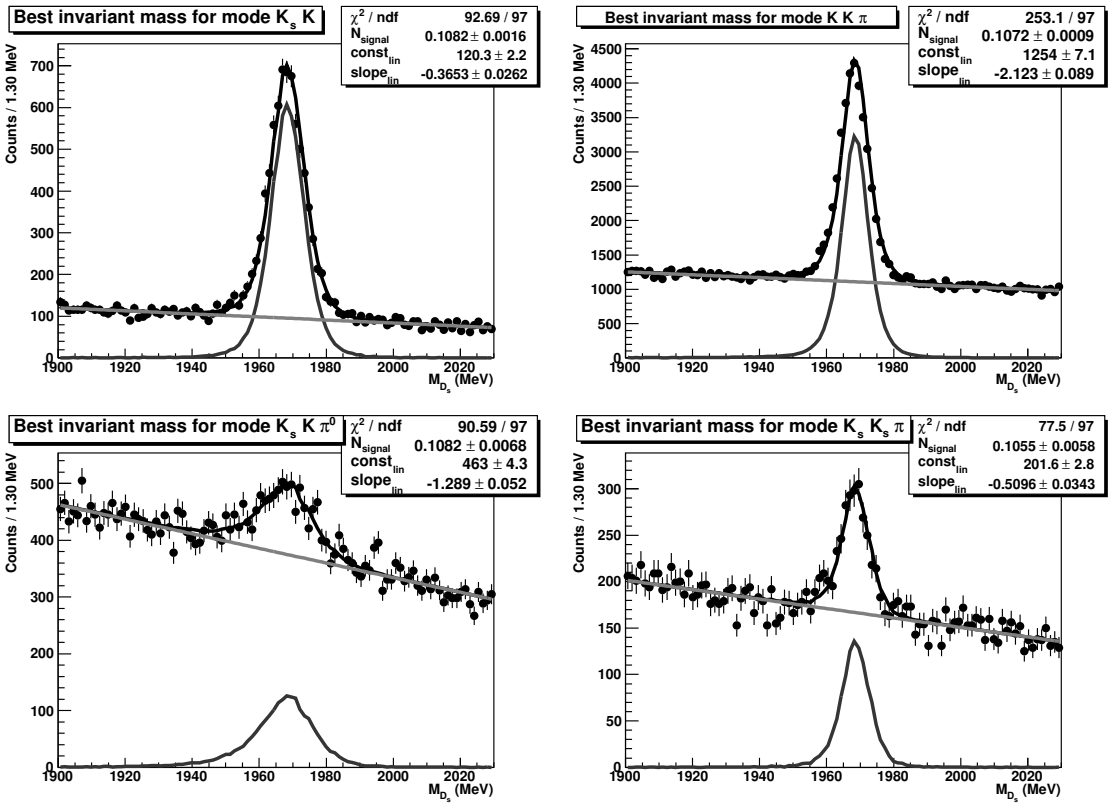

Figure G.16: D_s invariant mass fits in the data sample using a signal histogram from the truth-tagged Monte Carlo. These plots show our results for D_s to K_sK, $KK\pi$, $K_sK\pi^0$, and $K_sK_s\pi$.

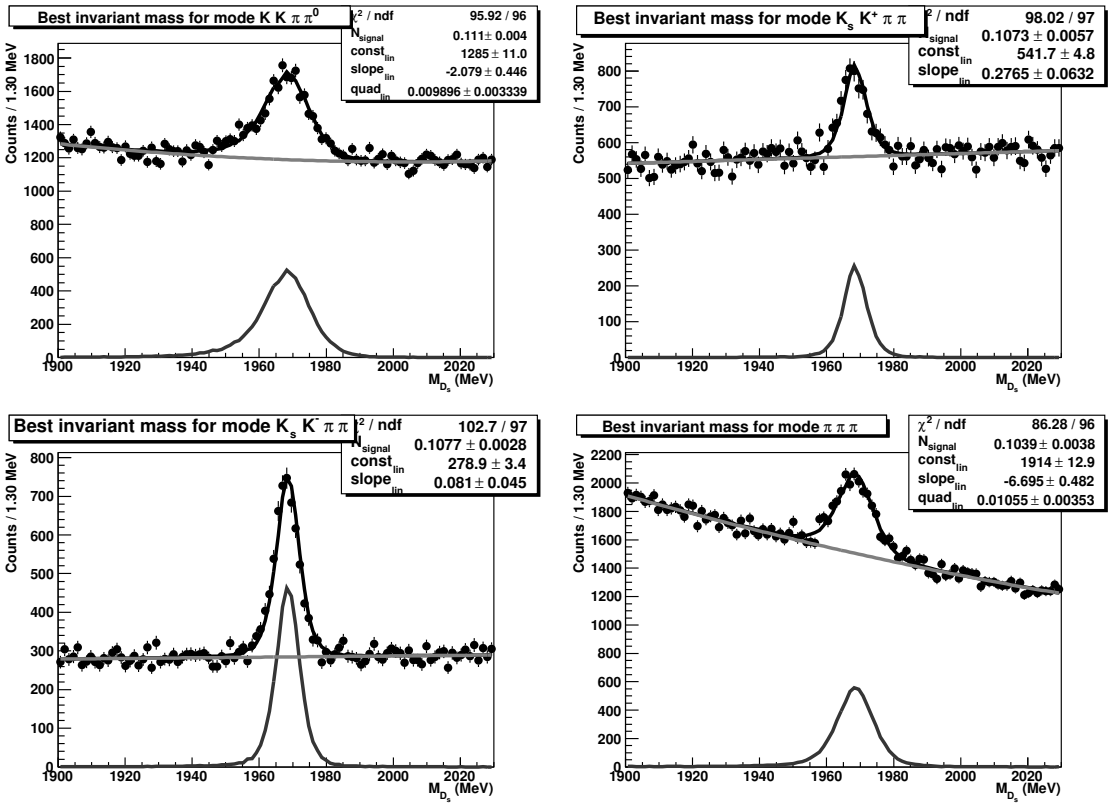

Figure G.17: D_s invariant mass fits in the data sample using a signal histogram from the truth-tagged Monte Carlo. These plots show our results for D_s to $KK\pi\pi^0$, $K_sK^+\pi\pi$, $K_sK^-\pi\pi$, and $\pi\pi\pi$.

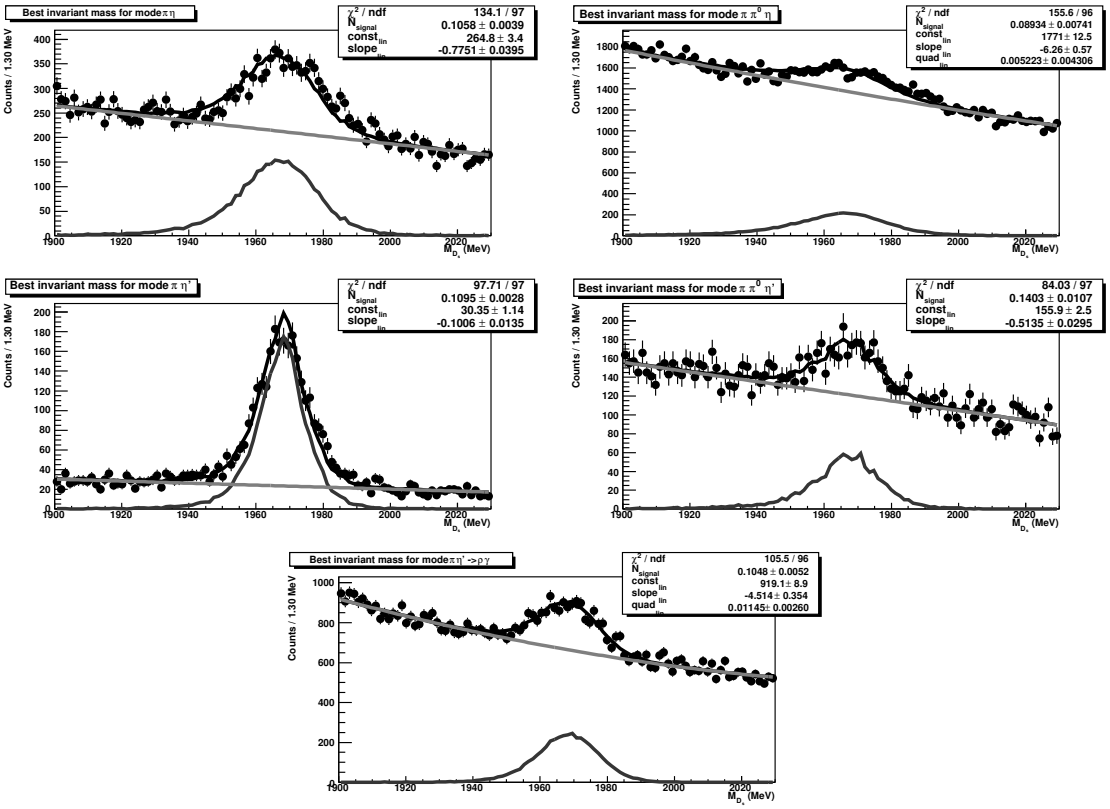

Figure G.18: D_s invariant mass fits in the data sample using a signal histogram from the truth-tagged Monte Carlo. These plots show our results for D_s to $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

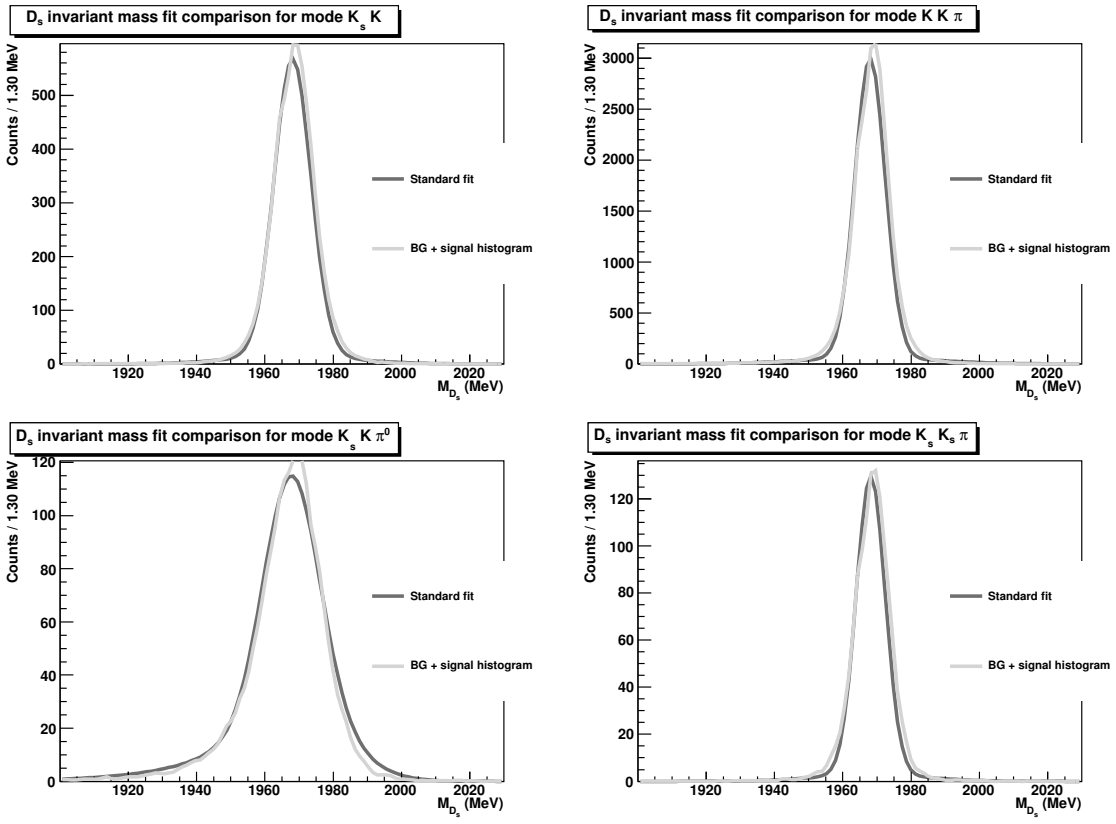

Figure G.19: D_s invariant mass fits for a double gaussian/gaussian+crystal ball signal shape compared to fits with a signal histogram for modes K_sK, $KK\pi$, $K_sK\pi^0$, and $K_sK_s\pi$.

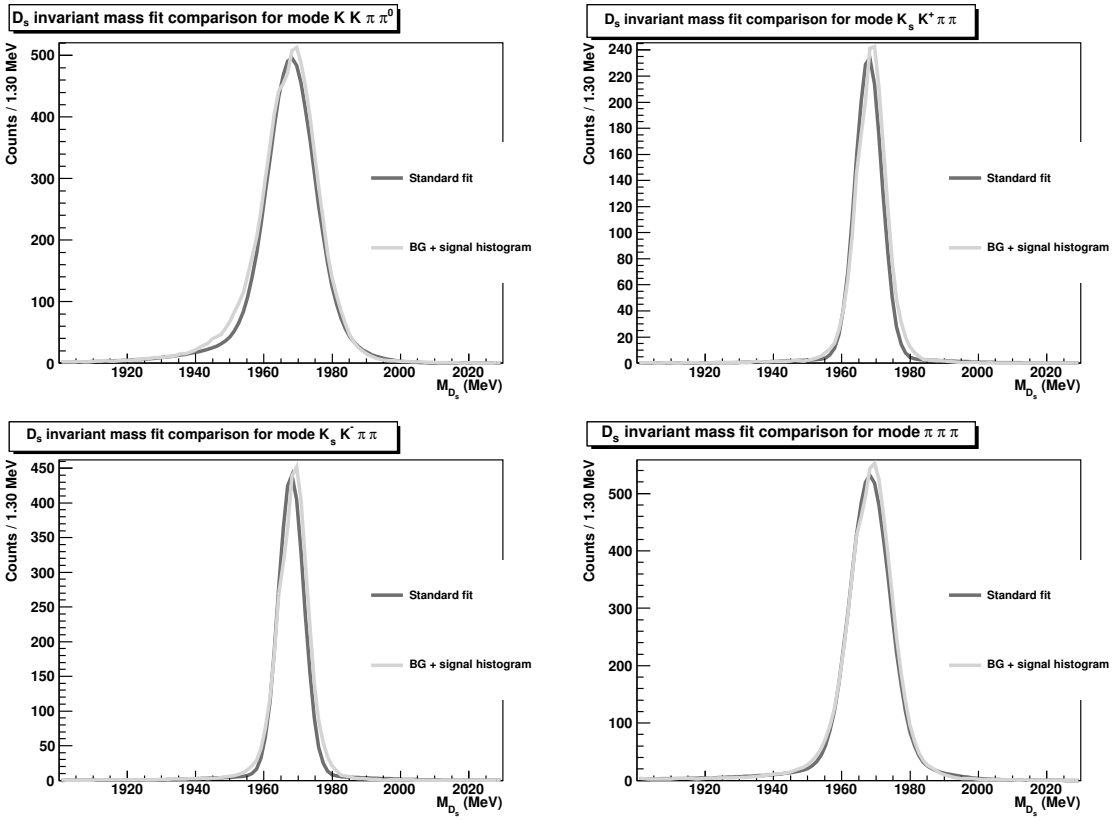

Figure G.20: D_s invariant mass fits for a double gaussian/gaussian+crystal ball signal shape compared to fits with a signal histogram for modes $KK\pi\pi^0$, $K_sK^+\pi\pi$, $K_sK^-\pi\pi$, and $\pi\pi\pi$.

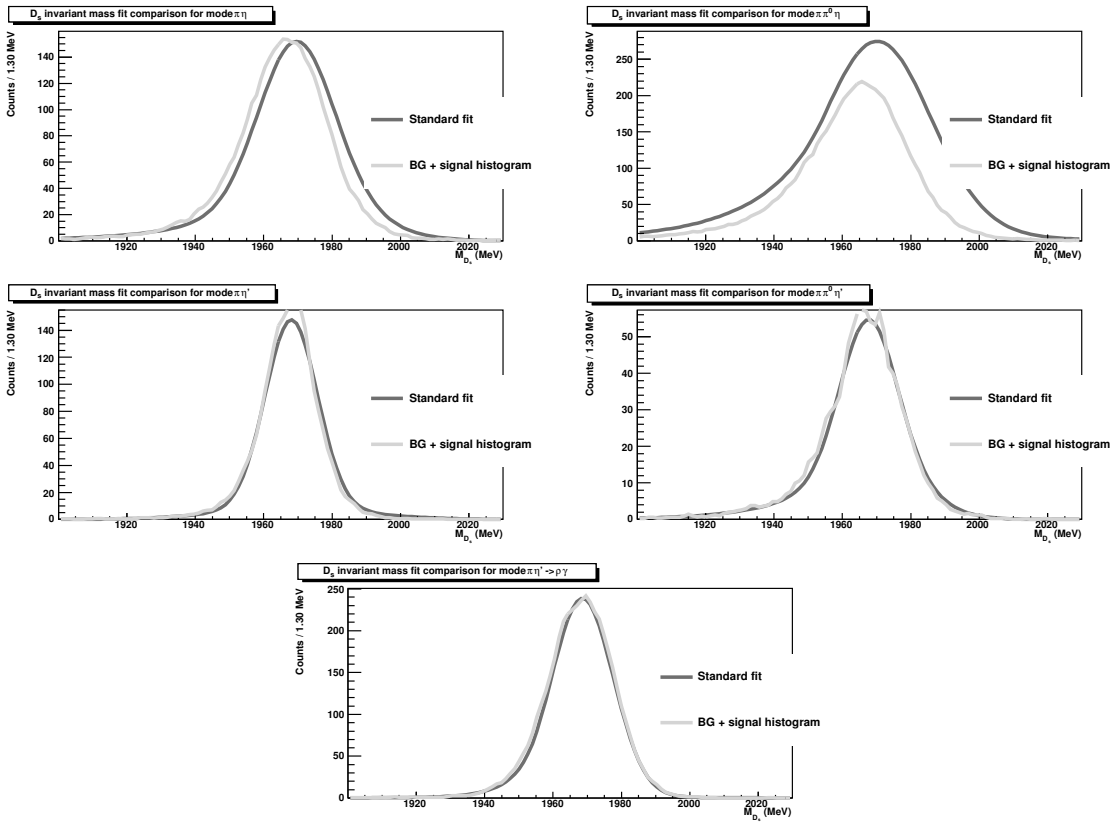

Figure G.21: D_s invariant mass fits for a double gaussian/gaussian+crystal ball signal shape compared to fits with a signal histogram for modes $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$.

