A MEASUREMENT OF *B* HADRON CORRELATIONS IN $p\bar{p}$ COLLISIONS AT $\sqrt{s}$ = 1.8 TeV

BY

KEVIN PATRICK LANNON

B.A., St. Norbert College, 1997

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
Illinois at Urbana-Champaign, 2003

Urbana, Illinois

A MEASUREMENT OF *B* HADRON CORRELATIONS
IN $p\bar{p}$ COLLISIONS AT $\sqrt{s}$ = 1.8 TeV

Kevin Patrick Lannon, Ph.D.
Department of Physics
University of Illinois at Urbana-Champaign, 2003
Professor Steven Errede and Professor Kevin Pitts, Advisors

B hadron correlations provide input into the significance of higher-order $b\bar{b}$ production mechanisms. We present a study of *B* hadron correlations in $p\bar{p}$ collisions at $\sqrt{s}$ = 1.8 TeV. Events containing a $b\bar{b}$ pair are identified by the presence of an energetic electron or muon and both *B* hadrons in the event are located by reconstructing their decay vertices using precision charged-particle tracking. The transverse opening angle distribution, $\Delta\phi$, measured in data is compared to prediction from the PYTHIA and HERWIG Monte Carlo programs and found to be qualitatively similar. The closest matching Monte Carlo model is used to extract the *B* hadron $\Delta\phi$ distribution from the raw tag distribution. For the combined electron and muon sample, 28.8 ± 1.0 (stat.) $\pm$ 3.1 (syst.)% of *B* hadron pairs are found to have $\Delta\phi < 90°$.

For Linnea, who has given so much and waited so patiently.

Acknowledgements

In my time at the University of Illinois, I have had the pleasure of learning from and working with a number of gifted professors. Special thanks are due to three who have had a particular impact on me. I want to thank Steve Errede, my advisor, for convincing me to attend the University of Illinois in the first place and for getting me started on the muon calibration system. I also would like the thank Kevin Pitts, my other advisor. When it became clear that the physics analyses of Run II were too far out of reach for my tenure as a graduate student, Kevin came to the rescue with more ideas for thesis topics than I could ever hope to pursue. Finally, I would like to thank Tony Liss for inviting me to work on the CMP detector, an endeavor which has taken me from the bottom (actually below the bottom) to the top of the CDF detector, as well as for allowing me to stop working on the CMP when it came time to finish off this thesis.

In addition, I want to thank a number of other professors, postdocs, staff, and graduate students at the University of Illinois who have served as mentors, colleagues, and friends. Lee Holloway and Nathan Eddy started me off in the muon business, testing CMU electronics (and, of course, taught me everything I know about CMU charge division). Trevor Vickey taught me to solder and spent countless long hours with me pulling cables and testing hardware in order to get the muon calibration system and the CMP detector ready for Run II. Hyunsoo Kim was my buddy on more controlled accesses than I would like to recall. I would also like to thank him for kicking the back of my chair and asking me if I had graduated yet nearly every day for the last three months. I have to thank Lucio Cerrito for taking the job as "CMP postdoc" so that I could spend more time on analysis and less time in access. Dave Lesny provided excellent computing facilities, exceptional technical support, and copious disk space for the billion or so Monte Carlo events I generated. Finally, I want to thank Tom Shaw, Shirley Rothermel, Donna Guzy, and Joyce Ruch for their steadfast help and support in dealing with financial matters and making sure I got paid on time no matter how badly I mangled my paperwork.

This thesis would not have been possible without the work of the hundreds of professors, postdocs, graduate students, engineers, and techs who make up the CDF collaboration. In particular, I am deeply indebted to the *B* group, and especially the *B* group conveners, Kevin

Pitts, Michael Schmidt, Marjorie Shapiro, and Christoph Paus. Without their support and willingness to make time in the agenda for a Run 1 analysis during the excitement of the start-up of Run II, I would never have been able to complete this analysis. I also owe many thanks to the Run 1 *B* physics student group, especially Tony Affolder and Michael Guenther, who served as an island of Run 1 calm in the turbulent seas of Run II frenzy. I am particularly thankful for Rick Field's theoretical insight into $b\bar{b}$ correlations, his work in generating interest in this subject and in my analysis in particular (I think more people learned about my results in his talks than in mine), and for all his guidance in tuning and running the Monte Carlo generators used in this analysis. Lastly, I would like to thank Bob Wagner (a.k.a. ArgoBob) who has been my mentor, collaborator, and friend for far longer than I have been a graduate student. It was Bob who taught me how to be a physicist.

Before graduate school there was St. Norbert College and many professors who helped get me started down this road in one way or another. In particular, I would like to thank John Spangler, Kyoko Mori, John Neary, John Pennington, and Rick Poss.

Life is not only physics and I have to thank a number of people who kept my life bright when the physics looked dark. In particular, I am grateful for the Fermilab Ultimate group for many afternoons of summer fun. It is refreshing how unimportant physics can seem when you're chasing down a long toss. I am also deeply thankful for the friendship of Ron Bishop, Matt Dusek, Steve Lute, Matt Mikulcik, Neil Saigh, and Justin Zamirowski, whom I have known since before I can remember. You guys are pretty much the only ones who can keep me up past my bedtime anymore.

I could never have written this thesis without the love and support of my family: my parents Jerome and Barbara, my siblings, Beth, Katie, Keegan, Kiernan, and Brianne, a host of aunts, uncles and cousins, and my grandmother. They have been my strength in hardship, have kept me humble, and have made me proud. To these, I have to add my family-in-law, especially Virgil and Patricia Schreiner who treated me as their son even before I married their daughter.

Finally, I am deeply thankful and infinitely indebted to my wife Linnea. Without your love and support, I would never have made it. Thank you.

This thesis was supported in part by the United States Department of Energy under grant DE-FG02-91ER-40677.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Particle physics is the study of the fundamental constituents of the universe, the basic building blocks of everything we experience in the world around us. This field of study is also referred to as high-energy physics because these fundamental constituents of matter only become manifest by studying interactions at energy scales much higher than those that define our everyday experience. To create these high-energy environments, particle physicists rely on huge particle accelerators that take the ordinary and easily accessible subatomic particles that make up the matter of our everyday experience, namely protons and electrons, accelerate them nearly to the speed of light, and bring them into collision.

This course of accelerator-based experiments coupled with observations from other sources like cosmic ray physics has produced the following picture of what makes up our universe (or at least, the portion of our universe with which we're most familiar). All matter seems to be composed of bound states of structureless, point-like, half-integer spin particles known as fermions. These fermions come in two types, leptons and quarks, depending on the interactions they experience, as discussed below. The fermions are also organized into three groups, or *generations*, of particles. Each generation contains a pair of leptons and a pair of quarks, where the individual elements of the pair differ in charge by one unit. The electron with charge $-e$ and neutral neutrino make up the first generation of leptons, while the first generation of quarks is composed of the up quark with charge $+2/3e$ and the down quark with charge $-1/3e$. (The symbol e is defined to be a unit of charge equivalent in magnitude to the charge of an electron or approximately 1.602×10^{-19} C.) The other two generations contain essentially heavier copies of the particles in the first generation. Table 1.1 lists some of the characteristics of the twelve fundamental fermions that make up matter. In addition, it should be noted that for each particle in Table 1.1, there is a corresponding antiparticle, identical to the particle in every respect except

	Leptons			Quarks		
Generation	**Flavor**	**Mass**	**Charge**	**Flavor**	**Mass**	**Charge**
1st	ν_e	< 3 eV/c^2	0	u	1.5–4.5 MeV/c^2	2/3e
	e	0.511 MeV/c^2	-e	d	5–8.5 MeV/c^2	-1/3e
2nd	ν_μ	< 0.19 MeV/c^2	0	c	1.0–1.4 GeV/c^2	2/3e
	μ	105.7 MeV/c^2	-e	s	80–155 MeV/c^2	-1/3e
3rd	ν_τ	< 18.2 MeV/c^2	0	t	174.3 GeV/c^2	2/3e
	τ	1.777 GeV/c^2	-e	b	4.0–4.5 GeV/c^2	-1/3e

Table 1.1 The fundamental fermions that make up matter. For each particle listed in this table, there is also an antiparticle with opposite quantum numbers, like charge.

with opposite charge and other quantum numbers. For example the antiparticle partner of the electron is the positron, which has the same mass as the electron but is charge +e.

In addition to revealing the spectrum of fundamental fermions that compose matter, high-energy physics experiments have also enabled us to classify all interactions between matter in terms of four fundamental forces known as the gravitational, electromagnetic, weak, and strong force. Each of these forces is mediated by the exchange of one or more type of integral-spin bosons between quarks and leptons. The mediator of the electromagnetic force is the familiar massless photon. A trio of massive gauge bosons known as W^-, W^+, and Z^0 mediates the weak force, while the exchange of eight massless gluons is responsible for the strong force. However, not all particles participate in the strong force, allowing a distinction between quarks, which couple to gluons, and leptons, which do not. Although not confirmed yet by experiment, it is believed that the carrier of the gravitational force is the massless, spin-2 graviton. Table 1.2 lists some properties of the four fundamental forces.

Arguably, the greatest accomplishment to date in high-energy physics has been the creation of a general theoretical framework to describe the strong, weak, and electromagnetic interactions among quarks and leptons. This framework is known as the Standard Model, and together with the calculational technique of perturbation theory, it dominates most of the ongoing research in particle physics.

Force	Boson	Boson Mass	Fermions Affected
Gravitational	Graviton	0	Quarks and leptons
Electromagnetic	Photon (γ)	0	Quarks and leptons
Weak	W^+, W^- Z^0	80.4 GeV/c^2 91.2 GeV/c^2	Quarks and leptons
Strong	Gluon (g)	0	Quarks

Table 1.2 The four fundamental forces.

1.1 The Standard Model

The Standard Model uses quantum field theory, which combines quantum mechanics and special relativity, to describe the dynamics of the fundamental fermions and their interactions under the strong, weak, and electromagnetic forces. Each kind of quark and lepton is described by a separate, four-component Dirac spinor field, and the dynamics of the particles are encapsulated in a Lagrangian containing terms involving the fermion fields and their derivatives [1]. Symmetries observed in nature are imposed on the Standard Model Lagrangian to constrain its form. For example, relativity requires the Lagrangian to be Lorentz invariant. A particularly powerful set of symmetries for dictating the form of the Standard Model Lagrangian are local gauge symmetries. A gauge transformation is a continuous transformation applied to the fields in a quantum field theory. It is usually characterized according to the symmetry groups, such as U(1) or SO(10), from group theory. The gauge symmetry is considered local if a different transformation can be applied at every point in space without ruining the symmetry. However, to keep the Lagrangian invariant under local gauge transformations, it is necessary to introduce additional integral-spin fields, known as gauge fields, that also obey certain transformation laws and couple to the fermion fields. If the correct gauge symmetries are imposed on the Lagrangian, the necessary fields introduced can be identified with the fundamental forces of the Standard Model [2]. For this reason, the bosons that mediate the different fundamental interactions are often referred to as *gauge bosons*. For example, requiring the Lagrangian to be invariant under U(1) phase shifts of the fermion fields necessitates the introduction of a vector field A_μ that couples to the charged fermions. Working through the consequences of the local U(1) gauge invariance allows one to identify the vector field with the massless photon, the gauge

boson of electromagnetism, and yields quantum electrodynamics (QED), the theory of relativistic charge particles interacting electromagnetically [3].

The full gauge structure of the modern-day Standard Model is said to be $SU(3)_{color} \otimes SU(2)_L \otimes U(1)_Y$ [4]. The $SU(2)_L \otimes U(1)_Y$ component of the Standard Model gauge symmetry generates a unified version of the weak and electromagnetic interactions, known as electroweak theory [5]. Under this theory, negative helicity or left-handed particles are grouped into SU(2) doublets of weak isospin, as shown in Table 1.3, while positive helicity or right-handed particles are left as singlets. In addition, the theory has a symmetry under U(1) hypercharge transformations, where hypercharge is defined as a linear combination of electric charge and the third component of the weak isospin, $Y \equiv I_3 + Q$. Taken by itself, the $SU(2)_L \otimes U(1)_Y$ gauge structure produces four massless gauge bosons, three for the $SU(2)_L$ component of the gauge symmetry and one for the $U(1)_Y$ part. This does not match with experimental observation of the electroweak interaction which involves three massive bosons (W^+,W^-, and Z^0) and one massless one (γ). In addition, the $SU(2)_L$ symmetry would be spoiled by mass terms for the quarks and leptons. However, both deficiencies in this theory may be remedied by the introduction of an SU(2) doublet of scalar fields, known as the Higgs field, that couple to both the massless $SU(2)_L \otimes U(1)_Y$ gauge bosons as well as the massless quarks and leptons [6]. If the appropriate potential is chosen for the Higgs field in the Standard Model Lagrangian, it obtains a non-zero vacuum expectation value that spontaneously breaks the SU(2) symmetry. As a consequence, the massless $SU(2)_L \otimes U(1)_Y$ gauge bosons obtain masses, although the linear combination of weak isospin and hypercharge corresponding to the electric charge remains a massless state. Thus, with the addition of the Higgs field, the familiar one massless electromagnetic boson and three massive weak bosons

	Left-Handed Helicity	**Right-Handed Helicity**
Leptons	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}, \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}, \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}$	ν_e, e, ν_μ, μ, ν_τ, τ
Quarks	$\begin{pmatrix} u \\ d \end{pmatrix}, \begin{pmatrix} c \\ s \end{pmatrix}, \begin{pmatrix} t \\ b \end{pmatrix}$	u, d, c, s, t, b

Table 1.3 The weak force groups fermions with left-handed helicity into SU(2) doublets of weak isospin, while right-handed particles from singlets under the weak SU(2) transformations.

emerge, as well as a new massive scalar Higgs boson. In addition, the terms generated by the coupling of the Higgs field to the quarks and leptons have the same structure as mass terms in the Lagrangian, allowing massive quarks and leptons to appear in the theory.

The $SU(2)_L \otimes U(1)_Y$ gauge symmetry of the electroweak theory, coupled with the spontaneous symmetry breaking caused by the Higgs field leads to a rich phenomenology. As a consequence of the non-Abelian nature of the $SU(2)_L$ symmetry, the electroweak gauge bosons couple to one another, unlike Abelian gauge theories, like QED, in which the photon has no self-coupling. In addition, because the weak interaction eigenstates are not identical to the mass eigenstates generated by the Higgs boson coupling, the weak interaction mixes the quarks and leptons from different generations, allowing interactions that covert quarks and leptons from one generation into quarks and leptons from a different generation. The matrix that describes this mixing is known as the Cabibbo-Kobayashi-Maskawa (CKM) matrix [7]. The CKM matrix can contain one complex phase and therefore allows for parity violation in weak interactions, leading to such interesting phenomena as mixing among neutral *K* and *B* mesons and parity violating decays of *K* and *B* mesons [8]. Many of the features of the electroweak sector of the Standard Model, with the notable exception of the Higgs boson, have been supported and confirmed by experimental results [9], making it one of the most successful theories in all physics.

The remaining $SU(3)_{color}$ gauge symmetry of the Standard Model is only relevant for quarks. Quarks are postulated to have an internal quantum degree of freedom labeled *color* consisting of three states: *red*, *blue*, and *green*. All other fermions in the Standard Model are assumed to be colorless. The quantum field theory that results from requiring the Standard Model Lagrangian to be invariant under $SU(3)_{color}$ transformations is known as quantum chromodynamics (QCD) [10]. In order to maintain invariance under local $SU(3)_{color}$ transformations, it is necessary to introduce eight bi-colored massless gauge bosons, known as gluons. Because of the non-Abelian characteristic of the SU(3) gauge symmetry of QCD, these gluons interact with one another and hence carry color charge. Just as in electroweak theory, this non-Abelian self-coupling leads to a richer and more complex phenomenology than is exhibited by Abelian gauge theories such as QED.

That gluons can interact with one another leads to an interesting consequence for the strength of the QCD coupling constant, α_S. In quantum field theories, the coupling "constant" that determines the strength of a force's interaction is not really constant. Rather, the strength of the coupling constant changes as the energy, and hence the distance scale, of the interaction changes.

For Abelian gauge fields like the electromagnetic field in QED, the coupling increases in strength for shorter distances (higher energies) and decreases at longer distances (lower energies). Qualitatively, this happens because virtual pairs of charged particles that are created and annihilated in the vacuum of the quantum field theory act like tiny dipoles, shielding the strength of charges at long distances, much like a charge is shielded in a dielectric material. A similar effect occurs in QCD except that in addition to virtual quark-antiquark pairs, there are also virtual gluon pairs. The effect of these gluon pairs is to cause the strength of the QCD coupling constant to increase for longer distances (lower energy) and decrease at shorter distances (higher energy). This behavior leads to two phenomenological features of QCD that make it quite different from other quantum field theories like QED: confinement and asymptotic freedom.

Confinement refers to the fact that quarks and gluons are never found as free particles. Another way to say this is that no free color charges exist. Instead, all quarks exist in either quark-antiquark bound states (known as mesons) or three-quark bound states, known as baryons. Collectively, meson and baryon bound states are referred to as hadrons. In mesons, the color of the quark is canceled out by the "anti-color" of the antiquark. For baryons, the equal mixture of colors from the three quarks adds up to give a "white" or "colorless" state. Another way to say this is that mesons and baryons must be in a color singlet state, much in the same way that two spin-1/2 particles can form a bound state of spin-0. This behavior is a consequence of the long distance (low energy) behavior of the QCD coupling constant. Consider a quark and an antiquark bound in a meson. Suppose we try to separate the quark from the antiquark. Because the strength of the strong force coupling grows with distance, the further we separate the quark from the antiquark, the greater the force of attraction between them becomes. To pull the quark free (in other words to separate it by an infinite distance from the antiquark) would require an infinite amount of energy. Before that point is reached, there will be enough energy stored in the color field to create a quark-antiquark pair from the vacuum. Thus, after separating the original quark from the antiquark, instead of having two free quarks, we are left with two mesons formed from the original quark and antiquark and the newly created quark-antiquark pair.

One consequence of confinement is that although QCD deals with interactions between quarks and gluons, experiments can only observe meson and baryon bound states. For example, in making calculations, a collision between a proton and an antiproton is treated as a QCD interaction between a quark or gluon from the proton and a quark or gluon from the antiproton.

The result of this QCD interaction is typically a set of two or more quarks, gluons, or a mixture of the two. However, in the lab, because of confinement, one does not observe quarks and gluons directly. Instead, as described above, the quarks and gluons hadronize and "dress" themselves into mesons and baryons. However, in a high-energy collision, with the quarks and gluons moving at high velocity, a large number of quark-antiquark pairs can be created during the hadronization process. Thus, the final result of a $p\bar{p}$ collision is a large collection of mesons and baryons, grouped into collimated collections called jets. Usually there is one jet of hadrons per quark or gluon created in the underlying QCD process and their angular distribution reflects that of the underlying quarks and gluons in the QCD interaction. The process of a single quark or gluon generating a jet of particles in a high-energy collision is referred to as fragmentation.

On the other hand, *asymptotic freedom* refers to the behavior of QCD in the opposite regime of short distances (high energies). Just as the QCD coupling constant grows strong when long distances are involved, it becomes weaker at short distances. At short enough distance scales, the quarks behave mostly like free particles, with the interactions of QCD manifesting themselves as small effects. In this regime, QCD can be treated with perturbation theory, as described below. Experimentally, the region of asymptotic freedom in QCD can be explored by observing the results of collisions between hadrons—for example, protons colliding with antiprotons—at high energies. Such collisions can be viewed as high momentum transfer interactions between the quarks and gluons contained in the proton and antiproton. Of course, because of confinement, the end result of such a collision must be a collection of hadrons. The theoretical approach to linking quark-level perturbative QCD calculations to experimentally observed hadron initial and final states is dealt with in greater detail in Section 2.1.

1.2 Perturbation Theory

The Lagrangian describing the Standard Model, in principle, provides all the information about the dynamics of quarks and leptons interacting through the strong, electromagnetic, and weak forces necessary to calculate such experimental observables as scattering cross sections and decay rates. However, in practice, making such calculations is less than straightforward. One of the most powerful and commonly used tools for extracting quantitative results from the Standard Model is known as *perturbation theory*. Perturbation theory is a standard calculational technique in which the answer to a complicated problem is approximated from the known answer to a

similar, simpler problem by treating the complicated parts of the problem as a small perturbation on the exact solution. Perturbation theory is applied to the Standard Model by treating the gauge interactions among the particles as a small perturbation on the free non-interacting theory. High-energy physics has honed the technique of applying perturbation theory to quantum field theories to a fine art through the use of a tool known as Feynman diagrams.

Feynman diagrams are pictorial representations of various terms in the perturbation series used to make calculations of the quantum mechanical amplitudes for processes of interest. The basic elements of a Feynman diagram include lines to indicate the space-time trajectories of fermions and bosons involved in the process, and vertices at which three or more lines intersect to represent interactions included in the quantum field theory. Standard techniques exist for determining what kinds of lines and vertices are required to describe the physics contained in a quantum field theory Lagrangian. Once the lines and vertices have been specified, prescriptions known as Feynman rules can be developed that allow one to translate the diagram into a term in the perturbative expansion of the theory. In general, the order of the term in the perturbative expansion can be determined by counting the number vertices present in the diagram: each vertex contributes a factor of the coupling constant, α, for the interaction, which is assumed to be small if perturbation theory is to be a valid approach. So, a Feynman diagram containing two vertices represents a second-order or $O(\alpha^2)$ term in the perturbative series. The strength of this approach is that the dynamics of a particular field theory can be codified as a general set of Feynman rule based on the Lagrangian of the theory once, and from then on the Feynman rules and Feynman diagrams can be applied to solve a wide variety of problems. The Feynman rules associated with the Standard Model for QED, electroweak theory, and QCD have been known for quite some time now. Most modern computations of Standard Model processes begin with writing down all allowed Feynman diagrams that link the initial state of the problem to the final state, to the order in the perturbation series necessary to achieve the desired accuracy. Once the contributions to the amplitude from the various diagrams have been calculated, a cross section or decay rate can be obtained from the square of the amplitude integrated over the relevant phase space with the necessary phase space and flux factors included.

The application of perturbation theory to electroweak theory has been very successful in producing quantitative results. This is a natural consequence of the small coupling constants of the electromagnetic and weak interactions, which allow good convergence of their perturbative series. The success of perturbation theory in electroweak physics can be seen in the stunning

accuracy of such calculations as the anomalous magnetic moment of the electron [11]. However, the situation for QCD is very different. Because of the features of QCD described in Section 1.1, namely confinement and the running QCD coupling constant, the coupling constant of QCD is not small enough in low-energy interactions to warrant a perturbative treatment. However, because of asymptotic freedom, at high energy scales the coupling constant of QCD becomes small and perturbation theory becomes a viable option. Although techniques have been developed to deal with the low-energy, strong coupling limit of QCD—for example, lattice QCD calculations [12]—the most developed tool in the high-energy physics remains perturbation theory, making it interesting to study QCD in the regime where perturbative approximations apply.

1.3 Heavy Flavor as a Test of Perturbative QCD

The energy regime in which perturbative techniques are applicable to QCD is described by the parameter Λ_{QCD}, which is defined to be the energy scale at which the QCD coupling constant α_s becomes of order unity. Λ_{QCD} can be measured experimentally and has a value of approximately 200 GeV [9]. Processes occurring at an energy scale significantly above Λ_{QCD} can be handled by perturbative means. One useful example of such a process is the pair production of heavy quarks through QCD interactions at hadron colliders, like the Tevatron (see Section 3.1 for more details on the Tevatron collider). Three quarks—charm, bottom, and top—are sufficiently massive that QCD interactions that pair-produce them are guaranteed to be in the perturbatively calculable regime. On the other hand, lighter quarks can be produced at scales where the perturbative approach fails, making perturbative calculations of light quark production more challenging.

Charm quarks, with a mass near 1.5 GeV/c^2 [9], are the most easily produced heavy flavor quarks at the Tevatron. However, there are a few theoretical and experimental difficulties that prevent them from being ideal tests of QCD in the perturbative regime. On the theoretical side, the charm quark mass, although greater than Λ_{QCD}, is also small enough that non-perturbative corrections could play an important role. In addition, charm quark signatures can be difficult to isolate from the vast background of light quark production. However, as discussed in Chapter 8, new experimental techniques developed for Run II at the Tevatron make QCD studies involving charm quarks far more feasible.

Top quarks, on the other hand, do not suffer from any of the drawbacks that plague charm quark measurements. The top quark mass, approximately 175 GeV/c^2 [9], is so large that non-perturbative effects may be neglected to a good approximation when calculating its production. In addition, the top quark decays so quickly that it does not hadronize, meaning theoretical uncertainties associated with that non-perturbative process are also completely absent. In addition, the large mass and distinctive decays of top quarks make it relatively easy to isolate a reasonably pure sample of top quark events. However, the top quark mass also makes it difficult to produce and detect top quarks in sufficient quantity to perform high-statistics tests of perturbative QCD calculations. Nonetheless, present indications are that, within statistical precision, the measured top quark cross section agrees well with perturbative QCD calculations [9].

Intermediate in mass between charm and top, the bottom quark, with a mass of approximately 4.5 GeV/c^2 [9], is a good compromise for making high statistics checks of perturbative QCD calculations. The *b* quark mass is sufficiently higher than Λ_{QCD} that non-perturbative effects should not be a serious consideration. At hadron colliders like the Tevatron, *b* quarks are primarily pair-produced through QCD processes like $q\bar{q} \to b\bar{b}$ or $gg \to b\bar{b}$. It is also possible to create *b* quarks either in pairs or singly through electroweak processes; however, the rates for such production are negligibly small compared to QCD production due to the much large coupling constant of QCD. Given reasonable assumptions about the size of the *b* quark production cross section, one expects on the order of hundreds of *b* quarks to be produced every second at a collider like the Tevatron. In addition, the weak decays of bottom hadrons (denoted *B* hadrons) have two key features that can be exploited to define efficient experimental signatures of *b* quark production. One feature of *B* decays useful for defining a *b* quark data sample is the relatively high semileptonic branching fraction. Roughly 20% of *B* decays involve an electron or muon [9]. In detectors like CDF, high momentum electrons and muons provide clean signals around which efficient triggers can be defined. The second feature of *B* decays useful for defining clean *b* quark samples is the long *B* lifetime. *B* hadron decays, much like charm hadron decays, proceed through the emission of a virtual *W* boson. Because the *W* is so much more massive than *b* or *c* quarks, the lifetime of this decay is fairly long, as opposed to the weak decay of the top quark, which is massive enough that it can decay through the emission of a real *W* boson with essential zero lifetime. The *B* lifetime is further enhanced over charm lifetimes because the dominant CKM matrix element relevant for the flavor-changing, weak

decay of a B hadron, $V_{cb} \approx 0.04$, while the CKM matrix element that dominates the weak decay of charm, $V_{cs} \approx 0.97$ [9]. Therefore, B hadrons decay with a lifetime of approximately 1.5 ps [9] and can travel on the order of millimeters in the detector with the relativistic boosts typical of detected B hadrons. This distance is sufficiently long that the precise tracking detectors employed at CDF are capable of resolving the separation of the decay point of the B hadron from the point at which the $p\bar{p}$ interaction occurred. In addition, the presence of a long-lived B hadron can be detected by reconstructing particle trajectories, or tracks, from the B decay that have large impact parameters with respect to the $p\bar{p}$ interaction vertex. Using these two standard techniques, it is possible to identify large, relatively low-background $b\bar{b}$ samples at hadron colliders like the Tevatron.

Given the advantages of testing perturbative QCD calculations using measurements of $b\bar{b}$ production at the Tevatron, it came as quite a surprise to the high-energy physics community when measurements of the b quark cross section from both the CDF and DØ detectors at the Tevatron were consistently higher than next-to-leading-order (NLO) QCD perturbative calculations. Figure 1.1, taken from ref. [13], summarizes the measurements of b quark production from CDF and DØ at the Tevatron, as well as measurements from the UA1 detector at the lower energy CERN $Sp\bar{p}S$ collider. These figures show that although the shape of the cross section distributions as a function of p_T is reasonably described by the theoretical calculation, the overall normalization of the data exceeds the theoretical calculations by a factor of two to three. Although the experimental measurements can be accommodated by somewhat extreme choices of the theoretical parameters, this discrepancy has lead many to search for modifications to the theory to allow a better description of the data.

One approach for improving the agreement between the theoretical calculation and the data is to improve the theoretical description of b fragmentation. Fragmentation refers to the processes by which bare quarks produced in the QCD interaction evolve into collimated collections—or jets—of hadrons (see Section 2.4.2 for more detail). Fragmentation is a non-perturbative process that must be described using phenomenological models. The most commonly used model for b quark production is due to Peterson *et al*. [14]; however, recent experimental evidence from e^+e^- colliders suggests that the Peterson model of fragmentation is not the best model available [15]. Cacciari and Nason have explored the effect of including an improved model of B fragmentation in comparing the theoretical calculation of the B^+ meson cross section to a particular measurement at CDF [16]. In addition to using an improved parameterization of the b

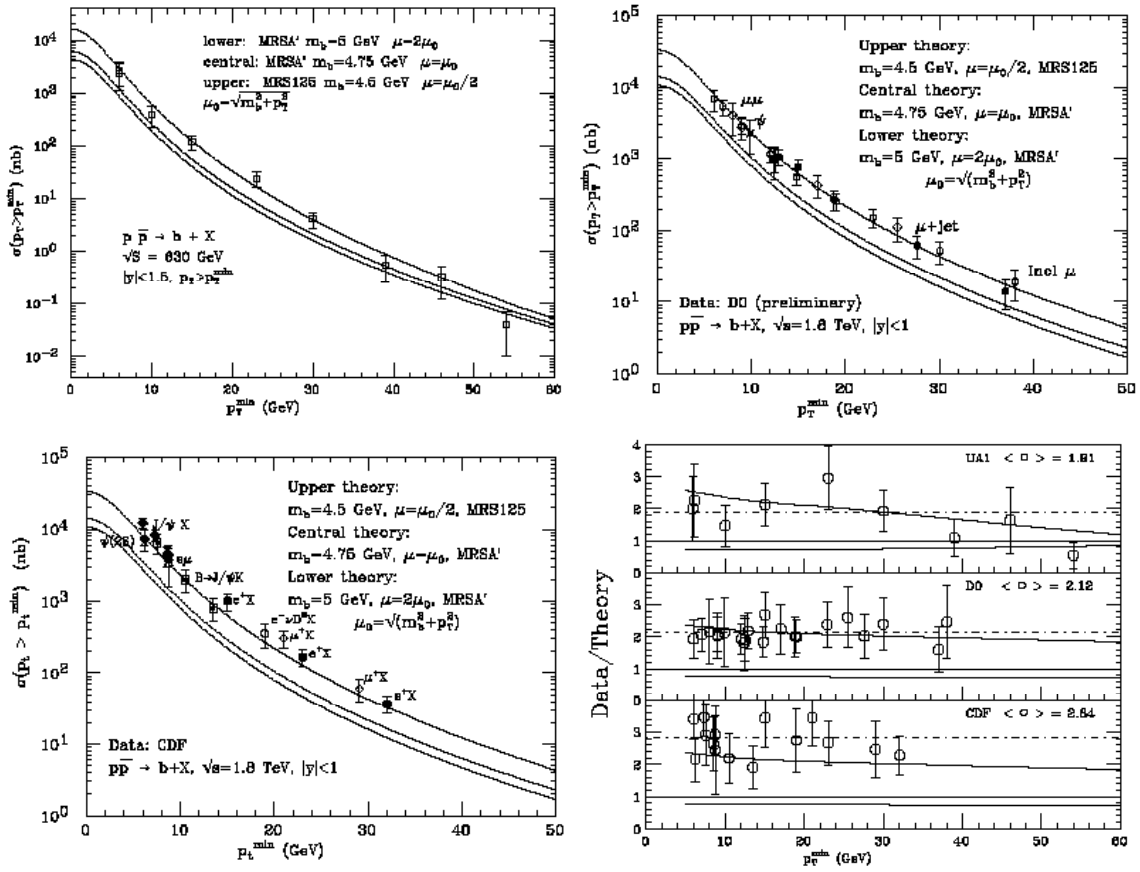

Figure 1.1 A comparison of measurements of the *b* quark production cross section at hadron colliders to the NLO prediction. In each plot, the *x*-axis shows the minimum *b* p_T to which the measurement was sensitive. The upper left plot shows the results from UA1. The upper right plot shows the DØ result. The lower left plot shows the CDF results. Finally, the lower right plot displays the ratio of the measured cross section to the prediction from NLO QCD. These plots are taken from [13].

fragmentation, they also account for logarithmic enhancements to the cross section coming from multiple soft gluon emissions. Their resulting calculation, which they call fixed-order, next-to-leading-log (FONLL), reduces the discrepancy between the data and theory from a factor of 3 to 1.7. Based on this improvement, further explorations of *B* fragmentation seem a promising route for better understanding the discrepancy in *b* production.

An alternative explanation for the measured *b* excess over theoretical predictions involves the production of supersymmetric particles. Supersymmetry is an extension to the Standard Model

that pairs each Standard Model fermion and boson with a partner particle differing by one-half unit of spin [17]. Berger has suggested that the production of a light gluino, which always decays to a bottom quark and a sbottom squark, could explain the observed *b* quark excess [18]. A gluino is the spin-1/2 supersymmetric partner of the Standard Model gluon while the sbottom squark is the spin-0 supersymmetric partner of the bottom quark. In this model, gluinos, with a mass between 12 GeV/c^2 and 16 GeV/c^2 are pair produced at the Tevatron and each decays with 100% branching fraction to a bottom quark and a sbottom squark. The sbottom squark is either stable and exits the detector without interacting or decays to a pair of hadronic jets, making it difficult to differentiate this production mechanism for *b* quarks from the conventional QCD one. If the gluino production cross section were sufficiently large, the bottom quarks from the gluino decay could explain the observed *b* excess at the Tevatron. Berger suggests that the definitive signature for this supersymmetric production mechanism would be an observed excess in the number of events containing either *bb* or $\bar{b}\bar{b}$ pairs. Such an excess could be detected by explicitly reconstructing specific final states, for example B^+B^+, or by finding an excess in like-sign leptons from semileptonic *B* decays over what is predicted from Standard Model *B* mixing. Conclusive evidence from Run 1 at the Tevatron is lacking to support this hypothesis; however, it is hoped that the larger statistics anticipated in Run II will either confirm or rule out this scenario.

These two proposed resolutions to the *b* quark production mystery represent opposite ends of the spectrum of resolutions. The first solution involving improved models of fragmentation suggests that the discrepancy comes simply from our inability to model certain non-perturbative effects that are part of the standard picture of *b* production. On the other hand, the second explanation involves a fundamental modification to the Standard Model. This broad range of

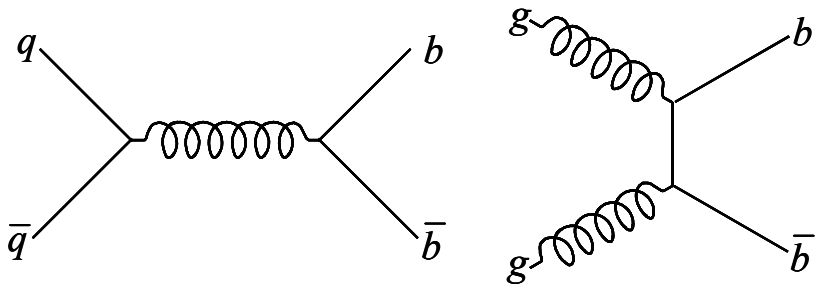

Figure 1.2 Leading-order Feynman diagrams for $b\bar{b}$ production.

possible explanations for the *b* production mystery is what fuels much of the interest in this subject.

Another possibility for understanding the *b* quark excess involves higher-order terms in the perturbative calculation of *b* quark production that are needed to describe the data accurately. On the surface, such a suggestion seems unlikely because in a typical perturbative series, if perturbation theory can be applied successfully, the leading terms will closely approximate the exact answer, and additional terms provide incremental improvements. However, despite the claim that bottom quark production can be calculated using perturbative QCD, it is not the case that the leading-order terms accurately represent the data and that higher-orders simply provide small corrections. In fact, the leading-order calculation of the *b* cross section lies another factor of two to three below the NLO prediction, which is itself a factor of two to three below the measured cross section [19]. This unexpected feature of the perturbative *b* production calculation can be explained by realizing that starting at next-to-leading order, two new mechanisms for $b\bar{b}$ pairs produced in association with a light quark or gluon appear. At leading order, the only QCD production mechanism for $b\bar{b}$ pairs includes processes like quark-antiquark annihilation and gluon-gluon fusion that yield a single $b\bar{b}$ pair with no other quarks or gluons in the final state, as shown in Figure 1.2. However, at higher-orders, it is possible to produce a $b\bar{b}$ pair in addition to a light quark or gluon by adding a $g \to b\bar{b}$ splitting to a lower-order QCD

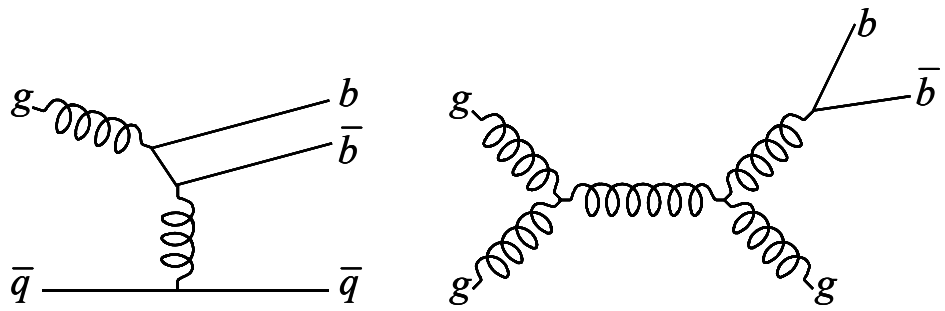

Figure 1.3 Higher-order $b\bar{b}$ created by adding a $g \to b\bar{b}$ splitting to a leading-order process containing only light quarks and gluons. The new production mechanisms are classified as to whether the $g \to b\bar{b}$ splitting occurs in the initial state, with one of the *b* quarks from the splitting scattering against a quark or gluon from the other beam (flavor excitation), or whether the $g \to b\bar{b}$ occurs in the final state (gluon splitting).

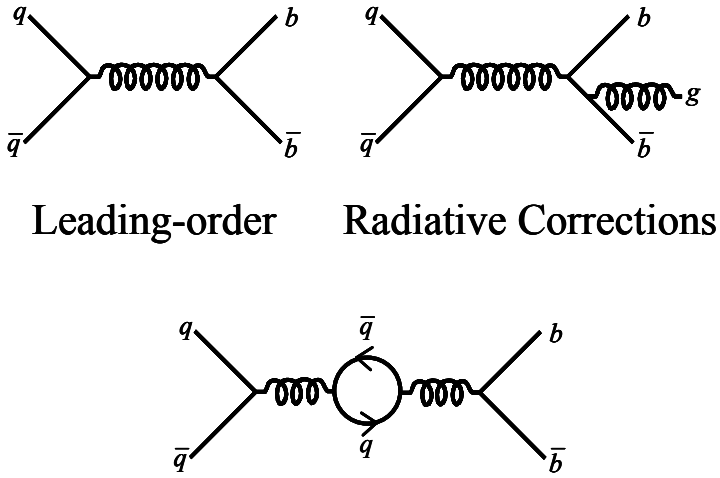

Figure 1.4 Diagrams contributing to flavor creation.

process involving only light quarks and gluons, as shown in Figure 1.3. If the gluon splits to a virtual $b\bar{b}$ pair in the initial state and one of the *b* quarks is scattered into the final state by interaction with a quark or gluon from the other beam, the process is called *flavor excitation*. On the other hand, if the $g \to b\bar{b}$ splitting occurs in the final state, the process is called *gluon splitting*. In addition, the term *flavor creation* is used to refer to the lowest-order $b\bar{b}$ production process, as well as higher-order corrections to the leading-order diagrams, as shown in Figure 1.4. It turns out that the higher-order flavor excitation and gluon splitting processes can contribute at approximately the same level as the leading-order diagrams. One way to understand this is to realize that the cross section for $gg \to gg$ is of order one hundred times larger than the cross section for $gg \to b\bar{b}$. Therefore, even after accounting for the extra $O(\alpha_s)$ suppression of flavor excitation and gluon splitting over the leading-order term, these higher-order production mechanisms are still numerically significant.

The perturbative calculations referenced above showing a factor of two to three discrepancy with experiment are performed with next-to-leading-order (NLO) accuracy because this is the lowest order that yields sensible results for *b* production. However it may be possible that higher orders make non-negligible corrections to the theoretical predictions that might help reduce the discrepancy. So far, the very difficult next-to-next-to-leading-order (NNLO) calculation of *b*

production has not been successfully performed. However, there is an approximation, known as the parton shower model, that is not exact to any particular order in perturbation theory, but captures the qualitative features of contributions from all orders. It is demonstrated in ref. [20] that calculations using the parton shower approach are able to achieve reasonable agreement with measurements of the b quark cross section, suggesting that perhaps corrections from orders beyond NLO are significant. On the other hand, the parton shower approach neglects certain interferences between the various $b\bar{b}$ production mechanisms (flavor creation, flavor excitation, and gluon splitting), making it possible that the improved agreement between the parton shower calculation and the data is merely the result of overcounting. For this reason, it is interesting to further compare data and parton shower predictions to ascertain whether the contributions from higher-order processes predicted by the parton shower model seem supported by the data.

1.4 Probing Higher-Orders with Correlations

One powerful probe of the contributions from higher-order $b\bar{b}$ production mechanisms is correlations. A $b\bar{b}$ correlation is quantity that combines information from both the b and $\bar{b}$ in the event, for example, the opening angle between the b and $\bar{b}$ trajectories or the asymmetry in the b quark transverse momenta, p_T, defined as

$$A = \frac{p_T^b - p_T^{\bar{b}}}{p_T^b + p_T^{\bar{b}}} \tag{1.1}$$

In general, p_T refers to the component of a particle's momentum perpendicular to the colliding particle beam direction. At leading-order, which is basically the lowest-order contributions to flavor creation, the only particles in the final state of $b\bar{b}$ production are the b and $\bar{b}$ quarks, so momentum conservation demands that the b and $\bar{b}$ quarks be produced back-to-back with equal p_T. At higher-orders, final states containing additional light quarks and gluons allow different correlations. Higher-order corrections to flavor creation in the form additional gluon radiation leads to a more smeared out opening angle distribution, although the opening angle still peaks at 180° and the p_T remains reasonably symmetric. Flavor excitation and gluon splitting give much more even opening angle distributions and are more asymmetric in b p_T.

Consider, for example, the transverse opening angle, $\Delta\phi$, which is defined as the angle between the b and $\bar{b}$ quarks in the plane perpendicular to the proton and antiproton beam axis.

Figure 1.5 gives the prediction from one particular parton shower implementation to the Δϕ distribution from each production mechanism. The flavor creation contribution is evident as the main source of the back-to-back peak in the overall distribution, while flavor excitation and gluon splitting fill in the distribution at smaller angles. The exact shape of the contributions from each production mechanism depends on the kinematic cuts applied to the data. For example, raising the p_T cut on the b and $\bar{b}$ quarks in Figure 1.5 would cause a peak in the gluon splitting contribution for low Δϕ. By measuring distributions like Δϕ in data, it is possible to gauge the importance of higher-order production mechanisms. In addition, by comparing the measurement specifically to predictions from the parton shower model, we can attempt to understand whether the large contributions from flavor excitation and gluon splitting predicted by that model are appropriate.

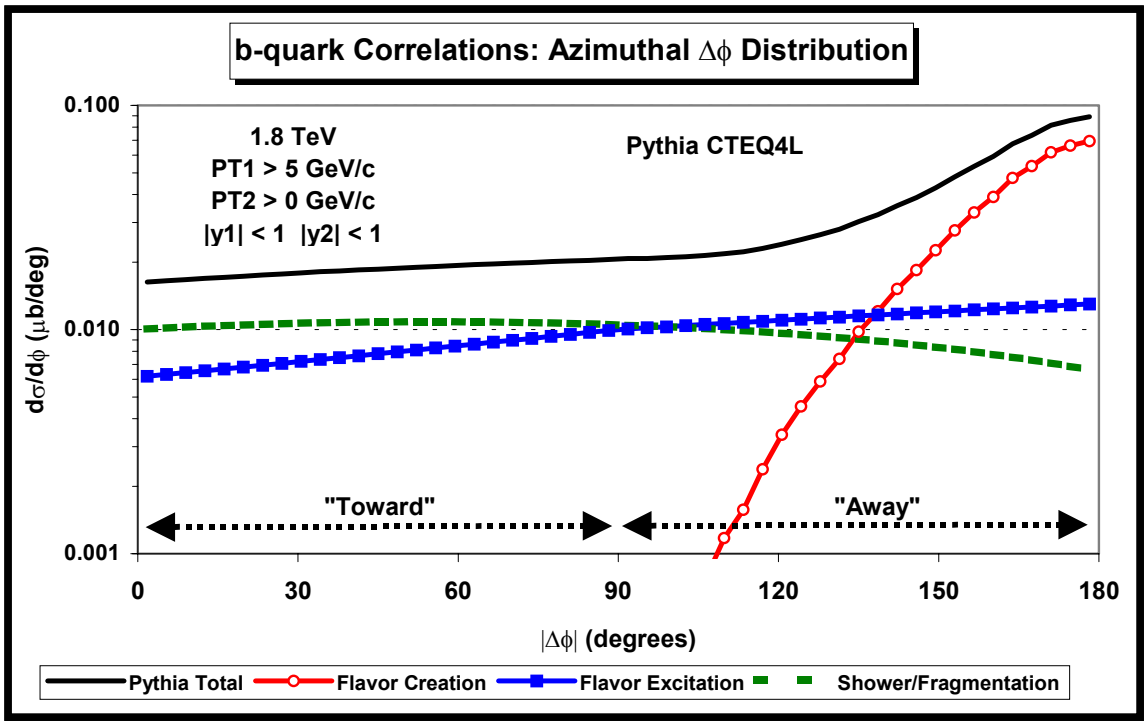

Figure 1.5 The Δϕ distribution predicted by the PYTHIA parton shower Monte Carlo for each of the three $b\bar{b}$ production mechanisms, as well as the total. This plot was taken from [20].

1.5 Previous Measurements of $b\bar{b}$ Correlations

A number of previous results have been published on $b\bar{b}$ correlations at hadron colliders. Each of these analyses take advantage of one of the two standard techniques for detecting or "tagging" the presence of *b* quarks in the events: namely using the semileptonic decay of a *B* hadron involving a lepton (in these cases a muon) or using the long-lifetime of *B* hadrons to detect the *B* decay point by reconstructing tracks with large impact parameters or explicitly resolving the secondary *B* decay vertex from the primary $p\bar{p}$ interaction vertex.

The first published measurements on $b\bar{b}$ correlations at a $p\bar{p}$ collider come from the UA1 collaboration using data taken at the CERN $Sp\bar{p}S$ collider with $\sqrt{s}$ = 630 GeV [21]. This analysis uses dimuon events from 4.7 pb^{-1} of data collected at the $Sp\bar{p}S$ from 1988 to 1989. Both muons in the events are required to have a $p_T > 3$ GeV/c and $|\eta| < 2.3$.[1] A cut is placed on the invariant mass of the dimuon pair of 6 GeV/c^2 < $m_{\mu\mu}$ < 35 GeV/c^2. The lower limit of this cut is intended to reduce the contribution from "sequential" *B* decays of the type $B \to D + l$, $D \to X + l$, as well as decays of the charmonium states, *J*/ψ and ψ′, to dimuons. However, it also limits the sensitivity to dimuons produced by $b\bar{b}$ pairs with small opening angle. The upper limit of the cut excludes dimuon pairs from Z^0 decays. From these events, a sample of approximately 1000 $b\bar{b}$ events are extracted and compared to both the order α_s^3, next-to-leading order (NLO) calculation from Mangano, Nason, and Ridolfi (MNR) [22] as well as the ISAJET parton-shower Monte Carlo program [23]. The cross section and $b\bar{b}$ angular distributions are found to be in reasonable agreement with both theoretical models. The contribution to $b\bar{b}$ production from higher-order processes—meaning processes not described by the leading-order perturbative calculation—is found to be about 30-40%.

Similar $b\bar{b}$ correlation measurements using dimuon events were also performed by the CDF [24] and DØ [25] collaborations at the Tevatron with $\sqrt{s}$ = 1.8 TeV. The CDF measurement uses dimuon pairs where the muons have $p_T > 3$ GeV/c and $|\eta| < 0.6$. A dimuon pair mass cut of $m_{\mu\mu} > 5$ GeV/c^2 is applied to reduce the contribution from sequential muon pairs and charmonium decays. From 17.4 pb^{-1} of data, a sample of approximately 2500 $b\bar{b}$ dimuon events is extracted. The DØ measurement uses dimuon events where the muons have $p_T > 4$ GeV/c and

[1] The symbol η refers to the quantity *pseudorapidity* which is related to the relativistic *rapidity* of a particle, as described in 3.2.

$|\eta| < 0.8$. In addition, each muon has to be located within $\Delta R = 0.8$ of a jet with $E_T > 12$ GeV.[2] The invariant mass of the dimuon pair is required to satisfy 6 GeV/c^2 < $m_{\mu\mu}$ < 35 GeV/c^2, just like the UA1 measurement, to reduce the contributions from sequential, charmonium, and Z^0 decays. In 6.5 pb^{-1} of data, a sample of approximately 200 dimuon events from $b\bar{b}$ production is collected. Both the CDF and DØ measurements of the $b\bar{b}$ angular correlations made using these analyses are found to be consistent in shape with the NLO QCD predictions, although it is worth noting that both analyses suffer from the same lack of resolution at small opening angles present in the UA1 analysis. In addition, both analyses found the measured $b\bar{b}$ cross sections to be higher than the NLO predictions by a factor of two to three.

In addition to the dimuon analysis, CDF also performed an analysis of $b\bar{b}$ angular correlations using events containing a muon plus a jet [26]. For this analysis, the muon is required to have $p_T > 9$ GeV/c and $|\eta| < 0.6$, while the jet must have $E_T > 10$ GeV/c and $|\eta| < 1.5$. In addition, the muon and the jet are required to be separated in η–ϕ space by at least ΔR, leading to a poor acceptance for events in which the b and $\bar{b}$ quarks are separated by a small opening angle. For each event, a jet probability algorithm [27] is used to determine the likelihood that the jet contains tracks from the decay of a long-lived particle. This jet probability is determined by looking at the impact parameter of charged particle tracks reconstructed within the jet. The contribution from $b\bar{b}$ events is determined by fitting the jet probability distribution for all events to templates from $b\bar{b}$ and $c\bar{c}$ Monte Carlo, plus an exponential function to represent jets containing no long-lived decays. The number of $b\bar{b}$ events extracted from the fit is approximately 2500. The $\Delta\phi$ distribution between the muon and the jets for these events is compared with the NLO QCD prediction. Although the shape of the measured $\Delta\phi$ distribution is similar to the prediction, the normalization of the data is larger than the prediction by almost a factor of two and the shape of the data shows some small discrepancies compared to the NLO prediction.

Finally, CDF has performed a measurement of $b\bar{b}$ rapidity (y) correlations using events that contain a muon associated with a jet in either the central or forward region of the detector, plus a jet tagged with a secondary vertex in the central region of the detector [28]. This measurement

---

[2] The quantity $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ defines an angular separation between two particles in terms of their separation in pseudorapidity and azimuthal angle. E_T is defined as the energy of a particle weighted by the sine of the polar angle between the particle and the colliding beam direction, $E_T = E\sin\theta$.

represents the ratio of events containing a b quark with $2.0 < |y| < 2.6$ ("forward" region) to events containing a b quark with $|y| < 0.6$ ("central" region), given that each event contains a second b quark with $|y| < 1.5$. The b quark in the central or forward region is identified by reconstructing a muon with $p_T > 6$ GeV/c associated with a jet with $E_T > 15$ GeV within $\Delta R = 0.7$ of the muon in the appropriate rapidity range. The second b quark in the event is detected as a jet with $E_T > 26$ GeV containing a secondary vertex reconstructed with the SECVTX algorithm [29]. The SECVTX algorithm, originally developed for the CDF top quark analysis, tags the decay of a long-lived B hadron by reconstructing the decay vertex using precision tracking information. In addition, a separation in transverse opening angle between the μ-jet and the SECVTX jet of $\Delta\phi > 60°$ was required to exclude a region of phase space having poor acceptance for reconstructing and tagging both b quark jets. This measurement uses 77 pb^{-1} of data, and after all cuts approximately 300 $b\bar{b}$ events were reconstructed in the forward region and about 5000 events in the central region. The ratio of forward to central rapidity $b\bar{b}$ production was found to be consistent within errors with the NLO QCD prediction.

Figure 1.6 summarizes the $b\bar{b}$ correlation results mentioned above form the Tevatron. Each of the above measurements indicates a general, at least qualitative agreement with theoretical predictions of shape of the $b\bar{b}$ distributions, although most show a disagreement in the normalization of the $b\bar{b}$ cross section. However, it should also be noted, that each of the above measurements includes either an invariant mass cut placed on the muon pairs or explicit angular separation requirements between the muon and the jet in the event. These requirements are necessary in order to reduce backgrounds coming from processes that give primarily small opening angle muon or muon-jet pairs. However, these requirements also reduce the sensitivity of the above measurements to $b\bar{b}$ pairs produced at small opening angles. The small opening angle region of $b\bar{b}$ production is interesting because the primary contribution in this region comes from higher-order production mechanisms. Therefore, it is worthwhile to attempt a measurement that does not exclude this region.

1.6 Analysis Overview

The main motivation of this analysis is to explore the region of $b\bar{b}$ production in which the b and $\bar{b}$ quarks have a small opening angle. To this end, we begin by collecting a large sample of B decays from 90 pb^{-1} of data taken at CDF during the 1994-1995 Tevatron run. To create a

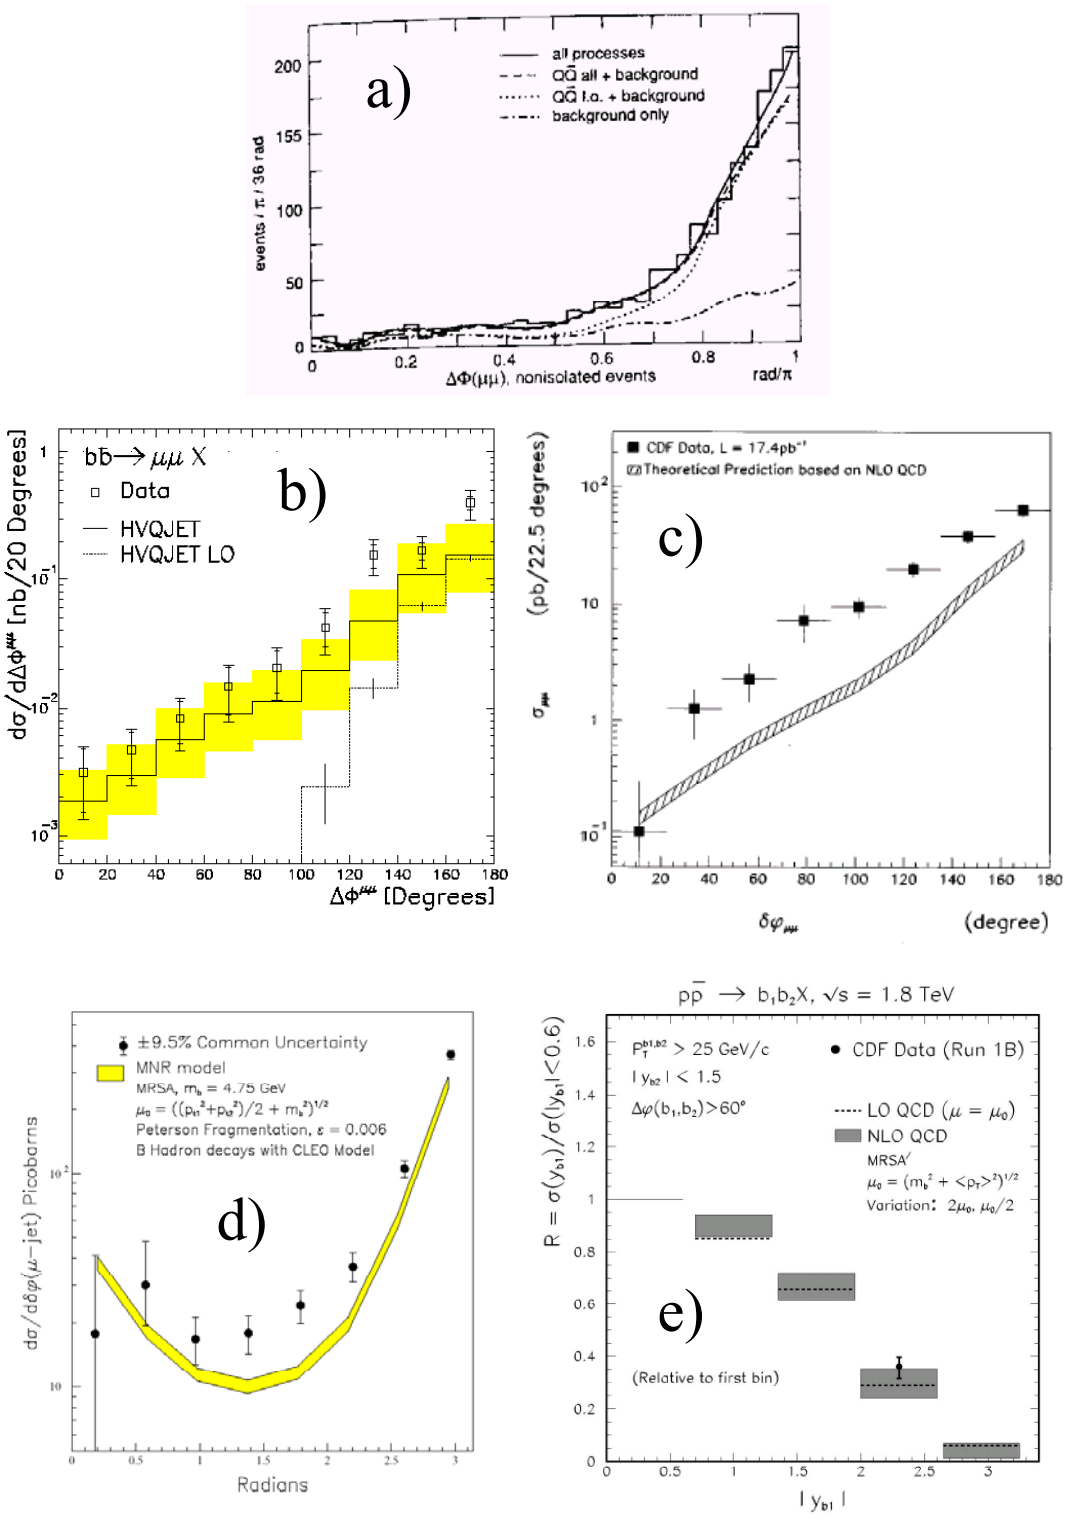

Figure 1.6 Plots from previous $b\bar{b}$ measurements. The plots correspond to the following measurements: (a) UA1 dimuon $\Delta\phi$ measurment, (b) DØ dimuon $\Delta\phi$ measurement, (c) CDF dimuon $\Delta\phi$ measurement, (d) CDF μ-jet $\Delta\phi$ measurement, (e) CDF μ + *b*-jet rapidity correlation measurement.

sample enhanced in *b* quark content, we demand that the events contain a lepton, either electron or muon, with $p_T > 8$ GeV/c. In addition, we use a vertex tagging algorithm similar to the SECVTX algorithm mentioned above, to tag the displaced decay vertices of both *B* hadrons in the event. The *b*-tagging algorithm used for this analysis has been specially optimized from previous versions to enable tagging of *B* hadron pairs with arbitrarily small opening angle. After background removal, we obtain a sample of approximately 17,000 double-tagged $b\bar{b}$ events, the largest sample ever collected at a hadron collider.

Using the sample above, we measure the $\Delta\phi$ distribution between the reconstructed secondary vertex tag pairs and predictions based on PYTHIA and HERWIG Monte Carlo programs. In addition, we use information from the Monte Carlo samples to correct the tag-pair $\Delta\phi$ for detector effects and backgrounds to allow an extraction of the *B* hadron $\Delta\phi$ distribution that can be directly compared to theoretical predictions from other models. We choose to measure the *B* hadron $\Delta\phi$ distribution rather than the $b\bar{b}$ distribution because our secondary vertex tags are more closely related to *B* hadrons than *b* quarks. Converting our measurement from the *B* hadron level to the *b* quark level would introduce a dependence on *b* quark fragmentation models that we wish to avoid.

Chapter 2

Theoretical Background

We would like to use measurements of $b\bar{b}$ production to make an experimental test of QCD. Because the energy required to generate a pair of massive *b* quarks is much greater than Λ_{QCD}, we expect perturbation theory to be able to make valid predictions regarding our measurements. However, perturbative QCD deals with interactions involving quarks and gluons. In the laboratory, we create *b* quarks in collisions between protons with antiprotons, not quarks and gluons. In addition, because of confinement, we don't directly observe the *b* quarks produced, but rather we see *B* hadrons. Bound states of quarks like protons and B^+ mesons involve energy scales far too low to allow perturbative calculations. In order to apply perturbative methods to the calculations of *b* quarks in the laboratory, it is necessary to relate the high energy interactions between quarks and gluons that create the *b* quarks to the relatively low energy physics of the hadrons available for use as experimental tools.

2.1 Factorization

The scheme for connecting high-energy QCD interactions between quarks and gluons to the hadronic initial and final states accessible in the laboratory is known as factorization [30]. The main idea behind factorization is that the low energy interactions at play in the hadron initial and final states do not have a significant effect on the hard QCD interaction that produces the *b* quarks. The hard interaction can be treated with perturbative methods and then the effects of the low energy physics can be added on after the fact using a phenomenological model. Furthermore, the soft physics is assumed to be universal in nature—in other words, not specific to the high-energy interaction under study. Thus, the soft physics models can be derived from other, experimental environments, like e^+e^- colliders, and then applied to $b\bar{b}$ production at the Tevatron.

The basic scheme for the calculation of $b\bar{b}$ production is as follows: We begin by assuming that the proton and antiproton are described as a collection of quarks, antiquarks and gluons, known collectively as partons. The hard QCD interaction that produces the $b\bar{b}$ pair is seen to occur not between the antiproton and proton as a whole, but only between one parton from each. Perturbative methods can be applied to calculate the cross section for producing *b* quarks in QCD interactions between these partons. Then the initial-state partons can be related to the initial beam particles through a set of parton distribution functions that describe the likelihood of finding a given parton with a certain fraction of the beam particle's momentum. See section 2.4.1 for more information on parton distribution functions. The final state partons are related to the *B* hadrons through a set of *B* fragmentation functions that provide a model for the soft physics involved in *b* quark hadronization. Section 2.4.2 presents more information on *b* quark fragmentation.

Algebraically, this scheme may be written as follows:

$$\sigma(AB \to CD) = \sum_{a,b} f_a^A(x_a) f_b^B(x_b) \otimes \hat{\sigma}(ab \to cd) \otimes F_c^C(z_c) F_d^D(z_d), \qquad (2.1)$$

where $\hat{\sigma}(ab \to cd)$ represents the parton level cross section for partons *a* and *b* to produce partons *c* and *d*. The functions $f_i^j(x_i)$ represent the parton distribution function for finding parton *i* in beam particle *j* with the fraction x_i of the beam particle's momentum. On the other hand, $F_k^l(z_k)$ represent the fragmentation functions, which parameterize the likelihood of producing a hadron *l* with fraction z_k of the momentum of the final state parton *k*. The sum is over all parton species present in the beam particles and the "⊗" denotes convolution integrals of the parton distribution and fragmentation functions with the parton level cross section.

In the next section we will discuss the various theoretical techniques for obtaining the parton level cross section for $b\bar{b}$ production from perturbative QCD. Then we will address the issue of the non-perturbative aspects represented by the parton distribution and fragmentation functions.

2.2 Leading-Order Perturbative Calculation

It is the job of perturbation theory to calculate the parton level cross section $\hat{\sigma}(ab \to cd)$. More specifically, we are usually interested not in the overall cross section for producing partons *c* and *d*, but in the cross section for producing these partons within a certain momentum range.

Therefore, what is usually of greatest interest is the differential parton level cross section $d\hat{\sigma}(ab \to cd)$ given by

$$d\sigma = \frac{1}{2E_a 2E_b |v_a - v_b|} |\mathcal{M}(ab \to cd)|^2 \, d\Pi \,, \tag{2.2}$$

where $\mathcal{M}$ is related to the quantum mechanical amplitude for the interaction between a and b producing final state partons c and d, and $d\Pi$ is the phase space interval for the final state particles. The factor $2E_a2E_b|v_a - v_b|$, which involves the initial parton energies, E_a and E_b, as well as the relative velocity between them, $|v_a - v_b|$, serves as an initial flux factor. Perturbative QCD supplies $\mathcal{M}$, while the rest of the formula comes from the kinematics of the experiment.

The Feynman diagrams used to calculate $\mathcal{M}$ to first order in perturbation theory are shown in Figure 2.1. In these diagrams, each vertex contributes a factor of α_s to the amplitude, and the contribution from the individual diagrams are summed to give the total amplitude calculated to leading order in perturbation theory. Therefore, the leading-order calculation of the amplitude $\mathcal{M}$ is proportional to α_s^2, and the cross section computed from this amplitude squared is proportional to α_s^4.

The only processes that contribute to $b\bar{b}$ production at leading order are $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$. Thus, at this order in the calculation, only flavor creation plays a role. In addition, the only partons in the final state at this order in the calculation are the b and $\bar{b}$ quarks. In order to conserve momentum, these quarks must be produced back-to-back in the collision center of mass frame, leading to a $\Delta\phi$ distribution described by a delta-function peaked at $\Delta\phi = 180°$. This relatively simple picture of $b\bar{b}$ production and correlations will become more complicated as the calculation is taken to higher orders.

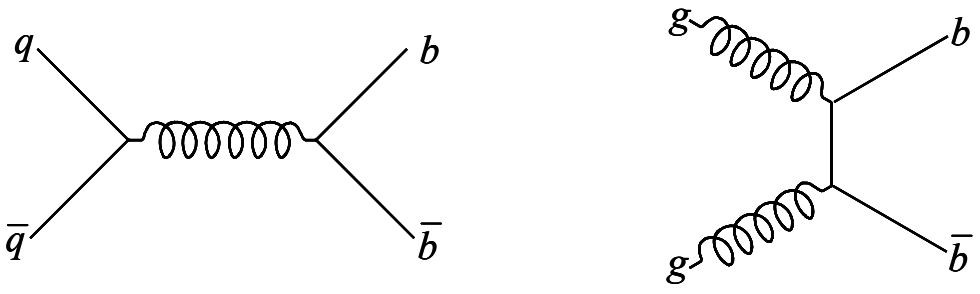

Figure 2.1 Leading-order $b\bar{b}$ production Feynman diagrams.

For many perturbative calculations, the leading-order is sufficient to give a reasonable qualitative and quantitative description of the physics in question. For $b\bar{b}$ production, this is not the case. The leading-order calculation differs significantly in its predictions for $b\bar{b}$ production rate and correlations from experimental measurements, suggesting that higher order processes, like flavor excitation and gluon splitting, play a significant role in $b\bar{b}$ production. Conceptually, this can be understood by considering the cross section for the process $gg \to gg$, which enters into $b\bar{b}$ production at higher orders through flavor excitation and gluon splitting. Flavor excitation and gluon splitting can be viewed as $gg \to gg$ processes where a gluon in the initial or final state splits into a $b\bar{b}$ pair. The cross section for $gg \to gg$ is approximately two orders of magnitude larger than the leading-order $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$ production mechanisms. Even after accounting for the suppression from the extra power of α_s and kinematics for one of the gluons splitting to a $b\bar{b}$ pair, the higher-order contribution from $gg \to gg$ is still comparable to the leading order $b\bar{b}$ production. Therefore, to get an accurate test of QCD through measurements of $b\bar{b}$ production, higher-order calculations are required.

2.3 Beyond Leading-Order

Going beyond the leading order in the perturbative calculation of $b\bar{b}$ production introduces a new level of complexity into both the calculation and the phenomenology predicted by the calculation. In addition to higher-order corrections to the flavor creation processes $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$ already present at leading order, processes with more particles in the final state, like $q\bar{q} \to b\bar{b}g$ or $qg \to b\bar{b}q$ enter into the higher order calculation. Some of these processes result from adding additional gluon radiation to the flavor creation process, while others enter because of contributions from flavor excitation or gluon splitting. These contributions increase the predicted rate of $b\bar{b}$ production, bringing it into closer agreement with measurement, and also yield more complicated $b\bar{b}$ correlations—for example, filling in the $\Delta\phi$ distribution for angles less than 180°.

Unfortunately, the higher-order corrections required to bring better agreement between theory and experiment are not straightforward to calculate. Evaluation of many of the higher-order Feynman diagrams leads to complicated integrals that are difficult to evaluate analytically and often contain divergences which must be carefully managed to produce sensible results from the calculation. Currently, there are two techniques available for taking the perturbative calculation

of $b\bar{b}$ production beyond leading order. One approach, known as the next-to-leading-order (NLO) approach calculates the exact $O(\alpha_s^3)$ contribution to the parton-level $b\bar{b}$ cross section. This approach has the advantage of being an exact calculation up to $O(\alpha_s^3)$ in perturbation theory, including correctly handling interference between the various production mechanisms. However, the NLO calculation only produces parton-level results for the *b* and $\bar{b}$ quarks, and it is difficult to relate these simple partonic final states to the complicated hadronic states observed at the Tevatron. An alternate technique for including the effects of higher-order corrections to $b\bar{b}$ production is the parton shower model. This model, also known as the leading-log approximation, starts with leading order matrix elements and the approximates the effects of higher order corrections to all orders in α_s using a probabilistic model for quarks to emit gluons and gluons to split to quark-antiquark pairs. The parton shower model has the advantage of more accurately modeling the complicated hadronic final states produced in the lab through the inclusion of multiple gluon emission. However, the parton shower model does not include the effect of interference between the various $b\bar{b}$ production mechanisms and is forced to rely more heavily on experiment to tune the model's free parameters. Recently, there has been a concerted effort to merge the two approaches to create a calculation that benefits from the strengths of both methods [31]. Unfortunately, the tools necessary for such a calculation are not yet finished, and thus won't be utilized in this thesis.

In the next sections, we will discuss some of the details of the two current approaches to higher order calculations.

2.3.1 Next-to-Leading Order

Figure 2.2 shows some representative examples of the diagrams that contribute to the next order of the perturbative calculation of $b\bar{b}$ production. These diagrams come in two categories. The first class consists of diagrams of order α_s^3 that involve the production of a $b\bar{b}$ pair in addition to a light quark or gluon. These processes include radiative corrections to the leading-order processes, as well as flavor excitation and gluon splitting, that do not contribute at leading order. The second class of diagrams include a set of virtual corrections to the leading order $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$. Although these diagrams are of order α_s^4, they contribute to the NLO cross section as cross terms when combined with the order α_s^2 terms in the squared amplitude $|\mathcal{M}|^2$. These cross terms give the same order contribution as α_s^3 diagrams.

Evaluation of the above Feynman diagrams leads to a number of divergent integrals. In the order α_s^3 diagrams, infrared divergences arise from gluons emitted with almost no momentum, while collinear divergences result from a radiated gluon being emitted along the same direction as a light quark. The order α_s^4 diagrams introduce divergences tied to the momentum of the virtual partons involved in the loop. Since these partons are not observed, they may take on any momentum value from zero to infinity. Divergences come from both the infinite momentum limit (ultraviolet) and the zero momentum limit (infrared). The integrals containing these divergences must be handled carefully and the divergent contributions isolated. When handled properly, the soft divergences in the order α_s^3 and α_s^4 diagrams cancel with one another. The remaining collinear and ultraviolet divergences must be carefully subtracted using the techniques of quantum field theory. As a result of these subtractions, the results of this calculation become dependent on two new energy scales: the renormalization scale (μ_R), which is used to determine

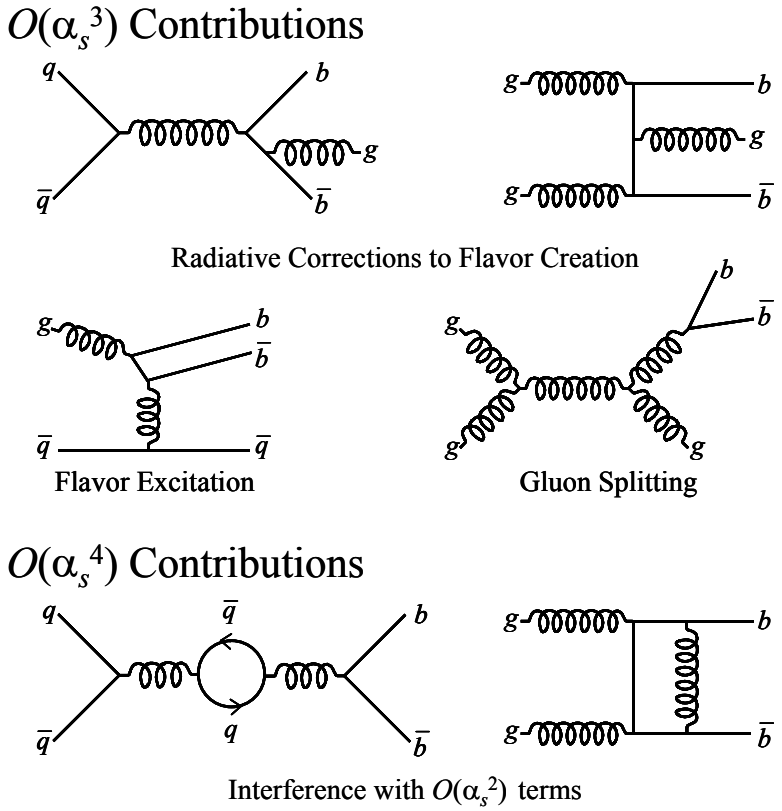

Figure 2.2 Representative Feynman diagrams for the NLO contribution to $b\bar{b}$ production.

the value of α_s, and the factorization scale (μ_F), which is used to determine the separation between the perturbative and non-perturbative regimes of QCD. Physically observable results should be independent of these scales, but finite-order perturbative calculations retain a residual dependence. The severity of this scale dependence can give an indication of the reliability of the perturbative calculation. For convenience the two scales are often set equal to one another. For $b\bar{b}$ production, a standard choice is $\mu_R = \mu_F \equiv \mu_0 = \sqrt{m_b^2 + (p_{Tb}^2 + p_{T\bar{b}}^2)/2}$.

Nason, Dawson, and Ellis (NDE) performed the first NLO calculation of $b\bar{b}$ production [32]. NDE calculated the cross section for producing single *b* quarks as a function of the rapidity and transverse momentum of the *b*, regardless of the kinematics of the other *b*. In performing this calculation, they integrated over the kinematics of the other *b* in the event, removing all information about the second *b* from their results. Thus, the NDE calculation cannot be used to predict $b\bar{b}$ correlations.

Mangano, Nason, and Ridolfi (MNR) obtained the first NLO production calculation that includes information about both the *b* and $\bar{b}$ produced [22]. The MNR calculation uses dimensional regularization to handle the soft and collinear divergences and numerical techniques to evaluate the necessary integrals. Care must be taken to organize the integrals so that the necessary cancellations occur among the divergences when evaluating the integrals numerically. The final result of the calculation is obtained with Monte Carlo techniques in which weighted parton level events are generated. Divergences are removed through the inclusion of events with large negative weights in certain regions of phase space to cancel out large positive contributions from the divergences. The delicacy of the cancellation of events with positive and negative weights makes this calculation difficult to use in certain circumstances. For this analysis, the presence of large negatively weighted events is undesirable, because we require realistic, fragmented final states and we apply a simulation of effect of the detector on the events, both of which may disturb the cancellation of positive and negative weights. For example, because fragmentation tends to smear out the momentum of the *b* quarks, events with large negative weights may be shifted so that they do not align with the positively weighted events, causing a distorted, or even negative result in certain part of the phase space. Also, an event with a large negative weight may be cut from the sample as part of the simulation of the detector efficiency and thus not be present to cancel out the positively weighted events. In general, Monte Carlo calculations that produce parton-level weighted events are not suitable for the kind of detailed, hadron level studies performed here.

2.3.2 Parton Shower Model

An alternative to the NLO approach for taking the calculation of $b\bar{b}$ production beyond the leading-order perturbation theory is known as the parton shower model. This model forms the basis for a number of popular Monte Carlo programs such as PYTHIA [33], HERWIG [34], and ISAJET [23]. These Monte Carlo programs are each capable of modeling a number of high-energy processes from both e^+e^- and $p\bar{p}$ collisions, including QCD $b\bar{b}$ production. For what follows, we will focus primarily on the features of the parton model implemented by PYTHIA and HERWIG, as they pertain to $b\bar{b}$ production.

The general approach of the parton shower model is to divide the QCD interaction into a hard component and a series of softer emissions known as initial- and final-state radiation. The hard scatter is identified as the part of the interaction with the largest momentum scale, Q^2_{Hard}. This part of the interaction involves only two incoming partons in the initial state, and two outgoing partons in the final state, and thus can be calculated using leading-order perturbative QCD. Added to this hard scatter is initial- and final-state radiation, in the form of gluons and photons radiated from quarks, and gluons splitting into quark-antiquark or gluon-gluon pairs. This radiation is approximated using a probabilistic model for gluon emission and splitting that can populate the final state of the interaction with a large number of quarks, antiquarks and gluons.

From the parton shower viewpoint, a typical collision between a proton and antiproton proceeds as follows, in approximately chronological order: First, a parton in each beam begins an initial state shower, branching into a number of quarks, antiquarks, and gluons. One parton from each of these showers participates in the hard 2-to-2 scattering, generating two final state partons. The final state partons shower, splitting into a number of quarks, antiquarks, and gluons. In addition, a "beam remnant" object is formed from each initial beam particle by subtracting the shower-initiating parton from the initial particle structure. This beam remnant has an overall color charge. Finally, confinement effects in QCD force the colored quarks, antiquarks, gluons, and beam remnants generated by the hard scatter and showers to fragment into color-neutral hadronic bound states. The structure of a typical event is illustrated in Figure 2.3. Of course, the actual Monte Carlo programs do not necessarily calculate these processes in this order. Typically, they start by calculating the 2-to-2 hard scatter, then proceeding to work forward from the hard scatter to add final state radiation and backwards from the hard scatter to add initial state radiation. Using the above steps, Monte Carlo programs like PYTHIA and HERWIG can approximate the rich and complicated structure of many-parton final states observed

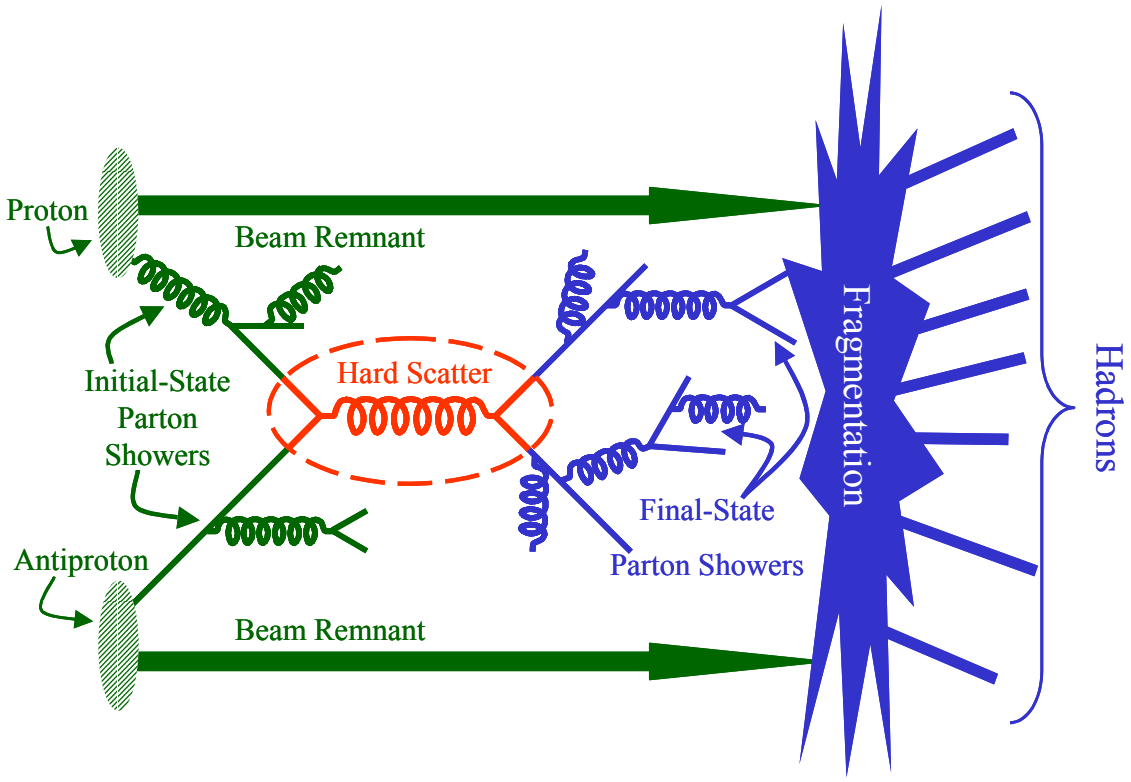

Figure 2.3 Typical event structure for a Monte Carlo event generated using the parton shower model.

in $p\bar{p}$ colliders like the Tevatron that would require unmanageably high-orders of a fixed-order perturbative calculation to handle.

A key element of the parton shower model is the showering algorithm. Shower development is modeled as a series of branchings such as $q \to qg$, $g \to q\bar{q}$, or $g \to gg$. Because of collinear divergences, small angle branchings receive a logarithmic enhancement, and therefore dominate the showering process. In addition, for suitably small angle splittings, the splitting process factorizes from the hard scatter process and can be computed separately as a multiplicative factor modifying the hard scatter cross section calculated using perturbative QCD. This is the justification that allows parton shower Monte Carlo programs to calculate initial- and final-state radiation using splitting probabilities appended to the quantum mechanical calculation of the hard scatter process. Also, for the case of multiple splittings, the degree of the logarithmic enhancement depends on the ordering of the splittings. Multiple splitting sequences that are most enhances are those ordered so that the splitting with the largest transverse momentum occur

closest to the hard scatter and the splitting p_T is smaller the farther one gets from the hard scatter. Other orderings are not as greatly enhanced. By keeping only diagrams with the largest logarithmic enhancement, the parton shower becomes a process of splitting ordered by p_T, or alternatively virtuality of the splitting partons. For initial state showers, the process progresses from some low virtuality scale up to the scale of the hard scatter. For final state showers, the evolution begins at the scale of the hard scatter and proceeds downward until some lower scale at which non-perturbative QCD effects become significant. Because the structure of shower evolution is determined by the prescription of considering only contributions from the most logarithmically enhanced diagrams at each order in the perturbation series, it is frequently referred to as a leading-log shower algorithm.

Although the algorithm used by parton shower Monte Carlo programs is very similar to analytic calculations made in the leading-logarithm approximation, the Monte Carlo programs incorporate a number of improvements. For example, the approximations necessary for the analytic calculation require one to view momentum conservation as a higher-order correction. However, in the Monte Carlo programs, because the detailed structure of the splittings in the shower is tracked, momentum conservation can be enforced at each step. In addition, sophisticated parton Monte Carlo programs like PYTHIA and HERWIG incorporate corrections from coherence effects from beyond leading-log order. Studies of higher-order effects in shower development [35] suggest that destructive interference in the shower suppresses splittings in which the splitting angle is greater than the angle of the preceding splitting. This effect is incorporated in parton Monte Carlo programs through angular ordering, the requirement that the splitting angle become smaller with each splitting. Through the incorporation of improvements like these, the results from parton shower Monte Carlo programs are superior to the approximations made with a simple leading-logarithmic analytic calculation, although they still lack the rigor of a complete higher-order treatment.

Unlike in the NLO calculation, the three production mechanisms—flavor creation, flavor excitation, and gluon splitting—can be rigorously separated from one another. Interference effects that intermingle these three production mechanisms do not appear until higher orders in this approximation. That is, the interference may be treated as a higher-order correction that is omitted at this order in the approximation. Each event can be classified as flavor creation, flavor excitation, or gluon splitting based on how many bottom quarks are produced in the hard 2-to-2

interaction compared to the number generated as part of the initial or final state. The classification scheme proceeds as follows:

Flavor Creation: Flavor creation events occur when both *b* quarks are produced as part of the hard scatter. The flavor creation process involves the leading-order matrix elements for the processes $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$. Initial- and final-state showers populate the final state with additional light quarks and gluons.

Flavor Excitation: In flavor excitation, only one of the *b* quarks is generated as part of the hard scatter. Flavor excitation events involve the leading order matrix elements for the following processes: $qb \to qb$, $q\bar{b} \to q\bar{b}$, $gb \to gb$, and charge conjugate processes. The bottom quark in the initial state of this hard scatter is generated through the evolution of the parton distributions functions of the beam particles. (See section 2.4.1 for more information on parton distribution functions.) The additional bottom quark is produced as part of the initial-state shower. The parton shower Monte Carlo program traces the initial state shower history back from the *b* quark involved in the hard scatter until it finds $g \to b\bar{b}$ splitting to account for the other *b* quark in the event.

Gluon Splitting: In gluon splitting events, neither the *b* nor the $\bar{b}$ participates in the hard scatter. The leading-order matrix elements relevant to this process include 2-to-2 processes that don't involve bottom quarks, although the process $gg \to gg$ gives the dominant contribution. The $b\bar{b}$ pair arises from a $g \to b\bar{b}$ splitting in either an initial-state or final-state shower. However, kinematics suppress the production of $b\bar{b}$ pairs from initial-state splittings, so the dominant contribution comes from final state showers.

It is interesting to note that it is possible for the parton shower model to generate events having more than one $b\bar{b}$ pair in the final state. These events occur when additional $g \to b\bar{b}$ splittings occur in the initial- or final-state showers that accompany any of the above processes. Thus, events containing both flavor creation and gluon splitting, or flavor excitation and gluon splitting do occur, although at a much lower rate than processes producing only one $b\bar{b}$ pair. These processes will be present in the data sample used for this analysis, and no explicit attempt is made to identify or remove these events. However, multi-$b\bar{b}$ QCD events are very interesting not only as an additional study of QCD, but also as a background for exotic processes, most notably Higgs production.

2.3.3 Combining the Parton Shower Model with the NLO Calculation

Although current parton shower Monte Carlo implementations limit themselves only to leading-order matrix elements for the hard scattering processes in $b\bar{b}$ production, there is no reason in principle that higher-order calculations could not be employed as well. In principle, the NLO calculation for the 2-to-2 and 2-to-3 hard processes involved in $b\bar{b}$ production could be combined with the parton shower model to generate addition quark, antiquark, and gluon radiation in the final state. Such a Monte Carlo calculation would have the advantages of incorporating the full interference structure of the NLO calculation with the more realistic jet structure of the parton shower model. However, care must be taken to avoid double-counting of final states accessible from either the leading-order 2-to-2 matrix element augmented with parton showering or the 2-to-3 NLO matrix element. For example, the $b\bar{b}g$ final state can be produced either from the leading-order matrix element for $b\bar{b}$ with an additional gluon coming from final-state radiation, or from the NLO matrix element for $b\bar{b}g$ production. Care must be taken in this case, to insure that the parton shower algorithm does not double-count production of events in the parts of phase space covered by the NLO matrix elements. At the time of this writing, there are no Monte Carlo programs available that merge the NLO result with the parton shower model. However, on-going effort is directed towards producing usable code in the near future [31].

2.4 Non-Perturbative Aspects

Every measurement of $b\bar{b}$ production involves low energy strong interactions that cannot be described within the framework of QCD. This is because the confining nature of QCD guarantees that in the laboratory, all quarks and gluons will be contained within hadronic bound-states that involve low energy strong force interactions. The factorization theorem states that this low-energy, non-perturbative QCD physics will not disturb the characteristics of the high energy interactions too much, but it is still necessary to have some sort of model to describe the non-perturbative physics in order to execute high precision tests of perturbative QCD.

In $p\bar{p}$ collisions at the Tevatron, it is necessary to use non-perturbative models to describe conditions in both the initial state and the final state of any high-energy QCD interaction. The initial state of each collision involves a proton and an antiproton. These two hadronic bound states can be viewed as complicated collections of quarks, antiquarks, and gluons, distributed according to a set of parton distribution functions that cannot be calculated from perturbation

theory but must instead be measured by experiment. On the other hand, the final state of any interaction involves not the quarks and gluons of QCD, but hadronic bound states. The process of transforming a set of quarks and gluons from a QCD interaction into final state hadrons is known as fragmentation. In the following section, we will discuss further the different phenomenological models available to describe the non-perturbative physics in these two regimes.

2.4.1 Parton Distribution Functions

The initial state for the high energy QCD interactions produced at the Tevatron involve not quarks and gluons, but protons and antiprotons. Unlike quarks and gluons, which appear to be fundamental, structureless particles at the energy scales probed so far, the proton is a complicated hadronic bound state with a rich structure that depends on the distance—or equivalently, momentum—scale at which it is probed. At low momentum (and therefore large distance scales), the proton appears to be a fundamental particle itself. As the momentum scale increases and the distance scale decreases, it becomes apparent that the proton is actually the bound state of three quarks (two up quarks and a down quark, called valence quarks) interacting via the strong force. Each of the quarks takes up a fraction of the proton's total momentum, although soft QCD interactions cause the exact fraction that a particular quark carries to shift from one moment to the next. In fact, at any given time, the gluon field itself carries a non-negligible fraction of the proton momentum. As the structure of the proton is probed at higher and high momentum (or shorter and shorter distance scales), an even richer structure is revealed. In addition to the three valence quarks and the gluons exchanged between them, the proton is seen to contain a number of other quarks and antiquarks generated by temporary fluctuation of a gluon into a virtual quark-antiquark pair. The complicated and momentum-dependent structure arises from QCD interactions that cannot be calculated using perturbative methods.

Although the absolute structure of the proton cannot be calculated perturbatively, the way in which that structure changes, or evolves, can be. Proton structure is parameterized according to a set of distributions $f_i(x)$ that describe the probability of finding a parton (quark, antiquark, or gluon) of type i with fraction x of the proton momentum. These functions are known as parton distribution functions, and they depend the momentum scale Q at which the proton structure is probed. Given a set of parton distribution functions at one momentum scale, their form at any other momentum scale can be calculated using perturbative methods very similar to those used to

generate parton showers in Monte Carlo programs like PYTHIA and HERWIG. The final result of such a calculation is a set of parton evolution equations known as the Alteralli-Parisi equations, given below:

$$\frac{d}{d\log Q} f_g(x,Q) = \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left\{ P_{g\to q}(z) \sum_f \left[f_f(\frac{x}{z},Q) + f_{\bar{f}}(\frac{x}{z},Q) \right] + P_{g\to g}(z) f_g(\frac{x}{z},Q) \right\}$$

$$\frac{d}{d\log Q} f_f(x,Q) = \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left\{ P_{q\to q}(z) f_f(\frac{x}{z},Q) + P_{q\to g}(z) f_g(\frac{x}{z},Q) \right\} \qquad (2.3)$$

$$\frac{d}{d\log Q} f_{\bar{f}}(x,Q) = \frac{\alpha_s(Q^2)}{\pi} \int_x^1 \frac{dz}{z} \left\{ P_{q\to q}(z) f_{\bar{f}}(\frac{x}{z},Q) + P_{q\to g}(z) f_g(\frac{x}{z},Q) \right\}$$

In the equations above, $f_g(x,Q)$ represent the parton distribution function for the gluons while $f_f(x,Q)$ and $f_{\bar{f}}(x,Q)$ represent the parton distribution functions for each flavor a quark and antiquark. The splitting functions $P_{x\to y}(z)$ parameterize the probability that a splitting of parton x will result in a parton y with a fraction z of x's momentum. For example, $P_{q\to g}(z)$ describes the probability that a gluon with momentum fraction z would be produced from a $q \to qg$ splitting. These splitting functions can be calculated using perturbative methods and the leading-order splitting functions are essentially the same functions used to describe the splittings in parton shower Monte Carlos. As can be seen from the form of the equations above, the splitting functions cause an interrelation between the parton distribution functions as they evolve from one momentum scale to another. The net effect of all this is that at higher and higher momentum scales, all parton distribution functions tend towards lower momentum fractions because it becomes more likely that the particular parton measured is the result of one or more splittings from a parton with higher momentum fraction.

All that remains to give a complete description of the proton structure in this framework is the exact form of the parton distribution functions for some particular value of Q^2. Since the solution to this problem relies on physics from the non-perturbative regime of QCD, answer must be determined by fitting phenomenological models to experimental data with the use of the Alterelli-Parisi equations to relate data taken with different Q^2 values. There are currently two main groups that perform global fits of a large number of experimental data points to determine the parton distribution functions for the proton: the CTEQ [36] and MRST [37] collaborations. Both groups use essentially the same data although they differ somewhat in their treatment of

individual data sets and in their fitting techniques. The main inputs to these global fits include the following:

- Deeply inelastic scattering (DIS): DIS experiments involve scattering a lepton from a target containing protons or neutrons. Data used includes electron-proton scattering, muon-nucleon scattering, and neutrino-nucleus scattering.

- Lepton pair production: Information about the proton structure can be obtained from lepton pair production through the Drell-Yan process in high-energy proton-proton or proton-deuteron collisions.

- Lepton charge asymmetry: Additional input comes from measuring the charge asymmetry of leptons from the decays of W bosons produced in high-energy $p\bar{p}$ collisions.

- High-p_T jet production: In the high transverse momentum region, where perturbative QCD provides a good description of the physics, measurements of jet production at hadron colliders can provide additional constraints on the proton parton distribution functions.

In general, DIS experiments determine the quark portion of the parton distribution functions. Studying DIS experiments with proton and neutron targets, as well as studying lepton charge asymmetries in Drell-Yan and W boson production helps determine the difference between the u and d quark distributions in the proton. High-p_T jet production at hadron colliders provides the primary determination of the gluon component of the parton distribution functions.

The $b\bar{b}$ production studied in this analysis is sensitive primarily to the low-x part of the parton distribution functions, although there is an indirect dependence on the high-x gluon component through the evolution of these high-momentum-fraction gluons into lower momentum $b\bar{b}$ pairs that participate in flavor excitation. For this analysis, the CTEQ5L parton distribution functions [36] are used for Monte Carlo generated using both PYTHIA and HERWIG. This parton distribution function set is fit using leading-order calculations for the parton evolution equations.

2.4.2 Fragmentation

In the laboratory, because of confinement in QCD, measurements of bottom quark production are made using *B* hadrons in the final state rather than *b* quarks. The process of going from quarks produced in high energy QCD interactions to a collection of hadronic jets in the final state is known as fragmentation. This process involves a transition through the low-energy regime of QCD that cannot be treated using perturbative methods. Therefore, a number of phenomenological models are used to parameterize the non-perturbative physics involved in fragmentation. Typically, these models contain one or more adjustable parameters that rely on experimental input to determine their values.

It is important to realize that fragmentation is modeled as a two-stage process. The first stage, which we will call perturbative fragmentation, involves the emission of gluons from quarks and the splitting of gluons into $q\bar{q}$ pairs at a momentum scale in which perturbative methods apply with reasonable accuracy. The exact form of the perturbative portion of the fragmentation model is intimately related to the nature of the perturbative methods used to calculate $b\bar{b}$ production. For fixed-order perturbative calculations, like the NLO calculation, the perturbative component of fragmentation is included in the diagrams containing radiative corrections to leading-order $b\bar{b}$ production. It is also possible to improve the perturbative description of fragmentation in fixed-order calculations by resuming diagrams involving collinear gluon emission to all orders. Such calculations are sometimes referred to as fixed-order, next-to-leading log (FONLL) calculations and have shown some promise in reducing the discrepancy between the NLO calculation of $b\bar{b}$ production and experimental measurements [16]. Finally, for parton shower Monte Carlo calculations of $b\bar{b}$ production, the perturbative component of fragmentation is encapsulated as part of the parton shower. The parton shower process in such Monte Carlo programs accounts for the effects of multiple gluon emissions and gluon splittings up to some lower momentum scale Q_0, at which point, non-perturbative fragmentation takes over.

The second component to the fragmentation model is the non-perturbative component. This component takes the form of a phenomenological model to approximate the effects of low energy QCD processes that cannot be computed using the usual perturbative methods. A number of approaches exist to handle this component of the fragmentation process, and they can be divided into two broad categories: independent fragmentation and string/cluster fragmentation. Independent fragmentation models approximate the affects of non-perturbative fragmentation processes of each final state parton separately. These models don't take into account information

about the colors of the various final state partons and how they are connected to one another through their color fields. In contrast, string or cluster fragmentation models utilize color flow information in an attempt to account for correlations between different colored objects in the process of creating color singlet final states. Because the NLO calculation does not retain color information for the quarks produced, only independent fragmentation models can be used for this calculation. Independent fragmentation can also be used in parton shower Monte Carlo programs (for an example, see ISAJET), but the programs used for this analysis, PYTHIA and HERWIG, use the more complicated string and cluster models respectively.

For the NLO calculation of $b\bar{b}$ production, the non-perturbative part of fragmentation for b quarks is described in terms of a fragmentation function $f(z)$ that describes the probability that a B hadron will have a fraction z of the b quark's energy and momentum, with z is defined as follows:

$$z = \frac{E(B) + p_{\parallel}(B)}{E(b) + p(b)} \tag{2.4}$$

The most commonly used fragmentation function is from Peterson, *et al*. [14]. This fragmentation function is derived from the assumption that most of the b quark energy and momentum is transferred to the B hadron created in fragmentation. Quantum mechanically, it can be argued that the amplitude for the transition from the state having a lone b quark to the state having a B meson ($b\bar{q}$ -bound state) and a light quark q is inversely proportional to the energy difference between those states $\Delta E = E_B + E_q - E_b$. From these assumptions, Peterson *et al*. derive the following form for their fragmentation function:

$$f(z) = \frac{N}{z\left(1 - \frac{1}{z} - \frac{\varepsilon}{1-z}\right)}, \tag{2.5}$$

where N is a normalization constant that depends on the number of B hadron states accessible to the B hadron, and ε is a parameter that must be determined experimentally. Traditionally, the parameter ε is determined from e^+e^- experiments where the clean environment allows a better measurement of the momentum fraction z. The traditional value of ε used for $b\bar{b}$ production at the Tevatron with the NLO calculation has been $\varepsilon = 0.006$, a value determined from fits to e^+e^- data involving leading-logarithmic QCD calculations. However, reference [38] indicates that the

value of ε depends of the type of perturbative QCD calculation for $b\bar{b}$ production used in its extraction. Recent studies indicate that a much smaller value of ε in the range of 0.0016 to 0.0033 are more appropriate for the NLO calculation of $b\bar{b}$ production [38].

In parton shower Monte Carlo programs, like PYTHIA and HERWIG, more sophisticated fragmentation models are used that keep track of the connection between different quarks and gluons in the final state in addition to the remnants of the proton and antiproton from which the partons involved in the hard collision are taken. The string fragmentation model used in PYTHIA begins by grouping the final state partons into color singlets by tracing through the color flow from the initial states to the final states. Each of these groupings is treated as if connected by a "string" representing a "color flux tube" whose energy increases linearly with the separation of the color charges. When the separation becomes large enough, the amount of energy stored in the string will be sufficient to allow the string to be broken by the formation of a $q\bar{q}$ pair. The breaking of the string through the formation of a $q\bar{q}$ pair is modeled as a quantum mechanical tunneling process, leading to a suppression of the production of higher mass quark flavors. This mass suppress effect leads to a ratio of quark flavors for the string-breaking pair of u:d:s:$c \approx 1{:}1{:}0.3{:}10^{-11}$; thus the production of charm and heavier quarks through non-perturbative fragmentation is treated as negligible in PYTHIA. The process of string breaking continues until the invariant mass of all strings is small enough to allow the strings to decay to either a pair of on-shell hadrons or a single hadron, with energy and momentum being conserved by shifting the necessary amounts to other strings. The probability distribution for generating a hadron with energy and longitudinal momentum z from the string is given by a fragmentation function $f(z)$, where $z = (E + p_{\parallel})_{hadron} / (E + p_{\parallel})_{string}$. The default fragmentation function for PYTHIA when one of the string endpoints is a massive quark, like a b quark, is the Bowler fragmentation function [39]:

$$f(z) \propto \frac{1}{z^{1+R_Q bm_Q^2}} (1-z)^a \exp(-bm_\perp^2), \tag{2.6}$$

where, m_Q is the mass of the b quark, R_Q for all quark flavors is assumed to be one, and a and b must be determined with experimental input, and $m_\perp$ is the transverse mass of the string, given by $m_\perp^2 = E^2 - p_\parallel^2$. It is possible to use other fragmentation functions with PYTHIA, such as the Peterson fragmentation function, but note that the variable z has a different meaning when used in Peterson fragmentation in the string model than it has when used with the NLO calculation.

The cluster fragmentation model of HERWIG is very similar to the string fragmentation model of PYTHIA, although it focuses on the one- and two-body decay of low invariant mass clusters rather than the iterative decay of high mass strings. Cluster fragmentation is based on the idea of "preconfinement" which states that during the evolution of a parton shower, the produced partons will become arranged in low mass, color neutral clumps. This idea is implemented in HERWIG as follows: At some low momentum scale, Q_0 during the parton shower, the perturbative process of gluon emission and gluon splitting is terminated and all remaining gluons are non-perturbatively split into $q\bar{q}$ pairs. The color information is traced from the initial state of the interaction to the final state, allowing the collection of quarks and antiquarks in the final state to be assembled into a collection of color neutral cluster that, according to the principle of preconfinement, have low invariant masses. Clusters are decayed into hadrons based on their invariant mass. Clusters whose mass is too small to decay to a pair of hadrons decay to a single hadron and momentum and energy are conserved by shifting the necessary amount to another cluster. The small fraction of clusters whose mass is above a certain threshold are split using an iterative procedure similar to string fragmentation until the resulting clusters are all below the splitting threshold. Clusters in between these two regimes are decayed to two hadrons. The flavor of the $q\bar{q}$ inserted into the cluster to allow it to decay to two hadrons is randomly selected and the kinematics of the decay are chosen using a probabilistic model based on the available phase space for the decay.

Chapter 3

The Tevatron Accelerator and CDF Detector

Nestled among the suburbs of Chicago, the Fermi National Accelerator Lab—abbreviated FNAL or Fermilab—has long been one of the premier accelerator facilities in the world. Since the mid-1980's, Fermilab has generated proton-antiproton collisions at the highest center of mass energy (often referred to in terms of the Mandelstam variable $\sqrt{s}$) in the world through the use of the Tevatron accelerator. The Tevatron is a circular accelerator, 6 km in circumference, that brings beams of 900 GeV protons and antiprotons into collision at two interaction regions where the CDF and DØ detectors are located. In its nearly two decades of operation, the Tevatron has produced a number of results in a variety of areas in high energy physics, from the measurement of charged particle multiplicities in $p\bar{p}$ collisions to the discovery of the top quark. This thesis reports on a measurement made using the CDF detector at the Tevatron.

Since the first collider tests in 1985, the Tevatron has operated in periods of data taking, known as "runs," separated by periods in which the accelerator and detector were upgraded, and fixed-target runs were performed. CDF took data during two short runs in 1987 and 1988-1989. After a period of time in which the DØ detector was completed and the CDF detector was upgraded, the first Tevatron run including both CDF and DØ, known as Run 1, began in 1992. This run was divided into parts, Run 1a, lasting from 1992 to 1993, and Run 1b, lasting from 1994 to 1995. The time between Run 1a and Run1b was used to, among other things, replace the silicon vertex detector at CDF. After a long shutdown in which the Tevatron accelerator complex, as well as the CDF and DØ detectors, underwent extensive upgrades, a new Tevatron run, Run II, began in 2001. A central part of the upgrade was the construction of a new accelerator, the Main Injector, which replaced the function of the Main Ring in proton and antiproton production. The Main Injector delivers greatly improved proton and antiproton production capabilities, as well as allowing simultaneous operation of the collider and fixed

target programs. In addition, as part of the Run II upgrades, the Tevatron beam energy was increased from 900 GeV to 980 GeV. Run II is expected to last most of this decade.

The analysis described in this thesis uses data taken at CDF during Run 1b of the Tevatron (1994-1995). The following sections will briefly describe the Tevatron and the CDF detector as they were configured during that time period. Emphasis will be placed on aspects of the accelerator and detector most relevant to this analysis.

3.1 The Tevatron

The Tevatron is the world's highest energy proton-antiproton collider, capable of accelerating beams of protons and antiprotons to energies of 900 GeV [40]. However, getting to that point requires a number of intermediate accelerators both to accelerate protons from low energies and to create antiprotons through collisions between a proton beam and a fixed target [41]. The complex of accelerators involved in generating beams for the Tevatron is shown in Figure 2.1.

Generation of a proton beam for the Tevatron begins with a Cockcroft-Walton electrostatic

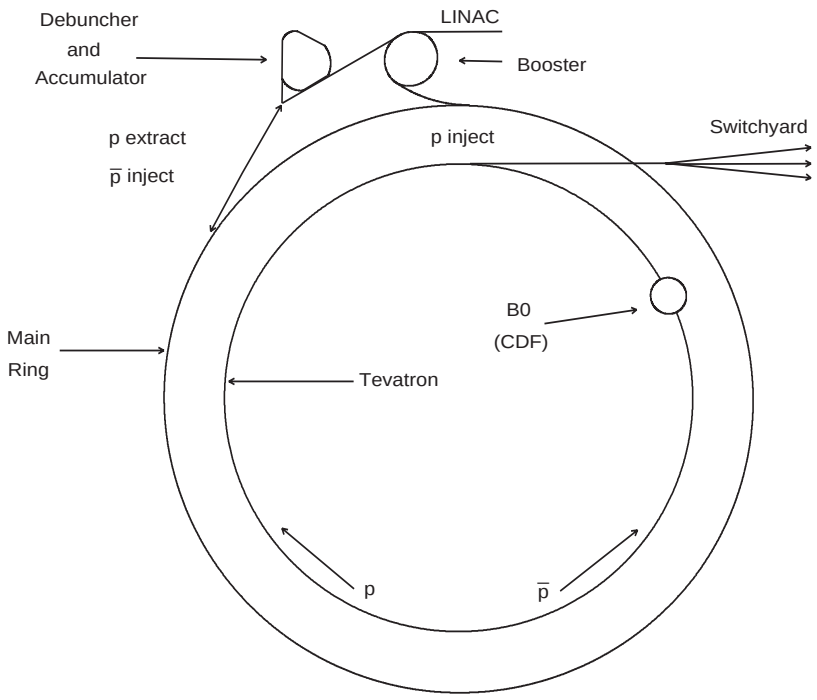

Figure 3.1 A schematic view of the Tevatron accelerator complex.

generator that accelerates H^- ions to an energy of 750 keV [42]. Next, these H^- ions are transferred to a 150 m long linear accelerator (Linac) that accelerates them to 400 MeV [43]. After the Linac, the H^- ions are transferred to a circular, alternating-gradient synchrotron accelerator with a 475 m circumference, known as the Booster Ring. There, a thin carbon film is used to strip the electron from the H^- ion, creating a proton beam which is then accelerated to an energy of 8 GeV, collected into bunches, and stored. Injecting H^- ions rather than protons into the booster allows the injection to proceed over multiple revolutions of the beam around the Booster Ring. If protons were instead injected, the magnetic field used to kick new protons onto orbit in the booster would also kick the existing protons out of orbit. Proton bunches are extracted from the Booster Ring and transferred to a second synchrotron called the Main Ring. The Main Ring, which was Fermilab's original high-energy, is located in the same tunnel as the Tevatron. The Main Ring accelerates protons to an energy of 150 GeV. From there, the protons can be extracted and injected into the Tevatron for their final stage of acceleration up to 900 GeV.

The process for generating an antiproton beam for the Tevatron is the same as the process for generating a proton beam, up to the stage in which a proton beam is accelerated to 150 GeV in the Main Ring. From that point, antiprotons are created by extracting protons from the Main Ring and colliding them with a nickel-copper target [44]. Every antiproton captured requires approximately 100,000 protons to be collided with the target. Antiprotons are extracted from the collision products using a mass spectrometer and directed to the Debuncher [45], a triangularly-shaped storage ring where the antiproton beam is "cooled" by converting narrow pulses of antiprotons with wide momentum spreads into broad pulses with small momentum spreads [46]. The resulting antiproton beam is transferred to a second, triangular storage ring, the Accumulator [47], which is co-centric with the Debuncher. Here antiprotons are stored until sufficient numbers have been produced for collider running. Once enough antiprotons have been accumulated, antiproton bunches are loaded into the Main Ring, just like protons, except traveling in the opposite direction because of their negative charge. From the Main Ring, the antiprotons are accelerated to 150 GeV and injected into the Tevatron (again moving in the opposite direction as the protons), where they are eventually brought into collision with the protons to produce the high-energy interactions that produced *b* quarks for this analysis.

A period of time in which colliding beams are present in the Tevatron is known as a "store." During a typical store for Run 1b, the Tevatron contained six bunches with 18×10^{10} protons and

six bunches with 6×10^{10} antiprotons, each traveling near the speed of light and separated from one another by 3.5 μs. The proton and antiproton beams are placed on helical orbits and "cogged" so that they only intersect and produce collisions at the designated interaction regions for CDF and DØ. In these regions, quadrupole magnets focus the beams down to a small cross-sectional area for maximum beam intensity at the point of collisions. The intensity of the colliding beams at the interaction point is called the "luminosity," and can be calculated from the number of protons and antiprotons per bunch, N_p and $N_{\bar{p}}$, the collision frequency of the bunches, f, and the transverse size of the beam in the horizontal and vertical directions, respectively σ_x and σ_y, according to the following formula:

$$\mathcal{L} = \frac{N_p N_{\bar{p}} f}{4\pi\sigma_x\sigma_y} \tag{3.1}$$

In the transverse plane, the beam is approximately circular in shape with a size of $\sigma_x = \sigma_y =$ 35 μm. Thus, the typical luminosity at the start of a store during Run 1b was around 2×10^{31} cm^{-2}s^{-1}. As a store progresses, the number of antiprotons and protons is depleted by two mechanisms: losses due to orbit variations and annihilations of proton-antiproton pairs.

The amount of data collected at CDF can be expressed as the luminosity integrated over the time during which the detector was taking data. During Run 1b, the total amount of data collected at CDF corresponds to an integrated luminosity of approximately 90 pb^{-1}, where a *barn*, b, is equal to 10^{-24} cm^2, and an *inverse picobarn*, pb^{-1} is 10^{12} *inverse barns*, b^{-1}.

3.2 The CDF Detector

The CDF detector [48] is a multipurpose particle detector designed to be flexible enough to allow a program of measurements in various areas of high-energy physics from precision QCD and electroweak measurements to searches for new physics. The main components of the CDF detector include the tracking systems, which make measurements of the trajectories of stable, charged particles, the calorimetry systems, which measure the energy of charged and neutral particles, and the muon systems, which identify the presence of muons by their ability to penetrate large amounts of matter. In addition, the readout of the CDF detector is controlled by a three-level trigger system, whose job is to reduce the data rate so that the limited bandwidth of the detector readout can be allocated to allow a variety of different physics measurements. Since

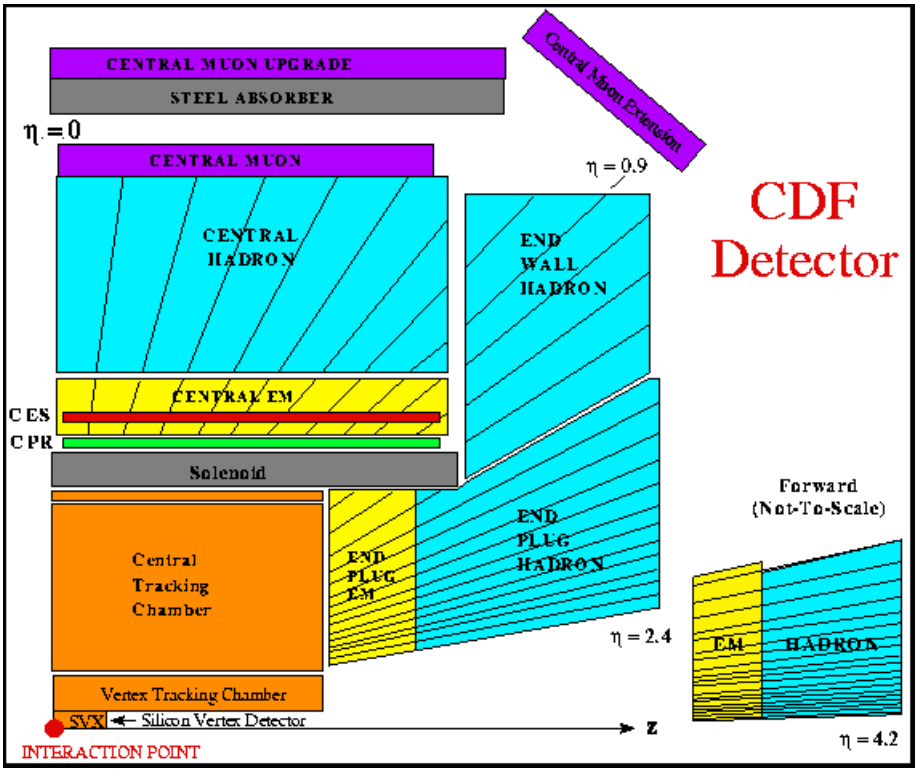

Figure 3.2 A view of one quadrant of the CDF detector in the *r*-*z* plane. Since the CDF detector is symmetric under rotations in ϕ and reflections about the *x*-*y* plane, the other three quadrants are the same.

its construction in the mid-1980's, the CDF detector has undergone a number of upgrades. The detector configuration we describe in the rest of this section was the one used during Run 1b, from 1994 to 1995, as shown in Figure 3.2. In the following description, we will place emphasis on those elements of the CDF detector used for this analysis.

Throughout the following discussion it is important to bear in mind the coordinate systems employed at CDF. In general, the CDF coordinate system is defined so that the *z*-axis is aligned in the direction of the proton beam, or east at CDF. The *x*-axis points outwards from the center of the Tevatron ring, which is north, while the *y*-axis points up. However, because of the cylindrical symmetry of the CDF detector about the *z*-axis, it is often useful to use a cylindrical coordinate system with *r* giving the distance from the *z*-axis, and ϕ giving the azimuthal angle, measured from the *x*-axis. Occasionally, coordinates are given in terms of θ, the polar angle,

measured from the *z*-axis. Another useful coordinate is *rapidity* because particle production per unit rapidity is constant. The rapidity, *y*, is defined by

$$y = \frac{1}{2}\ln\left(\frac{E + p_z}{E - p_z}\right). \tag{3.2}$$

Pseudorapidity, defined as $\eta = -\ln[\tan(\theta/2)]$, provides a purely geometric approximation to rapidity in terms of the polar angle θ. In the limit of a massless particle, pseudorapidity is exactly equal to rapidity. For high-energy massive particles, rapidity and pseudorapidity are roughly equivalent.

3.2.1 Tracking Systems

The purpose of the tracking systems at CDF is to measure the trajectory of charged particles using the charge liberated as the charged particle ionizes the active medium of the detector. The CDF tracking systems are placed within a nearly uniform 1.4 T solenoidal magnetic field [49], with the field direction oriented along the *z*-axis. Because of this field, particles travel in helical paths in the CDF tracking systems, with the radius of the helix related to the particle's momentum transverse to the beam direction by

$$p_T = \frac{c \cdot B}{2C} \times 10^{-9}, \tag{3.3}$$

where *c* is the speed of light in m/s, *B* is the magnetic field, approximately 1.4 T,[1] and *C* is the radius of curvature of the helical track expressed in units of m^{-1}. The resulting p_T is given in GeV/c. The *z*-component of the particle's momentum can then be calculated from the p_T using $p_z = p_T \cot\theta$. Thus the complete information about a particle's momentum and trajectory in space can be encapsulated in the following five numbers, known as tracking parameters:

- cotθ: The cotangent of the polar angle of the particles trajectory.
- *c*: The curvature of the particle's helical path. Curvature is also sometimes referred to as κ. Curvature is defined as 1/(2ρ) where ρ is the radius of the track's helical trajectory.

[1] The *z*-component magnetic field of the CDF solenoid has been measured to a precision of 0.5×10^{-4} T, and inhomogeneities in the *z*-component as well as radial components of the field have been mapped.

- z_0: The z position of the particle at its point of closest approach to the origin of the coordinate system.

- D_0: The impact parameter or distance of closest approach to the origin.

- ϕ_0: The azimuthal angle of the particle's trajectory at the point of closest approach to the origin.

The origin is intended to represent the location of the $p\bar{p}$ collision. Before the full event has been reconstructed, the default origin is assigned to be the center of the CDF detector and all tracks are measured with respect to this point. Later, when the event has been fully analyzed and a better determination of the $p\bar{p}$ interaction point is available, the track parameters are translated to reflect this new origin. In either case, the momentum and trajectory of the measured track remain the same; only the representation changes.

The CDF tracking systems are arranged as concentric cylindrical volumes. The innermost tracking system, the silicon vertex detector (SVX') is surrounded by the vertex time projection chamber (VTX), which is in turn surrounded by the central tracking chamber (CTC). The purpose of the various tracking systems is to measure the five tracking parameters for each particle produced with $|\eta| < 1$. Individually, the different tracking systems are capable of measuring these parameters with differing degrees of precision. The SVX' provides the most precise determination of the impact parameter, D_0. The VTX is optimized for measuring z_0. The CTC forms the starting point for all track reconstruction. In addition, it is especially important for measuring the curvature, c, and hence the particle's transverse momentum. Combining information from multiple systems when possible yields the best determination of a particle's momentum and trajectory.

This analysis relies heavily on information from the CDF tracking detectors. Information from all three detectors is used to determine the location of the $p\bar{p}$ collision point, known as the primary vertex, as well as the decay location of any long-lived particles, like B hadrons. Tracking information is also combined with signals from the calorimeter and muon systems for lepton identification. Finally, tracking information provides an estimate for the direction of travel of the B hadrons, allowing the measurement of the transverse opening angle, $\Delta\phi$.

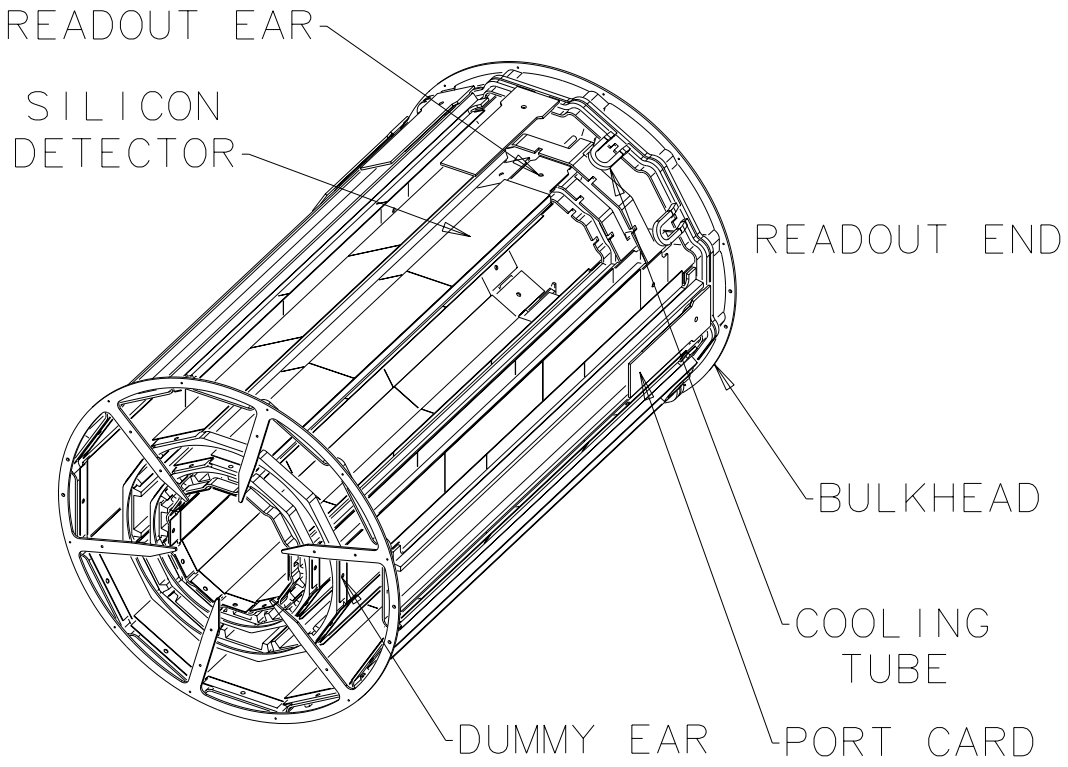

Figure 3.3 The SVX' detector.

3.2.1.1 Silicon Vertex Detector (SVX')

The SVX' [50] is a four-layer, silicon microstrip vertex detector designed to make precision tracking measurements close to the beamline to allow a high resolution determination of a track's impact parameter. The "prime" in SVX' denotes that this is the second silicon detector at CDF, a replacement for the original silicon detector installed in Run 1a. Below we will focus on a few details important for this analysis.

The SVX', the closest detector to the beam pipe at CDF, consists of four layers composed of 12 planes or "ladders" of silicon arranged in concentric cylinders around the beam pipe, as shown in Figure 3.3. The innermost layer is located at a radius of 2.94 cm from the center of the beam pipe while the outermost layer is at a radius of 8.07 cm. Each silicon ladder consists of a bulk layer of lightly doped n-type (n^-) silicon sandwiched between finely spaced strips of lightly doped p-type (p^+) silicon on one side and a thin, continuous layer of strongly doped n-type (n^+) silicon on the other. A positive voltage is applied to the n^+ layer, depleting the bulk layer of

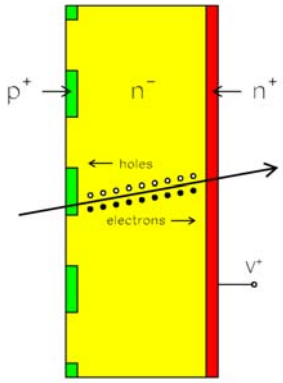

Figure 3.4 A schematic of how a hit is registered in a silicon detector. A charged particle traversing the bulk silicon layer creates electron-hole pairs. The holes drift in the electric field in the bulk silicon to the *p*-type strips, where they cause electrical signals allowing the detection of a hit.

electrons and creating an electric field in its volume. This causes the silicon ladder to act like a solid-state drift chamber.

When an ionizing particle traverses the bulk silicon wafer, it creates electron-hole pairs. Because of the electric field from the applied voltage, the electrons drift towards the p-type strips, where they cause an electric signal that can be read out to determine the position at which the particle passed through the silicon ladder. Figure 3.4 shows a schematic of this process. Because the strips are oriented along the direction of the beam, they provide only a measurement of the ϕ-position of the particle as it passes through each of the four layers at a different radius, allowing a determination of the r-ϕ trajectory of the track. No z position information is available from the silicon detector. The fine spacing of the silicon strips (about 60 μm) allows the track positions to be measured with a resolution of approximately 15 μm, making possible a high precision determination of a track's impact parameter.

3.2.1.2 Vertex Time Projection Chamber (VTX)

The VTX [51], which sits just outside the SVX and extends to a radius of 21 cm, is a time projection chamber designed to measure a particle's path in the r-z plane. It consists of 28 octagonal modules, segmented into two drift regions separated by a cathode plane. Adjacent modules are rotated by 11.25° to help resolve ambiguities in matching tracking information in

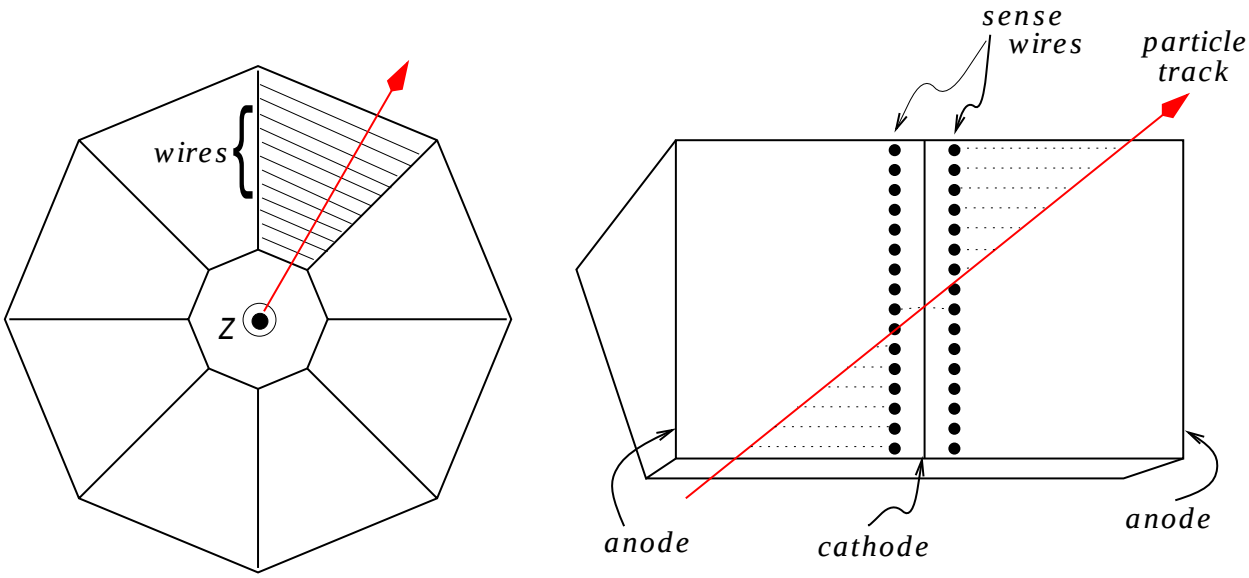

Figure 3.5 End and side views of one VTX drift module.

the VTX to information in the CTC, described below. Figure 3.5 illustrates the design of one VTX module. The drift regions are filled with a 50% mixture of argon and ethane gas and voltage is applied to the cathode plane to create an electric field parallel to the beamline. Near the cathode grid separating the two drift regions, sense wires are strung in the r-ϕ plane. A charged particle traversing the drift region ionizes the gas, causing electrons to drift along the direction of the electric field and cascade near the sense wires. The elapsed time between the beam crossing and the signal detected at the wires, known as the drift time, allows a determination of the drift distance for the charge in the z direction. From drift time measurements from wires at several different radii, a picture of the particle's trajectory in the r-z projection can be constructed.

3.2.1.3 Central Tracking Chamber (CTC)

The CTC [52] provides the starting point for tracking at CDF. Tracks are first reconstructed in the CTC in three dimensions, and then information from the VTX and SVX' is matched to the CTC tracking information to improve measurements of the track impact parameter and z position. In addition, the CTC measures the momentum of all tracks. The CTC consists of a cylindrical drift volume concentric to the beamline and 3.214 m long. Its active volume begins just outside the volume of the VTX at 31 cm and extends to a radius of 132 cm. The radial size

and length of the CTC restricts it to tracking particles primarily in the pseudorapidity range of $|\eta| < 1$, which is designated as the central region at CDF. Therefore all tracking in this analysis is confined to that region as well. Below we review a few details of the CTC pertinent to this analysis:

The CTC is a multi-wire drift chamber with 6,156 sense wires surrounded by a volume of argon-ethane gas. A charged particle passing through the active volume of the CTC ionizes the gas, causing charge to drift to the sense wires. Electric signals generated by charge drifting onto the sense wires allow measurements of the particle's trajectory. The sense wires are arranged in nine concentric "superlayers," each divided into a drift cell containing a number of anode sense wires surrounded by cathode field wires that create the drift field for the cell. Five of the superlayers, known as axial superlayers, contain drift cells with 12 sense wires that run parallel to the beam direction. The other four superlayers, known as stereo superlayers, contain 6 sense wires tilted by ±3° relative to the beam axis to allow a measurement of the *z* position of the track in those layers. Axial and stereo superlayers alternate. The drift cells in each superlayer are tilted by 45° in the *r*-ϕ plane relative to the radial direction to best match the Lorentz angle of the drift electrons moving in the combined electric and magnetic fields in the CTC. Figure 3.6 shows the arrangement of superlayers and drift cells in the CTC.

A particle passing completely through the CTC can have its trajectory sampled at a maximum of 82 locations, although chamber inefficiencies may prevent all tracks from being measured at that many points. Individual hits are measured in the CTC with a resolution of approximately 200 μm in the *r*-ϕ plane and 4 mm in the *r*-*z* projection. The resolution of the momentum measurement in the CTC depends on the momentum of the particle being measured. High momentum particles do not curve as much as lower momentum particles, making it difficult to measure momentum using the amount of curvature in the particle track. For the CTC alone, the momentum resolution can be parameterized as $\delta p_T / p_T = 0.002 \times p_T$, with p_T given in GeV/c. This resolution improves by around a factor of two if SVX' information is also used.

3.2.2 Calorimeter

The main purpose of the calorimeter is to measure energy deposited by particles. Unlike the tracking systems, the calorimeter is capable of detecting both charged and neutral stable particles, except for neutrinos that don't significantly interact in any detector system at CDF and muons that penetrate the calorimeter and are detected by the muon systems described in Section

3.2.3. The basic technique for measuring energy depositions in the calorimeter is to use a stack of alternating layers of dense material and active detector elements, such as lead or steel alternating with scintillator. Particles entering the dense material interact either through the electromagnetic or strong force and cause a particle cascade. The active detector material samples the number of particles generated as the cascade develops, and the energy of the incident particle is derived from the size of the cascade.

The calorimeter at CDF is divided into a thinner, electromagnetic portion at smaller radius, surrounded by a thicker, hadronic portion. Although CDF has calorimeter coverage out to $|\eta| < 4$, only calorimeters in the central region—the central electromagnetic calorimeter (CEM) and the central and wall hadronic calorimeters (CHA and WHA)—are used for this analysis. The calorimeter is contained in a cylindrical volume just outside the radius of the CTC, and is divided into two halves in the *z* direction. Each half consists of 24 wedges subtending 15° in ϕ. Each wedge contains both an electromagnetic and a hadronic module. Wedges are subdivided into towers with a width of about 0.1 in η, arranged in a projective geometry so that they point back towards the center of the CDF detector. Figure 3.7 shows the structure of a calorimeter wedge.

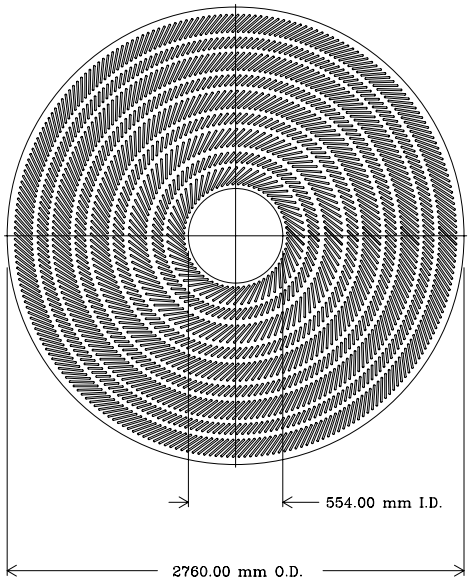

Figure 3.6 An end view of the CTC endplate, showing the superlayer structure. Each slot shown contained one drift cell of 12 sense wires for axial superlayer and 6 sense wires for stereo superlayers.

For this analysis, the calorimeter is used primarily to provide trigger signals for detecting electrons. Calorimeter information is also used offline to perform electron identification with more stringent criteria than used in the trigger. Although the calorimeter is also frequently used to identify jets for QCD analyses, that function is not used for this analysis.

3.2.2.1 Central Electromagnetic Calorimeter (CEM)

The CEM [53] is constructed of alternating layers of 1/8" thick lead plates and 5 mm thick layers of polystyrene scintillator. The CEM is optimized to detect electromagnetic showers generated by electrons and photons. An electron entering the CEM initiates an electromagnetic cascade in

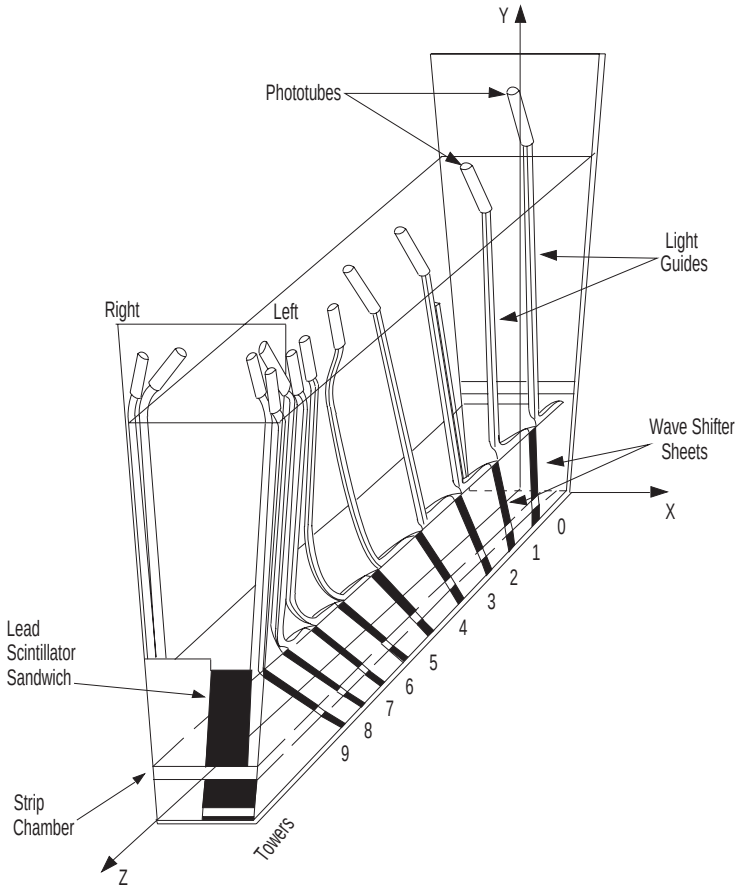

Figure 3.7 The layout of one calorimeter wedge, with the EM portion shown in detail. The EM portion of the calorimeter is closest to the interaction point. The hadronic portion (not shown in detail) has a similar structure and occupies most of the rest of the wedge.

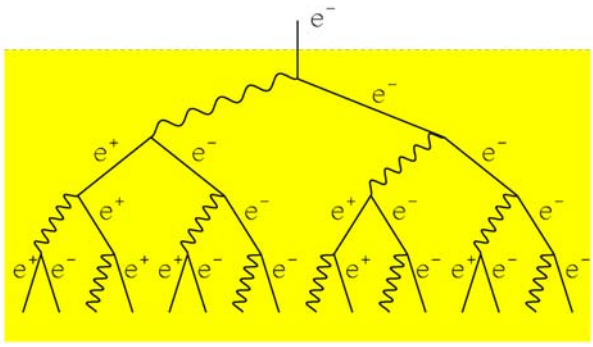

Figure 3.8 A schematic diagram of the development of an electromagnetic shower.

the lead plates by emitting a photon through bremsstrahlung. The emitted photon converts into an e^+e^- pair that then emits photons through bremsstrahlung, and so on. Photons entering the calorimeter begin the cascade by pair converting first, but otherwise the process is identical to the shower of an electron. The shower process will continue until the electrons and positrons reach approximately 10 MeV, at which point they loose energy through ionization rather than bremsstrahlung. An electromagnetic cascade is shown schematically in Figure 3.8. When electrons generated by the electromagnetic cascade enter the scintillator layers, they produce light, which is collected by a number of phototubes. The total amount of light collected is proportional to the energy of the electron or photon that entered the calorimeter. The depth of the CEM is chosen so that most electron and photon showers will be completely contained in its volume. For heavier particles, this is not true and the hadronic calorimeter, described next, comes into play.

In addition to the lead and scintillator layers, the CEM contains a gas proportional strip and wire drift chamber inserted at the depth where electron and photon showers on average attain their maximum energy deposition. These chambers, known as the central strip chambers (CES), are used to measure the location of the shower in the wedge to a higher precision than the segmentation of the calorimeter towers. These chambers contain a single layer of anode wires oriented parallel to the beam direction and cathode strips perpendicular to the beam direction. Charge collected on the wires and strips allows a ϕ-z measurement of the position and extent of the shower in the CEM. Such a measurement can be used to match an electromagnetic shower to a charged track reconstructed in the tracking system for electron identification.

The signature of an electron in the CEM is a cluster of energy with a charged track pointing to the position of the shower (as measured by the CES). Other information, such as the shape of the shower or whether the shower is well contained within the CEM, can also be used to discriminate true electrons from backgrounds. Conversely, a photon signature is typically defined to be a cluster of electromagnetic energy with no track pointing towards it.

3.2.2.2 Central and Wall Hadronic Calorimeters (CHA and WHA)

Particles heavier than electrons typically will not loose a significant fraction of their energy via bremsstrahlung within the CEM. However, hadrons also interact inelastically with nuclei via the strong force, and over longer distance scales will generate hadronic showers similar to the electromagnetic cascades described above. The CHA and WHA [54] are designed to contain the hadronic showers initiated by hadrons penetrating through the CEM. The CHA and WHA are constructed of 2.5 cm thick steel plates sandwiching 1 cm think plastic scintillator. The basic principles of operation are very similar to the CEM. For this analysis, the hadronic calorimeters are mainly used to veto electron candidates accompanied by significant deposition of energy in the hadronic portion of the calorimeter.

3.2.3 Muon Systems

Muons are too heavy to deposit much energy in the electromagnetic calorimeter through bremsstrahlung and lack the nuclear interaction necessary to shower hadronically in the hadronic calorimeters. Therefore, muons typically pass completely through the calorimeter system, allowing them to be detected using drift chambers placed beyond the calorimeter. CDF incorporates several muon detector systems based on this scheme. For this analysis, we use only muon systems that cover the central region at CDF, namely the central muon chambers (CMU) [55] and the central muon upgrade chambers (CMP). For this analysis, a muon is identified as any charged particle, reconstructed in the CTC that penetrates the calorimeter and the additional shielding of the CMP to leave hits in both the CMU and the CMP chambers. These so-called CMUP muons are used both as a trigger to define our sample and also used to select events during the offline analysis.

3.2.3.1 Central Muon Chambers (CMU)

The CMU is located within the calorimeter wedge just beyond the radius of the CHA, at R = 3.47 m. CMU chambers cover 12.6° out of the 15° in a wedge using 3 modules composed of four-by-four arrays of single-wire drift chambers with projective geometry, as shown in Figure 3.9. The wire in each chamber is oriented parallel to the beam line. The CMU chambers measure the trajectory of a muon at four points in the r-ϕ plane using the same principles as other drift chambers, such as the CTC. Track segments reconstructed in the CMU are referred to as "stubs." The CMU also measures the z position of hits left by muons using charge division. The r-ϕ resolution of the CMU is approximately 250 μm, while the z resolution is roughly 1.5 mm. In addition, the CMU can make a rough estimate of the muon p_T, despite being located outside the CDF solenoidal magnetic field, using the angle of the reconstructed track relative to the radial line. This is because infinite momentum tracks will emerge from the solenoidal field along the radial line, while lower momentum tracks are bent by the magnetic field and come out at an angle. Although the final muon momentum measurement is provided by the CTC, the rough estimate from the CMU is useful for triggering purposes.

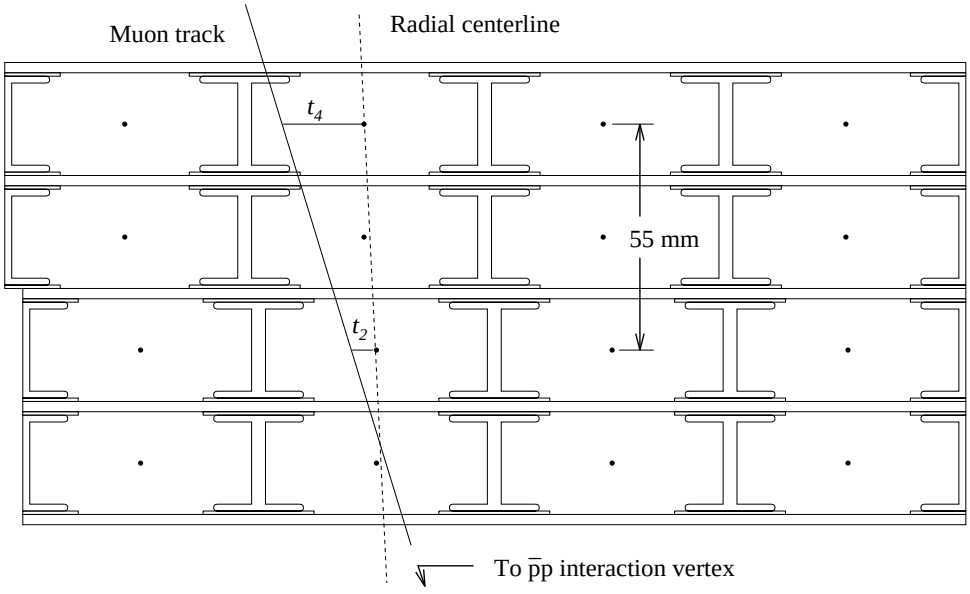

Figure 3.9 An end view (x-y plane) of one CMU module. Each rectangular drift cell contains one anode wire. Three modules like these provide coverage for each calorimeter wedge.

3.2.3.2 Central Muon Upgrade (CMP)

The CMP was added to the CDF muon system prior to Run 1 because of the significant background in the CMU of charged hadrons penetrating the CHA before losing all their energy and leaving tracks in the CMU. This background is called "punch-through." The CMP chambers virtually eliminate punch-through by being located behind an additional 60 cm of steel shielding. The CMP consists of a set of rectangular, single-wire drift chambers, four layers thick arranged in a roughly rectangular "box" around the CDF detector, as shown in Figure 3.10. Because of its rectangular geometry, and irregularities caused by the structural obstructions from

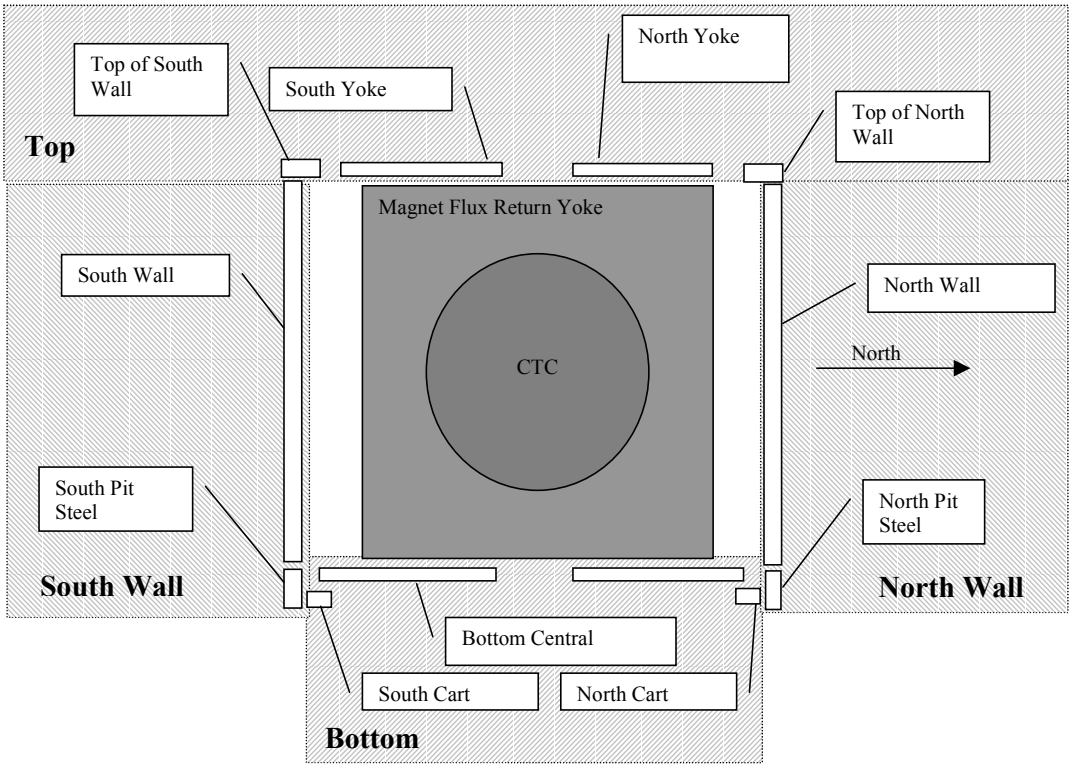

Figure 3.10 An end view of the CDF detector with the layout of the CMP chambers (denoted by white rectangles) shown. The CMP chambers are arranged in a rectangular geometry consisting of roughly four planes of chambers: the "North Wall," "Top," "South Wall," and "Bottom." Each of the four planes are further subdivided by the necessity of avoiding structural obstructions around the detector. Furthermore, supports at the bottom of the detector require that some of the chambers in the "Bottom" region be shorter than others, as reflected in the coverage map shown in Figure 3.11.

other parts of the detector, the CMP does not cover the entire acceptance of the CMU, as shown in Figure 3.11. The wire in each CMP chamber is oriented along the beamline direction. Since the CMP is not instrumented for charge division, it provides only an r-ϕ measurement of the muon trajectory at four points. The main purpose of the CMP is to confirm the validity of muon hits recorded in the CMU. However, because of the excess shielding, a muon must have a transverse momentum of approximately 2.8 GeV/c or more to penetrate all the way to the CMP and be detected. This sets a lower p_T threshold on the ability of CDF to cleanly detect muons.

3.2.4 Trigger

With proton-antiproton bunch crossings occurring every 3.5 μs at CDF, with a sufficient luminosity such that on average one to two $p\bar{p}$ interactions occur per bunch crossing, approximately 300,000 high energy collisions occur at CDF every second. However, the data acquisition hardware is only capable of recording detector data to tape at a rate of approximately five events per second. If data were written to tape indiscriminately as fast as possible, QCD interactions involving only light quarks (u, d, s) would dominate, preventing the accumulation of significant statistics for any other processes. Therefore, it is necessary to employ a trigger system at CDF capable of ensuring the mix of data collected allows for the study of a wide variety of high-energy physics processes regardless of the various rates of those processes.

There are two main strategies for accomplishing this goal: One approach is to give preference to certain detector signatures that herald the presence of a low-rate process in the event. Anytime such a signature is registered, the detector is always read-out, guaranteeing that the rare processes associated with this signature are collected with the maximum possible statistics. Alternatively, one can identify signatures for relatively common physics processes and require that such events be collected at a scaled down rate. For example, such a signature may only be read out once every 100 or 1000 times it is detected. The scale factor used to reduce the data rate for these types of triggers is known as a *prescale*. In practice, both strategies are used in combination for various physics signals in an attempt to strike the delicate balance that will satisfy the demanding needs of the diverse physics program at CDF.

The CDF trigger is implemented in the form of a three-level system. Each level in the trigger uses some amount of detector information to reduce the rate of data flow coming from the detector in manageable steps. Because each stage has a lower output rate than the preceding stage, more of the detector can be read out to facilitate the trigger decision. The first stage,

which must contend with the full ~300 kHz event rate can only use the fastest of detector signals to make its decision, while the third level of the trigger has the leisure to consider all the detector data and run a full but fast computer analysis of each event it considers.

3.2.4.1 Level 1

The Level 1 trigger [56] at CDF must make a decision on whether to accept (and pass on to the next trigger level) or reject an event within the 3.5 μs bunch crossing time. Therefore this trigger must be implemented in hardware using fast outputs from the detector systems and high-speed

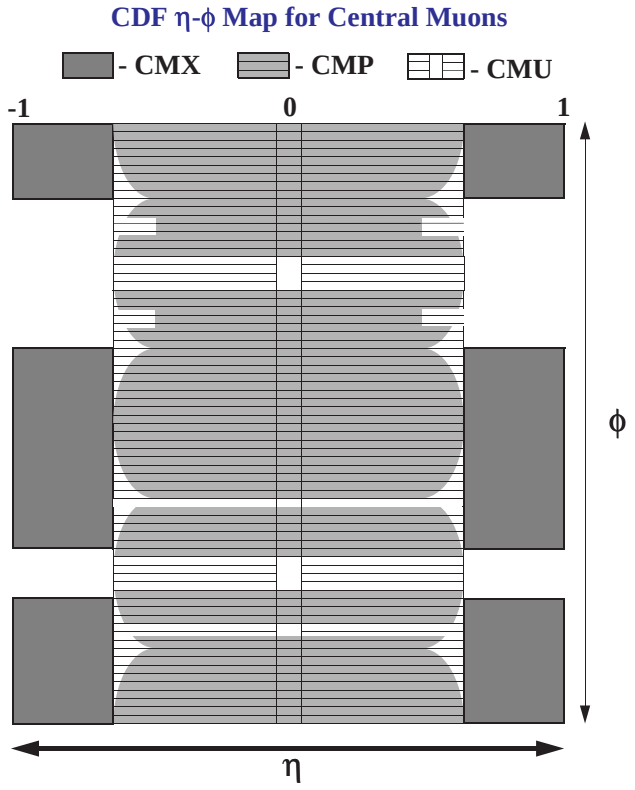

Figure 3.11 The η–ϕ coverage of the muon chambers. The CMP coverage is so irregular because of the projection of the rectangular CMP structure onto the cylindrical coordinate system at CDF.

electronics. The main Level 1 triggers for this analysis rely on energy deposition in the calorimeter and hits in the muon system. At Level 1, it is not possible to make use of tracking information or correlations between signals in different detectors. By carefully tuning trigger thresholds, the Level 1 trigger achieves a reduction of the rate of data flow from ~300 kHz rate of beam crossings to an output rate of ~1 kHz of events passed on to Level 2.

The main CDF Level 1 triggers that contribute to this analysis are specified below:

- **L1_CALORIMETER** (no prescale): This trigger accepts events with single tower calorimeter energy deposition above a given threshold. For CEM, the threshold is 8 GeV.

- **CMU_CMP_6PT0_HTDC_BBC** (no prescale): This trigger requires a $p_T > 6$ GeV/c central muon stub (as measured in the CMU) with strict requirements on the hit times to ensure that the muon was produced in association with a bunch crossing.

3.2.4.2 Level 2

Because Level 2 [56] sees a much lower input rate (~ 1 kHz) than Level 1, it can take more time to make a trigger decision and therefore can access more sophisticated detector information. For example, at Level 2, energy deposition in individual calorimeter towers can be clustered together to form jets. One component of the Level 2 trigger that plays a major role in this analysis is the central fast tracker (CFT) [57], a fast, hardware tracker that uses information from only axial CTC layers to detect tracks by recognizing pre-programmed hit patterns. The CFT provides a rough momentum estimate of the track with a resolution somewhat lower than the CTC p_T resolution, $\delta p_T / p_T = 0.035 \times p_T$. For tracks above p_T of 10 GeV/c, the CFT is over 90% efficient. The CFT also allows matching of tracks to calorimeter clusters and muon stubs, providing a powerful trigger for electrons and muons. The CFT plays a central role in defining triggers for this analysis.

Another relevant feature of Level 2 is its ability to implement *dynamic prescales*. A dynamic prescale adjusts itself depending on the output rate of Level 2. During high luminosity periods at the beginning of the store, the overall event rate coming into Level 2 may be so high that certain triggers must be heavily prescaled in order to maintain an acceptable rate out of Level 2. However, as the store progresses and the luminosity decreases, the Level 2 output drops and it is possible to regain some of this unused bandwidth by decreasing some of the trigger prescales.

Dynamic prescales work somewhat like a feedback mechanism to keep the output rate of Level 2 constant over a wide range of input rates or luminosities. Dynamic prescales are defined in three bins of Level 2 output rate: 0-12 Hz, 12-22 Hz, and 22+ Hz. For specific triggers listed below, the prescale for each bin is listed in order of low to high.

The main triggers that contribute to this analysis at Level 2 involve identifying electrons or muons by matching calorimeter or muon system information with tracks found by the CFT. The specific Level 2 CDF triggers are listed below:

- **CEM_8_CFT_7_5_XCES** (prescale = 1:2:8): This trigger requires a CEM energy cluster of 8 GeV, a matching CES cluster, and a 7.5 GeV/c CFT track pointing at both.
- **CEM_8_CFT_7_5** (prescale = 4:16:64): This trigger is the same as the above except without the CES requirement.
- **CMUP_CFT_7_5_5DEG** (prescale = 1:2:8): This trigger requires a 7.5 GeV/c CFT track within 5° in ϕ of both a CMU and a CMP stub.

3.2.4.3 Level 3

Level 3 takes the 20 Hz or less output rate from Level 2 and reduces it to the ~5 Hz rate required for output to tape. Because of the much larger time available for Level 3 to make a trigger decision, Level 3 can be implemented in software. Level 3 employs the same version of the analysis code used to reconstruct events offline, except that the calibration constants available to the Level 3 code are not as refined. At Level 3, a full reconstruction of tracks can be run as well as improved electron and muon identification algorithms. Because of corrections applied to tracking and calorimeter quantities, energy and momentum thresholds are better defined in Level 3 than at previous levels. When Level 3 accepts an event, it triggers a full readout of the detector and all detector data for that event is recorded to 8 mm tape for further analysis offline.

The Level 3 triggers contributing to this analysis deal with further refinements of the electron and muon triggers used at Level 2. The specific Level 3 CDF triggers are listed below:

- **ELEB_CEM_8_6** (no prescale): This trigger requires a reconstructed electron with clustered CEM $E_T > 7.5$ GeV matched to a track with $p_T > 6$ GeV/c.

- **MUOB_CMU_CMP_8** (no prescale): This trigger requires a track of $p_T > 8$ GeV/c matched to both CMU and CMP stubs.

Events passing the above triggers define the data sample on which this analysis is performed.

Chapter 4

Data Sample

Although $b\bar{b}$ production occurs at a high rate at the Tevatron, it is only a small fraction of the total activity generated by $p\bar{p}$ collisions at 1.8 TeV. One needs some criteria to reduce the volume of data to analyze and increase the *b* quark content in the data. For this analysis, we use electrons and muons produced by semileptonic *B* decays and secondary vertices produced by the decays of long-lived, weakly decaying *B* hadrons to help identify events resulting from $b\bar{b}$ production.

This analysis starts with the same basic data sample used for the measurement of time dependent *B* mixing at CDF [58]. This sample, characterized by the presence of at least one electron or muon plus at least one reconstructed secondary decay vertex, is known as the BVTX sample because the algorithm used to reconstruct, or tag, secondary vertices is known as the BVTX algorithm. For more information on the BVTX algorithm, consult Section 4.2. The BVTX sample was created from data collected in Run 1b using a variety of electron and muon triggers by requiring the presence of at least one displaced secondary vertex tag (BVTX tag). The original BVTX sample consists of over 480,000 electron-triggered events and over 430,000 muon-triggered events that have been collected together from the greater CDF data set and stored on 8mm tape.

Our basic event selection strategy is as follows: We start with the events from the BVTX sample, which is a loosely selected sample of events containing at least one electron or muon and at least one secondary vertex. We impose stricter cuts on the trigger lepton to ensure all *B* hadron decays producing the trigger lepton are kinematically similar. Next, we reprocess the data using a modified version of the BVTX tagging algorithm and demand that each selected event contain at least two secondary vertex tags. No requirements are placed on the properties of the tags at this point. Finally, we impose a few additional sample selection criteria to minimize

the contribution from background sources. Further details of the sample selection process are discussed below:

4.1 Trigger Lepton Selection

Because the rate of QCD interactions at the Tevatron involving only light quarks (u, d, s) is much larger than the production rate of events containing b quarks, any attempts to study b production without a trigger strategy to preferentially select b events over light-quark events would fail due to lack of sufficient statistics. During Run 1, the trigger took advantage of the roughly 20% branching fraction of B hadrons to electrons or muons to define data samples with enhanced b content. The BVTX sample uses these triggers as its starting point.

However, a number of electron and muon triggers were employed during Run 1b at CDF, with different lepton p_T thresholds, different lepton η requirements, and different numbers of leptons required (single-lepton versus dilepton triggers). The BVTX data sample is assembled from a variety of single electron, single muon, and dimuon triggers, each with different p_T and η ranges, as well as different dynamic or fixed prescales. Because of the link between the kinematics of the lepton and the kinematics of the B hadron that produced the lepton, the BVTX sample contains a non-trivial mix of subsamples with different B kinematics. The comparison of our B hadron correlation measurements to theoretical predictions is greatly simplified by restricting our analysis to triggers with similar geometric and kinematic acceptances.

4.1.1 Trigger Requirements

Which particular lepton trigger is used can have a significant impact on the kinematics of the B hadrons in the sample, both through the lepton p_T and η ranges allowed by the trigger, and also because inefficiencies in the trigger near the p_T threshold can sculpt the B hadron distributions. For this analysis, the triggers that have the biggest effect on the final sample kinematics are the Level 2 lepton triggers. Therefore, we require events selected from the BVTX sample to have been collected from specific Level 2 triggers by examining the trigger information in each event record and only accepting events that fired the specified Level 2 triggers. The Level 2 triggers required for this analysis are given below:

- CEM8_CFT_7_5_XCES and CEM8_CFT_7_5: The only difference between these two triggers is the CES matching requirement placed on the XCES trigger path, which improves the purity of that trigger sample. The effective luminosity, of this sample is roughly 74 pb^{-1}.

- CMUP_CFT_7_5_5DEG: The effective luminosity of this sample is approximately 48.4 pb^{-1}.

Details regarding these triggers are given in Section 3.2.4.2. This technique accounts for not only whether the event satisfied the trigger kinematically, but also whether the event was accepted or rejected based on the prescale factor. This is an important consideration because, for each of the above triggers, there is a similar trigger that has a higher p_T threshold and lower prescale factor. Events passing these higher threshold triggers will always satisfy the lower threshold triggers unless the prescale of the lower threshold trigger causes the event to be rejected. If not correctly accounted for, this could lead to an incorrect enhancement in the higher p_T end of the B spectrum.

These triggers were chosen to yield similar B hadron p_T spectra for both electron and muon data. However, because of slight differences in η acceptance, and to allow examination of any biases in the trigger requirements, we separate our data sample into an electron-triggered sample and a muon-triggered sample and perform the analysis on two samples separately.

4.1.2 Offline Requirements

In addition to the trigger requirements listed above, we apply certain requirements on trigger electrons and muons offline, after the full reconstruction code has been run on the event. These requirements are very similar to the requirements applied in the Level 3 trigger, although better calibration information can cause events that pass the Level 3 trigger to fail the cuts given below. In combination with the trigger requirements above, these cuts determine a sample of semileptonic B decays with fairly uniform kinematic properties for both electrons and muons.

4.1.2.1 Electron Requirements

Each event in the electron sample is required to have at least one electron satisfying the following cuts:

- E_T > 7.5 GeV: This cut, based on the energy measured in the calorimeter for the electron using offline calibrations which have a better resolution than quantities used in the trigger, corrects for any energy mismeasurements made in the trigger.

- p_T > 6.0 GeV/c: This cut is based on the electron track momentum as calculated from the fully reconstructed track. The momentum resolution from offline tracking is almost an order of magnitude better than the CFT measurement, so this cut creates a better defined threshold for electron acceptance. The p_T cut is allowed to be lower than the E_T cut because electrons frequently lose energy to bremsstrahlung passing through the SVX' and inner wall of the CTC. The momentum lost is not registered by the CTC, but is frequently detected in the calorimeter.

- $E_{Had}/E_{EM} < 0.04$: The quantity used in this cut is the ratio of the energy deposited by the electron candidate in the electromagnetic calorimeter to the energy deposited in the hadronic calorimeter. Because the CEM is designed to contain fully showers initiated by electrons at these energies, electrons are not expected to deposit significant energy in the hadronic calorimeter. In other words, the leakage of the EM shower into the hadronic calorimeter is expected to be small.

- $L_{shr} < 0.2$: The variable L_{shr}, where L here denotes *lateral*, parameterizes the transverse energy sharing of the electron candidate shower between the seed tower (the CEM tower with the highest energy deposition) and the two nearest towers within the same calorimeter wedge. Electron showers should be well contained within one calorimeter tower. Information about the primary interaction vertex and the CES position of the shower are used to calculate the energy sharing corrected for the electron trajectory into the wedge, which is then compared to expectations from electron test-beam data to determine a value for L_{shr}. A high value of L_{shr} indicates that the shower spread is too wide, inconsistent with the electron hypothesis.

- $\chi^2_{strip} < 10$, $\chi^2_{wire} < 10$: These χ^2 variables parameterize how well the transverse shower development as measured by the CES wires and strips match expectations based on test-beam data. There are eleven strips or wires used in the shower shape fit and the fit contains two free parameters, energy and position. Therefore, this χ^2 has nine degrees of freedom. These χ^2 variables are particularly useful for separating the background from

hadrons and π^0's, which decay to two photons, because the shape of these showers does not match the shape expected for electrons. In addition, the fine position resolution of the CES makes it possible to veto cases in which the energy deposit associated with the electron candidate appear to come from two separate particles.

- $|\Delta z(track - CES)| < 3$ cm, $|\Delta x(track - CES)| < 1.5$ cm: These matching variables indicate how well the electron track candidate matches the center of the shower as measured in the CES

In addition, the electron is required not to be consistent with coming from the conversion of a high-energy photon into an e^+e^- pair. This requirement is enforced by searching for an oppositely charged track that forms a good vertex with the electron in question. To be considered a valid conversion pair, the electron track and the conversion pair candidate track must be parallel within detector resolution at the point of intersection and must meet the following requirements:

- r-ϕ separation at point of tangency < 0.2 cm
- Difference in cotθ for the two tracks < 0.03
- z mismatch at point of tangency < 2.0 cm
- Radius at which the conversion took place between -5 cm and 50 cm
- $\Delta\phi$ at radius of conversion < 0.01
- Pointing residual at origin < 0.2 cm

See Section 4.4.4 for a discussion of the residual contribution from conversions not removed by these cuts.

4.1.2.2 Muon Requirements

Each event in the muon sample is required to have one muon satisfying the following cuts:

- $p_T > 6.0$ GeV/c: The momentum for this cut is measured using quantities from full offline track reconstruction. As in the electron case above, this p_T has a far sharper

resolution than the p_T measured in the CFT, and applying this cut gives a better defined muon momentum threshold.

- $\chi^2_x(CMU) < 9,\ \chi^2_z(CMU) < 12$: These χ^2 variables parameterize the quality of the matching between the muon track candidate and the muon stub measured in the CMU. The value of these χ^2 variables is determined from the expected error in track extrapolation to the CMU based on multiple scattering of the muon as it penetrates the electromagnetic and hadronic calorimeter. In each of these χ^2 variables, there is one degree of freedom.

- $\chi^2_x(CMP) < 9$: This χ^2 variable is similar to the ones above except that it is calculated for the extrapolation of the muon track candidate to a stub measured in the CMP. Note that the CMP stub position is not measured along the *z*-axis. The CMP matching χ^2 has one degree of freedom.

- $E_{Had} > 0.5$ GeV: Although the muon passes through the calorimeters without loosing significant energy, it does deposit a small amount of energy. As particles that deposit energy in the calorimeter primarily through ionization, muons are known as minimum ionizing particles. The ionization energy loss for minimum ionizing particles can be calculated, and the cut value above reflects expectations on muon energy loss in the calorimeter based on such calculations.

4.2 Secondary Vertex Tagging

In this analysis, *B* hadrons are identified by reconstructing their decay vertices. Because *B* hadrons have a long lifetime, their decay points, know as *secondary vertices*, are separated from the $p\bar{p}$ interaction point, known as the *primary vertex*. Using precision tracking information from the SVX', *B* hadron secondary vertices can be separated from the primary vertex. The algorithm used to identify and reconstruct these displaced secondary vertices is a slight variation on the algorithm used for the time dependent *B* mixing analysis [58], known as the BVTX algorithm. The BVTX algorithm itself is based on the SECVTX algorithm used to tag *b*-jets in the top analysis [29].

Secondary vertex finding begins by first locating the primary interaction vertex for the event using precision tracking information from the SVX' and VTX. Next the tracks in the event passing certain quality cuts are grouped into jets using a cone-based clustering algorithm. Each jet is then searched for the presence of one or more secondary vertices displaced from the primary. The secondary vertex finding is done in two passes for each jet. The first pass finds secondary vertices containing at least three tracks. When the first pass fails to find any more secondary vertices in a jet, a second pass is attempted in which the individual track cuts are made more stringent and two-track secondary vertices are accepted. Each secondary vertex found is required to be significantly displaced from the primary and not to be consistent with the decay of a K_S^0 or Λ.

Below we list the details of primary vertex finding, track selection, jet clustering, and secondary vertex finding.

4.2.1 Primary Vertex Finding

The first step to locating displaced secondary vertices is to find the primary interaction vertex. The primary vertex finding algorithm used for this analysis, known as VXPRIM, is a component of the CDF Run 1 offline software [59]. VXPRIM calculates a primary vertex location for each event by attempting to constrain the tracks in the event with $p_T > 400$ MeV/c to a common vertex. Trigger lepton tracks are explicitly excluded from this fit, however, because it is assumed that these leptons were produced at displaced secondary vertices. The average position of the CDF interaction point, known as the beamline position, is determined from the data for each run (an amount of data equal to approximately one to ten hours of continuous data taking in which accelerator and detector conditions remained reasonably stable). VXPRIM uses this run-averaged beamline position, with a beam spot size of 25 μm, and the best estimate for the z vertex figured from VTX data as a starting estimate for the primary vertex location. The fit then proceeds iteratively in two stages. In the first stage, called the non-steering stage, VXPRIM attempts to constrain all tracks to a common vertex. Tracks that contribute too much to the fit χ^2 are removed, and the fit is attempted again until no tracks are left that contribute too much. In the second stage, the track parameters are allowed to float within their errors as VXPRIM attempts to find a better fit. In the end, the fit is required to include at least 5 tracks and not to have an error greater than the beam spot. If a suitably high-quality primary vertex cannot be

found, VXPRIM uses the run averaged beamline position and best VTX z vertex instead, and assigns the primary vertex an uncertainty equal to the beam spot size estimate in x and y, 25 μm.

The danger in using an event-by-event calculation for the primary vertex rather than the run-averaged beamline estimate is that in $b\bar{b}$ events with low multiplicity, the primary vertex found by VXPRIM can be biased towards the location of one of the secondary vertices in the event. In this analysis, any bias in the calculated primary vertex location towards the location of a secondary vertex could lead to an inefficiency in the secondary vertex finding algorithm by reducing the significance of the separation between the primary and secondary vertices. To check that this is not the case, we compare the position of the vertex calculated from the estimate of the beamline position to the VXPRIM vertex. The beamline position is calculated from the run-averaged position of the primary vertex for many different events. Because of the tilt of the beam through the CDF detector, it is necessary to estimate the z position for the primary interaction vertex, which we taken to be the closest VTX z vertex to the highest p_T trigger lepton in the event. Figure 4.1 shows the difference in the separation in the x-y plane between the

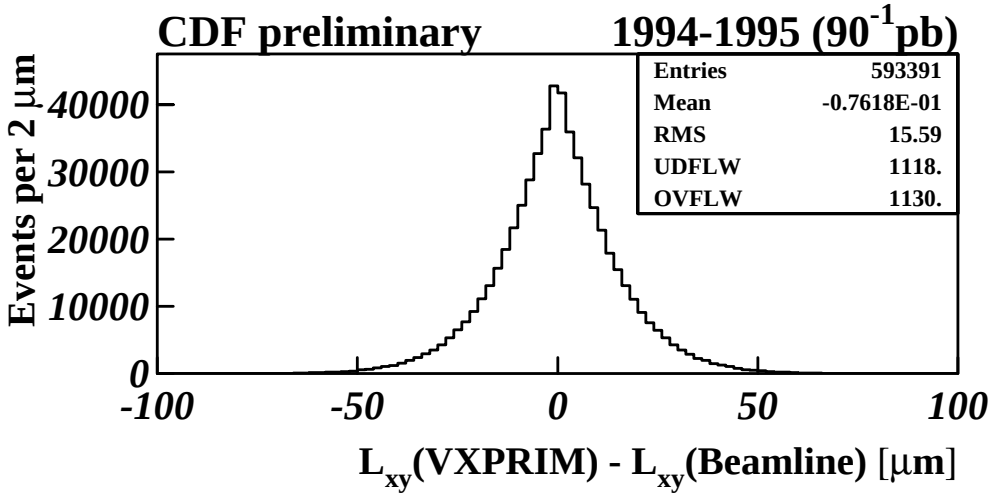

Figure 4.1 The difference between L_{xy} measured with respect to the VXPRIM primary vertex versus the beamline primary vertex estimate. The average difference is only slightly negative indicating perhaps only a very negligible bias of the VXPRIM vertex to be closer to the secondary vertex than the beamline vertex. In 97% of the events the VXPRIM measurement of the primary vertex agrees with the beamline estimate to within ±40 μm. For 55% of the events, the agreement is better than ±10 μm.

secondary vertex and the primary vertex found by VXPRIM versus the beamline estimate. This plot shows the two vertices tend to agree to better than 40 μm 97% of the time with no significant bias. However, because the VXPRIM vertex often has a smaller uncertainty than the beam spot size, the significance of track impact parameters and vertex displacements tend to be greater. This is far more important in allowing secondary vertex identification than a small potential bias in the primary vertex position towards the secondary vertex. Thus, for this analysis, VXPRIM is used to calculate the primary vertex for each event.

4.2.2 Track-Based Jet Clustering

Tracks are grouped into jets using the cone-based jet clustering algorithm. This algorithm groups tracks into jets based on the distance between the tracks in angular (η-ϕ) space. To be eligible for clustering in a jet, a track must satisfy the following requirements:

- The track must have at least five hits in each of at least two CTC axial superlayers, and at least two hits in each of at least two CTC stereo superlayers.
- The track must be reconstructed in three-dimensions using CTC tracking information. Some tracks with insufficient *z* information are only reconstructed in two-dimensions and thus not eligible for jet clustering.
- The *z* coordinate of the track at its point of closest approach to the primary vertex must satisfy $|\Delta z(track - primary)| < 5$ cm. Sometimes, more than one $p\bar{p}$ interaction will occur during a beam crossing, and this requirement reduces the contribution of tracks from additional $p\bar{p}$ interactions.
- The track impact parameter must satisfy $|D_0| < 1$ cm. The sign of the impact parameter, D_0 contains information about the way the track is curving, relative to the primary interaction point, and the track charge. This information is not useful for this cut, and so only the magnitude of the impact parameter is used.
- The track transverse momentum p_T must be greater than 0.4 GeV/c.

- The track must pass through the outer layer of the CTC before exiting the CTC detector volume. This requirement is satisfied by demanding that the z of the track extrapolated to the radius of the outer CTC layer be less than 160 cm.

- Tracks used for jet clustering are not required to be reconstructed in the SVX'. Only CTC tracking information is used at this stage.

From the list of tracks passing the above requirements, tracks with $p_T > 1.0$ GeV/c are chosen as jet seeds. Next, the algorithm tries to merge pair of seeds within $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} = 1.0$ of each other into clusters. Once all possible seed track pairs are merged into clusters, clusters are merged together. Two clusters are considered successfully merged if, all tracks in the merged cluster are within $\Delta R = 1.0$ of the cluster axis. Cluster merging continues until there are no clusters left to merge. Finally, all remaining tracks meeting the above criteria that were not selected as seed tracks are assigned to the nearest cluster in ΔR.

4.2.3 Secondary Vertex Algorithm (BVTX)

The secondary vertex finding algorithm used for this analysis is similar to the one used in the time dependent B mixing analysis, with some modifications to improve its performance for finding more than one secondary vertex per event. The details of the BVTX algorithm as used in this analysis are discussed in detail below:

The BVTX algorithm searches each jet in the event for displaced vertices, starting with the highest E_T jet. Jet E_T is determined by summing the energy of all tracks in the jet, assuming a pion mass for each track, since no attempt is made to differentiate charged particle species for individual tracks. The search for displaced vertices proceeds in two steps or passes. The first pass finds vertices made by the intersection of three or more tracks. The algorithm continues to try to find Pass 1 vertices until no further Pass 1 vertices can be found.[1] Then it moves on to the second pass, which accepts two-track vertices, but makes tighter quality cuts on the tracks. When no more Pass 1 or Pass 2 vertices can be found, the algorithm moves on to search the next highest E_T jet. Figure 4.2 shows a flow chart of the BVTX algorithm.

[1] One major difference between the BVTX algorithm used for this analysis and the one used for the B mixing analysis is that in the B mixing analysis, after one vertex was found of any pass in any jet, the search for secondary vertices stopped. In this analysis, we continue looking until no further vertices can be found.

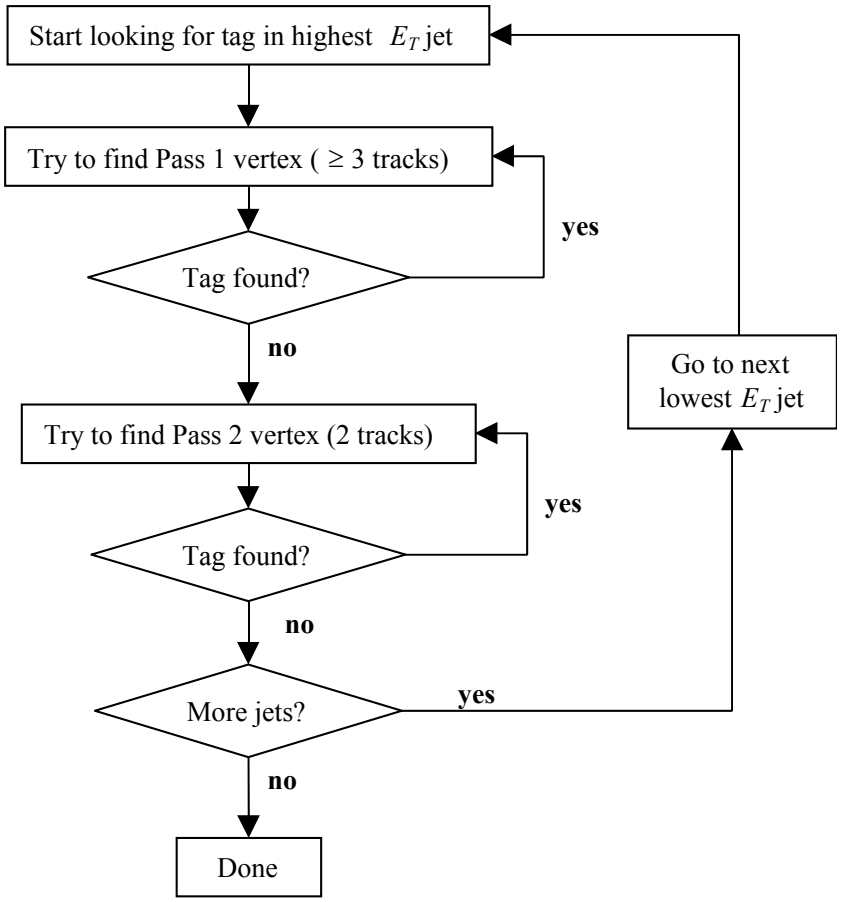

Figure 4.2 Flowchart of the BVTX algorithm.

The details of each pass are given below:

4.2.3.1 Track List

The first step in searching a particular jet for a displaced vertex is to make a list of tracks within a cone of $\Delta R = 1.0$ around the jet axis. This track list does not have to be the same as the list of tracks clustered to produce the jet, and tracks can appear in the list for more than one jet, although they can only be used to form one displaced vertex. To be placed on the list, the track must be reconstructed as a 3-D track in the CTC and matched to a track in the SVX'. It must also pass the standard CTC quality cuts: at least four hits in at least two axial superlayers and at least two hits in at least two stereo superlayers

Each track in the list is corrected for energy loss and multiple scattering from its passage through detector material, assuming a pion mass for all tracks. Each track is also required not to be consistent with being part of a "vee;" in other words, not consistent with originating at the decay vertex for either a $K_S^0 \to \pi\pi$ or $\Lambda \to p\pi$, produced at the primary vertex. To flag tracks as being part of a "vee," we consider all pairs of oppositely charged tracks where both tracks have $|D_0/\sigma| > 2.5$, and a pair mass calculated under either the $K_S \to \pi\pi$ or $\Lambda \to p\pi$ hypothesis within 0.1 GeV/c^2 of the K_S or Λ mass respectively. For the $\Lambda \to p\pi$ hypothesis, the track with the higher momentum in the pair is considered to be the proton, since the kinematics of a Λ decay favor that configuration. If a track pair meets the above requirements, the tracks are fit to a common vertex. The fit is required to converge with a χ^2, for one degree of freedom, less than 50. If the pair mass returned from the fit is within 0.01 GeV/c^2 of the K_S mass or within 0.006 GeV/c^2 of the Λ, and if the separation between the primary vertex and the fitted two track vertex is greater than 10σ along the direction of the two-track momentum vector, but less than 3σ perpendicular to it, the tracks are marked as a part of a "vee" and are not considered as a candidates for displaced secondary vertex tracks.

In addition, to be considered for inclusion in a secondary vertex, the tracks must meet the following requirements:

- The track must satisfy the standard CTC track cuts (described in Section 4.2.2).
- The track should not be flagged as part of a prompt "vee" (described above).
- The $\chi^2 / d.o.f$ of the SVX' track fit < 6.
- The track must include 3 or 4 SVX' hits, at least one of which must be defined as "high-quality", meaning that the hit comes from signals in a cluster of SVX' strips which is not shared with any other track, which uses information from no defective strips, and which includes no more than three strips.
- Track $p_T > 0.5$ GeV/c.

Finally, the list of tracks passing the above cuts is ordered from best to worst according to the number of good hits, the track p_T, and the impact parameter significance ($|D_0/\sigma|$). A diagram of this ordering scheme is shown in Figure 4.3. Tracks are assigned a "primary key" value as

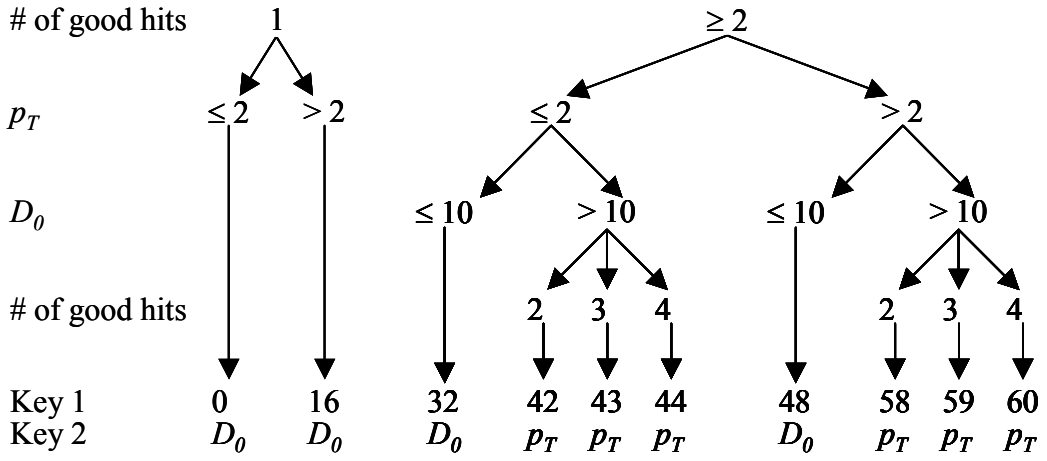

Figure 4.3 Ordering scheme for tracks used by the secondary vertexing algorithm. Tracks are assigned a primary key value based on the decision tree shown in the figure, and are arranged in order of descending key value. Tracks with equal values of the primary key are ordered according to the secondary key shown in the figure.

indicated in the diagram and are ordered with the highest key value first and the lowest last. Among tracks having equal primary key values, a second key, as indicated in Figure 4.3, is used to determine ordering. Tracks will be considered for inclusion in a secondary vertex in the order in which they appear in the list. The intention is to place tracks most likely to be from a displaced *B* hadron decay first in the list.

4.2.3.2 Pass 1

In the first pass through the track list, we attempt to find secondary vertices consisting of three or more tracks. To start, a list of "seed" tracks is made by considering every track from the track list which has $|D_0/\sigma| > 2.5$. For Pass 1 to proceed, at least three tracks must be placed on this list. However, if only one or two tracks qualify as seeds, additional tracks from the top of the track list that don't meet the criteria above will be allowed. If no tracks meet the seed criteria, Pass 1 automatically fails for this jet.

The next step is an attempt to fit pairs of seed tracks to a common vertex, starting with the first (best) tracks in the list. At least one of the two seed tracks is required to have $p_T > 2.0$ GeV/c. If two seed tracks are found to intersect, the algorithm then searches the general track list—seed

tracks plus non-seed tracks—to find all tracks that pass within 3σ of the seed vertex.[2] If at least one track is found that satisfies this requirement, we then attempt an iterative vertex fit of the seed plus additional tracks to a common vertex. Any tracks that contribute more than 50 to the χ^2 of the fit are removed, and the fit is repeated. This process of removing tracks that contribute too much to the fit χ^2 continues until no tracks in the fit contribute more that 50, or until there are fewer than 3 tracks remaining.

If the iterative fit successfully returns a vertex containing three or more tracks, the algorithm then demands that at least one track in the vertex pass the $|D_0/\sigma| > 2.5$ cut. However, if at least three tracks pass the cut, any extra tracks failing the cut are dropped and the vertex is refit.[3] Tracks dropped at this stage are still marked as used so that they cannot be used in any other secondary vertex fits. Removing these extra "sliding" tracks—that is tracks that fail the $|D_0/\sigma| > 2.5$ cut—yields three classes of Pass 1 vertex: vertices with no sliding tracks, vertices with one sliding track, and vertices with two sliding tracks. For each Pass 1 vertex, we calculate the separation between the secondary vertex and the primary vertex in the *x*-*y* plane, L_{xy}, as well as the uncertainty on this quantity σ_{Lxy}. L_{xy} is signed according to whether the vertex is consistent with a particle having been produced at the primary vertex and travelling to some distance away from the primary before decay, or with a particle decaying first, and then heading towards the primary vertex. See Section 4.4.1 for more details on how the sign of L_{xy} is determined.

To be accepted as a valid secondary vertex, a Pass 1 vertex must satisfy the following cuts:

- $|L_{xy}| < 2.5$ cm.
- Pseudo-cτ < 1 cm, where pseudo-cτ = $L_{xy} m_{vertex} / p_T$. The symbol τ refers to the proper lifetime of the *B* hadron. For a displaced vertex, the quantity is called pseudo-cτ because only the momentum and mass contributed by the tagged tracks is included in the calculation, rather than all the decay products of the *B* hadron.
- $|L_{xy}/\sigma_{L_{xy}}| > 2.0$ for vertices with no sliding tracks.
- $|L_{xy}/\sigma_{L_{xy}}| > 4$ for vertices with one or two sliding tracks.

[2] In the original BVTX algorithm, only tracks from the seed list were considered for inclusion in the vertex.

[3] This step was not included in the original BVTX algorithm because non-seed tracks weren't included in the vertex.

- Vertex fit $\chi^2 / d.o.f. < 12$.
- $\sigma_{L_{xy}} < 0.1$ cm.
- The invariant mass calculated from the tracks included in the vertex, $m_{vertex} < 7$ GeV/c^2

These requirements are intended to reduce the number of vertices that are poorly measured, unphysical, or not significantly separated from the primary vertex. A vertex that passes all these cuts is saved, and the tracks that are part of the vertex are marked as used so that they won't be considered for any further vertices.

If Pass 1 succeeds at tagging a displaced secondary vertex within a given jet, another attempt to find another Pass 1 vertex in the same jet is made. Attempts to find Pass 1 vertices continue for the same jet until no further Pass 1 vertices can be found.

4.2.3.3 Pass 2

Once Pass 1 fails to find any more displaced vertices within a given jet, BVTX moves on to Pass 2. Pass 2 looks for displaced vertices containing only two tracks, but makes tighter requirements on the tracks considered. Tracks from the track list that pass the following additional cuts are considered eligible for Pass 2:

- $p_T > 1.0$ GeV/c.
- $|D_0/\sigma| > 3.0$.
- SVX' tracks with hits in only three of the four SVX' layers must have 2 high-quality hits as defined in Section 4.2.3.1.
- 4-hit SVX' tracks need only 1 high-quality hit as defined in Section 4.2.3.1.

From the tracks in the list above, the algorithm attempts to find two tracks that intersect at a common vertex, considering the tracks as they were ordered in the track list. At least one track of the two must have $p_T > 2.0$ GeV/c. If such an intersection is found, an iterative fit is then attempted using these two tracks, plus all tracks that satisfy the Pass 2 requirements that are within 3σ of the two-track seed vertex. The iterative fitting procedure for Pass 2 is the same as the one used in Pass 1, except the iterations are allowed to continue as long as at least two tracks

are left. If a two-track vertex is found, it is required to pass the same L_{xy} and psuedo-$c\tau$ requirements as a Pass 1 vertex. Of course, because there are no sliding tracks in Pass 2, the vertex need only have $\left|L_{xy}/\sigma_{L_{xy}}\right| > 2$. In addition, the mass for the vertex returned from the fit is required not to be within 0.02 GeV/c^2 of the K^0 mass.

4.3 Raw Double-Tagged Sample

The data sample for this analysis includes all events from the BVTX sample that satisfy the lepton selection criteria from Section 4.1 in which we find at least two displaced secondary vertices (tags) using the modified BVTX tagging algorithm outlined above in Section 4.2. Of the 497,051 electron events and 430,270 muon events in the BVTX sample, 25,887 electron events and 19,944 muon events contain two or more secondary vertex tags. In addition, because these events are supposed to come from $b\bar{b}$ production in which one of the resulting B hadrons decays semileptonically, we demand that at least one lepton in each event which passes all trigger requirements be within a cone of $\Delta R = 1.0$ of a valid displaced secondary vertex passing all cuts. There are 24,107 electron events and 18,886 muon events that satisfy those requirements. This sample will be referred to as the "raw double-tagged sample."

Below we discuss the characteristics of these data. We plot distributions separately for the electron and muon samples. There is good agreement between the two samples, except for distributions related to pseudorapidity, where the difference in acceptance for electrons and muons can be seen.

4.3.1 Primary Vertex

Figure 4.4 shows plots of the primary vertex location for each event in x and y, as well as the errors on the primary vertex measurement in each direction. The variation in x and y for the primary vertex location is mainly due to the variation in beam position from run to run. As can be seen from the plots, the average error on the primary vertex position in x and y is $\sigma_x = 17.4$ μm and $\sigma_y = 16.8$ μm respectively. Note that the error distributions in x and y are truncated at 25 μm since the primary vertex finding algorithm uses the beam position and spot size if it cannot locate a primary vertex with smaller errors.