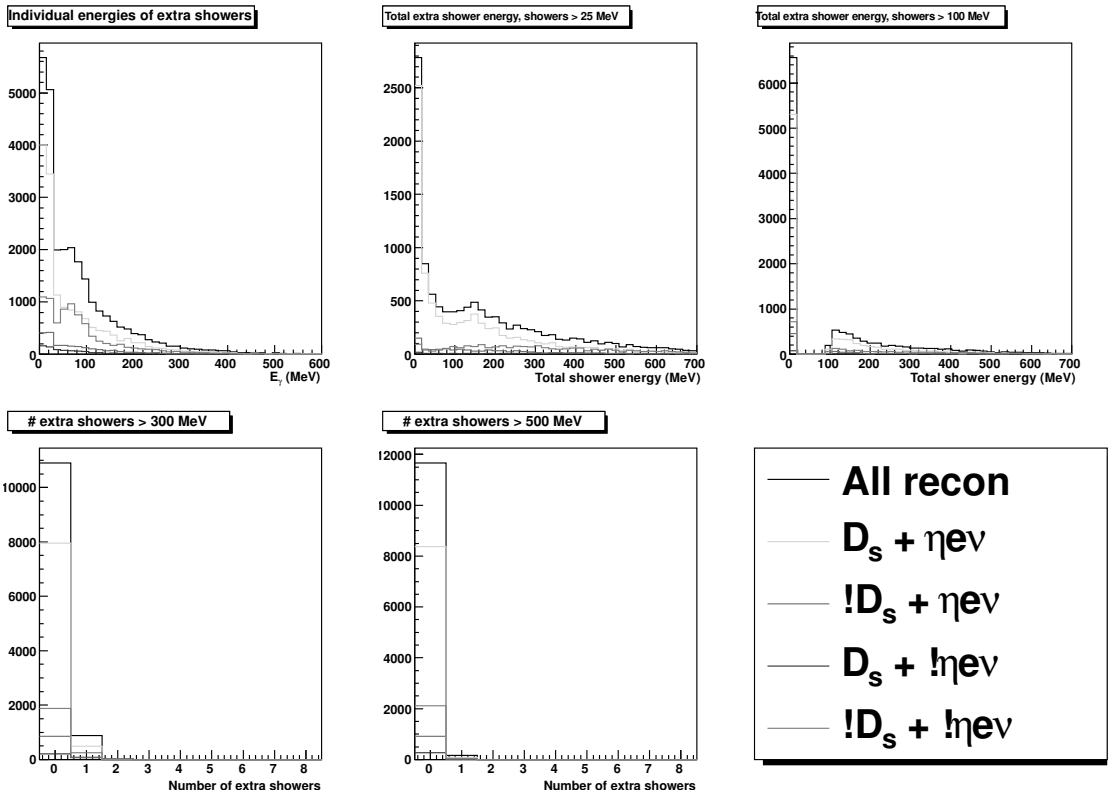

Figure G.22: Extra showers after finding the tagged D_s, the η, and the electron in $\eta e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays.

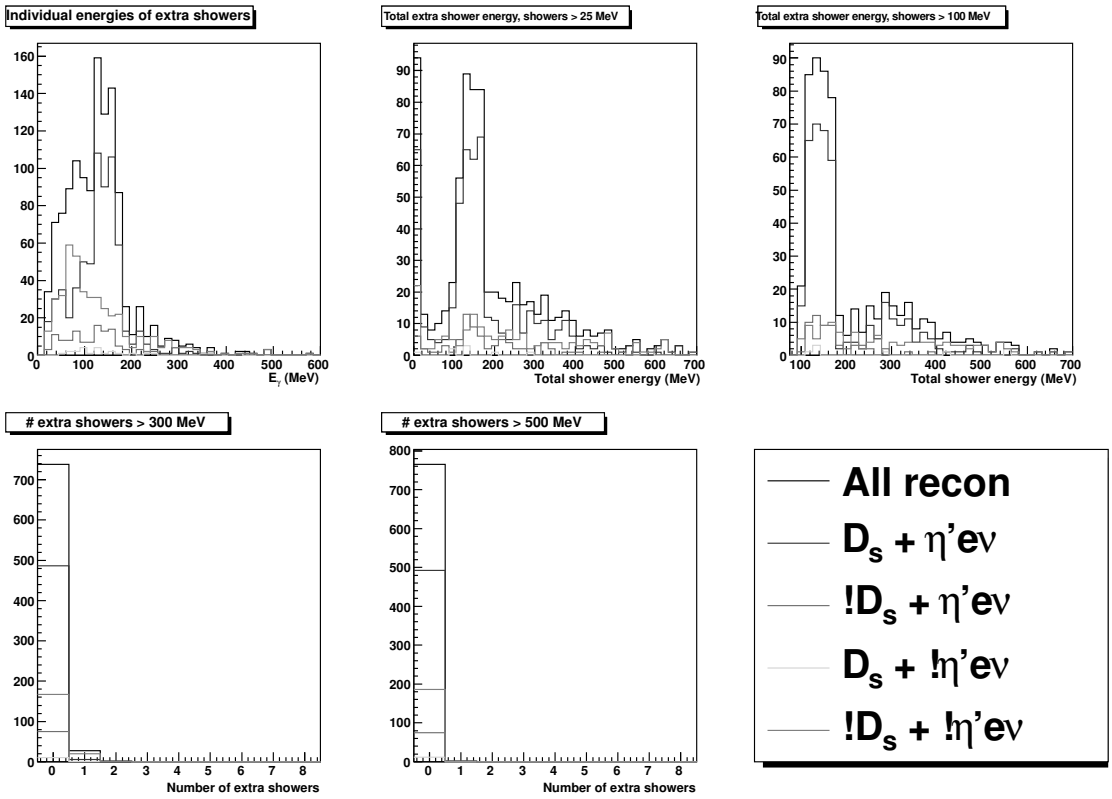

Figure G.23: Extra showers after finding the tagged D_s, the η', and the electron in $\eta' e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays.

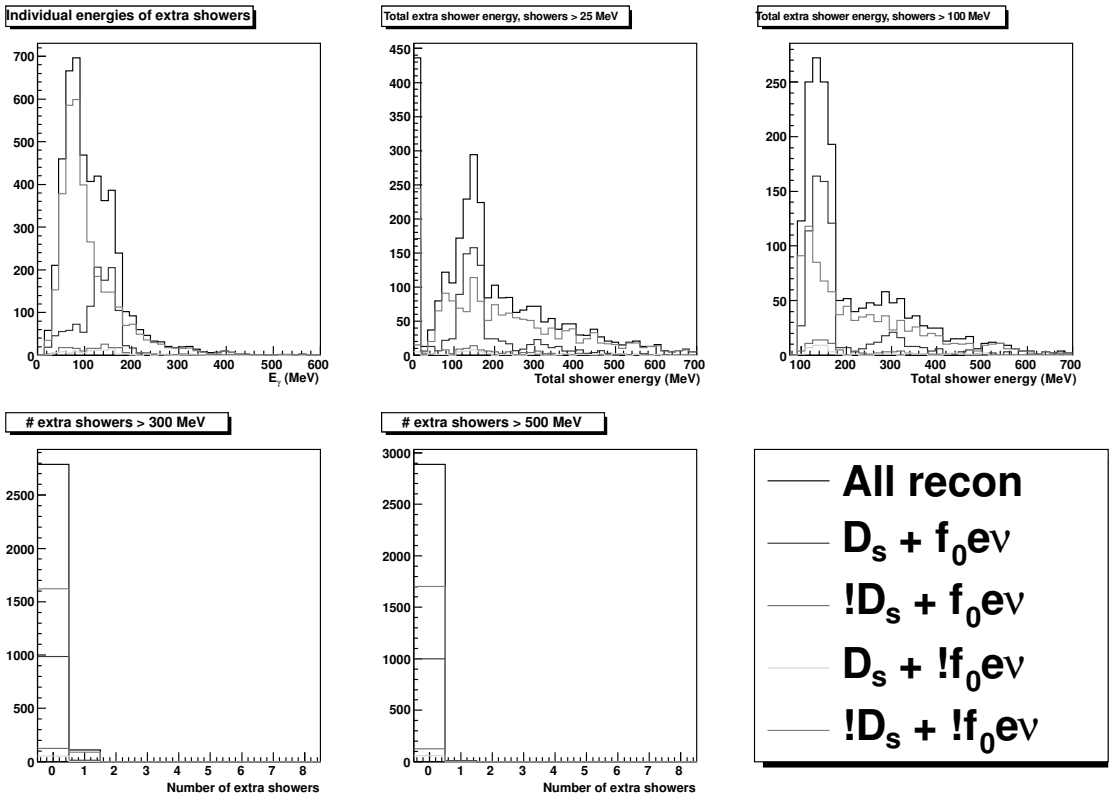

Figure G.24: Extra showers after finding the tagged D_s, the f_0, and the electron in $f_0 e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays.

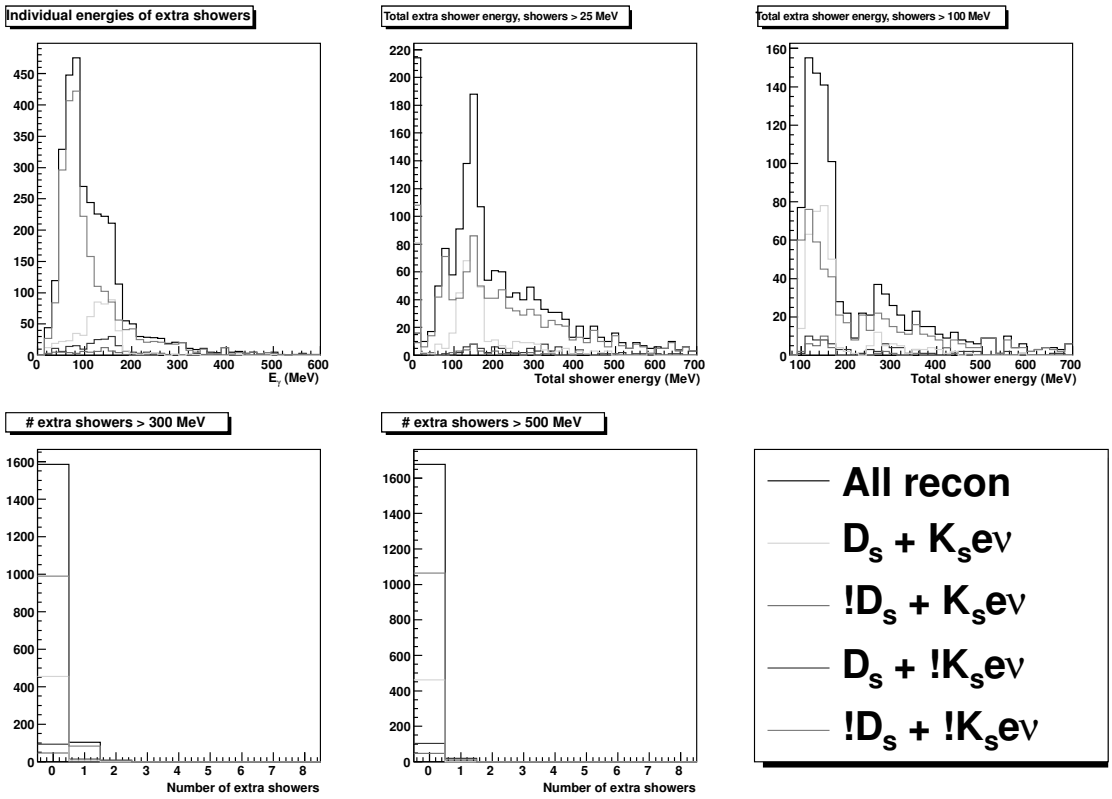

Figure G.25: Extra showers after finding the tagged D_s, the K_s, and the electron in $K_s e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays.

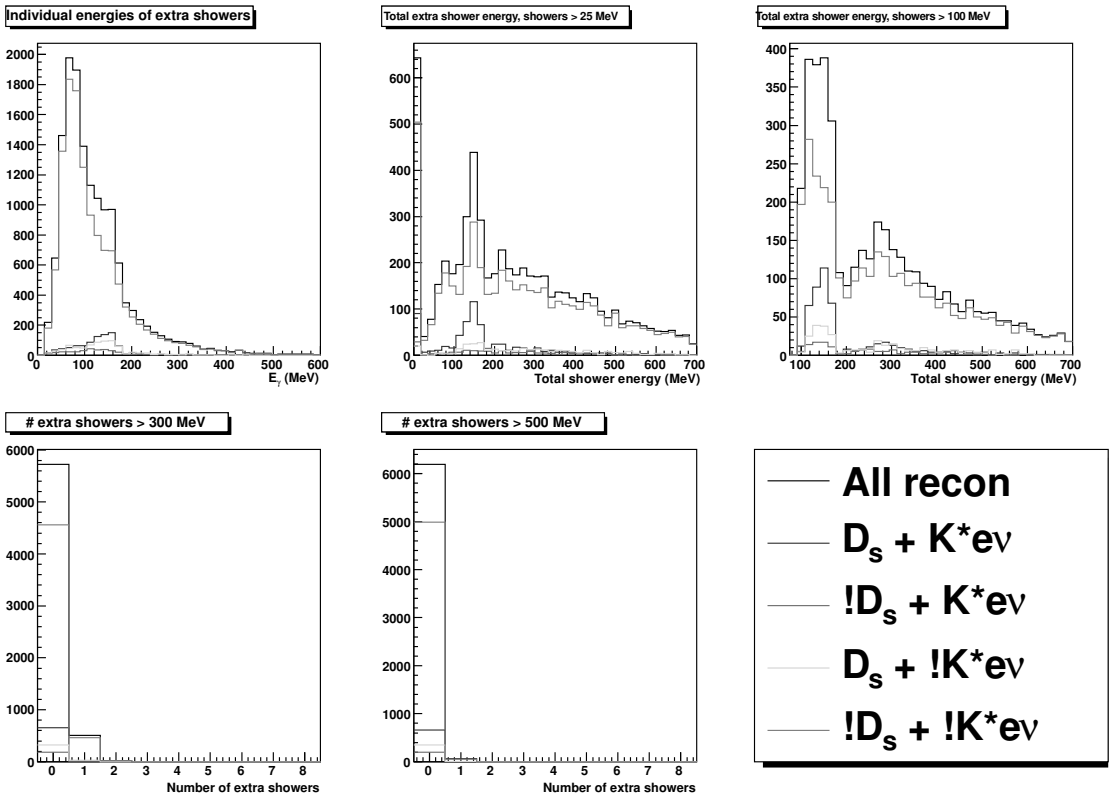

Figure G.26: Extra showers after finding the tagged D_s, the K^*, and the electron in $K^*e\nu$ (20× MC sample). Our shower quality selections include both $\frac{E_9}{E_{25}}$ O.K. and a splitoff rejection. The peak near 140 MeV is due to the γ from D_s^* decays.

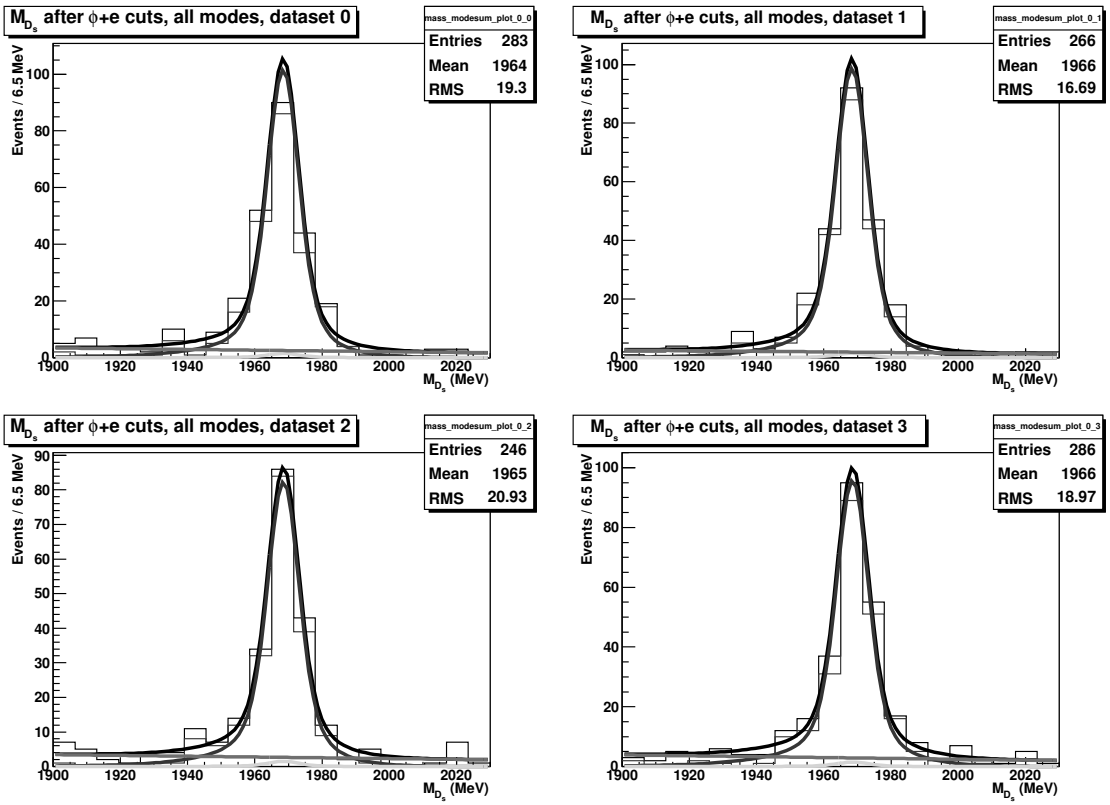

Figure G.27: Likelihood fit results for the D_s mass spectrum after all $D_s \to \phi e\nu$ semi-leptonic cuts in the first four data-sized Monte Carlo samples. The histograms show total events and Monte Carlo truth-tagged events, while the peaking fit line gives the signal part of our fit. The two solid background fit lines represent the non-peaking background fit and our peaking background subtraction.

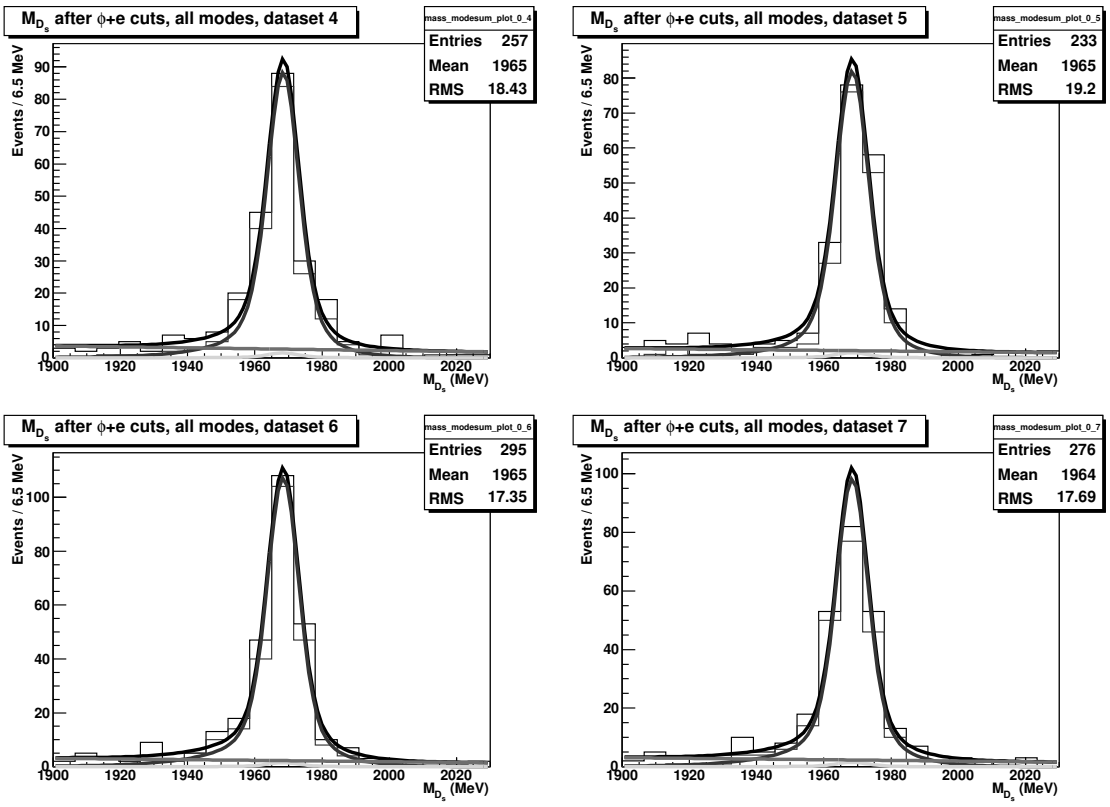

Figure G.28: $D_s \to \phi e \nu$ data-sized Monte Carlo results, second group of datasets.