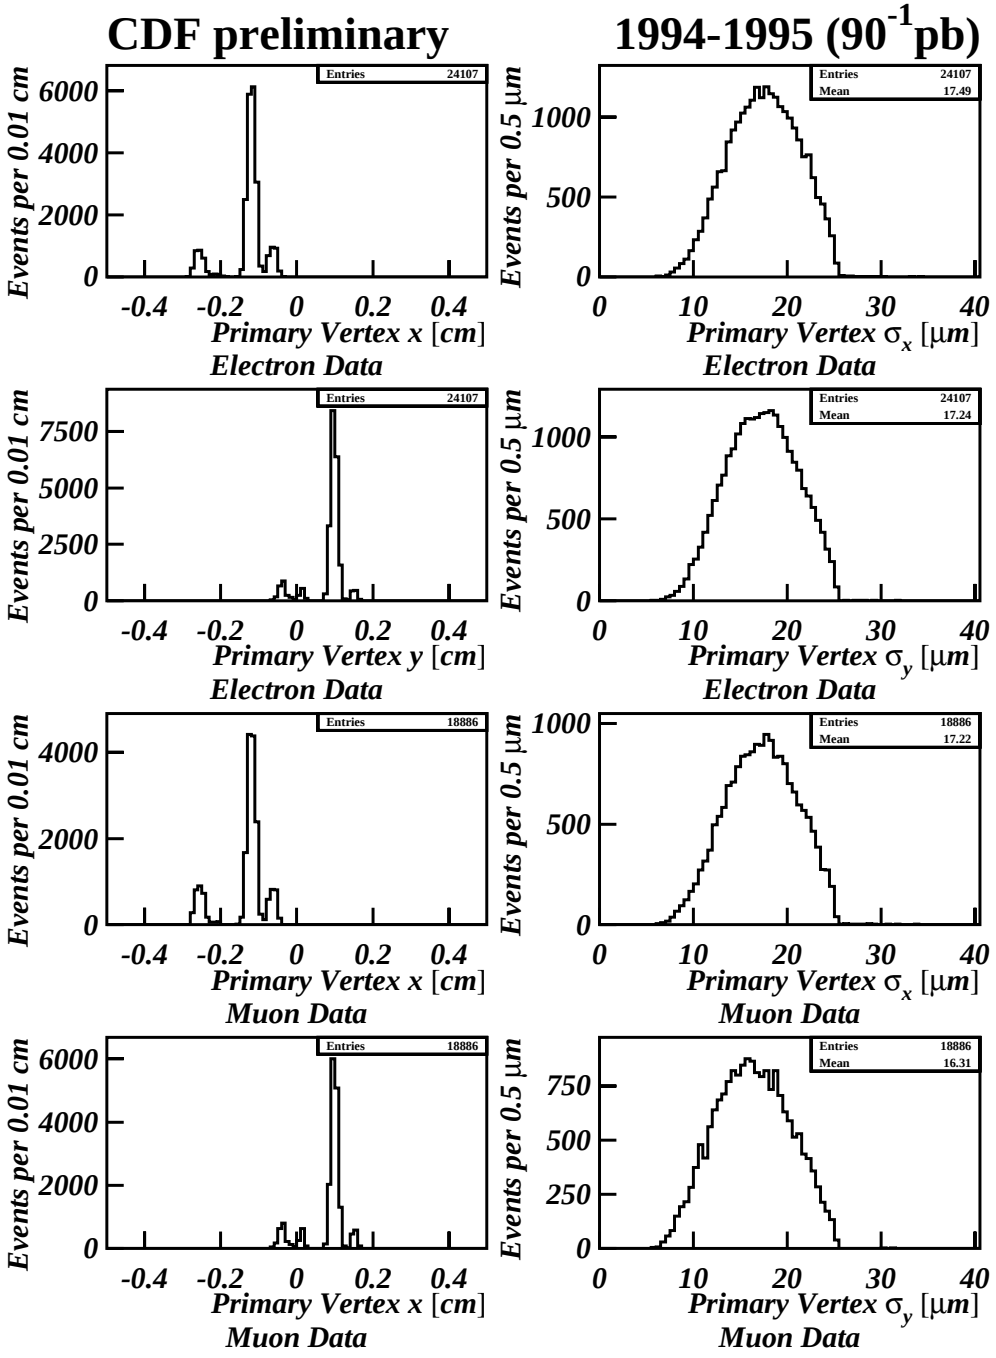

Figure 4.4 The primary vertex position and uncertainty for both electron and muon data. The variation in the *x* and *y* positions of the primary vertex is mainly due to the movement of the beam from run to run. Over the course of any given run, the beam position is stable to within a few microns [60]. The error on the primary vertex is truncated at the beam spot size, 25 μm.

4.3.2 Leptons

Figure 4.5 shows the number of trigger leptons and the separation between the lepton and the nearest secondary vertex tag in both the electron and muon samples of double-tagged events. Note that there are occasionally events in which more than one lepton passes the trigger requirements. For each sample, we count only leptons of the appropriate type as "trigger leptons." In other words, only electrons are counted as trigger leptons for the electron trigger sample, and the presence of extra muons are ignored, although the muon tracks are used in the secondary vertex search. Figure 4.6 shows the E_T distribution for the trigger electrons and the p_T distributions for both electrons and muons. The shapes at low E_T and p_T are determined primarily by the Level 2 trigger efficiency.

4.3.3 Jets

The number of jets found per event and the number of tracks per jet are shown in Figure 4.7. Figure 4.8 shows the p_T spectrum for clustered jets. The bump in the spectrum is an artifact of the p_T requirement on the trigger lepton. As can be seen from the plots, higher p_T jets are more likely to contain at least one secondary vertex tag.

4.3.4 Secondary Vertices

Figure 4.9 shows the number of tags per event and the number of tags of each pass. As can be seen, there are events containing more than two tags. Extra tags can come from mistags, which result from random combinations of tracks, or sequential double-tags, which result from tagging the same *B* decay twice, frequently because *B* decay involved the production of a long-lived *D* hadron which decayed at a separate vertex. There are also a small number of events in which more than two *b* quarks are produced, so it is possible that for some of the events containing more than two tags, all the tags do correspond to legitimate *B* vertices. From the plots in Figure 4.9 you can also see that most tags come from Pass 1.

Figure 4.10 shows the L_{xy} and $\sigma_{L_{xy}}$ distribution for the data. The average error on the two-dimensional decay distance is 145.9 μm. L_{xy} is signed so only vertices consistent with the decay of a particle produced at the primary interaction point have positive L_{xy}. The negative portion will be important for subtracting the mistag component from our data sample. See Section 4.4.1 for details. Figure 4.11 plots the p_T and mass distribution for the tags. Tag p_T is calculated from

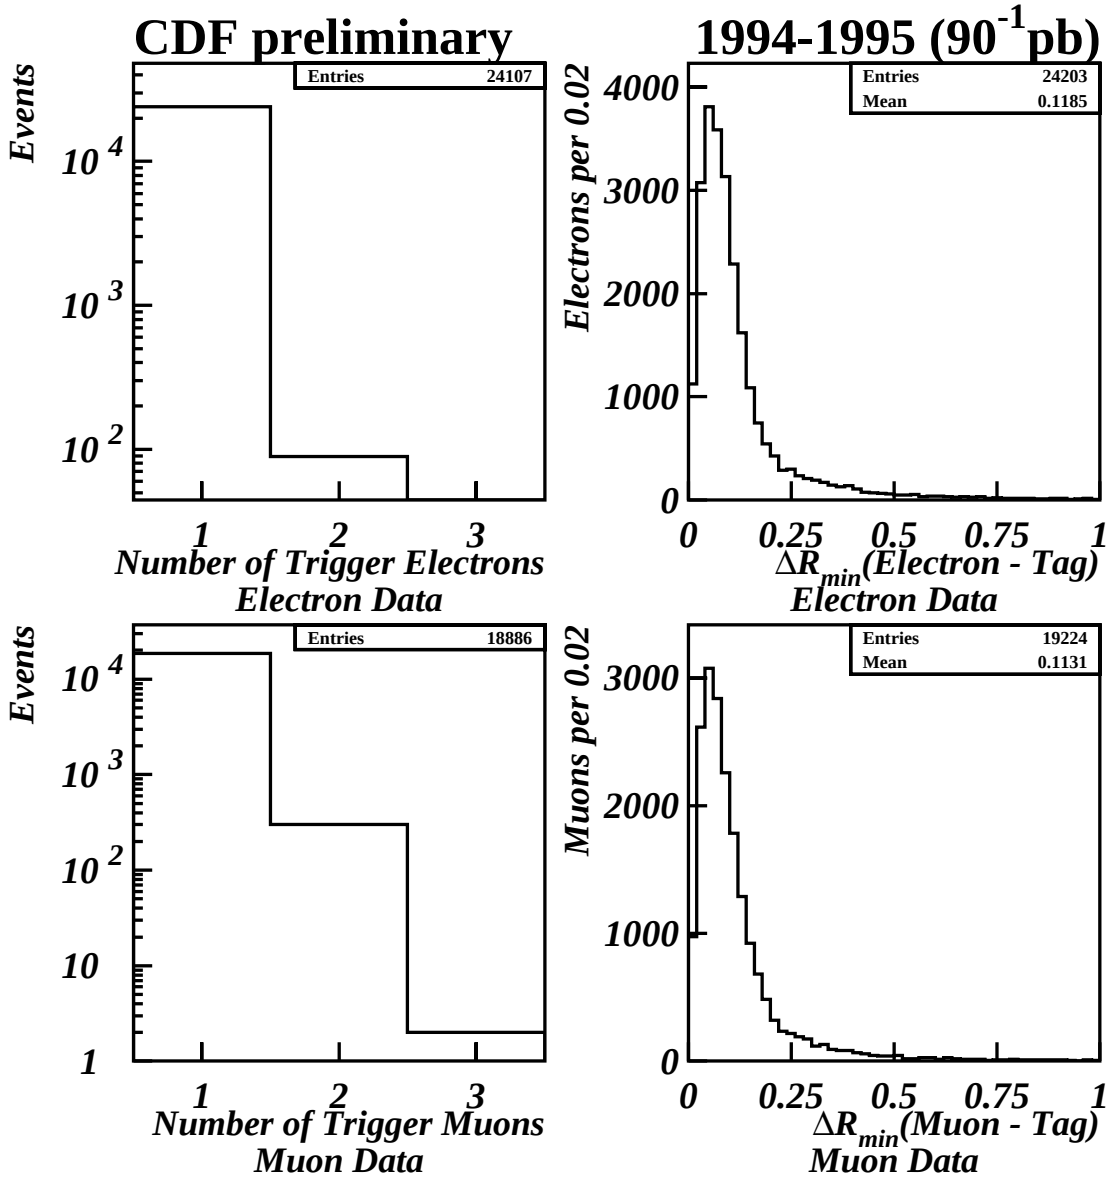

Figure 4.5 The number of leptons per event passing trigger cuts and the minimum separation between such lepton and the nearest secondary vertex tag. The direction of the tag is determined from the vector sum of the momenta for the tag tracks. Events with more than one trigger lepton make up 0.4% of the electron sample and 1.6% of the muon sample.

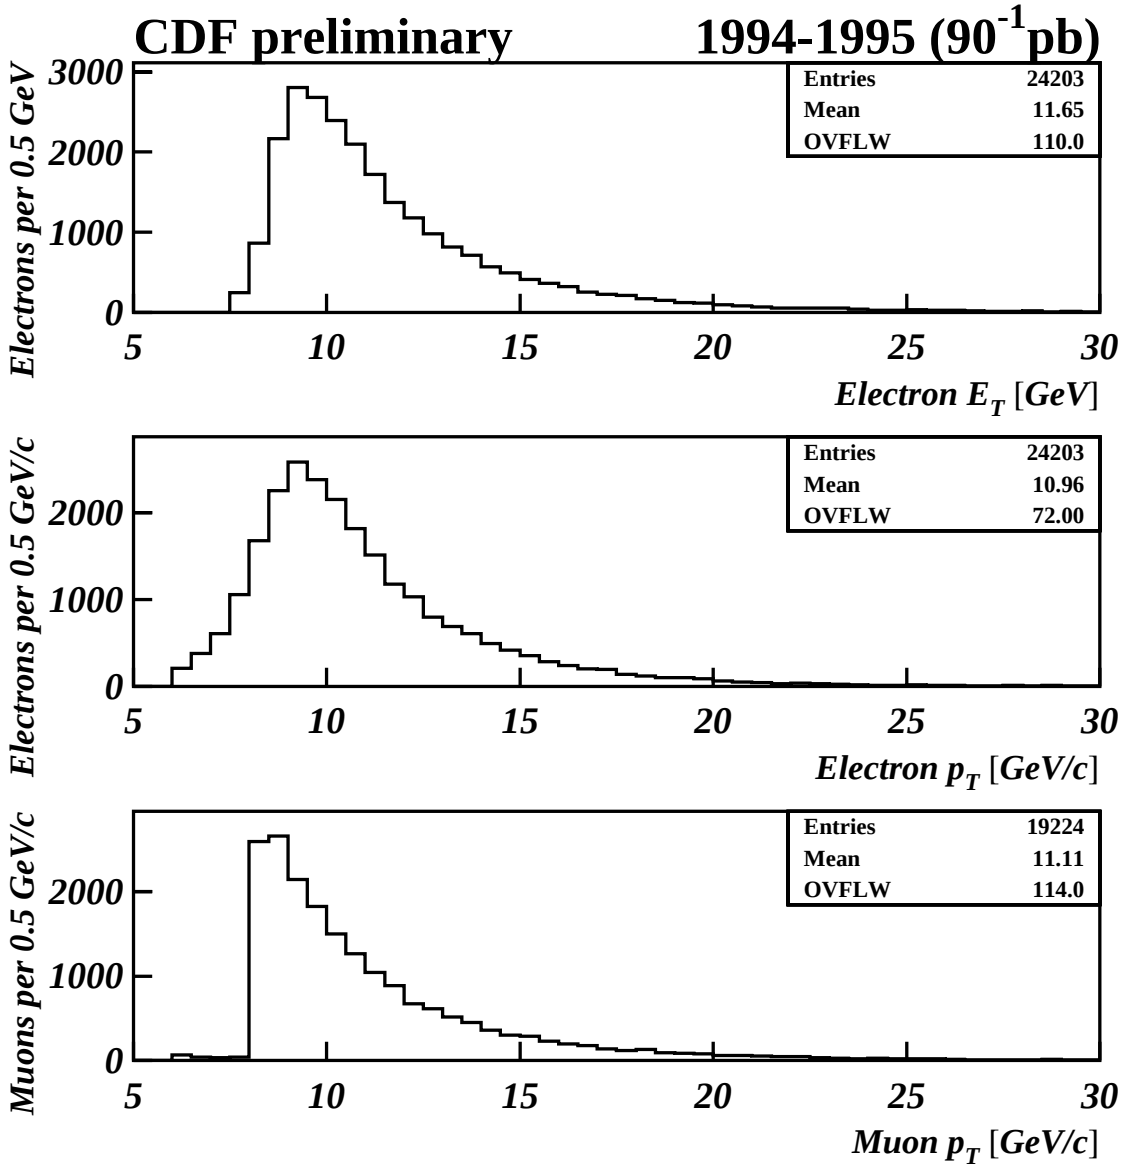

Figure 4.6 Electron E_T and p_T, and muon p_T. The shaping at the low end of the spectrum is primarily determined by the Level 2 trigger. Note that for the muon data, the Level 3 trigger imposes an 8 GeV/c p_T cut on trigger muons.

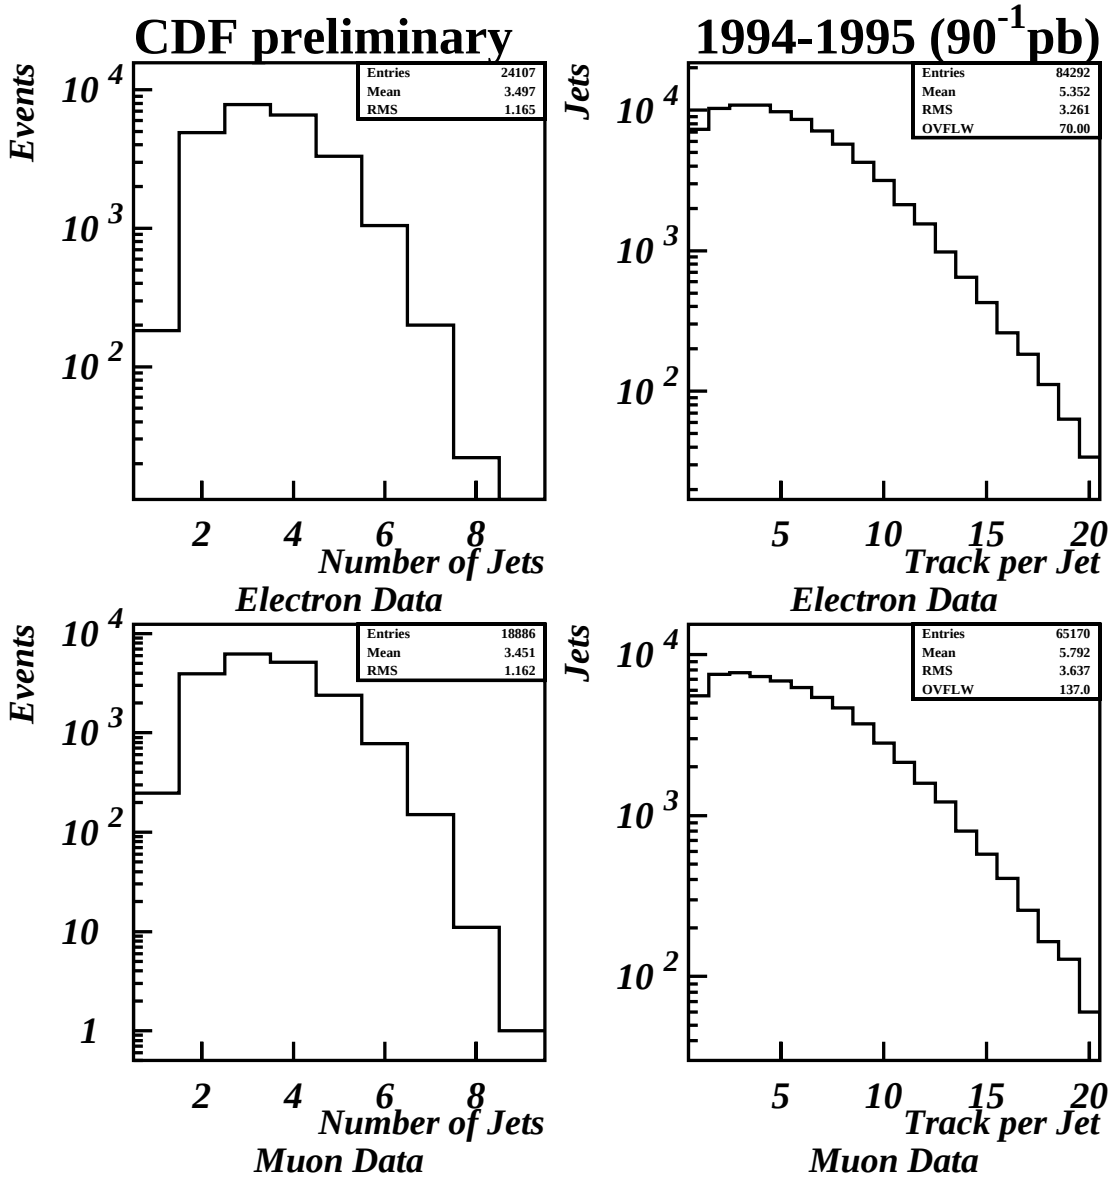

Figure 4.7 The number of jets per event and the number of tracks per jet for electron and muon data.

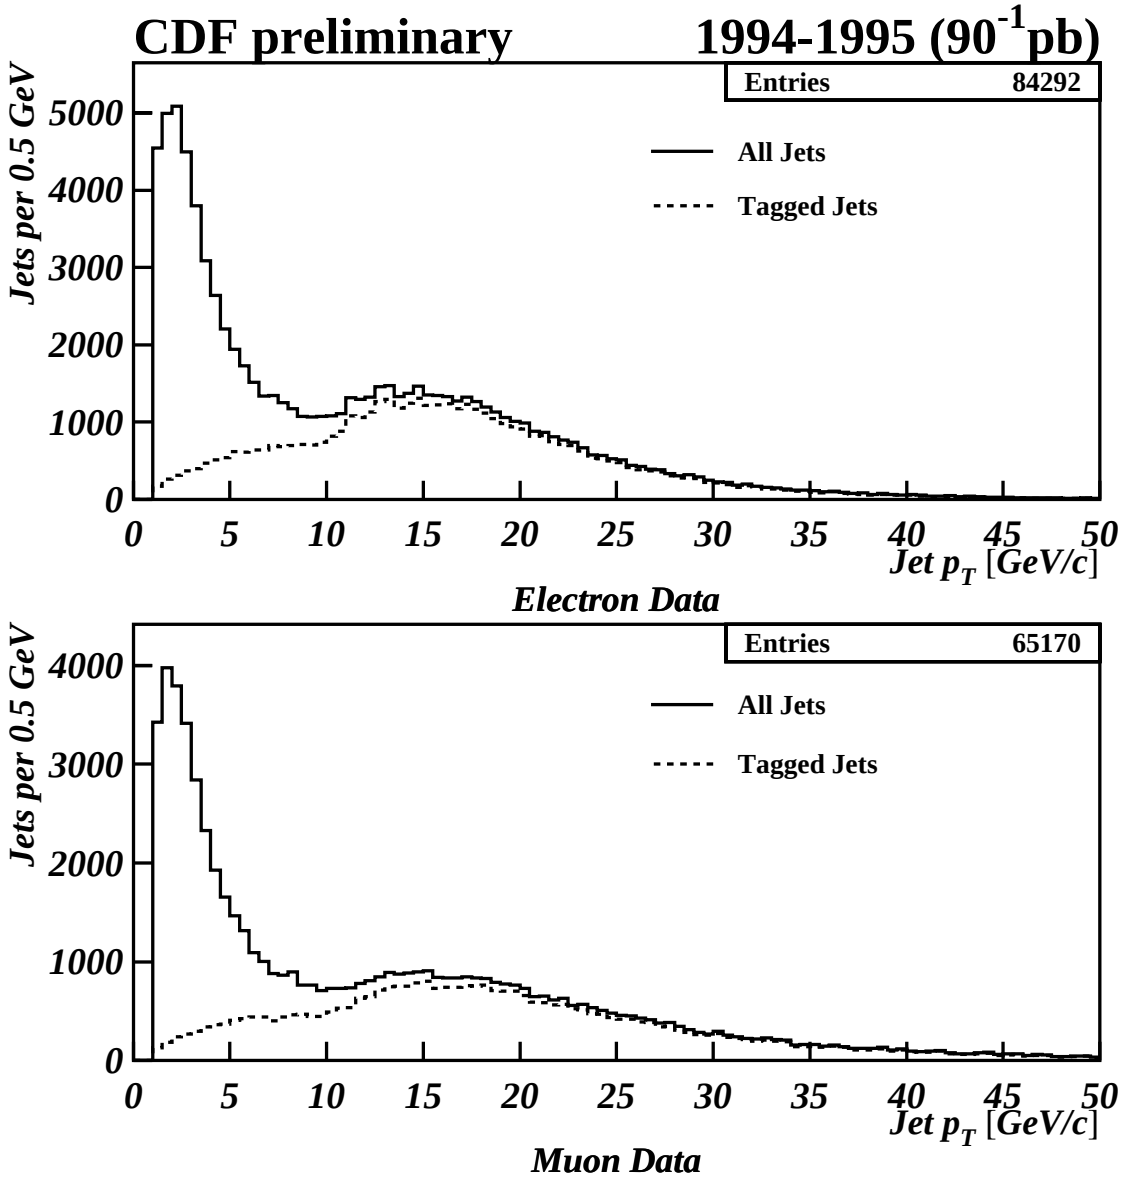

Figure 4.8 The jet p_T for electron and muon data. The p_T for all jets is shown by the solid line while the p_T for jets which contain a secondary vertex tag is shown by the dashed line.

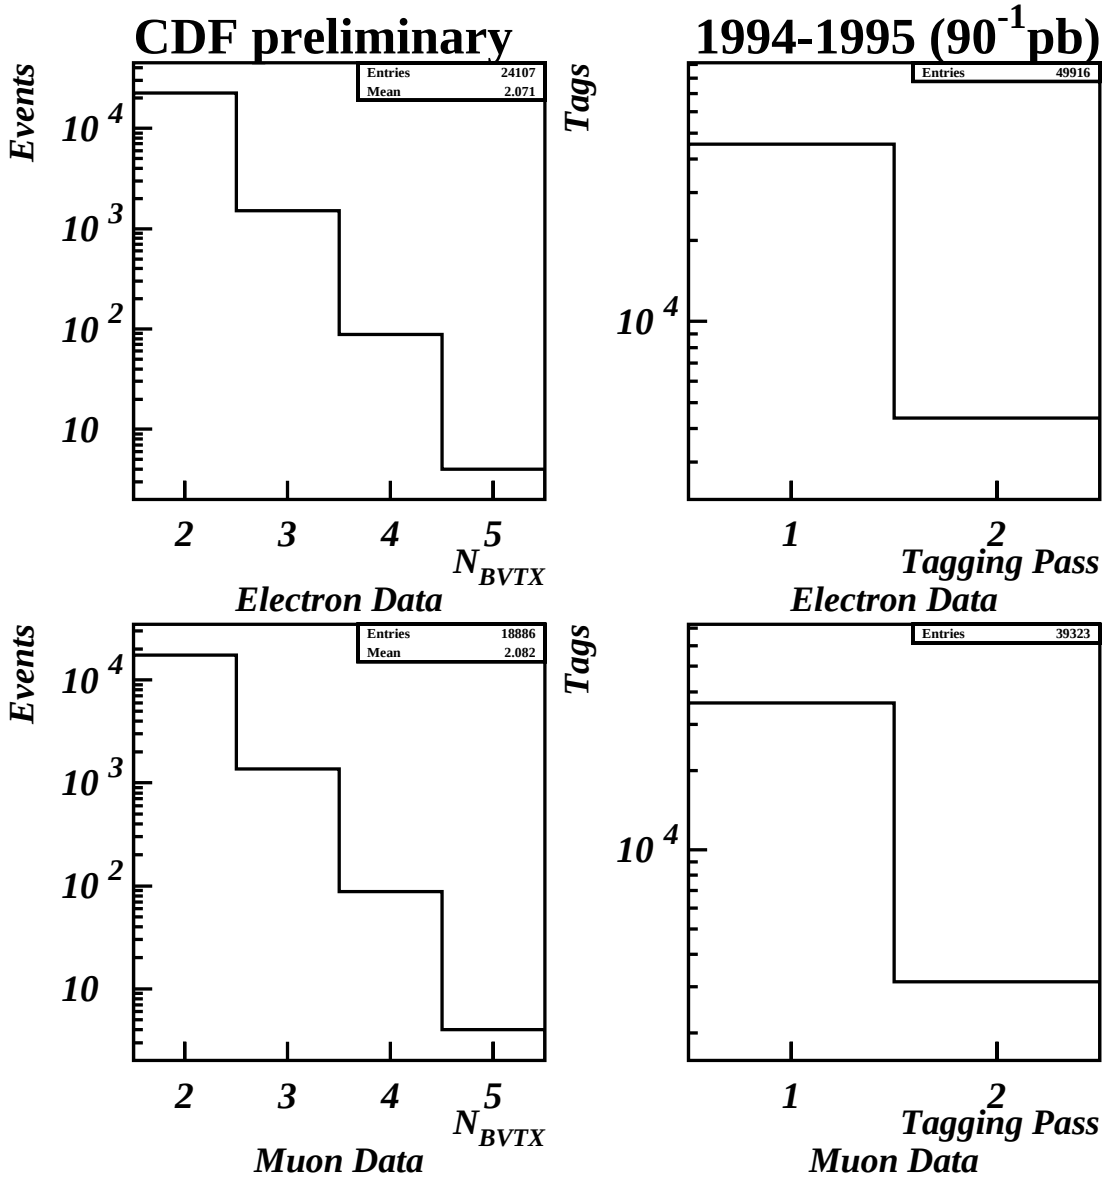

Figure 4.9 The number of secondary vertex tags per event and the tagging pass for each secondary vertex. In the electron sample, 6.7% of the events have more than two tags, while in the muon sample 7.7% of the events have more than two tags. Pass 1 tags account for 91.2% of the electron sample tags and 92.0% of the muon sample tags.

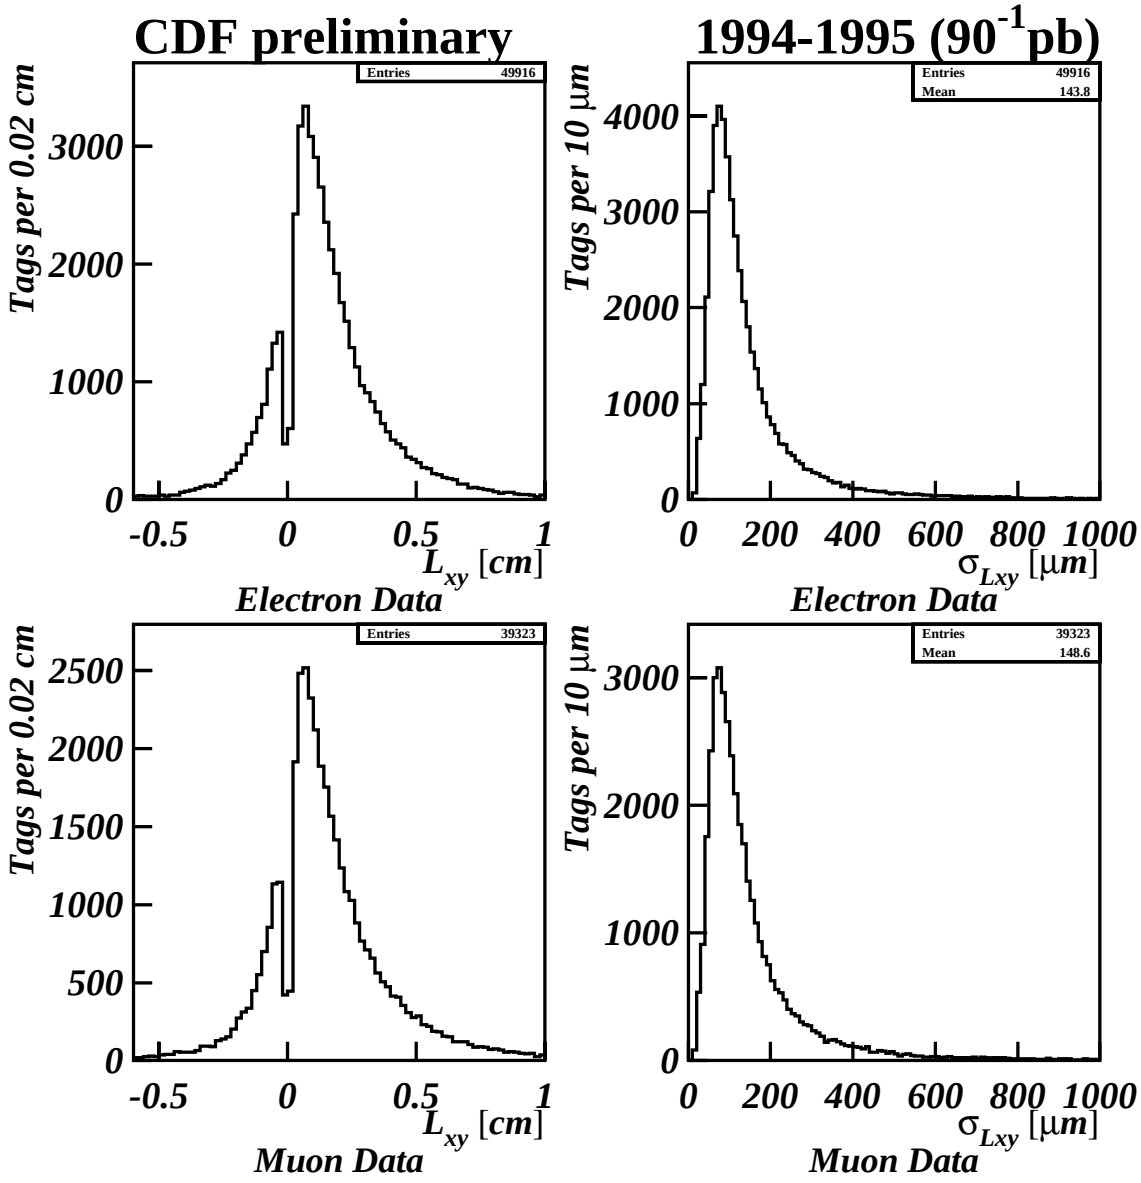

Figure 4.10 Secondary vertex tag L_{xy} and σ_{Lxy}. Tags with $L_{xy} < 0$ result from mistagging.

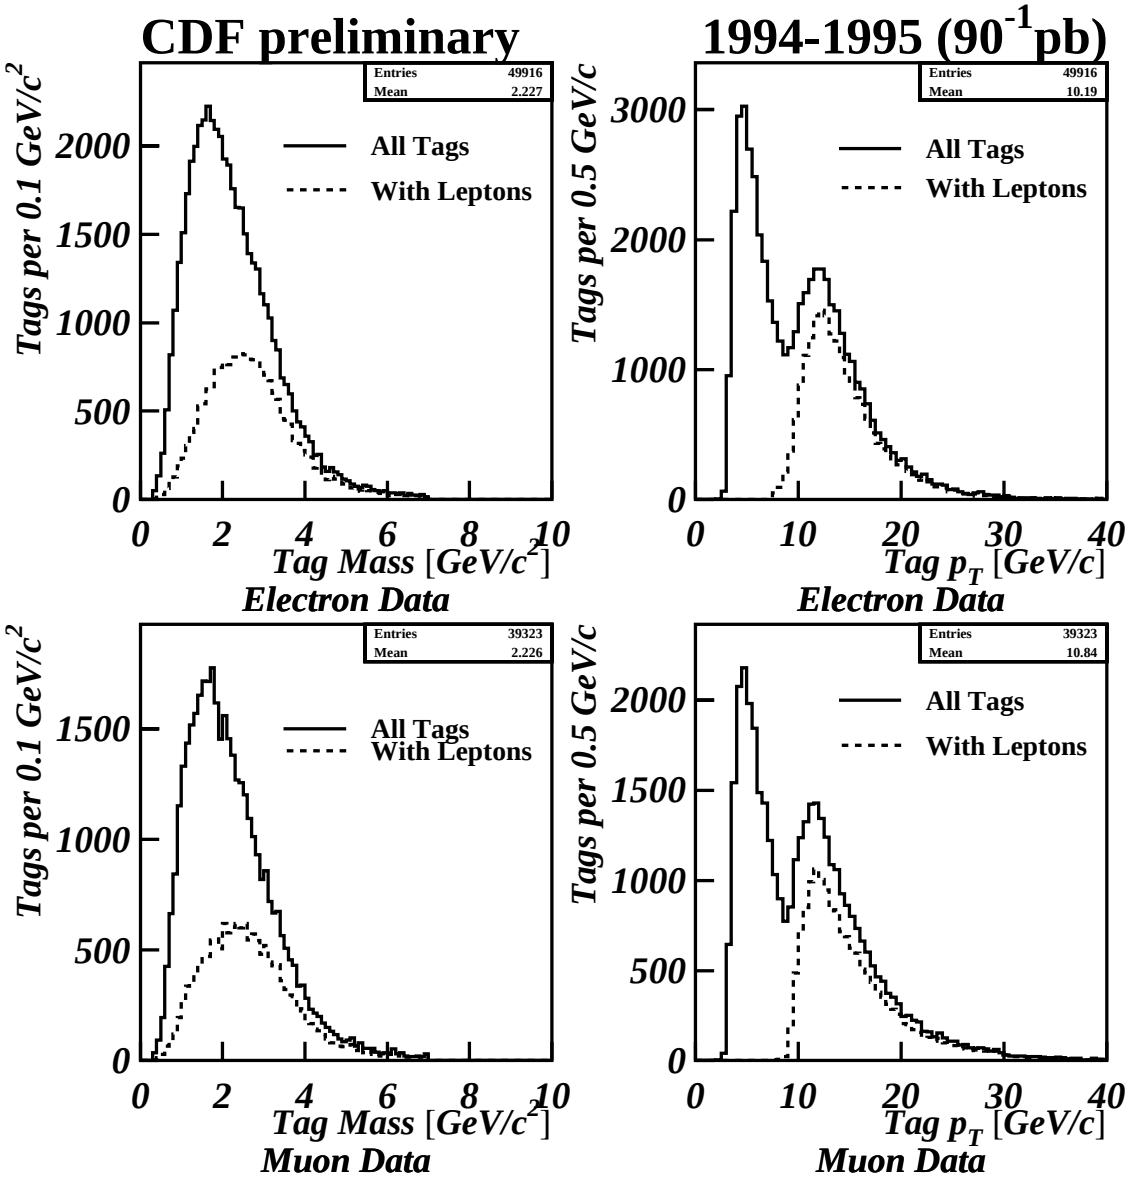

Figure 4.11 Secondary vertex tag mass and p_T. Tag p_T is defined as the vector sum of the p_T for tracks used in the tag. The tag mass is the invariant mass calculated from the tag tracks, assuming a pion mass for all tracks. The solid distributions are for all tags while the dashed distributions only include tags containing leptons. There is a marked difference, especially in the p_T distributions between tags that contain leptons and those that don't.

the vector sum of the momenta of the tracks used in the tag, and the tag mass is the invariant mass calculated from these tracks, assuming a pion mass for all tracks. There is a clear difference between the p_T distributions for tags containing leptons and for those that don't.

4.3.5 Raw Correlations

Figure 4.12 shows raw $\Delta\phi$, $\Delta\eta$, and ΔR distributions for tag pairs. No attempt beyond the quality cuts discussed in Section 4.2 has been made to remove sources of background such as mistags or sequential double-tags (discussed in Section 4.4.2). Since there is no way to know which tag pair, if any, in an event with more than two tags is the correct pair, every possible pairing of tags is included for multi-tag events. It will be up to background removal techniques to eliminate spurious tag pairs. Note that the peak in the $\Delta\phi$ distribution near $\Delta\phi = 0$ hints that the sample may include a number of sequential double-tag pairs.

4.4 Multi-Tag Sample Composition

The plots from the previous section suggest that $b\bar{b}$ production is not the only contribution to the double-tagged BVTX sample. In addition to $b\bar{b}$ production, one expects a number of other sources for double-tagged events. Certainly the largest source comes from mistags, which occur when the tagging algorithm accidentally associates into a vertex prompt tracks that do not come from the decay of a long lived particle. Also of concern are sequential double tags, which result from the multiple tagging of the same *B* decay. Sometimes, sequential double tags can occur because of confusion on the tagging algorithm's part, tagging the same decay point with two different sets of tracks because of tracking errors. However, it is also possible for one *B* decay to produce two secondary vertices because the *B* decay involves the production of a long-lived *D* hadron which decays at a point significantly separated from the *B* decay point. Both of these issues will have to be addressed. Another source of background in the raw double-tagged sample comes from direct $c\bar{c}$ production, either alone, or in events that already contain a $b\bar{b}$ pair (i.e. $b\bar{b} + c\bar{c}$ production). Although the background from prompt charm is a small contribution due to the smaller charm lifetime, it must be seriously considered because secondary vertex tags from charm decays are legitimate tags that cannot be removed in the same way as mistags and sequentials. Finally, one has to consider distortions arising from fake trigger leptons contributing events to the sample, although this last contribution proves to be inconsequential.

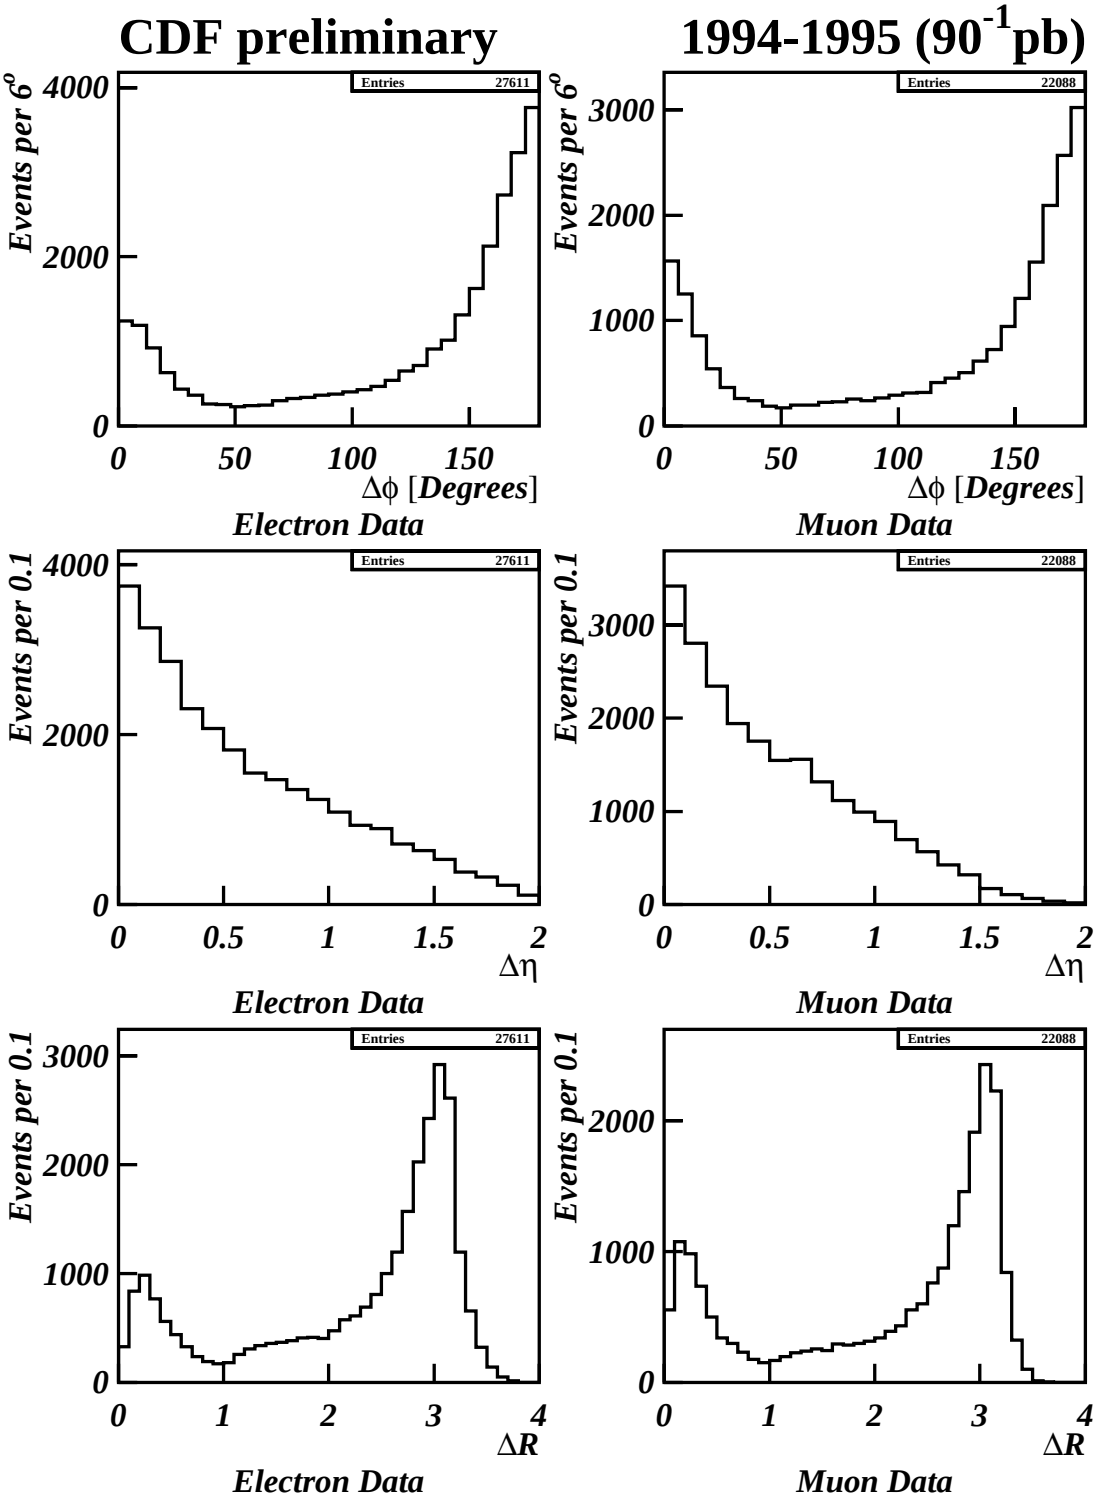

Figure 4.12 Raw correlations between tag pairs. No attempt has been made to remove backgrounds beyond the quality cuts detailed in Sections 4.1 and 4.2.

Table 4.1 gives a few examples of the kinds of tags that may be found in the data and how they are classified.

4.4.1 Mistags

A mistag occurs when the vertex finding algorithm fits a vertex from a set of tracks that don't physically originate from a common decay vertex. It is occasionally possible because of measurement errors or random chance to find a set of prompt tracks that seem to intersect at a vertex separated from the primary. Also, when *b* quarks are produced with a small opening angle, it is not hard to combine tracks from two different *B* hadron decays, sometimes with additional prompt tracks, accidentally to form a false vertex that doesn't correspond to the decay point of either *B*. These decay vertices distort the spectrum of correlations and must be removed in order to get a good quality measurement.

One way to identify mistags is by looking at the signed L_{xy} distribution. L_{xy} is signed based on the inferred direction of the particle that decayed to produce the secondary vertex tag relative to the primary vertex. A particle that seems to be moving out from the primary vertex at the time of decay is given a positive L_{xy}, while a particle that seems to have been moving towards the primary vertex when it decayed is assigned a negative L_{xy}. Figure 4.13 illustrates this concept.

In principal, a secondary vertex tag corresponding to the decay of real, long-lived particle will

Scenario	**Classification**
The tracks in the tag are from the same *B* decay (including any tracks from a secondary *D* decay)	Good Tag (Signal)
The tag contains random prompt tracks not associated with the decay of any long-lived particle	Mistag (Background)
The tracks in the tag are from a *B* decay (including secondary *D* decay) that has already been tagged with other tracks.	Sequential Double-Tag (Background)
The tag tracks are from a prompt *D* decay—in other words, a *D* not associated with the decay of a *B*.	Prompt Charm (Background)

Table 4.1 A list of the different tagging scenarios and how they are classified as signal or background for this analysis. In the table above, *secondary D* refers to any charmed-hadron resulting from the decay of a *B* hadron, while *prompt D* refers to a *D* hadron produced directly in the $p\bar{p}$ interaction.

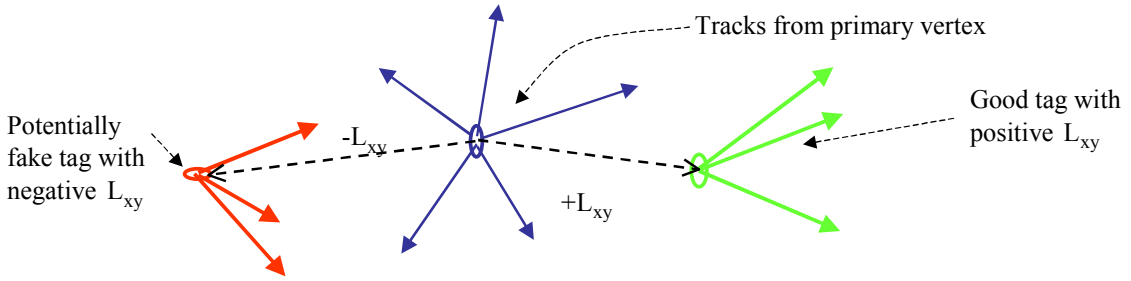

Figure 4.13 The relationship between negative L_{xy} and mistags.

have a positive L_{xy}. Resolution effects, however, make it possible to mismeasure a real secondary vertex so that it seems to have a negative L_{xy}. On the other hand, the random coincidence of tracks that leads to a mistag can occur anywhere relative to the primary vertex. It is just as likely that a mistag will be reconstructed with negative L_{xy} as with positive L_{xy}. One can use this feature of mistags to subtract them statistically from the data in a similar fashion to subtracting the appropriately normalized distribution from the side-bands of a mass peak to remove the background contribution. For single tag distributions, if we plot the distribution for all tags having $L_{xy} < 0$, to a good approximation, we will have the distribution for mistags only. The distribution for $L_{xy} > 0$ is a mixture of the distribution for good tags and the distribution for mistags. Since we assume mistagging is equally likely to produce tags with positive or negative L_{xy}, the mistag component in the positive L_{xy} sample is approximately the same size as the negative L_{xy} sample. Therefore, we can remove the mistag component by subtracting twice the distribution for tags with negative L_{xy} from the distribution for all tags. Figure 4.14 shows mistag-subtracted distributions for tag mass and tag p_T.

In the case of tag pairs, the technique remains the same, though the implementation becomes a bit more complicated. If we have a pair of secondary vertex tags, each tag in the pair may be either a good tag or a mistag. A plot of L_{xy}(tag 2) versus L_{xy}(tag 1) is useful in separating tag pairs containing mistags from tag pairs containing good tags. In such a plot, ignoring resolution effects, all secondary vertex tag pairs in which both tags are good tags should appear in quadrant I (L_{xy}(tag 1) > 0, L_{xy}(tag 2) > 0). However, if the first tag in the pair is a mistag, rather than a good tag, then half the time, the tag pair will appear in quadrant I, and half the time it will be in quadrant II (L_{xy}(tag 1) < 0, L_{xy}(tag 2) > 0). Likewise, pairs in which the first tag is a mistag and the second is good will appear half the time in quadrant I and half the time in quadrant IV.

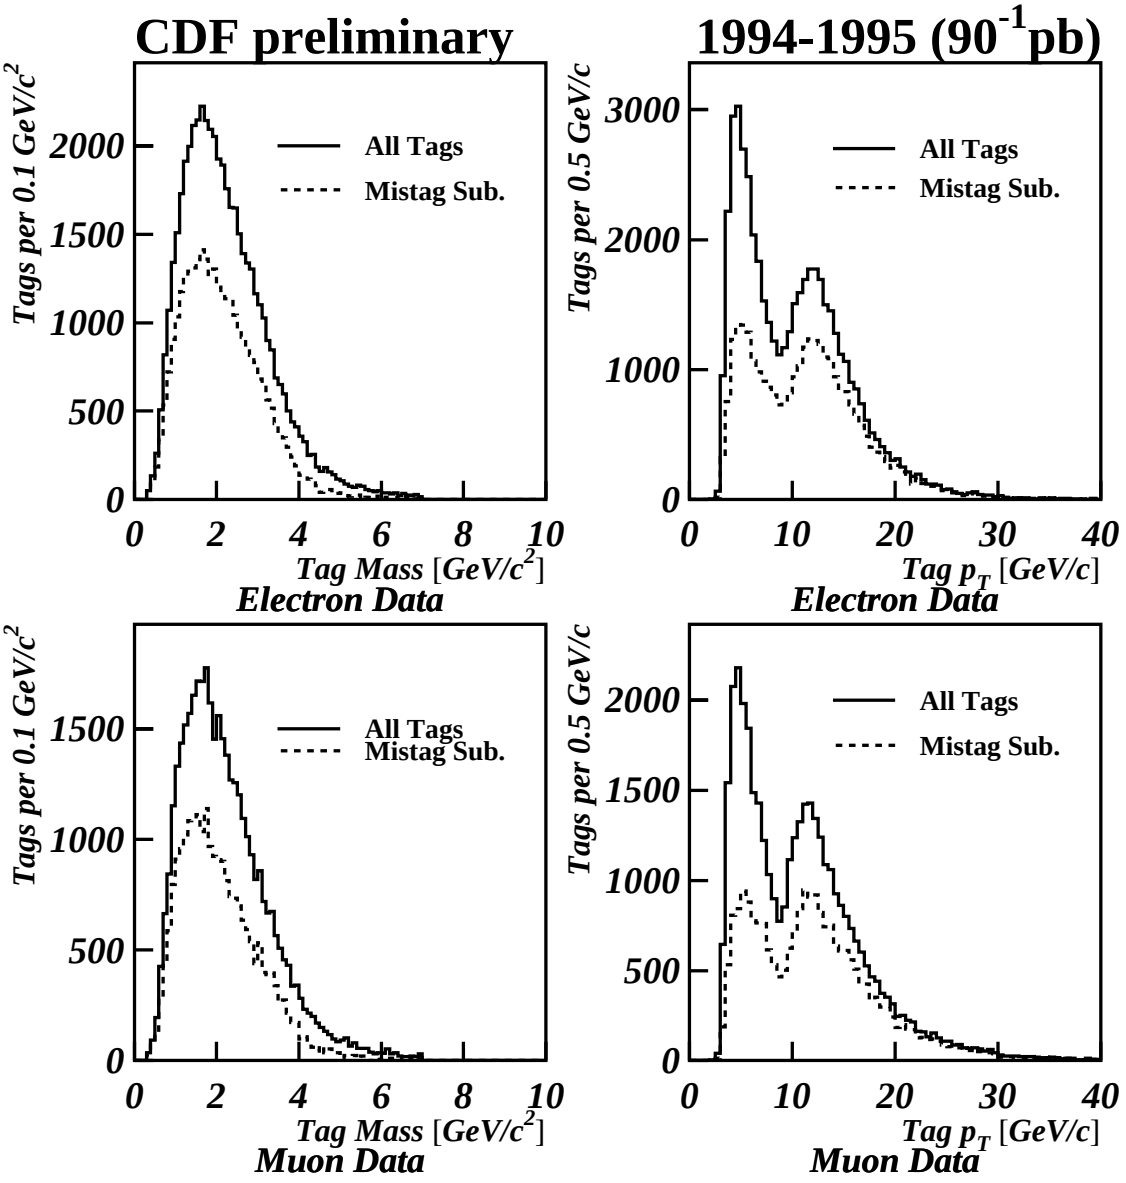

Figure 4.14 Mistag subtracted distributions. The solid histogram shows the raw tag distributions. The dashed line shows the distributions after mistags have been subtracted.

Finally, tag pairs consisting of two mistags will be equally likely to appear in all four quadrants. Figure 4.15 illustrates the reasoning above in algebraic notation.

We can solve the equations from Figure 4.15 for the number of tag pairs of each composition in terms of the number of tag pairs in each quadrant of the L_{xy}(tag 2) versus L_{xy}(tag 1) plot:

$$\begin{aligned} N_{GG} &= N_I - N_{II} + N_{III} + N_{IV} \\ N_{GM} &= 2(N_{IV} - N_{III}) \\ N_{MG} &= 2(N_{II} - N_{III}) \\ N_{MM} &= 4N_{III} \end{aligned} \tag{4.1}$$

Note that because jets are searched for secondary vertex tags in order of E_T, and because the tracks most likely to be from a displaced vertex are considered for tagging first, the two tags in the pair have an implicit ordering. Thus, it is not necessarily true that N_{GM} equals N_{MG}. Figure 4.16 shows a breakdown of tag pairs in the data sample by quadrant. From this plot, and using the equations above, we get 10,897 tag pairs in the electron data and 7,972 tag pairs in the muon data, in which neither tag is a mistag.

We can also create mistag-subtracted distributions by making plots of the distribution broken down by the quadrants of the tag pair. Then the above equations can be used to subtract these distributions from each other. Figure 4.17 and Figure 4.18 show this procedure applied to the $\Delta\phi$ distribution.

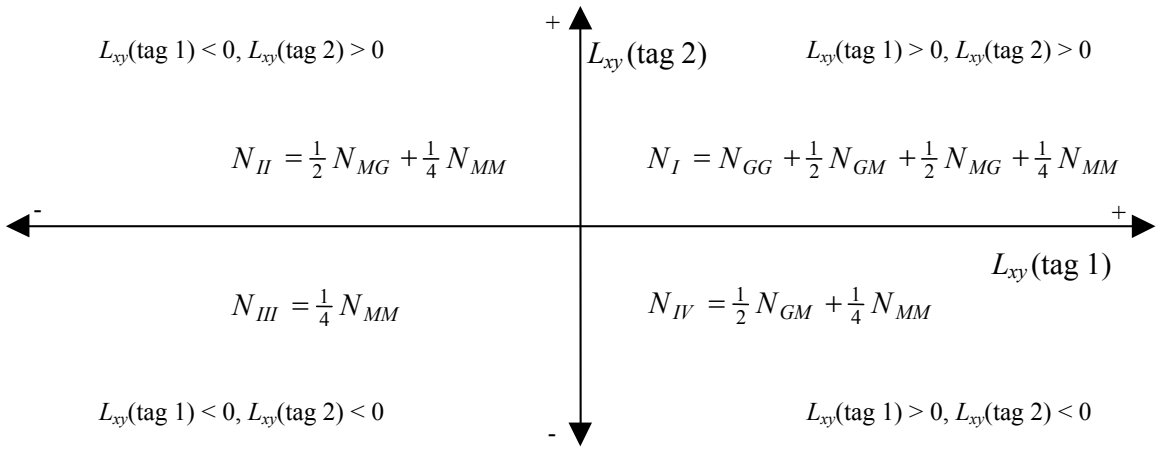

Figure 4.15 Schematic breakdown of tag pair types according to quadrant in the L_{xy}(tag 1) versus L_{xy}(tag 2) plot. Good tags are denoted by *G* and mistags by *M*.

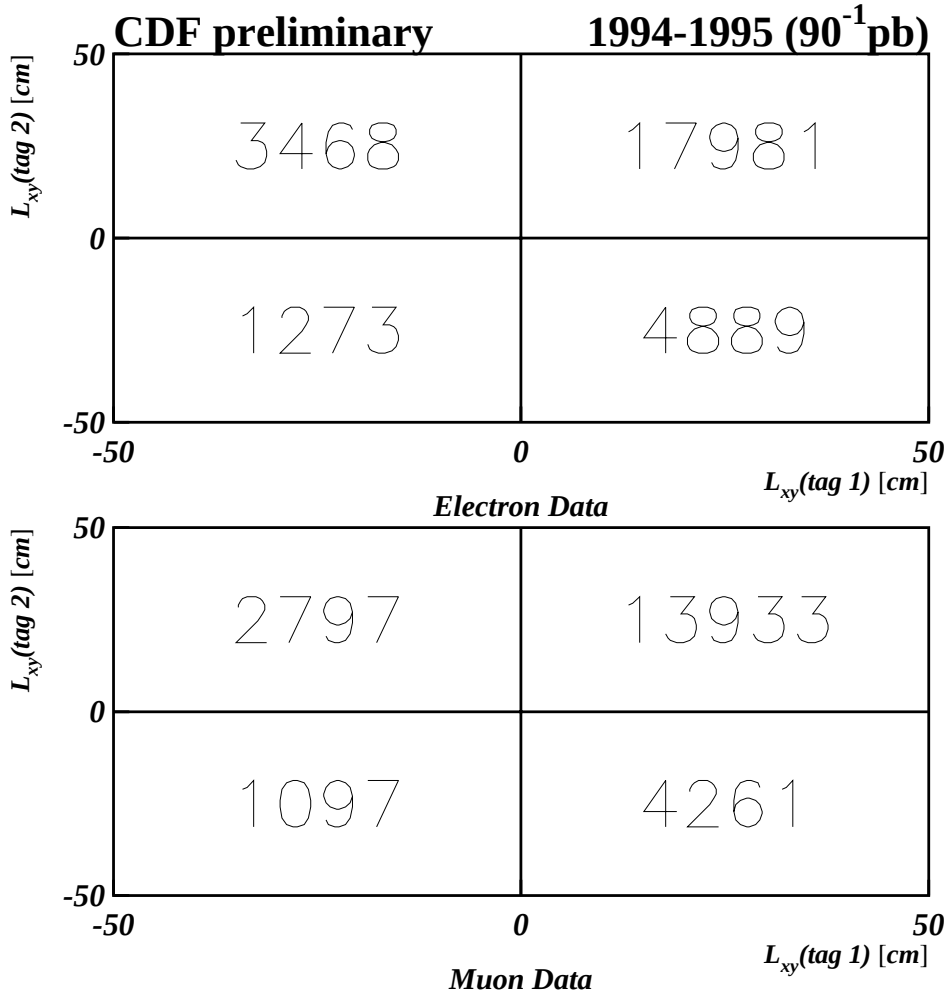

Figure 4.16 Tag pair breakdown by quadrant.

4.4.2 Sequential Double Tags

Another possible source of background involves tagging more than one secondary vertex from a single *B* decay. These tags, known as sequential tags, are most likely to occur when the *B* decay involves the production of a *D* hadron that travels a certain distance from the *B* decay vertex before itself decaying. It should be noted, however, that it is not necessarily true that in a sequential tag pair, one vertex tags the location of the *B* decay and the other tags the location of the *D* decay. It is also possible that at least one of the two tags contains a mixture of tracks from

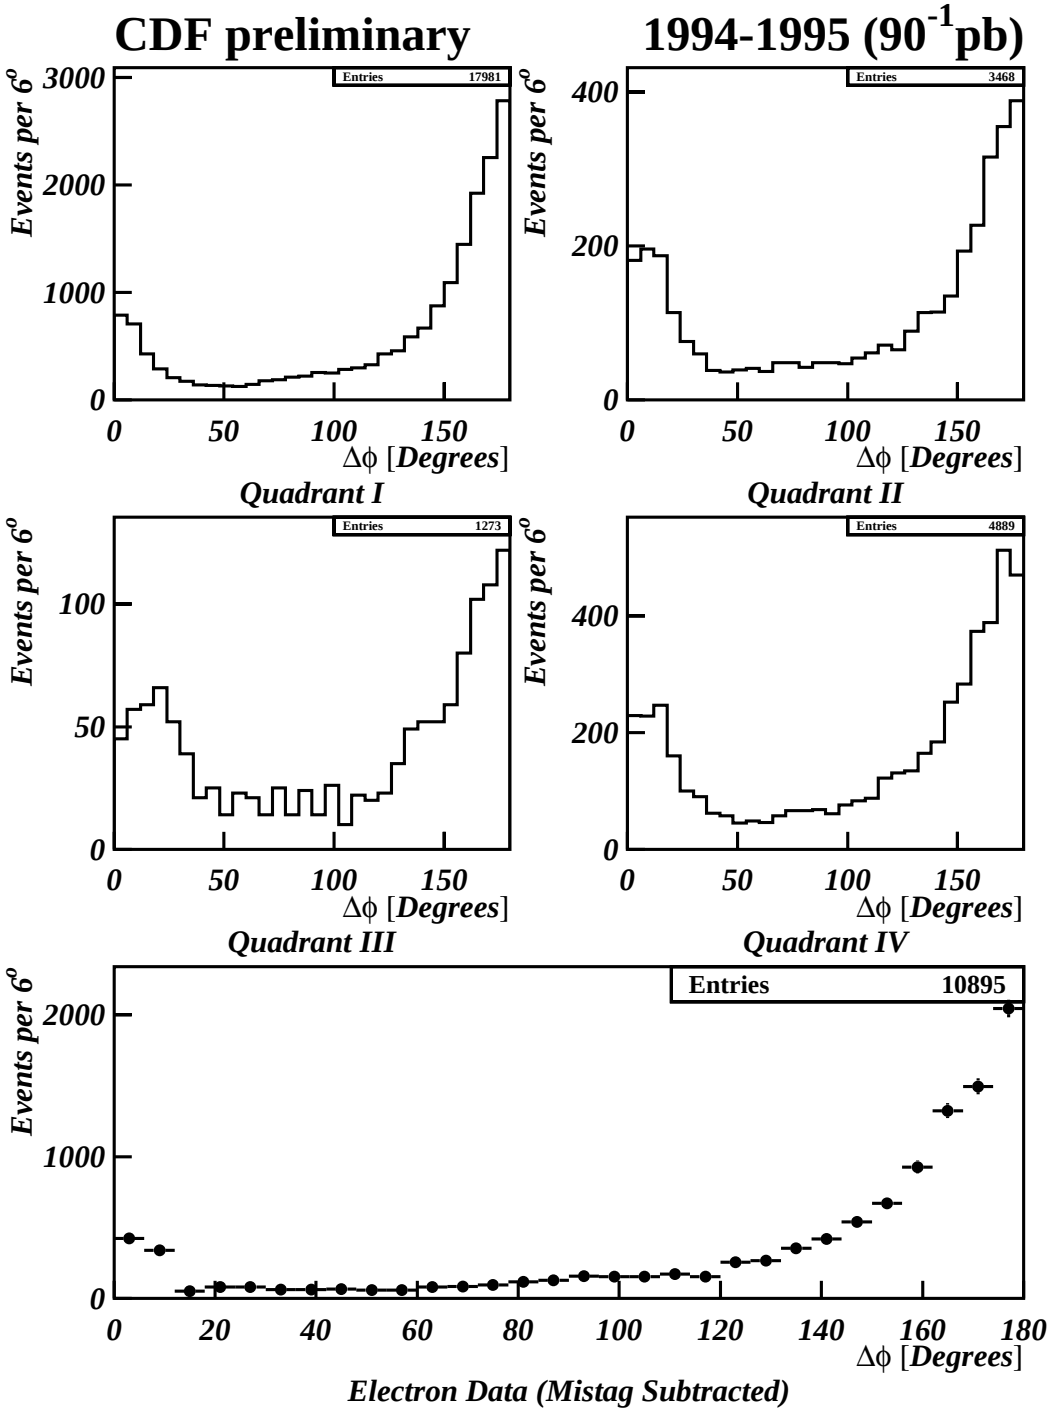

Figure 4.17 Mistag subtraction applied to the Δϕ distribution from electron data. The upper four plots show the raw distributions for tag pairs from each quadrant of the L_{xy}(tag 2) versus L_{xy}(tag 1) plot. In the bottom plot, the distributions from each quadrant are subtracted according to the equation from Section 4.4.1 to yield the mistag subtracted plot.

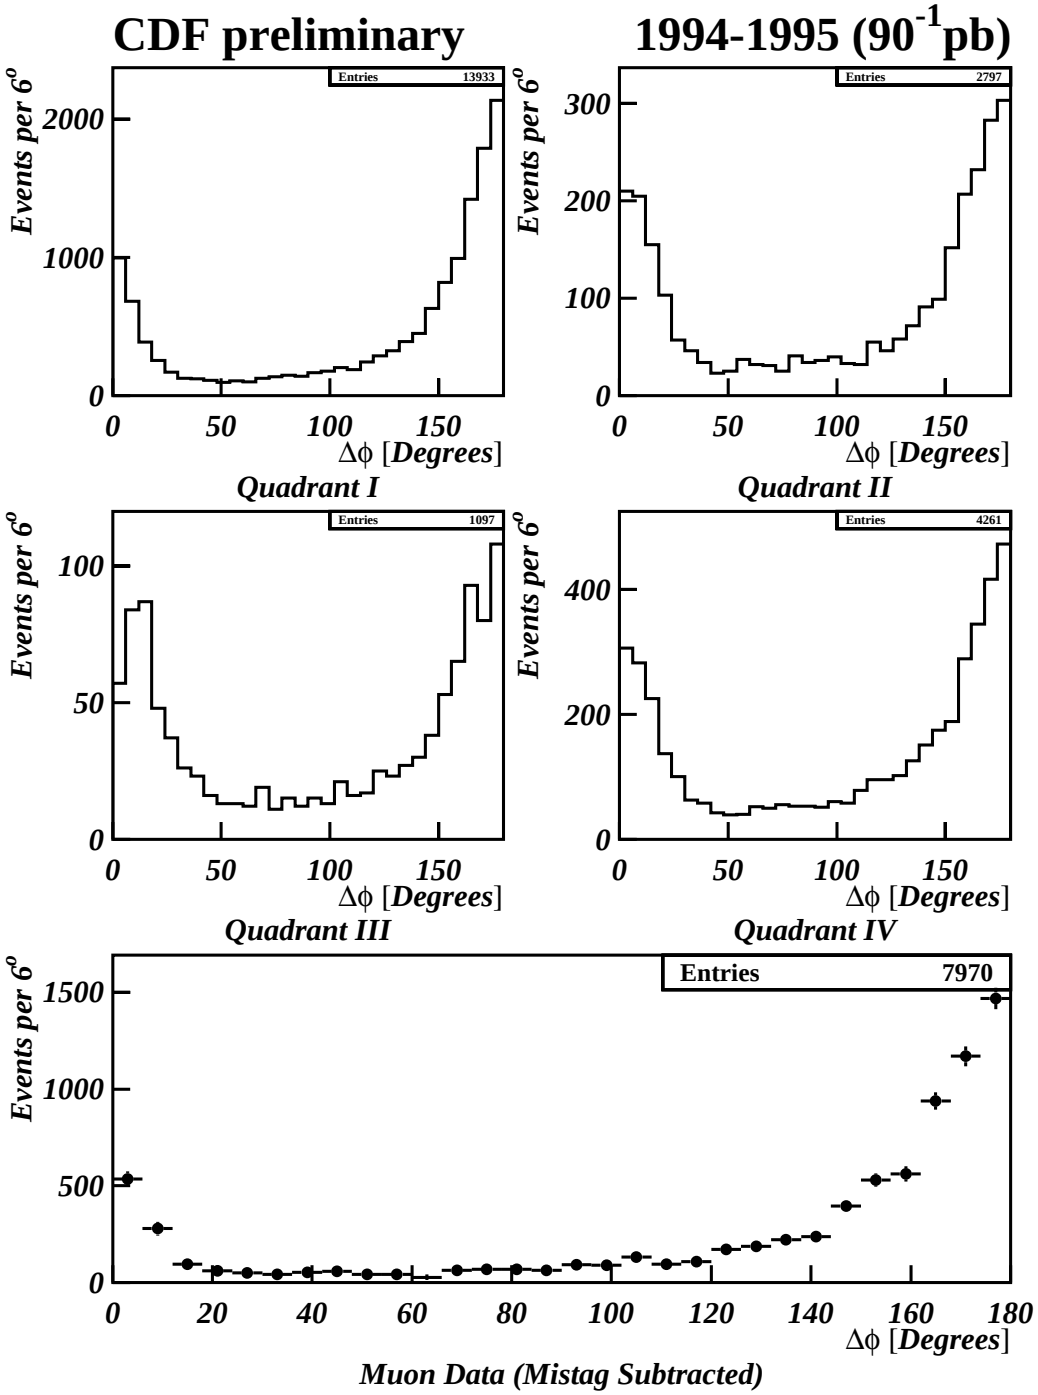

Figure 4.18 Mistag subtraction applied to the $\Delta\phi$ distribution from muon data. The upper four plots show the raw distributions for tag pairs from each quadrant of the L_{xy}(tag 2) versus L_{xy}(tag 1) plot. In the bottom plot, the distributions from each quadrant are subtracted according to the equation from Section 4.4.1 to yield the mistag subtracted plot.

the *B* and *D* decay vertices. Finally, it is possible that some sequential decay pairs arise from tracking errors that cause tracks actually originating from a common vertex to be reconstructed as coming from two vertices that are very close together.

Regardless of their origin, sequential tag pairs should share a few common characteristics. Because all the tracks from the two secondary vertex tags come from the same *B* hadron, the two vertices should be traveling in roughly the same direction, and thus have a small opening angle between them. This is the main reason for suspecting sequential tags as the source of the bump at low $\Delta\phi$ in Figure 4.12. Although we could potentially reduce the sequential tag contribution to our sample with a cut on $\Delta\phi$, this would hamper our ability to study correlations for $b\bar{b}$ pairs with small opening angle, one of the main goals of this analysis. Another feature of sequential tag pairs is that the invariant mass of the tag pair, in general, should be less than or equal to the mass of the *B* hadron that produced the tracks in the tag. A cut on invariant mass of the tag pair should be effective in eliminating sequential tags. Since most *B* hadrons have masses below 6 GeV/c^2, we call any tag pair with an invariant mass less than 6 GeV/c^2 a sequential tag pair. However, to preserve tagging efficiency, since presumably the two sequentials tags are tagging one valid *B* hadron, we only discard the tag in the sequential pair with the larger L_{xy}. Figure 4.19 shows the tag pair mass distribution before this cut.

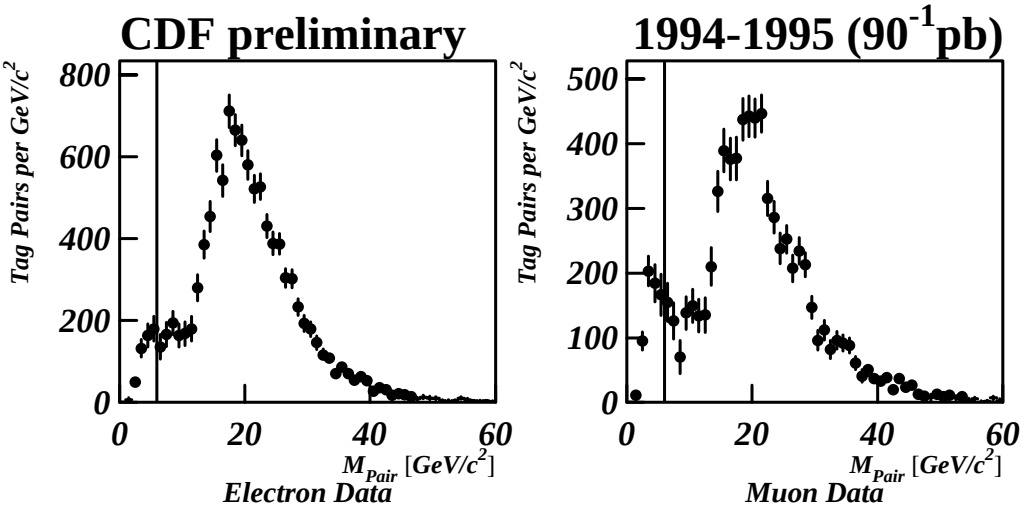

Figure 4.19 Invariant mass of the tag pairs. The cut to reduce sequentials is represented by the vertical line.

In principle, this cut would not have any effect on the distributions for secondary vertices originating from legitimate *B* hadron pairs if we were able to completely reconstruct the mass of each *B* hadron tagged, since the lower bound for the tag pair mass would be the sum of the *B* hadron masses. Because we only detect the charged decay products of the *B*, we only measure a fraction of the *B* hadron mass with the tagged tracks, as shown in Figure 4.11. However, the average tag mass is somewhat less than 3 GeV/c^2, so most legitimate tag pair originating from separate *B* hadrons should not be effected by our pair mass cut. Some valid tag pairs from *B* hadrons that are produced at small opening angles and that have small reconstructed tag masses, may be removed by this cut. This effect is modeled with Monte Carlo in the comparison between data and theoretical predictions.

Finally, to reduce the contribution from poorly measured tag pairs, tag pairs not sufficiently separated from one another are also removed. Tag separation is measured using the distance between the tags in the *x*-*y* plane, ΔL_{xy}, divided by the error on this quantity. ΔL_{xy} measures only the separation of the two tags from each other and makes no reference to the primary vertex location. All tags are required to be separated from each other by $\Delta L_{xy} / \sigma_{\Delta L_{xy}} > 2$. As in the case of the pair mass cut, only the tag with the larger L_{xy} is removed.

4.4.3 Prompt Charm

One background source of legitimate secondary decay vertices is $c\bar{c}$ production. In general, most *D* hadrons have a much smaller lifetime than *B* hadrons. However, those *D* hadrons that do live long enough to produce a secondary vertex capable of being tagged by BVTX will not be removed or accounted for by any of the methods mentioned above. In addition, it is possible to have events in which multiple heavy flavor pairs, such as $b\bar{b} + c\bar{c}$ are produced. For example, in a $b\bar{b}$ flavor creation event, an additional $c\bar{c}$ pair may be produced through gluon splitting. In such events it is possible for the $b\bar{b}$ to contribute one tag and the $c\bar{c}$ pair to contribute another. Although the rate of multiple heavy flavor production is much lower than single $b\bar{b}$ pair production, the opportunity to tag more displaced vertices in a given event can provide an enhancement in tagging efficiency, meaning such processes cannot be discounted outright.

It is important to point out that not all tags involving *D* hadrons should be considered background. Over 80% of *B* hadron decays involve the production of at least on *D* hadron [61]. Tagging these secondary *D* hadrons provides an important contribution to the legitimate $b\bar{b}$ production component of the double tag sample. In particular, *B* decays that produce a *D* hadron

plus only neutral final state particles can only be included by tagging the secondary *D* vertex. Therefore, whatever method is used to address the prompt charm contribution to the double tagged sample should not eliminate the contribution of secondary *D* decays.

For the *B* mixing analysis [58], it was determined that data from the BVTX sample containing at least one secondary vertex tag originated from $b\bar{b}$ production approximately 90% of the time. The other 10% of the data was consistent with being from $c\bar{c}$ production. This fraction was determined two different ways, using fits to the p_T(rel) and tag mass distributions. The quantity p_T(rel) is defined as the amount of lepton momentum perpendicular to the direction of the jet containing the secondary vertex tag. Because *B* hadrons are more massive than *D* hadrons, leptons resulting from *B* decays receive more of a transverse kick than *D* decay leptons. The $b\bar{b}$ and $c\bar{c}$ templates used for the p_T(rel) and tag mass fits were determined from Monte Carlo. It was also found that events containing more than one secondary vertex tag were even more likely to come from $b\bar{b}$ production. In [62] the fraction was determined to be approximately 95%.

The $b\bar{b}$ fraction determined for the *B* mixing analysis should be similar to the fraction for this analysis, although, since this analysis uses a slightly different version of the BVTX algorithm, differences are possible. Furthermore, since this analysis considers events with two secondary vertex tags, one has to consider processes where both a $b\bar{b}$ pair and a $c\bar{c}$ pair are produced, and we tag a *B* from the one pair and a *D* from the other. One way to estimate the size of the contribution from $c\bar{c}$ and $b\bar{b} + c\bar{c}$, is to see how much the tag pair distributions change if a mass cut of 2 GeV/c^2 is imposed on both tags in the pair. This cut should be almost completely effective in eliminating the *D* hadron contribution since almost all *D* hadrons have a mass less than 2 GeV/c^2 and the tag mass is typically only a fraction of the actual hadron mass. Unfortunately, since only a fraction of the tagged *B* mass is reconstructed, this cut also removes a good deal of the legitimate *B* tags. Only 2,234 electron events and 1,452 muon events remain if the tag masses for both tags in the pair are required to be greater than 6 GeV/c^2. On the other hand, 2,448 electron events and 1,867 muon events have both tags with a tag mass less than 6 GeV/c^2.

The results of a comparison between a sample enriched in prompt charm and a sample depleted in prompt charm is shown in Figure 4.20. This figure compares the $\Delta\phi$ distribution for a sample in which both tags have tag masses greater than 2 GeV/c^2 (sample depleted in prompt charm content) to a sample in which both tags have masses less than 2 GeV/c^2 (sample enhanced in charm content). As one can see, the distributions agree well within error bars over most of the

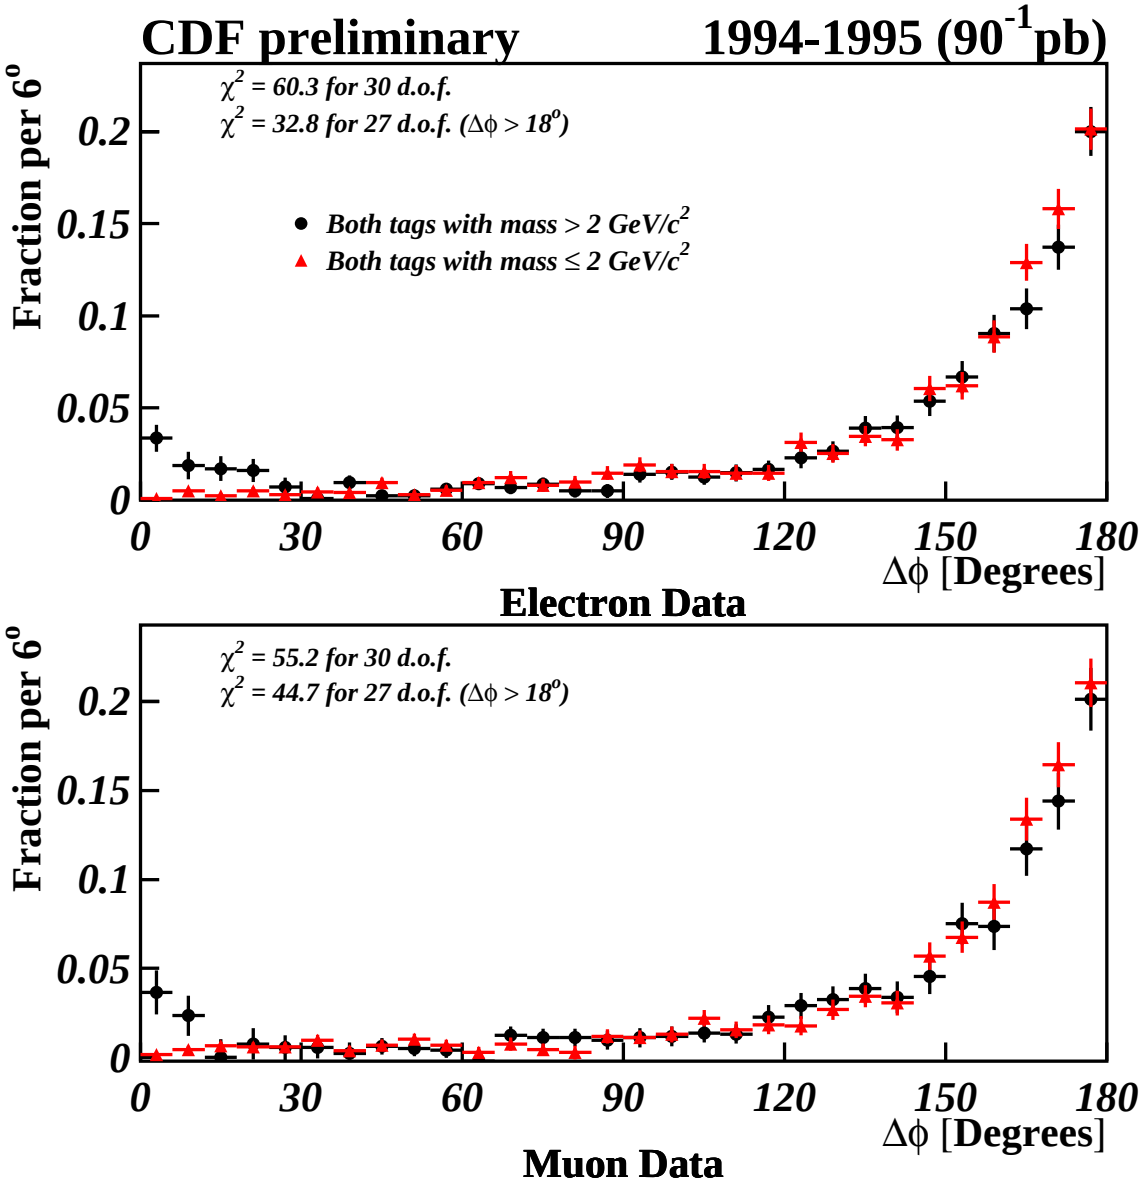

Figure 4.20 The $\Delta\phi$ distribution for events in which both tags pass the 2 GeV/c^2 mass cut (circles), and for event in which both fail the mass cut (triangles). The distributions are normalized to unit area to facilitate the comparison of their shapes. The χ^2 values on the plots indicate how well the shapes match, taking into account the uncertainties on both distributions. The χ^2 are given for both the entire distribution and the distribution excluding the first three bins where the effects of the 2 GeV/c^2 cut on each tag mass is most noticeable. The similarity between these distributions suggests that the contribution from charm is not large.

range, except in the region of small opening angle. The disagreement at small opening angle arises because of the pair mass cut imposed on the data to reduce sequential tags. For tags with small opening angle, this cut can eliminate tag pairs in which one or more of the tags would fail the 2 GeV/c^2 cut imposed in the comparison. Thus, fewer events are removed by the cut in this region, than in the regions with larger opening angle, leading to the disagreement shown in the plots. Based on the consistency shown in Figure 4.20, we conclude that the contribution to the double tag sample from charm is not significant.

Another method for estimating the approximate size of the contribution from prompt charm to the double-tag sample involves studying the distribution of tag masses for single-tag events versus the distribution for double-tag events. For example, consider the tag mass distribution for tags associated with the trigger lepton—the tag containing the trigger lepton or with the smallest ΔR from the trigger lepton if the trigger lepton track is not included in any tag. For all trigger lepton tags, regardless of the presence of a second tag, the tag mass distribution will contain a mix of tags from *B* decays and tags from prompt *D* decays. However, if only lepton tags from double-tagged events are considered, the contribution from prompt *D* decays is reduced because of the smaller chance that $c\bar{c}$ production will yield a second tag. If the double-tag mass distribution is normalized so that its shape matches the shape of the single-tag distribution for tag masses above 2 GeV/c^2, the excess in the single-tag distribution below 2 GeV/c^2 yields the approximate size of the prompt charm contribution removed by demanding the presence of a second tag. This comparison is shown in Figure 4.21.

Although the above studies do not establish the exact size of the prompt charm contribution to the double-tagged sample, they do suggest that this contribution should be relatively small. For this reason, we will ignore the prompt charm contribution for the time being. This issue will be revisited in Section 7.2.2 when Monte Carlo is used to estimate roughly the size of the correction required for prompt charm and the systematic uncertainty associated with this correction.

4.4.4 Fake Leptons

This analysis attempts to select events in which a *B* hadron underwent a semileptonic decay to produce a lepton satisfying the trigger selection criteria. However, it is possible that the object selected by the trigger was not a lepton from a semileptonic *B* decay. Leptons not from *B* decays can satisfy the trigger, although most such events should be eliminated by the double secondary vertex tag requirement. It is also possible for a non-leptonic particle from a *B* decay to fake a

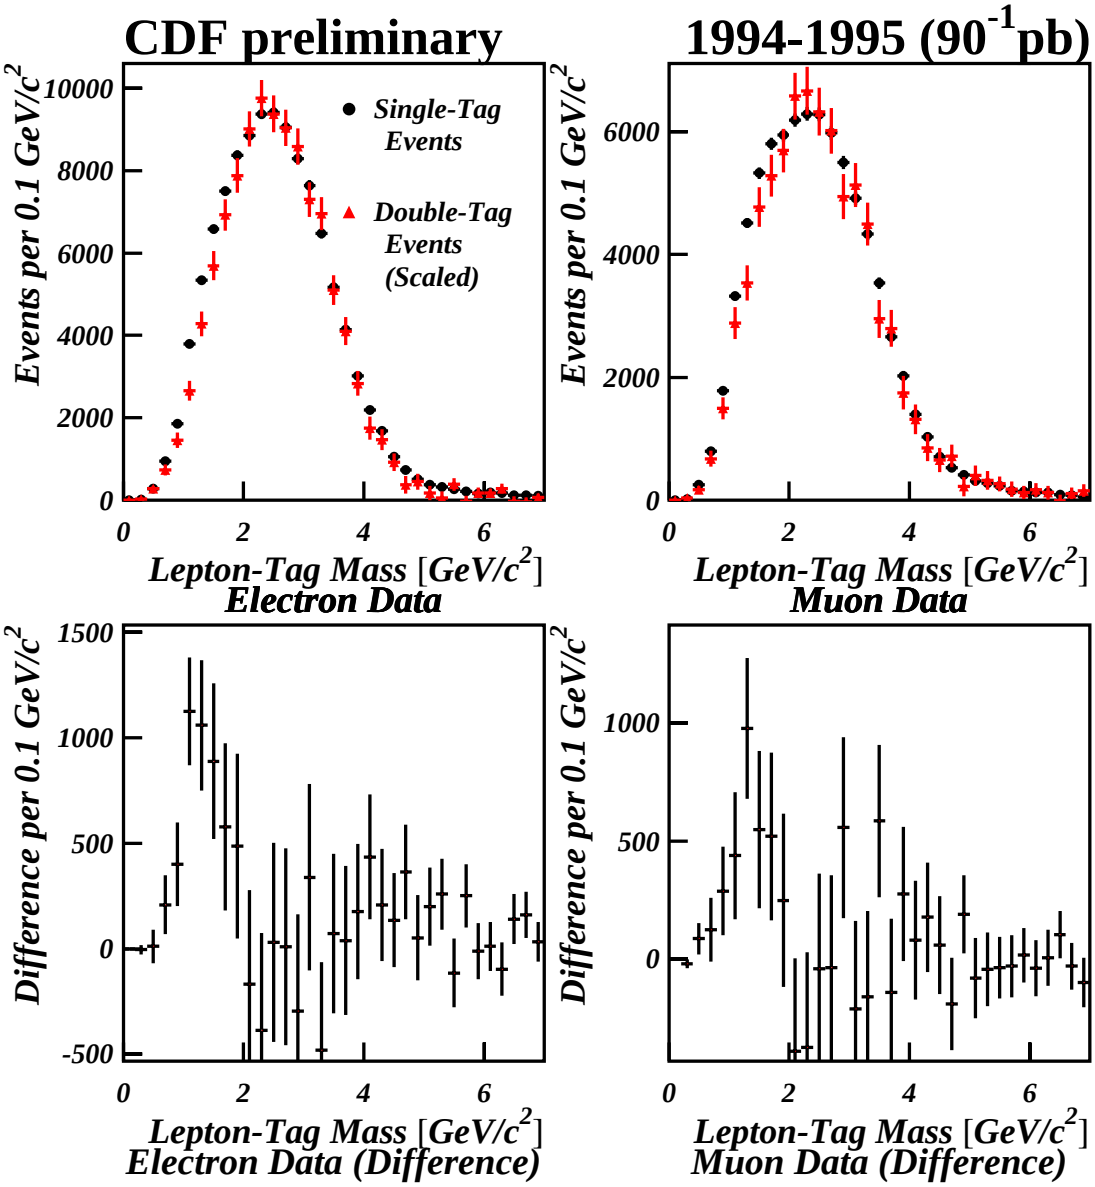

Figure 4.21 The upper plots show the tag mass for the lepton tag, where the lepton tag is defined to be the tag containing the trigger lepton, or the tag with the minimum ΔR separation from the lepton if no tag contains the trigger lepton. The black circles show the distribution for all events containing at least one BVTX tag. The red triangles show the distribution for events with at least two BVTX tags. The double-tag distribution has been scaled to match the single-tag distribution in the region with $m_{lepton\text{-}tag} > 2$ GeV/c^2. The lower plots show the difference between the single-tag distribution and the scaled double-tag distribution. The excess below 2 GeV/c^2 indicates the amount of prompt charm contribution removed by the two BVTX tag requirement.

lepton. The amount and type of fake trigger leptons can be different for the electron and muon data samples. In the time-dependent B mixing analysis [58], the issue of fake leptons was investigated for the BVTX data sample used for this analysis. The results are summarized below:

There are two main sources of fake trigger electrons: electrons produced in photon conversions, and hadronic tracks which fake electrons. In this analysis, photon conversions are removed using the same cuts as [58]. It is estimated that 1% of the electron data remaining in the BVTX sample after conversion removal is from conversions. The contribution of hadrons faking trigger electrons is determined using a fit to energy loss, dE/dx of the trigger electron candidates. Particles with different masses lose energy at different rates when traversing the detector material. Although the particle mass separation using dE/dx information at CDF is not adequate to differentiate particle species on a per particle basis, the mix of particle species in a sample of tracks can be statistically evaluated. The results of this fit for electron candidates indicate that 0.6% of the electron data result from hadrons that fake electrons. The contributions to the electron data from photon conversions and hadrons faking electrons are small enough that these contributions may be ignored.

The main source of fake muons that concerns this analysis are hadrons that generate a fake muon signature either through *decay in flight* or *punch-through*. Decay in flight results from the decay of a pion or kaon into a final state involving an actual muon. Although these decay-in-flight muons are real muons, they don't come from the semileptonic decay of a B and are therefore considered background for this analysis. Punch-through occurs when a hadron does not loose enough energy in the calorimeter and manages to penetrate into the muon chambers. Punch-through backgrounds cannot be reduced using dE/dx as in the electron case, because the separation between muons and pions is insufficient for the CDF detector. In [62], the contribution from fake muons is determined using the effectiveness of the flavor tagging for the muon sample. The contribution is determined to be $12 \pm 6\%$. Another source [63] uses fits to p_T(rel) and tag mass to determine the fake muon component. P_T(rel) is defined as the amount of momentum from the track believed to be the muon perpendicular to the jet direction. The fake muon distributions for p_T(rel) and tag mass were determined from a sample of fake muons. This sample was created in two ways. One was to look for muons that passed all the muon quality cuts but which don't have a defined z position because charge division procedure failed. When a hadron punches through the back calorimeter, the resulting shower of particles into the muon

chambers typically saturates the signals in the CMU making it impossible to determine the z position of the muon hits. Therefore, CMU muons for which the charge division fails are typically the result of hadronic punch-through. The other method for finding fake muons is to look for tracks which are matched to hits in the CMU chamber but which failed to leave hits in the CMP despite being contained within the CMP detector's geometric acceptance. Such a signal usually indicates that a hadron punched through the back of the calorimeter, leaving a signal in the CMU chamber, but failed to penetrate the shielding to CMP chamber. From fits to fake muon templates [63] determined the fake muon contribution to the BVTX data to be approximately 4 ± 1%.

However, comparisons made between secondary vertex tags found in the fake muon sample, versus distributions generated from $b\bar{b}$ Monte Carlo suggest that most of the secondary vertices in the fake muon samples actually come from heavy flavor (see Figure 4.22). As a result, it is safe to neglect the contribution from fake muons to the double-tagged sample since most of the secondary vertices found in the fake muon data actually arise from heavy flavor production.

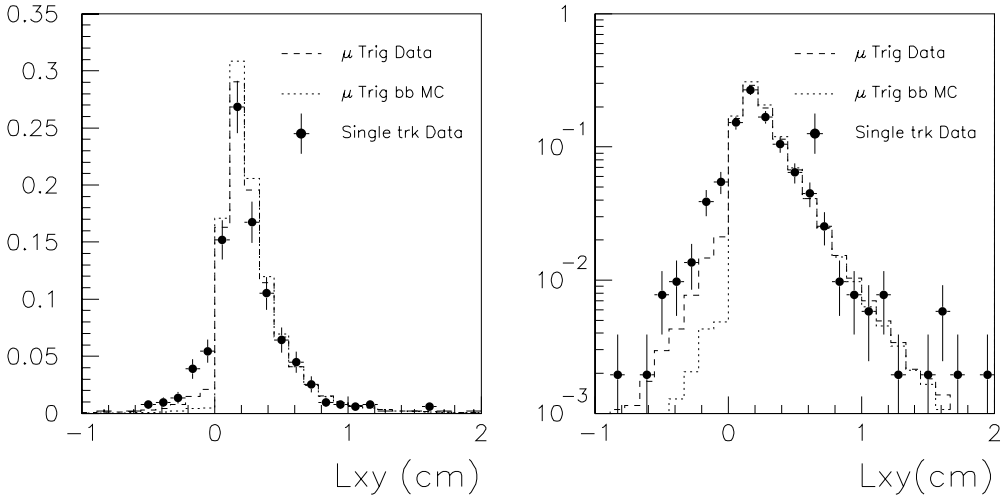

Figure 4.22 L_{xy} distribution of the fake muon data compared to the trigger muon data and $b\bar{b}$ Monte Carlo taken from [62]. The similarity among the distributions suggests that the fake muon data has a heavy flavor origin.

4.5 Final Sample Selection Cuts

In the end, the sample with the smallest contribution from background sources is determined to be the one satisfying the following criteria:

- Trigger lepton selection criteria, as defined in Section 4.1 applied to all events.
- All events required to contain at least two reconstructed secondary vertices passing all cuts described in Section 4.2.
- One tag from pairs failing the $m_{pair} > 6$ GeV/c^2 and the $\Delta L_{xy} / \sigma_{\Delta Lxy} > 2$ cut removed.
- At least one BVTX secondary vertex tag passing all cuts must be within $\Delta R = 1.0$ of a trigger lepton satisfying the above criteria.
- Mistag subtraction, as defined in Section 4.4.1 applied to all distributions of interest.

Figure 4.23 and Figure 4.24 show the tag pair distributions with all these cuts applied. These plots contain 10,097 electron events and 7,092 muon events.