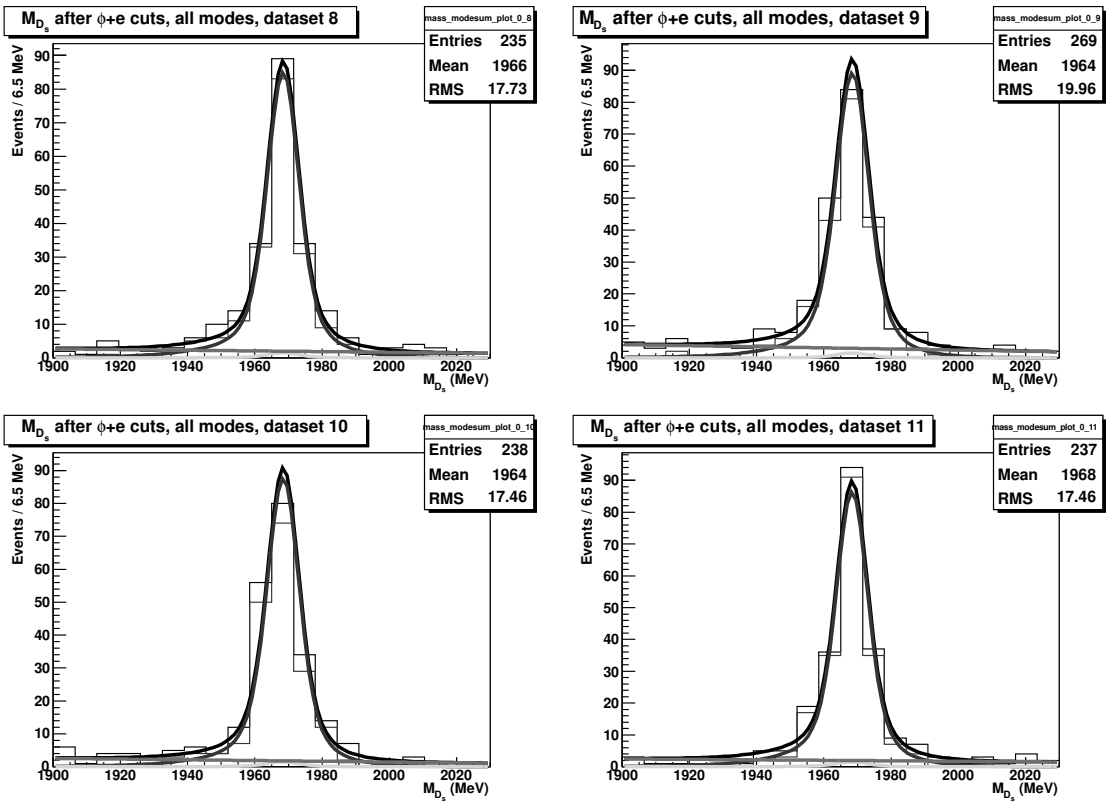

Figure G.29: $D_s \to \phi e \nu$ data-sized Monte Carlo results, third group of datasets.

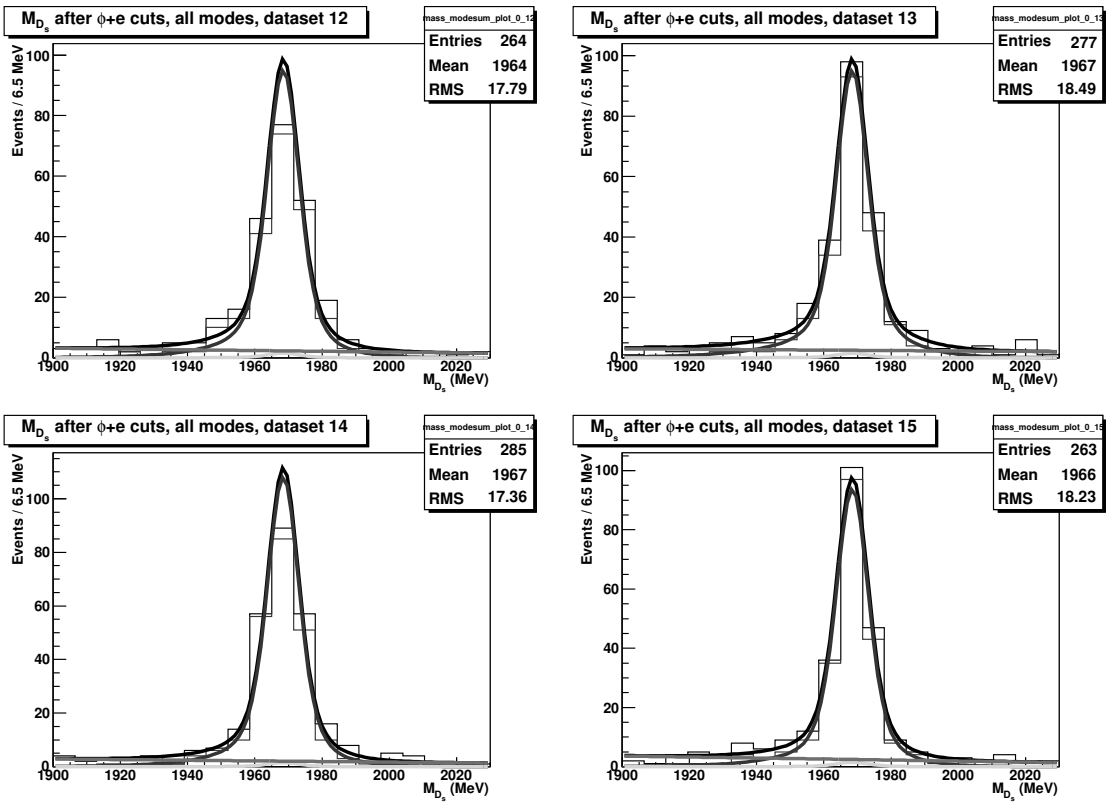

Figure G.30: $D_s \to \phi e \nu$ data-sized Monte Carlo results, fourth group of datasets.

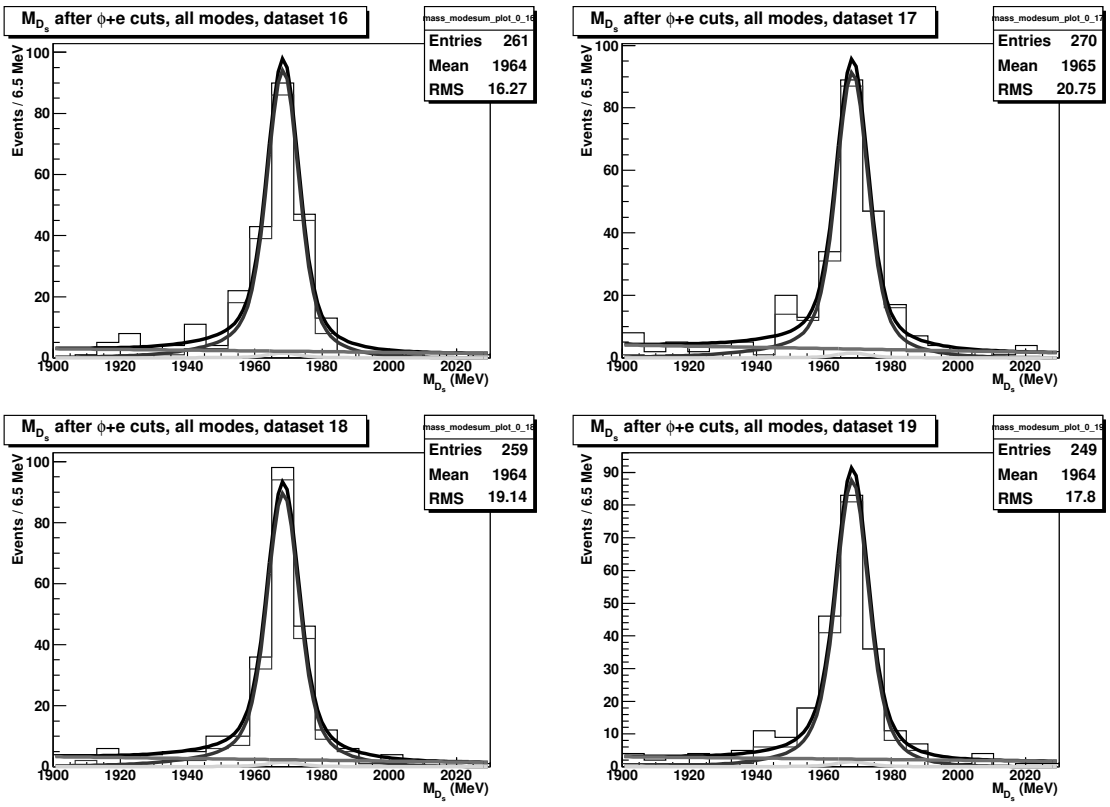

Figure G.31: $D_s \to \phi e \nu$ data-sized Monte Carlo results, fifth group of datasets.

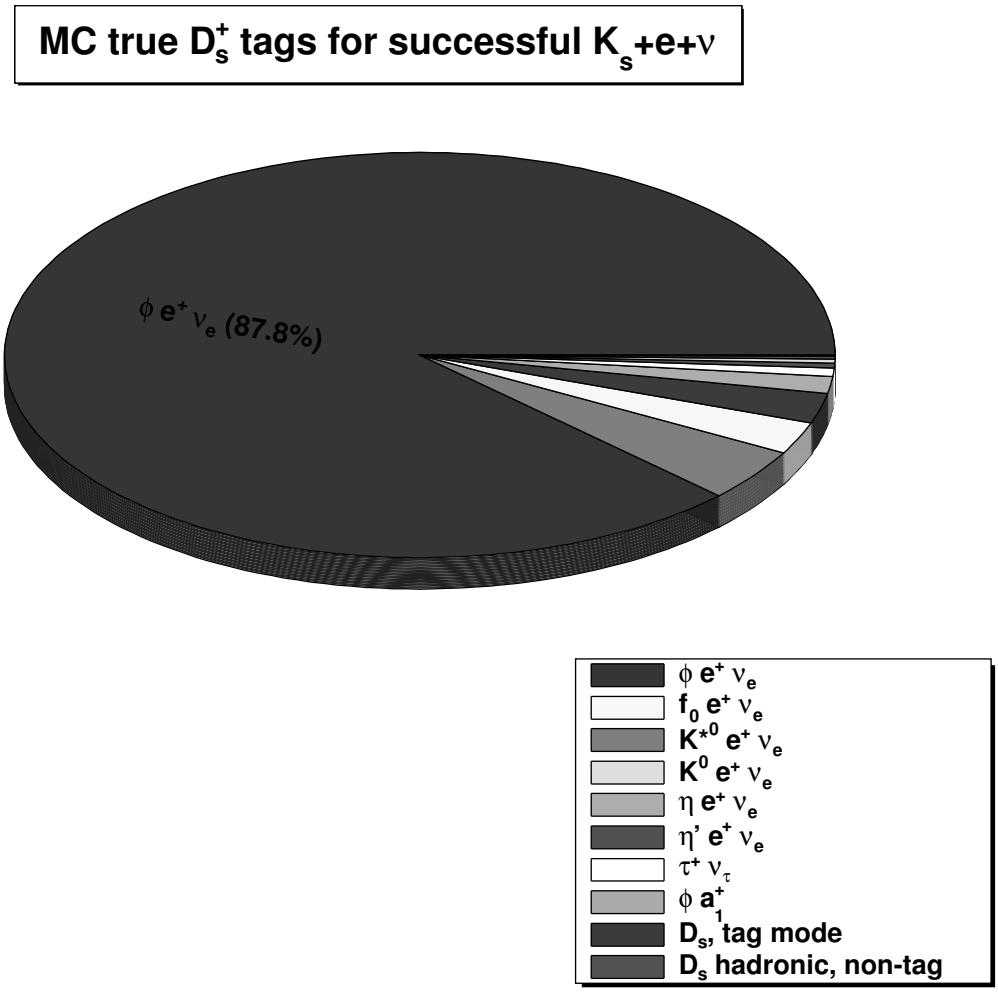

Figure G.32: $D_s \to K_s e\nu$ backgrounds with a true D_s tag (peaking background), from the 20× Monte Carlo. These remain after $K_s e\nu$ semileptonic cuts but before any missing mass cut or other, additional background restrictions. $\phi e\nu$ with $\phi \to K_L K_s$ dominates.

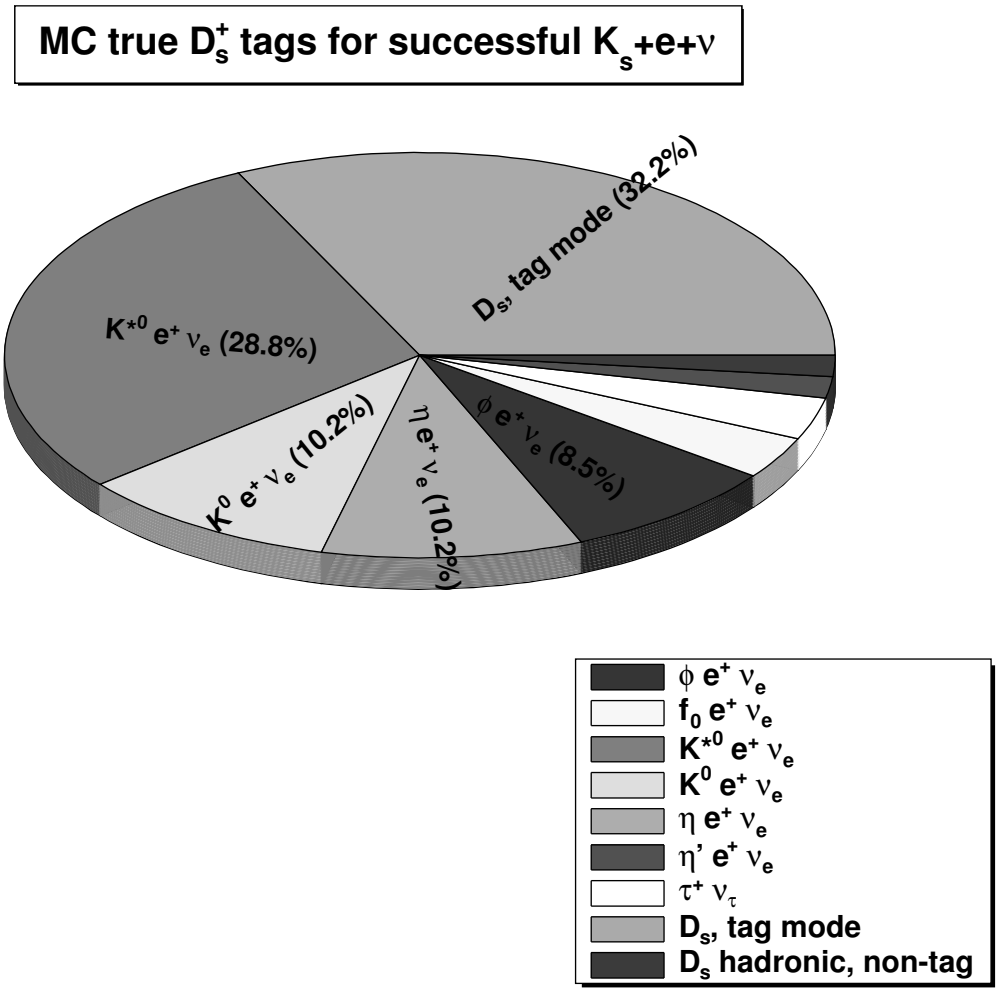

Figure G.33: $D_s \to K_s e\nu$ backgrounds with a true D_s tag (peaking background), after all cuts. The other semileptonic modes each give some fake events, while the dominant non-semileptonic contribution comes from D_s tag modes with a kaon faking the electron (e.g. $D_s \to KK_s$)

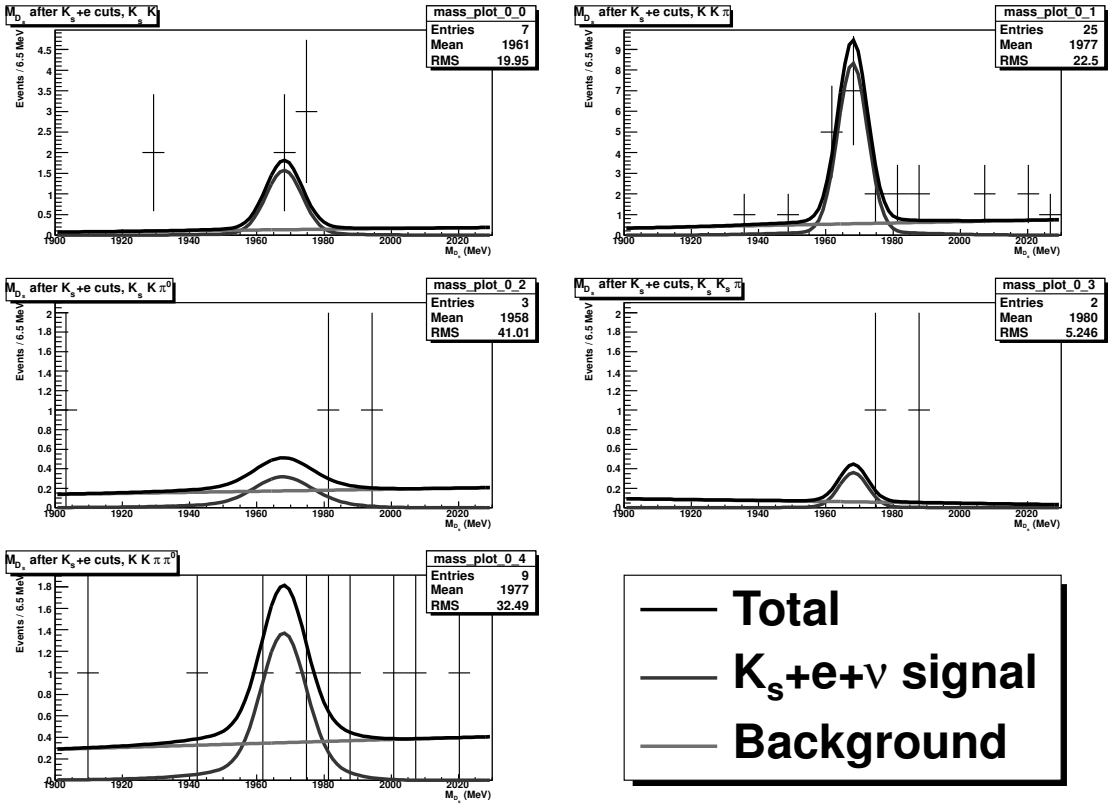

Figure G.34: $D_s \to K_s e\nu$ fit results in the data, after all semileptonic cuts, for tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

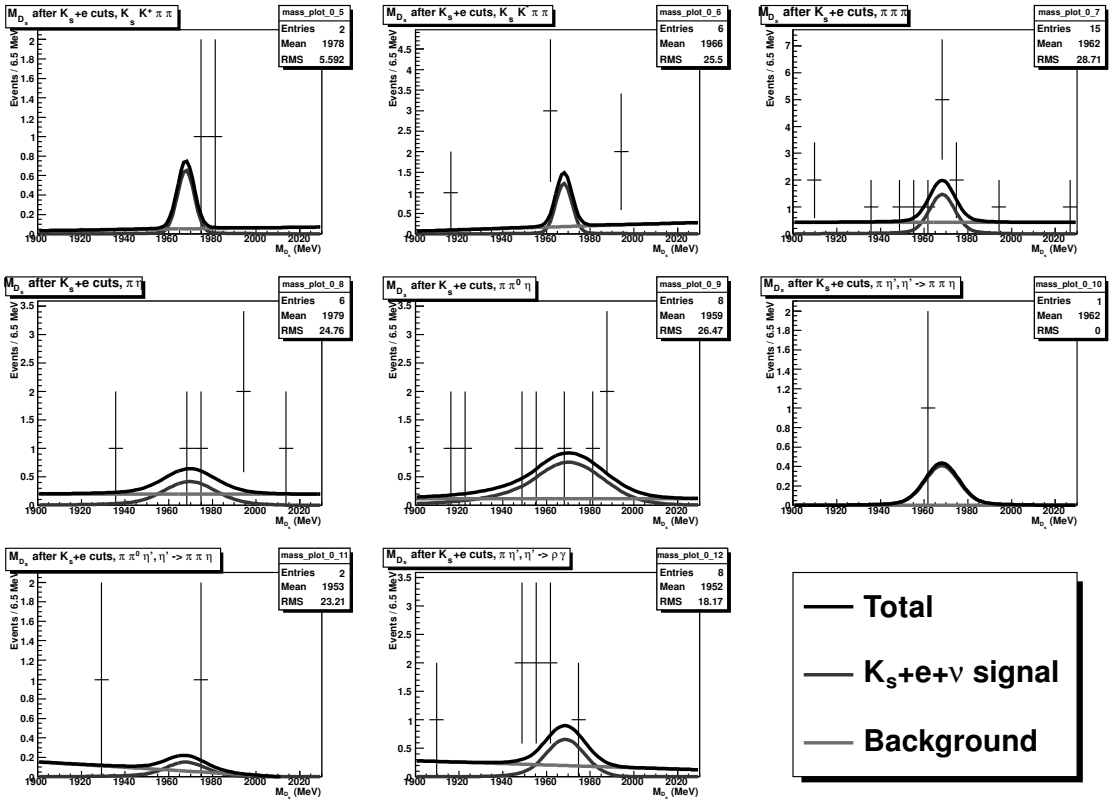

Figure G.35: $D_s \to K_s e\nu$ fit results in the data, after all semileptonic cuts, for tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