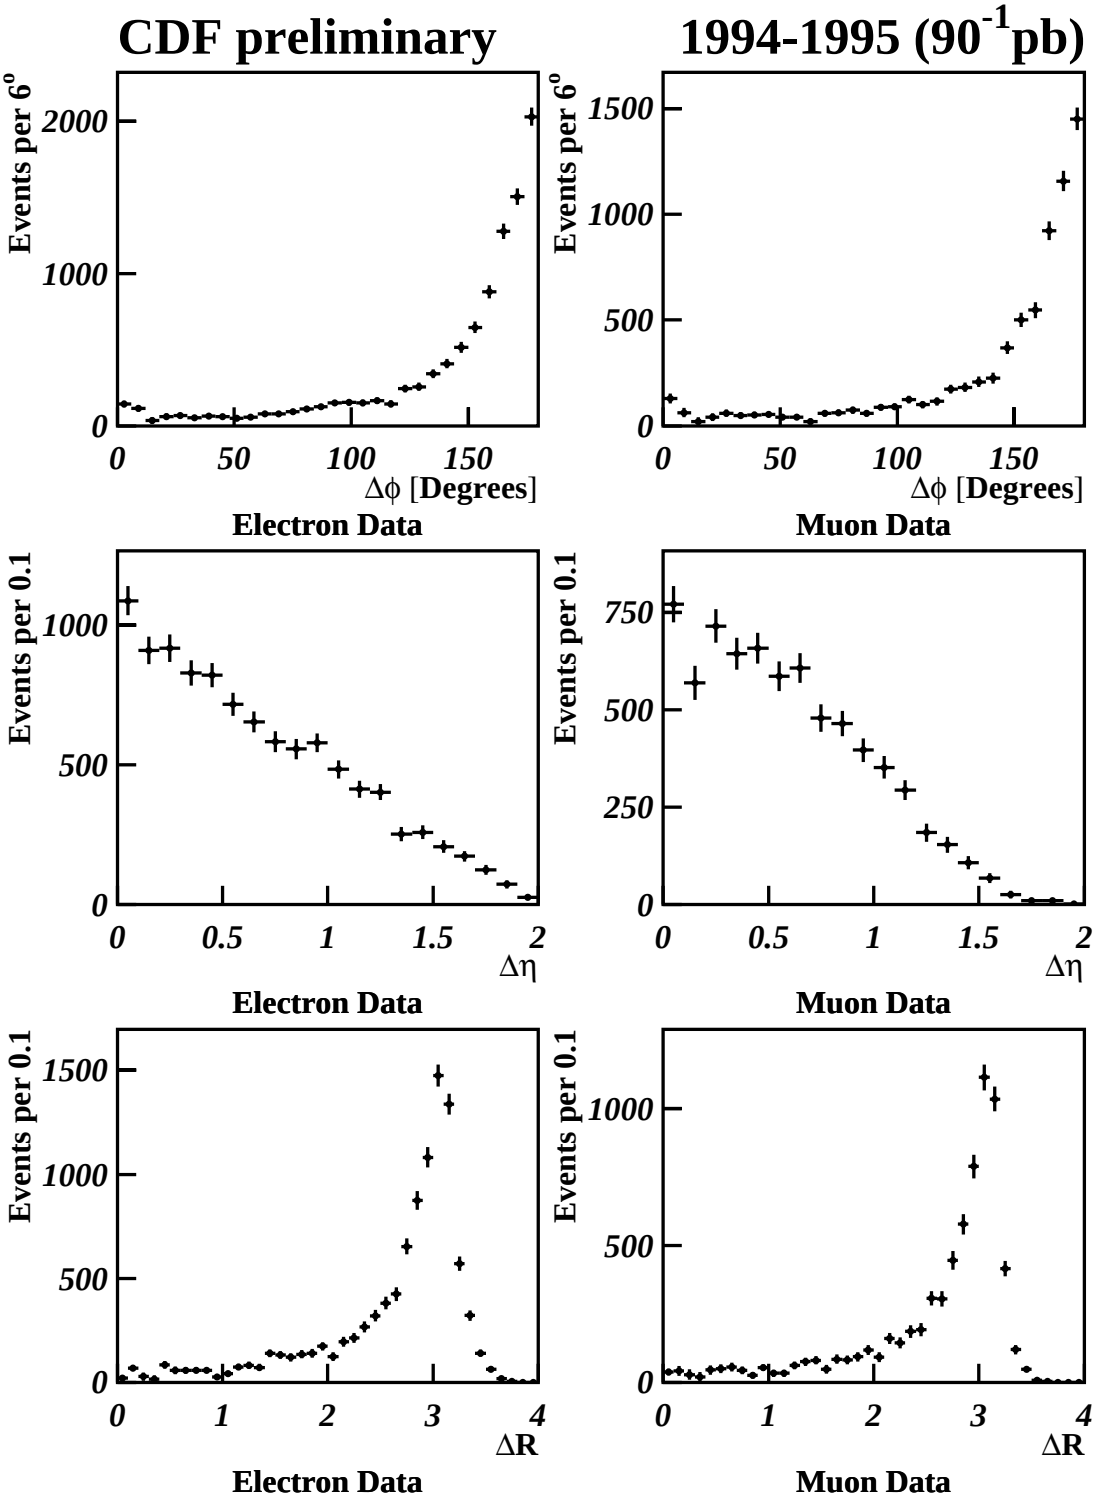

Figure 4.23 Tag pair angular correlation distributions with all cuts from Section 4.5 applied.

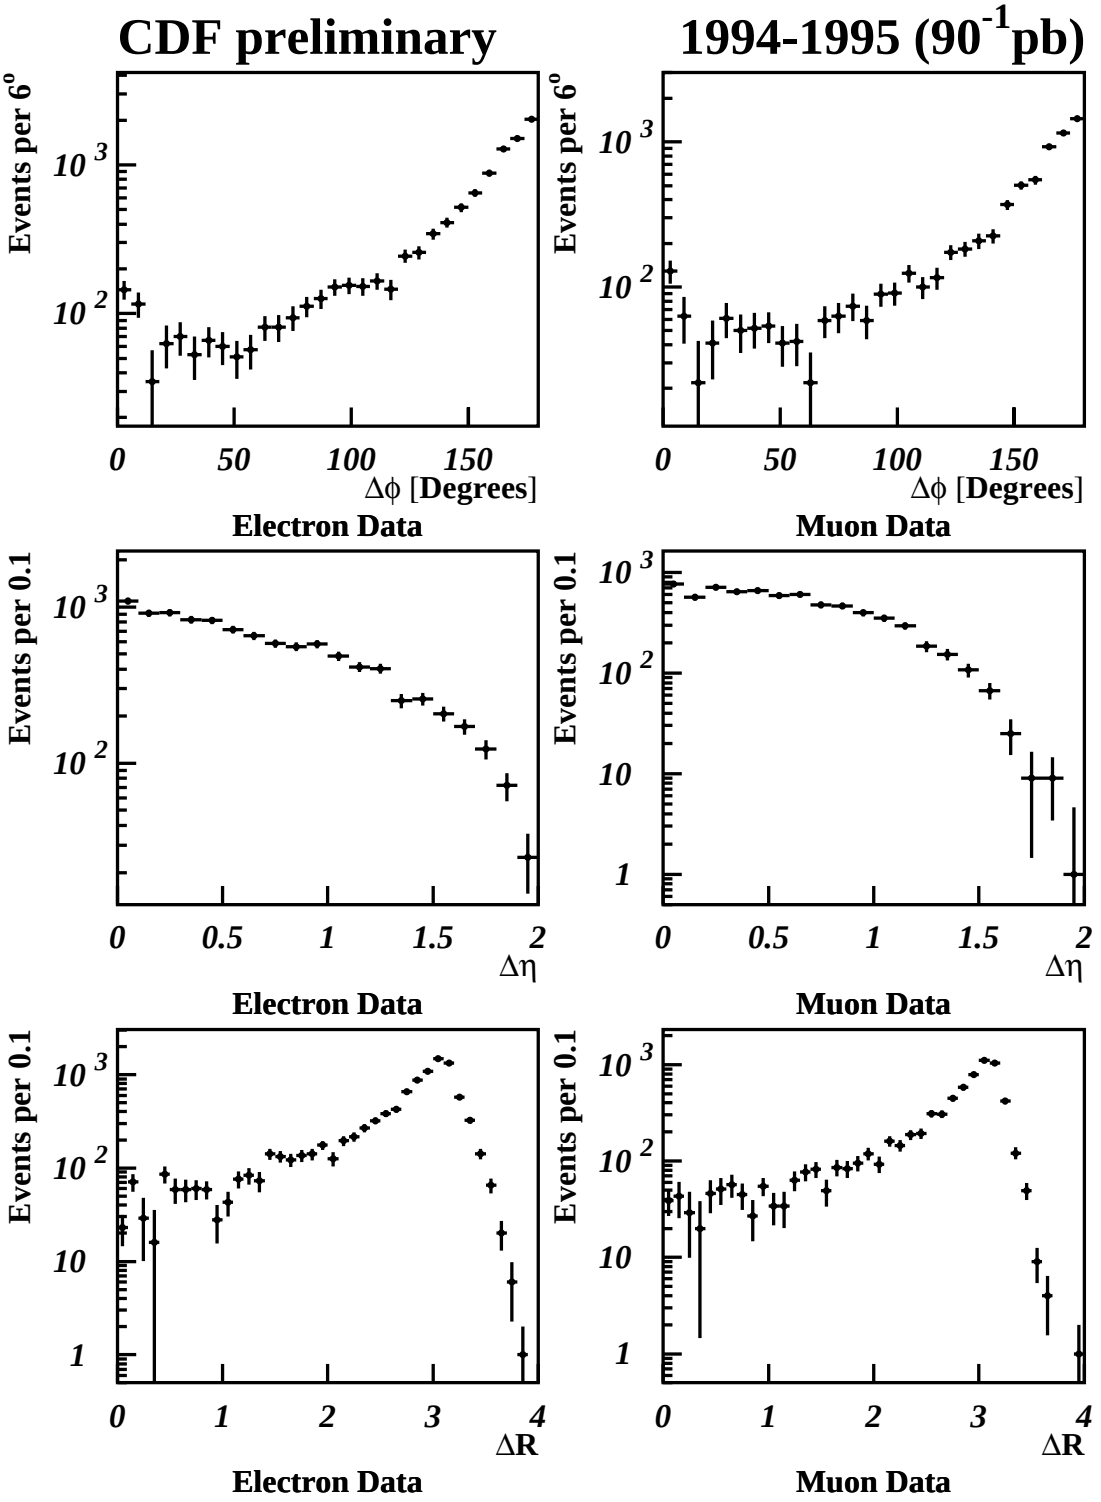

Figure 4.24 The same as the previous figure, except show with a semi-log scale.

Chapter 5

Monte Carlo Samples

Although analytic calculations of $b\bar{b}$ production exist, it is often difficult to apply such calculations to make predictions when nontrivial requirements are placed on the final state kinematics of the $b\bar{b}$ system. Furthermore, frequently one desires not only to know the kinematic distributions of the *b* quarks, or their associated hadrons, as predicted by QCD, but also to understand the effect that the detector and the analysis cuts will have on such distributions. In these cases, it is helpful to turn to Monte Carlo simulation.

Strictly speaking, the term "Monte Carlo" refers to the technique of performing difficult integrals using random sampling. In practice however, the term has come to represent the solution of any problem regarding physics processes or detector effects that involves randomly sampling various distributions. Typically, Monte Carlo simulations are broken into independent parts. Generation of a Monte Carlo sample begins with the simulation of one or more desired physics processes. Frequently, additional Monte Carlo simulations may be used to refine certain aspects of the physics, such as using a specialized simulation to handle the decays of certain particles. Finally, when it is necessary to account for the effect of the detector on the physics distributions, a separate detector Monte Carlo simulation is used to model the effects of detector responses, acceptances, and efficiencies.

For this analysis, $b\bar{b}$ events are generated using two commonly used, general-purpose Monte Carlo programs, PYTHIA and HERWIG. PYTHIA and HERWIG are FORTRAN programs that generate a wide variety of physics processes both from e^+e^- and $p\bar{p}$ colliders using the parton shower model described in Section 2.3.2. They both simulate physics processes using leading-order matrix elements, supplemented by initial- and final-state radiation generated in the parton shower model. In addition, both programs include the effects of fragmentation and hadronization of the final-state quarks and gluons as well as the effects of the beam remnants that are left when a parton from a particle in each colliding beam is removed to participate in the hard QCD interaction. As a result, PYTHIA and HERWIG produce realistic, multi-particle final states

that reproduce many of the characteristics observed in hadron collider data. Although PYTHIA and HERWIG share many common features, they have enough differences—particularly in the areas of parton shower implementation and fragmentation models—that it is worthwhile to compare both models to the data.

As discussed in Section 2.3.2, because only leading-order matrix elements are used, PYTHIA and HERWIG neglect interference effects that would be accounted for in Monte Carlo programs based on a full NLO calculation. Unfortunately, the currently available NLO Monte Carlo program generates weighted parton-level events rather than hadronic final states. A small number of events with large negative weights are used to cancel out effects from divergences in certain parts of phase space. Adding the necessary fragmentation and detector effects to these weighted parton-level events can disturb the delicate cancellation between positive and negative weights causing unphysical fluctuations in some distributions. Because it is important to account for hadronization, decay, and detector effects for this analysis, PYTHIA and HERWIG are used in favor of the available NLO Monte Carlo programs.

For this analysis, PYTHIA and HERWIG are used to generate large samples of events from $b\bar{b}$ production. In each program, special care must be taken to ensure that all three $b\bar{b}$ production mechanisms (flavor creation, flavor excitation, and gluon splitting) are included in generation. PYTHIA contains many options and parameters that users can vary, so three $b\bar{b}$ samples are generated using PYTHIA with different settings to compare to experiment. HERWIG, on the other hand, does not allow the user as much freedom in tuning parameters, so only one sample was generated using HERWIG. In addition, a small sample of $c\bar{c}$ events was generated using a particular set of parameters in PYTHIA, for the purposes of evaluating the possible effects of residual prompt charm as a background for this analysis. Details of the generation of Monte Carlo samples for this analysis are given below:

5.1 Basics

PYTHIA and HERWIG generate $b\bar{b}$ events using only leading-order matrix elements for the hard scatter in the event. This means that the hard scatter can only involve processes with two partons in the initial state and two partons in the final state, so called two-to-two processes. Processes involving three or more partons in the hard scatter final state (two-to-three processes) require higher order matrix elements.

The restriction of PYTHIA and HERWIG to two-to-two hard processes allows all $b\bar{b}$ events generated by these Monte Carlo programs to be classified unambiguously as flavor creation, flavor excitation, or gluon splitting based on the number of *b* quarks involved in the hard scatter process:

- **Two *b* quarks in the hard scatter** (see Figure 5.1): Processes involving two *b* quarks in the hard scatter, such as $q\bar{q} \to b\bar{b}$ or $gg \to b\bar{b}$ are classified as **flavor creation**. This classification includes the process $b\bar{b} \to b\bar{b}$ in which the initial-state *b* quarks come from the splitting of a gluon in each beam particle through the evolution of the parton distribution functions for the proton, as discussed in Section 2.4.1. Note, in this last case, there are actually two *b* quarks and two $\bar{b}$ quarks produced in the final state, one each from the hard process, and one each from initial-state parton showers (see **flavor excitation** below for more details).

- **One *b* quark in the hard scatter** (see Figure 5.1): QCD processes with only one *b* quark in the final state (and necessarily, one in the initial state as well), such as $qb \to qb$ or $g\bar{b} \to g\bar{b}$, are classified as **flavor excitation**. The initial-state *b* quark is selected from the initial state of one of the beam particles through the parton distribution functions. However, QCD must conserve quark flavor, meaning that *b* quarks cannot be produced singly in QCD processes. The second *b* quark in the event comes from the initial-state parton shower generated by PYTHIA or HERWIG for the beam particle supplying the *b* quark for the hard scatter. Because protons do not carry any bottom valence flavor, any *b* quark discovered within the proton at a given momentum scale, must come from the splitting of a gluon into a $b\bar{b}$ pair (see Section 2.4.1). PYTHIA and HERWIG take this into account by using their initial state parton showering models to trace the evolution back to locate this $g \to b\bar{b}$ splitting. It should be noted that the processes $bb \to bb$ and $\bar{b}\bar{b} \to \bar{b}\bar{b}$, which each produce two *b* quarks and two $\bar{b}$ quarks in the final state, are classified as **flavor excitation** for this analysis.

- **No *b* quarks involved in the hard scatter** (see Figure 5.1). In PYTHIA and HERWIG, $b\bar{b}$ events in which no *b* quark participates in the hard scatter are classified as **gluon splitting**. The *b* quarks in such events come from a $g \to b\bar{b}$ splitting. Although in

principle, such a splitting can occur in either the initial-state or final-state parton showers, kinematics favor splittings in final-state showers.

Technically, it is possible for an event to contain more than one $b\bar{b}$ pair and thereby fall into more than one category, particularly in the case where a gluon splitting $b\bar{b}$ pair is added to either a flavor creation or flavor excitation event, as depicted in Figure 5.2. For such events, only information about *b* quarks involved in the hard scatter is used to classify the event. In any case, production of multiple $b\bar{b}$ pairs in a single event occurs only at a low rate compared to single $b\bar{b}$ production.

One way to generate $b\bar{b}$ events from all three production mechanisms is to have PYTHIA or HERWIG generate all possible QCD two-to-two processes and select only events in which a $b\bar{b}$ pair is produced. However, because *b* quarks are so much heavier than light quarks (*u*, *d*, *s*), most events produced will only contain light quarks in the final state. Fortunately, there exist

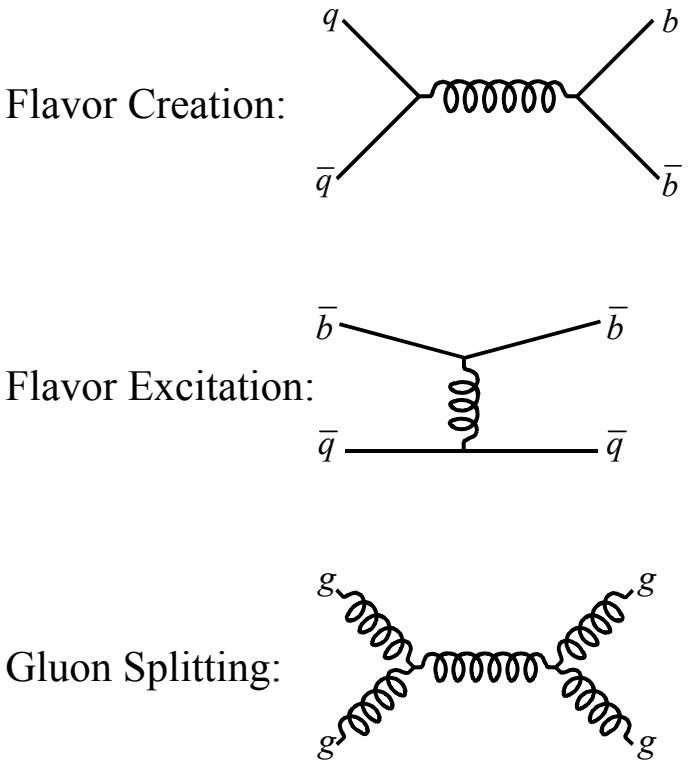

Figure 5.1 Feynman diagrams representing the leading-order matrix elements used by PYTHIA and HERWIG to represent the hard scatter part of $b\bar{b}$ production.

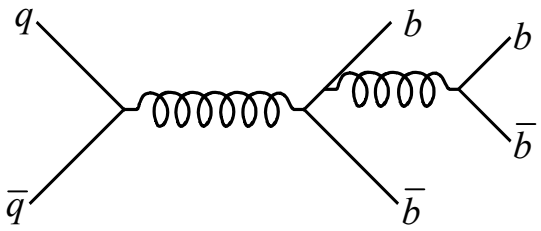

Flavor Creation + Gluon Splitting
(Classified as Flavor Creation)

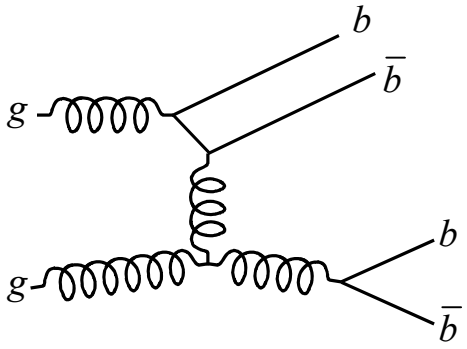

Flavor Excitation + Gluon Splitting
(Classified as Flavor Excitation)

Figure 5.2 Examples of $b\bar{b}$ production involving more than one production mechanism.

techniques for generating flavor creation and flavor excitation events with much higher efficiency. These techniques are specific to each generator and are discussed below. Unfortunately, the only way to generate gluon splitting $b\bar{b}$ production with PYTHIA or HERWIG is to use the inefficient technique of generating all possible QCD processes and retaining only events that contain $b\bar{b}$ pairs generated in the parton showers of such events.

5.2 PYTHIA Generation

Version 6.203 of PYTHIA is used for this analysis. However, this version of PYTHIA contained a bug that caused *b* quarks produced in flavor excitation events to be treated as massless in the final state half of the time [64]. For this analysis, the bug was fixed manually by editing the

PYTHIA source code. As of version 6.208, the bug fix has been implemented in the official PYTHIA code.

5.2.1 Production Mechanisms

It is most convenient to generate samples for each $b\bar{b}$ production mechanism separately using different PYTHIA configurations and to obtain the total sample by combining the individual subsamples, using the cross section reported by PYTHIA for each sample to determine the relative normalizations. In the leading-order matrix element plus parton shower approach used by PYTHIA, interference between these three production can be neglected mechanisms since it appears as a higher-order correction.

There are two QCD processes simulated by PYTHIA (selectable using the `MSEL` variable) that are relevant to $b\bar{b}$ production. The most basic process is PYTHIA's generic QCD process (`MSEL` = 1). This process uses leading-order matrix elements to calculate generic two-to-two QCD scattering in the approximation of massless quarks. Using this process, it is possible to generate the full set $b\bar{b}$ production mechanisms, although this approach is not very efficient because most events generated using PYTHIA's generic QCD process contain only light quarks. However, PYTHIA also has a special, heavy-flavor production process (`MSEL` = 5) that only generates events using the matrix elements for $q\bar{q} \to b\bar{b}$ and $gg \to b\bar{b}$, in other words, the flavor creation diagrams. In this process, the *b* quark mass is retained in the matrix elements, which can have a significant effect on the *b* quark production rate near the mass threshold. These two processes can be used together to generate the full set of $b\bar{b}$ production mechanisms with reasonably high efficiency as follows:

- **Flavor Creation:** In Pythia, flavor creation is generated simply by running the heavy flavor production process. Every event generated is a flavor creation $b\bar{b}$ event, and *b* quark mass effects are included in the matrix element calculations.

- **Flavor Excitation:** The generic QCD process must be used in PYTHIA to generate flavor excitation events. However, PYTHIA has a set of parameters (stored in the `KFIN` array), that allow the user to specify that only certain quark flavors be available in the beam particles to determine the hard scatter initial state. By turning off all flavors except bottom in one of the beams—say the proton beam—the generic QCD process generates

flavor excitation events with nearly 100% efficiency.[1] Using two runs, one with only b and $\bar{b}$ quarks available in the proton beam and one with only b and $\bar{b}$ quarks available in the antiproton beam, the entire flavor creation process can be generated with reasonable speed. However, note that for the purpose of the hard scattering matrix element, the b quark involved is treated as massless. In the final state, the b quark is given a mass and the momentum vector is scaled down in the direction of the b quark motion to conserve energy and momentum.

- **Gluon Splitting:** Gluon splitting events must be generated by running the generic QCD generation process, since $b\bar{b}$ pairs can arise from gluon splitting through a parton shower associated with any QCD event. The gluon splitting contribution is extracted by discarding events that don't contain a $b\bar{b}$ pair, or events in which one or both of the quarks in the $b\bar{b}$ pair participate in the hard scatter. This method for generating gluon splitting events is unfortunately very inefficient and requires significant processing power to achieve reasonable statistics.

In addition to the choice of process generated, PYTHIA provides users with a number of additional parameters that can be tuned to affect how well the Monte Carlo events match data. The PYTHIA program has literally hundreds of tunable parameters and switches; however, most are not relevant to b quark production. Below, we address parameters that were adjusted for this analysis. Other parameters were left at their default values [33].

5.2.2 Initial-State Radiation

One aspect of PYTHIA generation that has a marked effect on $b\bar{b}$ angular correlations is the amount of initial-state radiation PYTHIA generates. Increasing the amount of initial-state radiation broadens the back-to-back peak from flavor creation. Increased initial-state radiation can boost b quark, and hence B hadron, p_T, leading to more B hadrons above the minimum p_T threshold for detection. Finally, raising the amount of initial-state radiation directly effects the gluon splitting portion of the $b\bar{b}$ cross section because more energetic gluons in the event leads to a greater probability of producing a $b\bar{b}$ pair from a gluon branching.

[1] A small number of flavor creation events are generated through the process $b\bar{b} \rightarrow b\bar{b}$ that must be discarded to avoid double counting events generated in the flavor creation sample.

In the parton shower model, the correct scale at which to begin the perturbative evolution of the initial-state shower is not known from first principles. PYTHIA provides the users with a variable to tune this scale, known as `PARP(67)` The value of `PARP(67)` is multiplied by the Q^2 of the hard scatter to determine the maximum virtuality of the initial-state shower. Higher values of `PARP(67)` lead to more initial-state radiation. The default value of `PARP(67)` = 1.0 for the latest versions of PYTHIA (any after 6.138) comes from studies involving heavy quark production [65]. However, the older versions of PYTHIA had a higher amount of initial-state radiation, with a default value of `PARP(67)` = 4.0. In all likelihood, the best value lies somewhere between 1.0 and 4.0. For this analysis, multiple samples were generated with different amounts of initial-state radiation corresponding to `PARP(67)` = 1.0, 3.0, and 4.0.

5.2.3 Parton Distribution Functions

The choice of parton distribution function primarily affects the amount of flavor excitation predicted in PYTHIA. Changing the parton distribution set can yield higher or lower contributions from flavor excitation to the total amount of $b\bar{b}$ production. For this analysis, we use the CTEQ5L parton distribution functions. This choice was motivated by [20], which uses the CTEQ parton distributions to obtain reasonable agreement between the PYTHIA predictions for the inclusive *b* quark cross section and experimental measurements.

5.2.4 Underlying Event

The term "underlying event" refers to all the other activity associated with a proton-antiproton collision beyond the hard QCD interaction. At a minimum, this activity includes the breakup of the proton and antiproton remnants after the hard scatter, and any initial- and final-state radiation in the event. There is also good evidence that additional semi-hard parton interactions contribute to the underlying event [66]. In PYTHIA, this effect is labeled "multiple parton interactions" (MPI), or frequently, just "multiple interactions," not to be confused with the term "multiple interactions" referring to multiple proton-antiproton collisions in the same event. (In the PYTHIA documentation, events containing multiple proton-antiproton collisions are referred to as "pile-up events.") PYTHIA has a number of parameters that allow one to tune the underlying event to match the data. The default PYTHIA settings do not match CDF data very well. For this analysis,

we have used PYTHIA settings determined by [66] from studying minimum-bias and underlying event data from CDF.

The exact tuning of the underlying event in PYTHIA depends on the amount of initial-state radiation, as determined by `PARP(67)`, and the choice of parton distribution functions, in this case CTEQ5L. Table 5.1 shows the parameters used to tune the underlying event for the two different initial state radiation settings used to generate the PYTHIA samples. It is important to realize that a change in either the initial state radiation or the parton distribution functions used requires completely different values of the these PYTHIA parameters to match the underlying event.

5.2.5 Minimum Hard Scatter p_T

PYTHIA allows the user to specify a minimum p_T for the final-state partons to be generated in the parton-parton center of mass frame of the hard scatter. This parameter is called p_T(min). When

Parameter	**Meaning**	**PARP(67) = 4.0, 3.0**	**PARP(67) = 1.0**
MSTP(81)	Multiple-parton interaction switch	1 (Multiple Parton Interactions ON)	
MSTP(82)	Model of multiple parton interactions	3 (Varying impact parameter assuming a single Gaussian matter distribution)	
PARP(82)	p_T turn-off when using single Gaussian model of multiple interactions	1.7	1.6
PARP(85)	Probability that a multiple parton interaction produces two gluons with color connections to the "nearest neighbors"	1.0	
PARP(86)	Probability that an MPI produces two gluons either as described above or as a closed gluon loop. The rest of the MPIs produce quark-antiquark pairs	1.0	
PARP(89)	Determines the reference energy E_0	1800.	

Table 5.1 The table above shows the PYTHIA setting used to tune the underlying event to data for the CTEQ5L parton distribution set and three different initial-state radiation settings. For more details consult the PYTHIA manual [33].

the user specifies a value for p_T(min) greater than zero, no events are generated in which the hard scatter final state partons have less than p_T(min). This parameter does not actually change the nature of the physics generated by PYTHIA; it merely restricts the phase space available for the final state. For this analysis, this parameter is primarily useful in eliminating regions of phase space that are unlikely to yield events that will be successfully reconstructed. The appropriate value for p_T(min) depend on the $b\bar{b}$ production mechanism being generated as well as on the amount of initial state radiation. For flavor creation, p_T(min) is very closely related into the minimum b quark p_T in the lab frame because the b quarks are both produced directly in the hard scatter. However, in flavor excitation and gluon splitting, one or more of the b quarks are produced indirectly. In general, this means that p_T(min) for these processes can be set higher than for flavor creation. Finally, with more initial-state radiation, partons can receive a greater boost between the hard scatter center of mass frame and the lab frame. Therefore, events with lower b quark p_T will be reconstructed by the analysis code, requiring a lower value for p_T(min).

For this analysis, the value of p_T(min) is tuned to try to ensure the most efficient generation of $b\bar{b}$ events that will be reconstructed by the analysis code without distorting the generated distributions too much by omitting relevant portions of phase space. Table 5.2 shows the acceptable values for p_T(min) determined for this analysis as a function of different amounts of initial state radiation. Figure 5.3 and Figure 5.4 show the p_T distribution of the hard-scatter partons in the center of mass frame for flavor creation, flavor excitation, and gluon splitting events reconstructed with B hadron p_T in a range appropriate for this analysis.

5.2.6 Samples Generated

This analysis uses three separate PYTHIA $b\bar{b}$ samples, generated with differing amounts of

Process	**PYTHIA, `PARP(67)` = 1.0**	**PYTHIA, `PARP(67)` = 4.0, 3.0**
flavor creation	8 GeV/c	8 GeV/c
flavor excitation	10 GeV/c	8 GeV/c
gluon splitting	15 GeV/c	12 GeV/c

Table 5.2 The p_T(min) used to generate PYTHIA samples for this analysis. The values are lower for the cases with more initial-state radiation because the b quarks get an additional p_T kick over the case with less initial-state radiation.

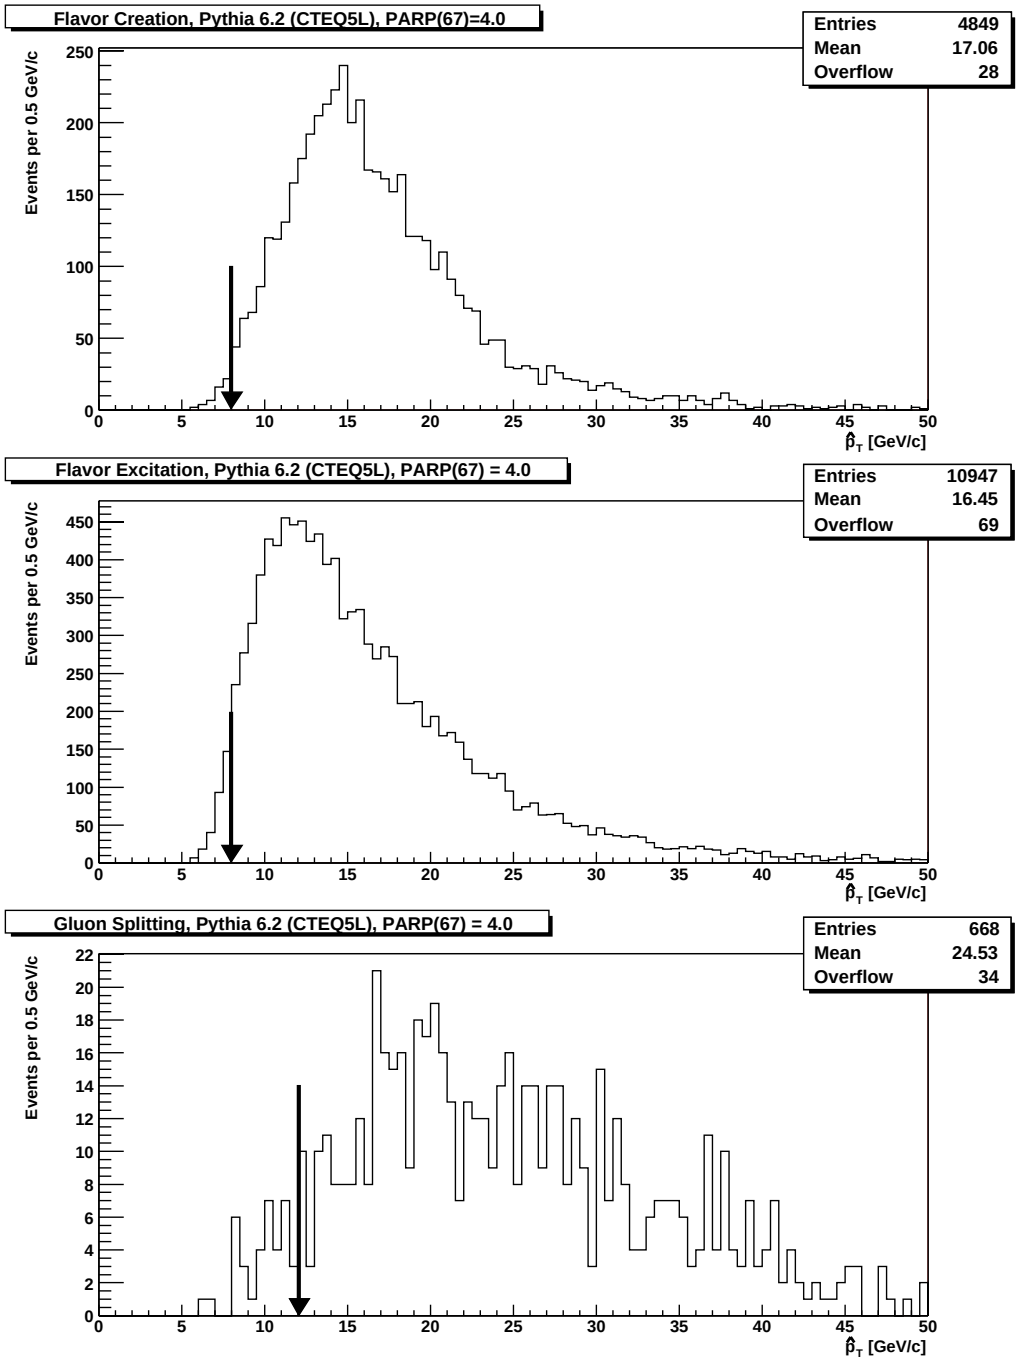

Figure 5.3 shows the minimum parton-parton center of mass p_T for the hard scatter, called p_T(min), for events from PYTHIA with `PARP(67)` = 4.0 that have one *B* hadron having $p_T > 13$ GeV/c and another *B* hadron with $p_T > 6$ GeV/c. The distributions for `PARP(67)` = 3.0 are similar. The *B* hadron p_T cuts were chosen as indicative of the p_T range for *B* hadrons reconstructed in this analysis. These distributions were generated using the lowest reasonable p_T(min) for each production mechanism. The arrows show the p_T(min) values used to generate the large samples used with the detector simulation.

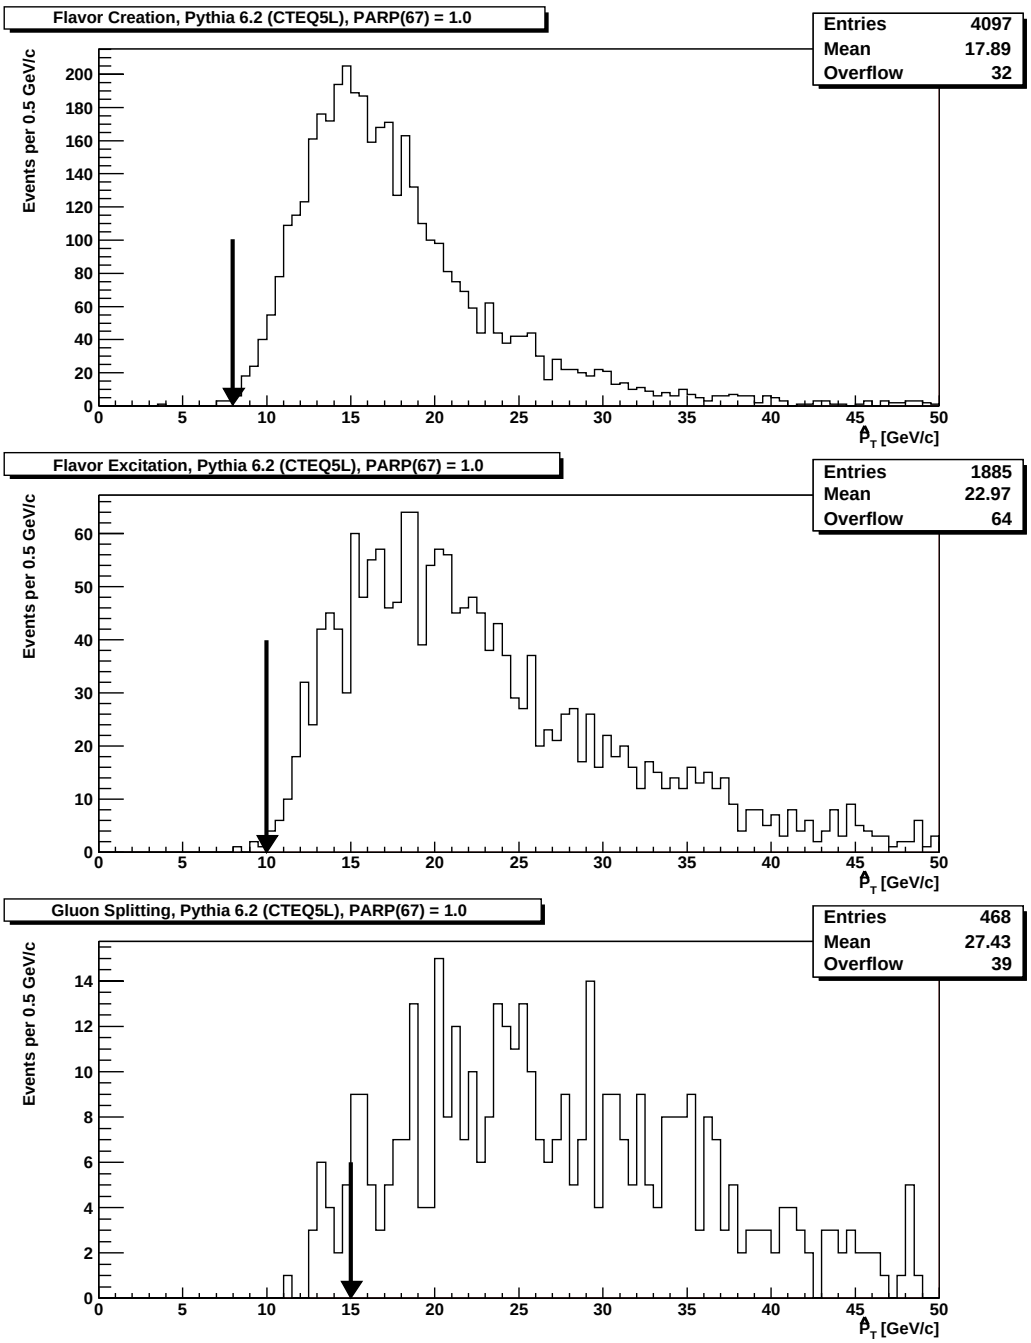

Figure 5.4 shows the minimum parton-parton center of mass p_T for the hard scatter, called p_T(min), for events from PYTHIA with `PARP(67)` = 1.0 that have one B hadron having $p_T > 13$ GeV/c and another B hadron with $p_T > 6$ GeV/c. The B hadron p_T cuts were chosen as indicative of the p_T range for B hadrons reconstructed in this analysis. These distributions were generated using the lowest reasonable p_T(min) for each production mechanism. The arrows show the p_T(min) values used to generate the large samples used with the detector simulation.

initial-state radiation as follows:

- PARP(67) = 1.0: Low amount of initial-state radiation.
- PARP(67) = 3.0: Intermediate amount of initial-state radiation.
- PARP(67) = 4.0: High amount of initial-state radiation.

The terms "low," "intermediate," and "high" simply refer to the value of the initial-state radiation parameter `PARP(67)` relative to the range from one to four and do not indicate any sort of absolute evaluation of the amount of radiation. The value of `PARP(67)` = 3 was chosen as an intermediate value rather than one closer to the center of the range because studies detailed in [66] prefer `PARP(67)` values near four to describe generic QCD data at CDF.

In addition to the $b\bar{b}$ sample generated in PYTHIA, one sample of $c\bar{c}$ events with `PARP(67)` = 4.0 was generated using PYTHIA to help evaluate backgrounds coming from prompt charm production.

5.3 HERWIG Generation

The version of HERWIG used for this analysis is 6.400. However, it was discovered that this version contains a small bug that had to be fixed specifically in the source code used for this analysis [64]. It is expected that the bug will be fixed in future HERWIG releases.

5.3.1 Production Mechanisms

Production of $b\bar{b}$ events in HERWIG is very similar to PYTHIA $b\bar{b}$ production, although there are some significant differences. Just like PYTHIA, HERWIG can generate two processes relevant to $b\bar{b}$ production, generic QCD and heavy flavor production. However, in HERWIG, the heavy flavor production includes not only flavor creation diagrams, but also flavor excitation as well. Just as in PYTHIA, the heavy flavor production option uses matrix elements that retain the *b* quark mass. This means that HERWIG treats flavor excitation with massive hard scattering matrix elements, rather than the massless ones used in PYTHIA. Whether it is preferable to use massive matrix elements for flavor excitation remains an open theoretical question. For this analysis, the different production mechanisms were generated as follows:

- **Flavor Creation and Flavor Excitation:** Both of these processes are generated at the same time in HERWIG with 100% efficiency by running HERWIG's heavy flavor production process (obtained by setting the HERWIG parameter `IPROC` = 1705). Events from the heavy flavor production process can be separated into separate flavor creation and flavor excitation samples by examining information about the partons involved in the hard scatter.

- **Gluon Splitting:** Just as in PYTHIA, gluon splitting events can only be generated by running the generic QCD process in HERWIG (`IPROC` = 1500), and separating gluon splitting events from the large number of other processes generated. Many generic QCD interactions must be generated to produce sufficient statistics in the gluon splitting sample.

5.3.2 Additional Parameters

Unlike PYTHIA, HERWIG does not expose many parameters for users to tune to obtain agreement with experimental observations. We use the same parton distribution functions (CTEQ5L) for HERWIG as were used for PYTHIA. However, HERWIG does not automatically set the value of Λ_{QCD} appropriate for the parton distribution functions chosen. The HERWIG samples generated for this analysis use Λ_{QCD} = 192 MeV. In addition, using HERWIG's default settings, its cluster fragmentation model does not produce any *B* baryons. To obtain a *B* baryon fraction compatible with observations at CDF, the cluster fragmentation parameter `CLPOW` = 1.26 must be used. Finally, it is worth noting that HERWIG does not contain a model for multiple-parton interactions. Therefore, it has no parameters to tune for this aspect of the data. This also means that HERWIG's simulation of the underlying event does not match CDF observations as well as PYTHIA's (after PYTHIA has been tuned as described above). However, the results of this analysis are not very sensitive to the details of the underlying event.

Like PYTHIA, HERWIG has a p_T(min) parameter. For this analysis, the HERWIG p_T(min) parameter is also adjusted to give the most efficient event generation possible with the least amount of distortion to generated distributions. The p_T(min) used to generate each production mechanism is given in Table 5.3, and the p_T of the final-state partons in the hard scatter center of mass frame for the different processes is plotted in Figure 5.5.

Process	HERWIG
flavor creation	8 GeV/c
flavor excitation	8 GeV/c
gluon splitting	10 GeV/c

Table 5.3 The p_T(min) used to generate the HERWIG sample for this analysis.

All other HERWIG parameters are left at their default values [34].

5.3.3 Samples Generated

Because HERWIG has fewer parameters to tune, only one sample of HERWIG $b\bar{b}$ events was generated for this analysis. No $c\bar{c}$ events were generated as the background contributions are evaluated using PYTHIA.

5.4 Heavy Flavor Decays

The decays of heavy flavor *B* and *D* hadrons have been studied extensively at e^+e^- colliders typically running at the Y(4S) resonance, which decays almost exclusively into *B* hadron pairs. Using inputs from these experiments, specialized models for heavy flavor decays have been developed and implemented in Monte Carlo programs. In general, these decay models are considered preferable to the less specifically tuned models in general purpose programs like PYTHIA and HERWIG. The heavy flavor decay Monte Carlo used for this analysis is the QQ program developed by the CLEO Collaboration [67].

There are two distinct advantages for this analysis, in addition to the better decay modeling, provided by using QQ for heavy flavor decays. The first is that QQ allows the user to specify particular decay channels for given *B* hadrons. In this analysis, this allows us to force one of the *B* hadrons to decay in a channel that involves an electron or a muon, enhancing the statistical power of the Monte Carlo samples. Without this option, approximately 60% of the $b\bar{b}$ events generated by PYTHIA or HERWIG would be discarded because they would not contain a trigger lepton from a *B* decay.

In addition, the statistical power of the Monte Carlo samples in this analysis is further increased by re-decaying $b\bar{b}$ events generated with PYTHIA or HERWIG multiple times. Each re-decayed event is treated as a separate event in the rest of the Monte Carlo chain and so each is

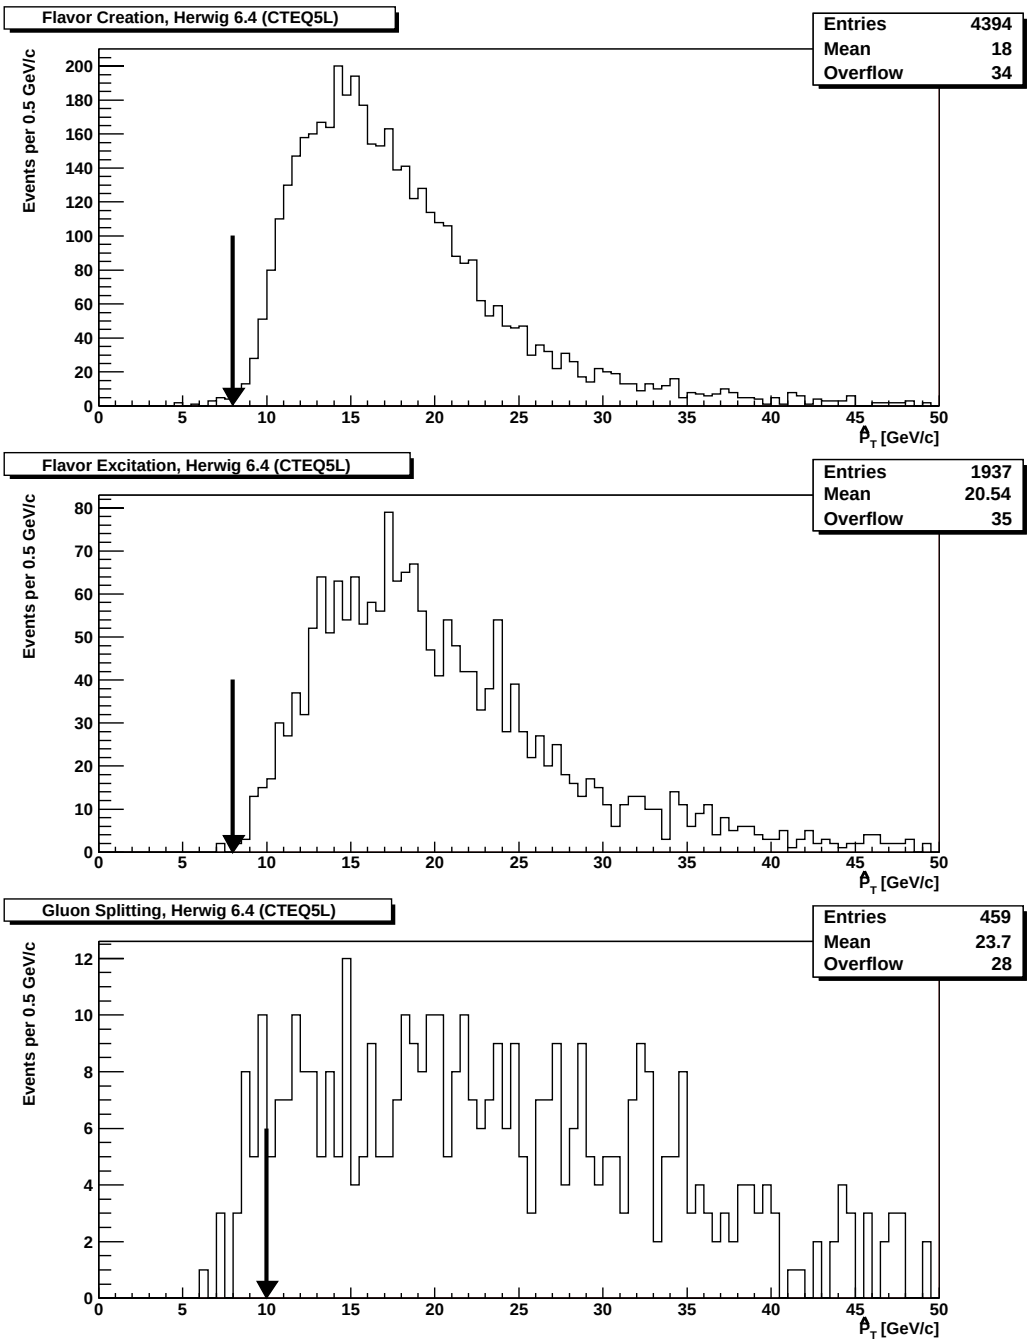

Figure 5.5 shows the minimum parton-parton center of mass p_T for the hard scatter, called p_T(min), for events from HERWIG that have one B hadron having $p_T > 13$ GeV/c and another B hadron with $p_T > 6$ GeV/c. The B hadron p_T cuts were chosen as indicative of the p_T range for B hadrons reconstructed in this analysis. These distributions were generated using the lowest reasonable p_T(min) for each production mechanism. The arrows show the p_T(min) values used to generate the large samples used with the detector simulation.

given a separate chance to satisfy detector, trigger, and reconstruction criteria to make it to the final sample. For this analysis, flavor creation and flavor excitation events from both PYTHIA and HERWIG are re-decayed ten times forcing all *B* hadrons containing a *b* quark to decay semileptonically while allowing the other *B* hadron to decay without constraint. The same events are then re-decayed another ten times, in this case forcing the semileptonic decay of the *B* hadron containing a $\bar{b}$ quark, for a total of twenty re-decays per generated event. Gluon splitting events from either generator are re-decayed twenty times each way, for a total of forty re-decays. After detector and trigger simulation, event reconstruction, and analysis cuts, no single event contributes more than 11 times for flavor creation or flavor excitation or 26 times for gluon splitting due to re-decay.

5.5 Detector, Reconstruction, and Analysis

The main goal of the Monte Carlo generation for this analysis is to produce a sample of simulated data that is as close to the actual data collected at CDF as possible. For this reason, after simulating the physics processes with PYTHIA and HERWIG, and simulating the decay processes with QQ, the Monte Carlo events are passed through a detector and trigger simulation to mimic the effects of the CDF detector on the data. After the detector and trigger simulation, the Monte Carlo events are treated identically to the real data. They are passed through the standard CDF event reconstruction code and subjected to the same analysis criteria used for the real data. As much as possible, including final quality cuts and mistag subtraction, the simulated events are treated the same as the real data.

5.5.1 Detector Simulation

The generated Monte Carlo events are passed though a detector simulation that models the data produced by the CDF detector in response to particles traversing the active regions of the detector subsystems. The main purpose of the detector simulation is to reproduce the acceptance and efficiency of the CDF detector as accurately as possible so that Monte Carlo can be compared directly with the actual data. This analysis uses the standard CDF fast detector simulation, known as QFL'. The “prime” in QFL' refers to the fact that this version of the CDF detector simulation includes the Run 1b silicon detector, the SVX', rather than the Run 1a SVX.

To begin, QFL' adjusts the location of the primary vertex for the Monte Carlo event to account for the misalignment between the proton-antiproton collision point and the center of the CDF detector. By default, PYTHIA and HERWIG generate collisions that occur exactly at the center of the CDF detector. QFL' translates the primary collision point for each event to reflect the beam offset, tilt, and spread, as described in the following formulae:

$$\begin{aligned} x &= x_0 + \frac{dx}{dz} z + \Delta x \\ y &= y_0 + \frac{dy}{dz} z + \Delta y \,, \\ z &= z_0 + \Delta z \end{aligned} \tag{5.1}$$

where x_0, y_0, and z_0 represent the offset of the primary interaction point with respect to the CDF detector, dx/dz and dy/dz represent the beam slope in the transverse directions, and Δx, Δy, and Δz represent the spread of the interaction region. For each event, QFL' chooses a random value for Δx, Δy, and Δz using a Gaussian probability distribution. However, the SVX' only covers a region extending approximately 35 cm in z from the center of the CDF detector. Because it is not possible to tag displaced vertices in events generated with primary vertices outside the coverage of the SVX', values beyond $|z| = 35$ cm are not allowed. The values used by QFL for the beam offset, slope and spread distributions, given in Table 5.4, are chosen to be indicative of the beam conditions at CDF during Run 1b.

Next, QFL' calculates the trajectory of each particle through the CDF detector, taking into account the translated interaction point, to determine which detector subsystems the particle traverses. Beginning with the innermost detector system, the SVX', QFL' determines the effect of the particle on the detector (in other words, whether the particle generates any hits in tracking chambers or deposits any energy in the calorimeter systems) as well as the effect of the detector material on the particles (for example, effects from multiple scattering and bremsstrahlung). However, rather than using a detailed simulation of the interactions and energy loss of the particle as it passes through the detector material, QFL' relies on parameterizations of these effects based on measurements from data and the results of more detailed simulations. For example, rather than calculating the amount of ionization generated by a particle traversing the argon-ethane of the CTC and determining how the ionized charge drifts to the CTC wires to generate electrical signals, QFL' uses a parameterization of the number of hits a particle traversing the CTC is likely to generate and whether those hits will be sufficient to reconstruct a

Parameter	Meaning	Value
x_0	The horizontal beam offset at z = 0 cm	-0.01 cm
y_0	The vertical beam offset at z = 0 cm	-0.1 cm
z_0	The offset of the center of the interaction region from the center of the CDF detector	0.0 cm
dx/dz	The slope of the beam in the horizontal plane.	5×10^{-4}
dy/dz	The slope of the beam in the vertical plane.	-5×10^{-4}
$\sigma_{\Delta x}$	The spread of the beam spot in the horizontal direction. This parameter determines the width of the Gaussian distribution QFL' uses to generate random Δx values.	25 μm
$\sigma_{\Delta y}$	The spread of the beam spot in the horizontal direction. This parameter determines the width of the Gaussian distribution QFL' uses to generate random Δx values.	25 μm
$\sigma_{\Delta z}$	The spread of the beam spot in the horizontal direction. This parameter determines the width of the Gaussian distribution QFL' uses to generate random Δx values.	33.0 cm

Table 5.4 The QFL' parameters that determine the location of the primary vertex generated for each event.

track. QFL' includes parameterized simulations for the SVX', CTC, electromagnetic and hadronic calorimeters, and the muon chambers.

The output of QFL' is a set of data roughly equivalent to what would be produced by reading out the CDF detector for an actual event. This simulated detector data is stored in the same format as the actual data so that the same reconstruction code used on the actual data may also be used to reconstruct events in the simulated data.

5.5.2 Trigger Simulation

Characteristics of the electron and muon triggers used to define the data sample for this analysis can have a significant affect of the kinematic distributions of the *B* hadrons in the final sample. To account for these effects in the Monte Carlo, a simulation of the trigger must be applied. The most significant effects come from the change in efficiency of the trigger over a certain range in p_T or E_T near the threshold for the trigger.

For the electron sample, the dominant effects come from the efficiency thresholds of the 8 GeV electromagnetic calorimeter cluster E_T and the 7.5 GeV/c CFT p_T requirements in the

Level 2 trigger. In addition, a slight E_T dependence is introduced by the CES matching requirement in the Level 2 trigger. These effects were studied in [68] using data collected by triggers with lower p_T and E_T thresholds, and without the CES matching requirement. Parameterizations of these trigger efficiencies were taken from [68] and applied to the Monte Carlo data for this analysis to account for effects of the electron trigger on the data.

The main influence on kinematic distributions in the muon sample comes from the effects of the 7.5 GeV/c threshold in the CFT at Level 2. There is also a p_T dependence in the efficiency for the Level 1 muon trigger. These effects have been studied in [69] and [70], and included in standard CDF software package known as DIMUTG [71] that can be used to simulate the effect of a variety of single muon and dimuon triggers. DIMUTG also includes the effects of the changing performance of the triggers as a function of time. For Monte Carlo data, this effect is included by randomly selecting a parameterization for the trigger efficiencies weighted by the total amount of luminosity collected under those trigger conditions. DIMUTG is used for this analysis to model the effects of the muon trigger requirements on the kinematic distributions of the *B* hadrons in this analysis.

5.5.3 Reconstruction and Analysis

Events that successfully pass the trigger simulation are then processed through the same event reconstruction code used for the data. This code performs basic event reconstruction, like taking hits—in this case, simulated hits—from the silicon detector and reconstructing track segments to attach to tracks reconstructed in the CTC. After reconstruction, the Monte Carlo events are processed through the same analysis code used for the data, including the modified BVTX vertex-tagging algorithm. Doubly tagged events are selected using the same cuts and criteria used for the actual data. In addition, mistags are removed from the Monte Carlo samples using the same mistag subtraction techniques used for the data. The final goal is to provide Monte Carlo samples of doubly tagged events that resemble as closely as possible the actual character of the data. This allows a direct comparison of the tag $\Delta\phi$ spectrum from the actual data to the Monte Carlo tag $\Delta\phi$ spectrum.

Chapter 6

Comparisons between Monte Carlo and Data

The previous chapter describes the efforts taken to make Monte Carlo $b\bar{b}$ samples using PYTHIA and HERWIG that are as comparable to the $b\bar{b}$ component of the data as possible. These efforts allow a direct comparison between the Monte Carlo and data, both for the purpose of establishing the validity of the PYTHIA and HERWIG models of *b* production, and also for allowing the measured tag pair correlations to be corrected back to the *B* hadron level. In generating the Monte Carlo samples, no attempt was made to simulate all possible sources of background present in the data. Only the $b\bar{b}$ component of the data is simulated. Nonetheless, after using mistag subtraction and cuts to reduce other backgrounds to negligible levels for both Monte Carlo and data, a direct comparison, neglecting residual backgrounds, is reasonable. Based on the success of this comparison, the Monte Carlo samples are used to derive corrections for the data to extract the *B* hadron level correlations.

6.1 Single *B* Distributions

Before evaluating the similarity between correlations in the Monte Carlo and the data, it is wise to consider the similarity in distributions involving the individual *B* hadrons in the events. Although distributions involving single *B* hadrons, rather than correlations between both *B* hadrons in the event, are not very sensitive to the mixture of flavor creation, flavor excitation, and gluon splitting production mechanisms, comparing the single *B* distributions between Monte Carlo and data reveals whether the detector and trigger simulation are adequately modeling the effect of the detector and trigger on the actual data. Since the single *B* distributions from each production mechanism are virtually indistinguishable from one another, after mistag subtraction, they are separately compared to the overall mistag-subtracted single *B* distributions from data.

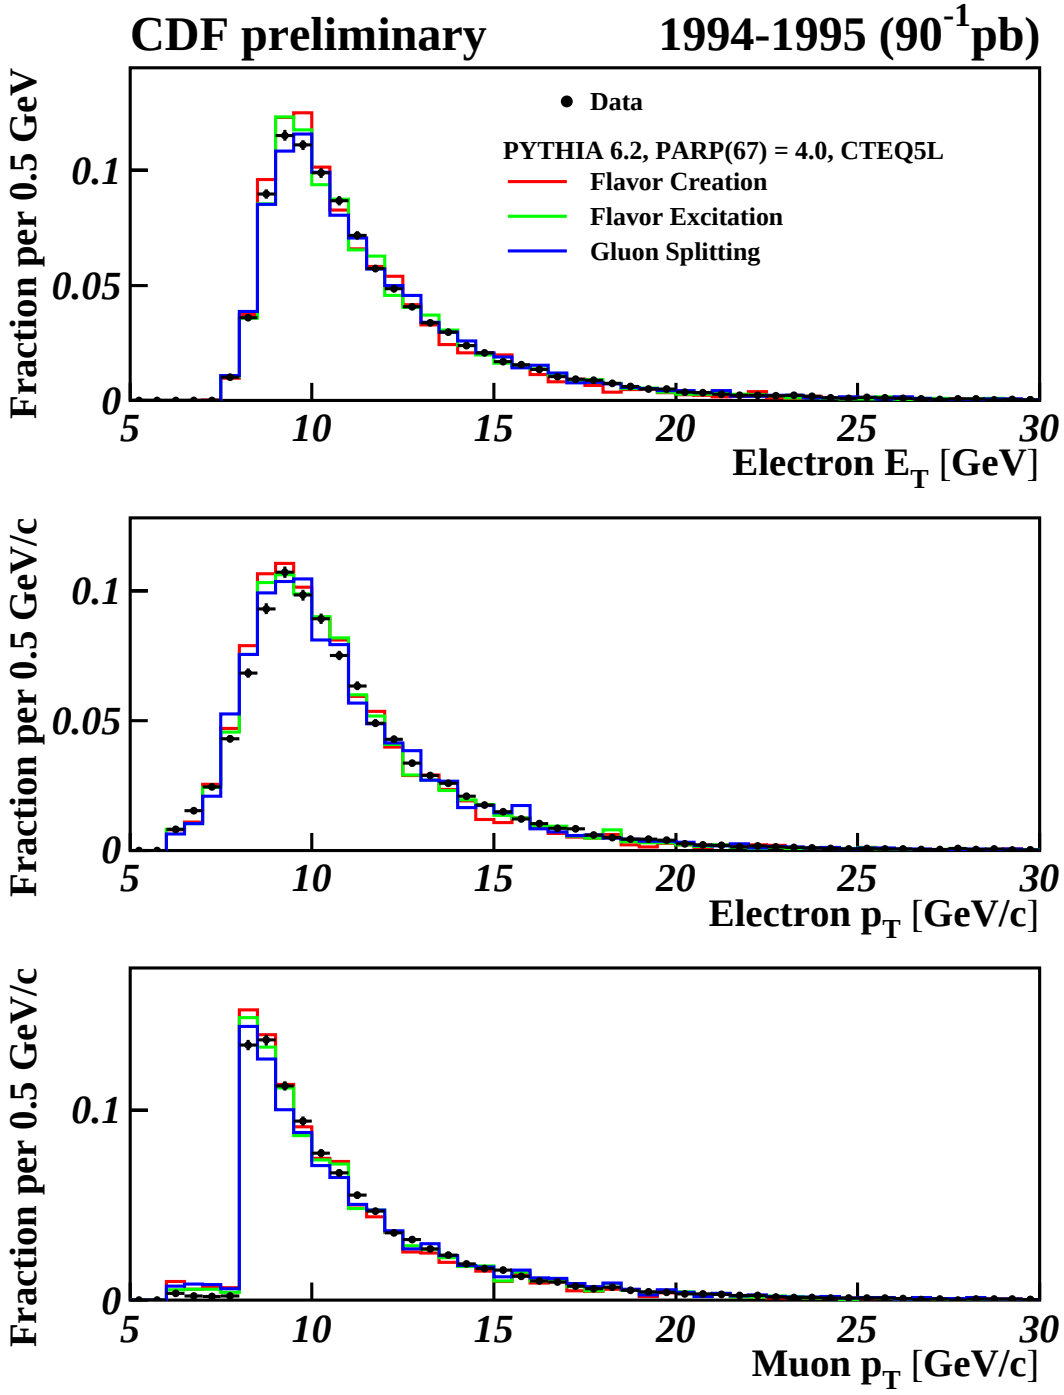

Figure 6.1 shows a comparison between PYTHIA with `PARP(67)` = 4.0 and data for the p_T and E_T spectra of electrons and the p_T spectrum of muons. The distributions are normalized to unit area.

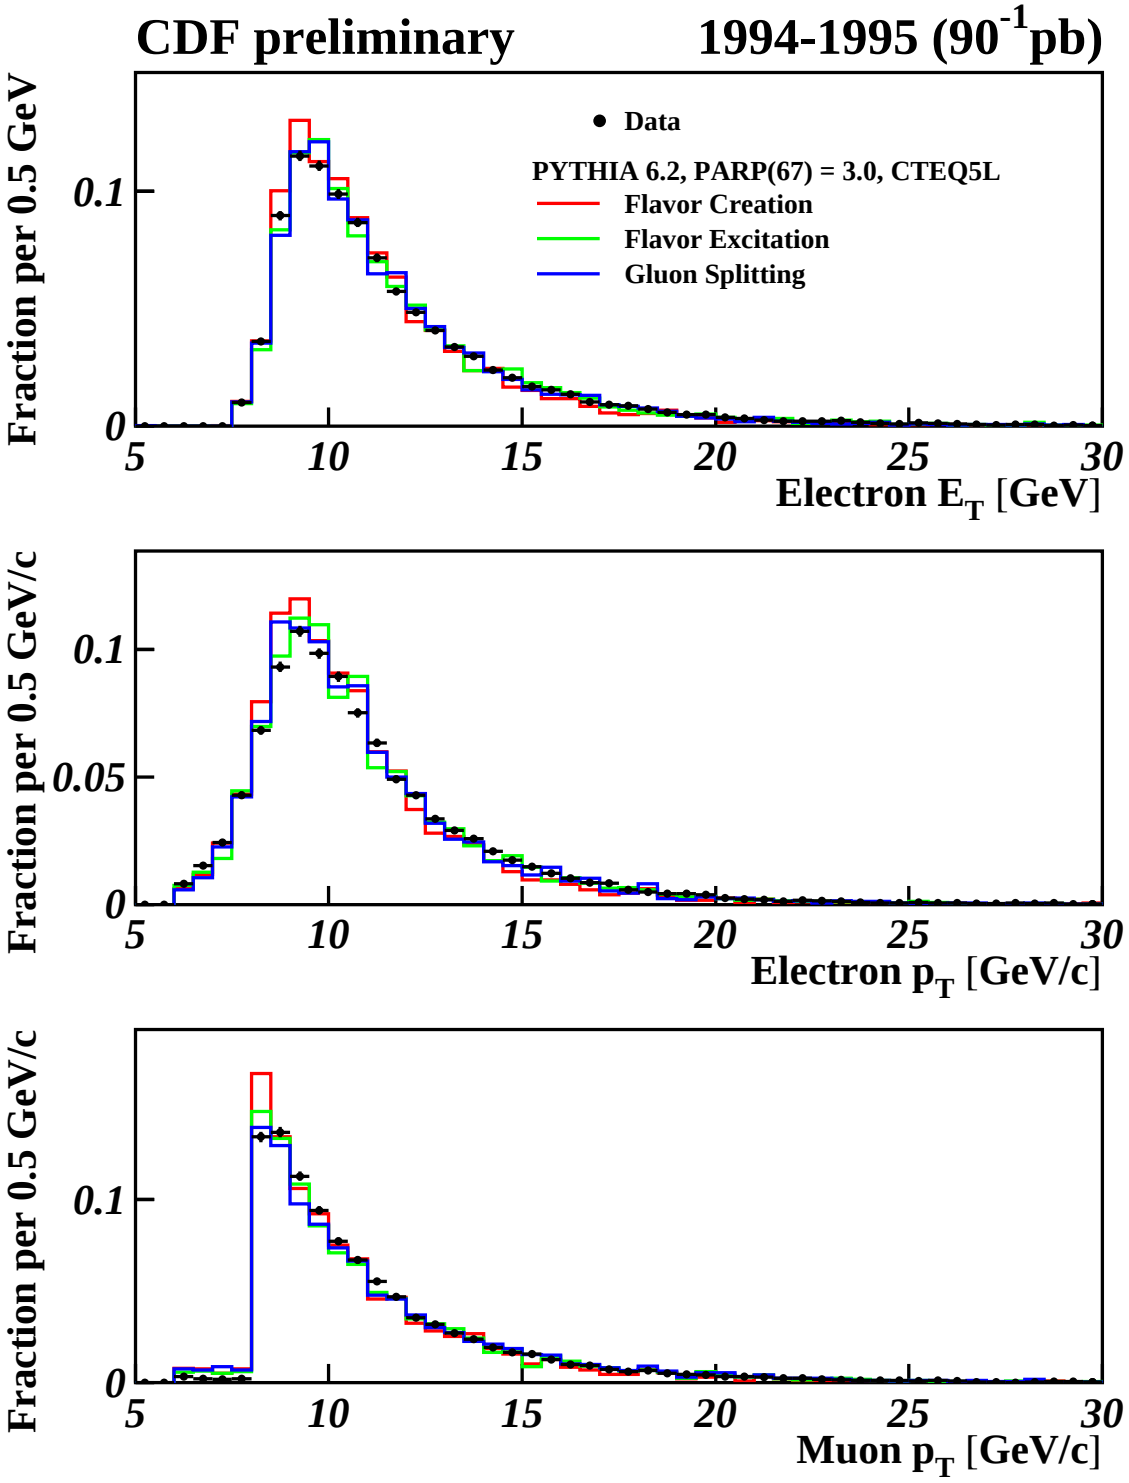

Figure 6.2 shows a comparison between PYTHIA with `PARP(67)` = 3.0 and data for the p_T and E_T spectra of electrons and the p_T spectrum of muons. The distributions are normalized to unit area.

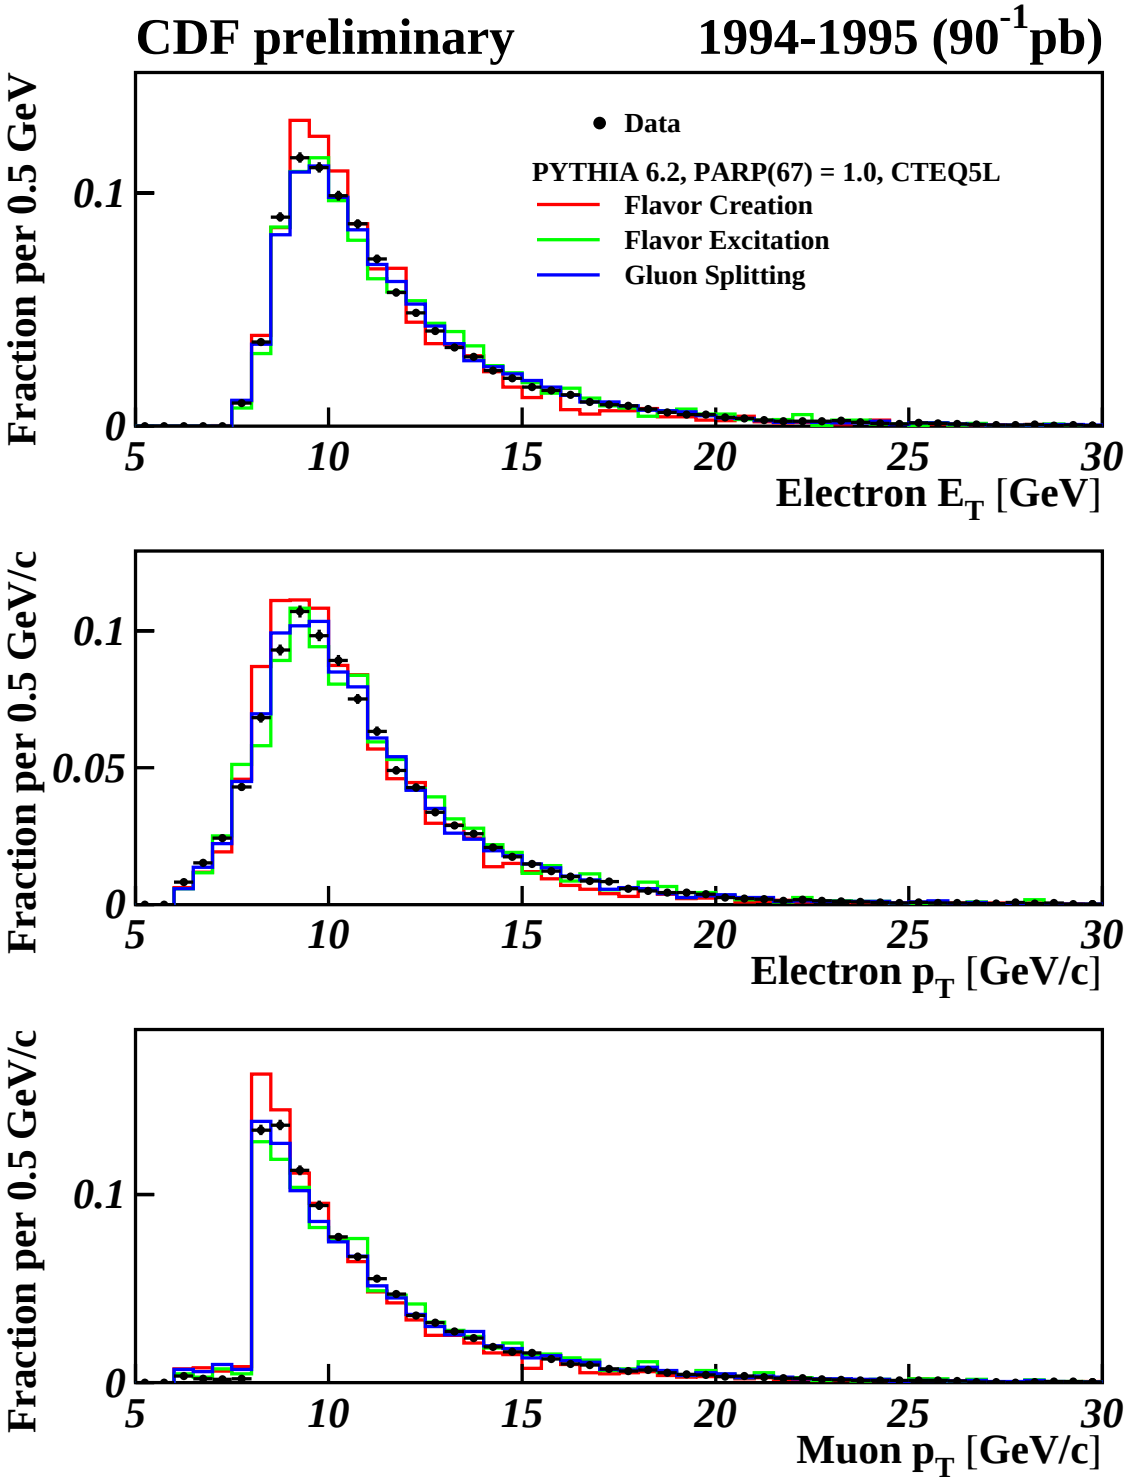

Figure 6.3 shows a comparison between PYTHIA with `PARP(67)` = 1.0 and data for the p_T and E_T spectra of electrons and the p_T spectrum of muons. The distributions are normalized to unit area.

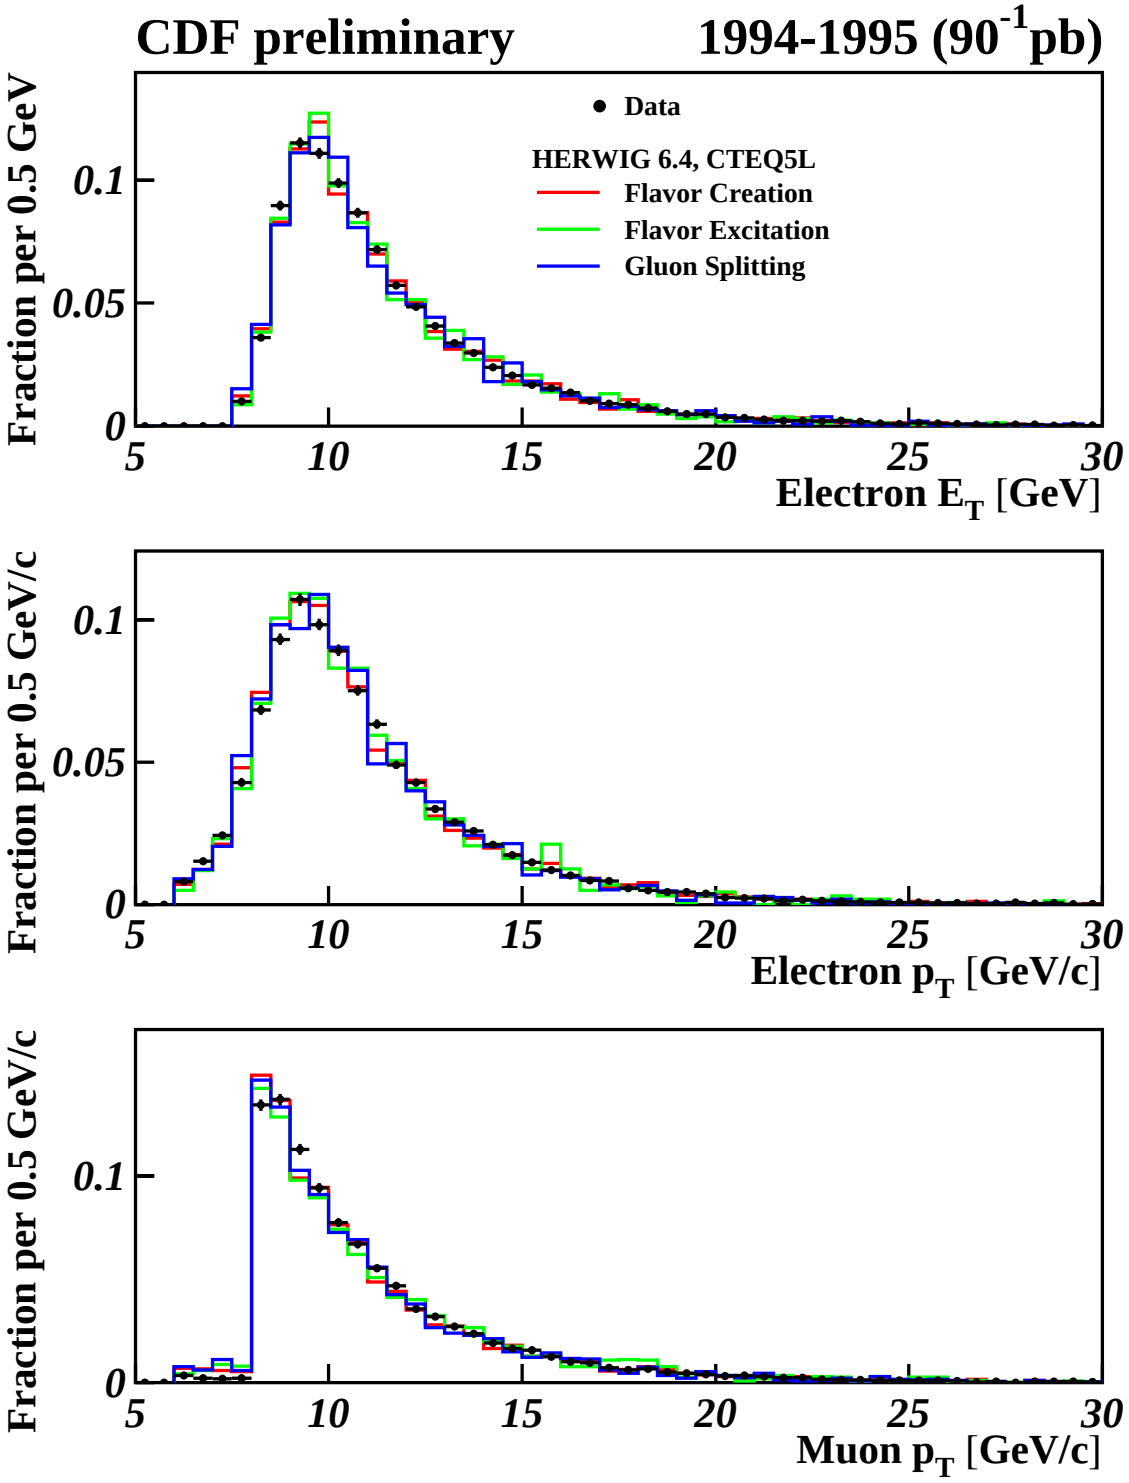

Figure 6.4 shows a comparison between HERWIG and data for the p_T and E_T spectra of electrons and the p_T spectrum of muons. The distributions are normalized to unit area.

6.1.1 Leptons

Figure 6.1 through Figure 6.4 show the comparison of the Monte Carlo distribution for lepton p_T and E_T compared to the distributions taken from data. Note that since muons don't deposit significant energy in the calorimeter, there is no muon E_T distribution plotted. The degree to which these distributions match indicates how well the trigger simulation models the actual trigger. It is important to model correctly the effects of trigger efficiency for different lepton energies and momenta because these efficiencies determine the effective minimum p_T B hadron to which this analysis is sensitive. The figures show that for each Monte Carlo the lepton distributions are in relatively good agreement with the data.

6.1.2 Tags

Distributions involving individual BVTX tags carry a number of different pieces of information. First, because these distributions are influenced by p_T dependences in the trigger, these distributions further confirm that electron and muon trigger effects are adequately modeled. In addition, agreement in these distributions between Monte Carlo and data indicates that the effects of the tagging algorithm itself on the data have been correctly taken into account. Finally, these distributions check the validity of the B production and decay models used to generate the Monte Carlo data. Also, because the mistag-subtracted Monte Carlo distributions contain contributions only from $b\bar{b}$ production, any discrepancy between Monte Carlo and data could signal a possible unaccounted-for background in the data that isn't remove by mistag subtraction or cuts.

Figure 6.5 through Figure 6.8 show comparisons between Monte Carlo and data distributions for individual BVTX tags. The agreement in the tag p_T, mass, η, and L_{xy} distributions indicate that both PYTHIA and HERWIG are reasonable models for the kinematics of B production and that detector, trigger, and reconstruction effects are correctly modeled in the Monte Carlo.

6.2 Momentum Acceptance

Because the exact shape of the measured B hadron $\Delta\phi$ distribution depends on the minimum momenta of the B hadrons included in the distribution, it is useful to understand the momentum acceptance of this analysis. Determining the minimum B hadron p_T to which this analysis is sensitive is necessary for future comparisons between measurements here and other theoretical models. It also helps in understanding how these results compare to results from other analyses.

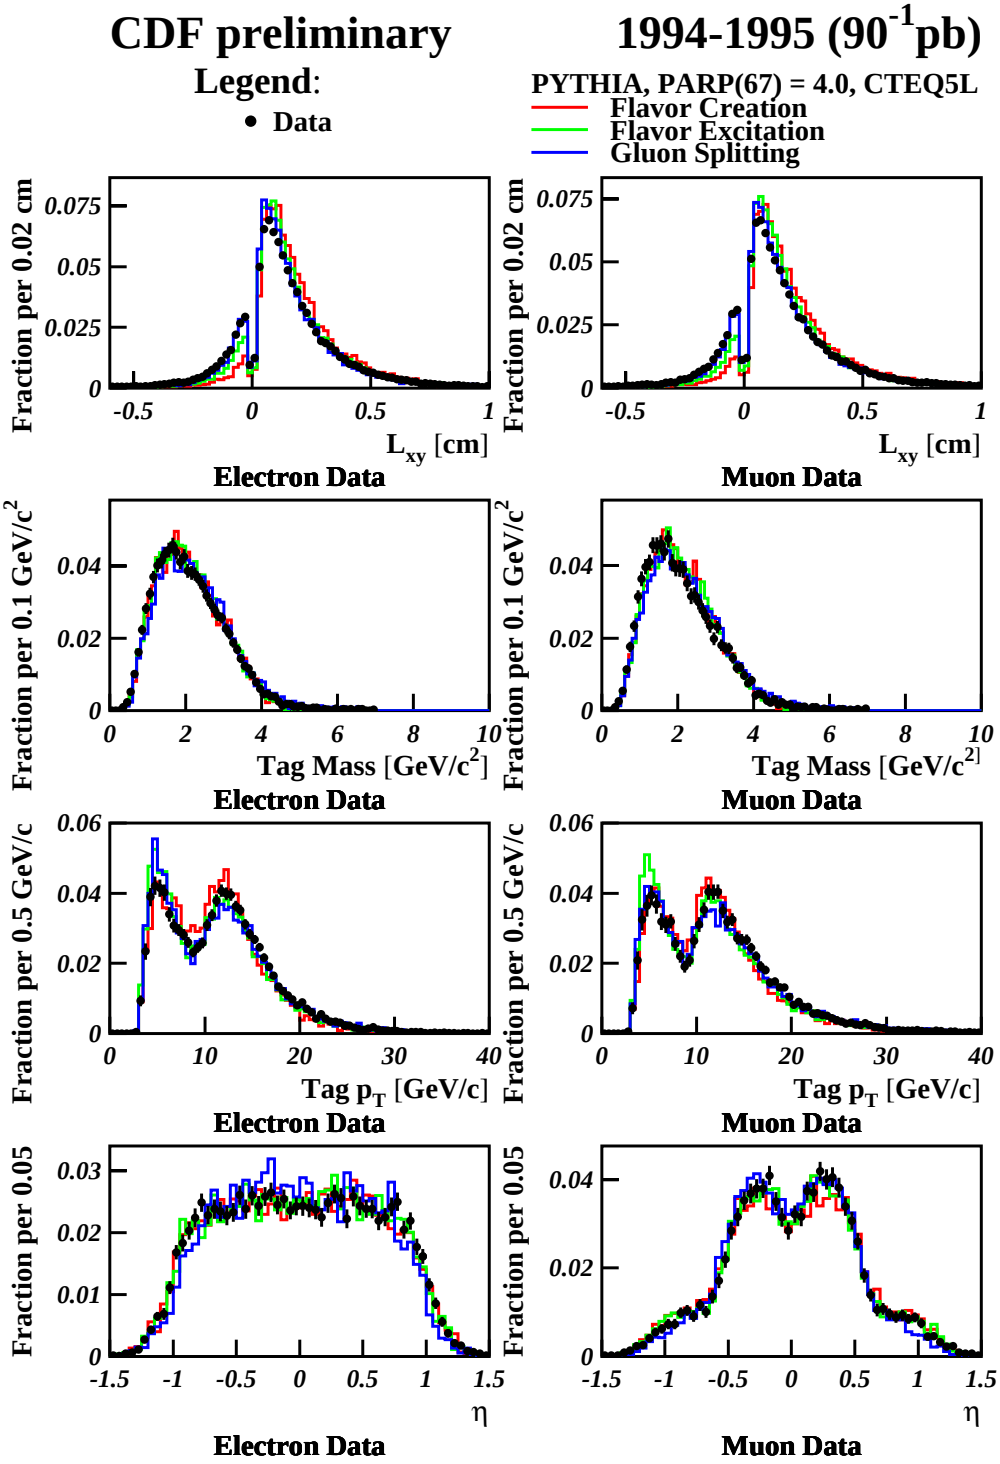

Figure 6.5 shows a comparison of several tag quantities between the data and PYTHIA with `PARP(67)` = 4.0. In each plot, all histograms have been normalized to unit area. The data is represented by the points with error bars, while the different Monte Carlo production mechanisms are shown as solid lines of different color.

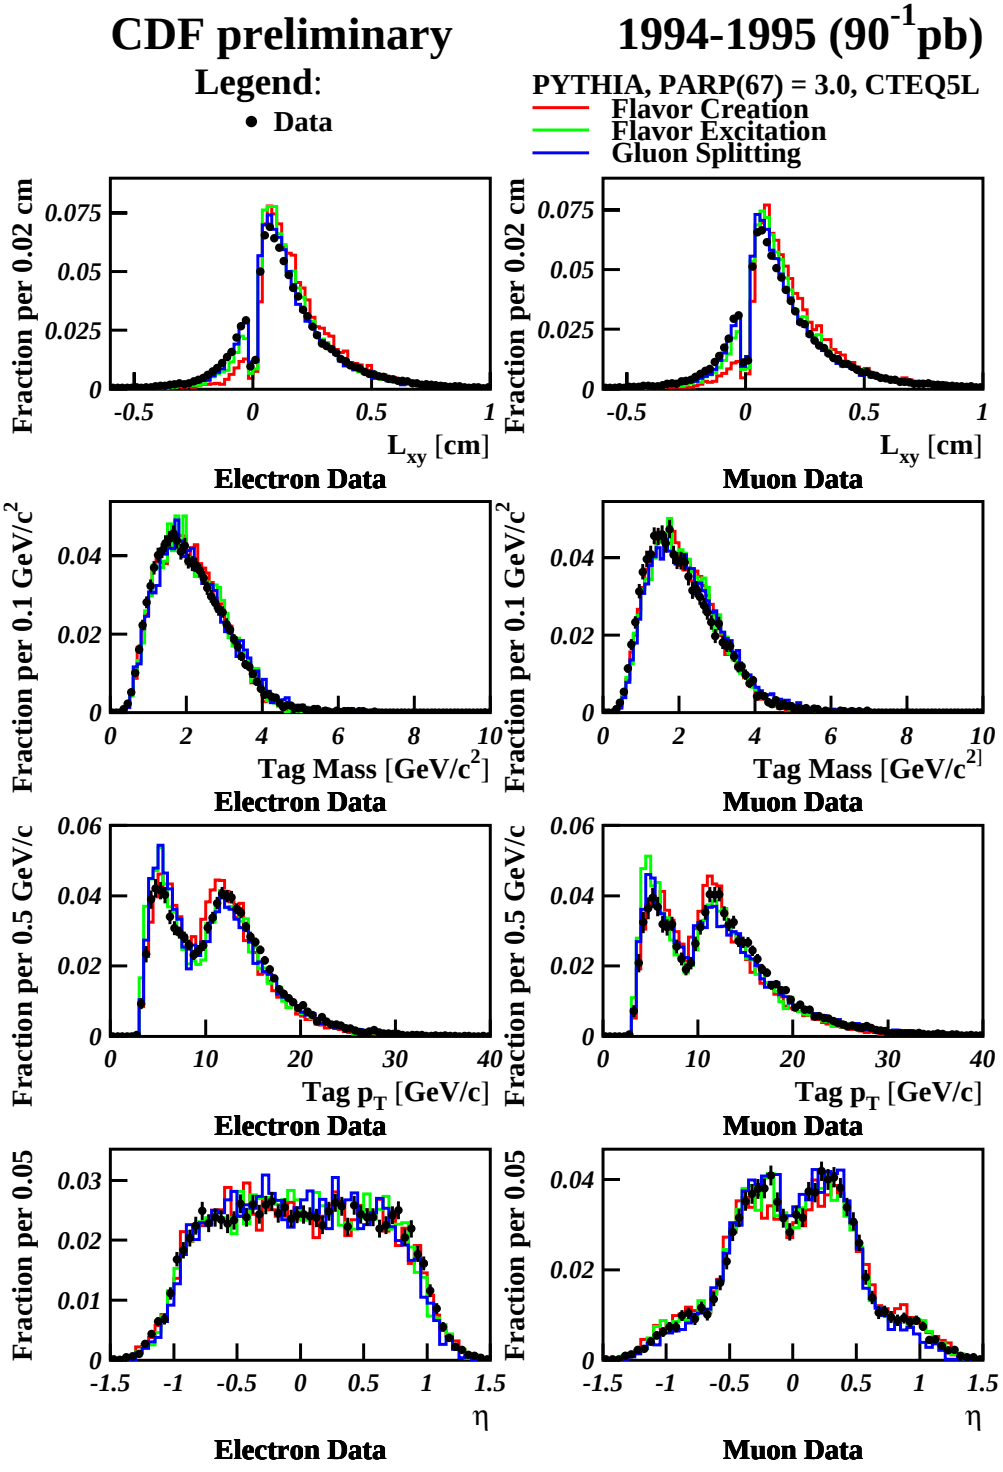

Figure 6.6 shows a comparison of several tag quantities between the data and PYTHIA with `PARP(67)` = 3.0. In each plot, all histograms have been normalized to unit area. The data is represented by the points with error bars, while the different Monte Carlo production mechanisms are shown as solid lines of different color.

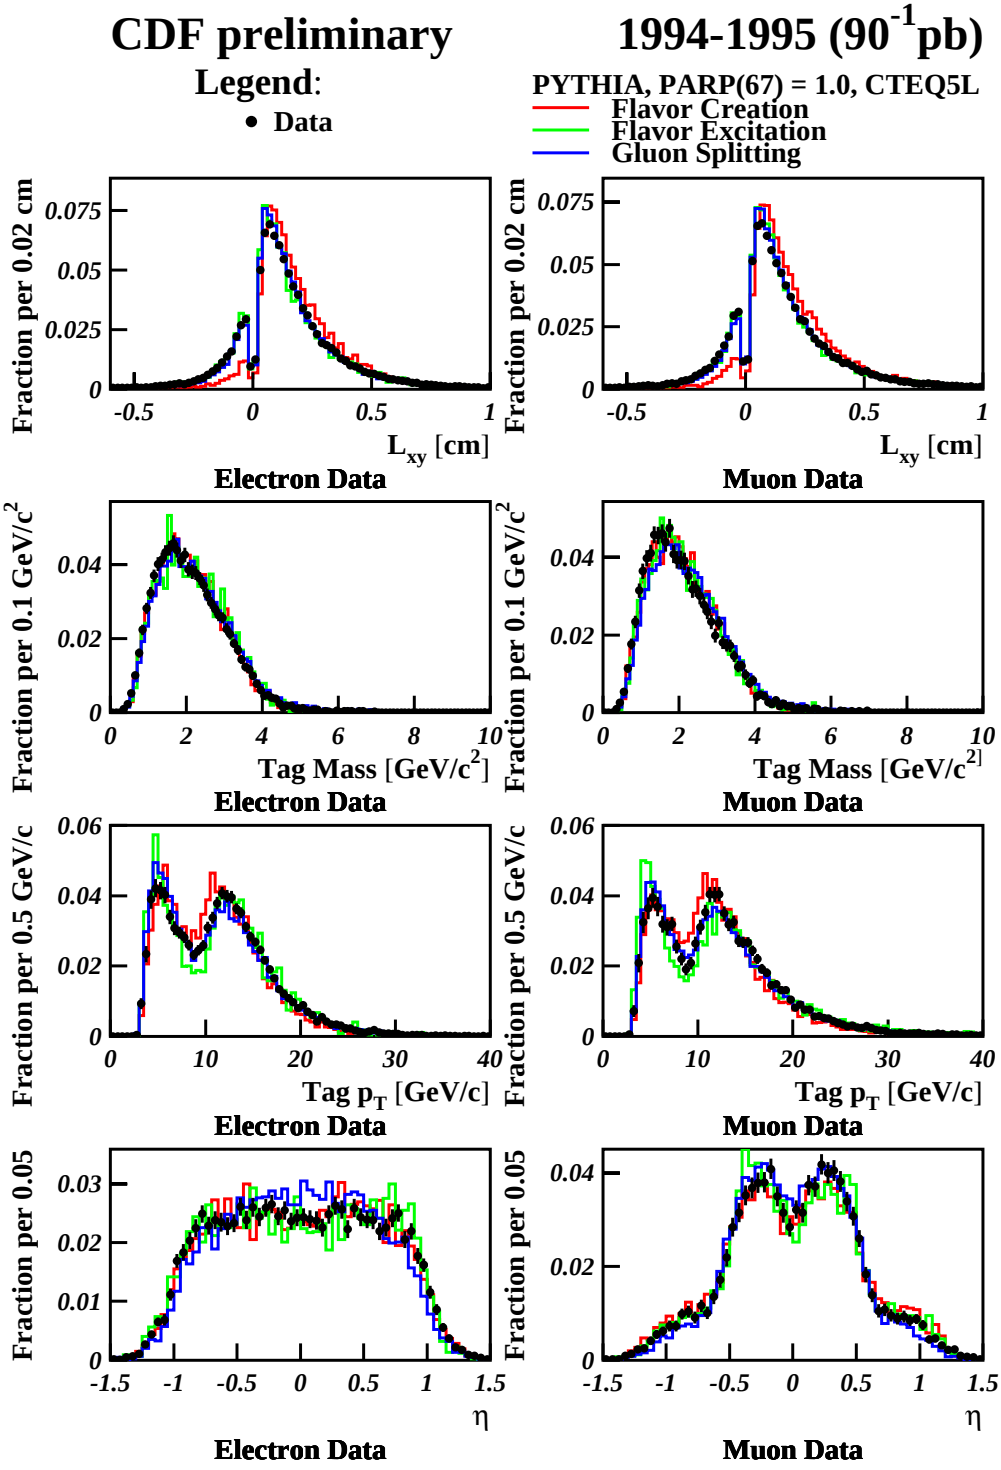

Figure 6.7 shows a comparison of several tag quantities between the data and PYTHIA with `PARP(67)` = 1.0. In each plot, all histograms have been normalized to unit area. The data is represented by the points with error bars, while the different Monte Carlo production mechanisms are shown as solid lines of different color.

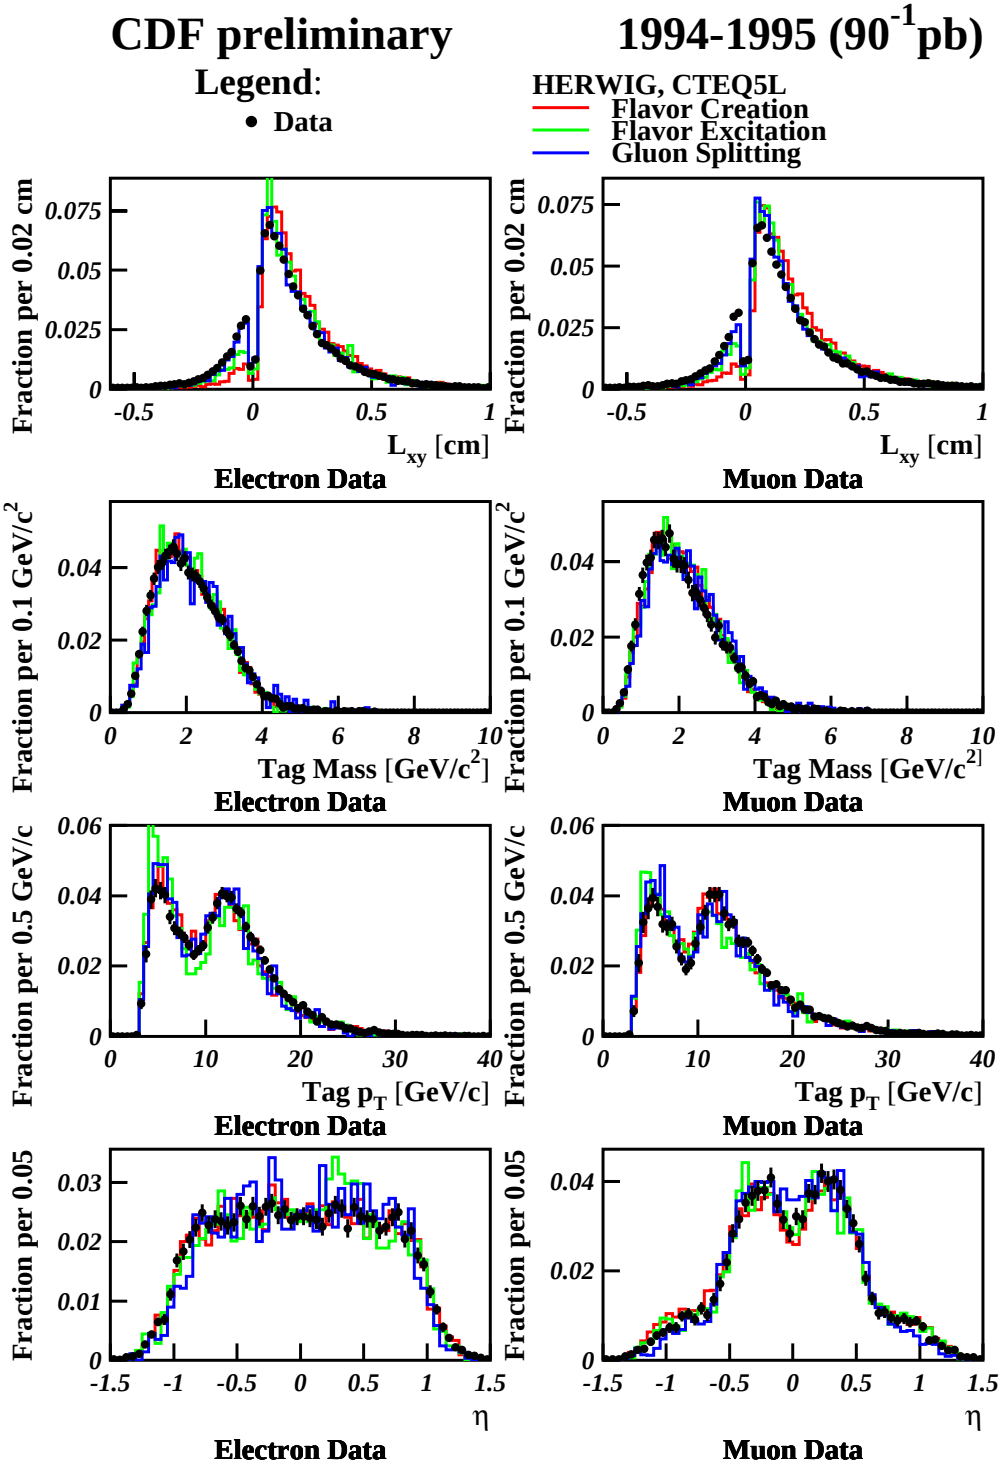

Figure 6.8 shows a comparison of several tag quantities between the data and HERWIG. In each plot, all histograms have been normalized to unit area. The data is represented by the points with error bars, while the different Monte Carlo production mechanisms are shown as solid lines of different color.

The momentum thresholds for this analysis are determined using the PYTHIA and HERWIG Monte Carlo samples. This is a reasonable approach given that the comparisons above show that these Monte Carlo samples reproduce the kinematics of the data reasonably well. The basic technique used is to match BVTX tags reconstructed in Monte Carlo events to the *B* hadrons generated in those events. From plots of the generator-level *B* hadron momentum distribution for these tagged *B*'s, the momentum threshold for the analysis is estimated as momentum above which 90% of the tagged *B* hadrons lie. The thresholds have to be determined separately for *B* hadrons that decay semileptonically to produce a trigger lepton because the trigger requirements set a higher minimum p_T than the BVTX tagging requirements alone. This approach can also be applied to *b* quarks by studying the p_T distributions for *b* quarks corresponding to tagged *B* hadrons. However, this threshold will depend not only on the Monte Carlo's *b* quark production model, but also on its hadronization model. Thus the *b* quark p_T thresholds given here for this analysis are only valid given PYTHIA's or HERWIG's fragmentation model.

Figure 6.9 through Figure 6.16 show the p_T distributions for tagged *B* hadrons and *b* quarks from the various Monte Carlo samples. The 90% thresholds for each of these distributions are summarized in Table 6.1, Table 6.2, Table 6.3, and Table 6.4. Differences between the different production mechanisms in the thresholds for the non-semileptonically decaying *B* hadron are related to correlations between the momenta of the two *B* hadrons. For example, for flavor creation, in which the *B* hadrons tend to have rather symmetric p_T values, the higher p_T requirement imposed on the semileptonic *B* decay by the electron or muon trigger translates into a higher threshold for the non-semileptonic *B* decay as well. On the other hand, for flavor excitation, which tends to produce more asymmetric *B* p_T's, the higher semileptonic *B* p_T threshold does not have as much of an effect on the other *B* p_T threshold. Since in the end, the minimum p_T thresholds for the *B* hadrons in this analysis should be a function of the leptonic trigger requirements and the BVTX tagging requirements, rather than the details of $b\bar{b}$ production correlations, the lowest p_T threshold is chosen as indicative of the actual threshold for the BVTX tagging algorithm. Also, since the thresholds for each Monte Carlo sample are reasonably similar, the common set of p_T thresholds, given below, is quoted for this analysis:

- The p_T threshold of the *B* that decays semileptonically to produce the trigger lepton (known henceforth as the *trigger B*) is given by $p_T > 14$ GeV/c.

- The p_T threshold of the other *B* is given by $p_T > 7.5$ GeV/c.

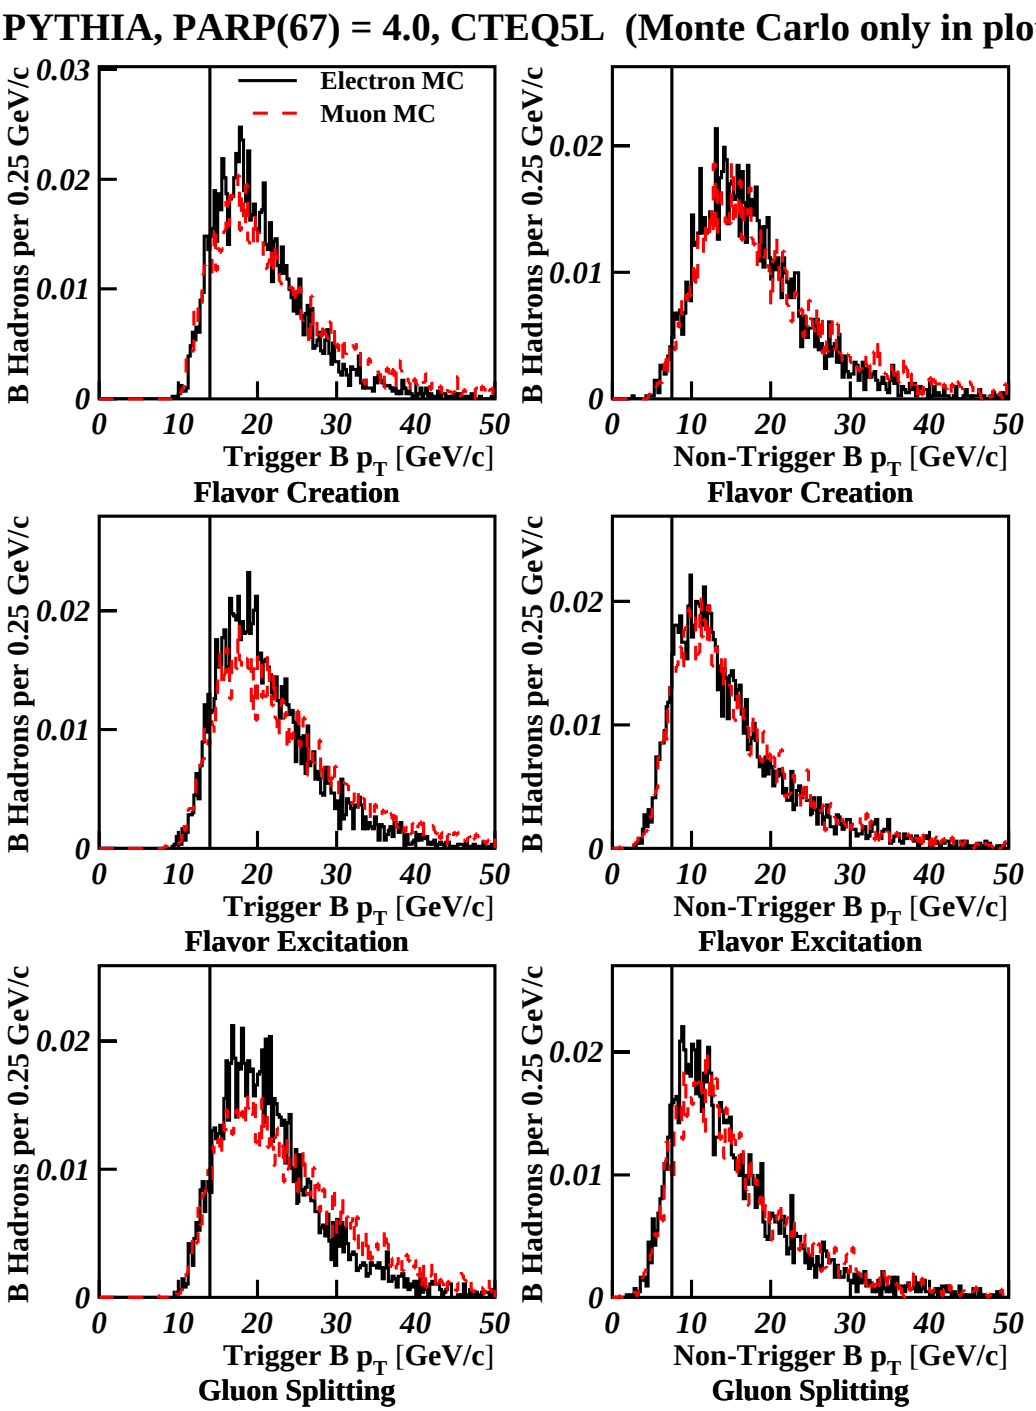

Figure 6.9 shows the *B* hadron p_T distributions for events from PYTHIA with `PARP(67)` = 4.0 where both *B* hadrons were tagged. The "Trigger *B*" is defined as the *B* whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *B*" is the other *B* in the event. The solid vertical line represents the *B* hadron threshold quoted for this analysis.

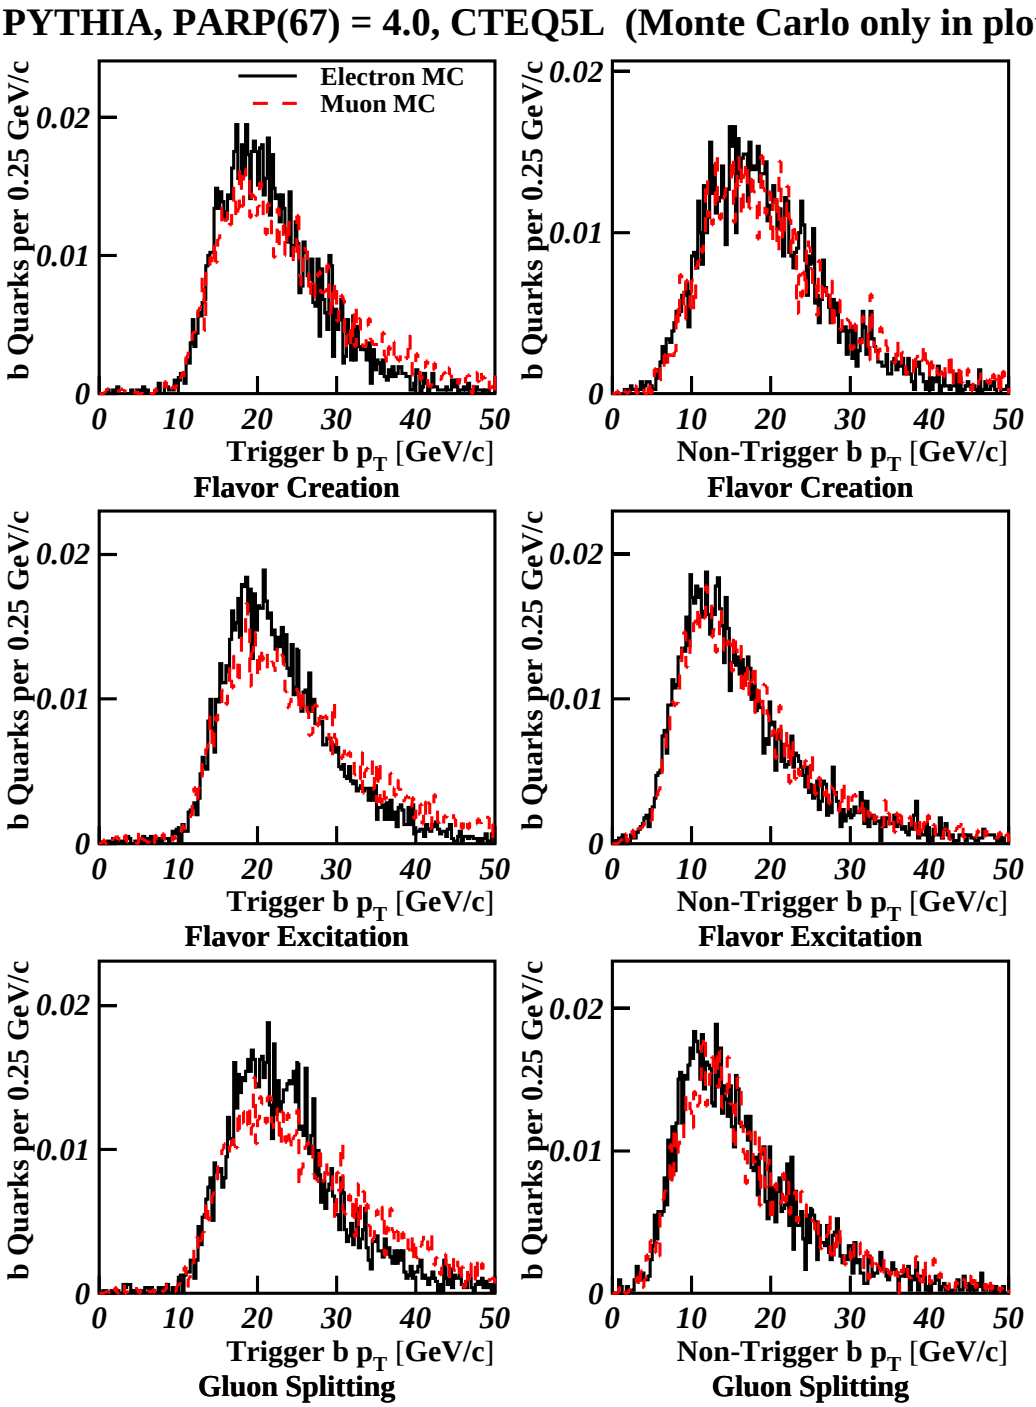

Figure 6.10 shows the *b* quark p_T distributions for events from PYTHIA with PARP(67) = 4.0 where both *B* hadrons were tagged. The "Trigger *b*" is defined as the *b* quark that hadronized to form a *B* hadron whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *b*" is the other *b* in the event.

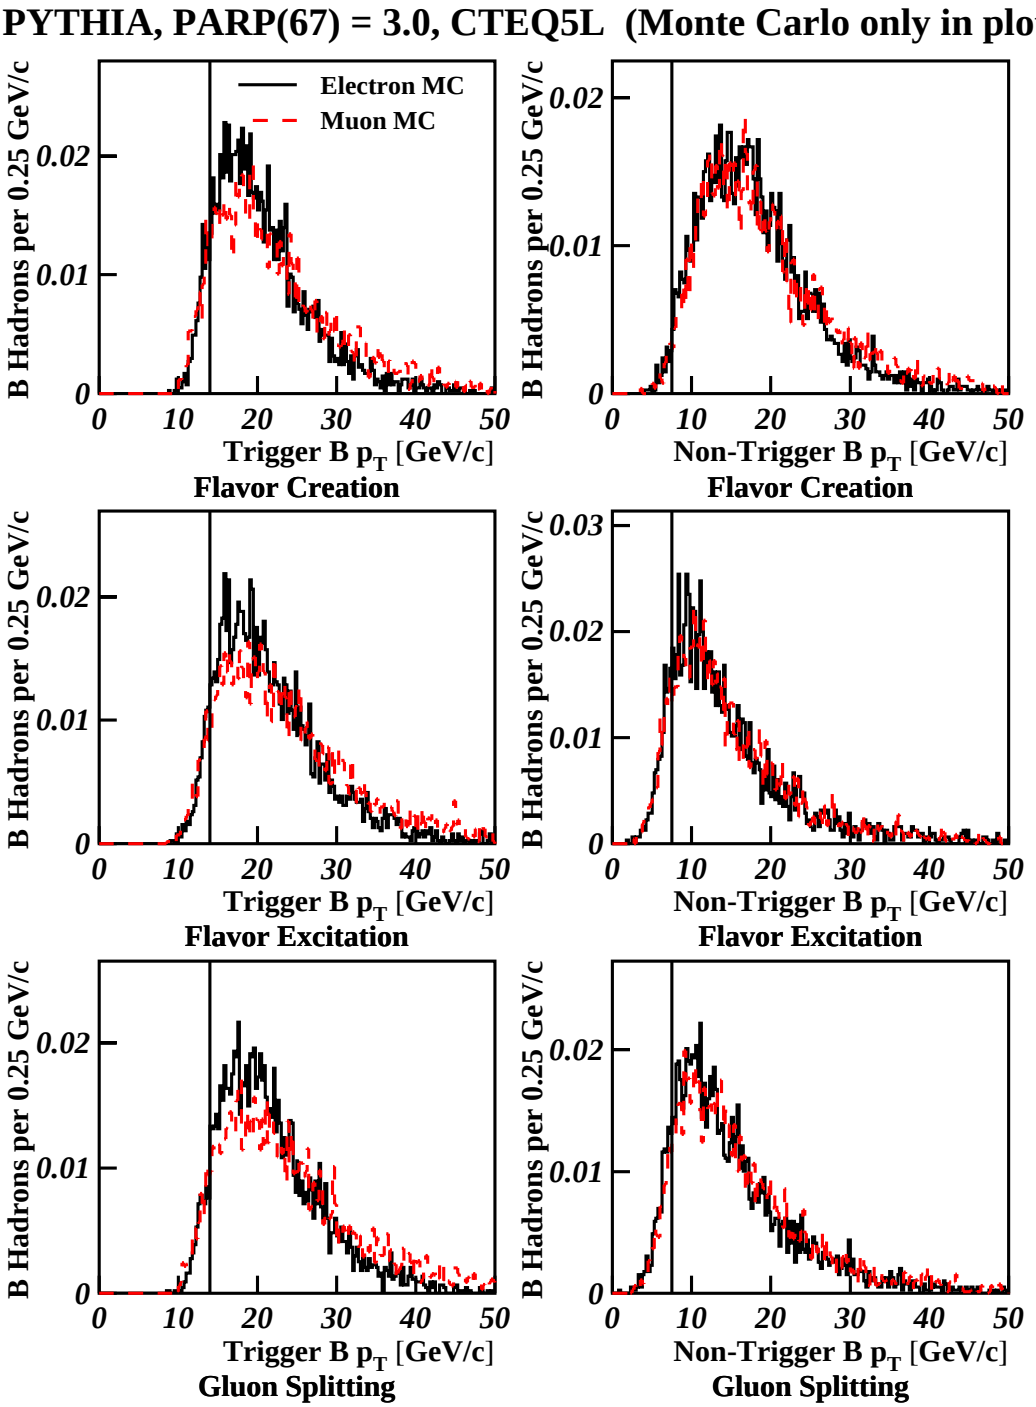

Figure 6.11 shows the B hadron p_T distributions for events from PYTHIA with `PARP(67)` = 3.0 where both B hadrons were tagged. The "Trigger B" is defined as the B whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger B" is the other B in the event. The solid vertical line represents the B hadron threshold quoted for this analysis.

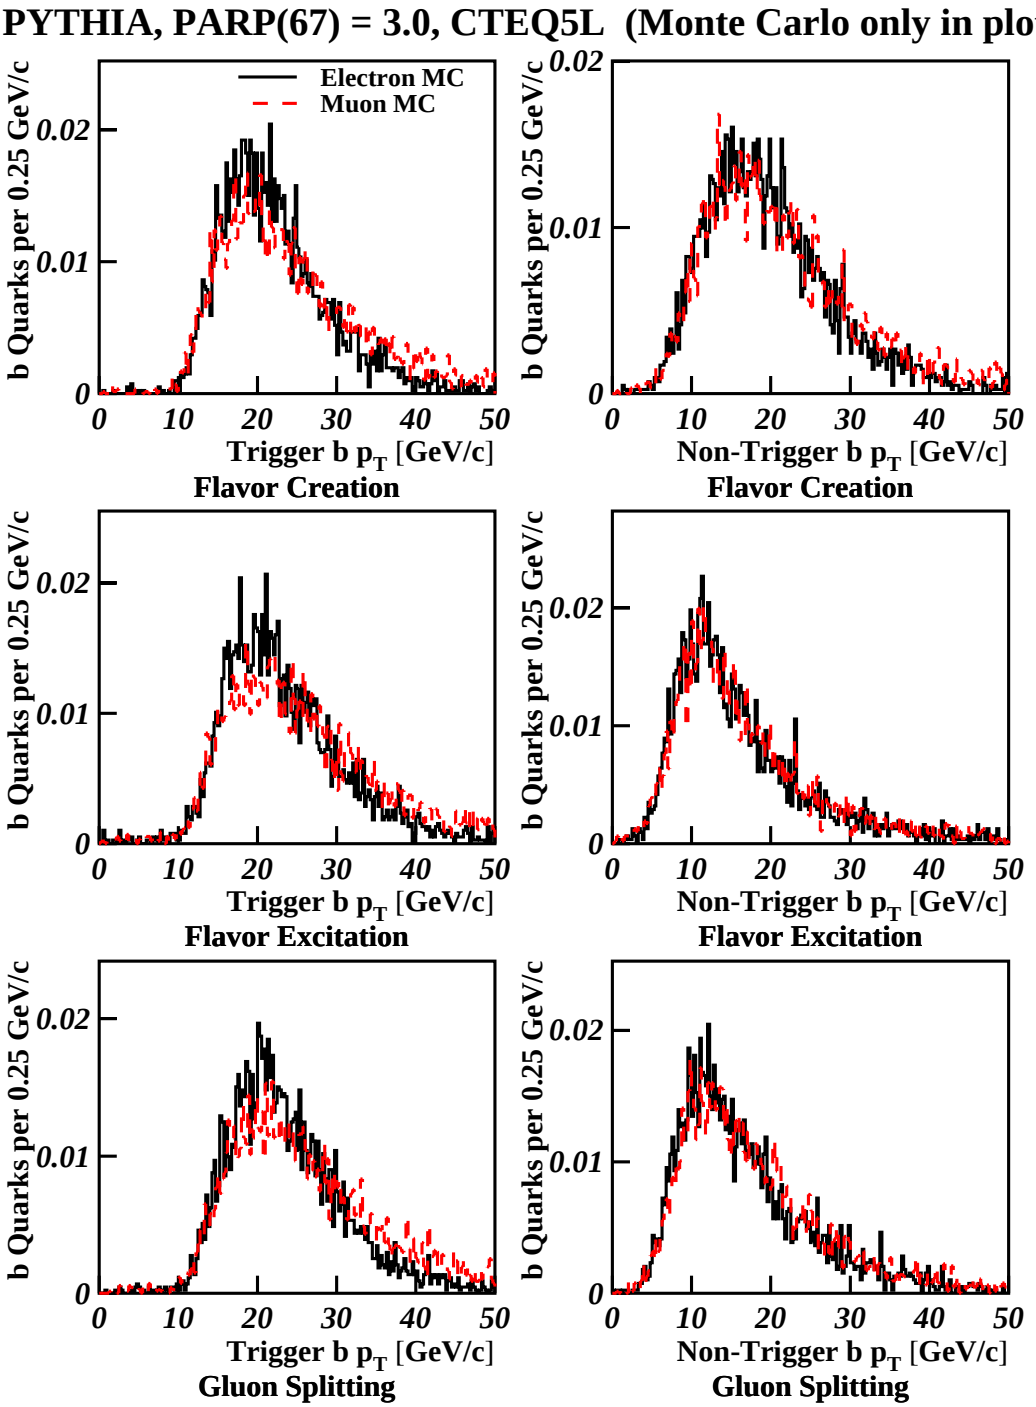

Figure 6.12 shows the *b* quark p_T distributions for events from PYTHIA with PARP(67) = 3.0 where both *B* hadrons were tagged. The "Trigger *b*" is defined as the *b* quark that hadronized to form a *B* hadron whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *b*" is the other *b* in the event.

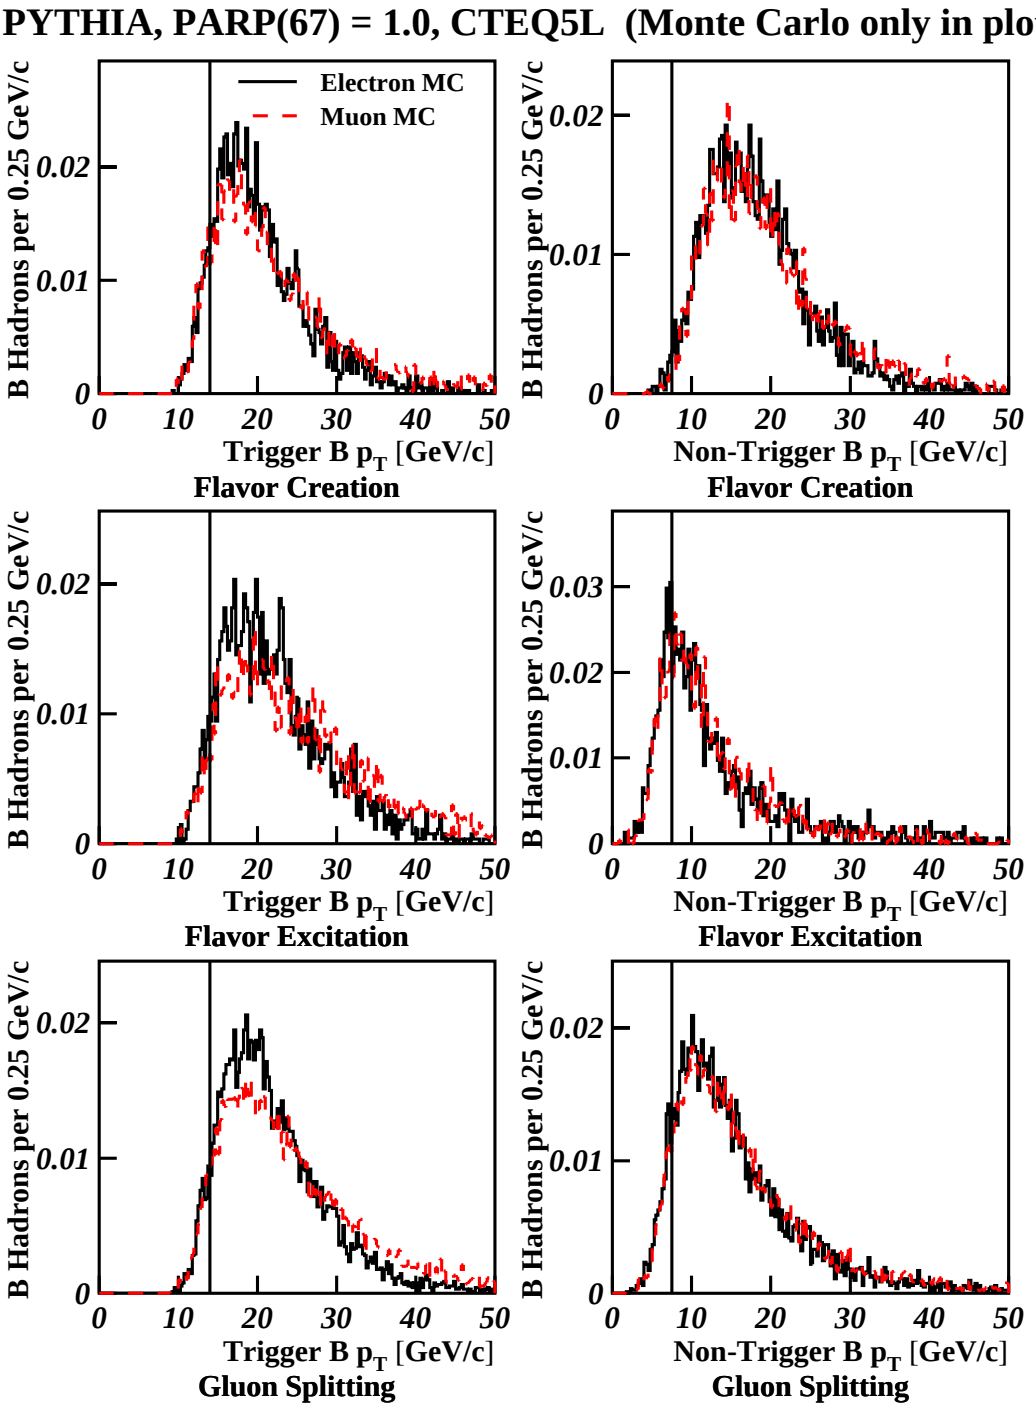

Figure 6.13 shows the B hadron p_T distributions for events from PYTHIA with `PARP(67)` = 1.0 where both B hadrons were tagged. The "Trigger B" is defined as the B whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger B" is the other B in the event. The solid vertical line represents the B hadron threshold quoted for this analysis.

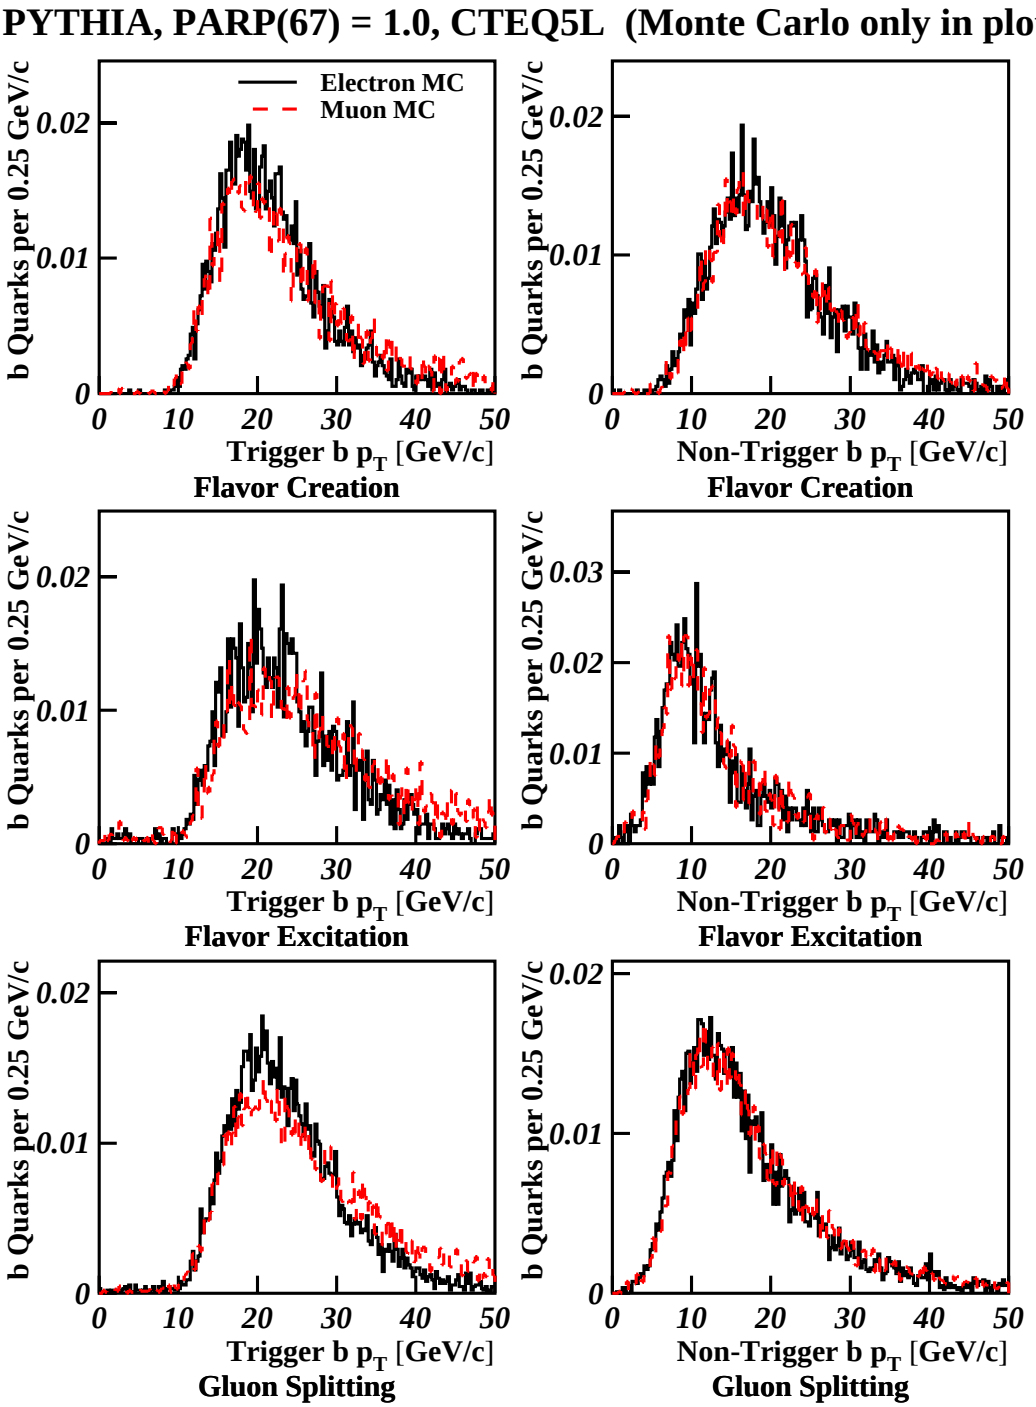

Figure 6.14 shows the *b* quark p_T distributions for events from PYTHIA with PARP(67) = 1.0 where both *B* hadrons were tagged. The "Trigger *b*" is defined as the *b* quark that hadronized to form a *B* hadron whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *b*" is the other *b* in the event.

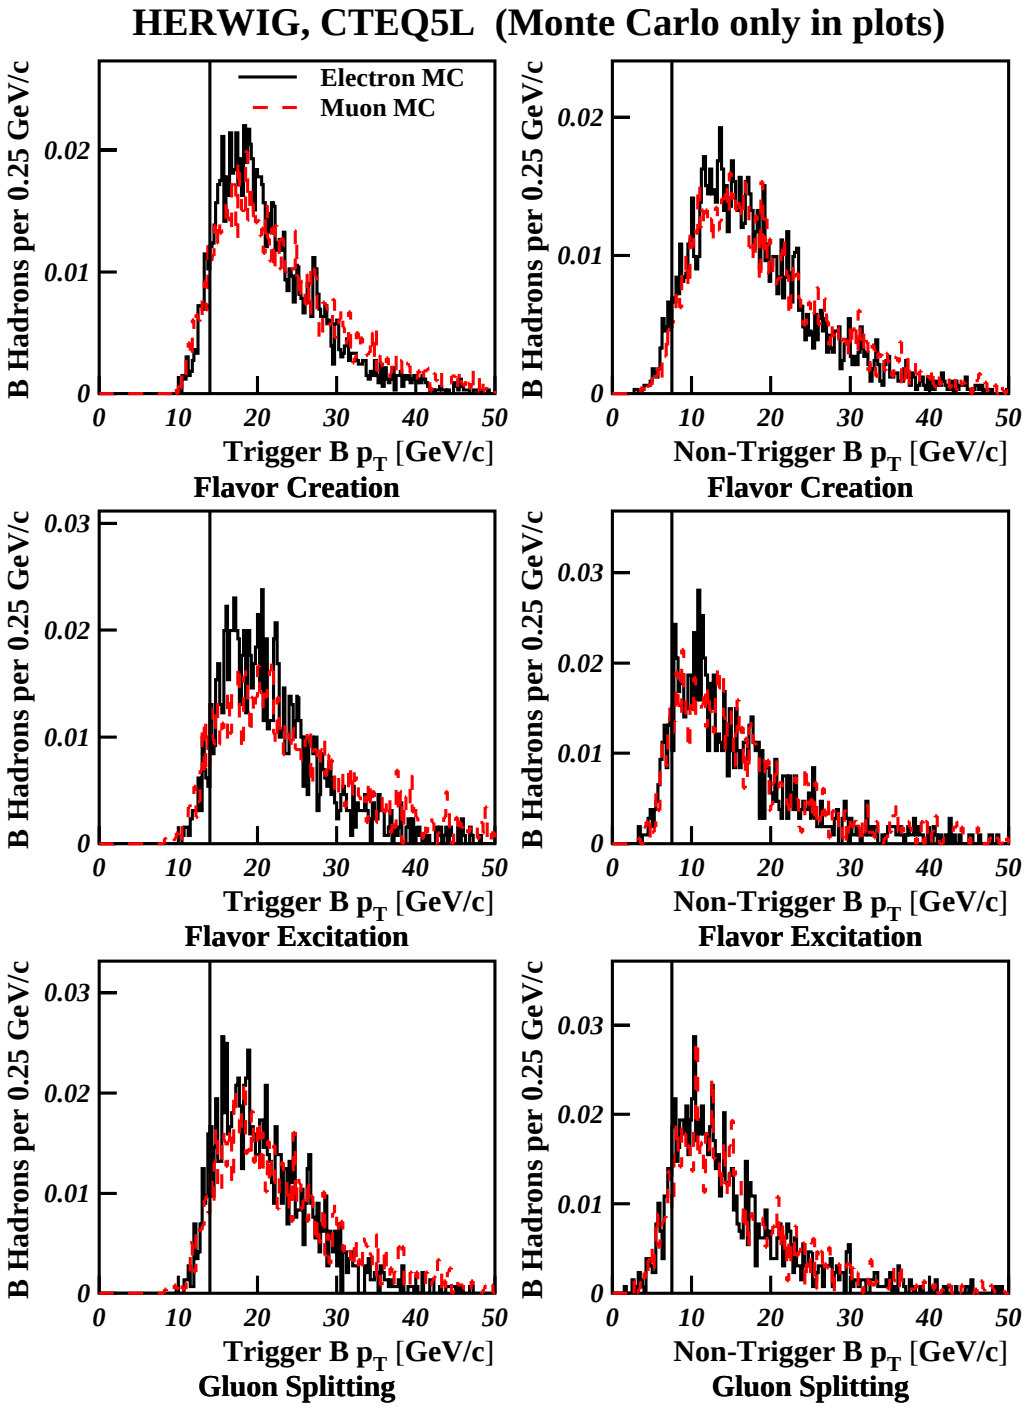

Figure 6.15 shows the *B* hadron p_T distributions for events from HERWIG where both *B* hadrons were tagged. The "Trigger *B*" is defined as the *B* whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *B*" is the other *B* in the event. The solid vertical line represents the *B* hadron threshold quoted for this analysis.

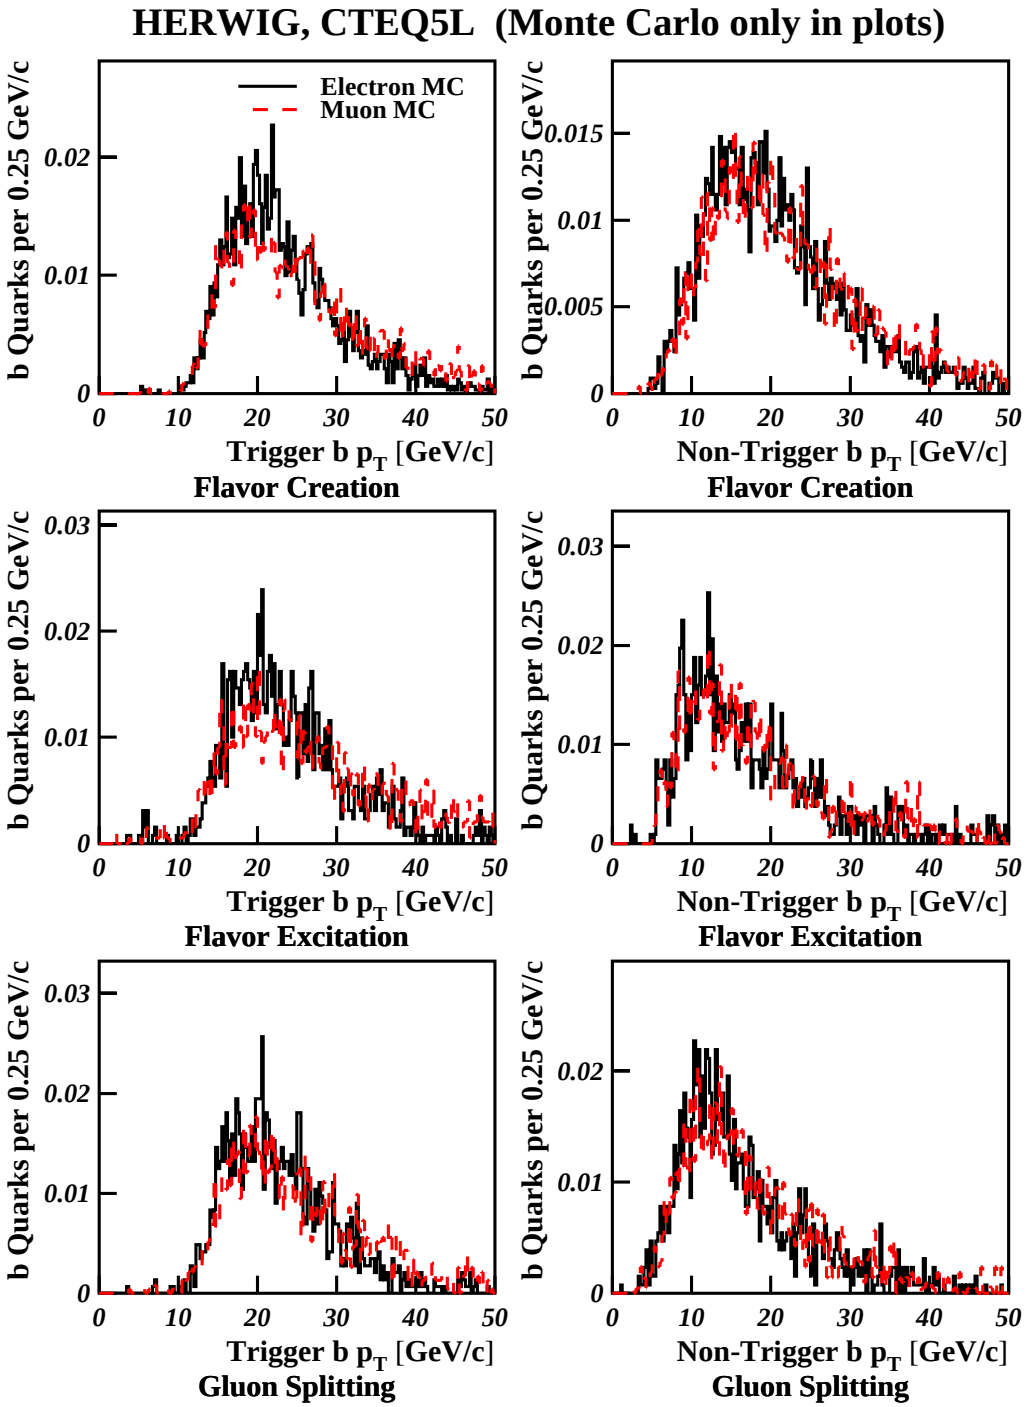

Figure 6.16 shows the *b* quark p_T distributions for events from HERWIG where both *B* hadrons were tagged. The "Trigger *b*" is defined as the *b* quark that hadronized to form a *B* hadron whose semileptonic decay provided an electron or muon for the trigger. The "Non-Trigger *b*" is the other *b* in the event.

Electron Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	20.5	14.0	17.8	10.2	21.9	14.6	19.5	10.7
fe	21.3	14.4	15.0	7.5	23.0	15.2	16.5	8.1
gs	21.9	14.6	15.0	7.4	23.7	15.4	16.5	8.1

Muon Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	22.1	14.0	19.2	10.4	23.7	14.6	20.9	11.1
fe	23.1	14.6	15.7	7.7	24.8	15.3	17.3	8.4
gs	23.8	14.7	16.0	7.7	25.6	15.7	17.5	8.3

Table 6.1 shows the average and the 90% threshold for the *B* hadron and *b* quark p_T distributions from PYTHIA with `PARP(67)` = 4.0. "Leptonic" *B* hadrons are *B* hadrons that decay semileptonically. "Non-leptonic" *B* hadrons and *b* quarks are associated with *B* hadrons that don't decay semileptonically.

Electron Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	20.6	14.1	18.0	10.1	22.2	14.7	19.8	10.7
fe	21.6	14.5	19.0	7.2	23.2	15.3	16.1	7.8
gs	21.7	14.6	15.1	7.4	23.6	15.5	16.8	8.1

Muon Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	22.2	14.0	19.2	10.6	23.9	14.6	21.0	11.0
fe	23.5	14.6	15.1	7.2	25.2	15.3	16.7	7.9
gs	23.7	14.6	16.1	7.7	25.6	15.8	17.8	8.6

Table 6.2 shows the average and the 90% threshold for the *B* hadron and *b* quark p_T distributions from PYTHIA with `PARP(67)` = 3.0. "Leptonic" *B* hadrons are *B* hadrons that decay semileptonically. "Non-leptonic" *B* hadrons and *b* quarks are associated with *B* hadrons that don't decay semileptonically.

Electron Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	20.4	14.0	18.5	11.0	22.0	14.5	20.4	11.6
fe	22.3	14.8	13.1	5.9	23.7	15.0	14.0	6.1
gs	21.9	14.8	15.5	7.6	23.7	15.6	17.0	8.3

Muon Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	21.9	13.9	19.6	11.3	23.5	14.5	21.4	12.0
fe	24.7	15.2	13.1	6.1	26.1	15.8	14.1	6.3
gs	23.8	14.9	16.1	7.8	25.7	15.9	17.7	8.6

Table 6.3 shows the average and the 90% threshold for the *B* hadron and *b* quark p_T distributions from PYTHIA with `PARP(67)` = 1.0. "Leptonic" *B* hadrons are *B* hadrons that decay semileptonically. "Non-leptonic" *B* hadrons and *b* quarks are associated with *B* hadrons that don't decay semileptonically.

Electron Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	21.4	14.5	18.0	9.4	23.4	15.5	20.6	11.0
fe	22.0	15.0	15.4	7.6	23.7	15.6	17.5	8.5
gs	21.1	14.4	14.8	7.5	22.9	15.4	16.5	8.3

Muon Monte Carlo (All numbers in GeV/c):

	***B* (leptonic)**		***B* (non-leptonic)**		***b* quark (leptonic)**		***b* quark (non-leptonic)**	
	Ave.	**90%**	**Ave**	**90%**	**Ave**	**90%**	**Ave**	**90%**
fc	22.9	14.5	19.3	9.9	25.0	15.6	22.0	11.4
fe	24.2	14.7	16.4	7.8	26.1	15.5	18.6	8.7
gs	23.0	14.7	15.6	7.6	24.8	15.7	17.4	8.5

Table 6.4 shows the average and the 90% threshold for the *B* hadron and *b* quark p_T distributions from HERWIG. "Leptonic" *B* hadrons are *B* hadrons that decay semileptonically. "Non-leptonic" *B* hadrons and *b* quarks are associated with *B* hadrons that don't decay semileptonically.

6.3 Tag Pair Correlations

Comparing tag pair correlations between the Monte Carlo samples and the data reveals whether PYTHIA or HERWIG provide an adequate model of the higher-order contributions to $b\bar{b}$ production. Disagreement between data and Monte Carlo may also suggest that background sources have not been adequately removed or accounted for. This analysis focuses on the transverse opening angle, $\Delta\phi$. For tag pairs, $\Delta\phi$ is defined as the angle between the p_T vectors determined by taking the vector sum of the p_T from all the tracks involved in the tag.

The $\Delta\phi$ distribution is interesting to study because it is sensitive to contributions from flavor excitation and gluon splitting. Also, the "broadness" of the back-to-back peak in $\Delta\phi$ is sensitive to the amount of initial-state radiation present in the Monte Carlo. Additional information can be gained by examining the p_T distribution of the two tags. The shape of the $\Delta\phi$ distribution and the relative contributions from the three production mechanisms depend on the p_T cuts placed each of the B hadrons. The lepton trigger requirement in this sample already places an effective, fairly high p_T cut on one B hadron in the event. However, more information can be gained about how well the Monte Carlo models the data by varying the p_T cut placed on the second B in the event and comparing the $\Delta\phi$ distribution of the Monte Carlo to that of the data.

There are two possible approaches to normalizing the relative contributions in Monte Carlo from flavor creation, flavor excitation, and gluon splitting. PYTHIA and HERWIG each provide predictions for the cross section of each production mechanism, and these cross sections can be used to normalize their contributions relative to one another. Alternatively, one could take the position that PYTHIA and HERWIG may not correctly model the amount of each contribution, and the relative contributions should be determined to provide the best match to data. In this analysis, both approaches will be examined. As described in the sections below, the data is compared to the Monte Carlo predictions in two ways. First, the Monte Carlo prediction for the cross section of each production mechanism is used to normalize the flavor excitation and gluon splitting components relative to the flavor creation contribution. In this "fixed normalization" scheme, the data is compared to the Monte Carlo using one arbitrary global normalization parameter. The arbitrary global normalization is included because this analysis attempts only a shape comparison, not an absolute cross section measurement. In addition, the Monte Carlo and data are compared using a "floating normalization" scheme. In this comparison, each production

mechanism is given an independent arbitrary normalization constant and the three normalizations are varied to yield the best match to data.

6.3.1 Fixed Normalizations

Both PYTHIA and HERWIG predict the relative normalizations of flavor creation, flavor excitation, and gluon splitting for $b\bar{b}$ production. The most straightforward test of agreement between the Monte Carlo models and the data would be to combine the individual $b\bar{b}$ production components from each Monte Carlo according to the cross sections predicted by the Monte Carlo generators and compare the shape of the $\Delta\phi$ spectrum between Monte Carlo and data. This comparison is shown in Figure 6.17. The contributions from the individual production mechanisms are shown in Figure 6.18 through Figure 6.21. Because this is a comparison of shape, an arbitrary overall normalization constant for the Monte Carlo predictions is fitted to get the best match between the Monte Carlo models and the data.

From these $\Delta\phi$ comparisons it can be seen that each Monte Carlo model matches the qualitative features of the data, although there are definite differences in shape, as reflected by the poor χ^2 values. For the PYTHIA sample with low initial-state radiation (ISR) (`PARP(67)` = 1.0), the peak in the back-to-back region is too narrow, while for the higher ISR samples (`PARP(67)` = 3.0 or 4.0), the back-to-back peak is too broad. Similarly, the HERWIG Monte Carlo sample also has a peak that is too broad at high $\Delta\phi$, perhaps even more so than in PYTHIA. However, aside from these discrepancies at high $\Delta\phi$, the rest of the $\Delta\phi$ distribution matches reasonably well between Monte Carlo and data using the normalizations predicted by the Monte Carlo generators for the different production mechanisms. The χ^2 values between the $\Delta\phi$ curves from Monte Carlo and data are listed in Table 6.5. On the basis of these χ^2 values, it appears that PYTHIA with `PARP(67)` = 3.0 provides the best match to data when using the Monte Carlo's default normalization for the three production mechanisms.

6.3.2 Floating Normalizations

Although the Monte Carlo generators predict the relative normalizations of the three components of $b\bar{b}$ production, it is possible that for one reason or another, the actual effective contribution from each component might be different. Since, in the parton shower approximation, the contributions from flavor creation, flavor excitation, and gluon splitting may be generated

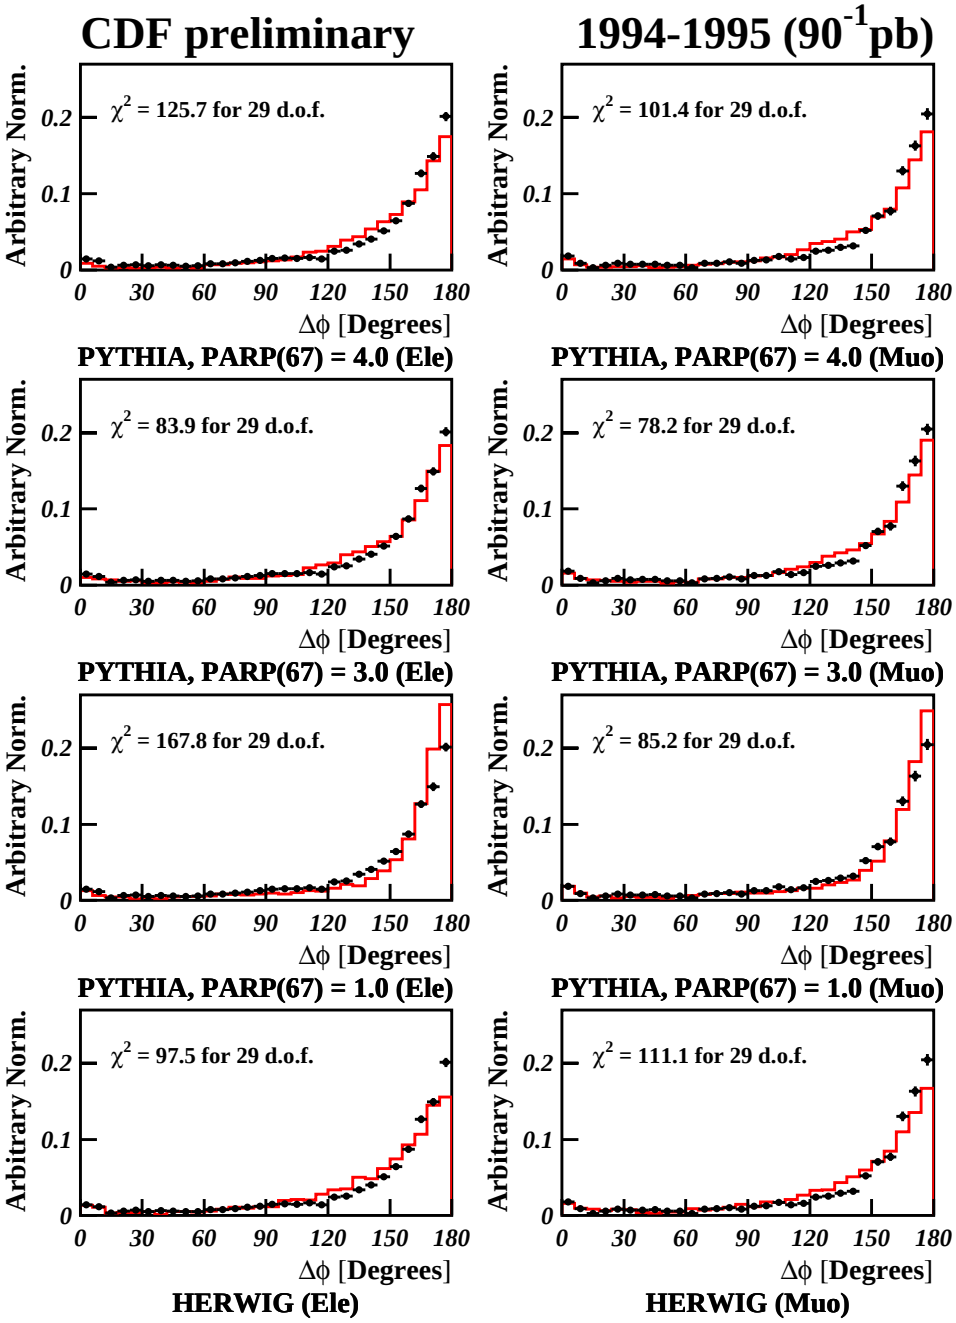

Figure 6.17 shows the comparison of the Δφ distribution from the data to the Δφ distribution from Monte Carlo when the three production mechanisms are combined according to the cross sections from each Monte Carlo. The data is shown as the points with statistical error bars only and Monte Carlo is the solid red line. The χ^2 value shown is for fit where only the overall normalization between Monte Carlo and data was allowed to vary. The fit χ^2 takes into account Monte Carlo statistics in addition to errors on the data.

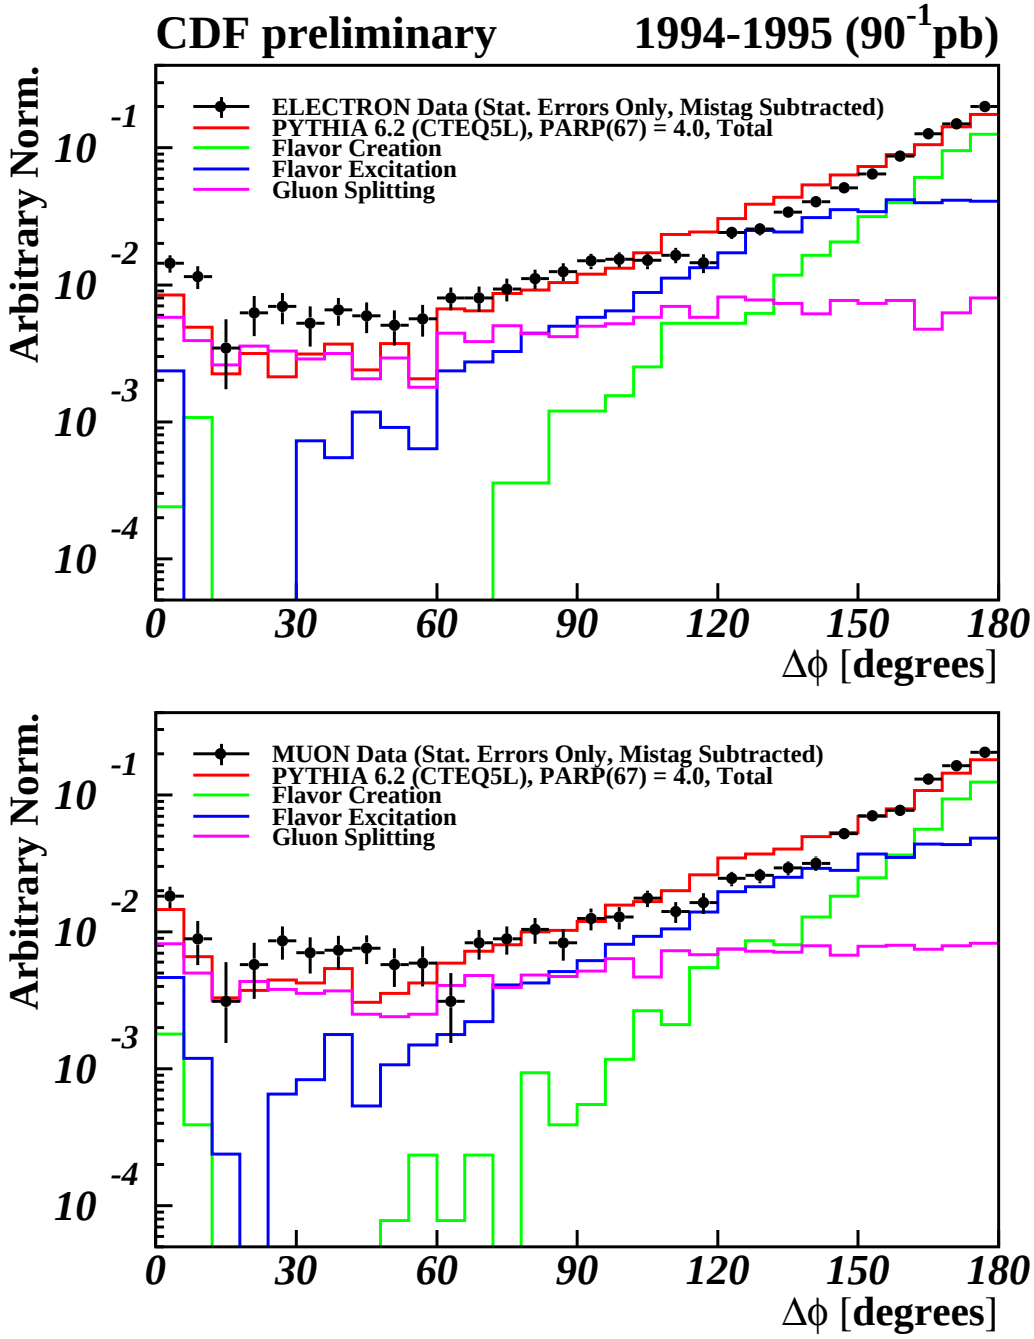

Figure 6.18 A detailed comparison between the $\Delta\phi$ distribution from data (points, statistical errors only) and the $\Delta\phi$ distribution from PYTHIA with `PARP(67)` = 4.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The contributions are normalized according to PYTHIA's cross section predictions and an arbitrary global normalization is used to give the best shape fit between data and Monte Carlo.

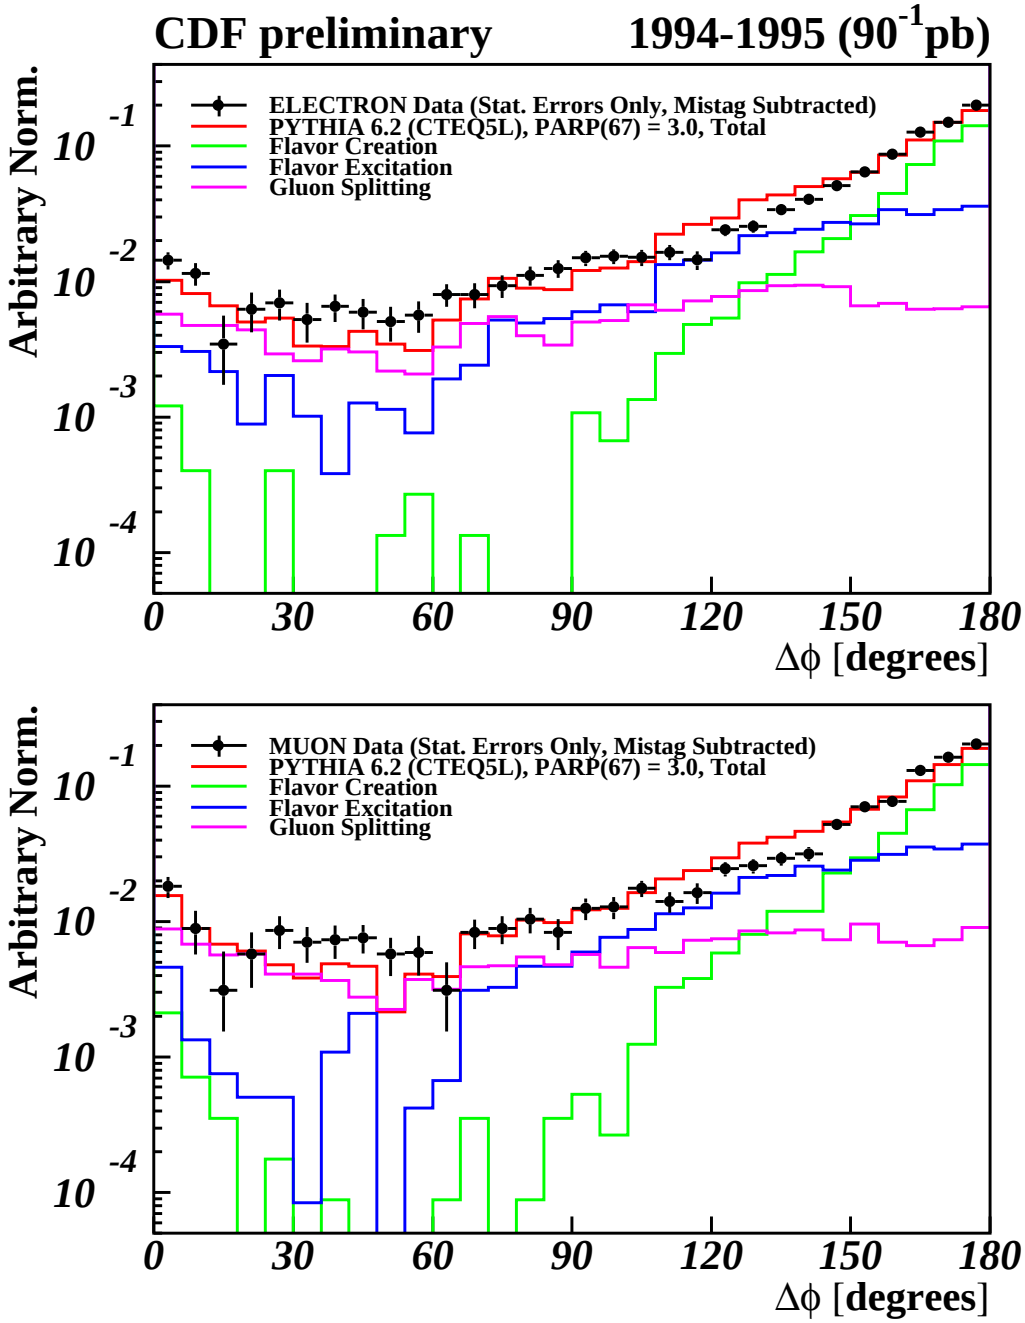

Figure 6.19 A detailed comparison between the Δϕ distribution from data (points, statistical errors only) and the Δϕ distribution from PYTHIA with `PARP(67)` = 3.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The contributions are normalized according to PYTHIA's cross section predictions and an arbitrary global normalization is used to give the best shape fit between data and Monte Carlo.

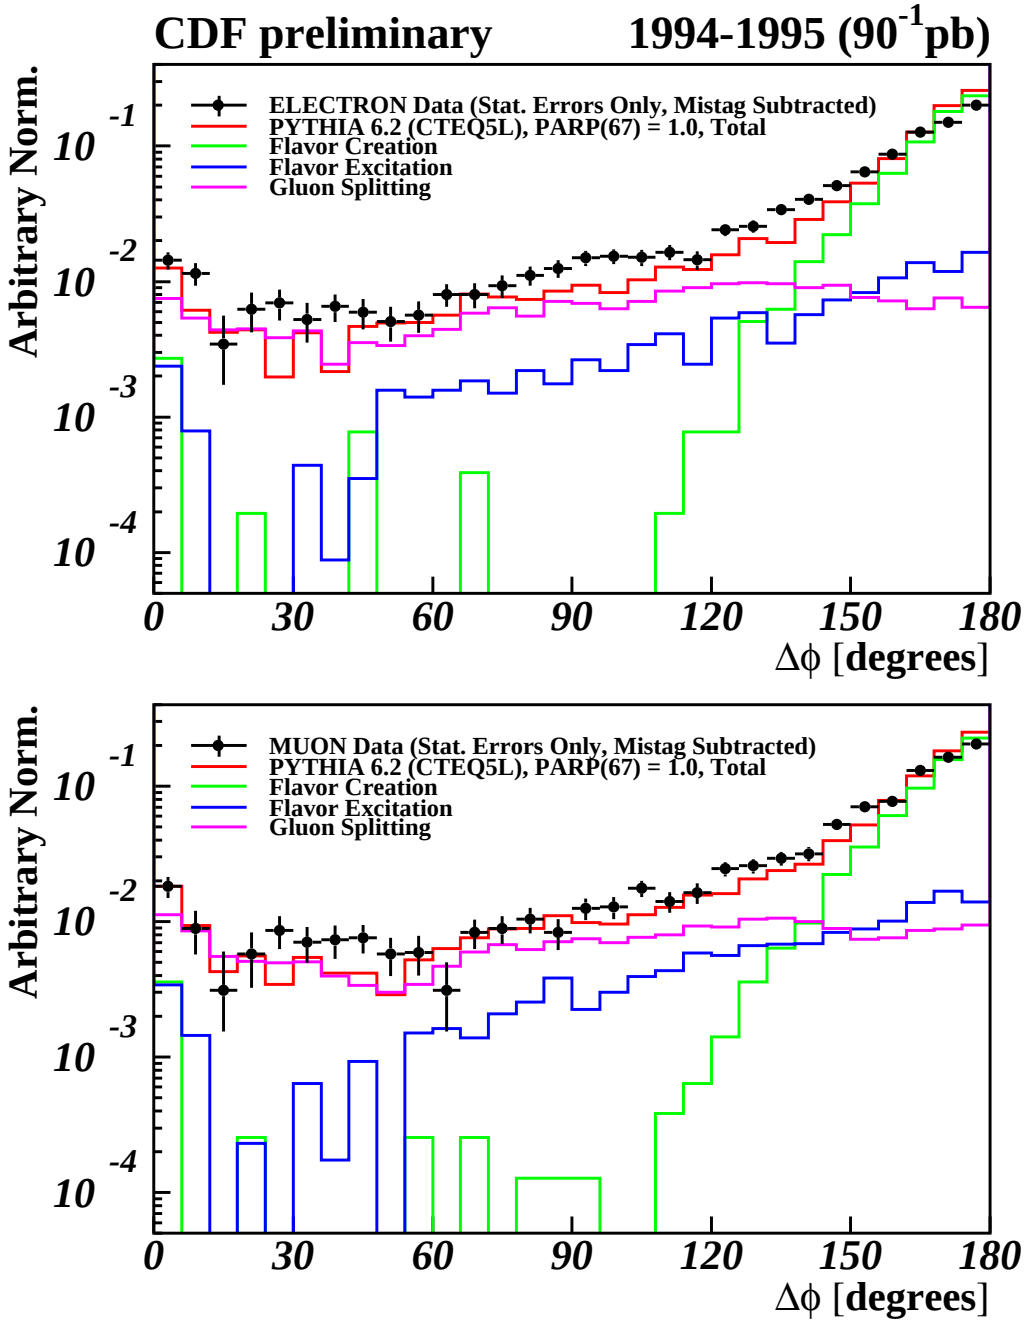

Figure 6.20 A detailed comparison between the $\Delta\phi$ distribution from data (points, statistical errors only) and the $\Delta\phi$ distribution from PYTHIA with `PARP(67)` = 1.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The contributions are normalized according to PYTHIA's cross section predictions and an arbitrary global normalization is used to give the best shape fit between data and Monte Carlo.

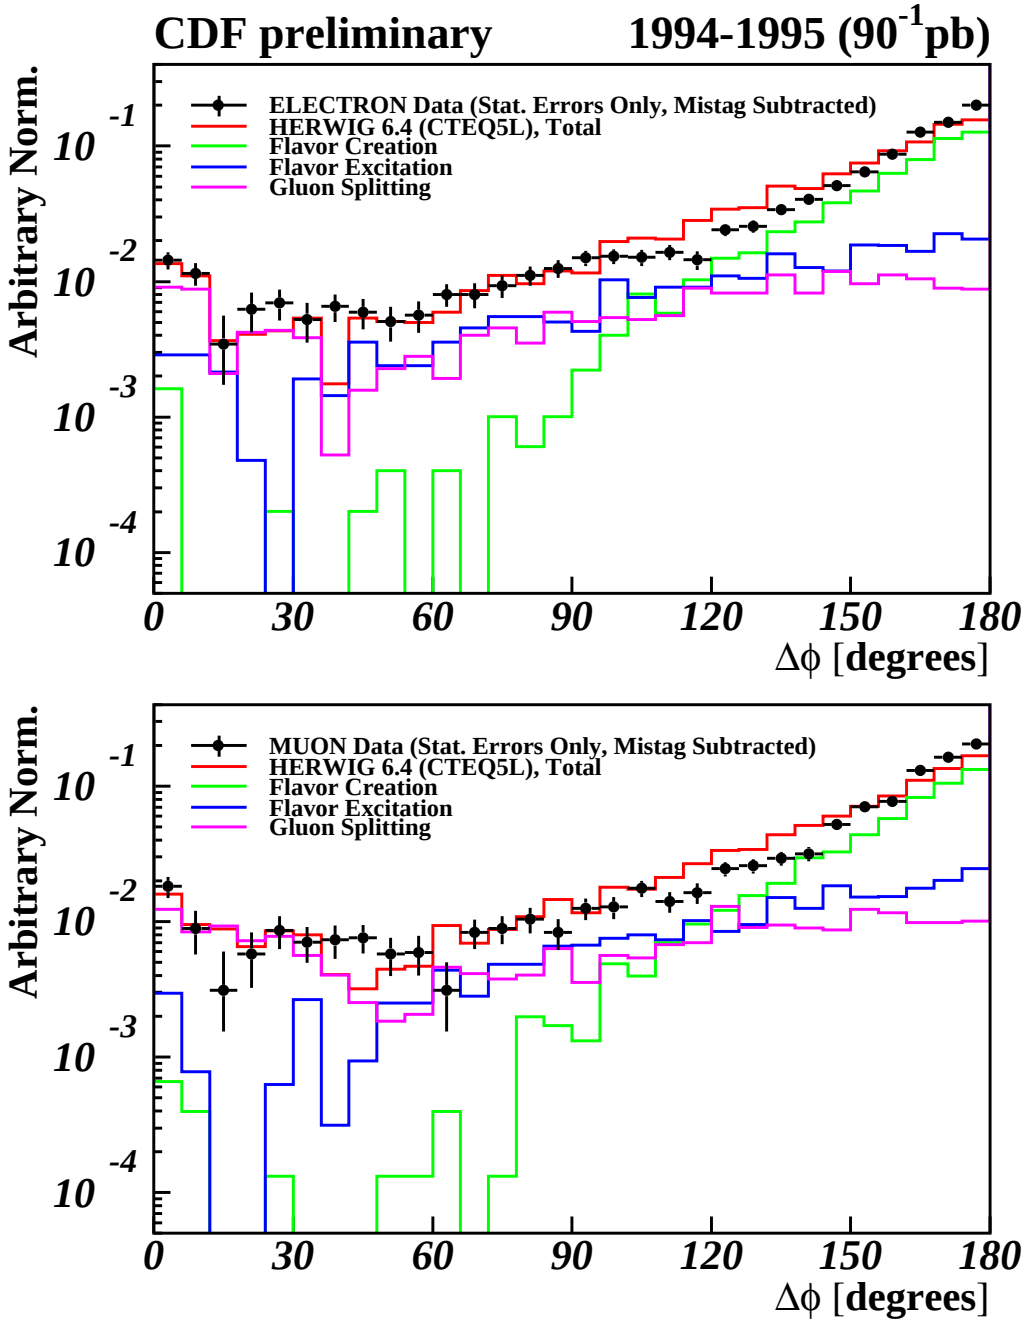

Figure 6.21 A detailed comparison between the $\Delta\phi$ distribution from data (points, statistical errors only) and the $\Delta\phi$ distribution from HERWIG (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The contributions are normalized according to HERWIG's cross section predictions and an arbitrary global normalization is used to give the best shape fit between data and Monte Carlo.

separately, each component can have a separate, arbitrary normalization and the three components can be fit for the combination of normalizations that gives the best match to the shape of the $\Delta\phi$ spectrum from data. These fits are shown in Figure 6.22. Figure 6.23 through Figure 6.26 show the relative contribution from each production mechanism yielded by the fits.

Surprisingly, when the normalizations of the individual components are allowed to float with respect to one another, one can obtain rather good agreement in shape between data and both the lower ISR and higher ISR PYTHIA samples. The fit of the lower ISR PYTHIA Monte Carlo to the data increases the broader contribution from flavor excitation to compensate for the narrowness of the back-to-back peak from flavor creation. For the higher ISR PYTHIA samples, the peak at high $\Delta\phi$ is made narrower to match the data by all but eliminating the contribution from flavor excitation. A comparison of the relative fractions of each production mechanism in the two PYTHIA fits is shown in Figure 6.27. The fit of the HERWIG sample to the data also tries to compensate for the excessive broadness of the HERWIG flavor creation peak at high $\Delta\phi$, but even after completely eliminating the flavor excitation contribution, the remaining contribution from flavor creation at high $\Delta\phi$ is too broad to model the data.

Table 6.5 compares the fit quality and effective contribution from flavor creation, flavor excitation, and gluon splitting in the fits of the various Monte Carlo samples to the data. That both low ISR and high ISR PYTHIA samples can be made to fit the data with approximately the same fit quality is unexpected, especially since the low ISR sample accomplishes this fit with a high flavor excitation content while the high ISR sample fits with almost no flavor excitation contribution. In the end, there seems to be an ambiguity in PYTHIA that allows a trade-off between initial state-radiation and the amount of flavor excitation.

6.3.3 Binned by p_T

Because of their differing flavor excitation content, the two PYTHIA fits may possess rather different B hadron p_T distributions and correlations. Adding p_T information to the fit, particularly information about the p_T of the B hadron in the event that does not decay semileptonically, should help to determine whether the lower or higher ISR PYTHIA sample better models the data. Therefore, we simultaneously fit the $\Delta\phi$ distribution of each Monte Carlo sample to the data separated into three bins in p_T of the non-leptonic tag. The p_T bins are defined as follows:

- Bin 1: 0 GeV/c $\leq$ tag p_T < 6 GeV/c

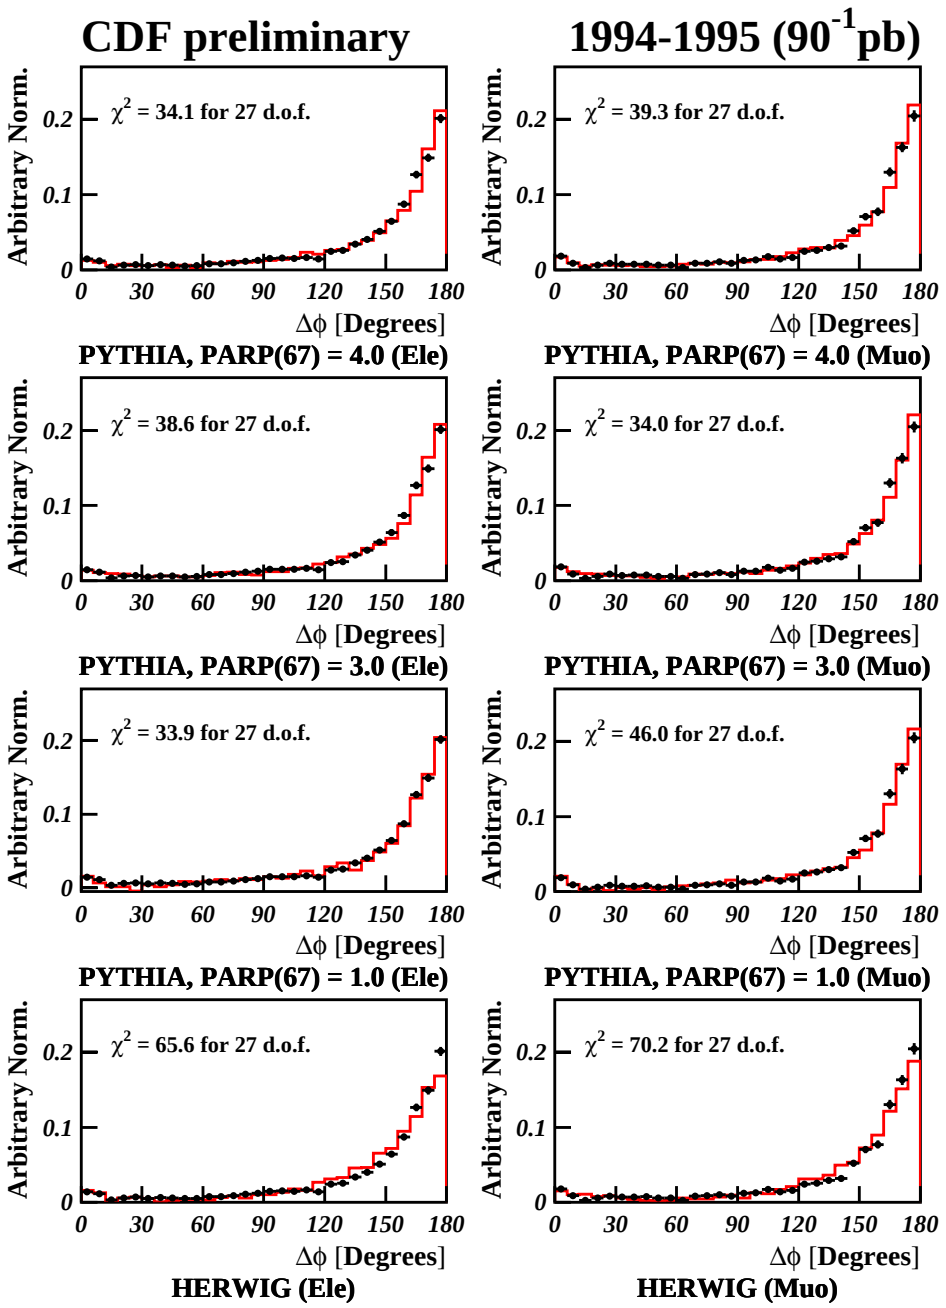

Figure 6.22 shows the comparison between the Δϕ distribution for Monte Carlo and data. The data is shown as the points with statistical error bars only and the Monte Carlo is the solid red line. In these comparisons, the normalizations of each production mechanism were allowed to vary independently and were chosen to give the best fit between the Monte Carlo and the data. Again, the fit χ^2 takes into account Monte Carlo statistics in addition to errors on the data.

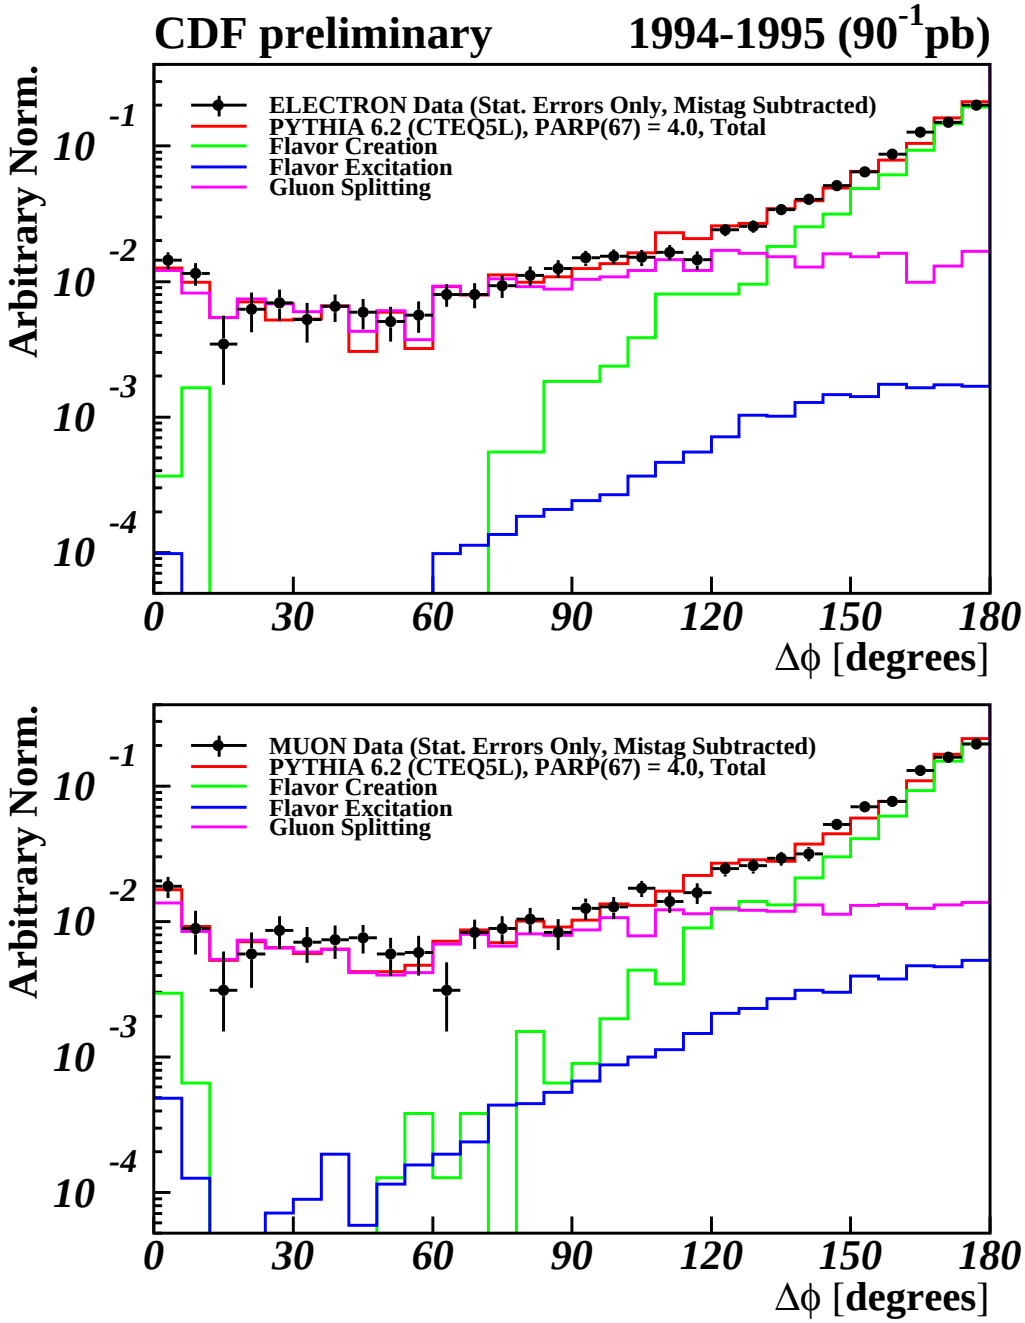

Figure 6.23 A detailed comparison between the $\Delta\phi$ distribution from data (points, statistical errors only) and the $\Delta\phi$ distribution from PYTHIA with `PARP(67)` = 4.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The normalization of each contribution is varied to achieve the best fit to the data.

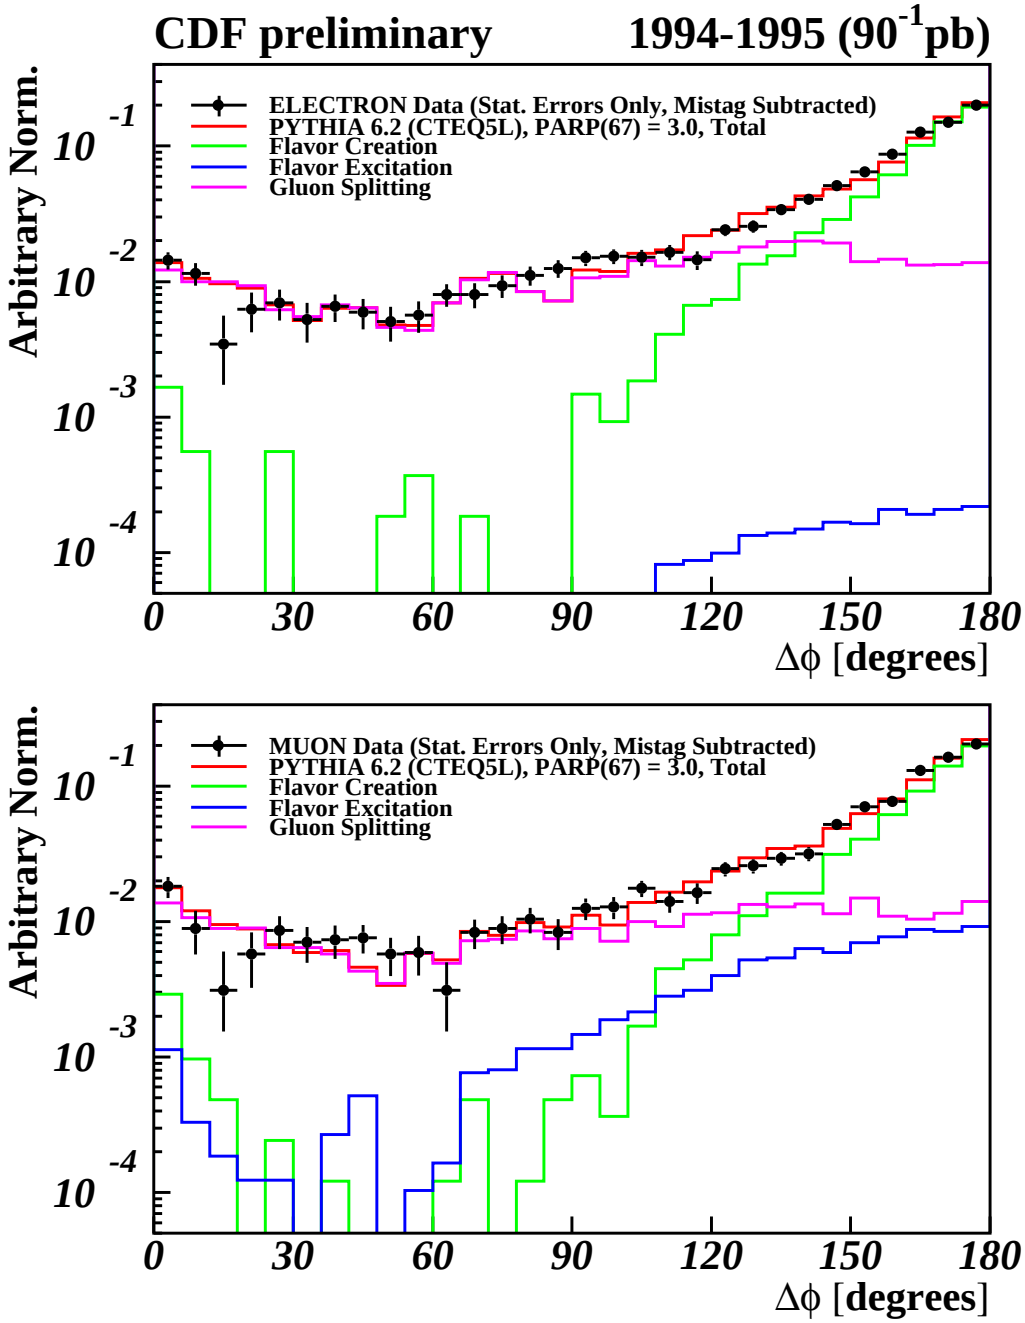

Figure 6.24 A detailed comparison between the Δϕ distribution from data (points, statistical errors only) and the Δϕ distribution from PYTHIA with `PARP(67)` = 3.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The normalization of each contribution is varied to achieve the best fit to the data.