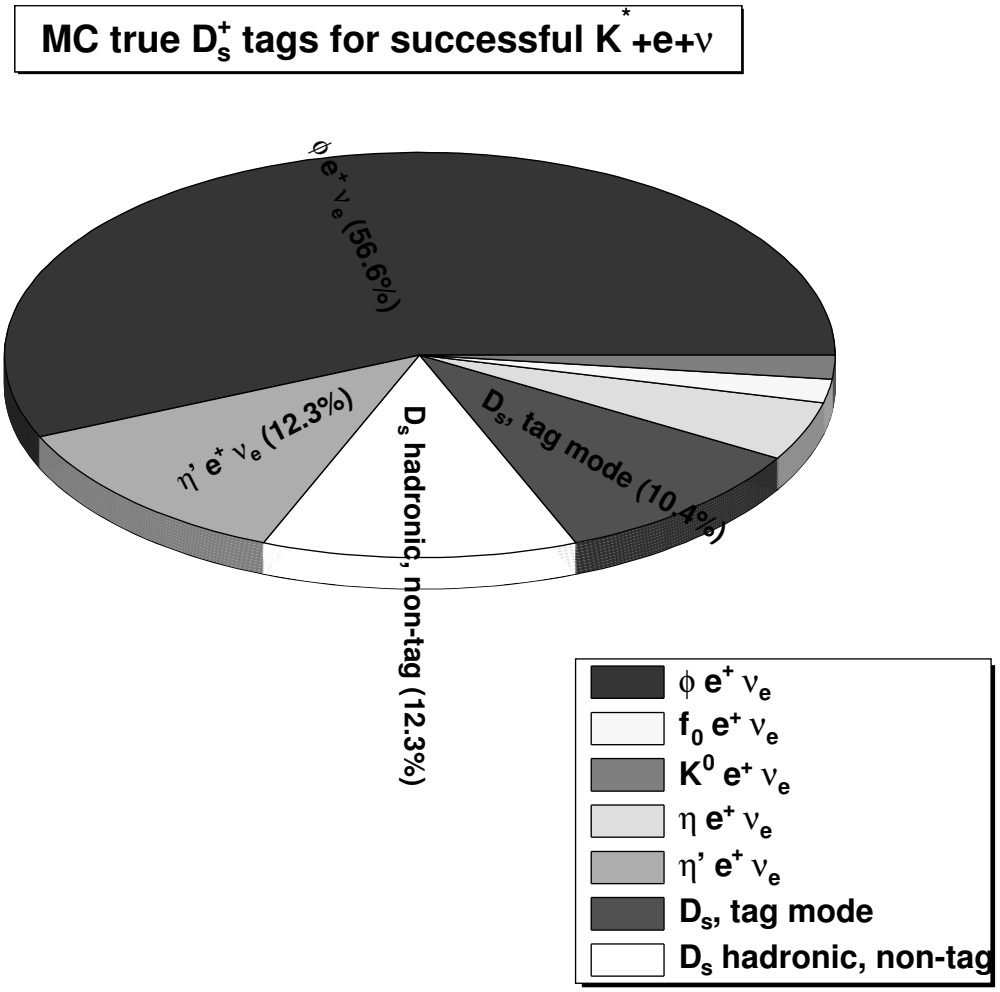

Figure G.36: $D_s \to K^* e\nu$ backgrounds with a true D_s tag (peaking background), before our specific $K^* e\nu$ cuts in the 20× Monte Carlo. Our best improvement in peaking background will come from reducing $D_s \to \phi e\nu$ where one kaon fakes a pion.

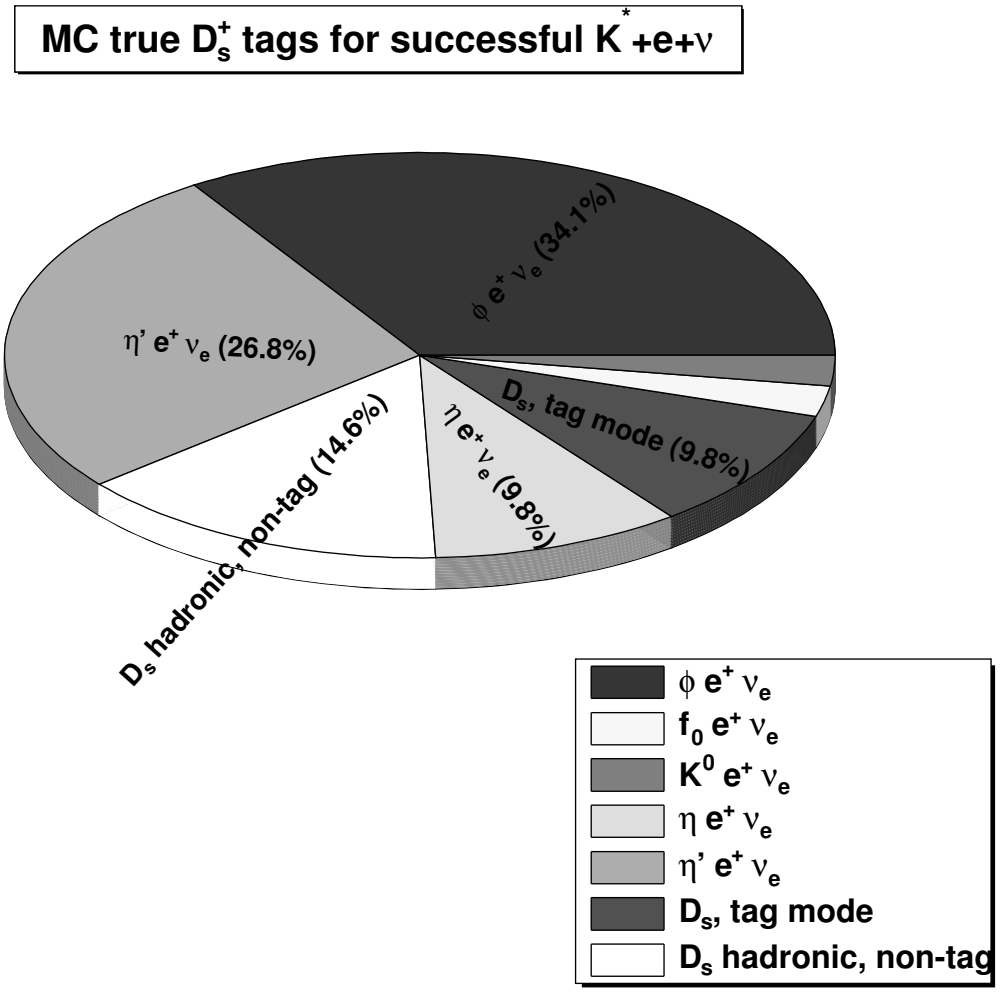

Figure G.37: $D_s \to K^* e\nu$ backgrounds with a true D_s tag (peaking background), after all cuts. The other semileptonic modes each give some fake events, while the dominant non-semileptonic contribution comes from D_s tag modes where a kaon fakes the electron (e.g. $D_s \to KK\pi$).

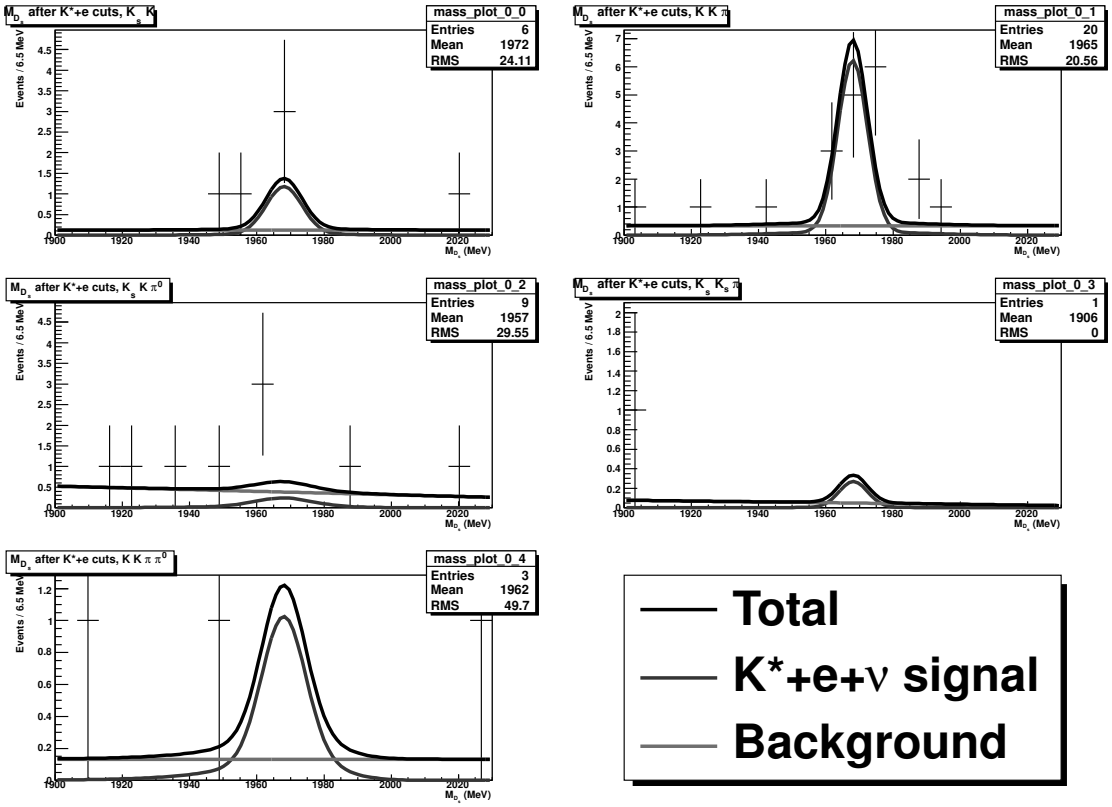

Figure G.38: $D_s \to K^* e\nu$ fit results in the data, after all semileptonic cuts, for tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode receives an independent background normalization.

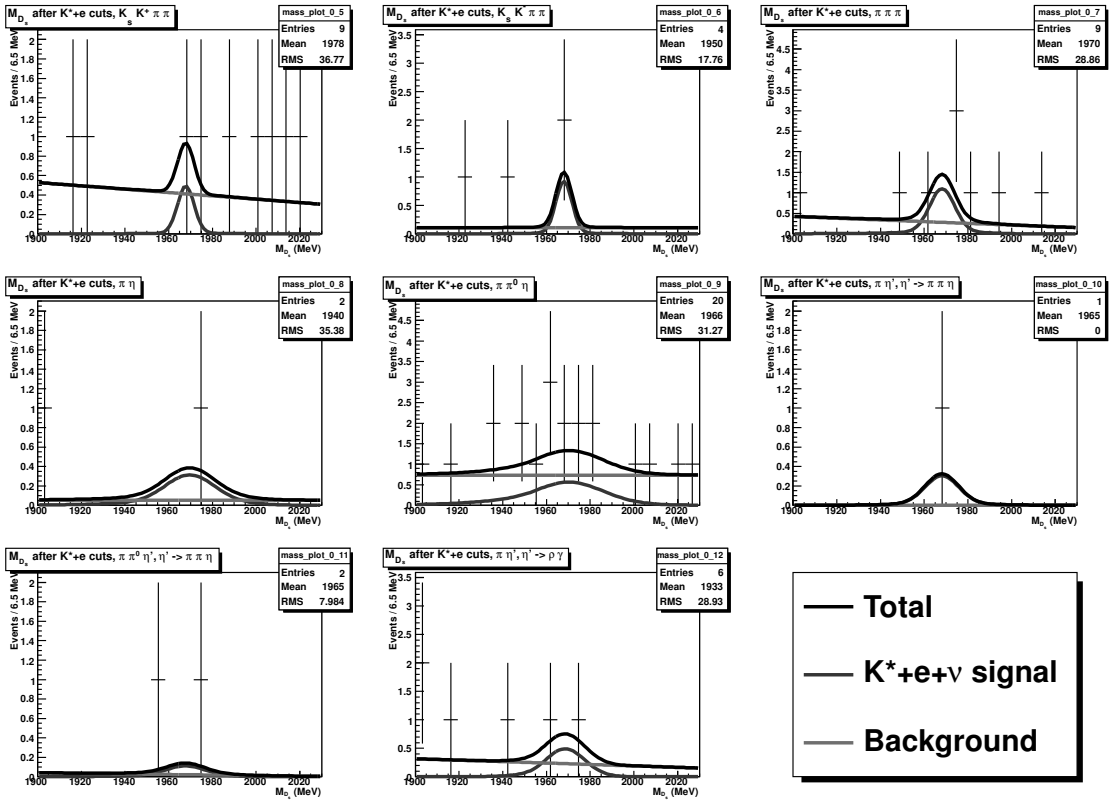

Figure G.39: $D_s \to K^* e\nu$ fit results in the data, after all semileptonic cuts, for tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

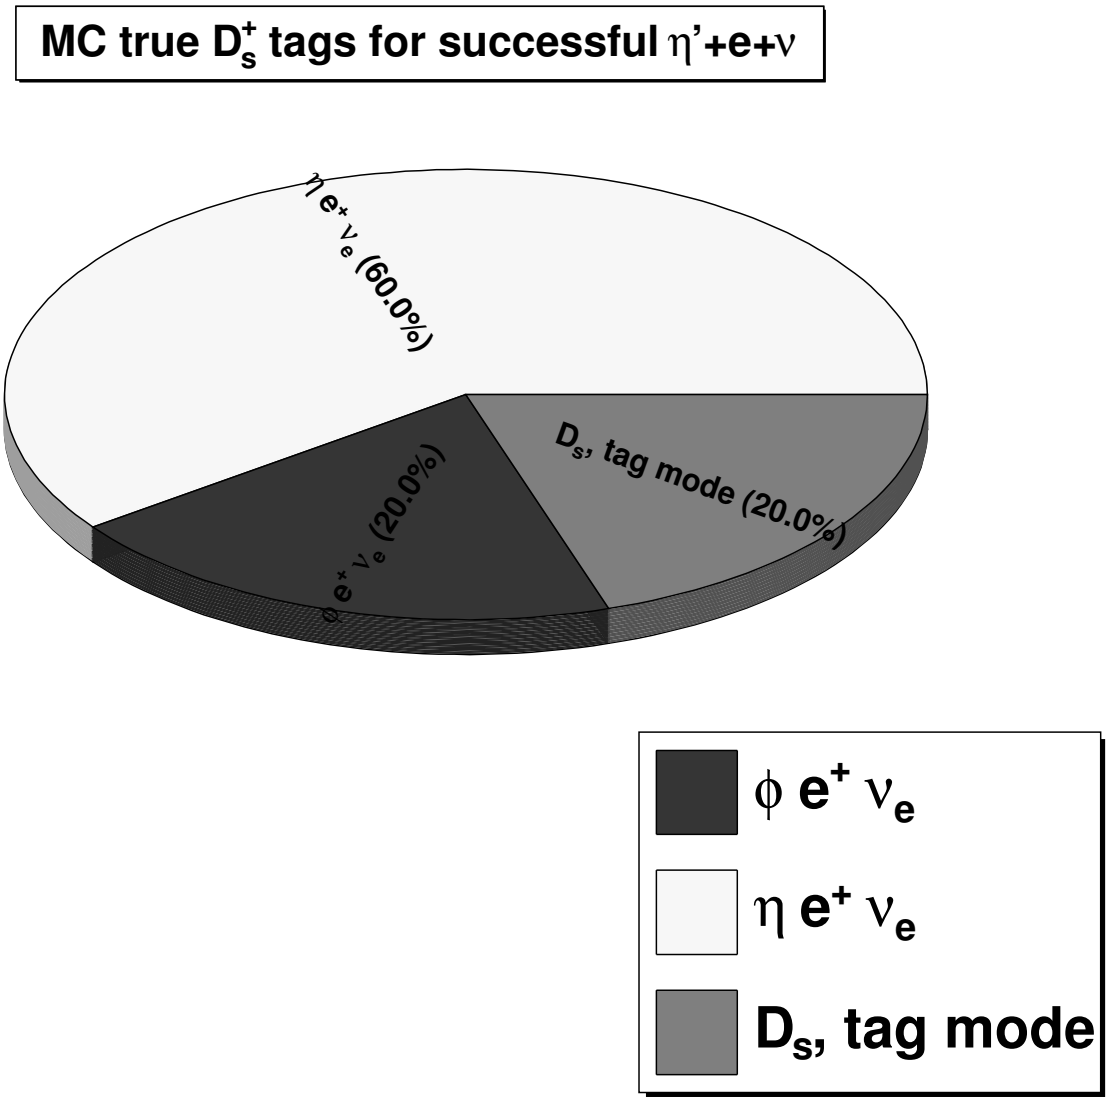

Figure G.40: $D_s \to \eta' e\nu$ backgrounds with a true D_s tag (peaking background) in the 20× Monte Carlo. $\eta e\nu$ with $\eta \to \pi\pi(\pi^0/\gamma)$ produces the most peaking background, while the dominant non-semileptonic contribution comes from D_s tag modes with a kaon faking the electron (e.g. $D_s \to KK_s\pi^0$)

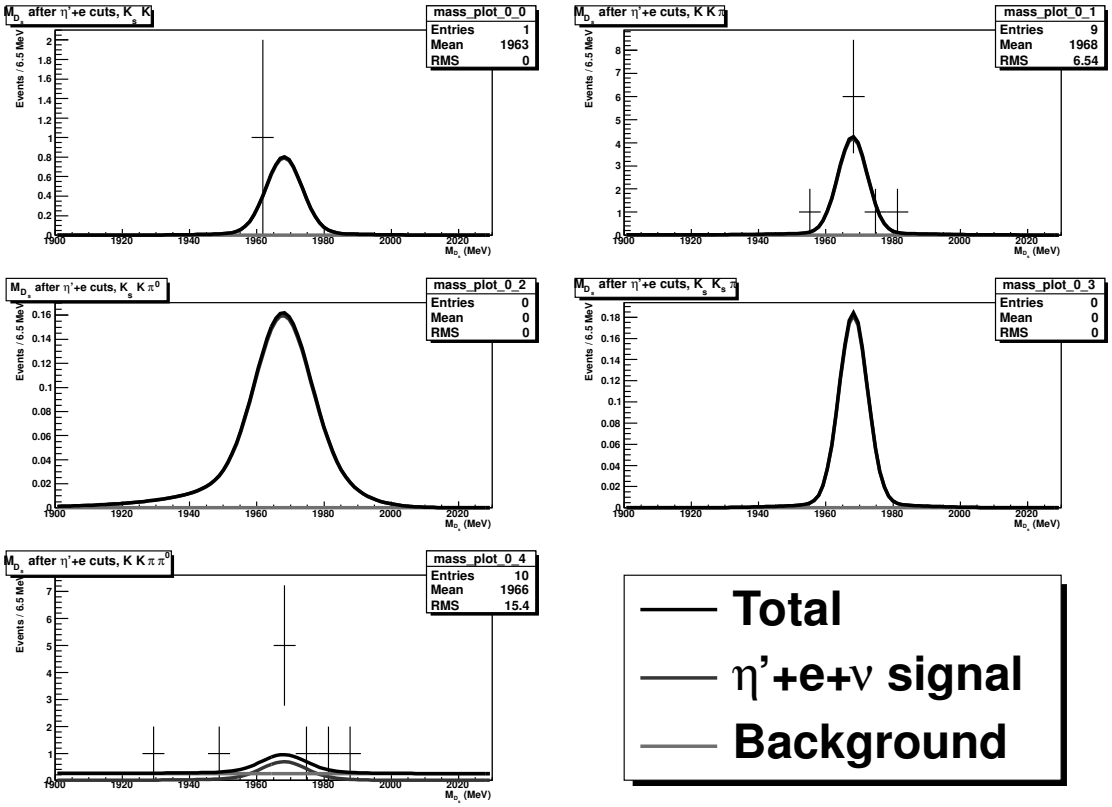

Figure G.41: $D_s \to \eta' e\nu$ fit results in the data, after all semileptonic cuts, for tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode receives an independent background normalization.

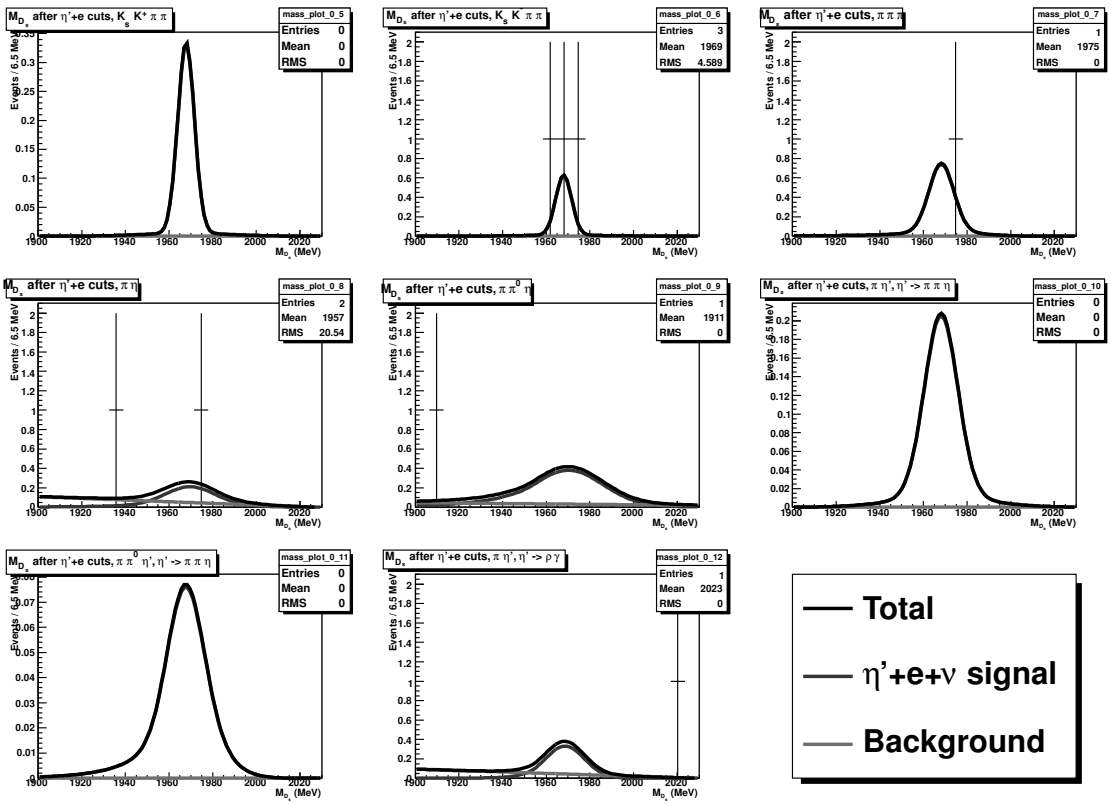

Figure G.42: $D_s \to \eta' e\nu$ fit results in the data, after all semileptonic cuts, for tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

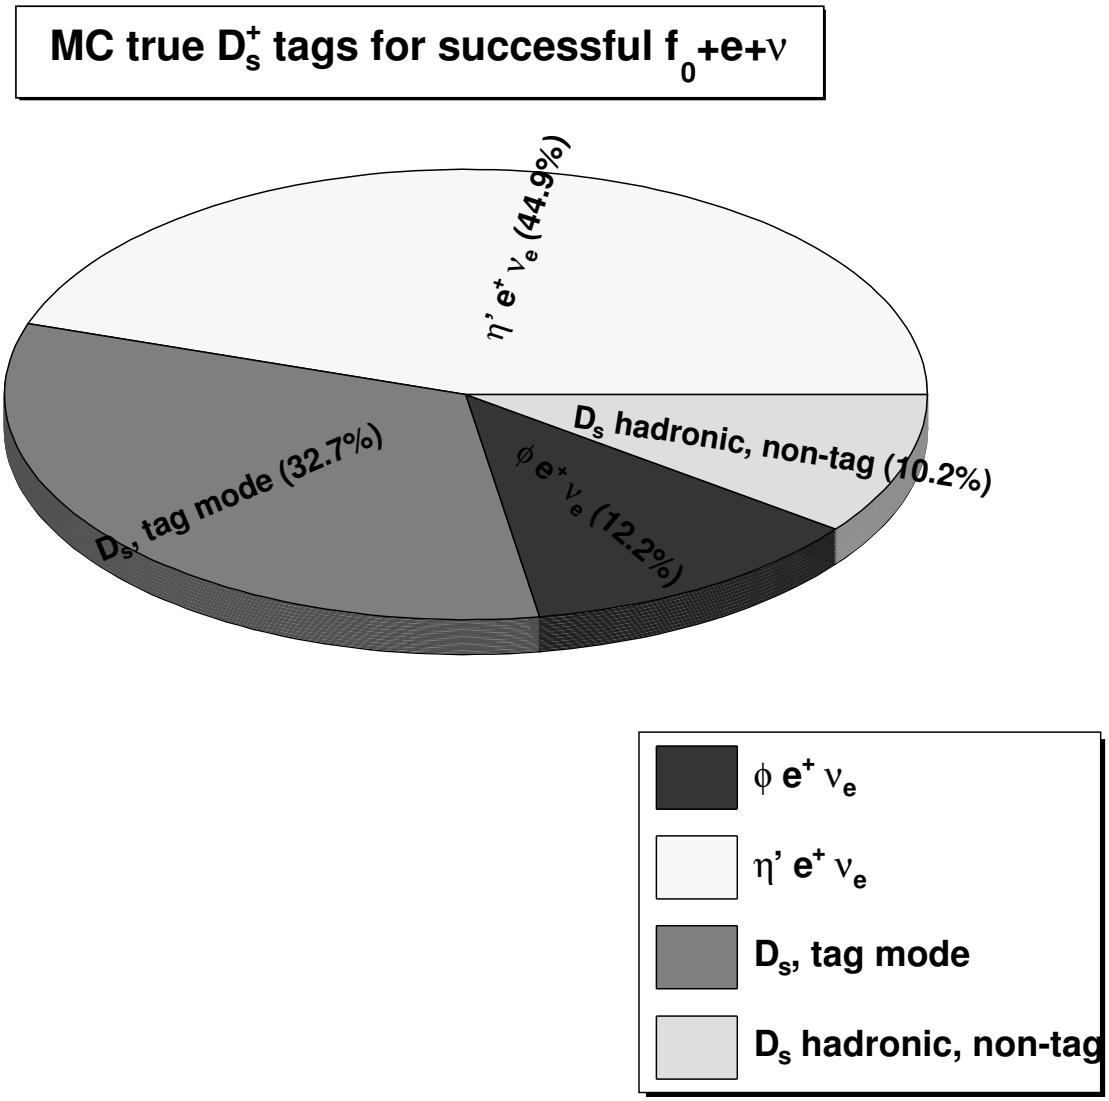

Figure G.43: $D_s \to f_0 e\nu$ backgrounds with a true D_s tag (peaking background) in the $20\times$ Monte Carlo. $\eta' e\nu$ with $\eta' \to \pi\pi X$ provides the plurality contribution, while the dominant non-semileptonic peaking background comes from D_s tag modes where a kaon fakes the electron (e.g. $D_s \to KK_s$).

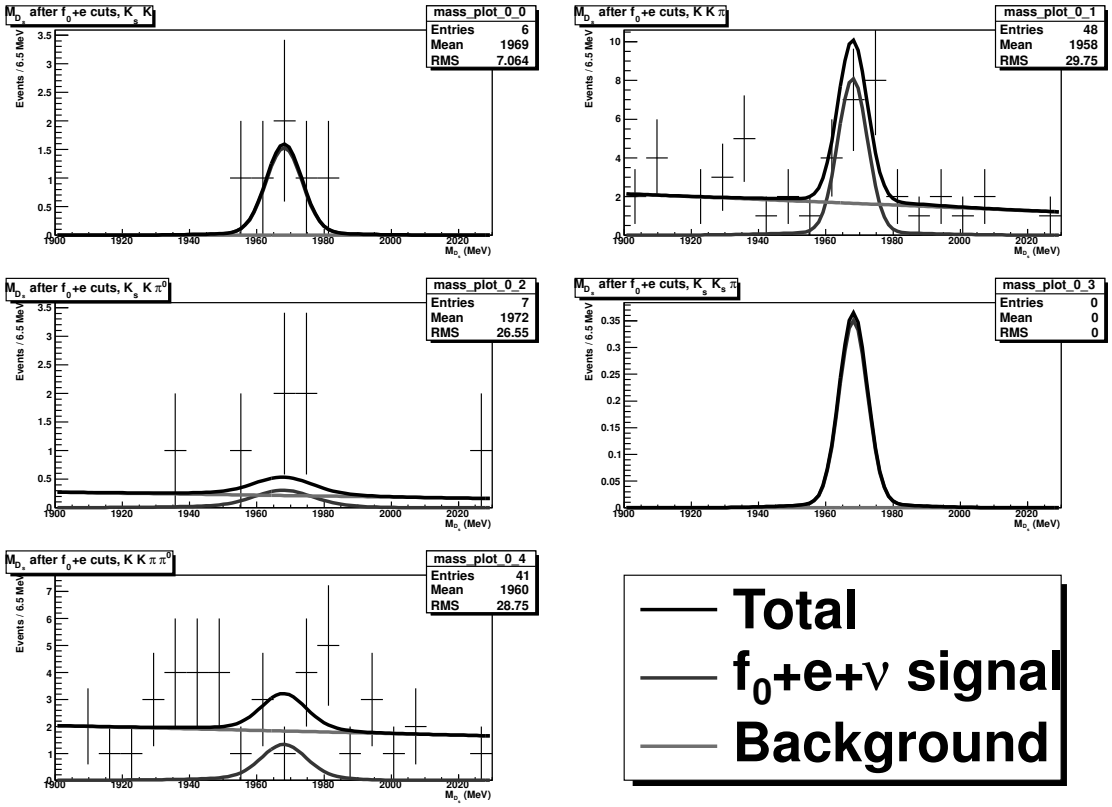

Figure G.44: $D_s \to f_0 e\nu$ fit results in the data, after all semileptonic cuts, for tag modes K_sK, $KK\pi$, $K_sK\pi^0$, $K_sK_s\pi$, and $KK\pi\pi^0$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode receives an independent background normalization.

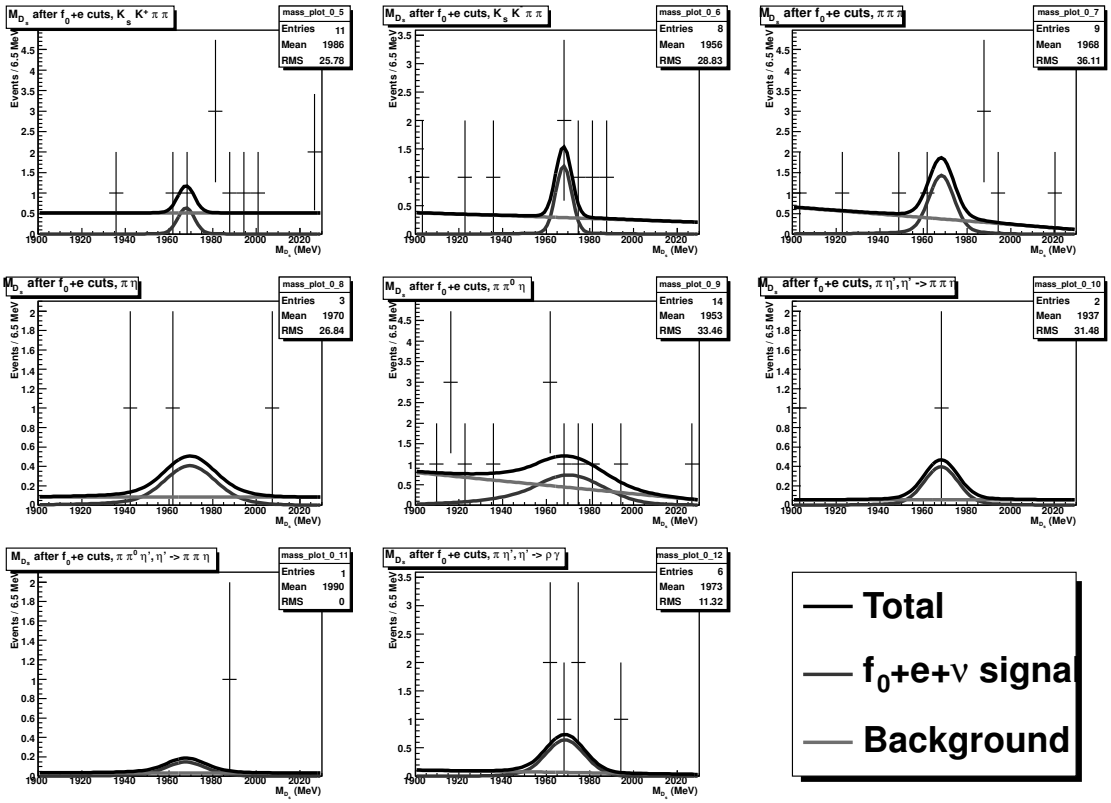

Figure G.45: $D_s \to f_0 e\nu$ fit results in the data, after all semileptonic cuts, for tag modes $K_sK^+\pi\pi$; $K_sK^-\pi\pi$; $\pi\pi\pi$; $\pi\eta$; $\pi\pi^0\eta$; $\pi\eta', \eta' \to \pi\pi\eta$; $\pi\pi^0\eta', \eta' \to \pi\pi\eta$; and $\pi\eta', \eta' \to \rho\gamma$. We fit the tagged M_{D_s} with a common signal normalization (branching ratio) for all 13 tag modes. Each mode does receive an independent background normalization.

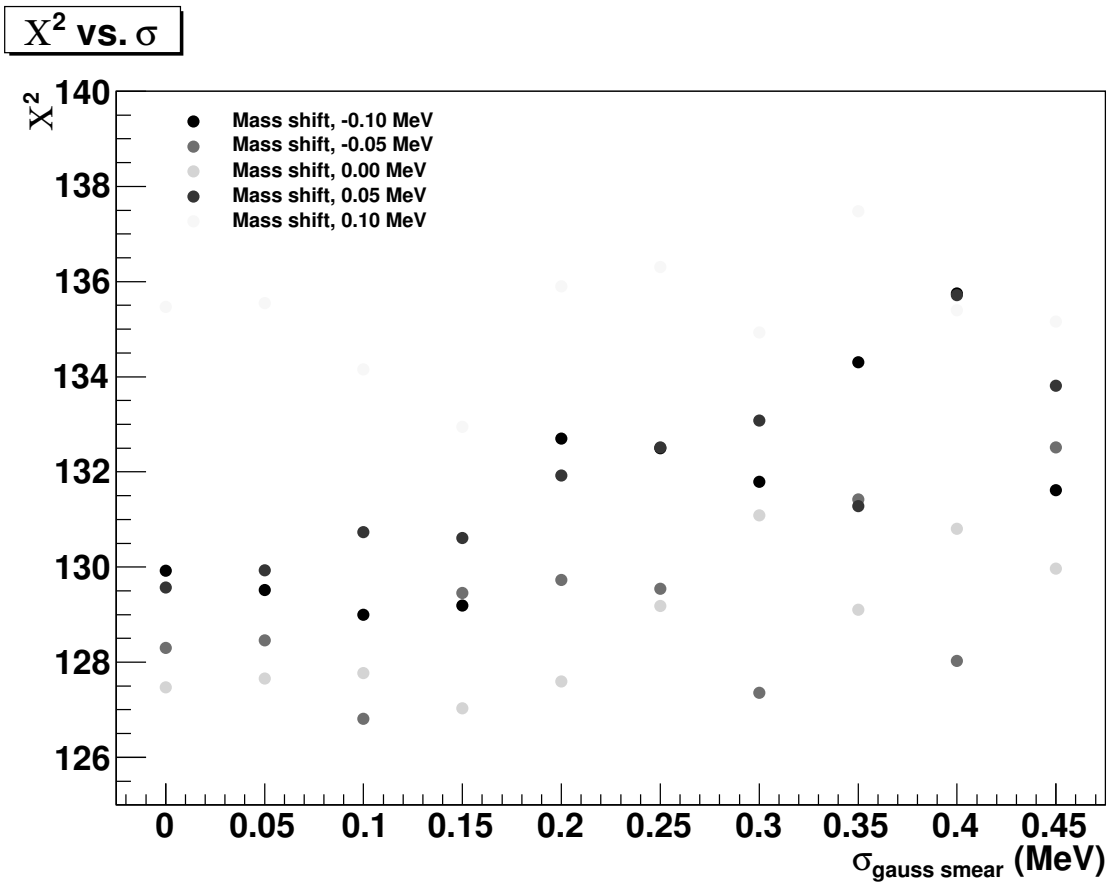

Figure G.46: χ^2 for data fits using various shifts and gaussian smears to the Monte Carlo's signal M_ϕ distribution. Smaller shifts and smears tend to be favored, implying a fairly accurate ϕ mass resolution in the Monte Carlo.