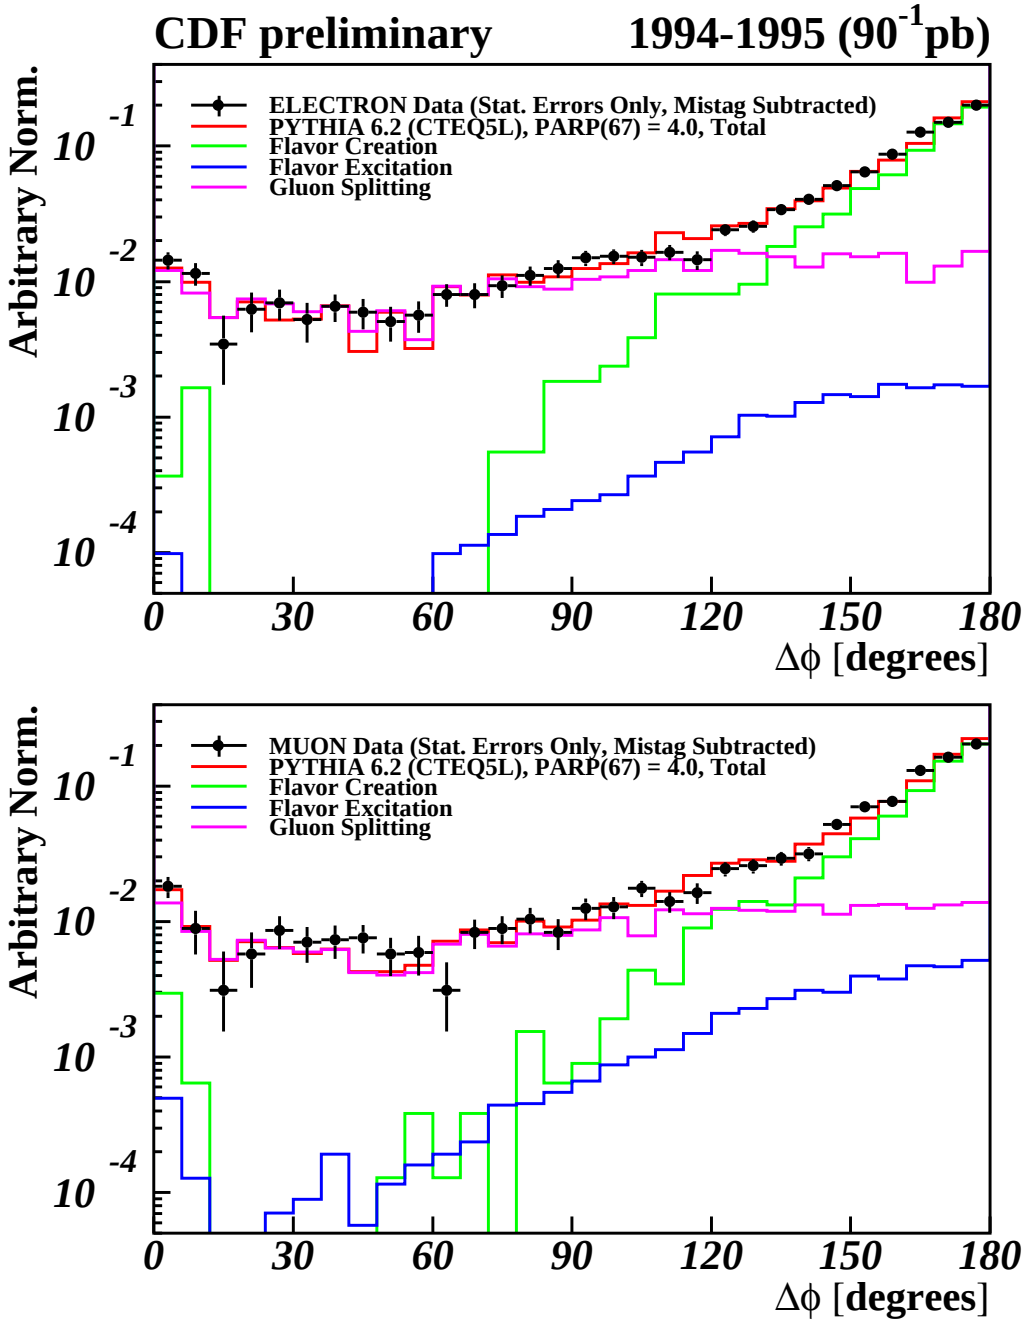

Figure 6.25 A detailed comparison between the $\Delta\phi$ distribution from data (points, statistical errors only) and the $\Delta\phi$ distribution from PYTHIA with `PARP(67)` = 1.0 (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The normalization of each contribution is varied to achieve the best fit to the data.

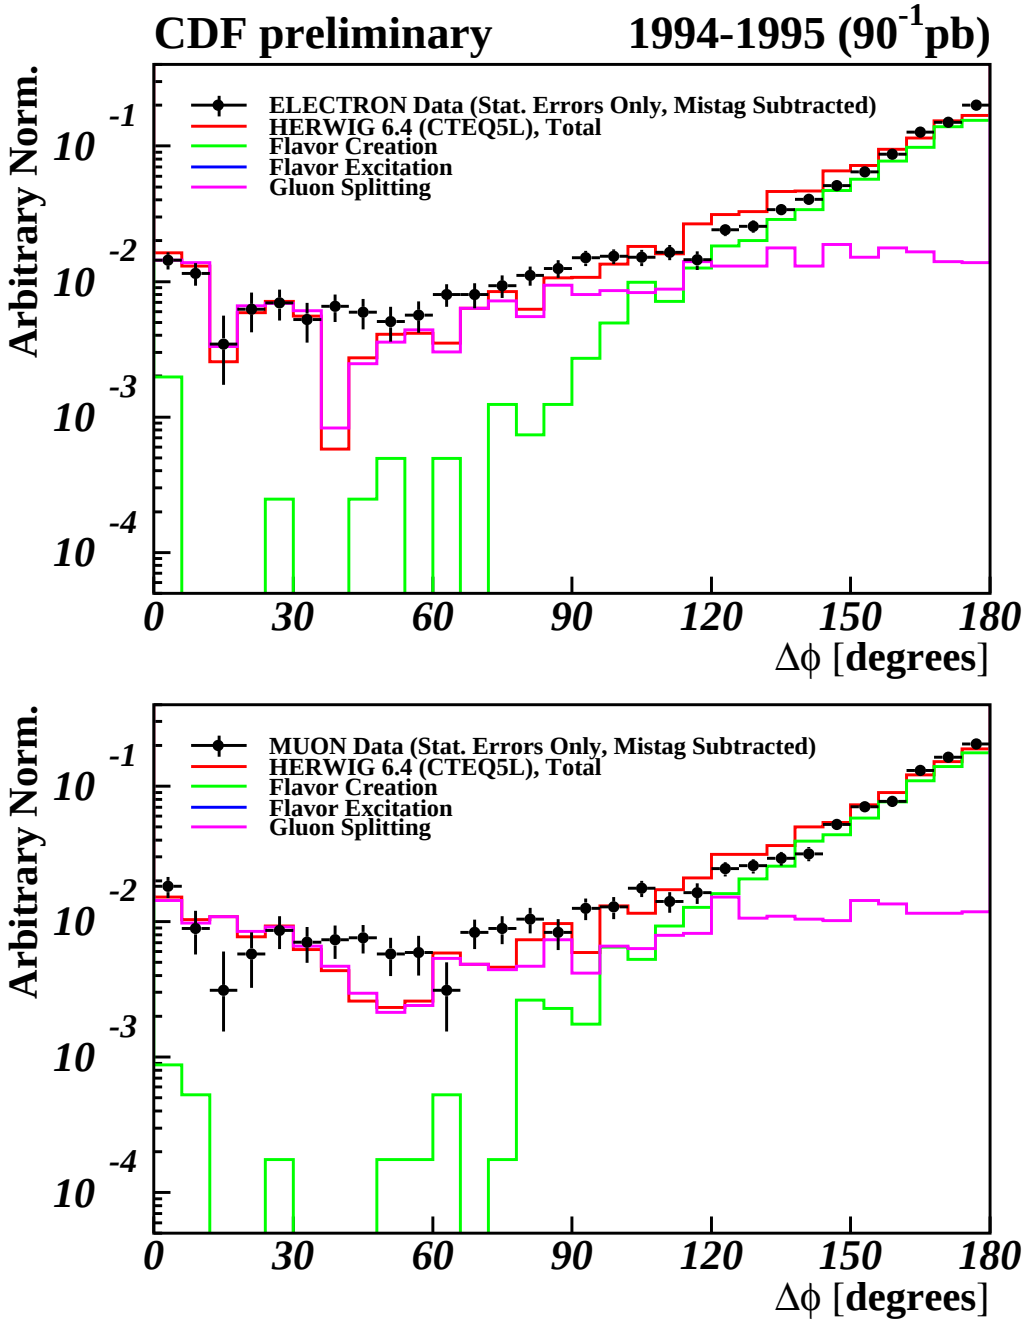

Figure 6.26 A detailed comparison between the Δϕ distribution from data (points, statistical errors only) and the Δϕ distribution from HERWIG (red line). In addition, the contributions from flavor creation (green), flavor excitation (blue), and gluon splitting (purple) are shown. The normalization of each contribution is varied to achieve the best fit to the data.

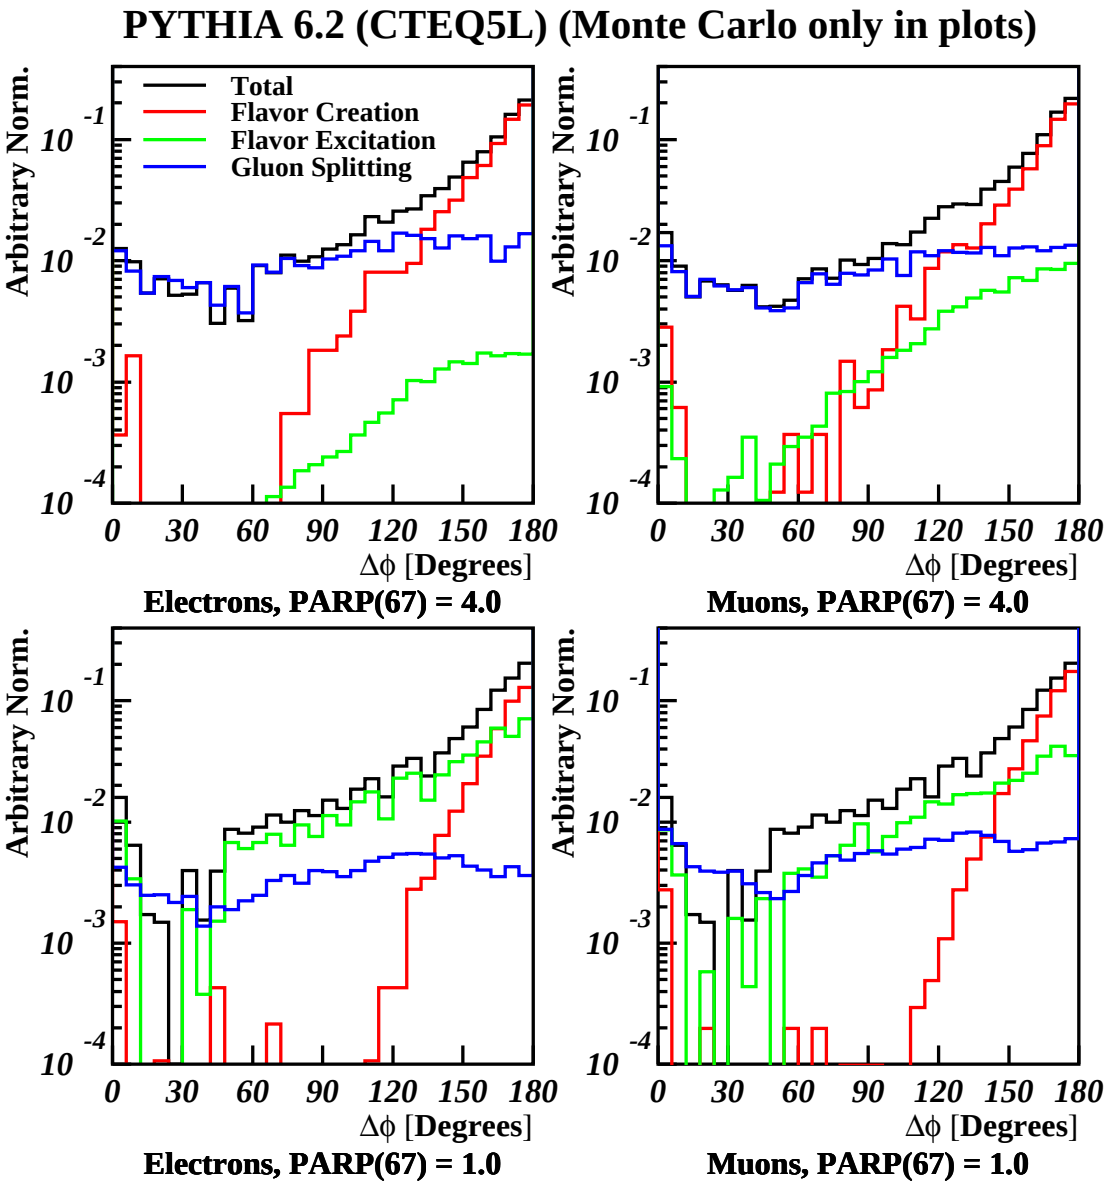

Figure 6.27 shows a comparison of the contributions from flavor creation (red), flavor excitation (green), and gluon splitting (blue) to the total $\Delta\phi$ shapes (black) for PYTHIA with `PARP(67)` = 4.0 (top) and `PARP(67)` = 1.0 (bottom). The normalization of each component is set by the best fit of the three components to the $\Delta\phi$ spectrum from data.

	Electrons		Muons	
	Fixed Normalization	**Floating Normalizations**	**Fixed Normalization**	**Floating Normalizations**
PYTHIA PARP(67) = 4.0	FC: 43.7% FE: 40.7% GS: 15.6%	FC: 66.1% FE: 1.7% GS: 32.2%	FC: 41.4% FE: 41.5% GS: 17.1%	FC: 64.5% FE: 8.1% GS: 27.4%
χ^2 / d.o.f.	125.7 / 29	34.1 / 27	101.4 / 29	39.3 / 27 d.o.f.
χ^2 probability	5.26×10^{-14}	0.163	5.83×10^{-10}	0.0595
PYTHIA PARP(67) = 3.0	FC: 47.7% FE: 35.8% GS: 16.5%	FC: 65.3% FE: 0.2% GS: 34.5%	FC: 46.5% FE: 35.3% GS: 18.2%	FC: 63.2% FE: 8.6% GS: 28.2%
χ^2 / d.o.f.	83.9 / 29	38.6 / 27	78.2 / 29	34.0 / 27
χ^2 probability	3.07×10^{-7}	0.0688	4.11×10^{-6}	0.166
PYTHIA PARP(67) = 1.0	FC: 68.3% FE: 12.0% GS: 19.7%	FC: 37.6% FE: 51.4% GS: 11.0%	FC: 63.9% FE: 13.9% GS: 22.2%	FC: 48.5% FE: 34.6% GS: 16.9%
χ^2 / d.o.f.	167.8 / 29	33.9 / 27	85.2 / 29	46.0 / 27
χ^2 probability	1.76×10^{-21}	0.169	1.96×10^{-7}	0.0127
Herwig	FC: 57.6% FE: 24.0% GS: 18.4%	FC: 70.9% FE: 0.0% GS: 29.1%	FC: 55.7% FE: 23.1% GS: 21.2%	FC: 74.8% FE: 0.0 GS: 25.2%
χ^2 / d.o.f.	97.5 / 29	65.6 / 27	111.1 / 29	70.2 / 27
χ^2 probability	2.45×10^{-9}	4.65×10^{-5}	1.52×10^{-11}	1.05×10^{-5}

Table 6.5 compares the effective contributions from flavor creation, flavor excitation, and gluon splitting to fits of the Monte Carlo $\Delta\phi$ to the data. The fit χ^2 takes into account Monte Carlo statistics in addition to errors on the data. The "χ^2 probability" entry refers to the probability of getting a worse fit, according to the χ^2 distribution.

- Bin 2: 6 GeV/c ≤ tag p_T < 10 GeV/c
- Bin 3: tag p_T ≥ 10 GeV/c

The fit results are summarized in Table 6.6. In this fit, the higher ISR PYTHIA samples fit with about the same relative contribution from each production mechanism as in the overall $\Delta\phi$ fit, although in each case, the fit quality is somewhat worsened. On the other hand, when p_T information is added to the fit for the lower ISR PYTHIA sample, a somewhat smaller contribution from flavor excitation is returned and the fit quality is drastically reduced. Figure 6.28 and Figure 6.29 compare the three different p_T slices of $\Delta\phi$ for the lower and higher ISR cases in PYTHIA. Adding p_T information to the HERWIG fit does not significantly alter the fit.

6.3.4 Best Fit

In general, any of the Monte Carlo samples compared to the data shows reasonable qualitative agreement. The Monte Carlo sample that best matches the data is PYTHIA with `PARP(67)` = 3.0 or 4.0 (higher ISR), when the individual normalizations of the flavor creation, flavor excitation, and gluon splitting are allowed to float separately to best fit the data. Although the fit using PYTHIA with `PARP(67)` = 1.0 is not so poor as to rule this model out completely, studies indicate that PYTHIA with higher initial state radiation does a better job of matching both the underlying event and minimum bias data at CDF [66]. Therefore, we select the PYTHIA sample, with `PARP(67)` = 4.0 and the relative normalizations of flavor creation, flavor excitation, and gluon splitting fixed by our fit to the $\Delta\phi$ distribution of the data, as the best Monte Carlo model of the data. Comparisons indicate that the differences between PYTHIA with `PARP(67)` = 4.0 and `PARP(67)` = 3.0 are minor. Figure 6.30 shows a comparison of other correlations between the data and PYTHIA with `PARP(67)` = 4.0.

It is interesting to note that before allowing the normalizations of each production mechanism to float in the fits, the agreement between HERWIG and the data is no worse than the agreement between the low ISR PYTHIA sample and the data. However, because the disparity between the data and the low ISR PYTHIA sample comes from the narrowness in the flavor creation peak at high $\Delta\phi$, when the normalizations are allowed to float, the fit can alleviate the disagreement by increasing the peak width through a higher contribution from flavor excitation. In contrast, for

HERWIG, once the contribution from flavor excitation has been reduced to zero, the fit has no way to make the width of the back-to-back flavor creation peak smaller, short of the unphysical situation of setting the flavor excitation normalization negative. If there were some other parameter for HERWIG, like PYTHIA's initial state radiation parameter, `PARP(67)`, that could be used to tune the width of the back-to-back flavor creation peak, it may be possible to achieve good agreement between HERWIG and the data as well.

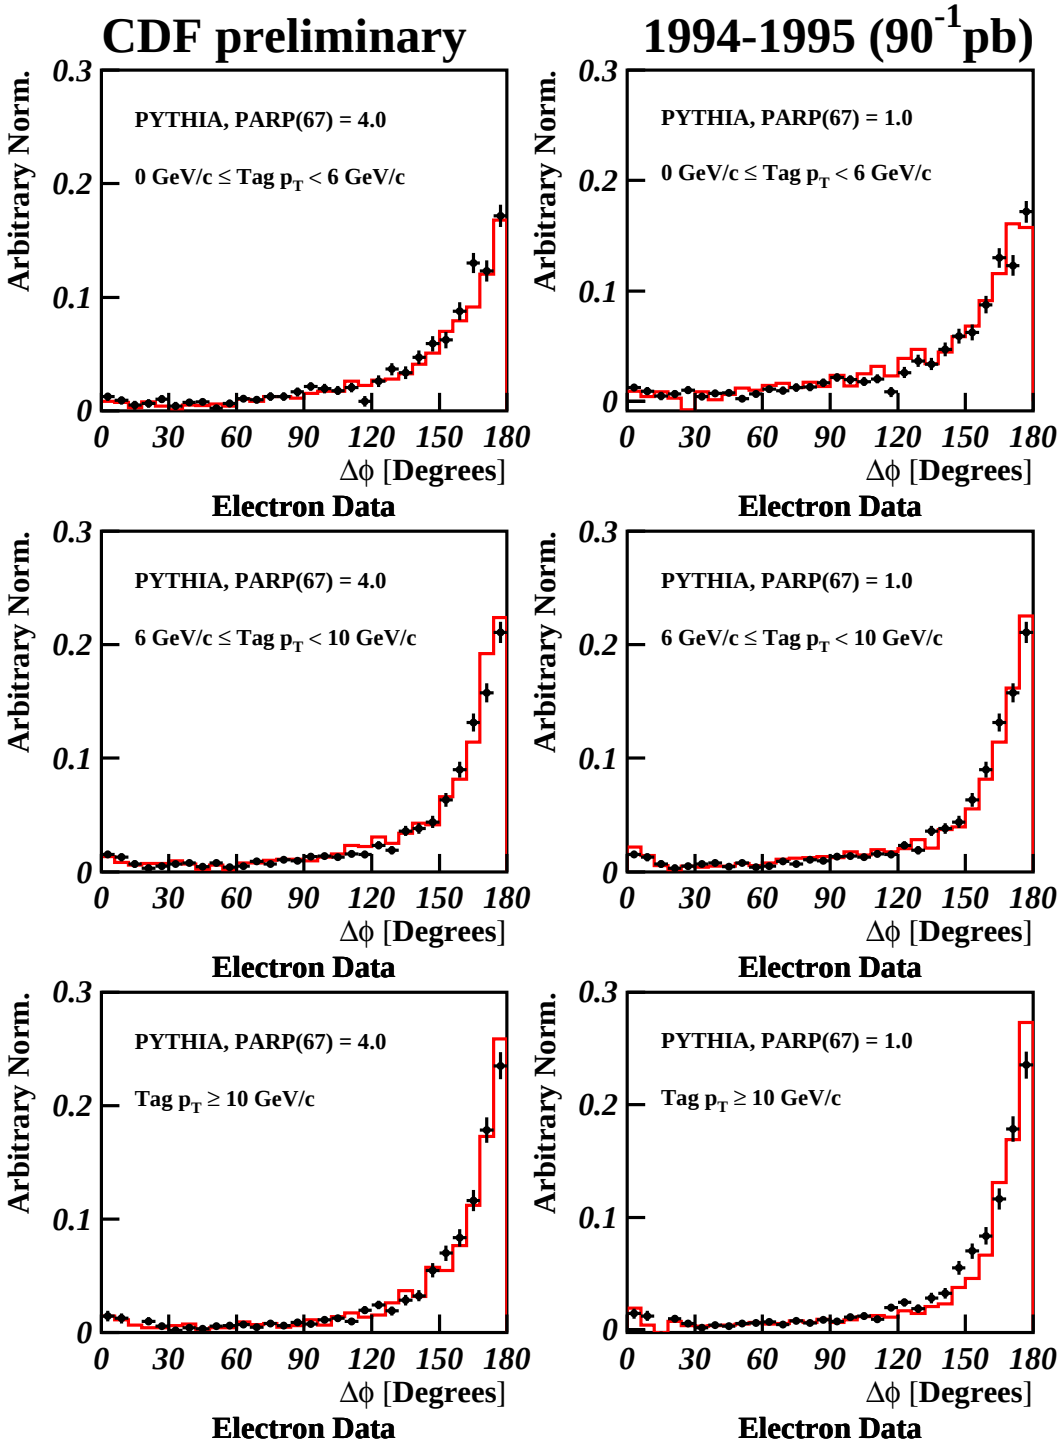

Figure 6.28 shows a comparison between the fit of the $\Delta\phi$ vs. p_T distribution for Monte Carlo to electron data for PYTHIA with `PARP(67)` = 4.0 and PYTHIA with `PARP(67)` = 1.0. The points with error bars (statistical only) are the data and the solid red line is the Monte Carlo.

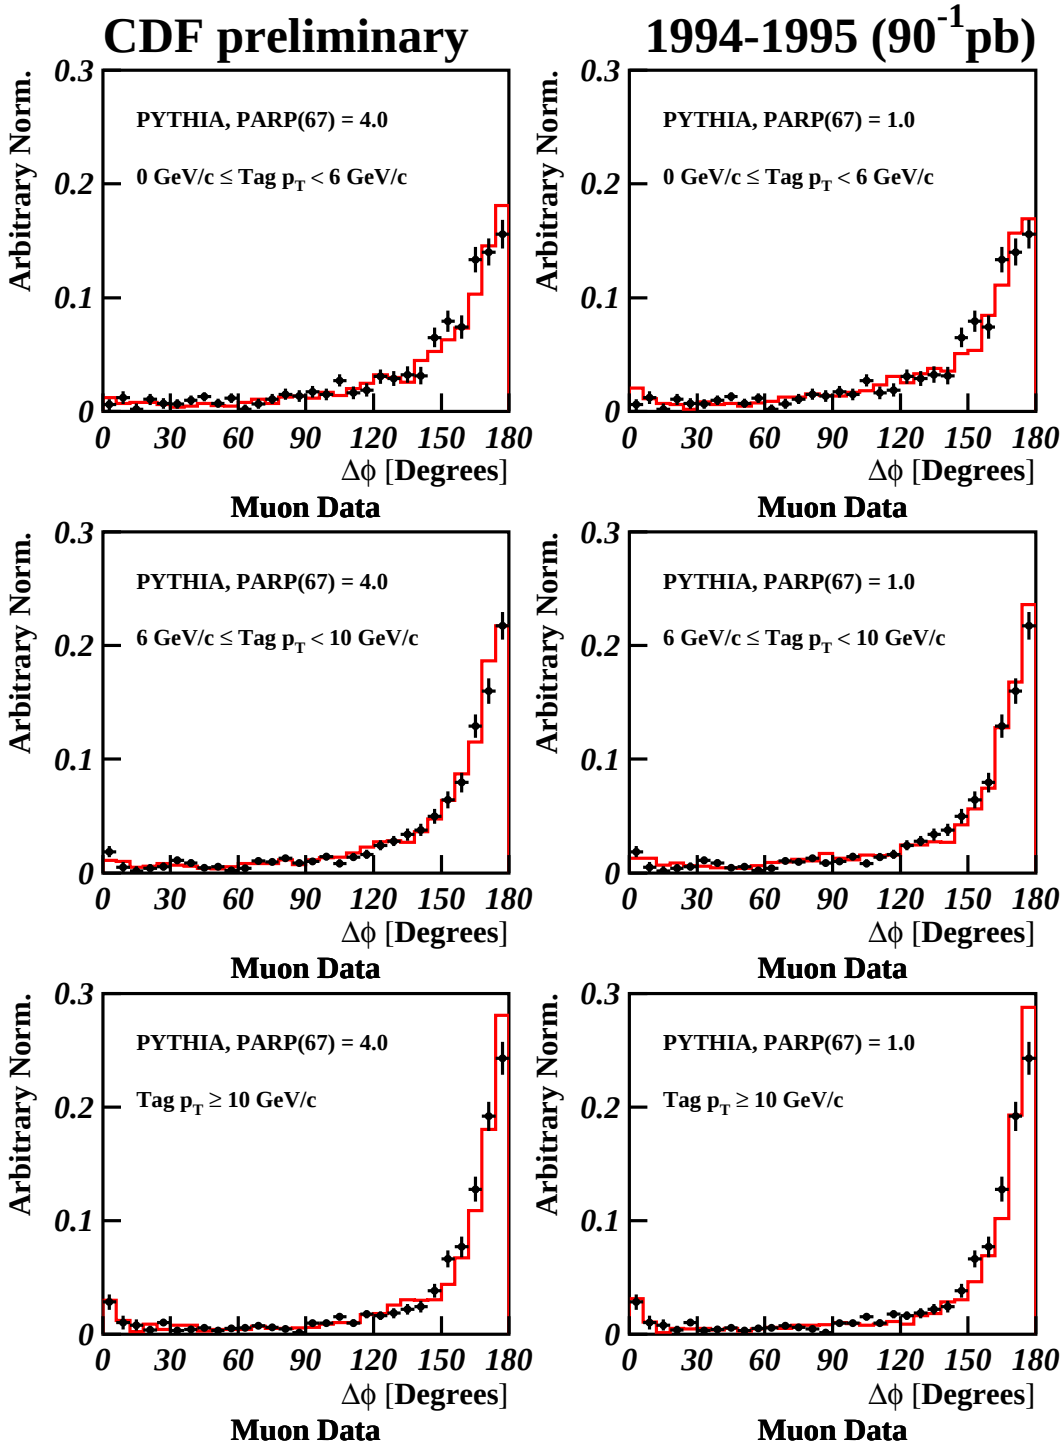

Figure 6.29 shows a comparison between the fit of the $\Delta\phi$ vs. p_T distribution for Monte Carlo to muon data for PYTHIA with `PARP(67)` = 4.0 and PYTHIA with `PARP(67)` = 1.0. The points with error bars (statistical only) are the data and the solid red line is the Monte Carlo.

	Electrons		**Muons**	
	$\Delta\phi$ Fit Only	**$\Delta\phi$ vs. p_T Fit**	**$\Delta\phi$ Fit Only**	**$\Delta\phi$ vs. p_T Fit**
PYTHIA `PARP(67)` = 4.0	FC: 66.1% FE: 1.7% GS: 32.2%	FC: 66.4% FE: 2.5% GS: 31.0%	FC: 64.5% FE: 8.1% GS: 27.4%	FC: 67.4% FE: 4.4% GS: 28.2%
χ^2 / d.o.f.	34.1 / 27	36.0 / 27	39.3 / 27	33.6 / 27
χ^2 probability	0.163	0.115	0.0595	0.177
PYTHIA `PARP(67)` = 3.0	FC: 65.3% FE: 0.2% GS: 34.5%	FC: 65.4% FE: 2.6% GS: 32.0%	FC: 63.2% FE: 8.6% GS: 28.2%	FC: 65.0% FE: 6.7% GS: 28.3%
χ^2 / d.o.f.	38.6 / 27	35.8 / 27	34.0 / 27	36.1 / 27
χ^2 probability	0.0688	0.120	0.166	0.113
PYTHIA `PARP(67)` = 1.0	FC: 37.6% FE: 51.4% GS: 11.0%	FC: 44.2% FE: 34.8% GS: 21.0%	FC: 48.5% FE: 34.6% GS: 16.9%	FC: 54.8% FE: 19.0% GS: 26.2%
χ^2 / d.o.f.	33.9 / 27	44.6 / 27	46.0 / 27	50.9 / 27
χ^2 probability	0.169	0.0179	0.0127	0.00358
HERWIG	FC: 70.9% FE: 0.0% GS: 29.1%	FC: 70.8% FE: 0.0 GS: 29.2%	FC: 74.8% FE: 0.0 GS: 25.2%	FC: 73.5% FE: 0.0% GS: 26.5%
χ^2 / d.o.f.	65.6 / 27	64.6 / 27	70.2 / 27	72.5 / 27
χ^2 probability	4.65×10^{-5}	6.38×10^{-5}	1.05×10^{-5}	4.90×10^{-6}

Table 6.6 compares the $\Delta\phi$ only fit of the Monte Carlo to data, to the fit incorporating p_T information as well. The fit χ^2 takes into account Monte Carlo statistics in addition to errors on the data. The "χ^2 probability" entry refers to the probability of getting a worse fit, according to the χ^2 distribution.

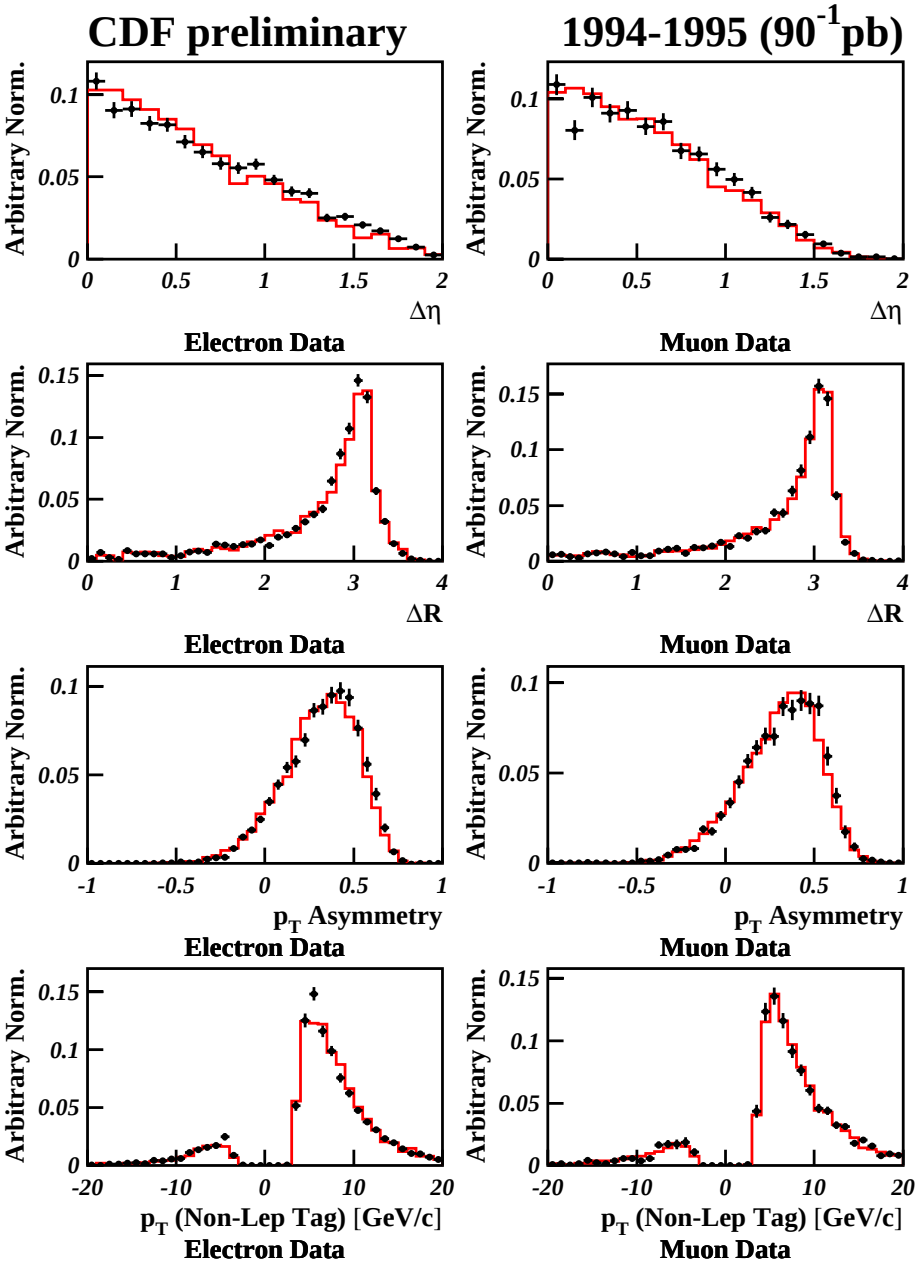

Figure 6.30 shows a comparison of PYTHIA with `PARP(67)` = 4.0 to the data for several different correlations. The normalizations for the three production mechanisms in PYTHIA have been determined by the fit of the PYTHIA $\Delta\phi$ vs. p_T distribution to data. The p_T asymmetry is given by $A_{p_T} = \left(p_T(lep) - p_T(non-lep)\right)/\left(p_T(lep) + p_T(non-lep)\right)$. In the p_T(Non-Lep Tag) plot, the sign of the p_T is determined by the opening angle between the lepton-tag and the non-lepton tag: negative for tag pairs with $\Delta\phi < 90°$, positive otherwise. The data are shown as points with statistical error bars only. The red line is PYTHIA with `PARP(67)` = 4.0.

Chapter 7

Corrections and Systematic Errors

The correlations examined so far in the data involve pairs of BVTX tags, rather than pairs of *B* hadrons. There are detector effects, such as the tagging efficiency for pairs of *B* hadrons as a function of $\Delta\phi$, that distort the shape of the measured tag pair correlations from the true *B* hadron distribution. In addition, residual contributions from background can affect the shape of the tag pair distribution. For the comparison between Monte Carlo and data, the detector effects are accounted for by using a detector and trigger simulation to adjust the Monte Carlo to match the conditions in the data, while the backgrounds are assumed to be negligible. However, since the Monte Carlo models examined in Chapter 6 match the data reasonably well, Monte Carlo events can be used to determine the relationship between the measured tag pair distribution and the actual *B* hadron distribution. In the sections below, two kinds of corrections to the tag pair $\Delta\phi$ distribution are considered: a correction for the relative tagging efficiency, which is a detector effect, and a correction for the contributions from mistags, prompt charm, and sequentials that remain in the data after the steps taken in Section 4.4 to remove backgrounds. In addition, the Monte Carlo can be used to estimate the systematic uncertainties associated with correcting for the relative tagging efficiency and removing background events. These corrections and systematic errors are evaluated using several different Monte Carlo samples to account for uncertainties involved in the Monte Carlo model itself.

7.1 Relative Tagging Efficiency

The BVTX tagging algorithm is not equally effective for all topologies of $b\bar{b}$ production. In particular, as the opening angle between the two *B* hadrons decreases, the BVTX algorithm has a more difficult time of reconstructing both displaced secondary vertices. This effect becomes especially severe when the two *B* hadrons are both contained within the cone of a single jet for track clustering purposes. Furthermore, correlations between opening angle and *B* p_T for the

various $b\bar{b}$ production mechanisms can lead to differences in the relative efficiency for reconstructing tag pairs at different opening angles. These effects distort the shape of the $\Delta\phi$ distribution measured for tags from the true B–B $\Delta\phi$ distribution.

We correct for these relative efficiency effects using the Monte Carlo that best matches the data, as determined in Section 6.3.4. Because we are only examining the shape of the $\Delta\phi$ distribution, our goal in making this correction is only to account for differences in the relative efficiency of the tagging algorithm, as a function of $\Delta\phi$. We do not attempt to correct for effects that impact all parts of the $\Delta\phi$ spectrum equally. For example, an overall shift in the muon trigger efficiency would not affect this correction. To determine the correction for each bin we take the ratio of the number of tag-pairs reconstructed in the Monte Carlo to the number of pairs that could have been reconstructed if the tagging algorithm had perfect efficiency. The number of tag pairs that would have been reconstructed assuming perfect efficiency is determined by looking at the generator level B hadron $\Delta\phi$ distribution. For electron Monte Carlo, to simulate the electron trigger, we require one B hadron in the event to have a $p_T > 14.0$ GeV/c and $|\eta| < 1.0$. For the muon Monte Carlo, we demand one B hadron with $p_T > 14.0$ GeV/c and $|\eta| < 0.6$. For both cases, we require a second B hadron with $p_T > 7.5$ GeV/c and $|\eta| < 1.0$. The cuts placed on the generator-level Monte Carlo were determined by examining the p_T and η distributions for B hadrons from Monte Carlo events in which two BVTX tags were reconstructed. The p_T and η values were chosen by determining the cuts for which 90% of the B hadrons in the double-tagged Monte Carlo events would pass. We take the $\Delta\phi$ distribution resulting from the event selection above and convolute it with a Gaussian resolution function with a width of 0.1086 radians, characteristic of the $\Delta\phi$ resolution of the BVTX tagging algorithm (see section 7.1.1 below). Finally, we scale the smeared $\Delta\phi$ distribution up by the average number of times each Monte Carlo event was redecayed.

Figure 7.1 shows the reconstructed tag pair $\Delta\phi$ and the smeared generator-level B hadron $\Delta\phi$ distribution for PYTHIA with `PARP(67)` = 4.0 and the different production mechanisms normalized according to best $\Delta\phi$ fit. Figure 7.2 shows the ratio of these two distributions. The curve on this plot shows the fit of the efficiency distribution to an empirical function described in section 7.1.1 below.

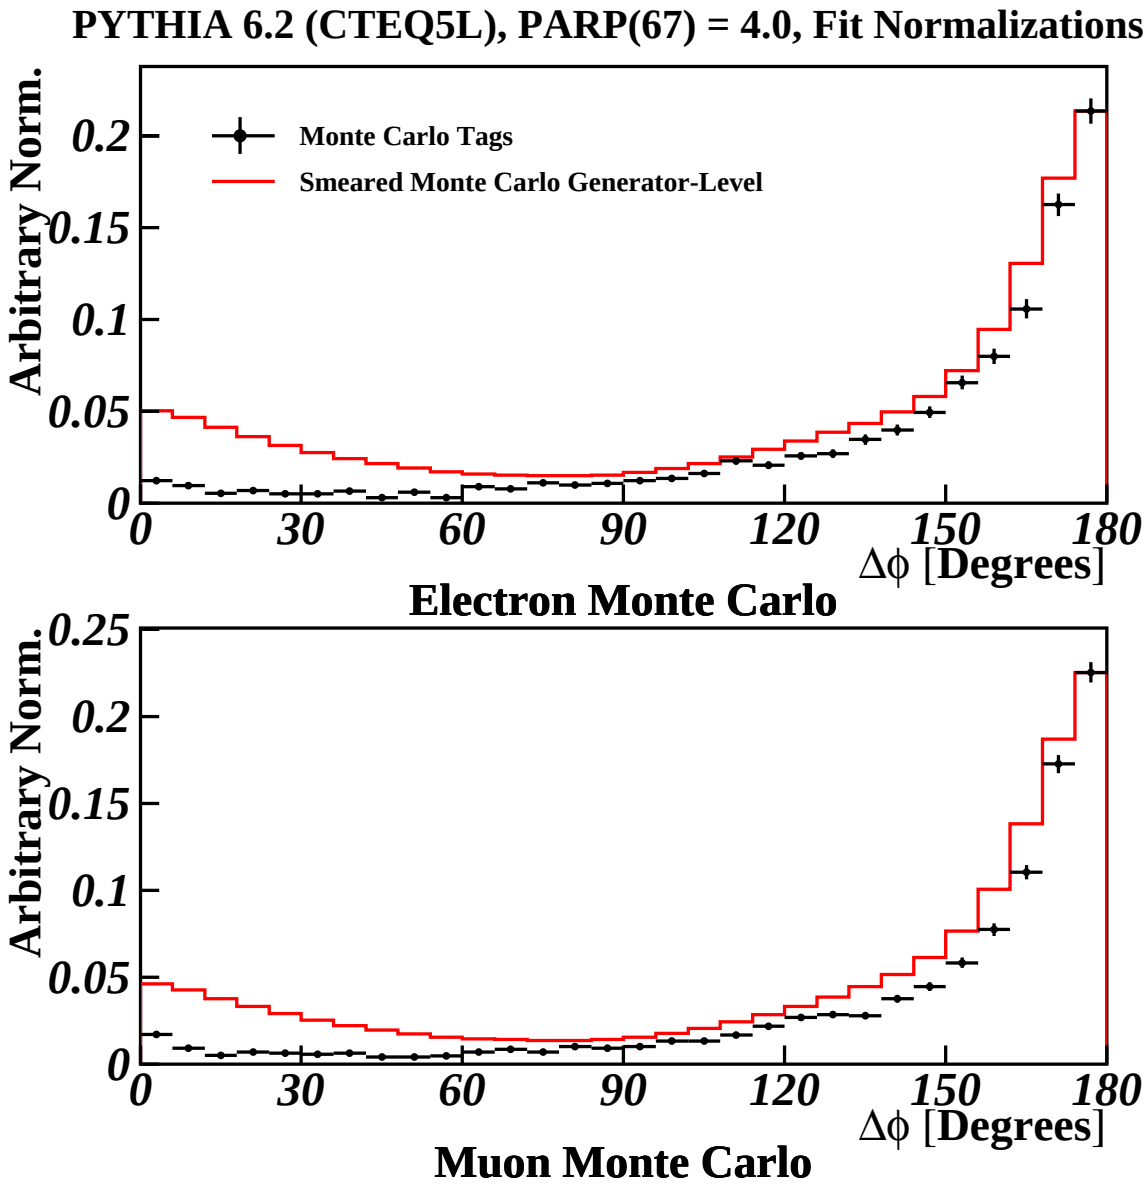

Figure 7.1 A comparison between the shape of the Monte Carlo tag distribution and the smeared, generator level B distribution. The smearing is done with a Gaussian resolution function with $\sigma = 0.1086$ radians or 6.22°. The points with error bars (statistical only) are the Monte Carlo tag distribution while the red line is the smeared, generator level curve. The two curves are normalized so that they have the same value in the last $\Delta\phi$ bin.

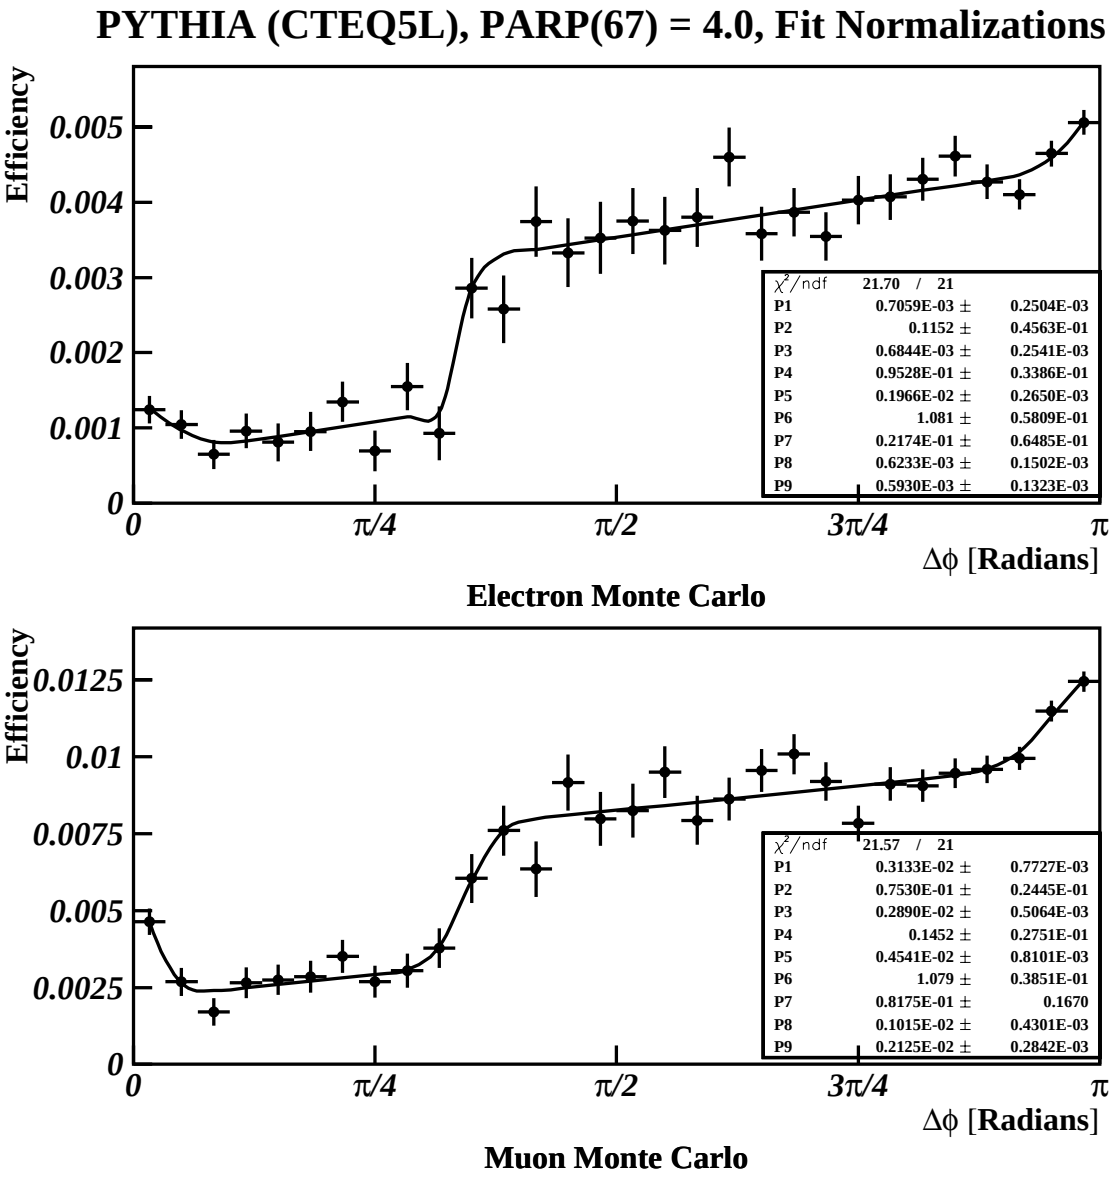

Figure 7.2 The relative efficiency distribution calculated from Monte Carlo. The points are the bin-by-bin ratio of the two curves in Figure 7.1. The curve is a fit to these points using the function described in 0.

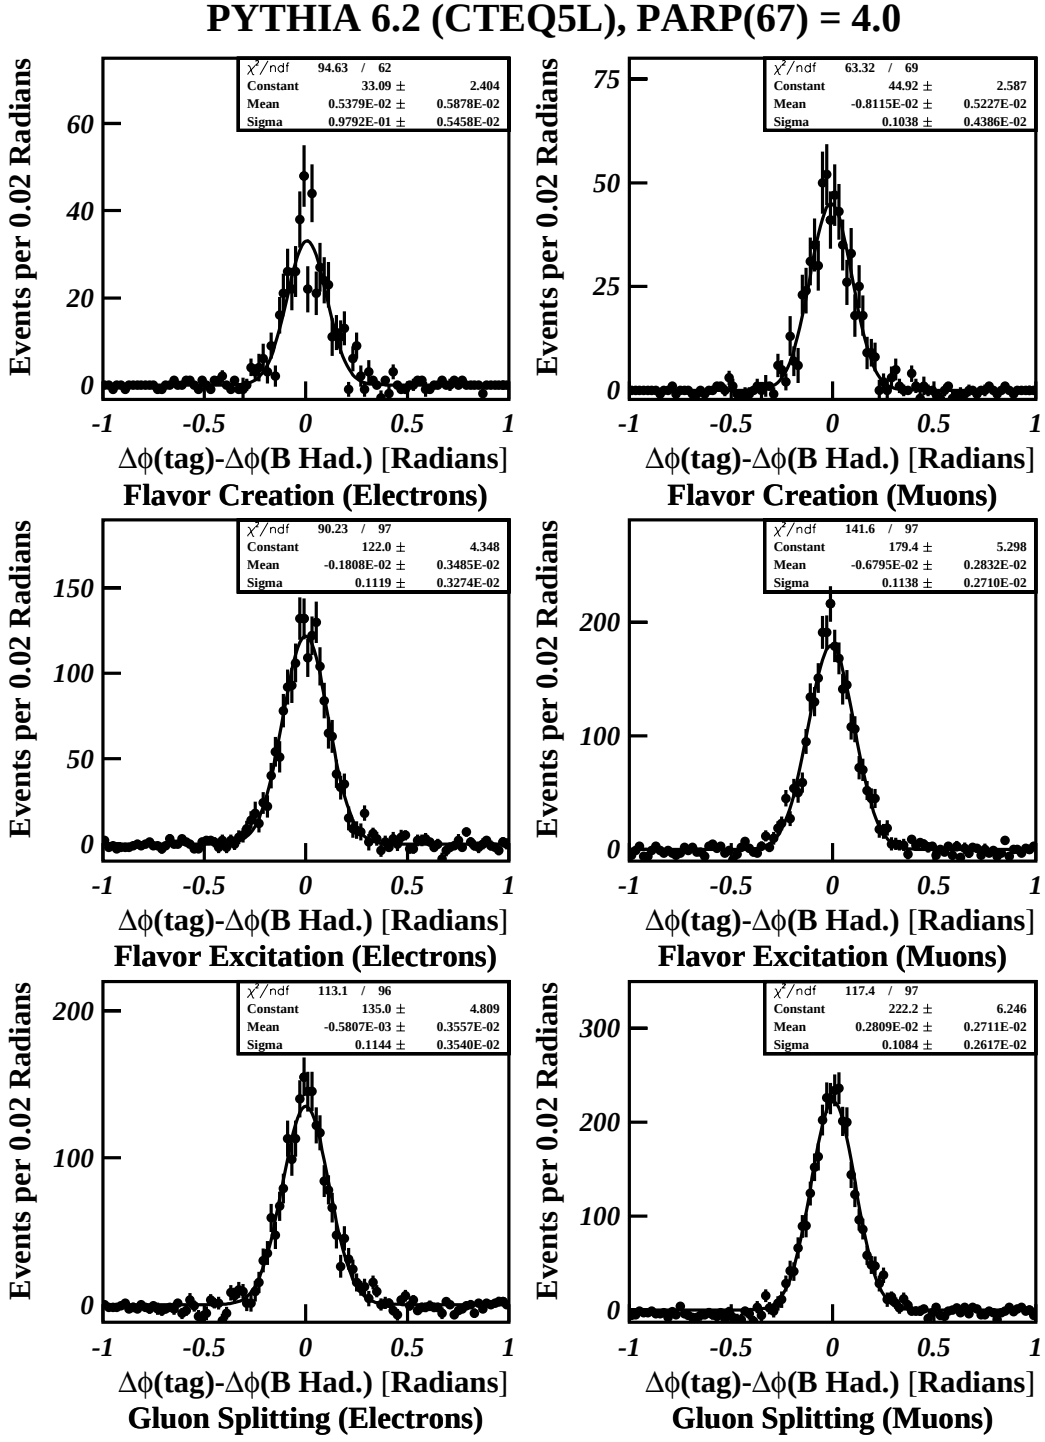

Figure 7.3 shows the matching in Monte Carlo events between the tag Δϕ and the generator-level *B* hadron Δϕ for PYTHIA with `PARP(67)` = 4.0. A Gaussian fit to this distribution gives the Δϕ resolution. The fit excludes *B*-*B* pairs that have Δϕ near 0° or 180° since these events have an asymmetric matching distribution.

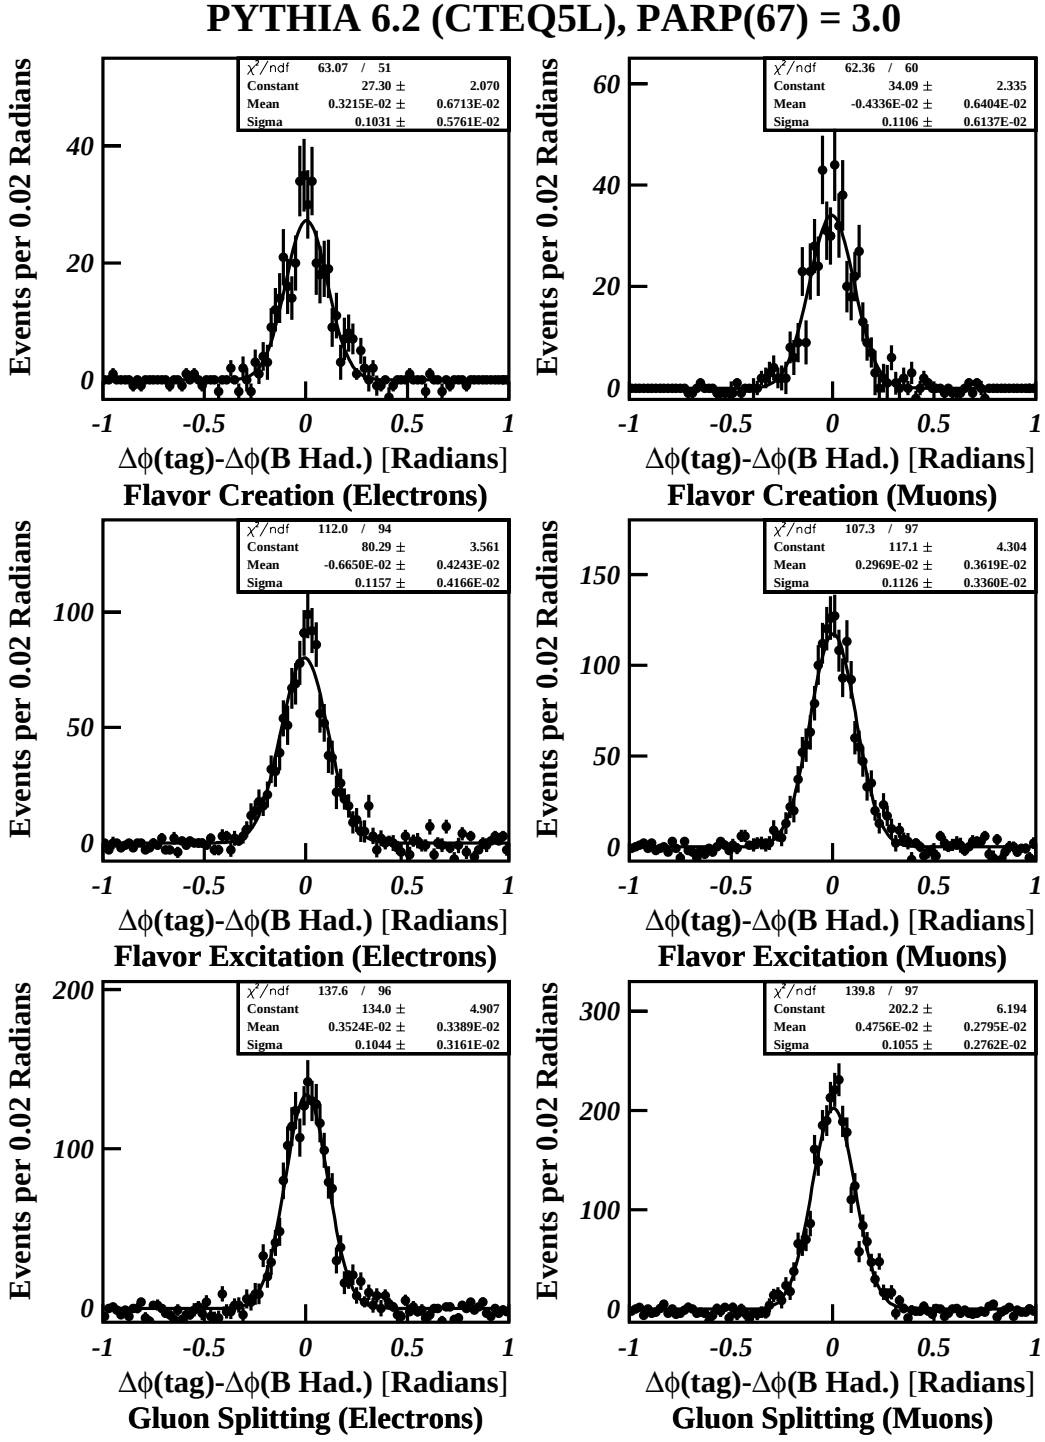

Figure 7.4 shows the matching in Monte Carlo events between the tag Δϕ and the generator-level *B* hadron Δϕ for PYTHIA with `PARP(67)` = 3.0. A Gaussian fit to this distribution gives the Δϕ resolution. The fit excludes *B*-*B* pairs that have Δϕ near 0° or 180° since these events have an asymmetric matching distribution.

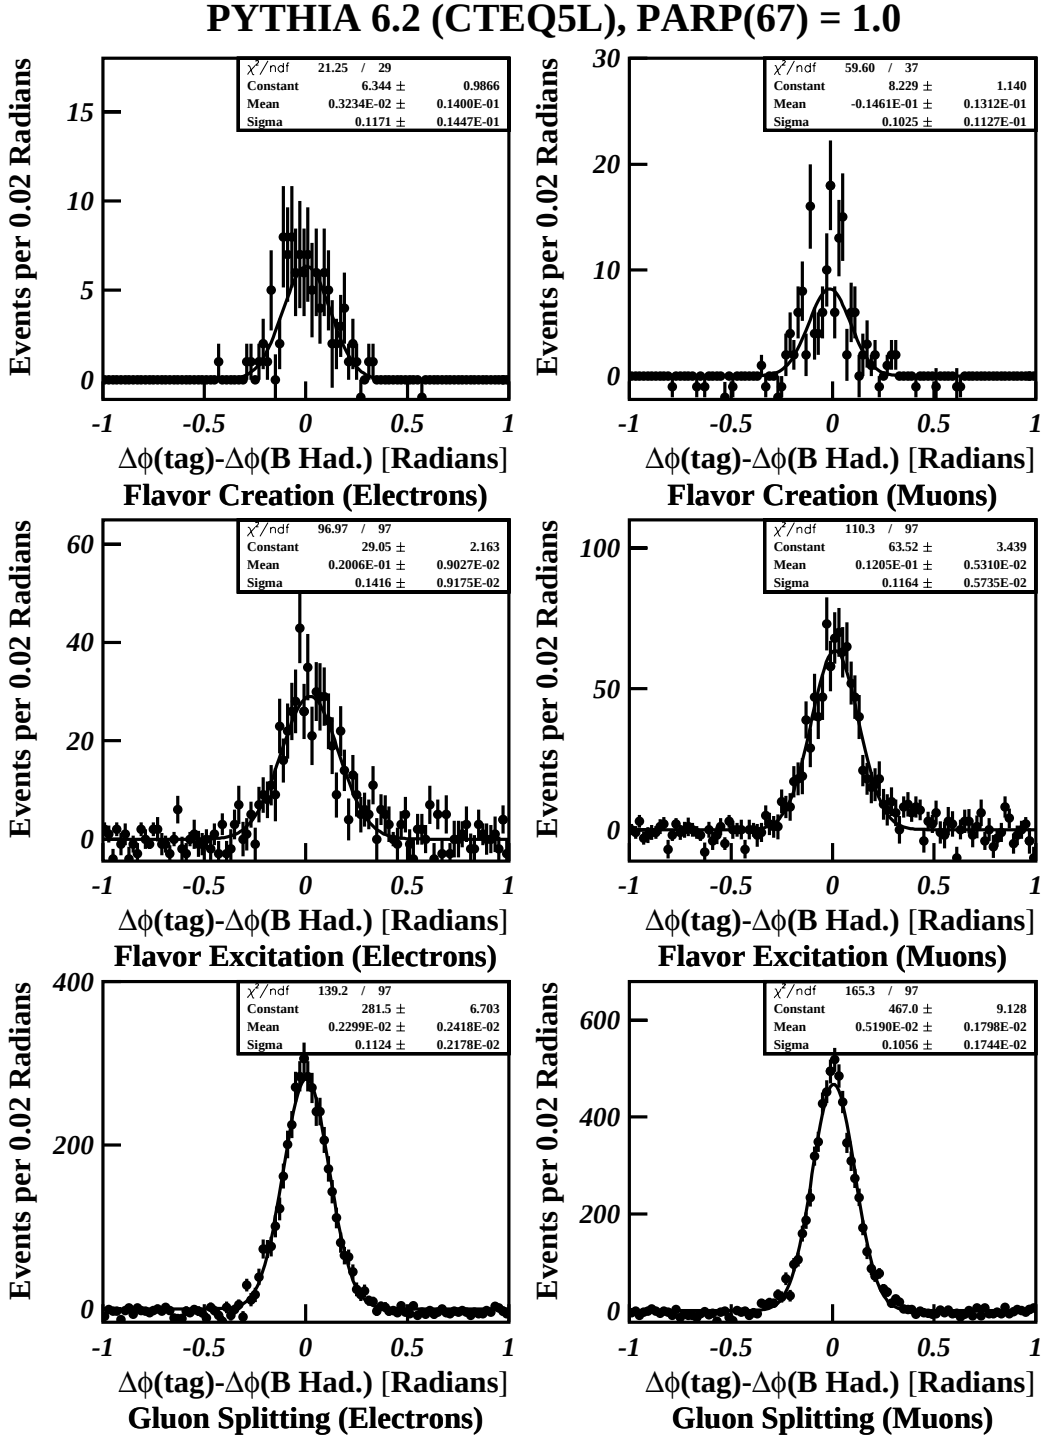

Figure 7.5 shows the matching in Monte Carlo events between the tag Δϕ and the generator-level *B* hadron Δϕ for PYTHIA with `PARP(67)` = 1.0. A Gaussian fit to this distribution gives the Δϕ resolution. The fit excludes *B*-*B* pairs that have Δϕ near 0° or 180° since these events have an asymmetric matching distribution.

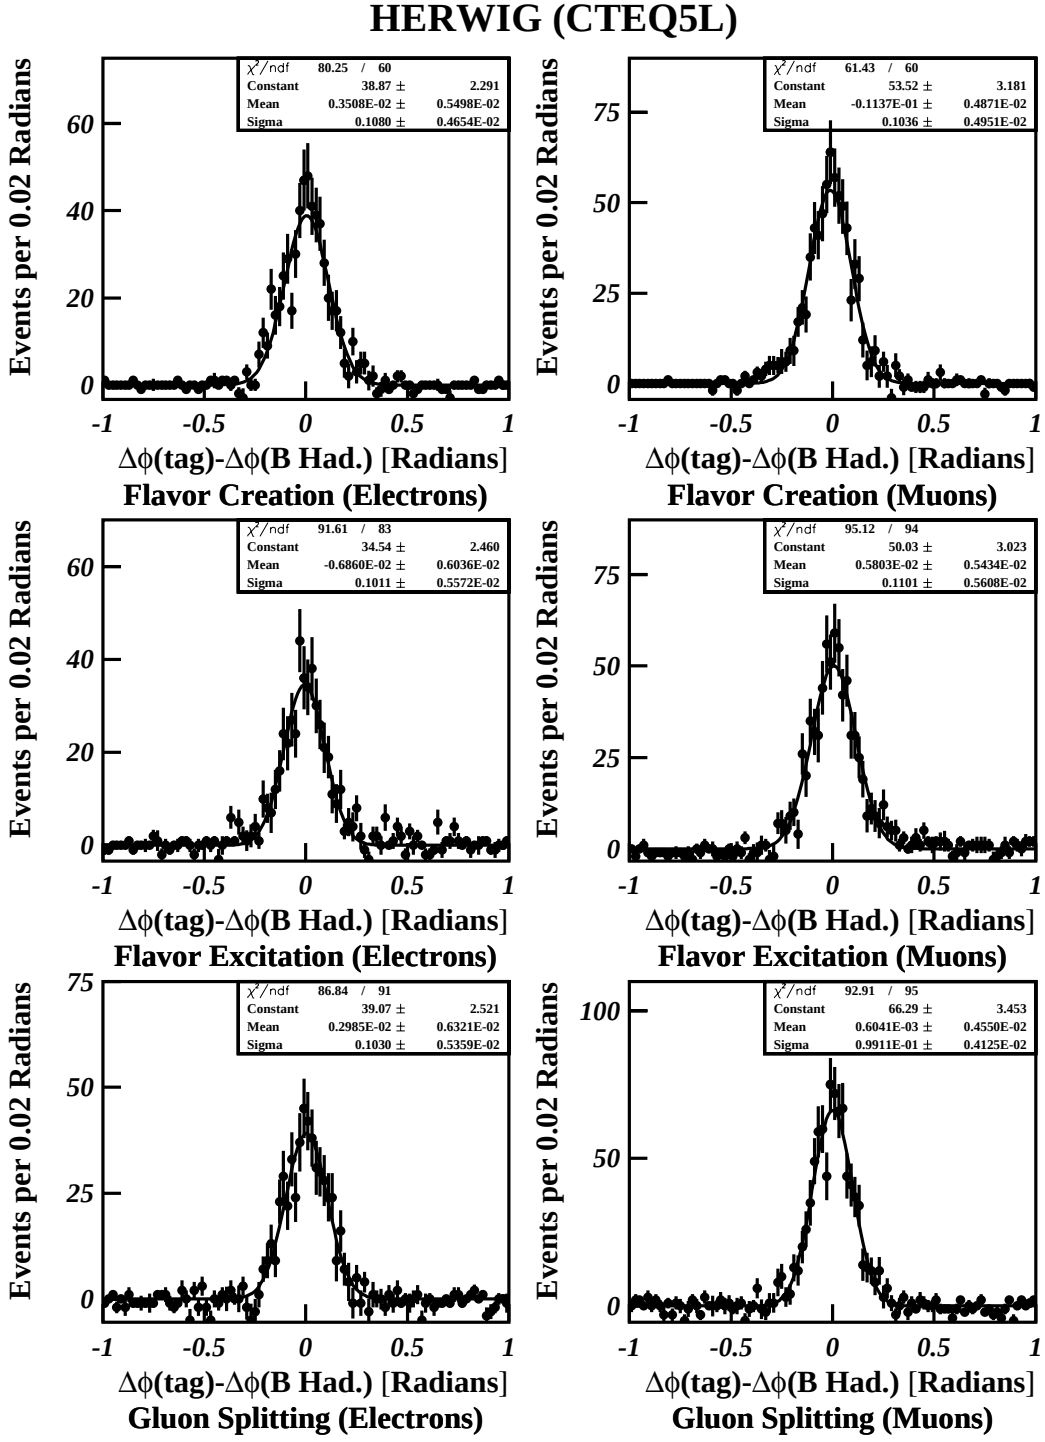

Figure 7.6 shows the matching in Monte Carlo events between the tag $\Delta\phi$ and the generator-level B hadron $\Delta\phi$ for HERWIG. A Gaussian fit to this distribution gives the $\Delta\phi$ resolution. The fit excludes B-B pairs that have $\Delta\phi$ near 0° or 180° since these events have an asymmetric matching distribution.

7.1.1 $\Delta\phi$ Resolution

In order to find the "perfect efficiency" tag-pair $\Delta\phi$ distribution, we need to know the $\Delta\phi$ resolution of the BVTX tagging algorithm. To determine this, we look at Monte Carlo and compare the generator level $\Delta\phi$ of the *B* hadrons in the event to the measured tag $\Delta\phi$. The distribution of the difference between the generator level and tag $\Delta\phi$, with mistags subtracted as in the data, is then fit separately for each production mechanism in each sample. The results of the fits are shown in Figure 7.3 through Figure 7.6. We then use the weighted average of fit results to determine the actual σ to use in the convolution of the generator level $\Delta\phi$ distribution with a Gaussian resolution function. Table 7.1 summarizes the combined fit results for σ. Figure 7.7 demonstrates how the tag $\Delta\phi$ distribution can be duplicated by smearing the generator level $\Delta\phi$ distribution for tagged events.

7.1.2 Parameterizing the Relative Efficiency

In order to minimize the effect of statistical fluctuations in the tagging efficiency determined from Monte Carlo, we fit the tagging efficiency to an empirical function of the following form:

$$\varepsilon(\Delta\phi) = P_1 \exp\left\{-\frac{1}{2}\left(\frac{\Delta\phi^2}{P_2^2}\right)\right\} + P_3 \exp\left\{-\frac{1}{2}\left[\frac{(\Delta\phi-\pi)^2}{P_4^2}\right]\right\} + P_5\,\mathrm{freq}\left(\frac{\Delta\phi - P_6}{P_7}\right) + P_8(\Delta\phi) + P_9, \quad (7.1)$$

where the linear terms represent the general trend in the efficiency curve, the two Gaussian terms at low $\Delta\phi$ and high $\Delta\phi$ model the rise at either end of the $\Delta\phi$ range, and the step at $\Delta\phi = 1$ is modeled by $\mathrm{freq}\left[(\Delta\phi - P_6)/P_7\right]$, where P_6 represents the location of the step, P_7 represents the width of the step region, and *freq*(*x*) is known as the error function:

PYTHIA, `PARP(67)` = 4.0 Average	0.1102 ± 0.0014
PYTHIA, `PARP(67)` = 3.0 Average	0.1082 ± 0.0015
PYTHIA, `PARP(67)` = 1.0 Average	0.1093 ± 0.0013
HERWIG Average	0.1038 ± 0.0020
Combined Average	0.1086 ± 0.0007

Table 7.1 The measured $\Delta\phi$ resolution for each Monte Carlo sample, as well as the combined average. The $\Delta\phi$ resolutions are given in radians.

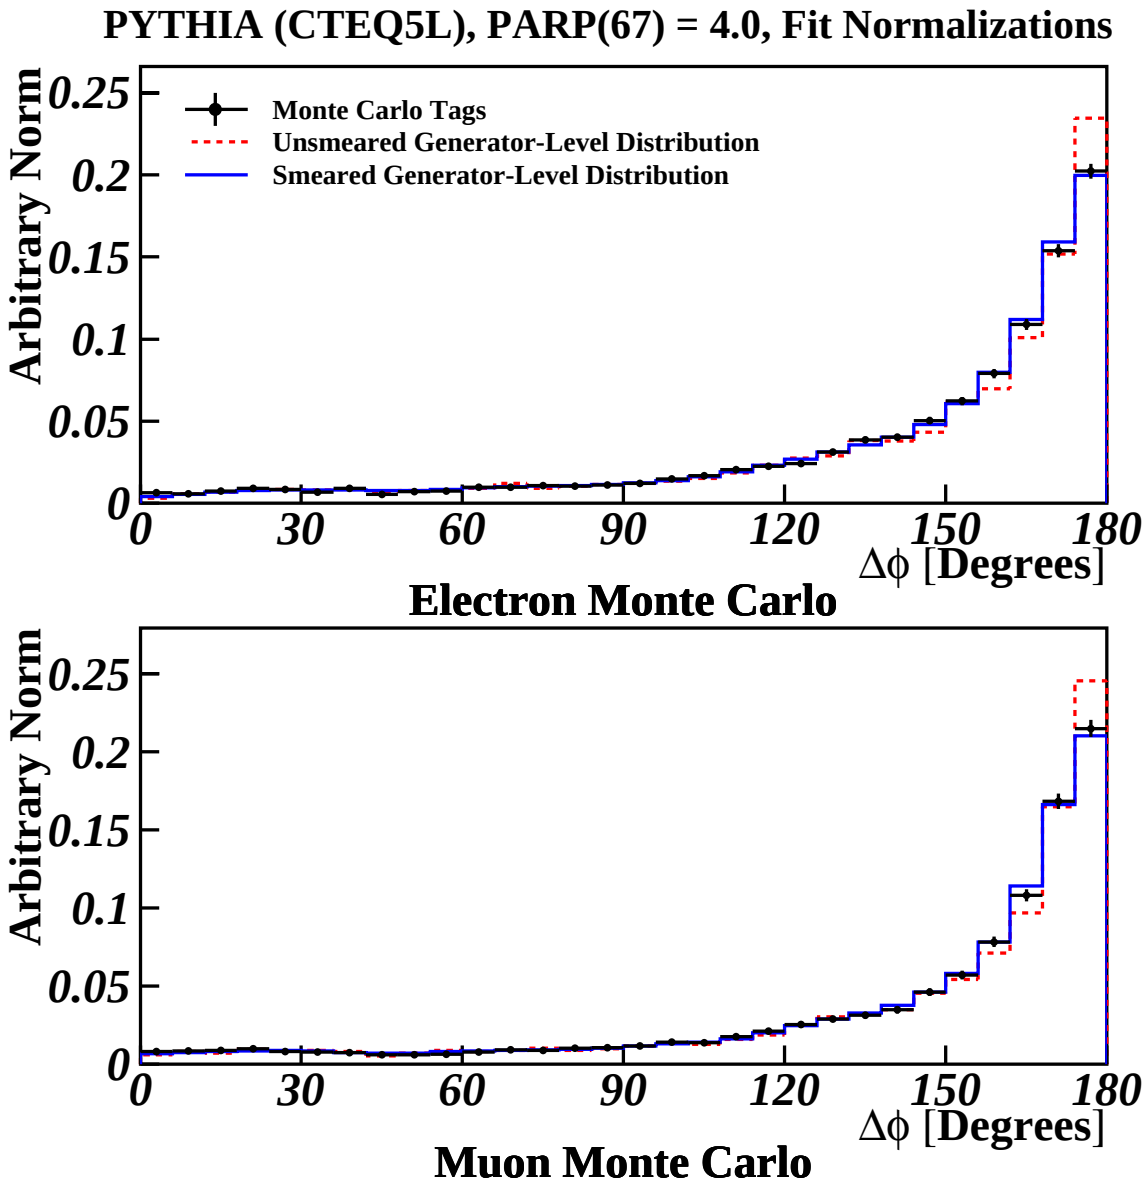

Figure 7.7 demonstrates the effects of tagging resolution on the Δϕ distribution. The points with error bars show the Monte Carlo tag Δϕ distribution. The red line shows the generator-level *B* hadron Δϕ distribution for only the tagged events. The blue line shows that distribution smeared according to the Δϕ resolution measured in Section 7.1.1. The extent to which the blue line matches the points shows how well the smearing procedure replicates the Δϕ of the tagging algorithm.

$$freq(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{1}{2}t^2} \, dt \tag{7.2}$$

The result of the fit for PYTHIA, with `PARP(67)` = 4.0 and the components normalized according to the $\Delta\phi$ fit, is displayed in Figure 7.2. The relative efficiency curve resulting from this fit is shown in Figure 7.8. Since we are only interested in the effect of the efficiency on the shape of the $\Delta\phi$ distribution, and not on its absolute normalization, we have rescaled the curve in Figure 7.8 so that the relative efficiency in the last $\Delta\phi$ bin is defined to be one. Thus this curve shows the effect of the BVTX tagging efficiency for a given bin relative to the last $\Delta\phi$ bin.

7.1.3 Systematic Error

There are two main contributions to the systematic uncertainty associated with the relative tagging efficiency correction: First, the statistical errors on the fit value for the relative efficiency correction factor should be propagated into systematic uncertainties on the corrected $\Delta\phi$ distribution. There is an additional systematic uncertainty that comes from the model used to calculate the relative efficiency correction. The PYTHIA Monte Carlo sample, with `PARP(67)` = 4.0 and with the normalization of the different production mechanisms taken from the fit to the $\Delta\phi$ distribution in the data, is used as our baseline for the relative efficiency correction. However, other models, like the lower ISR PYTHIA sample or HERWIG also match the data to varying degrees and so could also have been used. To account for this ambiguity, we compare the relative efficiency corrections from other Monte Carlo models to our baseline model. Figure 7.9 shows the worst comparison. In the worst case, the difference for the bin-by-bin relative efficiency correction factor is approximately equal in magnitude to the statistical error from the fit. Therefore, to account to modeling uncertainties in the relative efficiency correction, we increase the systematic error associated with the correction by a factor of $\sqrt{2}$.

7.2 Backgrounds

Although this analysis uses cuts to reduce the backgrounds in the double-tagged sample to negligible levels, it is still possible that the cuts leave a small residual background contribution that distorts the shape of the $\Delta\phi$ spectrum in a noticeable way. In addition, uncertainties involved in the techniques used to limit and account for background contributions lead to an

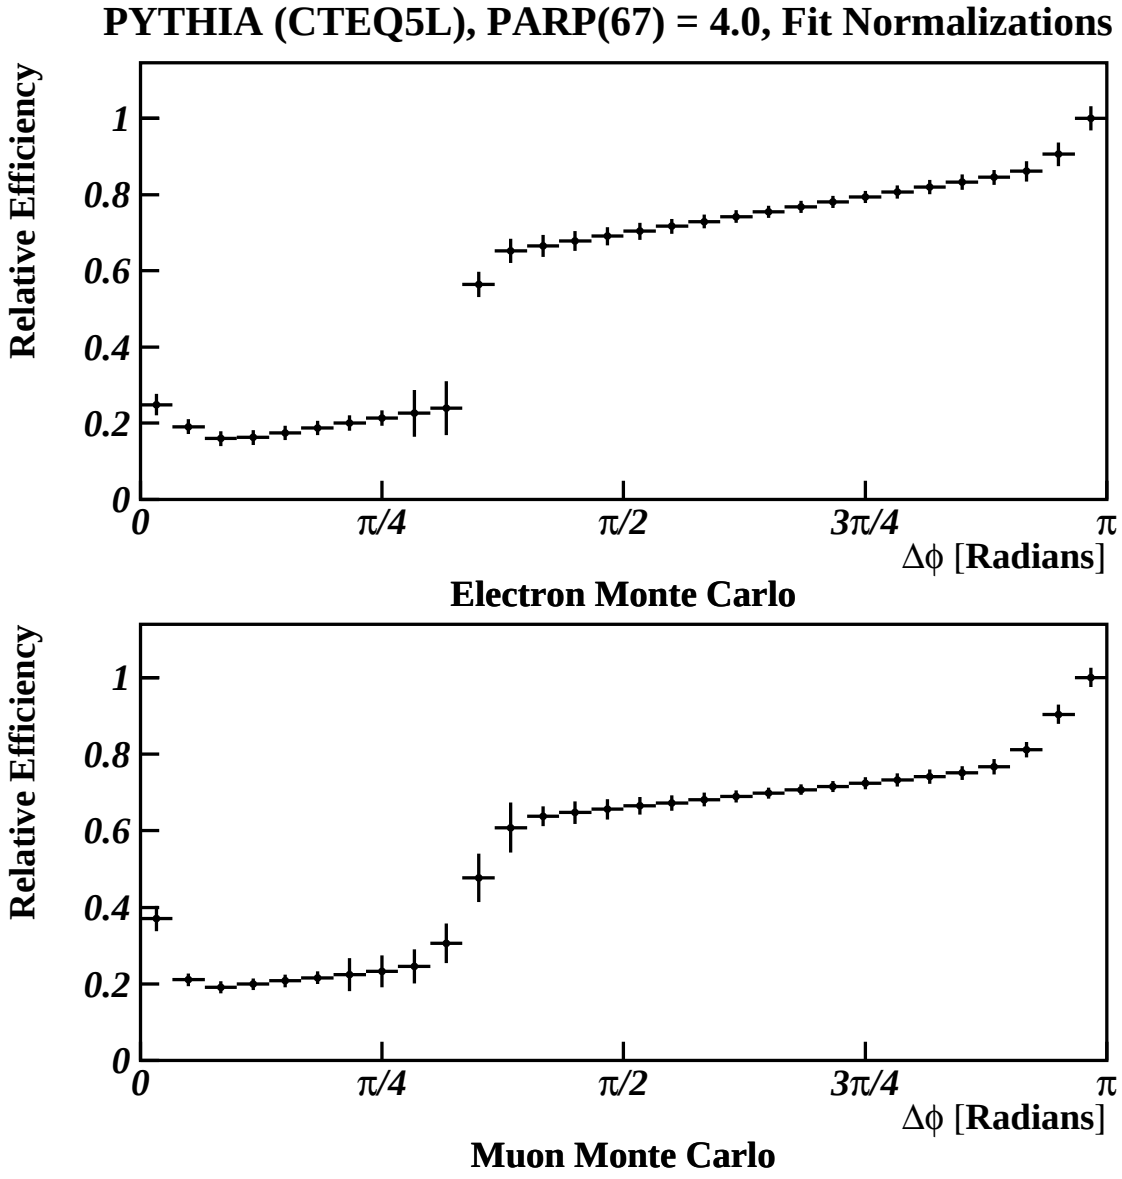

Figure 7.8 The bin-by-bin values for the relative efficiency returned by the fit. The curves have been normalized so that the last $\Delta\phi$ bin has a value of one by definition. The error bars on these curves indicate the statistical error on the bin values returned from the fit. The statistical errors for the fit are correlated from bin to bin.

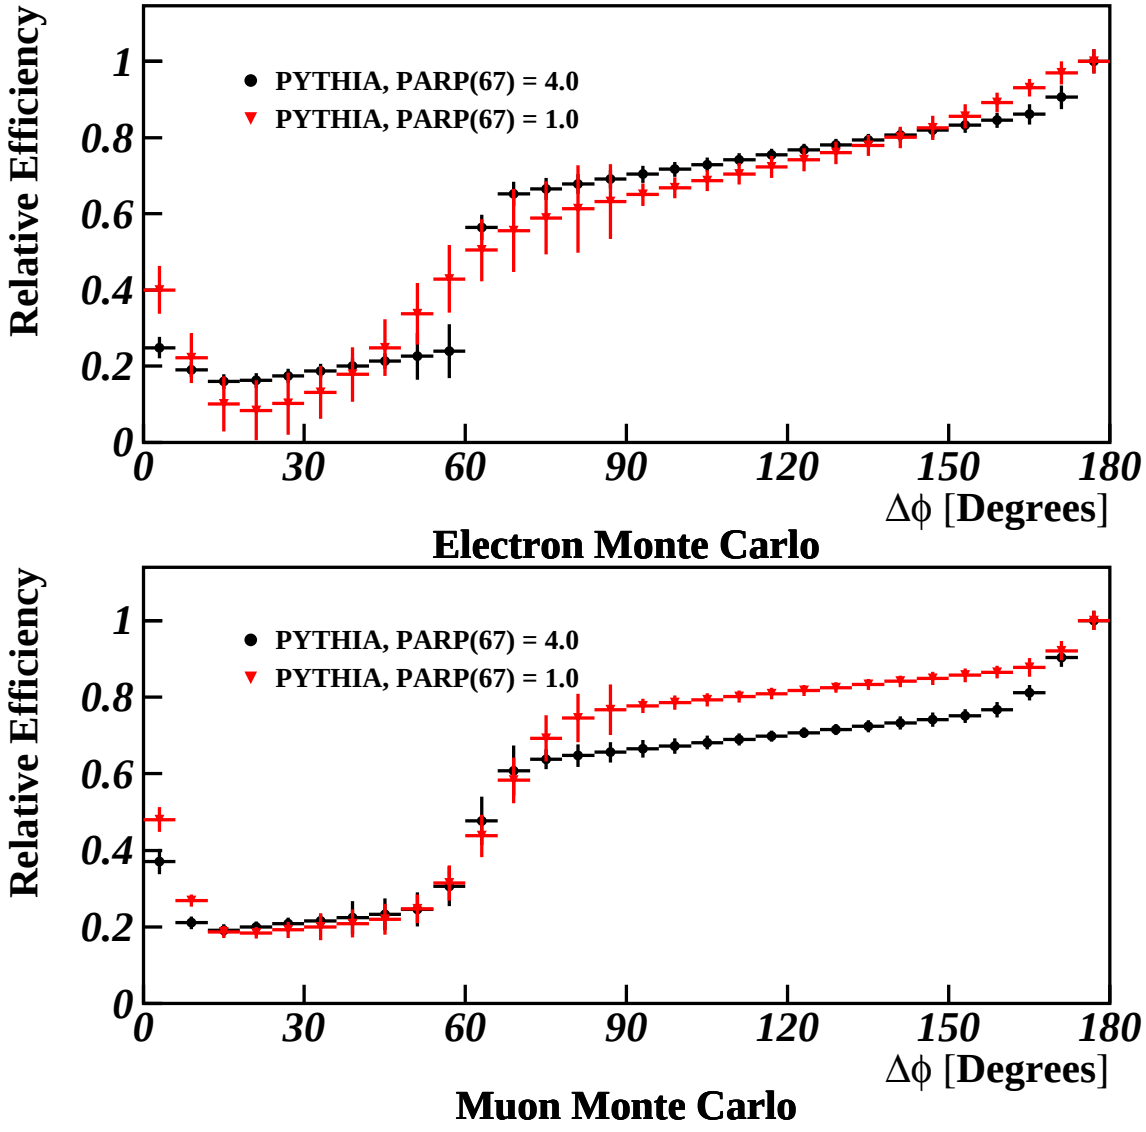

Figure 7.9 compares the relative efficiency curve derived from PYTHIA with `PARP(67)` = 4.0 to the curve derived from PYTHIA with `PARP(67)` = 1.0. The normalizations are set so that the last $\Delta\phi$ bin for both curves is set to one by definition. The error bars indicate the statistical error on the relative efficiency for a given bin coming from the statistical errors on the fit.

uncertainty in the final B hadron $\Delta\phi$ extracted from the data. In the sections below, the possible residual contributions from various background sources are considered and systematic uncertainties associated with these backgrounds are evaluated.

7.2.1 Mistag Subtraction

The mistag subtraction scheme used for this analysis relies on the assumption that 100% of legitimate tags and 50% of mistags have positive L_{xy}. The true fraction may be somewhat different. For example, if most of the events contain at least one B hadron, then the L_{xy} distribution of mistags may be biased towards positive values by the presence of actual displaced tracks in the events. Furthermore, the bias in L_{xy} may depend on the topology of the event. To investigate any possible bias in the L_{xy} distribution of mistags, we examined Monte Carlo events containing mistags identified by matching tracking information to Monte Carlo truth information.

From Monte Carlo sample to Monte Carlo sample, the fraction of legitimate secondary vertex tags that have positive L_{xy} varies from 0.97 to 1.0. For mistags, the positive L_{xy} fraction varies from 0.45 to 0.55. To estimate the possible effect of using the wrong fractions when performing mistag subtraction, we redo the mistag subtraction in the data using different assumptions about the positive L_{xy} fraction for good tags and mistags. The mistag subtraction formula from Section 4.4.1, generalized for an arbitrary fraction p of good tags with $L_{xy} > 0$ and an arbitrary fraction q of mistags with positive L_{xy} is given by

$$N_{GG} = \frac{(q-1)^2}{(q-p)^2} N_I + \frac{q(q-1)}{(q-p)^2} N_{II} + \frac{q^2}{(q-p)^2} N_{III} + \frac{q(q-1)}{(q-p)^2} N_{IV} \qquad (7.3)$$

The change in the total number of events for several different values for p and q is shown in Table 7.2. Changing the positive L_{xy} fractions from mistag subtraction affects both the normalization and the shape of the $\Delta\phi$ distribution. However, for this analysis, we are only concerned about the shape. Therefore, before we compare the shape of the $\Delta\phi$ distribution using the standard mistag subtraction scheme to the shape obtained using alternative values for the positive L_{xy} fractions, we normalize the distributions to unit area. Changing the good tag positive L_{xy} fraction from 1.0 to 0.97 has a negligible effect on the $\Delta\phi$ shape. On the other hand, varying

Good Tag Positive L_{xy} Fraction	Mistag Positive L_{xy} Fraction	Electron Data	Muon Data
1.0 (Default)	0.5 (Default)	10,097.0	7,092.0
1.0	0.45	11,006.4	7,831.9
1.0	0.55	9,086.6	6,268.4
0.97	0.5	11,427.1	8,026.2
0.97	0.45	12,312.9	8,761.6
0.97	0.55	10,431.0	7,195.9

Table 7.2 The number of good tag pairs in the data estimated using different positive L_{xy} fractions for the mistag subtraction.

the mistag positive L_{xy} fraction from 0.45 to 0.55 makes a noticeable difference, as shown in Figure 7.10.

To estimate the systematic error from mistag subtraction, we take the $\Delta\phi$ distributions calculated using the various mistag subtraction schemes and fit them to the functional form given below:

$$f(\Delta\phi) = P_1 \exp\left[-\frac{1}{2}\left(\frac{\Delta\phi}{P_2}\right)^2\right] + P_3 \exp\left[P_4(\Delta\phi - \pi)^2 + P_5(\Delta\phi - \pi)\right] + P_6\Delta\phi + P_7 \qquad (7.4)$$

We then calculate the maximum deviation between the result from the default mistag subtraction scheme and the results obtained from varying the positive L_{xy} fractions. This maximum deviation is assigned as the systematic error on the $\Delta\phi$ shape from mistag subtraction.

7.2.2 Prompt Charm

In Section 4.4.3, we used the characteristics of the data to establish that after requiring two secondary vertex tags, the contribution from prompt charm is reasonably small and doesn't have a large effect on the shape of the tag $\Delta\phi$ distribution. In this section, we will use Monte Carlo and data to estimate the approximate size of the residual prompt charm contribution as well as to assess its effect on the systematic errors.

To perform studies of the prompt charm contribution, we generated samples of $c\bar{c}$ Monte Carlo. Just as in $b\bar{b}$ production, flavor creation, flavor excitation, and gluon splitting diagrams

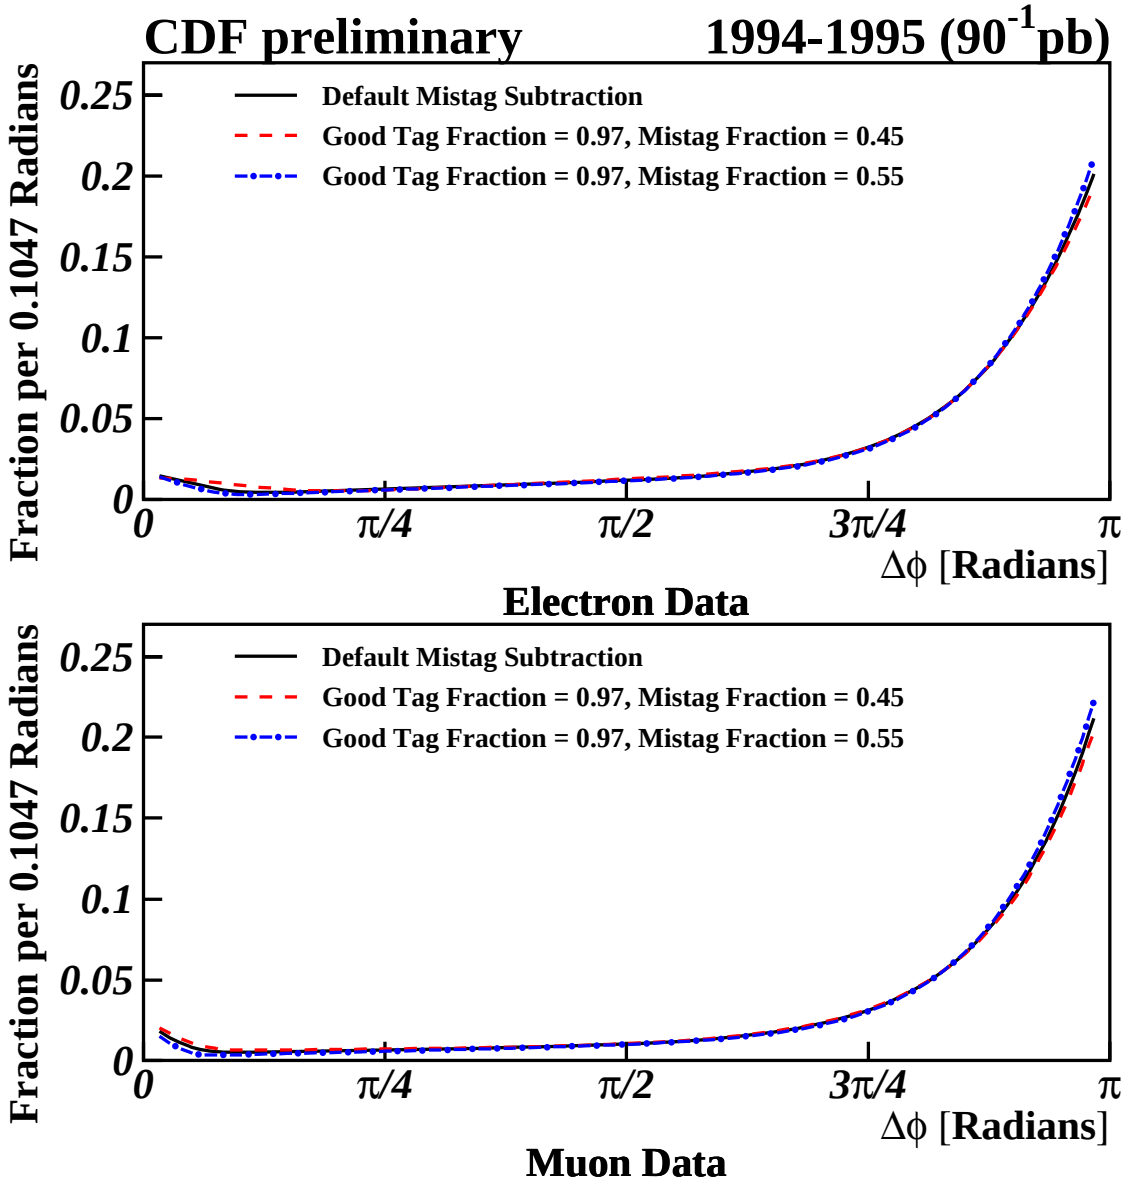

Figure 7.10 illustrates the change in shape for the $\Delta\phi$ distribution as two different extreme variations in positive L_{xy} fractions are used for mistag subtraction. The curves show the fits to the mistag subtracted data using the functional form given in Section 7.2.1.

all provide significant contributions to $c\bar{c}$ production. Unfortunately, it would be too lengthy an undertaking to perform a full simulation—including detector simulation, reconstruction code, and analysis code—of all three contributions. Instead, we generated two different $c\bar{c}$ samples: one larger flavor creation $c\bar{c}$ sample that we passed through the detector simulation and full analysis chain, and a smaller generator-level sample involving all three production mechanisms. Both samples were generated using PYTHIA with `PARP(67)` = 4.0 and the same tuning as the $b\bar{b}$ sample with the same amount of ISR, as presented in [66]. The larger, flavor creation sample consists of 18 million $c\bar{c}$ events generated with p_T(min) = 8 GeV/c. This sample was passed through the QFL' detector simulation, as well as the same reconstruction and analysis code as the data, in order to produce a sample of double-tagged events from $c\bar{c}$ production. The smaller, generator-level only sample consists of 2 million flavor creation events generated with p_T(min) = 0 GeV/c, 4 million flavor excitation events generated with p_T(min) = 3 GeV/c, and 100 million generic QCD events generated with p_T(min) = 3 GeV/c, of which approximately 2.6 million contained $c\bar{c}$ pairs from gluon splitting. These events are then combined using PYTHIA's predictions for the cross sections of each of the three processes. To create a Monte Carlo sample to model multiple heavy flavor production, we selected those events from our full PYTHIA $b\bar{b}$ Monte Carlo in which a $c\bar{c}$ pair was also produced.