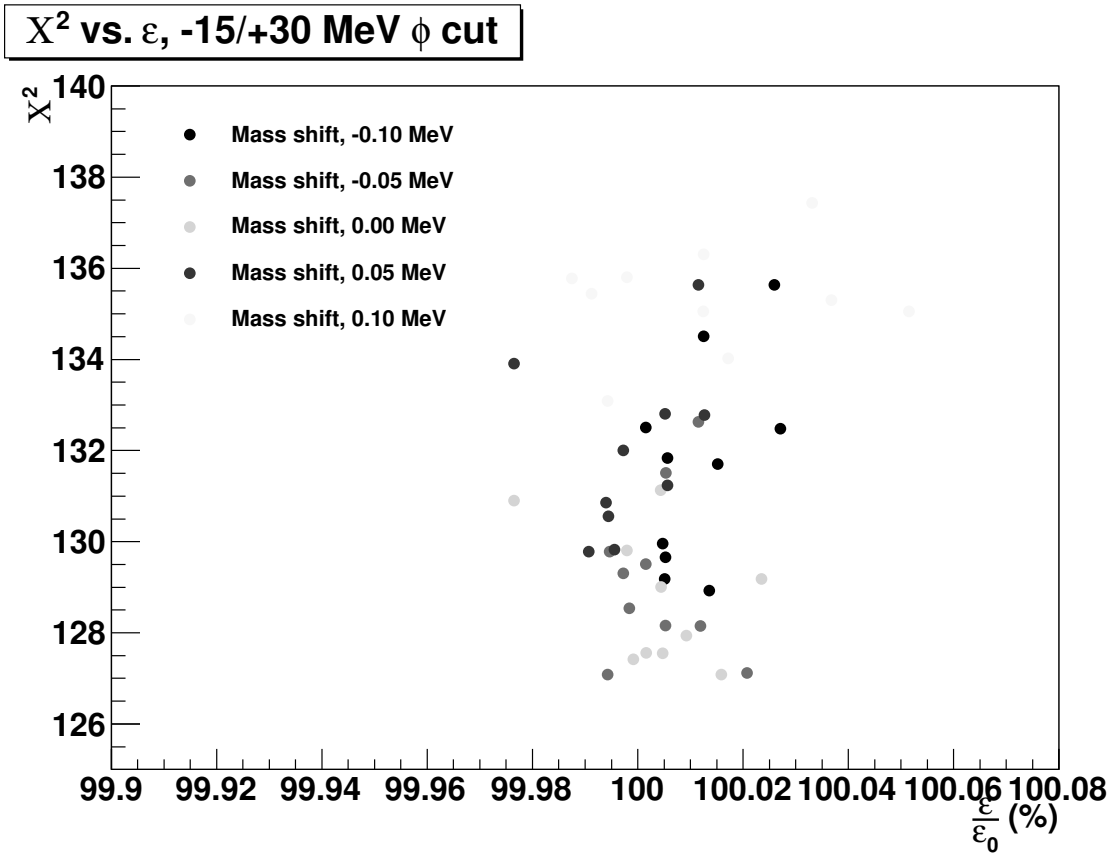

Figure G.47: Our large M_ϕ cut window means that even ϕ lineshapes that don't fit the data particularly well still have a relative efficiency difference from predicted ($\frac{\Delta\varepsilon}{\varepsilon_0}$) of less than 0.1%.

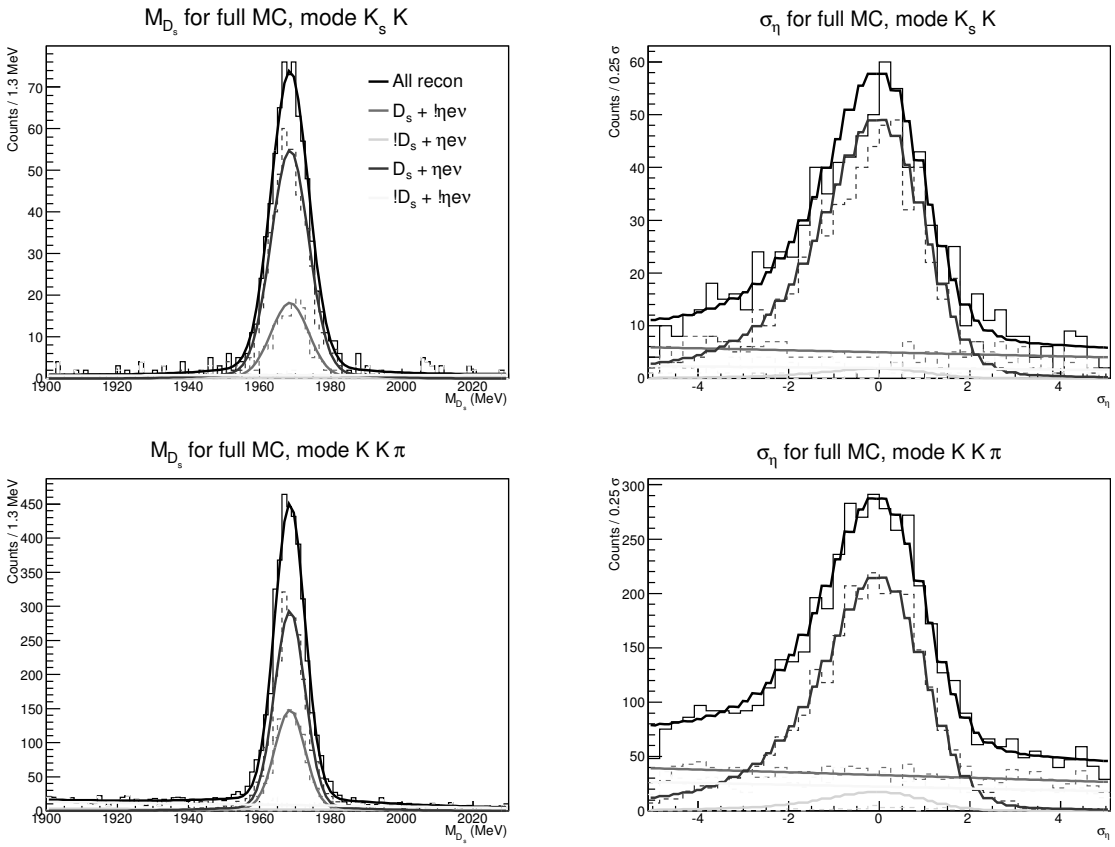

Figure G.48: The top two plots show the projections from our 2D $\eta e\nu$ fit to D_s invariant mass (left) and η pull mass (right) over the full 20× Monte Carlo when we tag $D_s \to K_s K$. The bottom two plots do the same for $D_s \to KK\pi$.

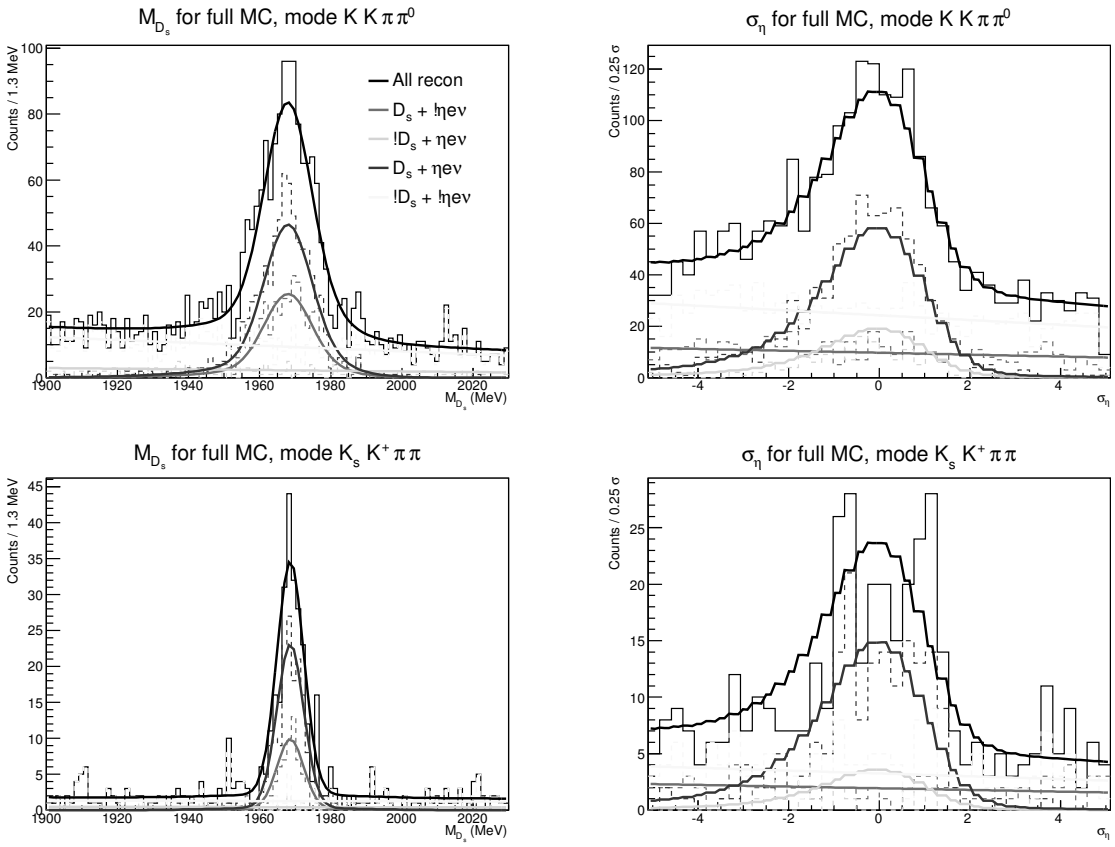

Figure G.49: The top two plots show the projections from our 2D $\eta e\nu$ fit to D_s invariant mass (left) and η pull mass (right) over the full 20× Monte Carlo when we tag $D_s \to KK\pi\pi^0$. The bottom two plots do the same for $D_s \to K_s K^+\pi\pi$.

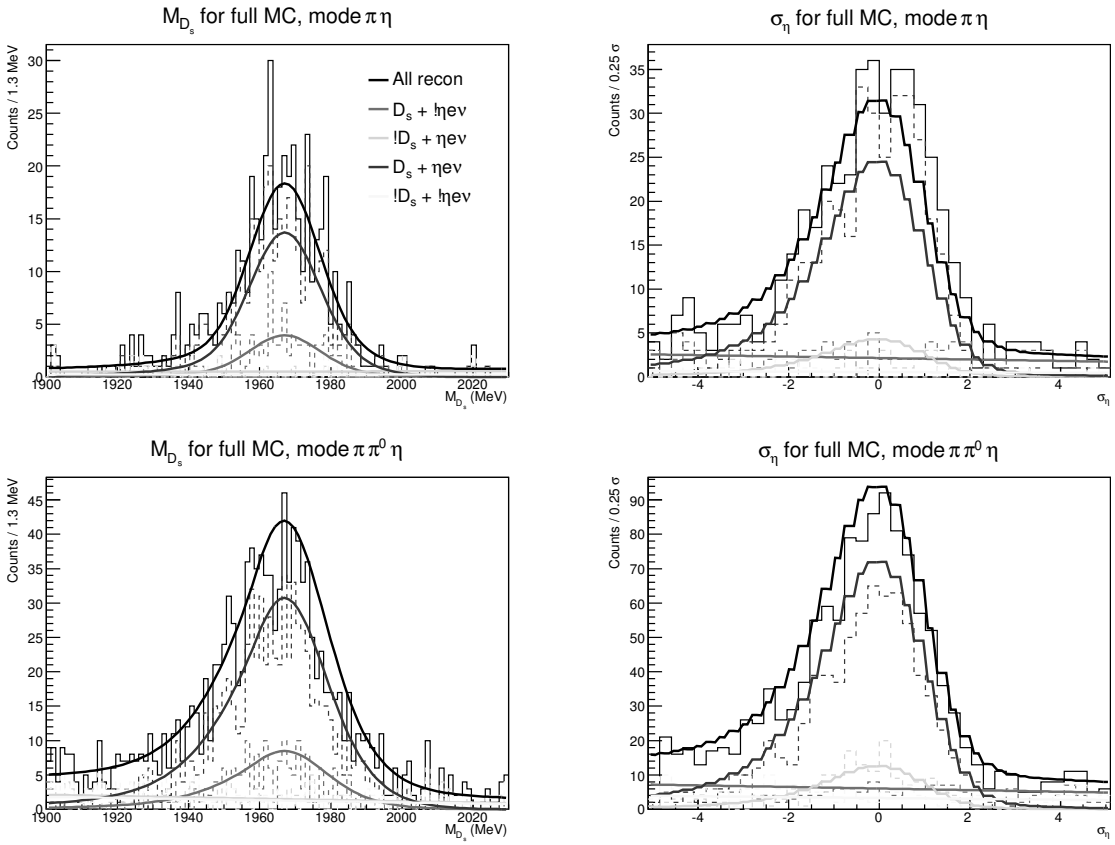

Figure G.50: The top two plots show the projections from our 2D $\eta e\nu$ fit to D_s invariant mass (left) and η pull mass (right) over the full 20$\times$ Monte Carlo when we tag $D_s \to \pi\eta$. The bottom two plots do the same for $D_s \to \pi\pi^0\eta$.

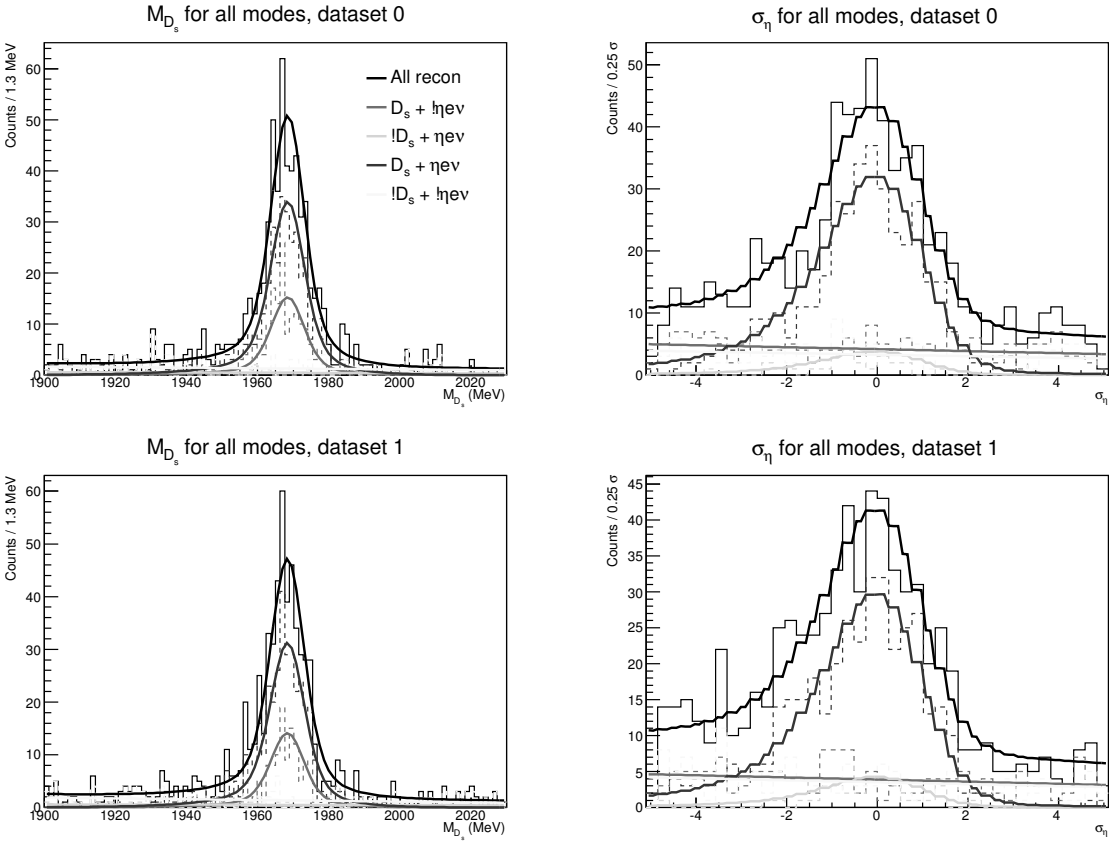

Figure G.51: The top two plots show the projections from our 2D $\eta e\nu$ fit to D_s invariant mass (left) and η pull mass (right) over all D_s tag modes for one data-sized Monte Carlo sample (dataset 0). The bottom two plots give the projections for a different data-sized Monte Carlo sample (dataset 1).

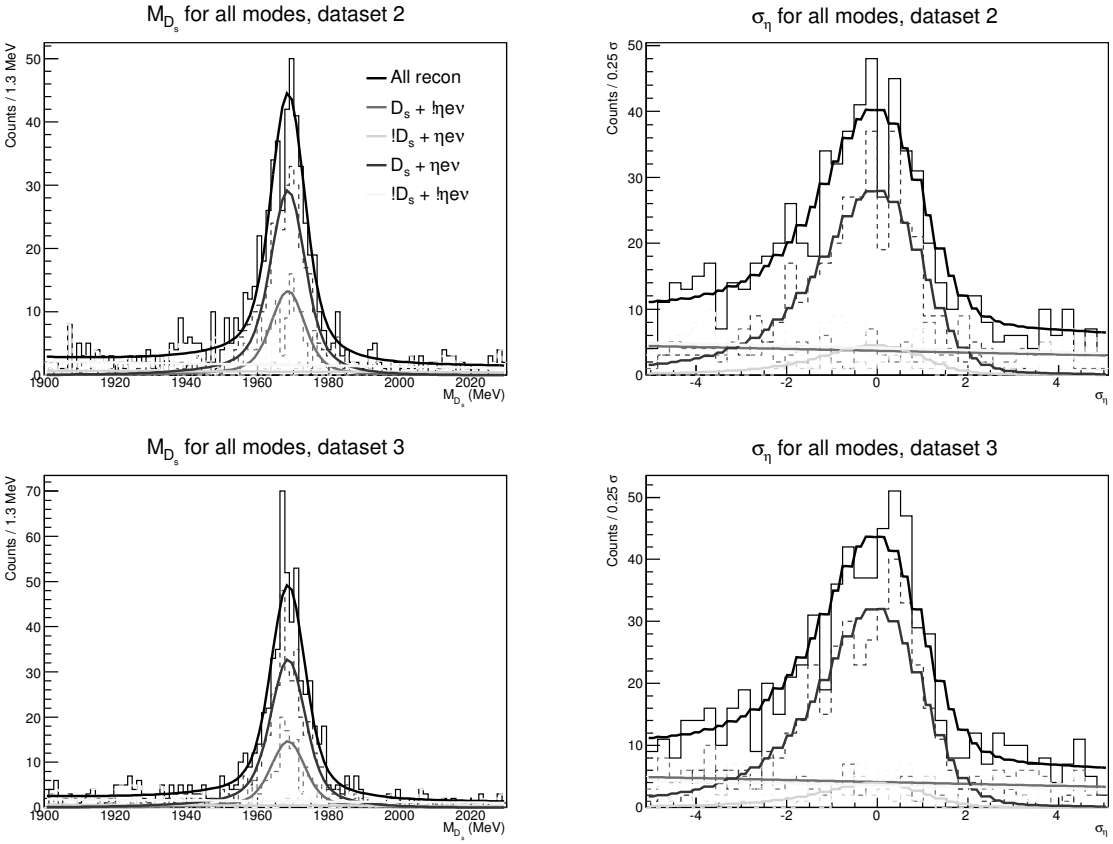

Figure G.52: The top two plots show the projections from our 2D $\eta e\nu$ fit to D_s invariant mass (left) and η pull mass (right) over all D_s tag modes for another data-sized Monte Carlo sample (dataset 2). The bottom two plots give the projections for a fourth data-sized Monte Carlo sample (dataset 3).