7.2.2.1 Relative Rates from Monte Carlo

One approach to estimating the contribution from charm to our double tag sample is to compare the relative rates for double tagging events from the $c\bar{c}$ and $b\bar{b} + c\bar{c}$ Monte Carlo to the rate from $b\bar{b}$ Monte Carlo. To do this, we count the number of double-tagged events reconstructed in each Monte Carlo sample and adjust for the equivalent luminosity of each sample. This relative rate takes into account both the relative tagging efficiency for each sample as well as the relative production rates. However, because we only fully simulate the flavor creation component of $c\bar{c}$ production, we underestimate the relative tagging rate for this component. To correct for the missing contributions from flavor excitation and gluon splitting, we use the generator-level Monte Carlo to estimate the cross section for producing two D hadrons with $|\eta| < 1$ and enough p_T to be tagged, compared to the cross section from flavor creation alone. From this, we determine that the tagging rate estimated from the flavor creation $c\bar{c}$ sample should be scaled up by a factor of 4.3 to account for flavor excitation and gluon splitting. Note that this

	Electrons		Muons	
	Relative Rate	**Fraction of Events**	**Relative Rate**	**Fraction of Events**
$c\bar{c}$ **(Flavor Creation only)**	0.0071	0.7%	0.015	1.5%
$c\bar{c}$ **(All Contributions)**	0.031	2.9%	0.065	6.0%
$b\bar{b} + c\bar{c}$	0.019	1.8%	0.019	1.8%

Table 7.3 The prediction of the tagging rates for $c\bar{c}$ and $b\bar{b} + c\bar{c}$ events relative to the tagging rate for $b\bar{b}$ events. The first line in the table shows the $c\bar{c}$ rate before we scale it up to account for the $c\bar{c}$ production mechanism not generated. The "fraction" columns indicate what fraction of double-tagged events come from each source assuming that only $b\bar{b}$, $c\bar{c}$, and $b\bar{b} + c\bar{c}$ production contribute to the double-tag sample.

estimate does not take into account the difference in relative tagging efficiency between flavor creation and the other production mechanism, so it is likely an over estimate.

Table 7.3 shows the tagging rates for $c\bar{c}$ and $b\bar{b} + c\bar{c}$ Monte Carlo relative to the tagging rate for $b\bar{b}$ Monte Carlo for both electrons and muons. If we assume that, after mistag subtraction, all tag pairs in the data come from one of these three sources, then the relative tagging rates can be used to calculate the fraction of the total tags from each process. Using this procedure, we estimate that 2.9% of double-tagged events in the electron sample and 6.0% of double-tagged events in the muon sample come from $c\bar{c}$ production. In both samples, we estimate 1.8% of the events come from $b\bar{b} + c\bar{c}$ production.

7.2.2.2 Mass Template Fits

An alternative approach to estimating the contribution from prompt charm to the double-tag sample is to examine the tag mass distribution in the data compared to the tag mass distributions predicted by the Monte Carlo samples. Using our fully simulated and reconstructed $b\bar{b}$ and $c\bar{c}$ Monte Carlo samples, we create mass templates for *B* tags and prompt *D* tags. We then fit these templates to the tag mass distributions from double-tagged events in the electron and muon data. The fit results are shown in Figure 7.11. From these fits we obtain prompt charm fractions of 7.1% for the electrons and 13.3% for muons. However, because these numbers represent the

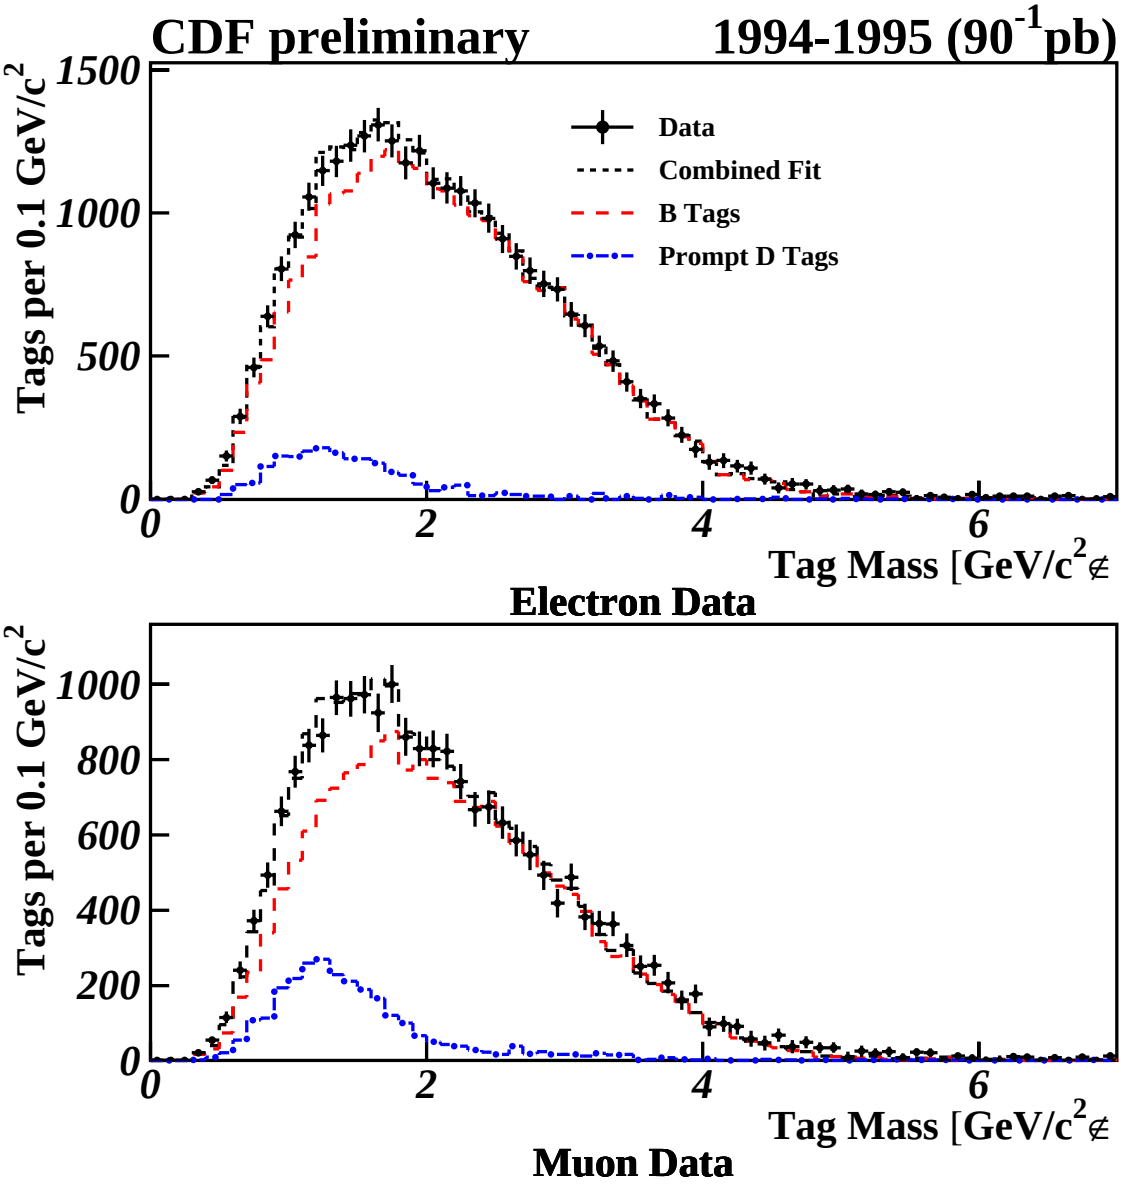

Figure 7.11 shows the fits of the tag mass templates from *B* tags and prompt *D* tags to the data. The points with error bars are the data. The dashed red line is the *B* tag distribution. The dashed-dotted blue line is the prompt *D* tag distribution. The dashed black line shows the sum of the two contributions.

fraction of individual tags from prompt charm, the do not tell us directly what fraction of tag pairs contain at least one tag from prompt charm. To determine that number, we need to know the relative proportions of tag pairs from $c\bar{c}$ production to tag pairs from $b\bar{b} + c\bar{c}$ production. Table 7.4 shows a comparison between the prompt charm fraction obtained from the fit to the fraction predicted by the Monte Carlo tagging rate study above. In general, the fits show approximately 1.5-2 times as much contribution from prompt charm to the double-tag sample. This uncertainty in the small fraction that prompt charm tags contribute the measured $\Delta\phi$ distribution is reflected in the systematic uncertainty on the prompt charm contribution.

7.2.2.3 Estimate of the $\Delta\phi$ Shape

Although the contribution from prompt charm appears small and does not seem to have a drastic impact on the shape of the $\Delta\phi$ distribution in the data, we would still like to estimate the contribution from prompt charm to each bin in the $\Delta\phi$ histogram. To do this, we need a prediction of the $\Delta\phi$ shape of tag pairs from $c\bar{c}$ events as well as from $b\bar{b} + c\bar{c}$ events. The $\Delta\phi$ shape for $b\bar{b} + c\bar{c}$ events can be gotten directly from the $b\bar{b}$ Monte Carlo by selecting those events in which one tag matches a *B* hadron while the other tag matches a prompt *D*. However, because we do not have a fully simulated sample of $c\bar{c}$ Monte Carlo that includes contributions from all three production mechanisms, we cannot take the $c\bar{c}$ $\Delta\phi$ shape directly from the Monte Carlo. Instead, we estimate the shape of this contribution as follows. We begin by constructing the generator-level $\Delta\phi$ spectrum for *D* hadrons. We then apply the relative efficiency correction derived from the $b\bar{b}$ Monte Carlo in Section 7.1 to estimate the effect that the detector would have on this distribution. Finally, as noted in Section 4.4.3, the invariant mass cut on the tag

	Fit Results	**Monte Carlo Prediction**	**Ratio**
Electrons	7.1%	3.8%	1.87
Muons	13.3%	6.9%	1.92

Table 7.4 shows a comparison between the mass template fits for the prompt charm contribution and the prediction from the Monte Carlo study of relative tagging rates. The "ratio" column gives the ratio between the two estimates of prompt charm contamination. The fits predict a factor of two more contribution from prompt charm tags than does the Monte Carlo study.

pairs to reduce the contribution from sequentials affects the $\Delta\phi$ spectrum for tags differently depending on the individual tag masses. Figure 7.12 shows a comparison of the effect of this cut on $b\bar{b}$ compared to $c\bar{c}$ events. To account for this, we apply an addition correction factor to the $c\bar{c}$ Monte Carlo $\Delta\phi$ distribution.

Finally, to estimate the contribution from prompt charm to each bin, we need an appropriate normalization of the $c\bar{c}$ and $b\bar{b} + c\bar{c}$ $\Delta\phi$ shapes. The studies detailed above and in Section 4.4.3 indicate that the prompt charm fraction in the data is anywhere from negligible to 7–13%. As an estimate of the actual normalization, we choose the lower values predicted by the Monte Carlo studies in Section 7.2.2.1—namely 2.9% (6.0%) of the tag pairs coming from $c\bar{c}$ production for electron (muon) data, and 1.8% of the tag pairs coming from $b\bar{b} + c\bar{c}$ production. The estimated contributions are shown in Figure 7.13. To account for this fraction in the data, we then subtract off the predicted contribution from prompt charm in each bin. We set the systematic error from this procedure equal to the amount subtracted off for the correction. Thus, the systematic error from this contribution covers the range of possibilities from having no prompt charm contribution, as suggested by comparing the $\Delta\phi$ spectrum with different tag mass cuts, to having twice as much prompt charm contribution, which roughly matches the results from the mass template fits.

7.2.3 Sequentials

To estimate the size of the contribution from residual sequentials, we turn to the Monte Carlo. To study sequential tagging, we select events from the Monte Carlo in which a single *B* hadron contributes tracks to two separate secondary vertex tags. To increase the statistics of our Monte Carlo sample, we combine sequential events from all Monte Carlo samples—both PYTHIA and HERWIG. This is not such a bad approximation in that the dominant issue for sequential tagging is the *B* decay model, which was the same for all Monte Carlo samples, and otherwise, the Monte Carlo samples share generally the same features.

From this Monte Carlo study, we determine that after mistag subtraction, 25.9% of the tags removed by the 6 GeV/c^2 mass cut were from sequential tag pairs. Furthermore, the 6 GeV/c^2 cut has an efficiency of approximately 97.6%, meaning that for every 100 sequential tags removed by the invariant mass cut, about 2.4 will remain. In the data, after mistag subtraction, the 6 GeV/c^2 invariant mass cut removes 471 tags from the electron sample and 598 tags from the muon sample. Using the numbers derived from the Monte Carlo above, this means that of

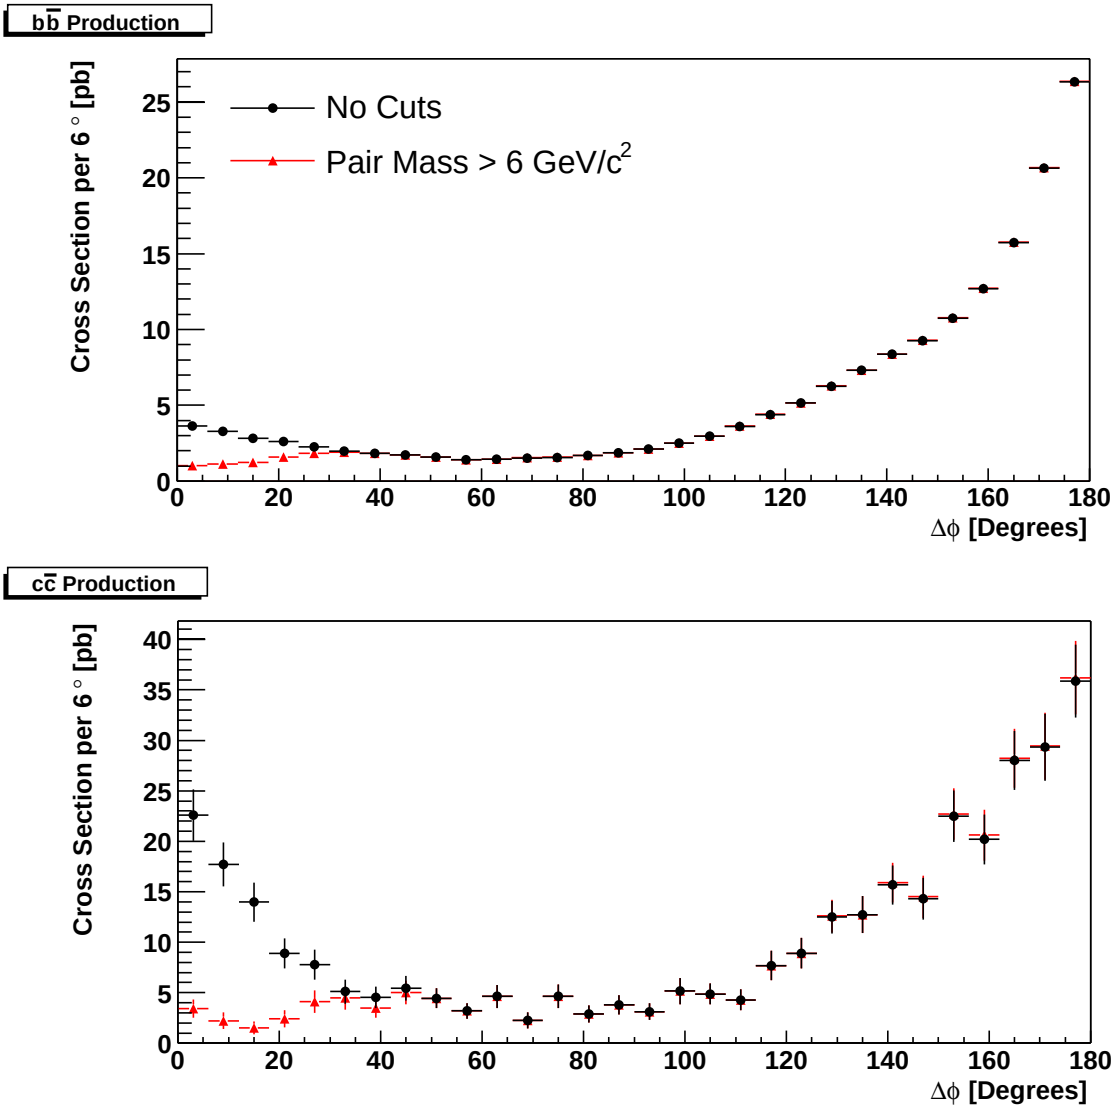

Figure 7.12 estimates the effect that the invariant mass cut on tag pairs has on the $\Delta\phi$ distribution for tags from $b\bar{b}$ and $c\bar{c}$ production. The distributions are taken from the generator-level PYTHIA Monte Carlo with `PARP(67)` = 4.0. The black points show the generator-level distribution for *B* or *D* hadrons with no cuts. The red points show the distributions with "tag pair" invariant mass cuts. To simulate the effect of reconstructing only a portion of the *B* or *D* mass and momentum, the generator-level hadron mass and momentum were scaled back. For the *B* hadrons, we assume that the BVTX tagging algorithm reconstructs 45% (35%) of the *B* mass for tags and 65% (55%) of the *B* momentum containing (not containing) leptons. For *D* hadrons, the fractions are 80% (75%) for both *D* mass and momentum for tags containing (not containing) a lepton. These fractions were derived from studying the Monte Carlo that was passed through the detector simulation, analysis code, and tagging algorithm.

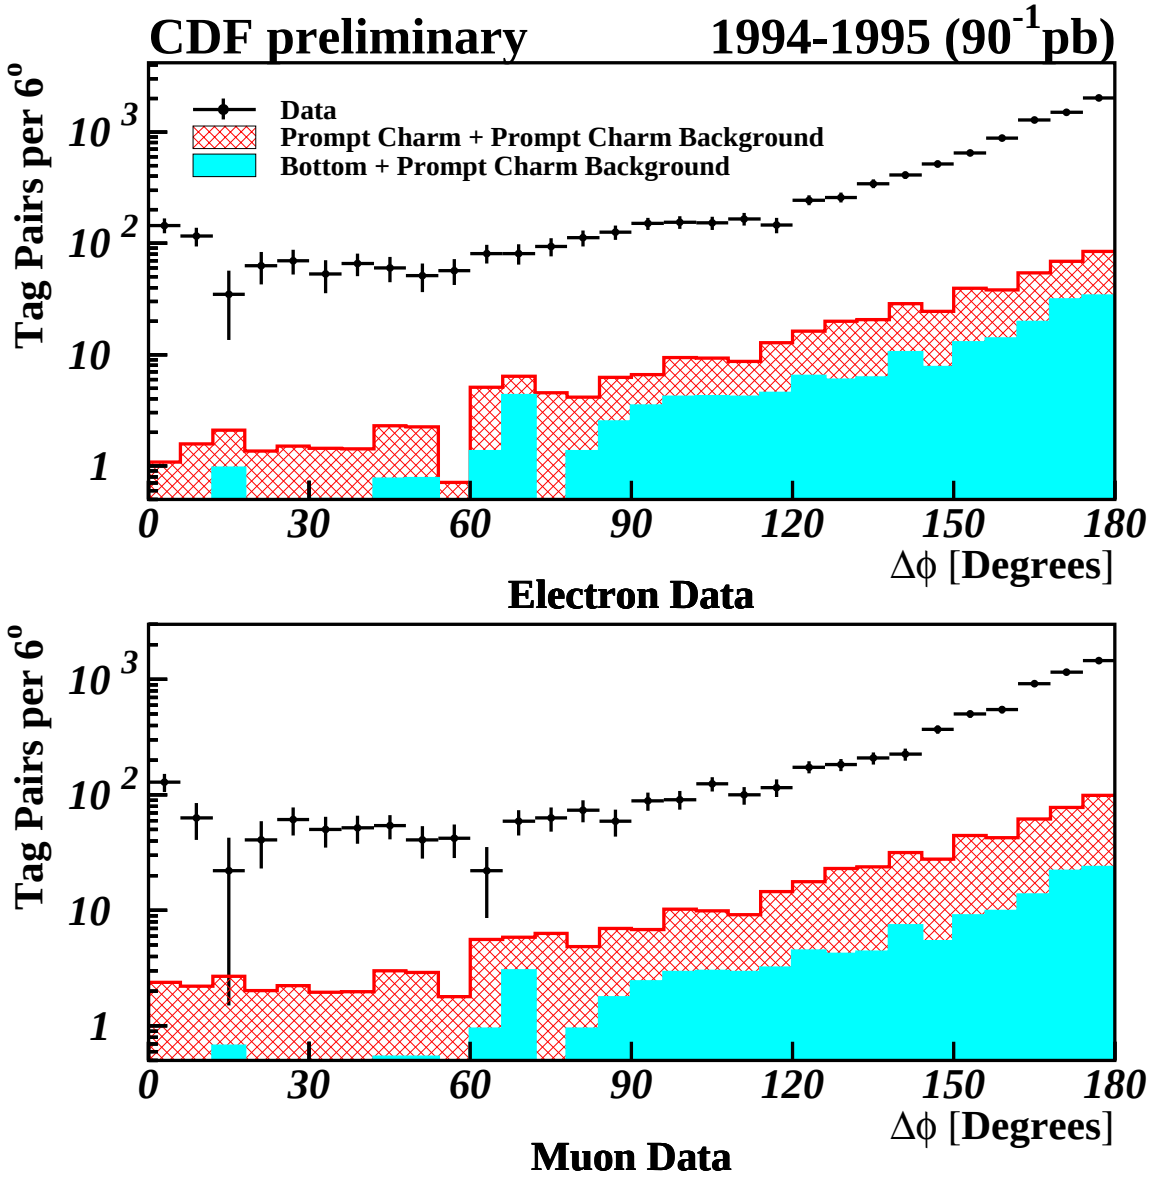

Figure 7.13 The estimated shape of the background from "prompt charm + prompt charm" ($c\bar{c}$ production, the red hatched area) and "bottom plus prompt charm" ($b\bar{b} + c\bar{c}$ production, the solid cyan area). The points with error bars show the data.

the tags removed by the invariant mass cut, 122.1 electron and 155.0 muon tags come from sequential double tag pairs. This leads to an estimate that 2.9 electron sequential tag pairs and 3.7 muon sequential tag pairs remain in the data.

The Monte Carlo study also indicates that the $\Delta\phi$ distribution of the sequential tag pairs is well described by a Gaussian distribution with a mean of zero and a width of 0.122 radians. To correct for the sequential double tag contribution in the data, we take the estimated number of sequential double tags, with a Gaussian distribution as described above, and subtract them from the $\Delta\phi$ bins in the data. We also assign the amount subtracted as a systematic error from the sequential double-tag subtraction.

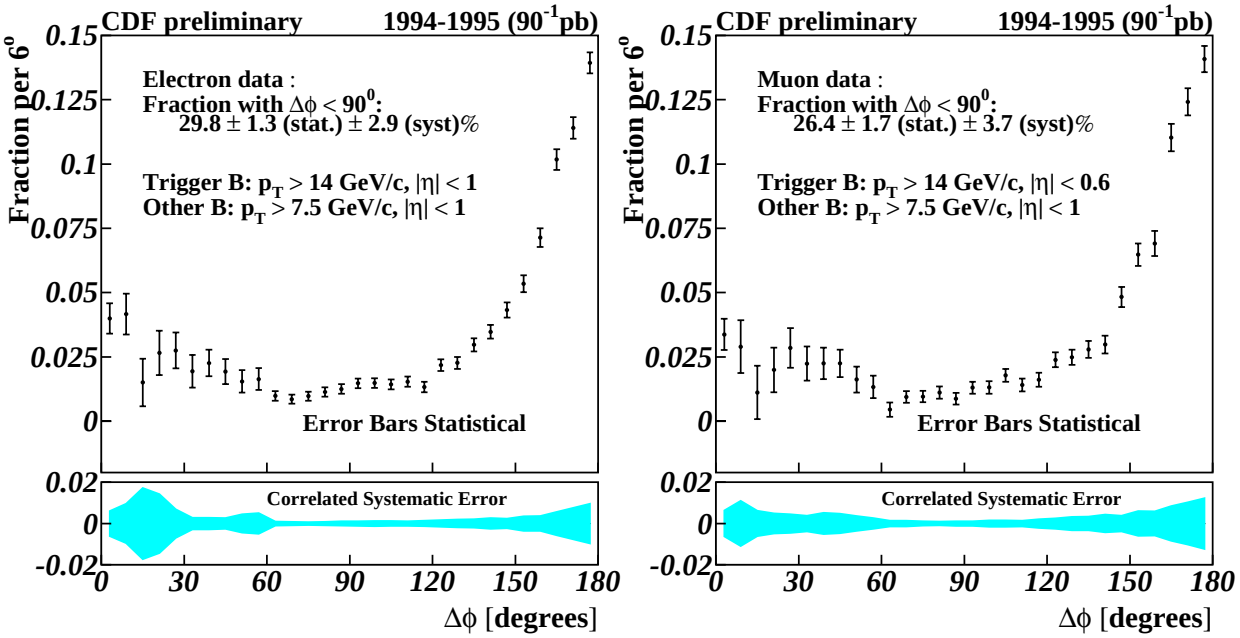

Figure 7.14 The final, corrected $\Delta\phi$ distribution for electron (left) and muon (right) data. The corrections made to the data include mistag subtraction, sequential removal, prompt charm subtraction, and the relative tagging efficiency correction. The error bars display the statistical error on the points. The cyan region at the bottom indicates the systematic errors. The systematic errors are correlated from bin to bin. Mistag subtraction provides the dominant contribution to the systematic errors.

7.3 Corrected Distribution

Figure 7.14 shows the final, corrected tag $\Delta\phi$ distribution, including systematic errors. To obtain this distribution, the contributions from residual sequentials and prompt charm are removed from the mistag-subtracted distributions. Then the relative efficiency corrections derived in Section 7.1 are applied. Systematic errors from the various corrections are combined appropriately to give the total systematic error. Mistag subtraction gives the largest contribution to the systematic error. The final corrected tag $\Delta\phi$ distribution provides a measurement of the $B - \overline{B}$ $\Delta\phi$ distribution where the B providing the trigger electron (muon) has $p_T > 14.0$ GeV/c and $|\eta| < 1.0$ (0.6), and the other B has $p_T > 7.5$ GeV/c and $|\eta| < 1.0$, with a $\Delta\phi$ resolution of 6.22°. This distribution can be compared to generator-level $B - \overline{B}$ $\Delta\phi$ distributions from Monte Carlo that have been convoluted with a Gaussian resolution function to account for our $\Delta\phi$ resolution. Finally, ignoring the small difference in η acceptance between the electron and muon samples, these two distributions can be combined to give the overall B hadron $\Delta\phi$ distribution, shown in Figure 7.15. Table 7.5 specifies the corrected fraction in each $\Delta\phi$ as well as the breakdown of the systematic errors for each bin.

From the corrected data, we can also calculate the fraction of tag pairs in the "towards" region, defined by $\Delta\phi < 90°$. This fraction is of interest because $b\overline{b}$ production in the "towards" region is dominated by the higher order production diagrams. The towards fraction provides a single figure of merit to indicate relative size of the contribution from flavor excitation and gluon splitting. To account for correlated systematic errors, we calculate the towards fraction for our data by essentially repeating the analysis with two $\Delta\phi$ bins instead of thirty, and then taking the ratio of the "towards" bin over the total. For the electron data, we obtain a towards fraction of 29.8 ± 1.3 (stat.) ± 2.9 (syst.)%. For muon data, we obtain a towards fraction of 26.4 ± 1.7 (stat.) ± 3.7 (syst.)%. The electron and muon samples can be combined to give a towards fraction of 28.8 ± 1.0 (stat) ± 3.1 (syst)%. Table 7.6 shows the uncorrected number of tag pairs in the "towards" and "away" bins in the data and gives the corrections applied to obtain the final number. Table 7.7 breaks down the contributions to the systematic uncertainty on the towards fraction.

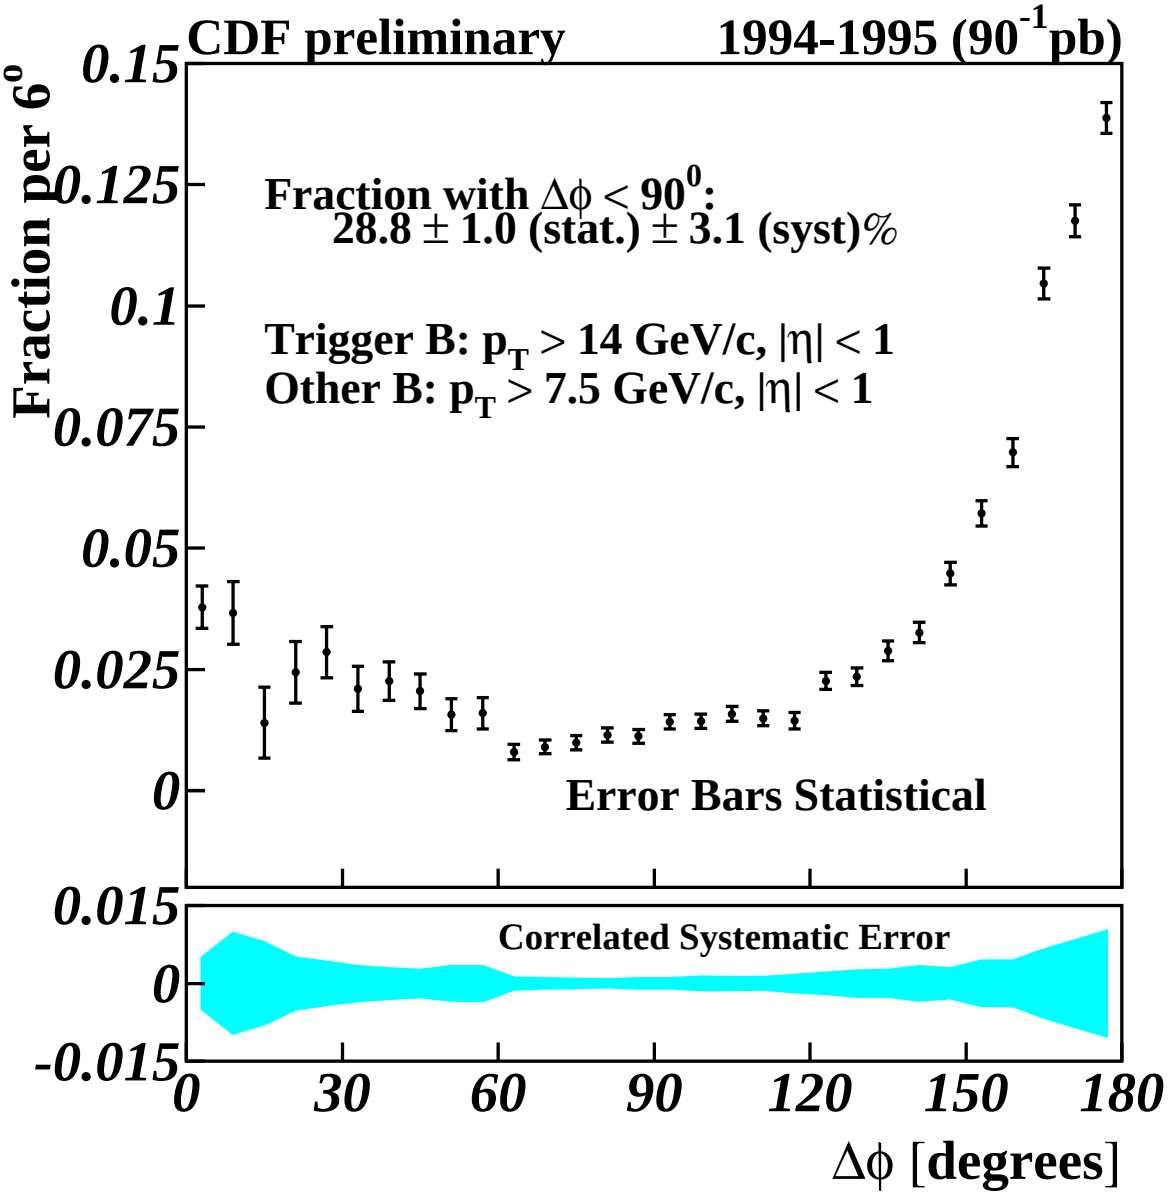

Figure 7.15 The combined, corrected electron and muon $\Delta\phi$ distribution. In making this plot we have ignored the difference in η acceptance between the electron and muon triggers. The corrections made to the data include mistag subtraction, sequential removal, prompt charm subtraction, and the relative tagging efficiency correction. The error bars display the statistical error on the points. The cyan region at the bottom indicates the systematic errors. The systematic errors are correlated from bin to bin. Mistag subtraction provides the dominant contribution to the systematic errors.

				Systematic Error Components			
Bin	**Fraction**	**Statistical Error**	**Systematic Error**	**Sequential**	**Prompt Charm**	**Mistag Subtraction**	**Relative Efficiency**
0°–6°	0.03901	0.00462	0.00411	0.00060	0.00051	0.00421	0.00411
6°–12°	0.03765	0.00684	0.00336	0.00044	0.00082	0.00982	0.00336
12°–18°	0.01347	0.00774	0.00149	0.00013	0.00125	0.00810	0.00149
18°–24°	0.02498	0.00674	0.00257	1.5×10^{-5}	0.00084	0.00472	0.00257
24°–30°	0.02942	0.00561	0.00279	8.6×10^{-7}	0.00087	0.00370	0.00279
30°–36°	0.02152	0.00493	0.00194	2.4×10^{-8}	0.00074	0.00309	0.00194
36°–42°	0.02323	0.00420	0.00206	3.2×10^{-10}	0.00069	0.00256	0.00206
42°–48°	0.02077	0.00379	0.00185	2.1×10^{-12}	0.00101	0.00211	0.00185
48°–54°	0.01568	0.00349	0.00380	6.8×10^{-15}	0.00093	0.00171	0.00380
54°–60°	0.01651	0.00344	0.00438	1.1×10^{-17}	0.00043	0.00137	0.00438
60°–66°	0.00751	0.00167	0.00049	0	0.00087	0.00054	0.00049
66°–72°	0.00869	0.00151	0.00044	0	0.00084	0.00037	0.00044
72°–78°	0.00973	0.00153	0.00045	0	0.00073	0.00030	0.00045
78°–84°	0.01156	0.00156	0.00047	0	0.00059	0.00025	0.00047
84°–90°	0.01100	0.00155	0.00040	0	0.00085	0.00022	0.00040
90°–96°	0.01423	0.00157	0.00046	0	0.00084	0.00023	0.00046
96°–102°	0.01395	0.00160	0.00040	0	0.00121	0.00026	0.00040
102°–108°	0.01559	0.00162	0.00040	0	0.00117	0.00033	0.00040
108°–114°	0.01474	0.00163	0.00034	0	0.00106	0.00044	0.00034
114°–120°	0.01370	0.00177	0.00029	0	0.00159	0.00057	0.00029
120°–126°	0.02203	0.00187	0.00045	0	0.00195	0.00071	0.00045
126°–132°	0.02244	0.00193	0.00045	0	0.00242	0.00082	0.00045
132°–138°	0.02813	0.00213	0.00059	0	0.00246	0.00088	0.00059
138°–144°	0.03128	0.00223	0.00069	0	0.00328	0.00080	0.00069
144°–150°	0.04471	0.00249	0.00106	0	0.00279	0.00051	0.00106
150°–156°	0.05622	0.00275	0.00137	0	0.00444	0.00018	0.00137
156°–162°	0.06983	0.00306	0.00169	0	0.00419	0.00113	0.00169
162°–168°	0.10516	0.00341	0.00319	0	0.00583	0.00267	0.00319
168°–174°	0.11783	0.00346	0.00391	0	0.00688	0.00457	0.00391
174°–180°	0.13944	0.00336	0.00419	0	0.00779	0.00662	0.00419

Table 7.5 The corrected fraction of combined electron and muon data in each bin as well as a breakdown of the components of the systematic errors on each bin. The total systematic error is the sum in quadrature of the individual components.

	Electrons		Muons	
	Towards	**Away**	**Towards**	**Away**
Mistag-Subtracted Data	1210	8887	832	6260
Charm Contamination	42.1	442.6	52.9	500.3
Sequential Contamination	2.9	0.0	3.7	0.0
Relative Efficiency Correction Factor	0.326	1.0	0.376	1.0
Corrected Data	3573.6	8444.4	2062.2	5759.7

Table 7.6 The number of events in the "towards" and "away" regions before and after applying corrections to the data.

	Electrons	**Muons**
Towards Fraction	29.8%	26.4%
Statistical Error	±1.3%	±1.7%
Mistag Subtraction Systematic Error	±2.0%	±2.8%
Sequential Removal Systematic Error	±0.05%	±0.09%
Charm Subtraction Systematic Error	±1.3%	±2.2%
Relative Efficiency Correction Systematic Error	±1.6%	±1.1%
Total Systematic Error	±2.9%	±3.7%

Table 7.7 The break-down of the systematic errors by contribution. The total systematic error is the quadrature sum of the individual components.

Chapter 8

Conclusions

This analysis was motivated primarily by the discrepancy observed between measurements of the single-inclusive *b* quark cross section and predictions from NLO QCD and Monte Carlo models using the parton shower or leading-log approximation. The factor of two to three discrepancy between the measurements and the NLO prediction remains of great concern, and it is unclear why the less formally exact parton shower Monte Carlo models seem to agree with the measurements better. The main goals of this analysis were two-fold: to compare a measurement of the *B* hadron $\Delta\phi$ spectrum to prediction from HERWIG and PYTHIA to evaluate whether these Monte Carlo programs provide reasonable estimates of the contributions to $b\bar{b}$ production from the higher-order flavor excitation and gluon splitting production mechanisms, and to provide a measurement of the $\Delta\phi$ distribution with high statistics and reasonable sensitivity down to arbitrarily small opening angles that can be compared to additional theoretical predictions from fixed-order calculations (NLO, etc.) and improved parton shower approaches (NLO + parton shower, for example). On both fronts this analysis has provided useful insight, although some questions still remain.

This analysis has established that the $\Delta\phi$ distribution predicted by PYTHIA and HERWIG match qualitatively, if not quantitatively with the distribution measured from the data. Both Monte Carlo models provide reasonable matches the data, although there are minor differences in some regions, particularly the width of the peak at high $\Delta\phi$. The agreement between PYTHIA and HERWIG and the data establishes that the higher-order production mechanisms, particularly flavor excitation and gluon splitting, provide significant contributions to the total $b\bar{b}$ production cross section at the Tevatron. These contributions result in a non-negligible fraction of $b\bar{b}$ pairs being produced with small opening angles. For *B* measurements that are sensitive to such correlations—for example, some *B* mixing analyses—it is necessary to consider $b\bar{b}$ production from all three sources, as modeled by PYTHIA, HERWIG, or a NLO calculation, to account for the

effects of small opening angle B production. In general, running the lowest order $b\bar{b}$ production mechanisms in PYTHIA (MSEL = 5) and HERWIG (IPROC = 1705) may not be sufficient for all analyses.

Despite the conclusion that PYTHIA and HERWIG provide adequate qualitative models for higher-order $b\bar{b}$ production, this analysis is not able to resolve whether individually, the size of the flavor excitation and gluon splitting contributions from PYTHIA and HERWIG are justified. For example, much of the agreement in cross section between PYTHIA and HERWIG comes from the large size of the flavor excitation component. However, depending on the amount of initial-state radiation used for the PYTHIA prediction, there is reasonable agreement between the data and PYTHIA for cases that involve both little flavor excitation and much flavor excitation. PYTHIA seems able to trade off initial-state radiation for the size of the flavor excitation component when matching the data. Although HERWIG does not expose its initial-state radiation parameters to the user, it is possible that by adjusting the amount of initial-state radiation in HERWIG a similar effect would be observed. So, in the end, these studies are unable to resolve separately the exact size of the flavor excitation and gluon splitting contributions that should be included in PYTHIA and HERWIG to match the data. Indeed, it is probably worth realizing that even though the production mechanisms do not interfere with one another in the parton shower approximation, since the production mechanisms refer to quantum mechanical amplitudes, some of which have identical final states, in reality interferences make experimental separation of the contributions from each production mechanism impossible.

In the end, this analysis establishes that PYTHIA and HERWIG provide adequate qualitative models for $b\bar{b}$ production, including contributions from higher-order production. However, because of the approximations involved in the parton shower model, the predictions from these Monte Carlo programs, especially for flavor excitation and gluon splitting, should not be considered quantitatively accurate. For higher precisions studies of $b\bar{b}$ production, a more exact model, like a NLO calculation that uses the parton shower approach to account for the effect of multiple light quark and gluon emissions, is desirable.

However, since the agreement between the data and the parton shower Monte Carlo programs is reasonable, the Monte Carlo samples can be used to extract the B hadron $\Delta\phi$ distribution from the measured tag pair distribution. This B hadron distribution provides a convenient reference against which future $b\bar{b}$ correlation predictions can be compared. However, in making such comparisons, one has to keep in mind the p_T and η acceptances of the measurements made here,

as well as the $\Delta\phi$ resolution, as detailed in Section 7.3. In addition, because the kinematics of this analysis are similar to the kinematics of several *B* mixing analyses, insights into the $\Delta\phi$ correlations measured here can be applied to issues involved with *B* mixing. In particular, it is important to remember that roughly 30% of the $b\bar{b}$ production measured in this analysis results in an opening angle between the *b* quarks of less than 90°.

Beyond the motivation of understanding *b* quark production, the measurements in this analysis provide input to several other interesting areas of high-energy physics. As already mentioned, certain techniques for measuring *B* mixing benefit from an improved understanding of the angular correlations between *B* hadrons. An example application of how higher-order *B* hadron production mechanisms can impact *B* mixing measurements is given in Appendix A. In addition, the secondary vertex tagging techniques employed in this analysis are also useful for detecting events that contain more than two secondary vertex tags. As shown in Figure 4.9, a small fraction of the events used for this analysis contain more than two secondary vertex tags. It is interesting to study such events, not only because of the QCD mechanisms for generating multiple heavy flavor pairs in a single interaction, but also because Higgs production and other exotic processes have multiple *b* signatures that can be detected by searching for an excess of events with multiple *B* tags. This issue and some preliminary results from the data are discussed in Appendix B.

The future, and in particular Run II at the Tevatron, holds much promise for understanding the unresolved questions surrounding $b\bar{b}$ production at hadron colliders. During Run II, it is expected that CDF will collect large samples of *B* and *D* decays, not only because of the increase in luminosity over Run 1, but also because of the improved techniques available for triggering on heavy flavor decays. In particular, at CDF in Run II, a new element, the secondary vertex trigger (SVT) has been added [72]. The SVT allows CDF to trigger on the presence of tracks with large impact parameters in the Level 2 trigger. Not only does the SVT provide a convenient trigger for high purity $b\bar{b}$ samples when combined with a lepton requirement, it also, when used alone, allows for the collection of large samples of fully hadronic *B* and *D* decays. These decays can be fully reconstructed, since they do not involve neutrinos, which cannot be detected effectively at CDF. At the time of this writing, the SVT has already made possible the collection large fully reconstructed charm samples with larger samples promised for the future [73]. Because these samples contain fully reconstructed heavy flavor decays, they allow an exploration of a wider range of correlations such as momentum correlations or even flavor correlations. Studying

flavor correlations is of particular interest to *B* mixing and *CP* violation measurements. The assumption for Standard Model QCD $b\bar{b}$ production is that that *b* and $\bar{b}$ fragment independently of one another so that the type of *B* hadron formed by the *b* has no effect on the type formed by the $\bar{b}$. However, this feature has never been experimentally tested at hadron colliders. Any correlation between the flavors of the two *B* hadrons formed in $b\bar{b}$ production could lead to a distortion in certain mixing and *CP* violation measurements. With the new heavy flavor data currently being taken at the Tevatron, and improvements in the understanding for theoretical issues involving $b\bar{b}$ production and *b* fragmentation, it seems likely that many of the unresolved questions in $b\bar{b}$ production will be answered. In addition, many new avenues of exploration will open up, perhaps leading to additional surprises and mysteries.

Appendix A

B Flavor Tagging and Acceptance

As described in Section 1.1, since the mass eigenstates of the Standard Model are not equal to the electroweak eigenstates, electroweak interactions introduce a mixing among quark flavors, parameterized by the CKM mixing matrix, given below.

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \tag{A.1}$$

The CKM matrix is interesting for a number of reasons, not the least of which is that it allows for an imaginary phase that leads to *CP* violation in electroweak interactions. The amount of *CP* violation allowed through this mechanism is not known from first principles and so must be measured experimentally. *CP* violation is of interest in the field of cosmology because it is one of the necessary conditions for generating the net matter excess (over antimatter) observed in our universe starting from a matter-antimatter symmetric initial state [74]. Although the amount of *CP* violation allowed by the CKM matrix does not appear large enough to explain the asymmetry between matter and antimatter, additional *CP* violation from non-Standard model sources, such as neutrino mixing or additional Higgs bosons may be able to make up the difference [75]. Multiple, overlapping measurements of the CKM matrix elements and *CP* violation allow us to over-constrain the parameters of the CKM matrix. Any sign of inconsistency among the measurements would indicate physics beyond the Standard Model. For this reason, measurements of the CKM matrix and *CP* violation are at the forefront of interest for the world high-energy physics community.

Because the decay of a *B* hadron involves a weak transition of the *b* quark to a lower mass quark, *B* hadron decays provide a useful opportunity to study the CKM matrix and *CP* violation [76]. Two of the most prevalent techniques involve measuring *B* mixing and studying decays of

B hadrons to *CP* eigenstates. *B* mixing occurs when a neutral *B* meson, like the B^0 composed of a $\bar{b}$ quark and a *d* quark transforms into a $\bar{B}^0$ containing a *b* quark and a $\bar{d}$ quark, through the weak process shown in Figure A.1. Note that although any "up-type" quark, *u*, *c*, or *t*, may participate in the loops shown in Figure A.1, the top quark contribution dominates because the amplitude for the loop diagram increases with the mass of the quark in the loop. This process causes neutral *B* mesons, namely B^0 and B_s^0 to oscillate into $\bar{B}$ mesons and vice versa. The oscillation frequency is proportional to the combination of CKM matrix elements $\left|V_{tx}V_{tb}^*\right|^2$, where *x* is either *d* or *s* depending on the type of neutral *B* meson. Since this measurement is only sensitive to the amplitude of the CKM matrix elements, it does not give direct information about *CP* violation, although it can be used, in conjunction with other measurements, to check the consistency of the *CP* violating parameters of the CKM matrix.

Another approach to measuring the parameters of the CKM matrix using *B* physics is to study the decay of *B* mesons to a *CP* eigenstate that is accessible both to the *B* and $\bar{B}$ states. Because of mixing for neutral *B* mesons, the decay of a *B* meson into such a *CP* eigenstate can proceed either directly, or via mixing to the $\bar{B}$ state before decaying, as shown in Figure A.2. These two

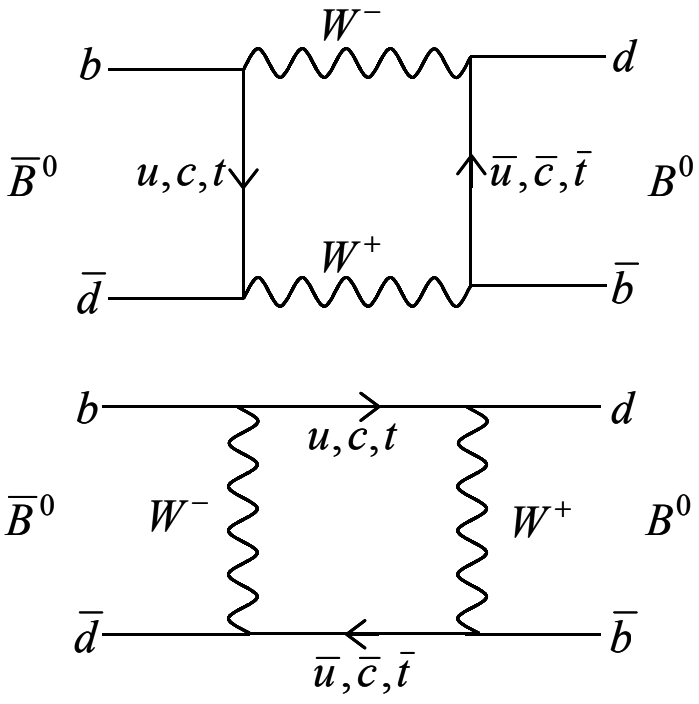

Figure A.1 Feynman diagrams for B^0 mixing.

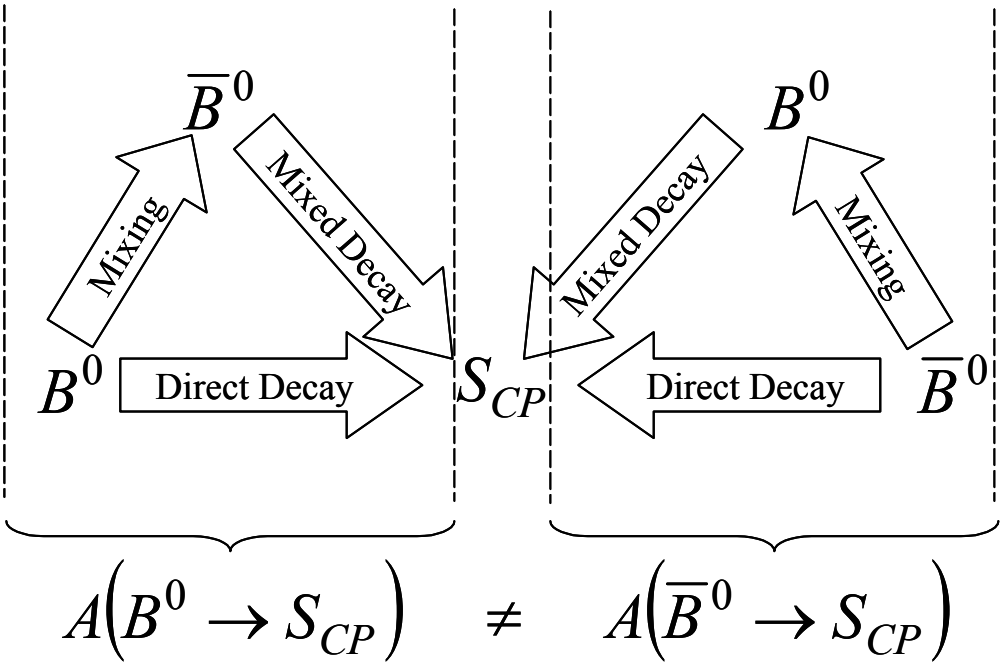

Figure A.2 When the *CP* eigenstate S_{CP} is accessible to the decay of both the B^0 and $\overline{B}^0$, interferences between the direct and mixed decays cause the amplitude for $B^0 \to S_{CP}$ not to be equal to the amplitude for $\overline{B}^0 \to S_{CP}$.

decay paths interfere and the interference is different between the *B* and $\overline{B}$ decays, leading to different decay rates. Studying the difference in these decay rates gives information about the amount of *CP* violation in the Standard Model.

A.1 Flavor Tagging

Measurements of the CKM matrix parameters using *B* physics, like *B* mixing and *B* decays to *CP* eigenstates, rely on an experimental technique known as *flavor tagging*. Flavor tagging refers to a suite of methods for determining the flavor of a neutral *B* meson—that is whether it was a *B* or $\overline{B}$ —at a particular time. Flavor tagging in *B* mixing analyses is used to determine the flavor of the *B* both at the time of production and at the time of decay. In studies of *B* hadrons decaying to *CP* eigenstates, flavor tagging is used to determine whether the decay that produced the *CP* eigenstate of interest was of a *B* or $\overline{B}$. Often, flavor tagging relies on an indirect determination of the *B* flavor, either because the *B* was not explicitly reconstructed in the measurement, or

because the *B* decayed to a state, like $J/\psi K_S^0$, that contains no direct information about the flavor of the *B* decay.

Flavor tagging techniques may be classified either as *same-side tagging* or *opposite-side tagging* depending on whether they rely on information from the *B* whose flavor is being determined, or the other *B* in the event.

A.1.1 Same-Side Tagging

Any flavor tagging technique that uses information directly related to the *B* is considered a same-side tagging method. A trivial example of same-side tagging is using the charge of a lepton from the semileptonic decay of a *B* to determine the flavor of the *B*: positive charge indicates a *B* while a negative charge indicates a $\overline{B}$. This technique does not require that the entire *B* decay be reconstructed. For example, it could be used to determine the flavor of a *B* hadron from the experimental signature of a lepton contained within a jet.

In some cases, it is not possible to determine the desired flavor information directly from the *B* decay. For example the decay $B^0 / \overline{B}^0 \to J/\psi K_S^0$ contains no direct information about the *B* flavor. In addition, for studies of *B* mixing, it is necessary to know not only the *B* flavor at decay, but also at production, to determine if mixing has occurred. In these cases, to use same-side tagging, one must rely on additional input, such as information from the fragmentation of the *b* quark that produced the *B* meson. For example, a *b* quark can fragment to a $\overline{B}^0$ through the acquisition of a $\overline{d}$ quark from a $d\overline{d}$ pair popped from the vacuum. If the remaining *d* quark combines with a $\overline{u}$ quark to form a π^- meson, the charge of the pion can be used to determine the flavor of the *B* meson produced. It has been demonstrated that the statistical correlation between the pion charge from *B* fragmentation and the *B* flavor can be used as a same-side tagging method, and CDF has used this approach as part of its measurements of *CP* violation in the *B* system [77].

A.1.2 Opposite-Side Tagging

The overwhelming majority of *B* hadrons produced at the Tevatron result from $b\overline{b}$ pair production through the strong interaction. That means for every $\overline{B}$ produced containing a *b* quark, there is a *B* somewhere in the event, containing a $\overline{b}$ quark. Opposite-side tagging techniques use the flavor correlation between the two *B* hadrons resulting from $b\overline{b}$ production to

determine the flavor of one B from information about the other B. Frequently, in order to maintain reasonable efficiency for accepting B events, the other B decay used for the opposite-side tag cannot be explicitly reconstructed. Instead, some assumptions are made about where in the event the other B lies.

Two examples of opposite-side tagging methods used at CDF are *soft lepton tagging* and *jet charge tagging*. Soft-lepton tagging is used for cases in which there is a lepton in the event that is not associated with the B hadron whose flavor is to be determined. This lepton is labeled *soft* because often the lepton p_T is below the lepton trigger threshold. The charge of this lepton can be used to determine the flavor of the "opposite-side" B, as describe in Section A.1.1 above, and then the flavor of the B in question is taken to be opposite. In jet charge tagging, a jet in the event, typically the highest E_T jet not associated with the B hadron decay in question, is assumed to contain the other B decay in the event. A momentum-weighted sum of the charge in the jet is used to determine the flavor of the B decay in that jet. The flavor of the B hadron of interest is then assumed to be opposite.

A.1.3 Modeling with Monte Carlo

In nearly every flavor tagging method described above, the B hadron decay is not fully and explicitly reconstructed. As a result certain assumptions must be made—for example, whether or not a second B is present in the detector—and the results of the flavor tagging become more complicated to interpret. In addition, certain flavor tagging methods, such as jet charge tagging, give the correct results on average, but the quality of the results for individual events may vary. Understanding the quality of the flavor tag result, or even when the flavor tagging method is applicable, is important. Finally, under certain circumstances, a flavor tag method may seem applicable yet give completely false results. For example, in the case of a hadronic B decay involving a D meson that decays semileptonically, the charge of the lepton from the D decay gives exactly the wrong answer for the flavor tag. Again, to correctly interpret the results of any flavor tagging method, it is important to understand the circumstances under which the flavor tagger may report wrong results.

For these reasons, it is helpful to have a $b\bar{b}$ Monte Carlo generator for use in designing, optimizing, and interpreting flavor tagging algorithms. Because some flavor tagging methods rely on correlations from b fragmentation, the Monte Carlo used should incorporate a realistic model for b quark fragmentation. In addition, because flavor tagging performance can depend

on how much activity is present in the underlying event, it is useful to have a Monte Carlo program that generates realistic hadron collider final states. These qualities are present in parton shower Monte Carlo programs like PYTHIA and HERWIG, and these programs are frequently used to study flavor tagging techniques.

In addition to the considerations above, it may also be important, when studying certain kinds of flavor tagging methods, to make sure that the Monte Carlo correctly accounts for the higher-order $b\bar{b}$ production mechanisms. Correctly modeling $b\bar{b}$ correlations is important for both same-side and opposite-side tagging methods. For same-side tagging, the presence of a second B decay in the region of interest may distort the same-side tagger's results. For example, if the pion whose charge is used to define the same-side flavor tag comes from the other B in the event, the wrong B flavor will be inferred. On the other hand, for opposite-side tagging, understanding $b\bar{b}$ correlations is useful for determining whether there is a second B in the detector acceptance and if so, where in the detector that B is likely to be found. Depending on the flavor tagging algorithm, getting the correct answer to the above questions may involve correctly accounting for the higher-order $b\bar{b}$ production mechanisms. In particular, many studies of flavor tagging use only the flavor creation contribution to $b\bar{b}$ production because this is the easiest contribution to generate with PYTHIA or HERWIG. In some cases, this practice may give misleading results.

A.2 Monte Carlo Acceptance Study

Consider a question of particular interest for opposite-side flavor tagging methods: given a $b\bar{b}$ event in which one b quark produces a B hadron that decays within the detector acceptance above a certain momentum threshold, what is the chance the other B will also be contained in the detector acceptance? This question is especially relevant for opposite-side tagging because this method is only effective if the other B in the event decays within the detector. At best, opposite-side flavor taggers will not be able to return any flavor information from events in which the second B was not contained within the detector. At worst, such a tagger will use information from particles not related to the other B decay and draw random, erroneous conclusion.

The basic idea behind this study is as follows: We generate a number of $b\bar{b}$ Monte Carlo events using PYTHIA and HERWIG. Using only generator-level information, meaning no attempt has been made to simulate the effects of the detector on the data, we count the number of events, N_{trig}, containing at least one B hadron with $|\eta| < 1$ and $p_T > p_T(\text{trig})$. The pseudorapidity range

chosen is representative of the main instrumented region of the CDF detector, known as the central region, which is used for most *B* physics analyses. From the event sample defined above, we then identify a number of events, N_{other}, in which the other *B* hadron in the event is also in the central region ($|\eta| < 1$) and which has $p_T > p_T(\text{min})$. We report the results as an acceptance, defined by $A = N_{other} / N_{trig}$, for $b\bar{b}$ events to have a second *B* in the detector acceptance, given one *B* in the detector acceptance, as a function of different $p_T(\text{trig})$ and $p_T(\text{min})$ values.

A.2.1 Monte Carlo Samples

For this study, we generated PYTHIA and HERWIG Monte Carlo samples using the same techniques described in Chapter 5. However, we did not use a detector or trigger simulation. This study includes PYTHIA Monte Carlo generated with lower initial-state radiation, PARP(67) = 1.0, and higher initial-state radiation, PARP(67) = 4.0. Also, for both PYTHIA and HERWIG, we used the lowest minimum parton-parton center of mass p_T cutoff reasonable for the production mechanism. For PYTHIA, we used a minimum parton-parton center of mass p_T of 0 GeV/c for flavor creation and 5 GeV/c for flavor excitation and gluon splitting. For HERWIG, we used 0 GeV/c for flavor creation and flavor excitation and 5 GeV/c for gluon splitting. To combine events from the different production mechanisms, we used the PYTHIA and HERWIG predictions for the cross section of each production mechanism. The cross sections given by PYTHIA and HERWIG for the different production mechanisms are shown in Table A.1.

A.2.2 Acceptance and Production Mechanism Plots

In Figure A.3 through Figure A.14, we plot the acceptance as defined above. For the trigger *B*, we choose the *B* with the highest p_T in the central region ($|\eta| < 1$). The acceptance is plotted as a function of the minimum required p_T for the other *B* in the central region, $p_T(\text{min})$. In each plot, the trigger *B* is required to be above a minimum p_T as noted. The curve corresponding to each contribution shows what the acceptance would be if $b\bar{b}$ production consisted only of that contribution. Each plot also shows the acceptance for all contributions to $b\bar{b}$ production added according to the cross sections predicted by the Monte Carlo program.

Although the *B* acceptance plots show the fraction of the events that have at least two *B* hadrons in the acceptance—the trigger *B* and one other—it is possible to produce more than one $b\bar{b}$ pair in an event and therefore to have more than two *B* hadrons in the acceptance. An

HERWIG

	Cross section with no cuts	Cross section with leading *B* having $\|\eta\| < 1$ and $p_T >$ (below):			
		5 GeV/c	**10 GeV/c**	**15 GeV/c**	**20 GeV/c**
Flavor creation	28.7 μb	4.29 μb	0.757 μb	0.159 μb	46.4 nb
Flavor excitation	34.8 μb	8.01 μb	1.57 μb	0.351 μb	104 nb
Gluon splitting	19.9 μb	3.45 μb	0.611 μb	0.139 μb	41.1 nb
Total	83.4 μb	15.75 μb	2.94 μb	0.649 μb	192 nb

PYTHIA, `PARP(67)` = 4.0

	Cross section with no cuts	Cross section with leading *B* having $\|\eta\| < 1$ and $p_T >$ (below):			
		5 GeV/c	**10 GeV/c**	**15 GeV/c**	**20 GeV/c**
Flavor creation	39.1 μb	7.29 μb	1.10 μb	0.210 μb	54.7 nb
Flavor excitation	42.3 μb	12.4 μb	3.05 μb	0.664 μb	190 nb
Gluon splitting	22.2 μb	4.67 μb	1.02 μb	0.249 μb	77.8 nb
Total	103.6 μb	24.4 μb	5.17 μb	1.123 μb	323 nb

PYTHIA, `PARP(67)` = 1.0

	Cross section with no cuts	Cross section with leading *B* having $\|\eta\| < 1$ and $p_T >$ (below):			
		5 GeV/c	**10 GeV/c**	**15 GeV/c**	**20 GeV/c**
Flavor creation	39.1 μb	6.66 μb	0.894 μb	0.155 μb	41.1 nb
Flavor excitation	42.2 μb	11.1 μb	2.22 μb	0.450 μb	129 nb
Gluon splitting	14.8 μb	3.22 μb	0.707 μb	0.175 μb	58.5 nb
Total	96.1 μb	21.0 μb	3.82 μb	0.780 μb	229 nb

Table A.1 shows the cross sections predicted by PYTHIA and HERWIG for the various contributions to $b\bar{b}$ production before requirements are placed on the second *B* hadron. The leftmost column shows the cross sections when no cuts are placed on the *B* hadrons. The other columns show the resulting cross sections when cuts are placed on only the leading *B* hadron in the event. The leading *B* hadron is defined as the *B* with the highest p_T in the central region.

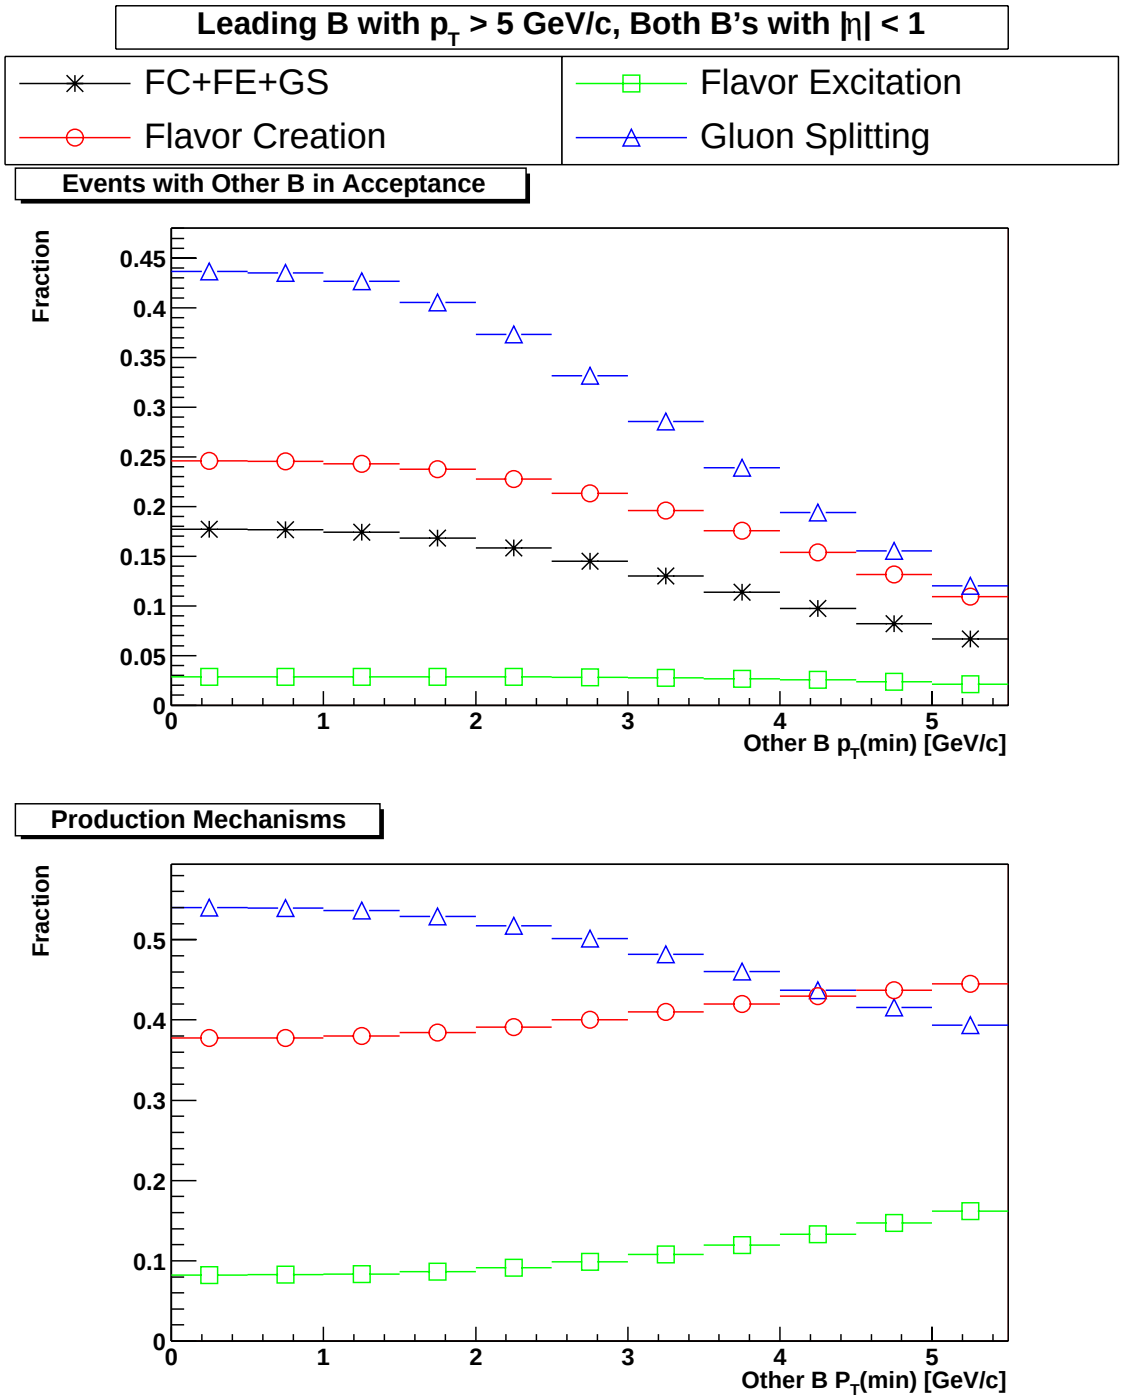

Figure A.3 shows plots made from HERWIG. The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 5$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

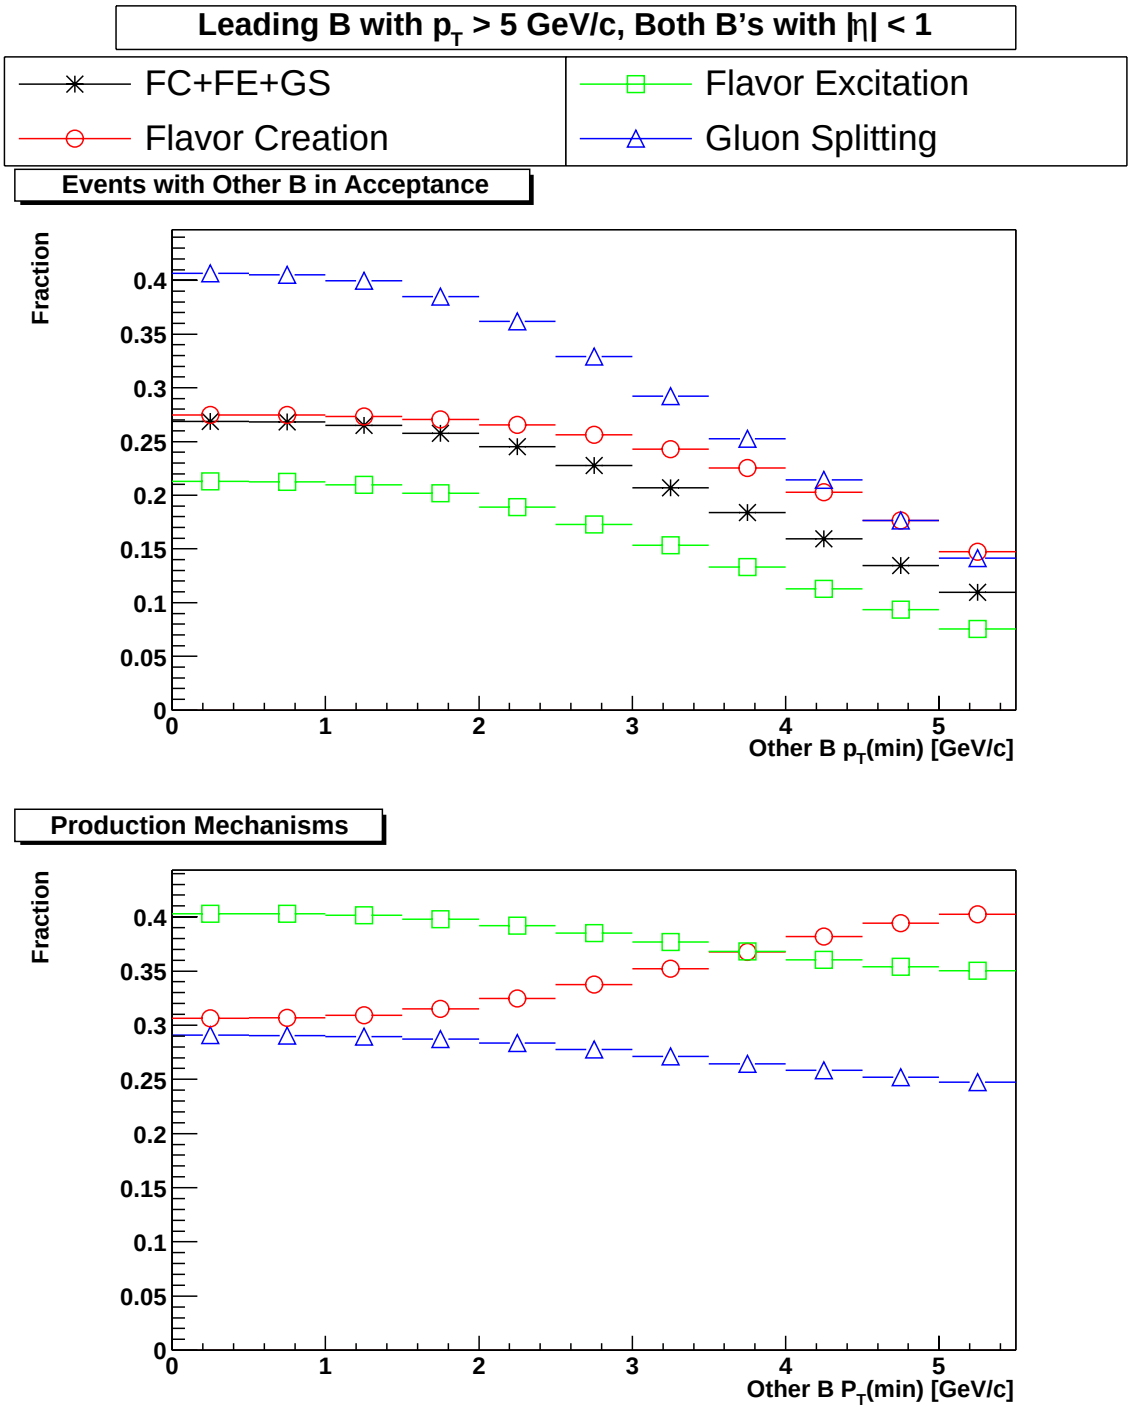

Figure A.4 shows plots made from PYTHIA with `PARP(67)` = 4.0 (more initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 5$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

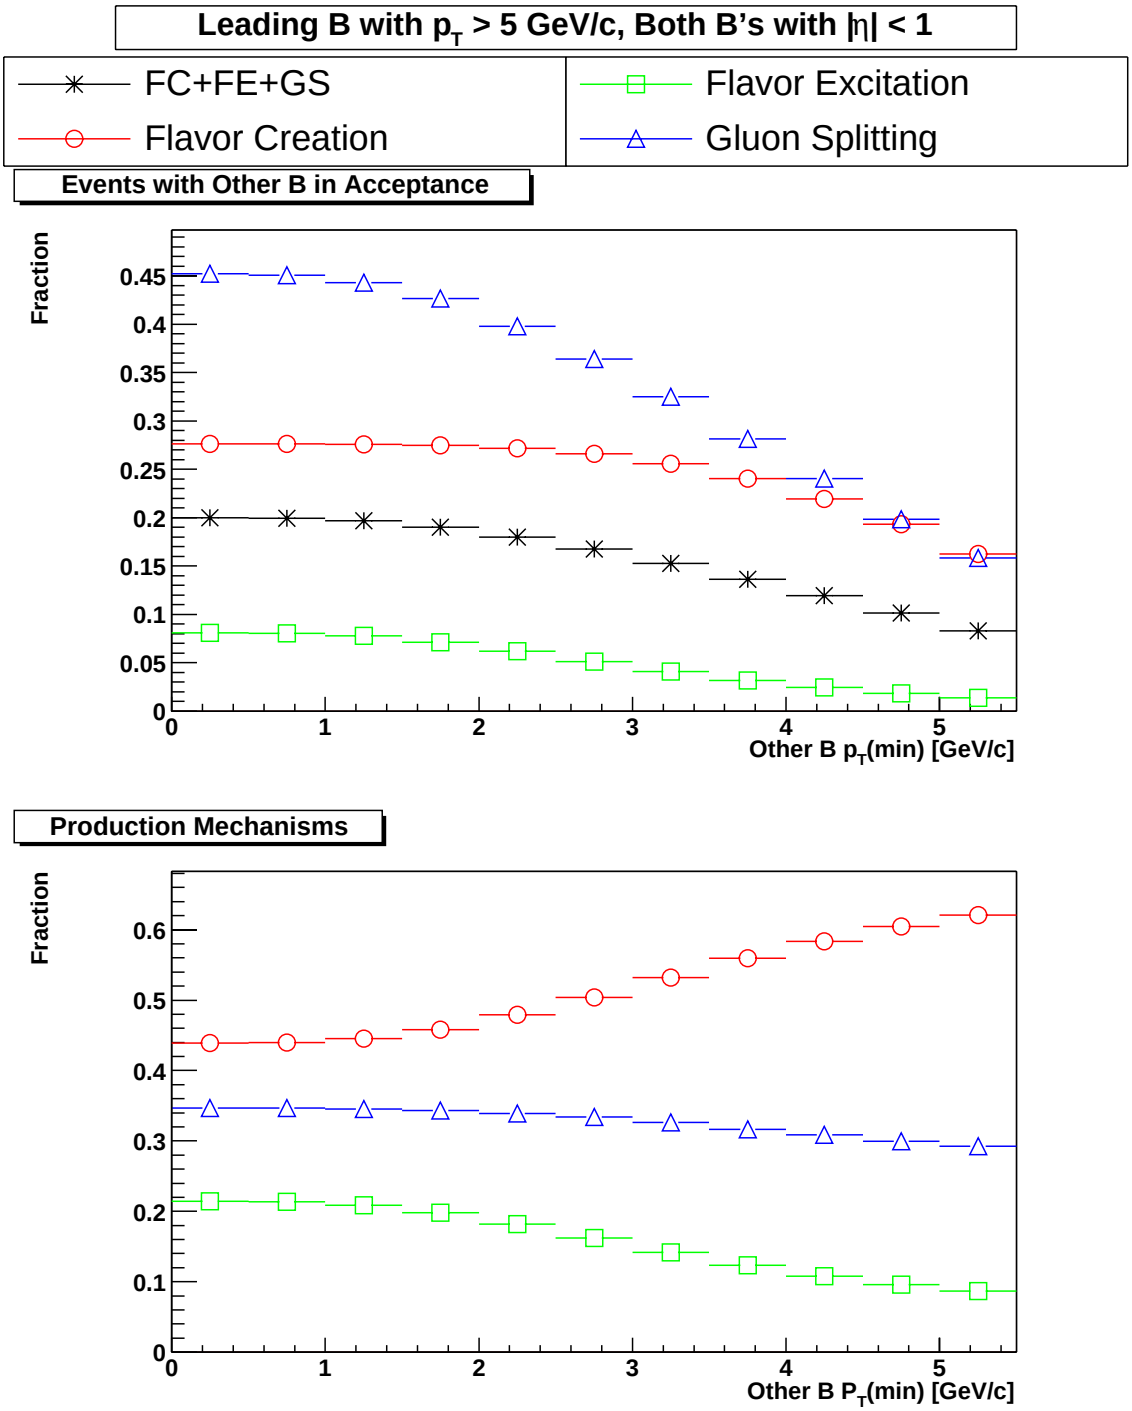

Figure A.5 shows plots made from PYTHIA with `PARP(67)` = 1.0 (less initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 5$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

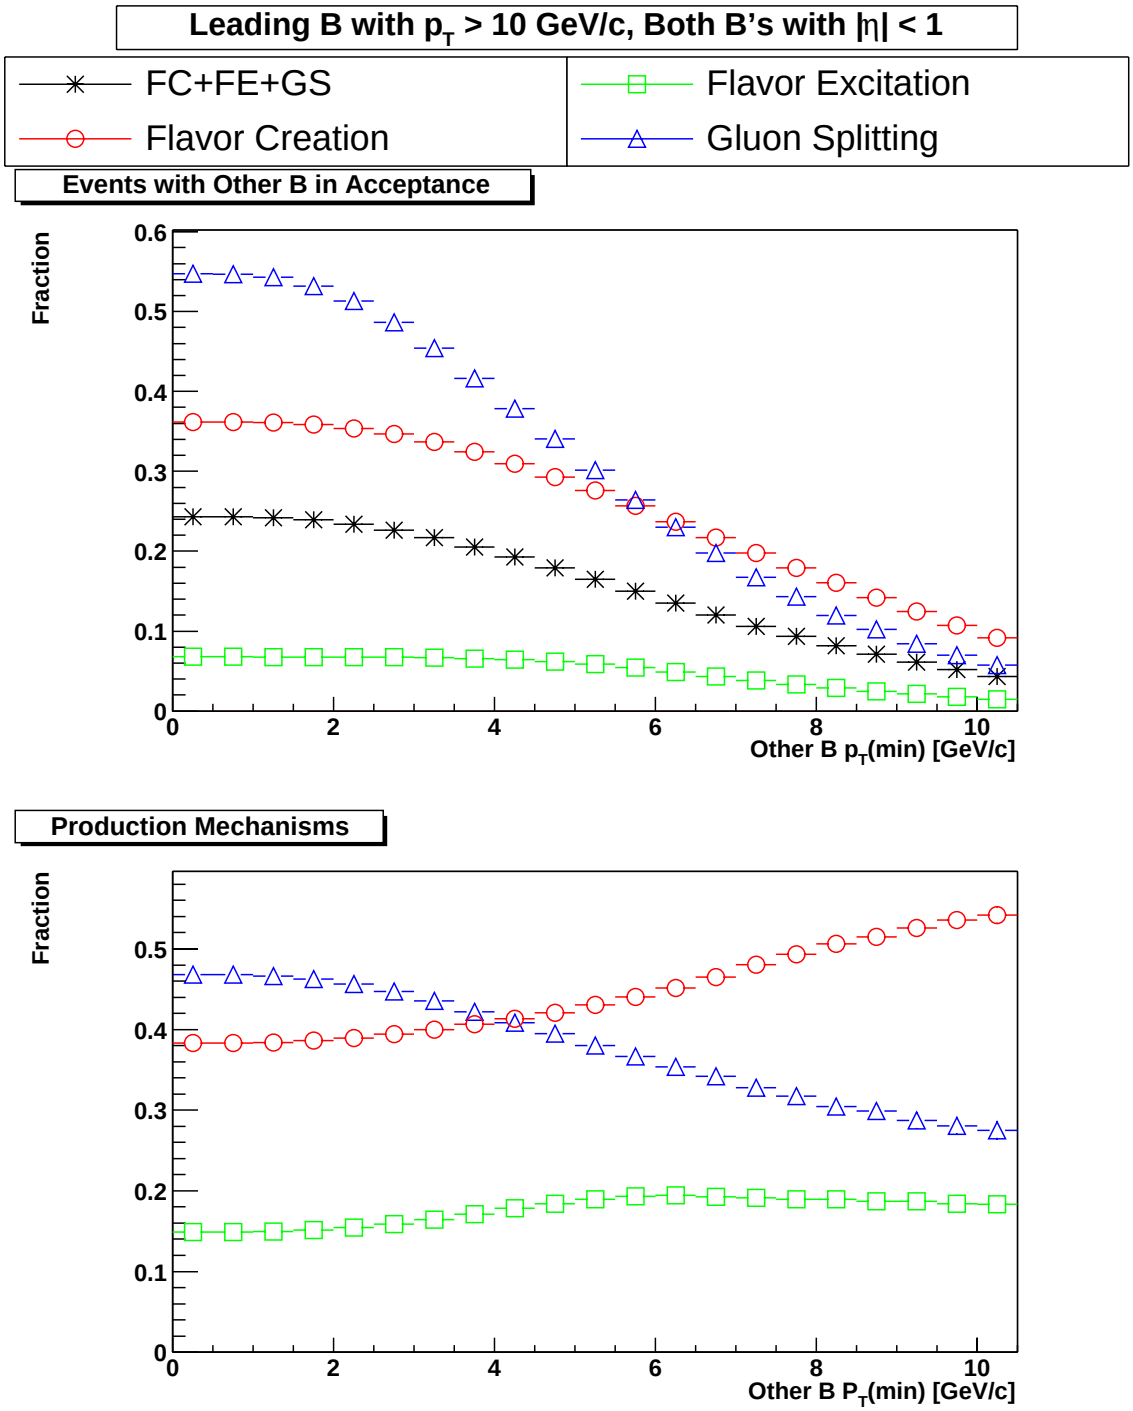

Figure A.6 shows plots made from HERWIG. The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 10$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

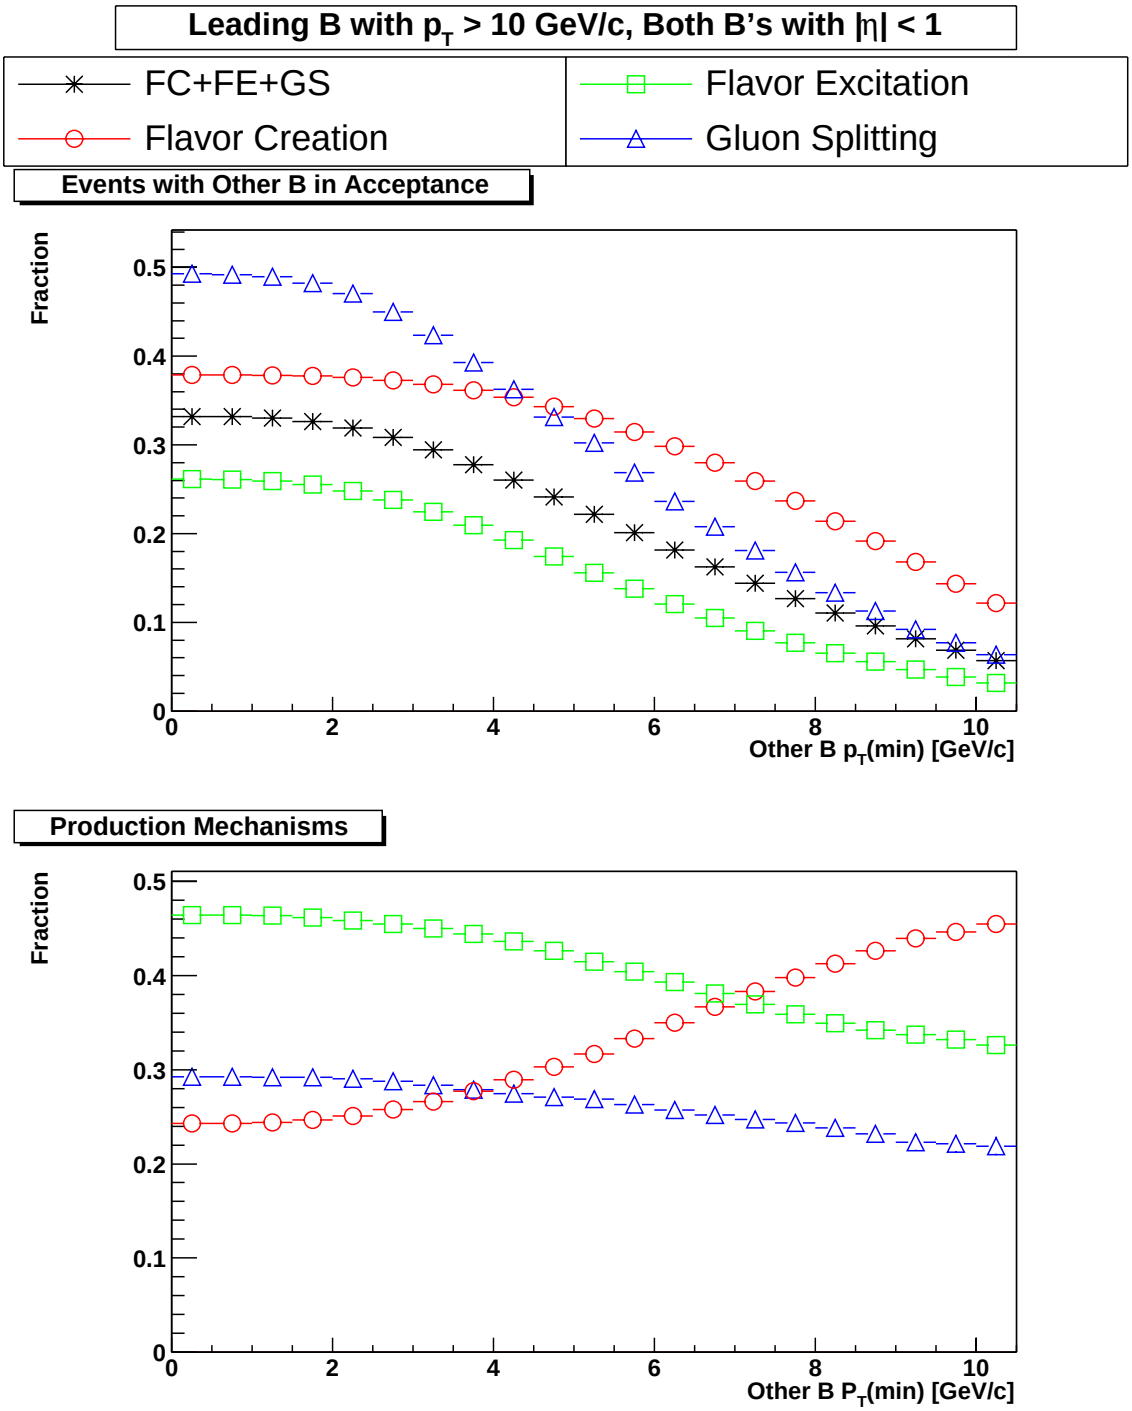

Figure A.7 shows plots made from PYTHIA with `PARP(67)` = 4.0 (more initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading *B* hadron in the central region (|η| < 1) with p_T > 10 GeV/c also has a second *B* in the central region, as a function of the p_T of the other *B*. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

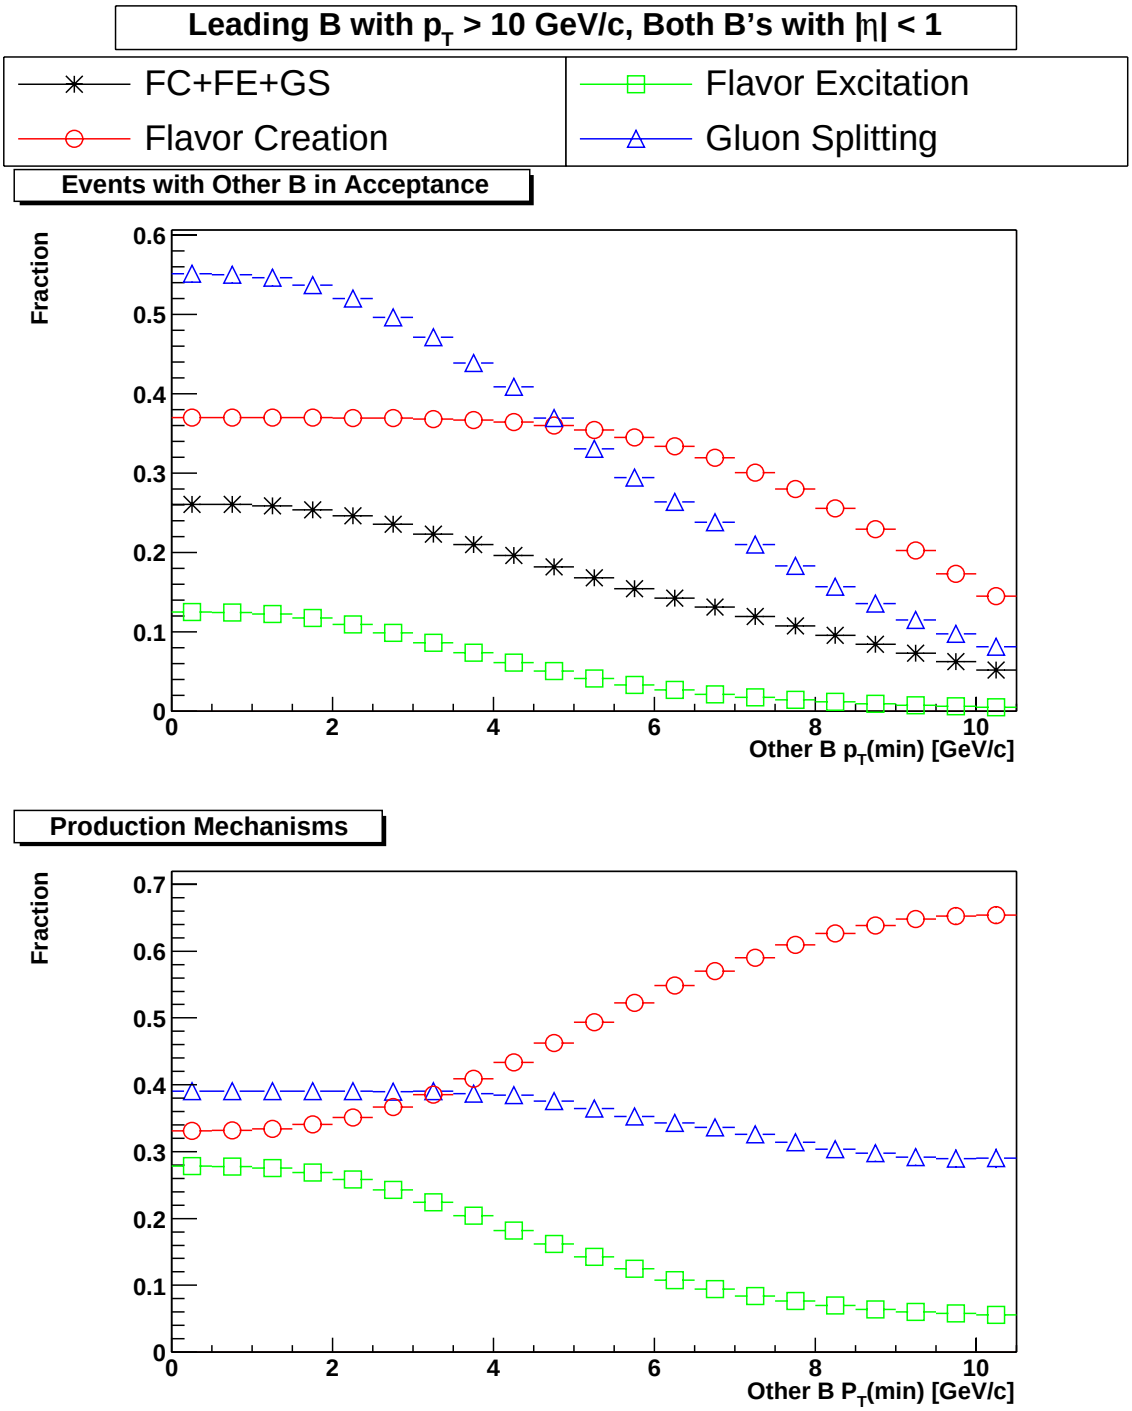

Figure A.8 shows plots made from PYTHIA with `PARP(67)` = 1.0 (less initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 10$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