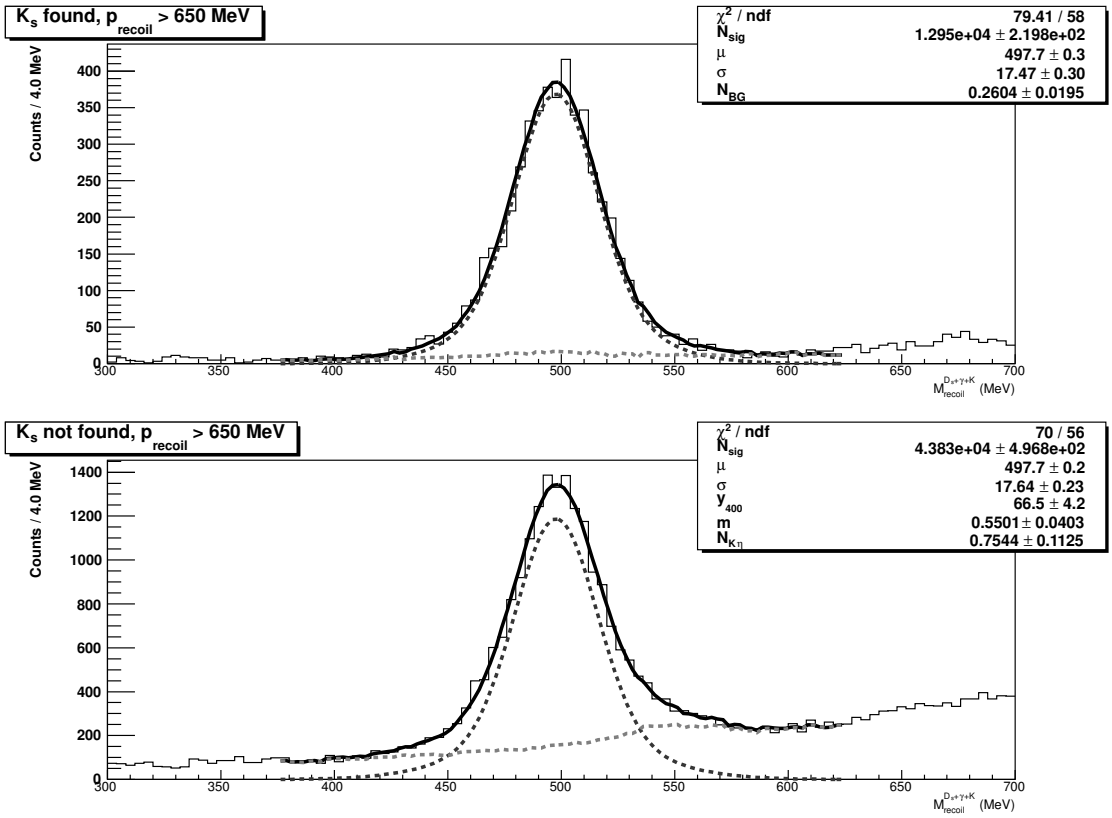

Figure G.53: $D_s + \gamma + K$ recoil mass in $K_s K$ events for "found" and "not found" K_s, from the Monte Carlo.

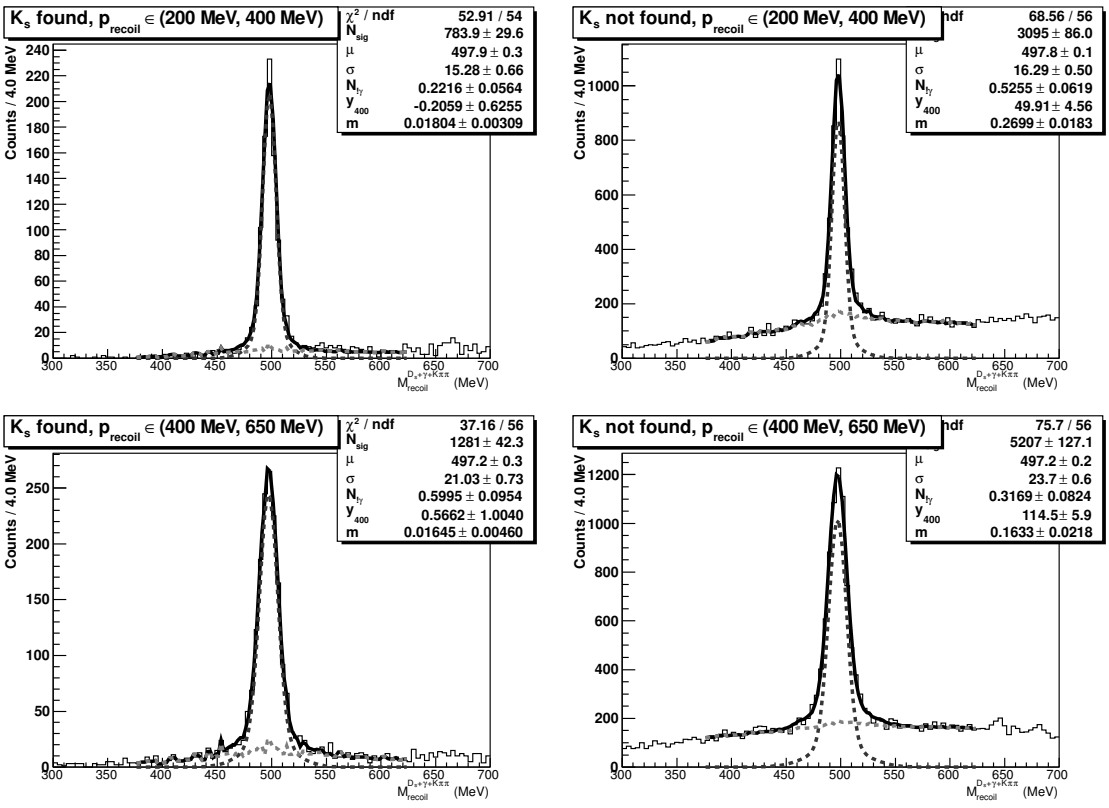

Figure G.54: $D_s + \gamma + K\pi\pi$ recoil mass in K^*K^* events for "found" and "not found" K_s, from the Monte Carlo. The top row corresponds to low momentum K_s while the bottom row corresponds to our medium K_s momentum region, as determined by the recoil momentum.

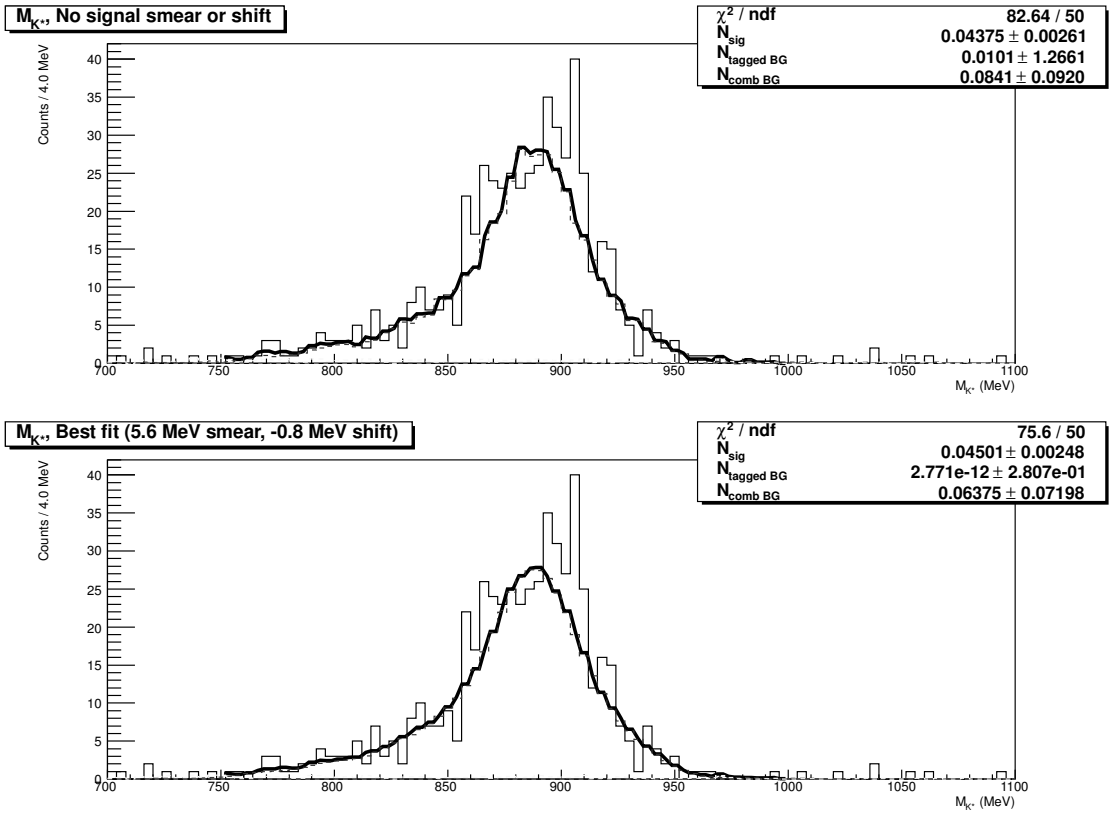

Figure G.55: Top: Signal shape fit to the K^* mass in K^*K. Bottom: K^* mass fit after allowing the M_{K^*} signal shape to shift left or right and convoluting it with a variable width gaussian.

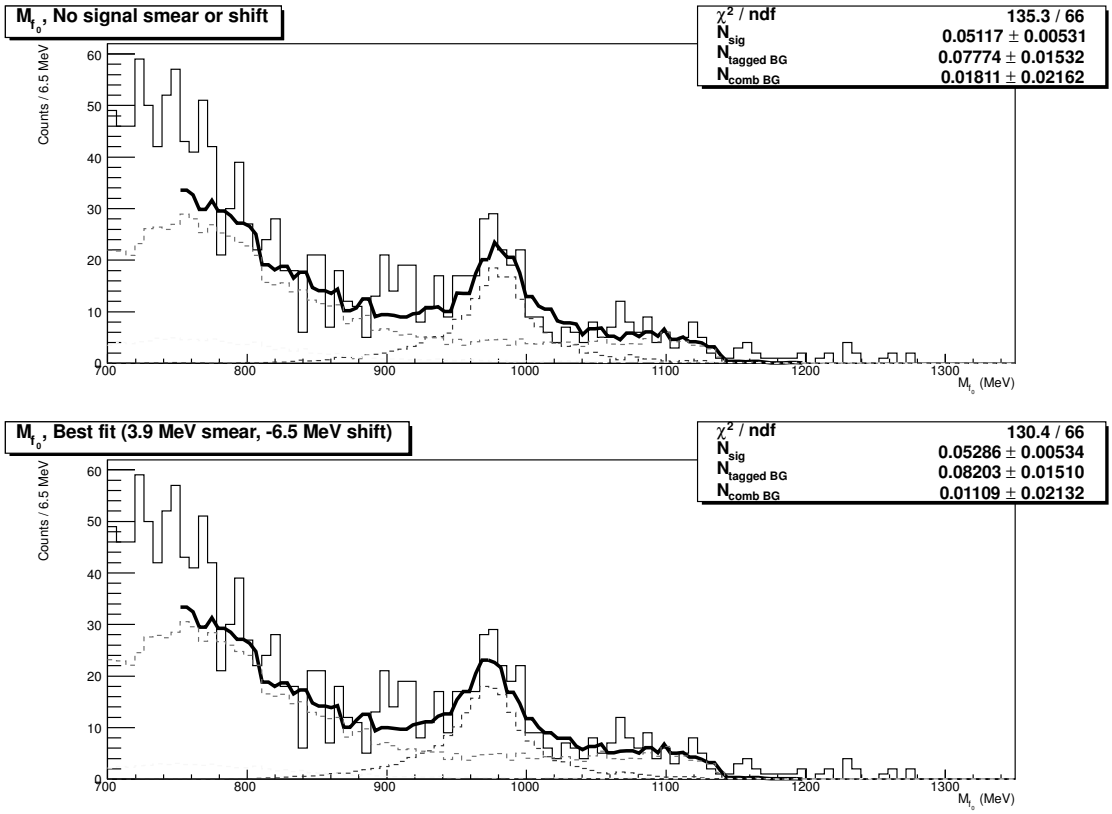

Figure G.56: Top: Signal shape fit to the f_0 mass in $f_0\pi$. Bottom: f_0 mass fit after allowing the M_{f_0} signal shape to shift left or right and convoluting it with a variable width gaussian.

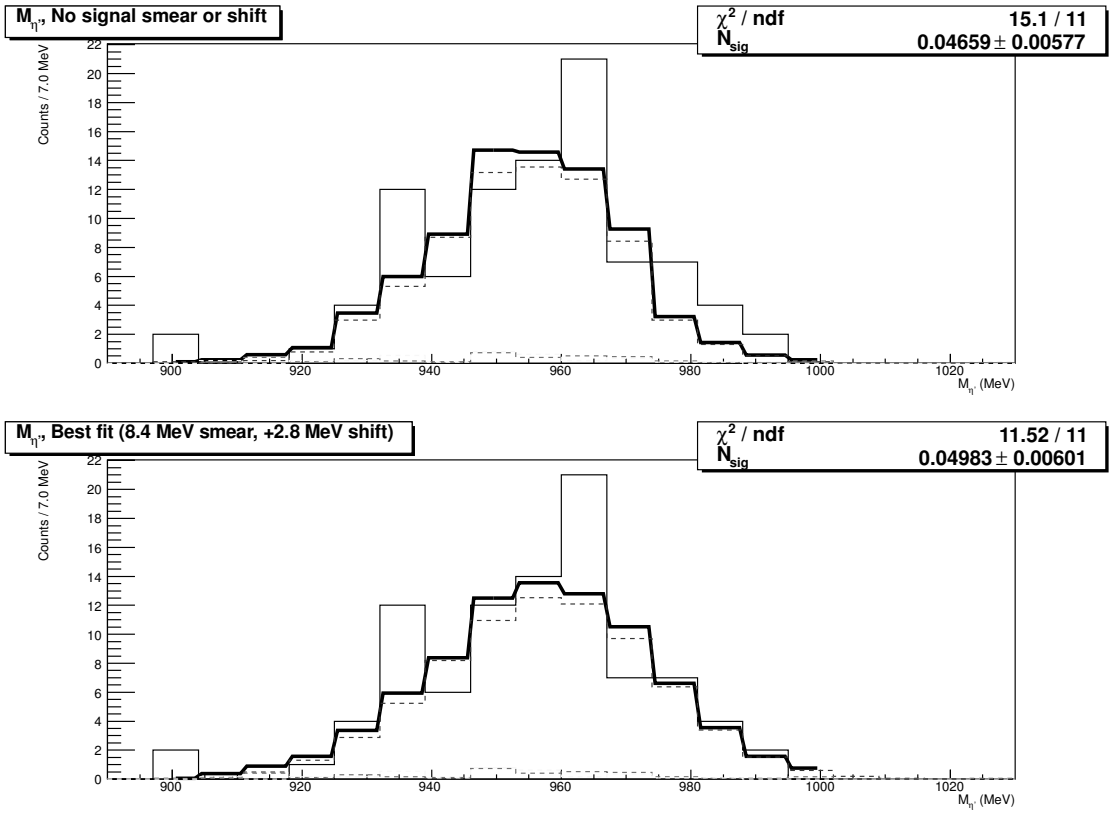

Figure G.57: Top: Signal shape fit to the η' mass in $\pi\eta', \eta' \to \pi\pi\eta$. Bottom: η' mass fit after allowing the $M_{\eta'}$ signal shape to shift left or right and convoluting it with a variable width gaussian.

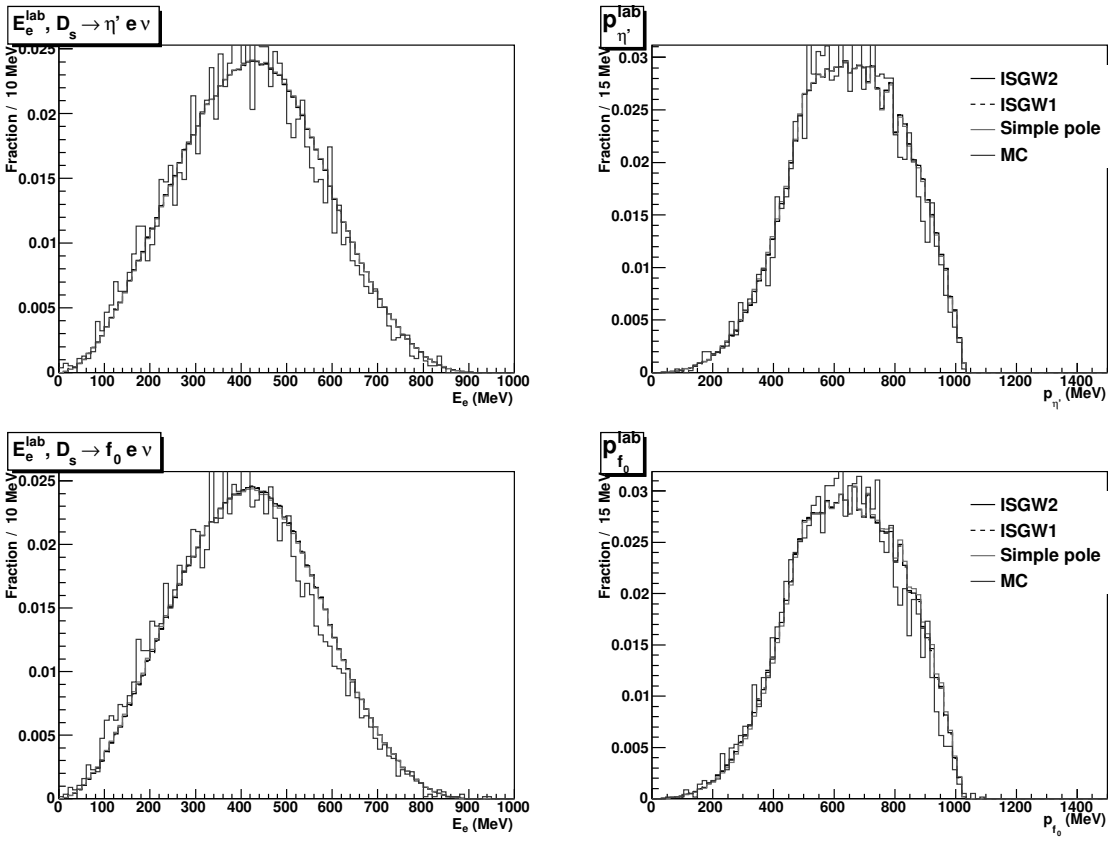

Figure G.58: Top: Lab frame electron energy (left) and η' momentum (right) in $D_s \to \eta' e\nu$ for the ISGW2 and pole models. Bottom: Lab frame electron energy and f_0 momentum in $D_s \to f_0 e\nu$.

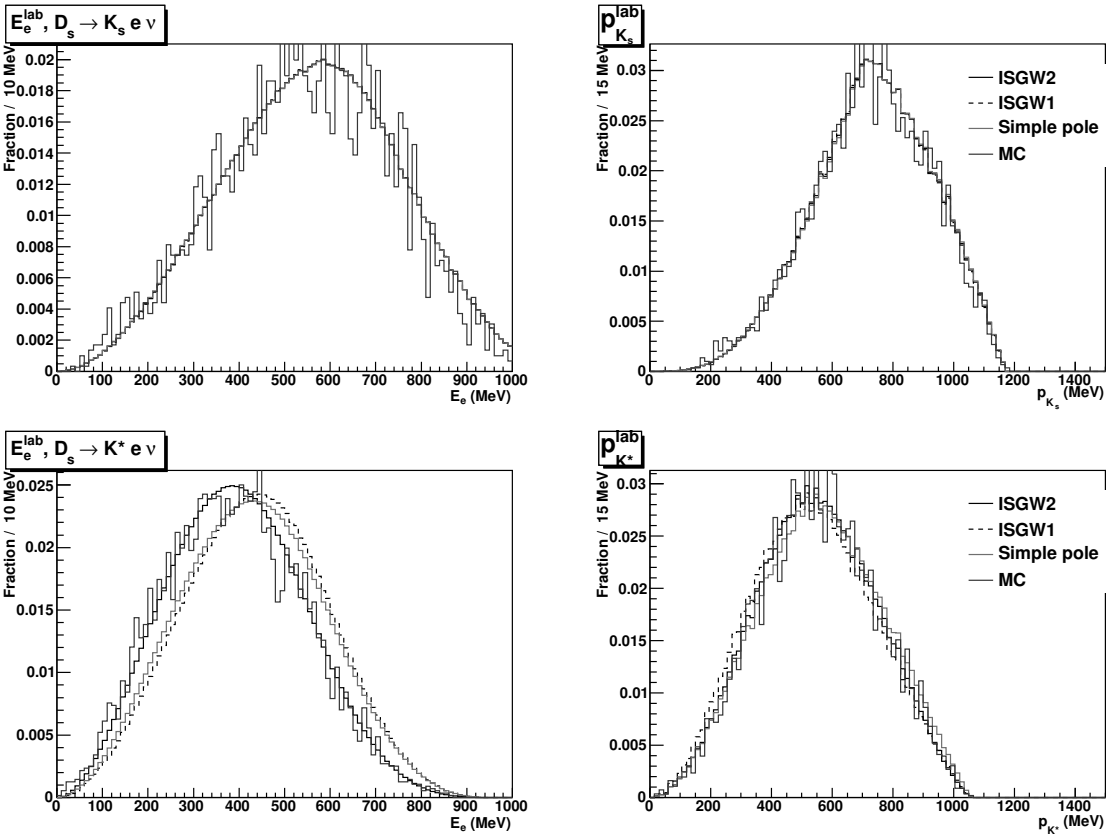

Figure G.59: Top: Lab frame electron energy (left) and K_s momentum (right) in $D_s \to K_s e\nu$ for the ISGW2 and pole models. Bottom: Lab frame electron energy and K^* momentum in $D_s \to K^* e\nu$.

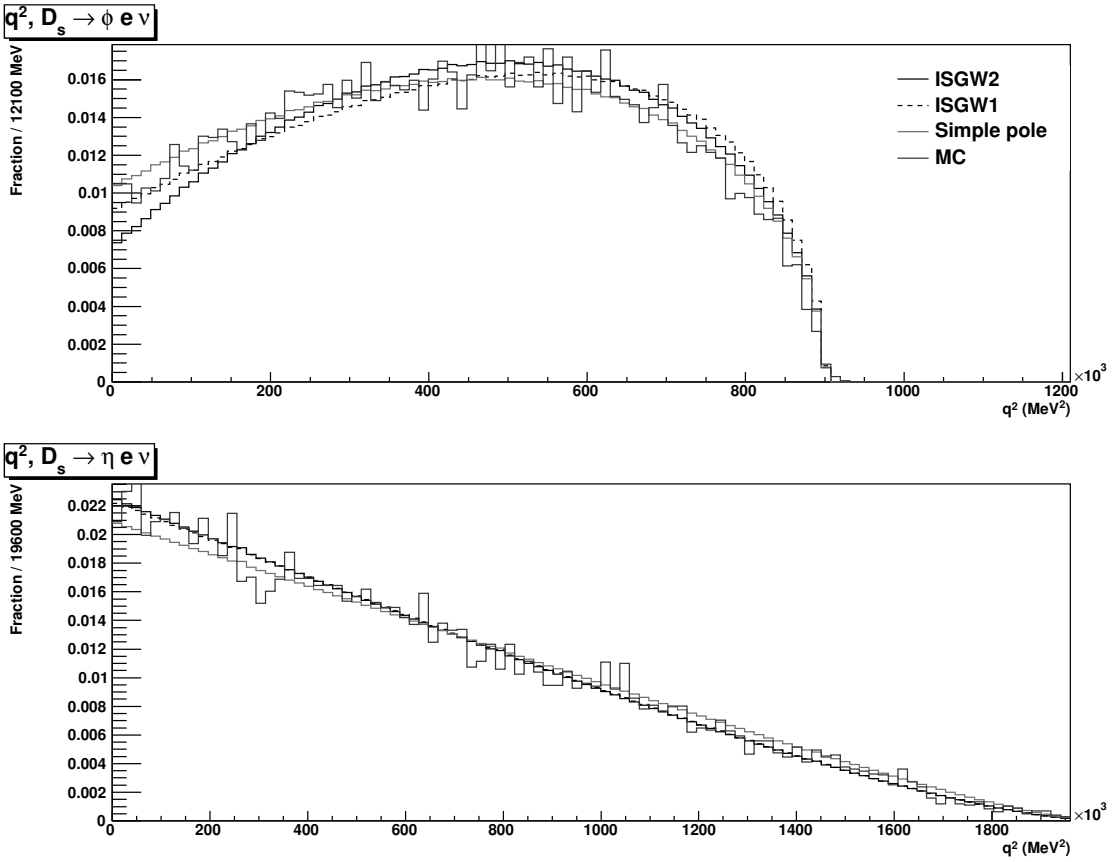

Figure G.60: q^2 distribution under different models for decays with a vector hadron ($D_s \to \phi e\nu$) and a pseudoscalar hadron ($D_s \to \eta e\nu$). The difference between the Monte Carlo and ISGW2 for low q^2 in $\phi e\nu$ comes from a correction we make to the Monte Carlo's masses.

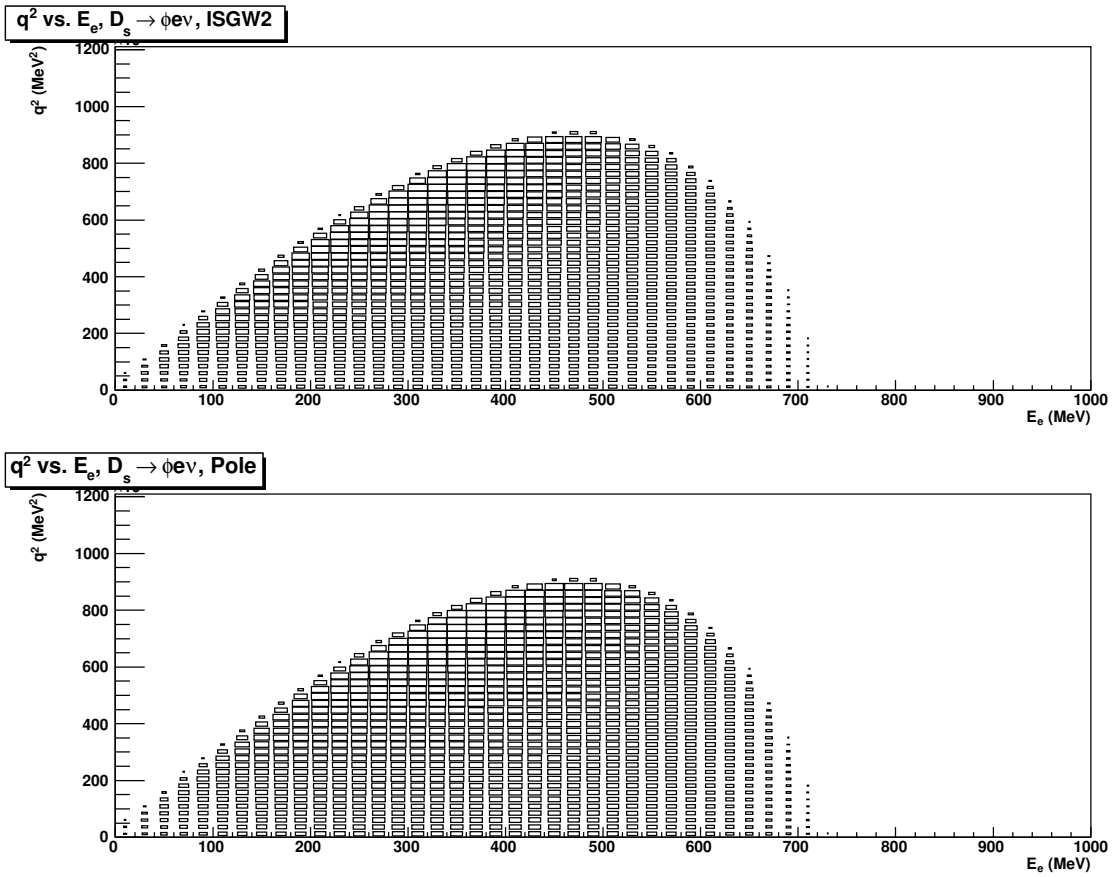

Figure G.61: Relationship between the q^2 and electron energy in the D_s rest frame for $D_s \to \phi e\nu$. Top: ISGW2 model. Bottom: Pole model.

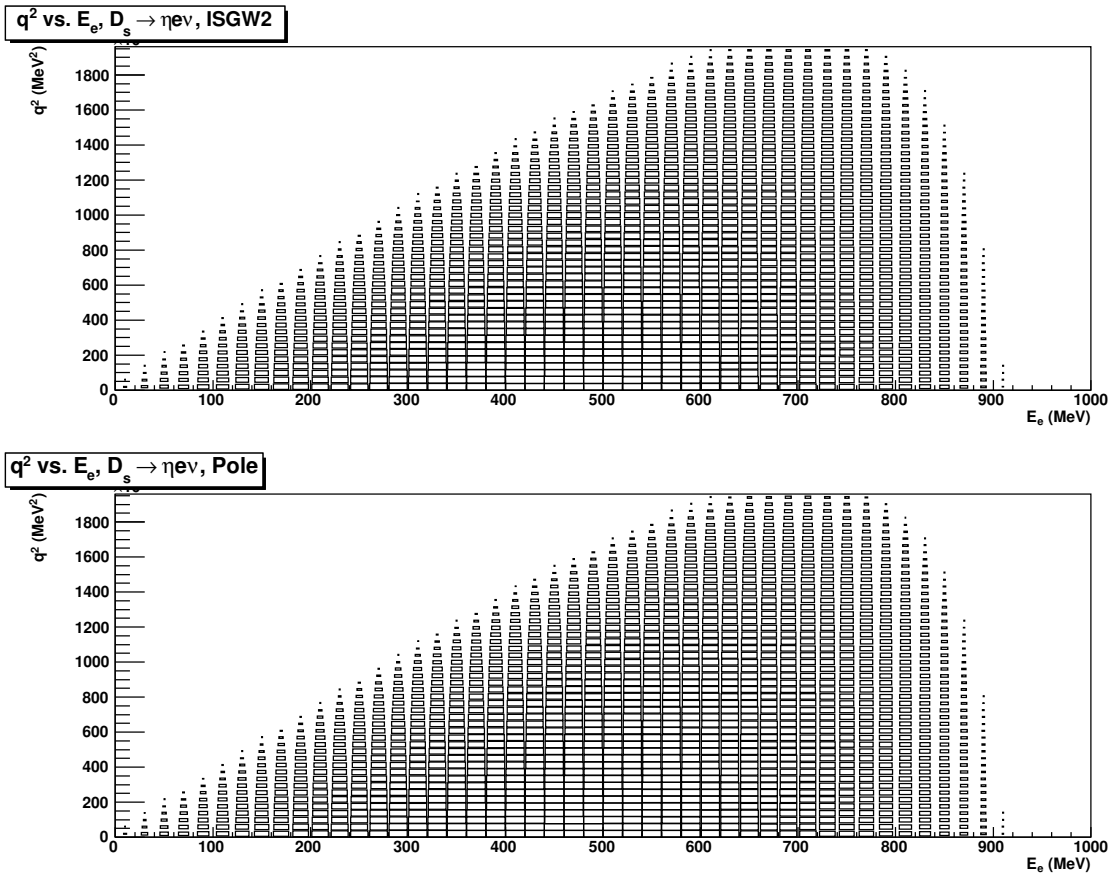

Figure G.62: Relationship between the q^2 and electron energy in the D_s rest frame for $D_s \to \eta e \nu$. Top: ISGW2 model. Bottom: Pole model.

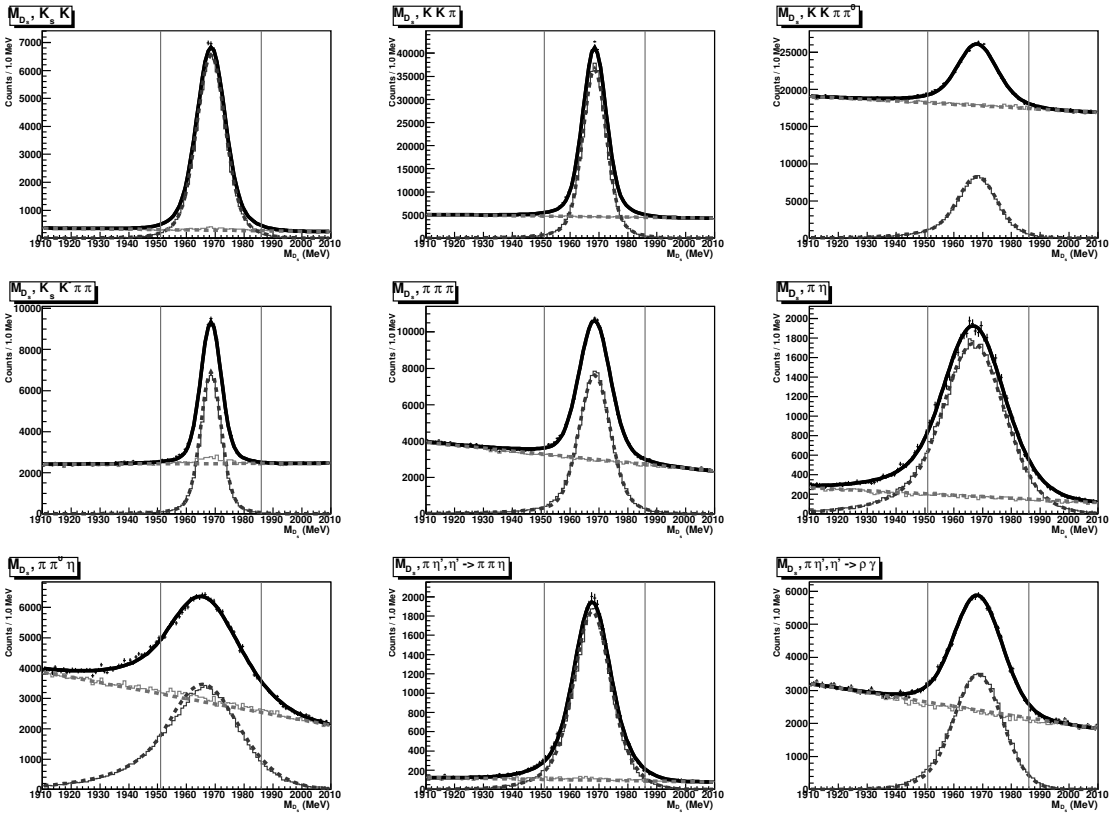

Figure G.63: D_s invariant mass fits after making a $D_s + \gamma$ recoil mass cut in the Monte Carlo. The dotted lines give our signal and background fit functions, while our solid lines give the truth-tagged signal and background.

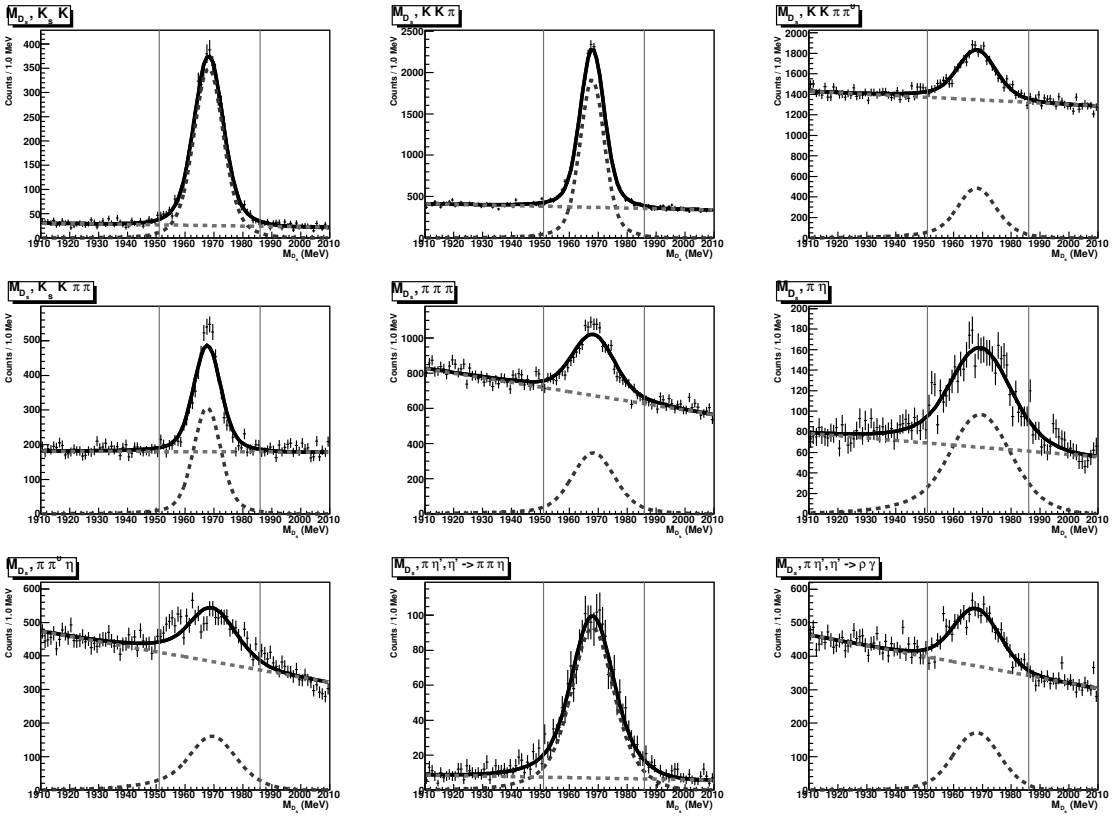

Figure G.64: D_s invariant mass fits after making a $D_s + \gamma$ recoil mass cut in the data. The dotted lines give our signal and background fit functions.

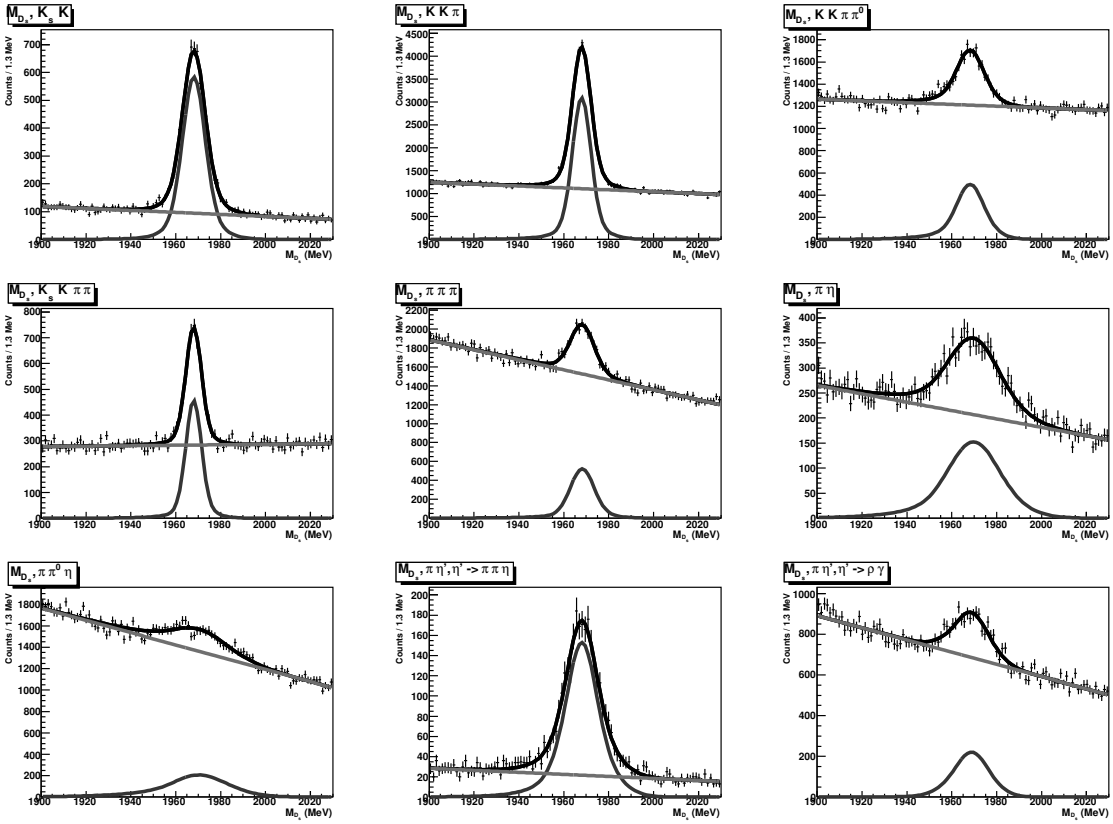

Figure G.65: D_s invariant mass fits using a gaussian+crystal ball signal shape and a linear background fit function.