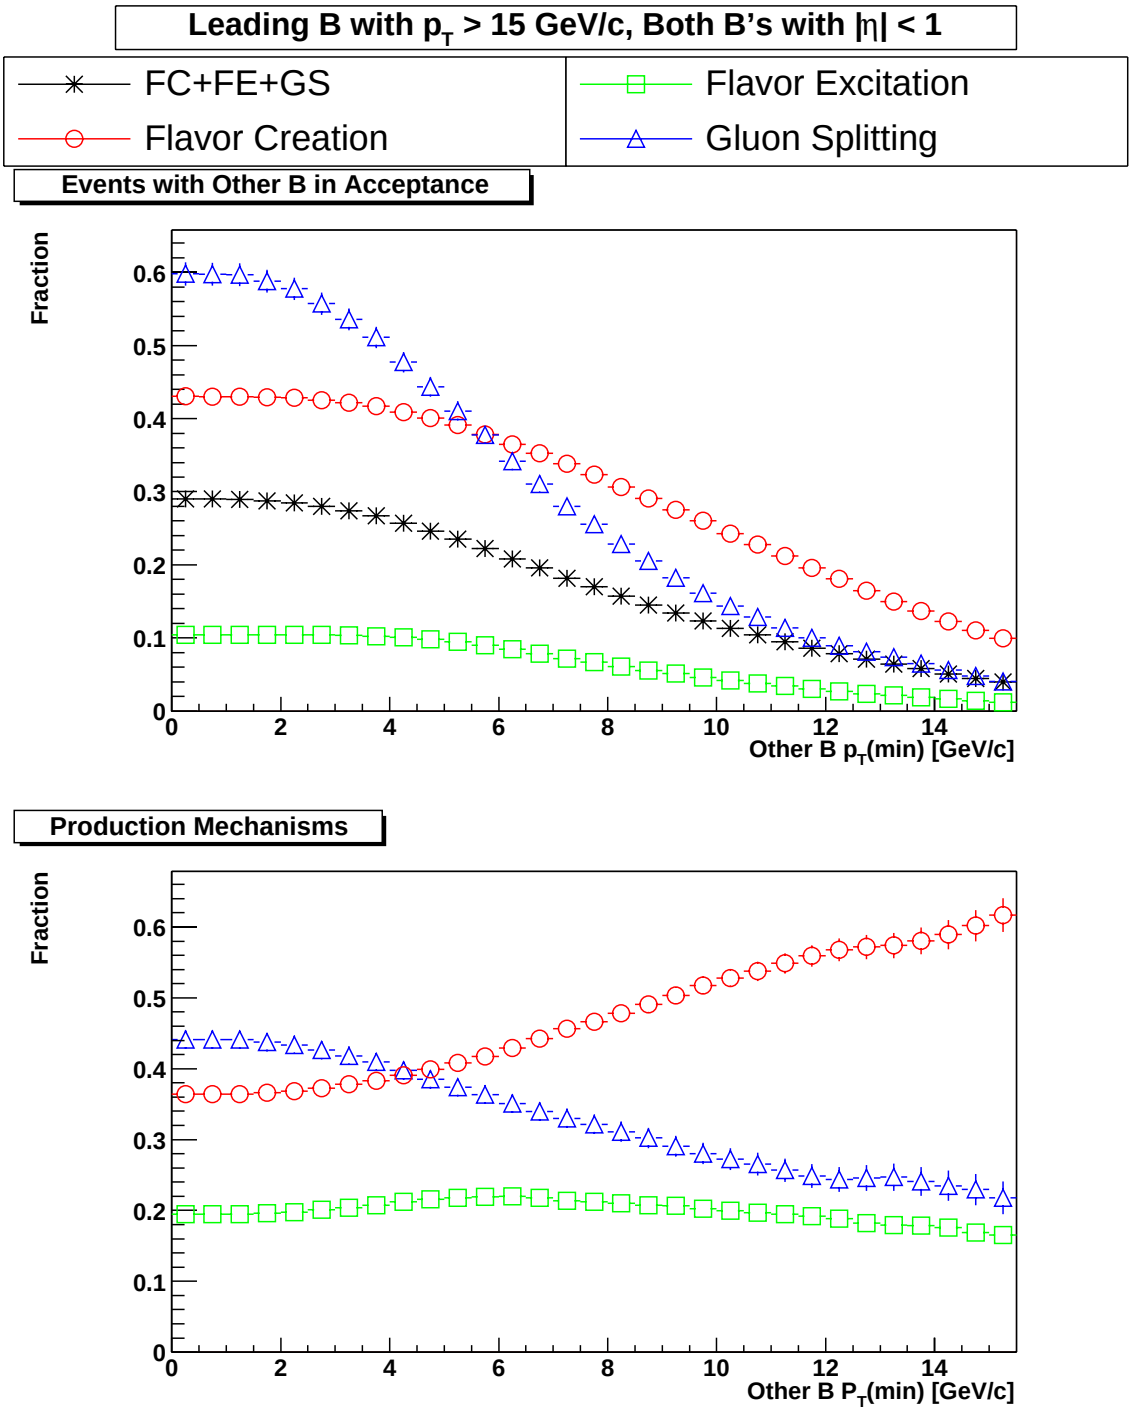

Figure A.9 shows plots made from HERWIG. The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 15$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

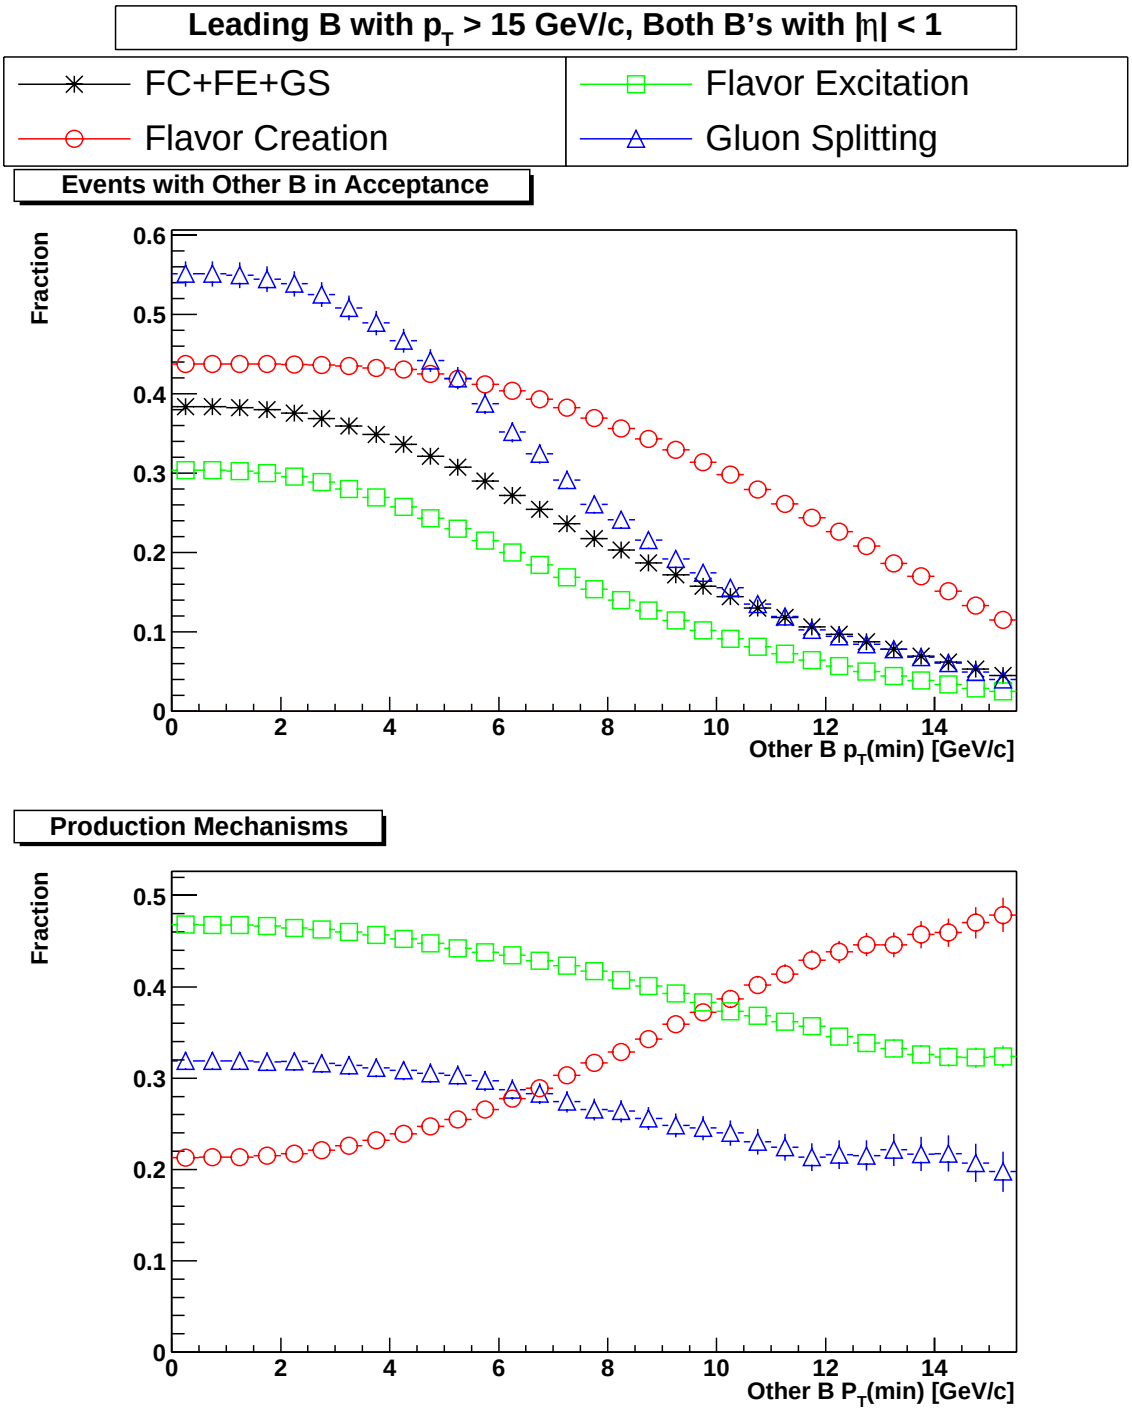

Figure A.10 shows plots made from PYTHIA with `PARP(67)` = 4.0 (more initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 15$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

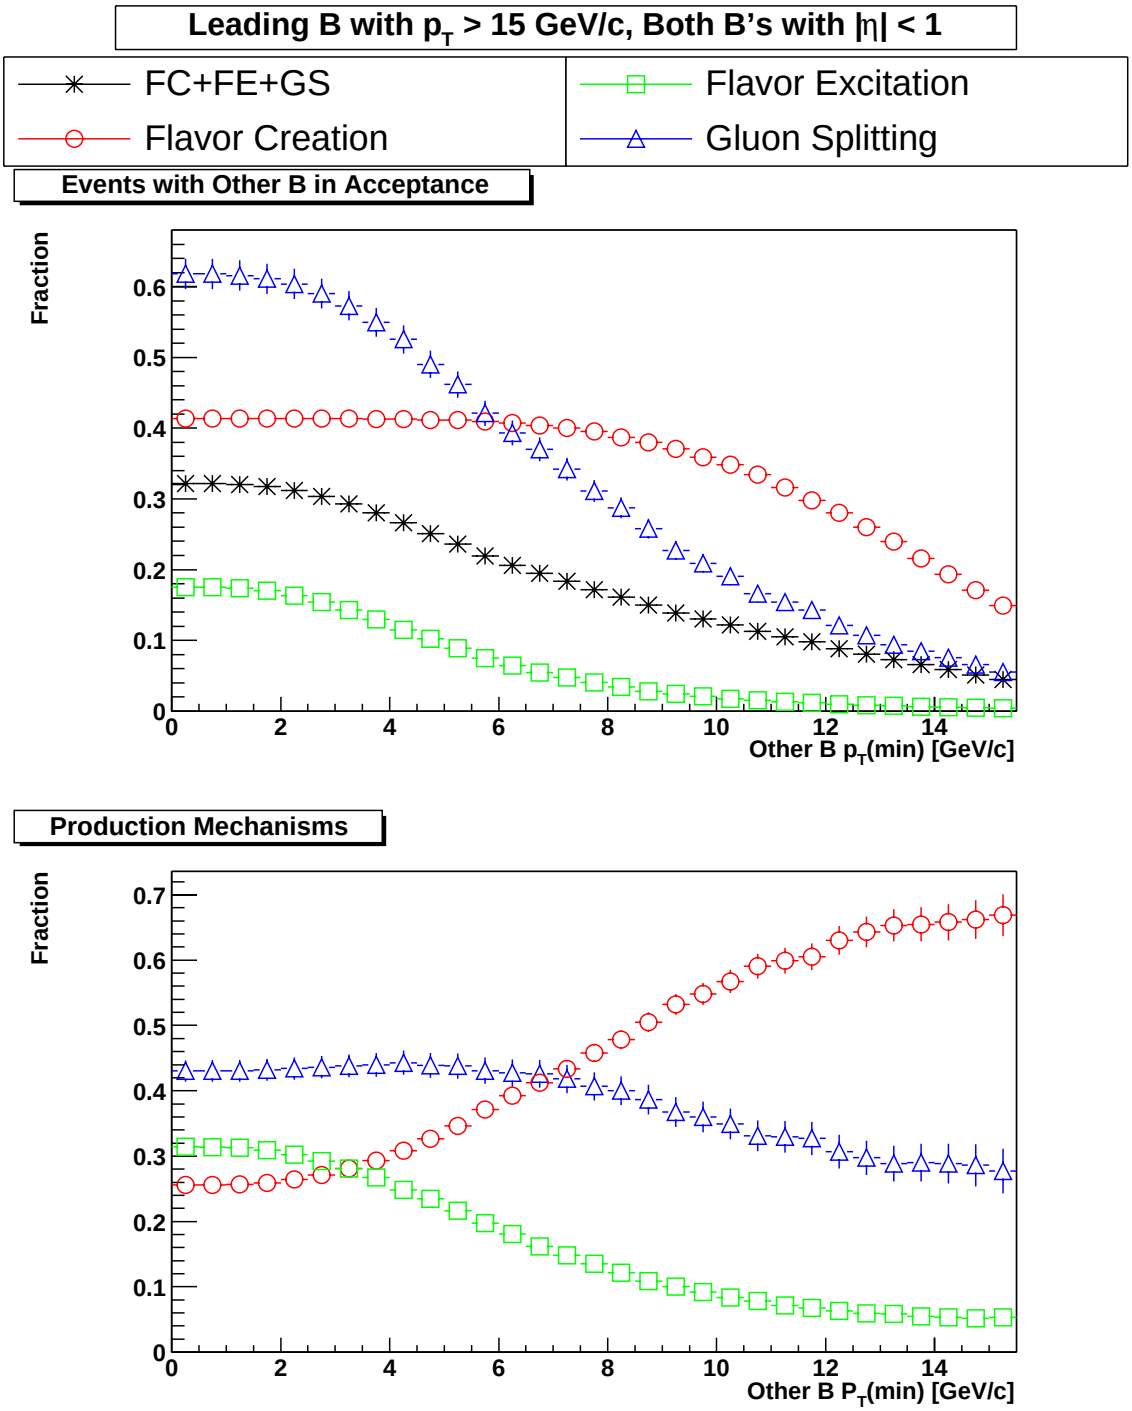

Figure A.11 shows plots made from PYTHIA with `PARP(67)` = 1.0 (less initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 15$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

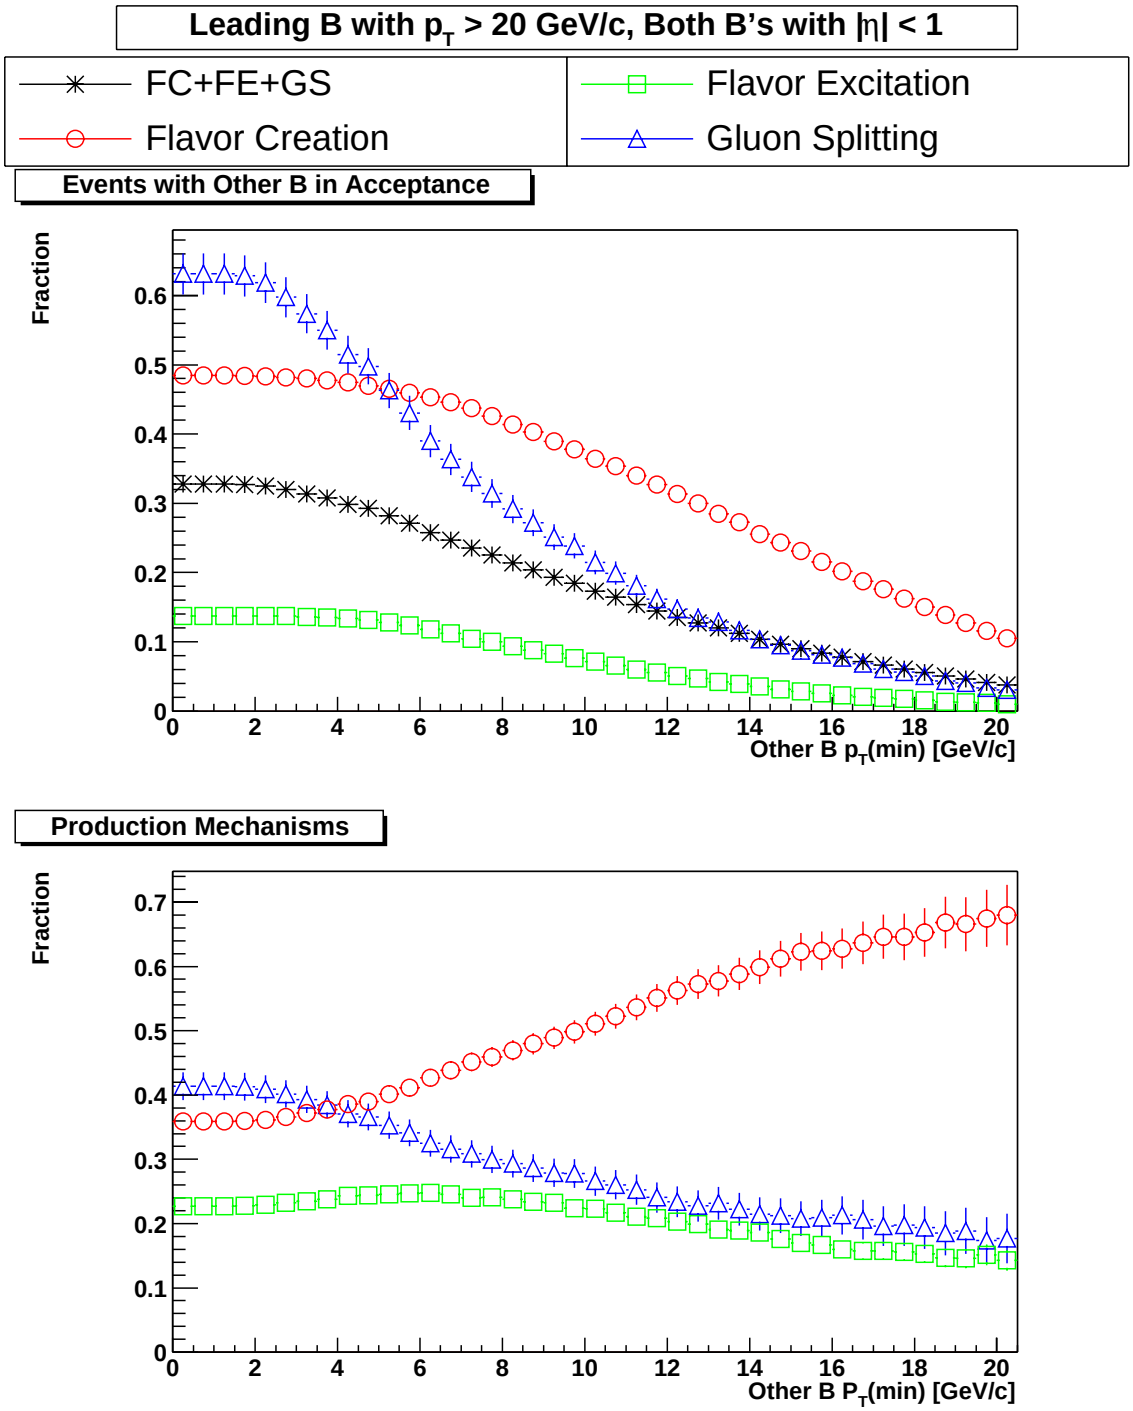

Figure A.12 shows plots made from HERWIG. The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 20$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

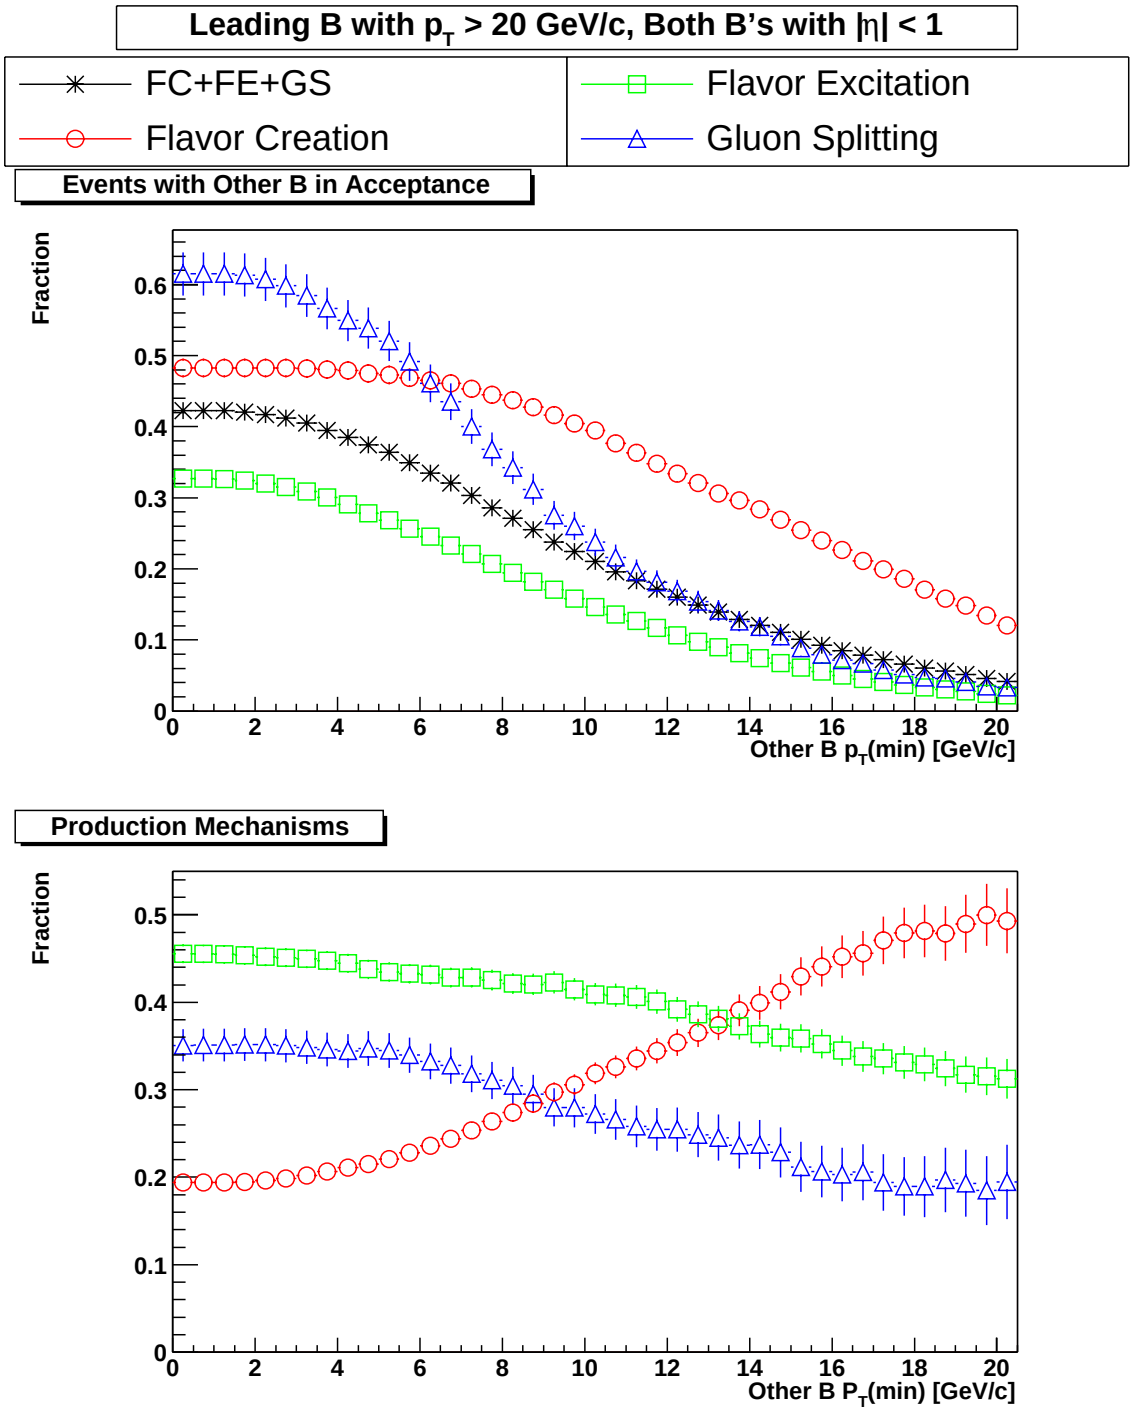

Figure A.13 shows plots made from PYTHIA with `PARP(67)` = 4.0 (more initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 20$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

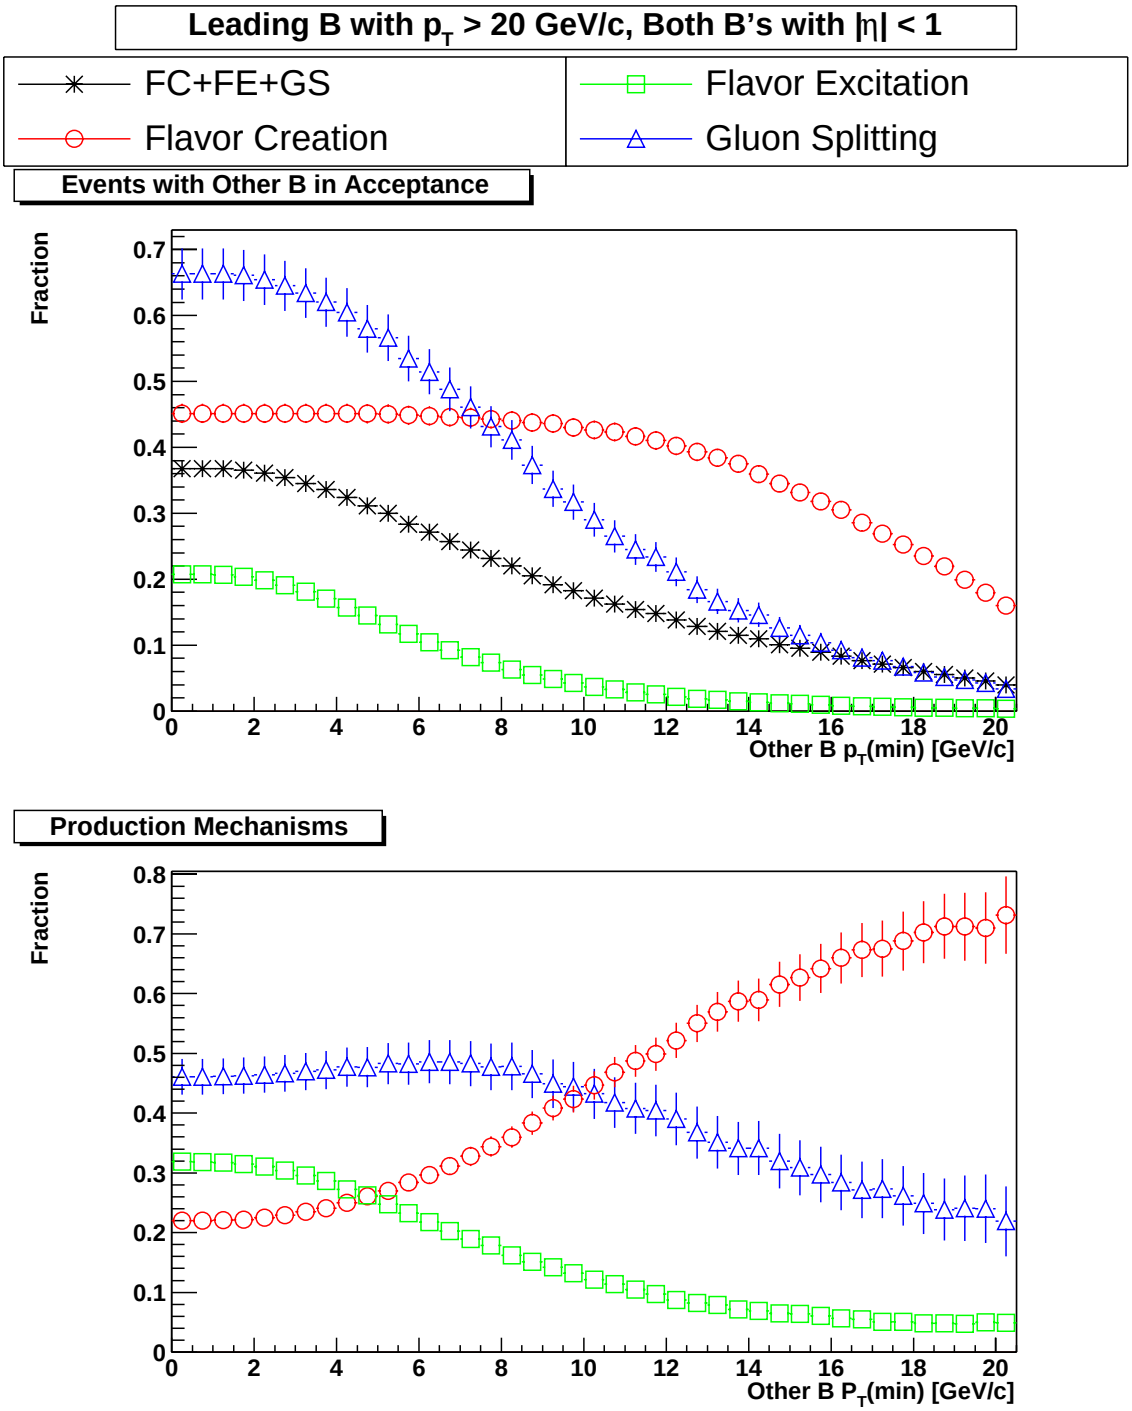

Figure A.14 shows plots made from PYTHIA with `PARP(67)` = 1.0 (less initial-state radiation). The upper plot shows the fraction of the time that an event in which the leading B hadron in the central region ($|\eta| < 1$) with $p_T > 20$ GeV/c also has a second B in the central region, as a function of the p_T of the other B. The lower plot shows the fraction of the events in the acceptance from each production mechanism.

attempt was made to study the fraction of the events in which more than one additional B was present in the acceptance as a function of B p_T; however, statistics were poor. The general trend was that approximately 0.5% of the events with one additional B in the acceptance had more than one B in the acceptance. In some cases, especially at higher p_T cutoffs for the trigger B, this number rose as high as approximately 1%.

Figure A.3 through Figure A.14 also contain plots showing the fraction of the events with at least two B hadrons in the central region satisfying the thresholds p_T(trig) and p_T(min) from each production mechanism. At any point on the curves, the sum of all three fractions should be one. Knowing the composition of the events in the detector acceptance can be useful since the different production mechanisms tend to produce events with different topologies and correlations. For example, the flavor creation process tends to produce b quarks with large azimuthal opening angle and symmetric p_T. Flavor excitation events tend to have a flat azimuthal opening angle distribution and asymmetric p_T and $\Delta\eta$ distributions. Gluon splitting events also have a flat opening angle distribution at low p_T , but as p_T increases, the opening angle distribution tends to peak at smaller values.

The following is an example of how to interpret these plots: Suppose we have a B physics sample for which the trigger was sensitive to B hadrons in the central region of CDF with a minimum p_T of 10 GeV/c. The relevant plots from this note would be Figure A.6, Figure A.7, and Figure A.8. Suppose we also have a method for opposite-side tagging which is sensitive to B hadrons with at least 5 GeV/c p_T. If we wanted to estimate the fraction of the events in our B physics sample for which our algorithm would find at least two B hadrons in the central region, we would read off the value from the above plots corresponding to an "other B p_T(min)" of 5 GeV/c. In this case, we would get a value of ~17% for HERWIG, ~23% for PYTHIA with `PARP(67)` = 4.0, and ~17% for PYTHIA with `PARP(67)` = 1.0. The HERWIG and PYTHIA with `PARP(67)` =1.0 samples would be dominated by approximately equal amounts of flavor creation and gluon splitting, with flavor creation having the slightly larger fraction. For the PYTHIA sample with `PARP(67)` = 4.0, flavor excitation is the dominant contribution. If one were to model this hypothetical B physics sample with only flavor creation Monte Carlo, one would get acceptance values of ~28% for HERWIG, ~34% for PYTHIA with `PARP(67)` = 4.0, and ~36% for PYTHIA with `PARP(67)` = 1.0.

A.3 Conclusions

It is important to have good models of $b\bar{b}$ production for B physics involving flavor tagging. Although qualitatively similar, HERWIG and PYTHIA do not agree completely on all details of $b\bar{b}$ production. In fact, conclusions drawn from PYTHIA alone vary with the amount of initial-state radiation included in the simulation. Therefore, it may be necessary to evaluate more than one model of B production to understand the uncertainty coming from the $b\bar{b}$ production model. Also, in general, flavor excitation and gluon splitting make significant contributions to $b\bar{b}$ production at the Tevatron. Care should be taken not to assume these processes may be neglected without carefully evaluating their contribution. For example, this study shows that using flavor creation only, rather than the full set of production mechanisms leads to an overestimate of the fraction of events containing at least two B hadrons in the acceptance by approximately a factor of 1.5 to 2, depending on the Monte Carlo. This is important to realize because opposite-side flavor tagging techniques are ineffective on events in which both B hadrons don't decay in the detector. In some cases, the lack of a second B in the detector will cause the flavor tagging technique to fail to return any result, such as when there is no lepton from a semileptonic B decay because the second B decayed too far forward to be detected. In other cases, the flavor tagger may return a wrong or random answer if the second B is not in the detector acceptance. For example, in jet charge tagging, if the second B does not decay in the detector, a jet unrelated to the production of either B may be chosen for the tagging. In this case, the result is not likely to be correlated at all with the actual flavor of the B in question.

Having a good model of $b\bar{b}$ production also allows the design of more effective flavor tagging strategies. For example, opposite-side tagging techniques would be most effective in a sample of B hadrons coming primarily from flavor creation, because these events would be more likely to contain two well-separated B hadrons in the detector acceptance. It may be possible to enhance the contribution from flavor creation relative to flavor excitation and gluon splitting by only considering events containing two jets that more or less balance one another in p_T, since flavor excitation and gluon splitting should contain more than two jets. Using a model of $b\bar{b}$ production that more or less accurately replicates the correlations in the data due to higher-order production allows one to evaluate whether the loss in statistics by choosing cuts to enhance the relative contribution from flavor creation is justified by the gain in tagging efficiency and reliability.

Appendix B

Multiple-Tag Events

As seen in Figure 4.9, there are a number of events in our data sample that have three or more tags. While it is certainly true that many of these events contain one or more mistags, possibly combined with genuine bottom or prompt charm tags, in principle, some of these events contain three or more legitimate heavy flavor—that is *B* or *D* hadron decay—tags. These events are potentially very interesting because of the mechanisms for producing multiple heavy flavor quark pairs.

Through higher-order processes, it is possible to generate multiple heavy flavor pairs in high-energy QCD interactions, such as those shown in Figure B.1. Typically, multiple heavy flavor production occurs in QCD when a $b\bar{b}$ or $c\bar{c}$ pair generated through the usual mechanisms of flavor creation, flavor excitation, and gluon splitting, is supplemented by another heavy quark pair created through additional gluon splitting. However, it is also possible to have processes like double-flavor excitation, in which a heavy flavor quark found in the parton distributions of

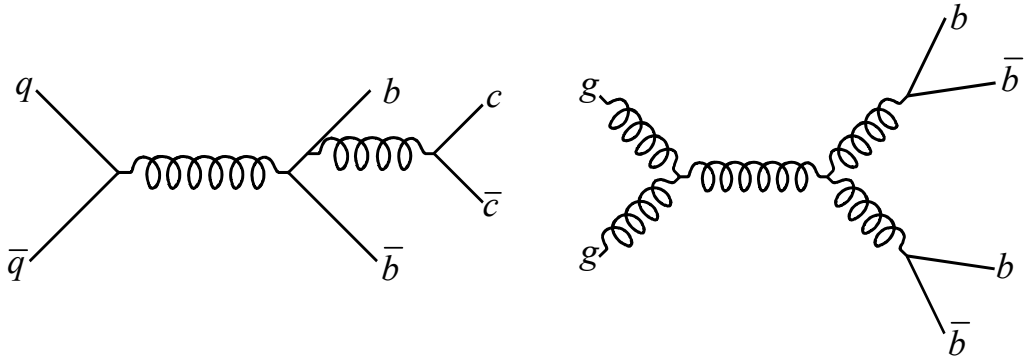

Figure B.1 Representative Feynman diagrams for multiple-heavy flavor production through QCD processes.

each of the two beam particles interact, as shown in Figure B.2. In this case, two of the heavy flavor quarks come from the hard QCD interaction, and the other two are present as part of the initial-state radiation in the event. These processes yield events containing multiple pairs of heavy flavor quarks, such as $b\bar{b} + b\bar{b}$, $b\bar{b} + c\bar{c}$, or $c\bar{c} + c\bar{c} + c\bar{c}$. Such multiple heavy flavor interactions are valuable for testing higher-order perturbative QCD calculations involving heavy flavor. In addition, these processes serve as backgrounds for more exotic physics with multiple heavy flavor final states.

In addition to QCD production of multiple heavy flavor pairs, there are also possibilities for production mechanisms involving new physics. For example, one possible signature of the Higgs boson at a hadron collider is the production of a Higgs in association with a pair of bottom quarks [78]. If, in addition, the Higgs boson decays to pair of bottom quarks, this yields a Higgs signature that could be detected by searching for events containing three or more secondary vertex tags. Production of supersymmetric particles provides another avenue to create multiple heavy flavor pairs. Supersymmetry is an extension to the Standard Model that postulates that every Standard Model fermion and boson is matched to a *superpartner* differing from the regular Standard Model particles by one-half unit of spin [17]. For example, all spin-1/2 quarks are matched to spin-0 *squarks*, while the superpartner of the spin-1 gluon is the spin-1/2 gluino. Currently, there are a number of different scenarios for the realization of supersymmetry that have not been ruled out by experiment. Some of these models include particles light enough to be produced copiously at colliders like the Tevatron, such as Berger's light gluino hypothesis [18], mentioned in Section 1.3. Depending on the model, gluinos produced at the Tevatron can

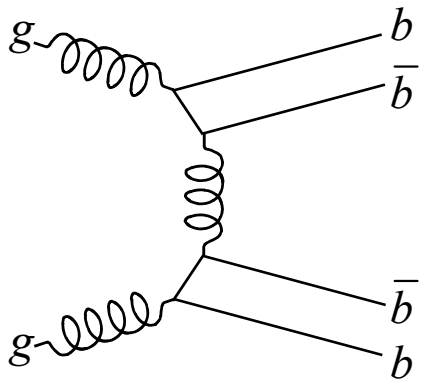

Figure B.2 Double flavor excitation.

decay to various combinations involving top, bottom, or charm quarks as well as their superpartners, stop, and sbottom [79]. In addition, depending on the model, these stop and sbottom squarks can decay to bottom and charm quarks. Therefore, many of these theories could be investigated through a study of events containing several secondary vertex tags.

The sample of events with three or more secondary vertex tags is unfortunately, not large enough to allow detailed studies of any of the multiple heavy flavor mechanisms described above. Nonetheless, we detail some of the properties of multi-tag events here as a reference point, for further investigation into the subject, perhaps with a higher-statistics Run II dataset.

B.1 Mistag Subtraction

Just as for two tag events, mistag subtraction is necessary to remove the background from events containing one or more mistags. As in Section 4.4.1, mistag subtraction for events with three or more BVTX tags relies on the numbers of events with different combinations of tags with positive and negative L_{xy}. However, as the number of tags grows, the number of possible combinations of positive and negative L_{xy} tags quickly becomes unmanageable. Therefore, to determine the mistag subtraction formulae for events with more than three tags, we take a somewhat different approach than the one used in Section 4.4.1.

Consider first the case in which events have a single secondary vertex tag. Suppose the probability that a good tag will have $L_{xy} > 0$ is p and the probability that a mistag will have $L_{xy} > 0$ is q. The number of tags with positive (negative) L_{xy} is given by N_+ (N_-), while the number of good tags (mistags) is given by N_G (N_M). Then we have

$$\begin{aligned} N_+ &= pN_G + qN_M \\ N_- &= (1-p)N_G + (1-q)N_M \end{aligned}, \tag{B.1}$$

which can be solved to yield

$$N_G = \frac{(1-q)}{(p-q)} N_+ - \frac{q}{(p-q)} N_- . \tag{B.2}$$

Now, suppose we have events with two tags. To find the number of events in which both tags are good, N_{GG}, we simply need to use Eq. (B.2) above to subtract the number of events in which the first tag is good and the second tag has positive L_{xy} from the number of events in which the

first tag is good and the second tag has negative L_{xy}, where the number of events in which the first tag is good is also give by Eq. (B.2):

$$\begin{aligned} N_{GG} &= \frac{(1-q)}{(p-q)} N_{G+} - \frac{q}{(p-q)} N_{G-} \\ &= \frac{(1-q)}{(p-q)}\left[\frac{(1-q)}{(p-q)} N_{++} - \frac{q}{(p-q)} N_{-+}\right] - \frac{q}{(p-q)}\left[\frac{(1-q)}{(p-q)} N_{+-} - \frac{q}{(p-q)} N_{--}\right]. \\ &= \frac{(1-q)^2}{(p-q)^2} N_{++} - \frac{q(1-q)}{(p-q)^2}\left(N_{+-} + N_{-+}\right) + \frac{q^2}{(p-q)^2} N_{--} \end{aligned} \quad \text{(B.3)}$$

The above equation reduces to the familiar two-tag mistag subtraction formula from Section 4.4.1 with $p = 1.0$ and $q = 0.5$.

Through another application of Eq. (B.2) to Eq (B.3) above, the number of three tag events in which all tags are good tags can be obtained from the data. In this case, it is useful to define the notation N_{i+}, in which the index i denotes the number of positive L_{xy} tags. For example, for three-tag events, N_{2+} are all possible tag combinations with two having $L_{xy} > 0$ and one having $L_{xy} < 0$, namely $N_{2+} = N_{++-} + N_{+-+} + N_{-++}$. Using this notation, the number of three-tag events with all good tags is given by

$$N_{GGG} = \frac{(1-q)^3}{(p-q)^3} N_{3+} - \frac{q(1-q)^2}{(p-q)^3} N_{2+} + \frac{q^2(1-q)}{(p-q)^3} N_{1+} + \frac{q^3}{(p-q)^3} N_{0+}. \quad \text{(B.4)}$$

In fact, through repeated application of Eq. (B.2), it becomes clear that the general formula for n tags is

$$N_{nG} = \frac{1}{(p-q)^n} \sum_{i=0}^{n} (-q)^{n-i} (1-q)^n N_{i+}. \quad \text{(B.5)}$$

Table B.1 shows a breakdown of the combinations with different numbers of positive L_{xy} tags for events with three or more tags, as well as the statistical uncertainty on this number. The numbers in the table assume $p = 1.0$ and $q = 0.5$. For events with more than two tags, it is clear that only the three tag events have a statistically significant number of events with all good tags. For three tag events, these calculations indicate that there are 142 ± 39.4 (stat.) events in which all three tags are not mistags. In the next section, we examine these events in greater detail. If we vary

Number of positive tags						Total	Mistag Subtracted	Stat. Error
0	**1**	**2**	**3**	**4**	**5**			
Two-Tag Events								
1570	10632	25417	--	--	--	37619	16355	194.0
Three-Tag Events								
38	290	669	559	--	--	1556	142	39.4
Four-Tag Events								
1	6	23	27	7	--	64	-2	8.0
Five-Tag Events								
0	0	1	0	1	0	2	-2	1.4

Table B.1 The break-down of the number of tag pairs with different combinations of positive and negative L_{xy}. The first five columns show the number of tags with zero, one, two, etc. number of positive tags. The "Total" column gives the total number of events with the indicated number of tags, regardless of the sign of L_{xy}. The "Mistag Subtracted" column indicates the numbers of events in which all tag pairs are good, based on the mistag subtraction formulae derived in this chapter. The final column shows the statistical error on this mistag subtracted number of events.

the p and q values within the reasonable range determined in Section 7.2.1, the number of good three tag events changes within the range of 105.2 to 218.9.

B.2 Distributions

Figure B.3 shows the separation between pairs of tags in the three-tag events in the plane perpendicular to the beam, ΔL_{xy}. This quantity is illustrated in Figure B.4. The two dimensional separation is calculated for the three possible combinations of the tag pairs and the plots show the minimum separation on the y-axis and the maximum separation on the x-axis. Also shown is the significance of the separation which is defined as the two-dimensional distance between a pair of tags divided by the uncertainty on that distance, $\Delta L_{xy}/\sigma_{Lxy}$. These plots show that after mistag subtraction, most of the events have a minimum separation between pairs of tags of at least 10σ.

In Figure B.5, we plot quantities involving pairs of tags. The upper plots show the minimum $\Delta\phi$ separation between any two of the three tags versus the maximum $\Delta\phi$ separation. Most tags

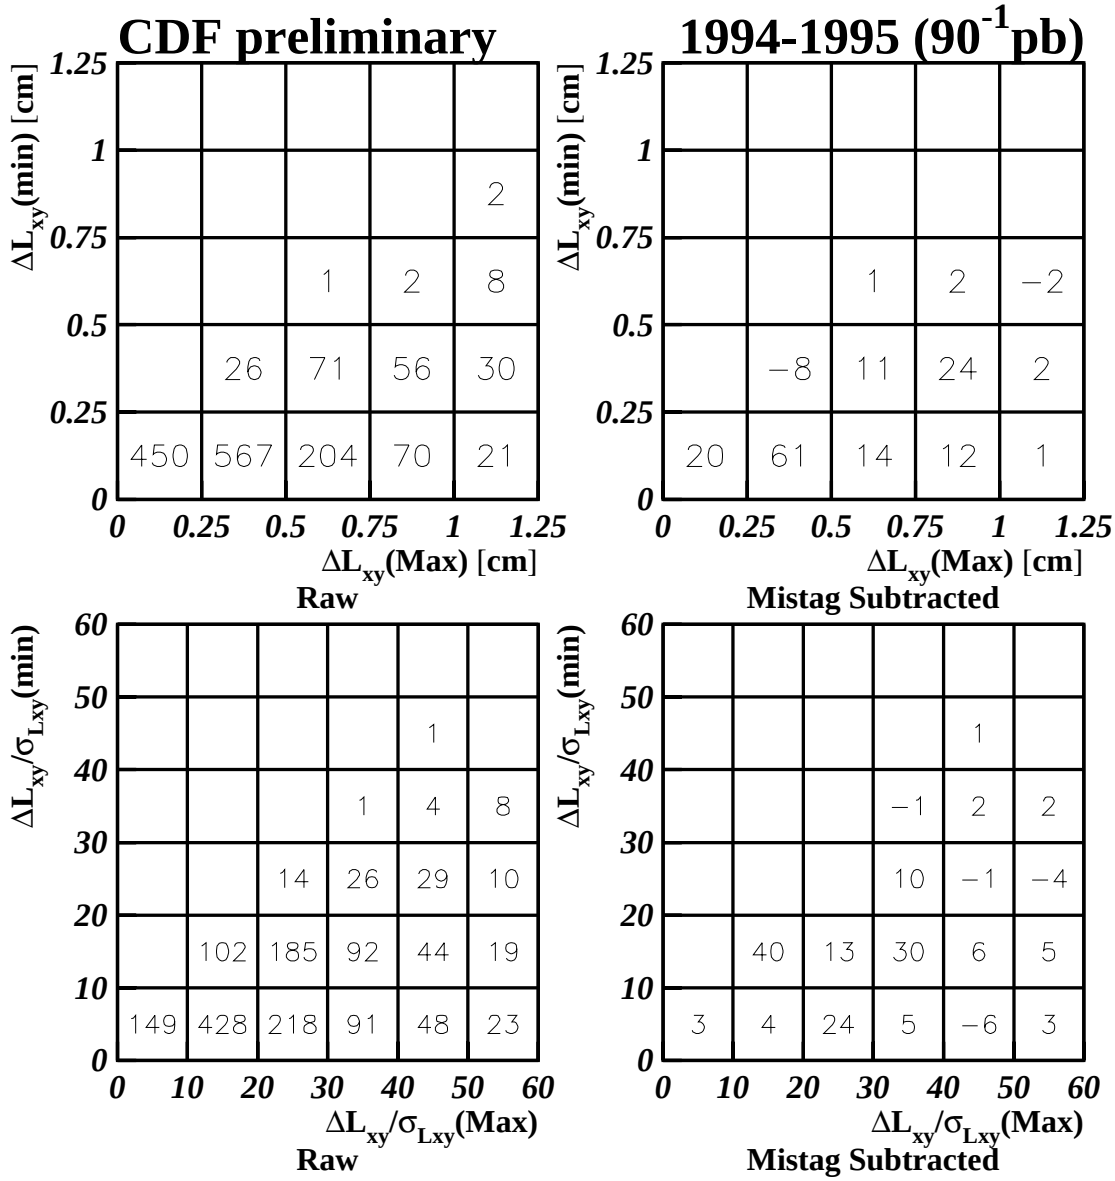

Figure B.3 The 2D separation between tag pairs, as defined by Figure B.4, before and after mistag subtraction. The numbers on the plot show the number of events in each bin.

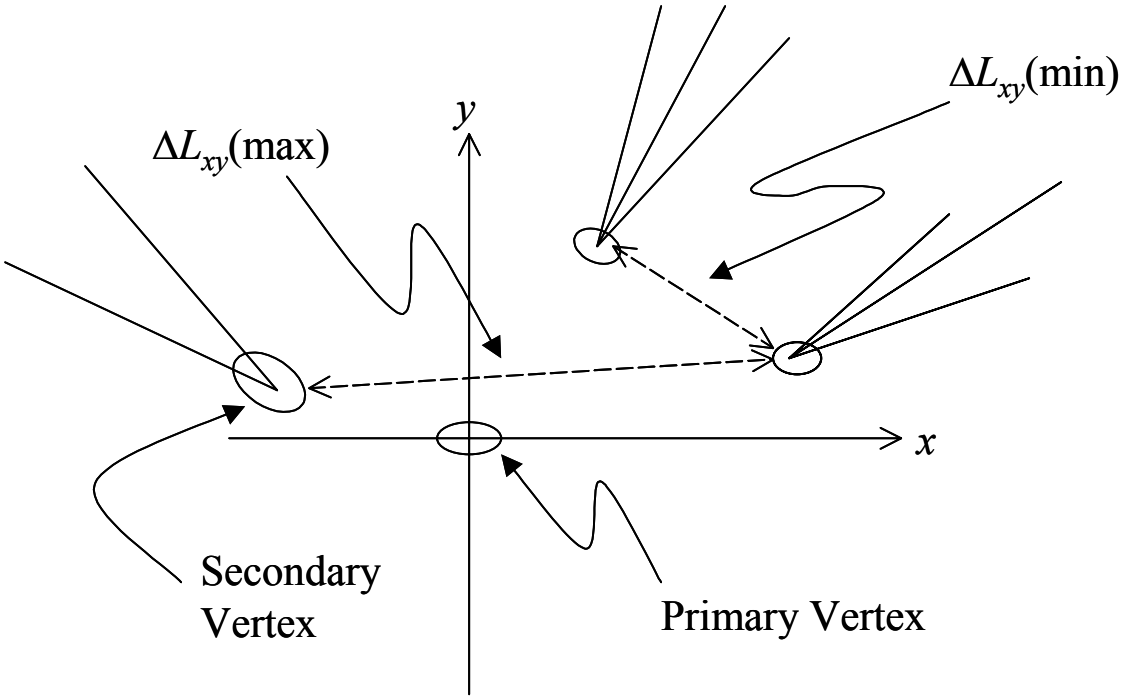

Figure B.4 shows the definitions of ΔL_{xy}(min) and ΔL_{xy}(max) from Figure B.3.

appear in the region where there is a large $\Delta\phi$ separation between two tags and a small angle between the third tag and one of the two back-to-back ones. However, there are a small number of in which the third tag is more than 60° away from either of the back-to-back tags. The bin in which all three tags are near to one another (no tag more than 60° from any other) shows no statistically significant excess over the mistag background.

Figure B.5 also shows the jet masses for pairs of tags. The jet mass is figured from the track-based jets used to organize the BVTX tagging algorithm, as described in Section 4.2.2. Since the jets are based on tracks, only information about the charged-particle component of the jets are available. For a pair of tags, the jet mass is considered to be the invariant mass of the jet or jets containing the tags. In the case that both tags come from the same jet, only information from this jet is used to calculate the mass. The plots show the minimum jet mass for any pair of tags versus the maximum jet mass for any pair of tags.

The bin in the $\Delta\phi$ plots in which one pair of tags has $\Delta\phi > 120°$, while the third tag is at least 60° from either of the other two is an interesting region. In this case, all three tags have good angular separation. Figure B.6 shows the jet masses for these events calculated in three different

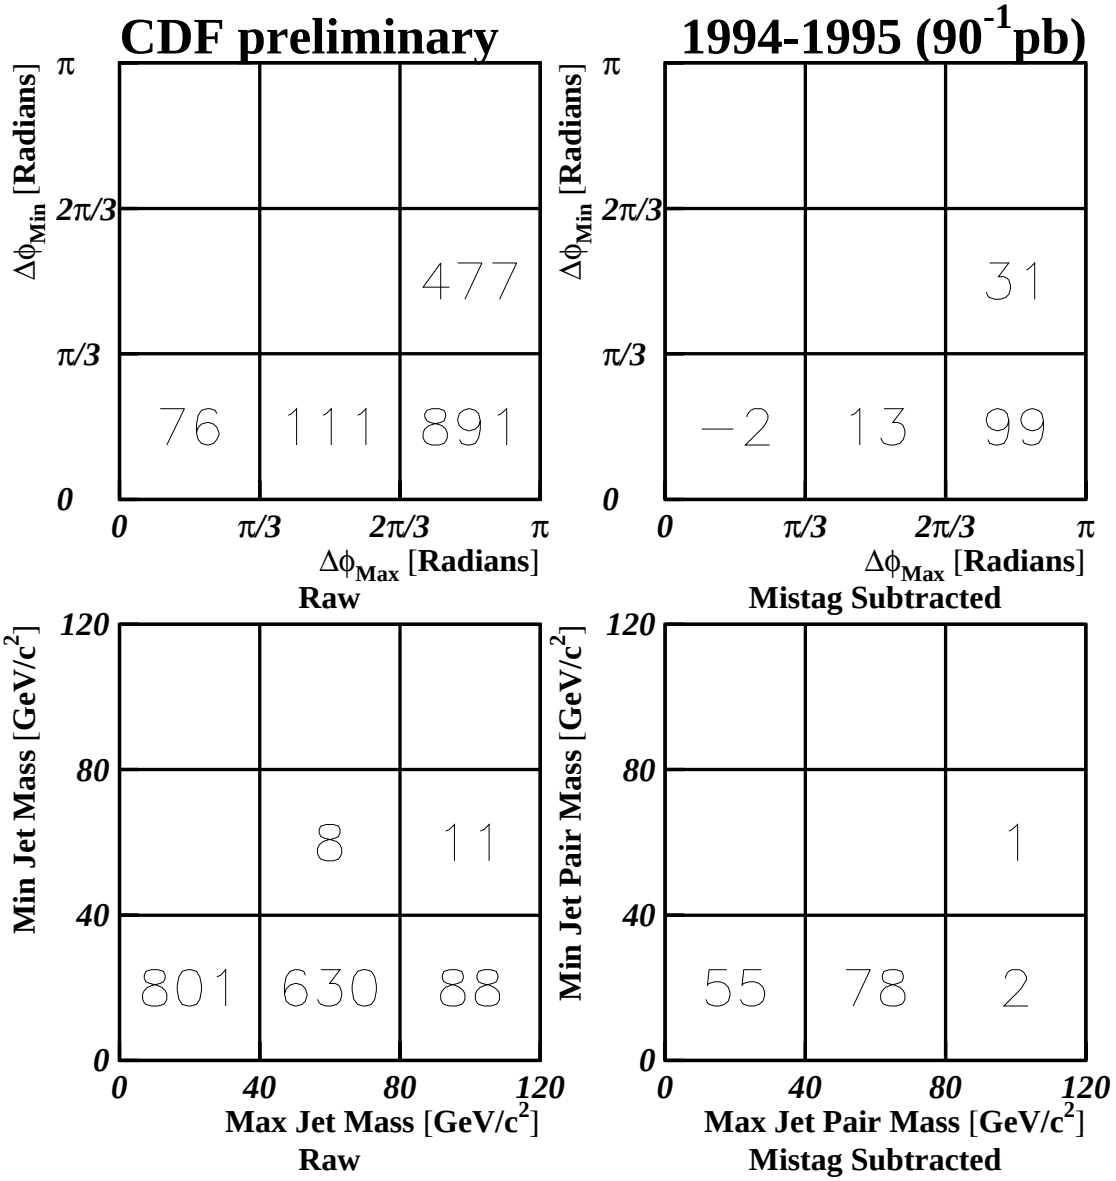

Figure B.5 The angular separation between pairs of tags in three-tag events, as well as the jet mass for tag pairs. The jet mass is defined as the invariant mass of the track-based jet or jets containing both tags in the tag pair. Since the jet mass is based on tracking quantities, it only contains information about the charge-particle component of the jet. The numbers in the plots show the number of events in each bin.

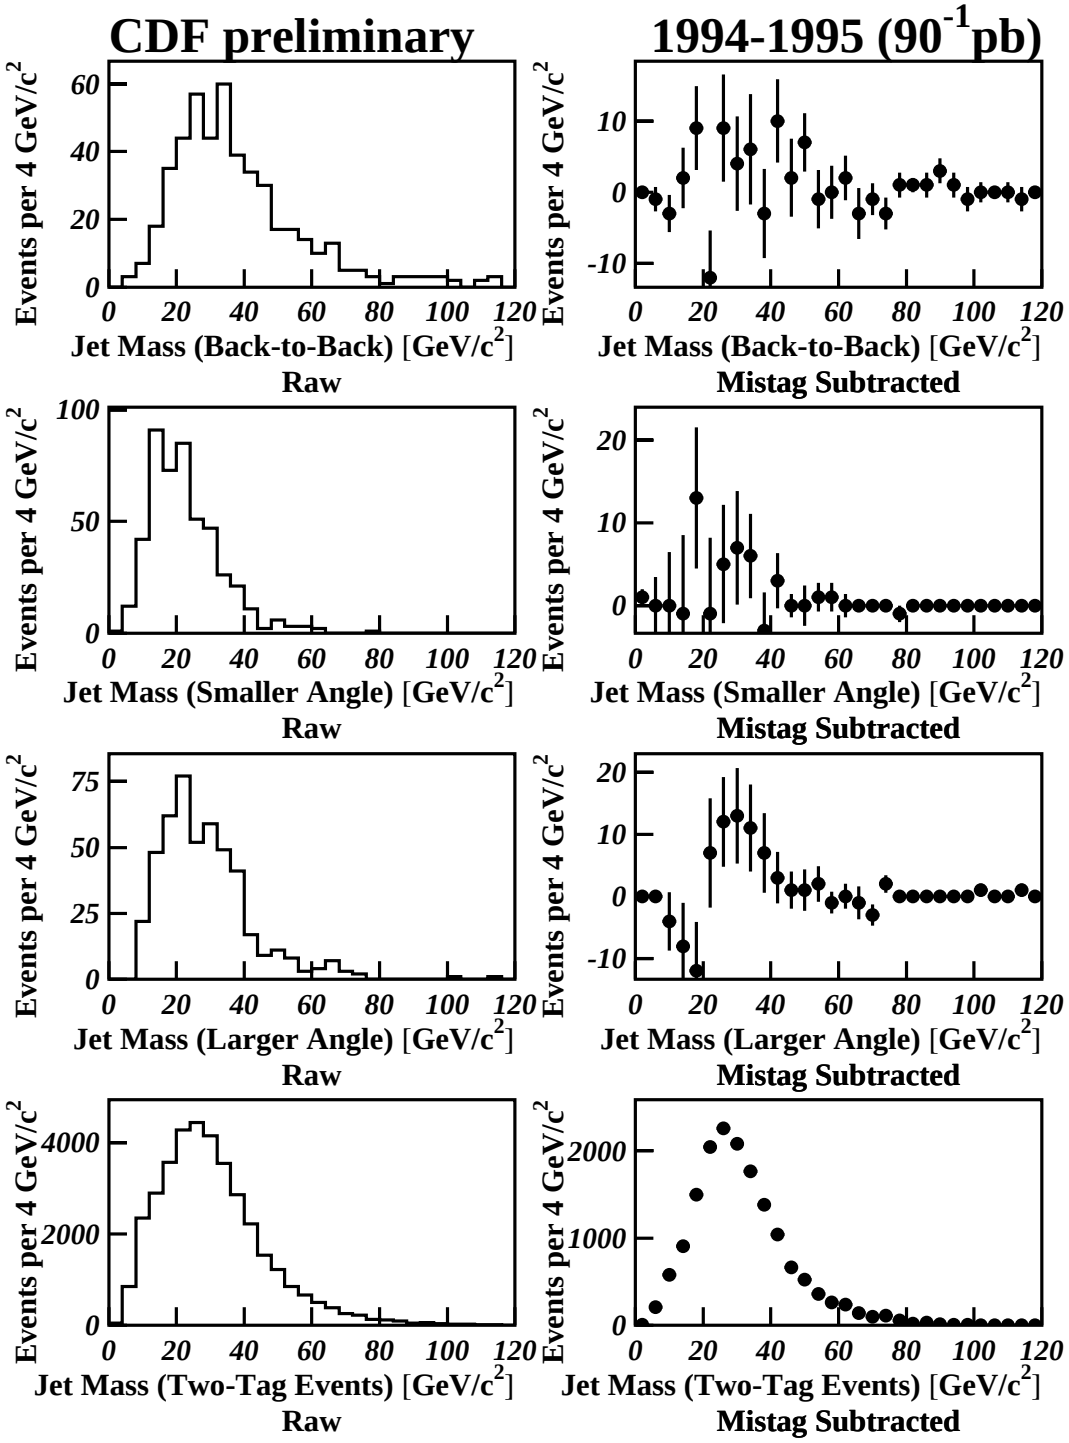

Figure B.6 The jet mass for different combinations of tag pairs, as shown in Figure B.7, for events with one pair of tags that is relatively back-to-back ($\Delta\phi > 120°$) and a third tag that is relatively well separated from the other two ($\Delta\phi > 60°$). The bottom plots show the jet mass distribution for two-tag events.

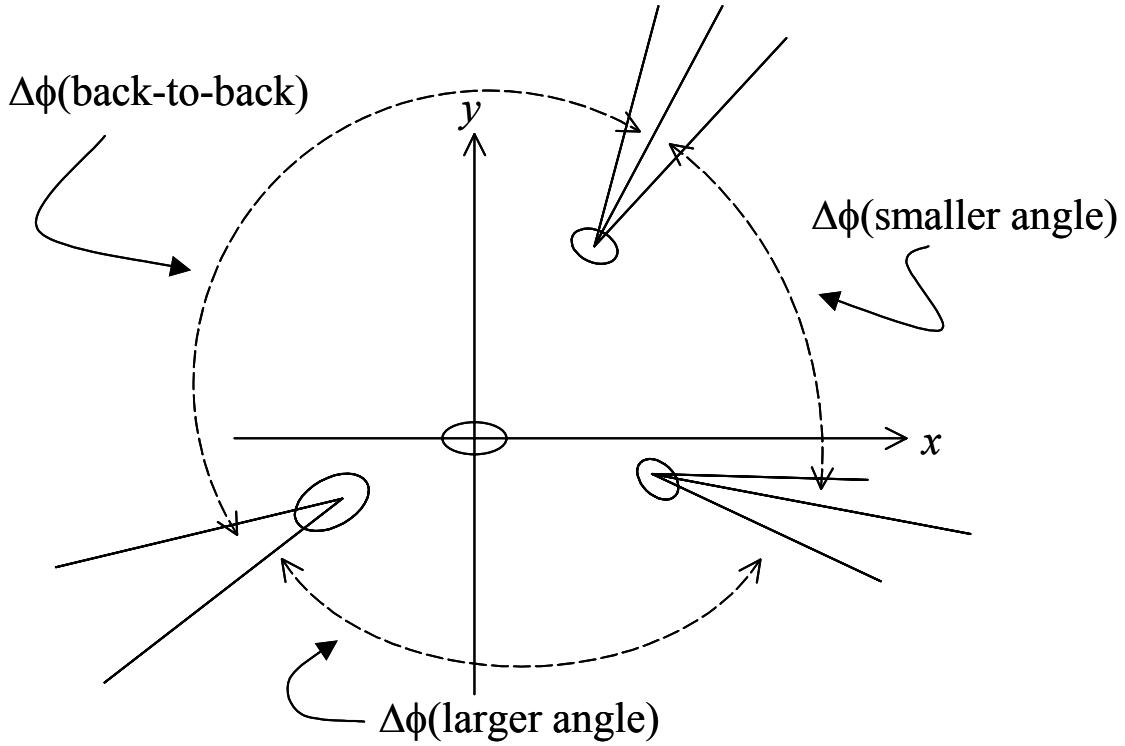

Figure B.7 The tag-pair configurations used for Figure B.6.

tag pair combinations: the back-to-back pair ($\Delta\phi > 120°$), as well as the third tag paired with the tag giving the smallest opening angle and the largest opening angle. These tag pair configurations are sketched in Figure B.7. The last plot in Figure B.6 shows the jet mass distribution for two-tag events for comparison.

In Figure B.8, comparisons are made between distributions involving three-tag events and those involving only two tags. The upper plots in the figure show the number of track-based jets that are present in three-tag events compared to two-tag events. After mistag subtraction, it appears that the three tag events may favor a slightly higher number of jets per event, although the statistical uncertainty makes this difficult to discern. Additionally, the scalar sum of the p_T for all tracks in the event with $p_T > 1$ GeV/c is compared between three- and two-tag events. Again, the three-tag events seem to favor a slightly higher p_T although it is difficult to tell for this size sample. Because the requirement for three tags may simply bias our selection towards events with more jets and more tracks, giving more opportunities to find tags, mistag or otherwise, we also look at the scalar sum of track p_T divided by the number of tracks in the event.

This quantity compares the average p_T per track for three-tag and two-tag events, as shown in Figure B.8. On average, two-tag events have approximately 3.6 GeV/c per track, while three-tag events have 4.1 GeV/c.

The plots in this section give a rough idea of the characteristics of three-tag events in this data sample. Although the statistics are not sufficient to attempt to understand the composition of this three-tag sample, the plots here give some rudimentary insight into the nature of the three-tag data collected during Run 1.

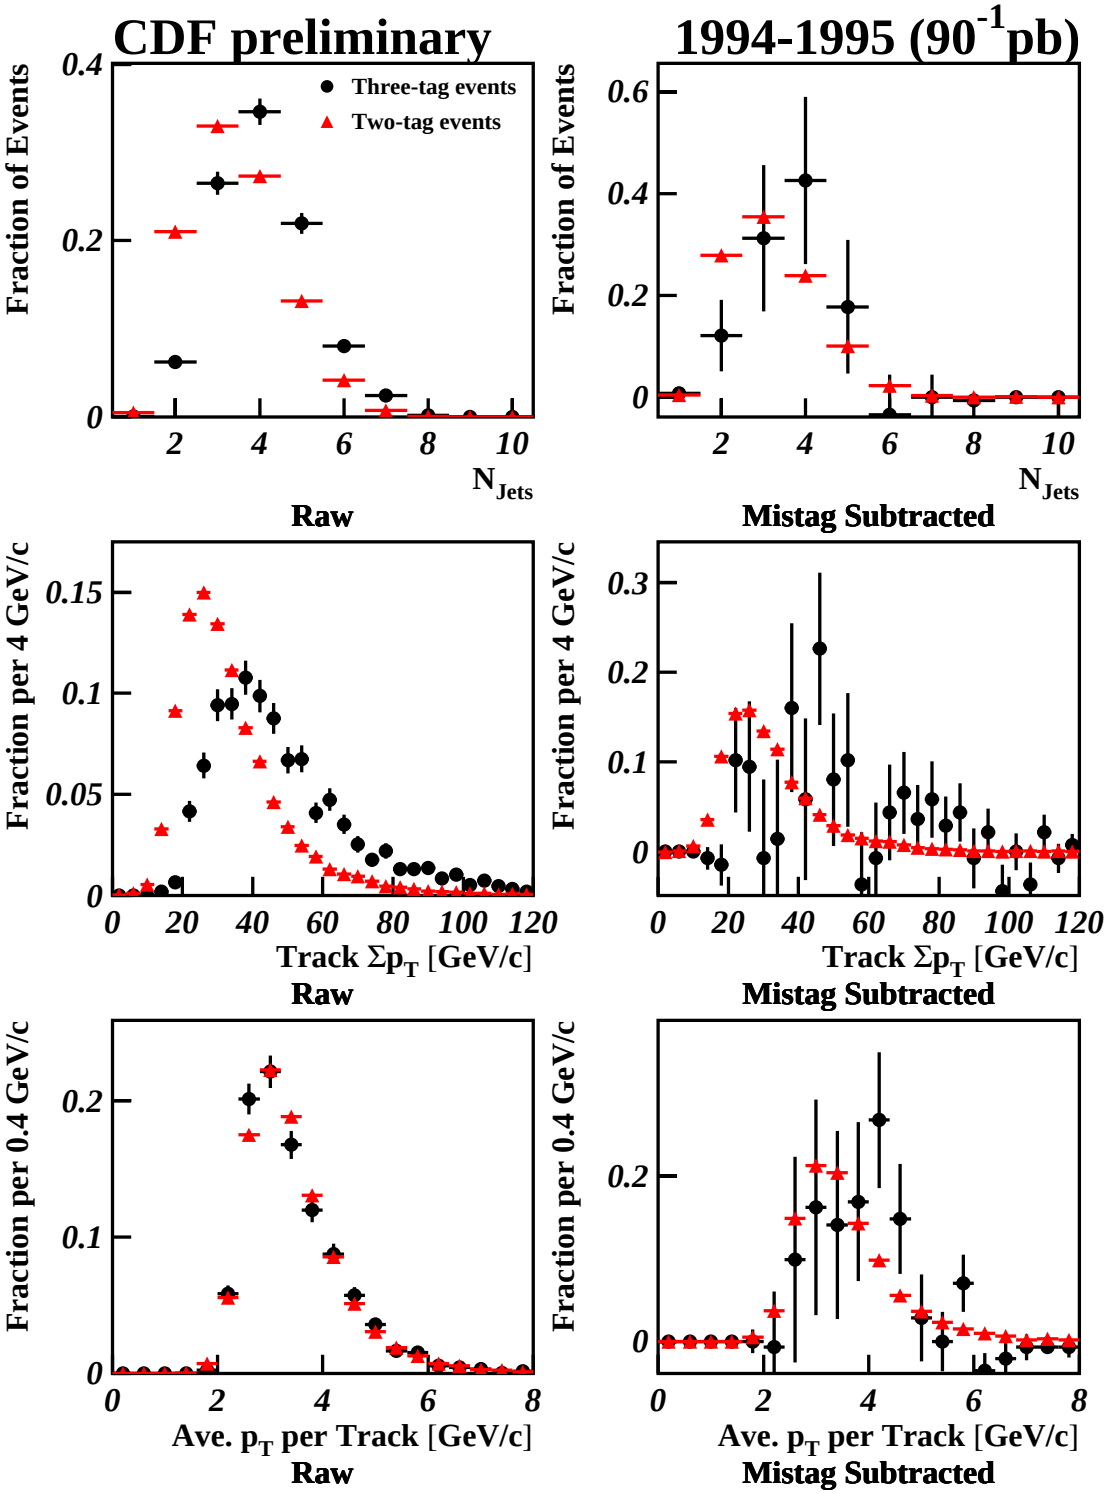

Figure B.8 A comparison of the number of jets per event, the scalar sum of the track p_T, and the average p_T per track for three-tag (black circles) and two-tag (red triangles) events.

Appendix C

Muon Calibration System for Run II

Because the typical cycle of design, construction, operation, and analysis of a given high-energy physics experiment can be longer than the time a graduate student spends obtaining a degree, it is not always possible for a graduate student to analyze data for his or her thesis from an experiment that he or she helped to build and commission. As a result, high-energy physics graduate students frequently find themselves helping to design or upgrade one experiment while analyzing data from an earlier experiment. This thesis details an analysis involving data collected at the CDF detector during the 1994–1995 run, known as Run 1. However, during my time as a graduate student, I also worked on upgrades to the muon system for the Run II CDF detector. This section details work I did on the calibration systems for the Run II muon detectors.

The same principles are applied to muon detection in Run II as were used in Run 1. Muons are identified by their ability to penetrate the CDF calorimeters and additional shielding without losing significant energy. There are two types of muon detectors used at CDF in Run II: chambers and scintillators. Muon chamber detectors consist of arrays of single-wire drift tubes filled with argon-ethane gas. High-voltage is applied to create a potential difference between the anode wire and cathode strips on the chamber walls. A muon (or any other charged particle) that traverses the drift chamber ionizes the argon-ethane gas. The electrons liberated by this ionization are quickly accelerated to terminal velocity by the chamber potential and drift at a constant speed towards the anode wire. Charge reaching the anode wire creates electric signals that are amplified, digitized, and recorded. The time elapsed between the $p\bar{p}$ interaction that created the muon and charge deposition on the muon chamber wire, known as the *drift time*, determines how close the muon trajectory passed to the known location of the muon chamber wire. Signals from the same muon as it passes through several chambers can be used to reconstruct the muon trajectory. Because the time required for the ionized charge to drift to the anode wire, as much as 2 μs, is much longer than the 396 ns $p\bar{p}$ bunch crossing time, there is a

potential ambiguity in associating charge deposition—also known as a hit—in the muon chambers with a particular beam crossing. The muon scintillator detectors can be used to resolve this discrepancy. The muon scintillators are composed on scintillating plastic that emits light when excited by the passage of a muon. This light is collected by photomultiplier tubes (PMT's) that convert the collected light to electrical signals. The muon scintillator detectors have broader granularity than the muon chambers, but much faster response times, making them useful for resolving beam crossing ambiguities for muon hits, reducing noise hits from chambers, and providing fast trigger information.

The muon detector systems used in Run II are for the most part, upgrades of the original Run 1 muon detector systems [80][81]. The configuration of the Run II muon detectors is shown in Figure C.1. The central muon detector (CMU), located at the back of the CDF calorimeter wedges, is composed of a number of rectangular drift cells that cover $|\eta| < 0.6$. Pairs of adjacent CMU chambers are ganged together to allow a determination of the z position of CMU hits using charge division. The Run II CMU detector is largely the same in Run 1, although the chamber

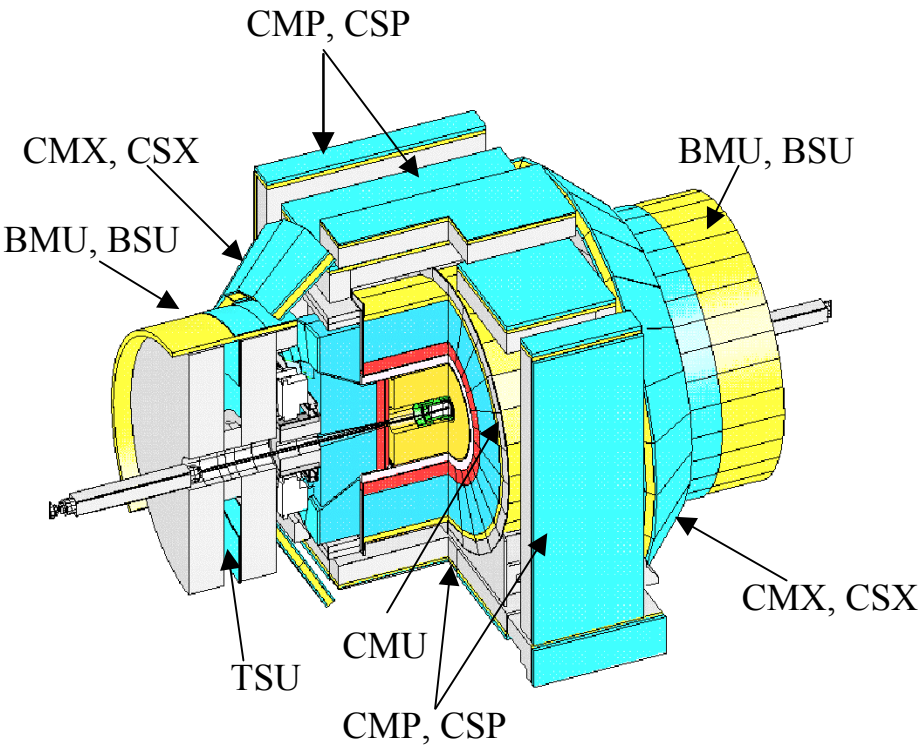

Figure C.1 The muon detector configuration at CDF.

and front-end electronics are upgraded. The central muon upgrade (CMP) detector is composed of a rectangular array of drift tubes that also cover $|\eta| < 0.6$. In addition to upgraded chamber electronics, additional chambers were added to the CMP to close some of the ϕ gaps in CMP coverage, as indicated in Figure C.2. The central muon extension (CMX) detector covers from $0.6 < |\eta| < 1.0$ with a set of rectangular drift chambers arrayed along the surface of a conic section. The primary change in the CMX from Run 1 is the addition of chambers to close some of the CMX ϕ gaps. In addition to the central muon chambers described above, there are also two sets of muon scintillators in the central region for Run II. The central muon upgrade scintillators (CSP) provide scintillator coverage for the CMP chambers. The CSP counters for the CMP walls were installed in Run 1 while the counters for the top and bottom of the detector are new for Run II. The central muon extension scintillators (CSX) sandwich the CMX chambers and remain largely unchanged from Run 1 with the exception of additional scintillators added to cover the new CMX chambers. Replacing the forward muon detector (FMU) of Run 1 is the new IMU detector, which covers the region $1.0 < |\eta| < 1.5$. The IMU detector uses the toroid magnets from the Run 1 FMU detector to provide a support structure and shielding. However, the toroid magnet is not powered up and provides no magnetic field. The IMU consists of a set of rectangular drift chambers in a cylindrical arrangement along the outside of the toroids known as the "barrel" muon detector (BMU). The "barrel" muon scintillators (BSU) are arranged on the outside of the BMU chambers, while the toroidal muon scintillators (TSU) are arranged perpendicular to the beam on the toroid face.

Although different sets of electronics are used to readout the muon chambers and scintillators, there are similarities. For each muon chamber system, the signals generated by charge deposition on the anode wire are amplified by a preamplifier, before being further amplified again and converted into differential emitter-coupled-logic (ECL) signals by amplifier-shaper-discriminator (ASD) boards. The digital signals from the ASD boards are converted into drift times in the time-to-digital converter (TDC) boards. Each muon system has a different preamplifier card design (although some designs use essentially the same components). The CMU uses ASD boards that are newly designed for Run II, while the other muon chambers use older ASD boards reused or recycled from Run 1. All the muon systems use the same TDC boards newly designed and built for Run II. For the muon scintillators, signals from the PMT's are discriminated and the logic signals are sent to the same TDC boards used by the muon chambers.

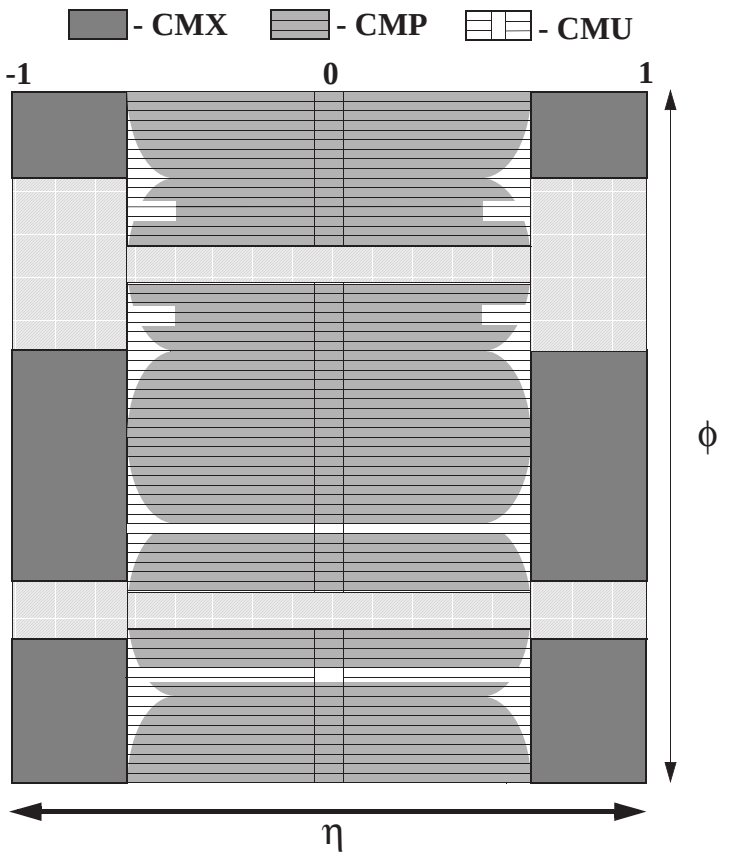

Figure C.2 Upgraded CMP and CMX detector coverage for Run II. New chambers added to the CMP and CMP for Run II with light shading.

Despite the differences in muon detector location and geometry, the electronics and operating principles of the muon detectors are all similar enough that one calibration system may be used to monitor performance, debug problems, and calibrate the detector. The upgraded front-end electronics added to the CMU chambers possess built-in calibration functionality and so the CMU has a separate calibration scheme.

The basic principles for calibrating the muon chambers and scintillator are very similar. The idea is to manipulate the muon chamber and scintillator electronics to generate a signal similar to the signals caused by actual muons traversing the detector, and to read out this calibration signal as if it were real data to test the operation of the muon systems. In the muon chambers, this is

accomplished by injecting charge either onto the muon chamber wires, as in the CMU, or into the preamplifier electronics, as for the other muon detectors. For many of the scintillator detectors, a fake muon signal is generated via an LED that flashes light into the scintillators. This light is read out as if it were produced by a muon traversing the scintillating plastic.

The primary purpose of the muon calibration system for Run II is to monitor the chamber performance to detect new dead and malfunctioning channels that develop over the course of the experiment. In addition, the muon calibration system is useful for debugging problems with the muon chamber electronics because it can provide a steady and controlled source of signals that can be monitored with an oscilloscope. Despite its name, the muon calibration system is not extensively used for muon calibrations. Most of the calibrations for muon detectors are taken from collider or cosmic ray data. However, in some situations, the calibration system provides better calibration information than the data. For example, calibrating CMU charge division requires more control over the amount and location of charge deposited on the CMU wires than is available using collider or cosmic ray data. In addition, the statistics available from data may not be sufficient to determine small timing offsets on individual channels due to differences in cable lengths.

C.1 Hardware

At CDF, calibrations are coordinated and controlled by a piece of hardware known as the trigger supervisor (TS). The main purpose of the TS is to coordinate the timing and readout of the entire CDF detector in response to trigger decisions. During calibrations, the TS controls the timing of the calibration input, such as the signal that fires charge injection or the LED's, with readout of the calibration data. The TS delays the readout of the detectors being calibrated by a long enough period after the calibration signal to account for propagation of the calibration signal to the various detectors and propagation of the detector response back through the data acquisition system (DAQ). The TS communicates with the CDF detector hardware by sending signals to a special "trigger and clock + event readout module" (TRACER) board located in each VME crate. The TRACER accepts signals from the TS and propagates the necessary information to the rest of the detector electronics via the VME backplane. One of the main purposes of the muon calibration hardware is to take calibration signals sent by the TRACER and transmit them to electronics not housed in VME crates, like the old Run 1 ASD boards.

C.1.1 CMU Calibration Hardware

All necessary calibration hardware for the CMU is built into the ASD board, know as the ASDII, designed specifically for the CMU in Run II. The ASDII boards are VME boards capable of receiving calibration signals directly from the TRACER. Upon receiving a calibration signal, the ASDII injects charge onto ganged-pairs of CMU wires at one of four charge-injection points. The charge-injection point and the amount of charge can be programmed by communicating to the ASDII board along the VME backplane. The charge injected on the CMU wire pairs simulates a hit from an actual muon and is read out through the same DAQ chain used for regular data taking, including the chamber preamplifier modules, the ASD cards and the TDC cards. Data taken from charge injection at different points along the ganged wire pair and with different amounts of charge injected allows a calibration of the CMU *z* position measurement through charge division.

C.1.2 Other Chamber Calibration Hardware

The ASD cards used for the CMP, CMX, and BMU are all reused –or in the case of the BMU, recycled—from Run 1. Although the BMU ASD cards are slightly different from those of CMP and CMX, the differences are minor and don't affect most of this discussion. The Run 1 ASD cards for these detector systems reside in a custom crate design that contains no processing or communications functionality. As a result, the muon calibration system must provide VME hardware that can receive calibration signals passed on through the TRACER from the TS, and send the appropriate inputs directly to cards in the Run 1 style ASD crates.

On the VME side, the main component of the calibration hardware for CMP, CMX, and BMU is a commercial digital delay generator (DDG) seated in a custom-designed calibration interface card (CIC). The digital delay generator is capable of providing logic signals at a programmable delay after a trigger signal. Varying the delay programmed into the DDG allows calibration pulses to be sent to the muon chambers at differing times within the fixed window between the calibration signal from the TS and the readout of the muon TDC cards. The pulse width of the DDG signals is also programmable. The CIC interprets signals from the TRACER card sent along the VME backplane and upon receive a TS calibration signal, it triggers the DDG to fire. The pulse from the DDG is fanned-out using a custom-built VME fanout card with one copy being sent to each of the twelve Run 1 style ASD crates used for the CMP, CMX, and BMU.

Calibration pulses sent from the DDG are received in the ASD crate by the calibration card, which serves as a programmable fanout. The calibration card in each ASD crate fans out the calibration signal to each card in the crate. Upon receiving the calibration signal, the ASD card discharges a capacitor to inject charge into the preamp circuits. This charge behaves like charge accumulated on the anode wires from muon ionization and is read out exactly like the data through the regular DAQ chain (preamp, ASD, and TDC). The calibration card is capable of controlling which muon detector channels receive charge injection and the smallest block of channels that can be selected through this means is a group of channels corresponding to one-half of an ASD card. In addition, the calibration card controls the size of the voltage used to charge the charge-injection capacitors on the ASD cards. The channels selected for pulsing and the amount of charge to be injected is communicated to the calibration card using RS-232 signals sent from an adapter card mounted on a VME crate controller. In addition, a VME-based ADC module monitors the voltage setting for charge injection to insure the correct value is used.

The BMU ASD's have a special feature that differentiates them from the CMP and CMX ASD cards. It is possible to inject charge directly into the BMU ASD cards, bypassing the preamp portion of the readout. This feature is useful to help determine whether a problem noticed by the calibration system in the BMU is the result of a malfunction in the preamp or the ASD card. Whether pulses are sent to the BMU preamps or just to the BMU ASD cards directly is also determined by the calibration card and is programmed using the same RS-232 line as the other calibration card functions.

C.1.3 Scintillator Calibration Hardware

Like the CMP, CMX, and BMU, the muon scintillator systems at CDF require hardware to transmit the appropriate calibration signals in response to commands from the TS to the non-VME electronics used for the scintillator readout. The same DDG is used to send calibration pulses to the scintillator electronics in response to calibration signals from the TS. What happens to these DDG signals depends on which category the scintillator electronics falls into: the new style designed and built for Run II or the old style reused from Run 1.

Readout and calibration of the new-style scintillator is handled by control and concatenation units (CCU's). The scintillator systems using CCU's in their readout include the BSU, TSU, and parts of the CSP. The CCU's handle setting the high voltage for the scintillator PMT's. In addition, the CCU is responsible for gathering the discriminated signals from the PMT,

amplifier, and discriminator units (PAD's) on the new scintillator and sending them up to the TDC's. Finally, the CCU triggers the calibration LED flashers in response to calibration signals from the DDG. Light generated by the LED flashers is registered by the PAD and read out through the CCU and TDC just like regular data.

The older-style scintillator systems include the CSX and part of the CSP. These systems do not use CCU's or PAD's in their readout. Instead, raw signals from the scintillator PMT's are sent to NIM-based commercial discriminator units. The digital output from these units is then passed on to the TDC cards. The older-style scintillators have no way of simulating muon hits in the PMT. Instead, the calibration signals from the DDG are sent to directly to the discriminators causing them to fire and send signals to the TDC, which are read out like normal data.

C.2 Software

A variety of software is required to control the various elements of the muon calibration system. RUN CONTROL software provides the interface between the calibration user and the calibration hardware. The run configuration determined in RUN CONTROL is transmitted to front-end code residing on the crate controllers of the VME crates where it is used to determine the exact settings to be programmed into the calibration hardware. After the calibration has been performed, a calibration consumer program is used to analyze the raw data and to store the results of the calibration in a database. Finally, a program known as DBANA allows users to view graphically results in the calibration database and also to analyze that data to determine whether any new dead or malfunctioning channels have been detected. The sections below describe each of these pieces of software in more detail.

C.2.1 RUN CONTROL

RUN CONTROL [82] is the general-purpose software used for controlling the entire CDF detector. It provides users with a graphical user interface (GUI) for selecting which detector components should be part of a run, what kind of run should be taken (collider physics, calibration, etc.), and how long the run should last. Running muon calibrations is only a small part of RUN CONTROL's functionality. RUN CONTROL is written in the JAVA programming language and can be run on any PC with internet access to the CDF online network.

To begin a muon calibration run using RUN CONTROL, the user selects from a collection of predefined muon run configurations. These include default configurations of muon calibration hardware needed for various kinds of runs (see Section C.3 below). Once the run configuration is selected, RUN CONTROL downloads default setting for all the muon calibration hardware from the run configuration and hardware databases. The user then has an opportunity to edit these settings in order to create a non-standard configuration before the information is passed on to the front-end code in the individual crates.

C.2.2 Front-end Code

Front-end refers to any code that runs on the VME crate controller processors. This code is responsible for the actual low-level details of configuring and operating the CDF DAQ electronics. For example, this code writes configuration information into the registers on the various boards used to readout the detector and determines what happens in each crate during the various stages of any given run. Configuration information is communicated to the front-end code from RUN CONTROL using Ethernet. This code is written using the C programming language and makes use of a modified version of the VISION VME programming library. The front-end code runs under the VXWORKS operating system, which is a paired down, real-time version of the UNIX operating system designed for use on VME crate controllers.

When RUN CONTROL transmits the muon calibration configuration information to the front-end code residing in the muon crates, a number of different things happen. In the CMU crates, the front-end code configures that ASDII cards as necessary for the calibration. It configures the calibration card to inject the desired amount of charge into the specified injection point and sets the ASDII to begin injecting charge upon receiving calibration signals from the TRACER. The front-end code in the crate containing the muon calibration DDG configures the DDG with the necessary delay values and pulse widths for this calibration run. It also instructs the RS-232 adapter to send the appropriate configuration commands to the calibration cards in the CMP, CMX, and BMU ASD crates. It triggers the ADC to check the capacitor voltage for charge injection on the calibration cards to make sure the correct value has been set on each run. Finally, the front-end code in each TDC crate resets the muon TDC's and prepares them to receive data from the calibration run.

At points during the actual calibration run, as determined by the run configuration, the front-end code will pause the data taking and adjust some of the parameters of the calibration run. For

example, during certain run configurations, every so many calibration events, the calibration pauses and the delay value for the DDG is changed or the amount of charge and charge injection point for the CMU is adjusted. Also, during these pauses, the front-end code takes advantage of the opportunity to instruct the ADC to readout the capacitor voltages on the calibration cards to make sure the expected amount of charge is still being injected into the CMP, CMX, and BMU preamps. Once all the adjustments are complete, the front-end code reinitiates data taking with the new settings.

Finally, as the calibration data is recorded by the TDC cards, the front-end code is responsible for packaging this data into a format that can be stored on disk. Data from all the TDC cards in a given crate is packaged together, formatted, and sent off to the CDF event builder hardware. The event builder assembles data from all the detector crates into one collection to be written to a file.

C.2.3 Calibration Consumer

After the data from a calibration run has been collected, a calibration consumer program analyzes it. The muon calibration consumer is written in C++ employing the same software framework used for physics analyses at CDF. It converts the raw calibration data into meaningful statistics, for example the number of TDC hits recorded per calibration pulse. The calibration consumer can be run online automatically by run control directly on the data as it is collected, or it can be used offline on calibration data stored in a file. The exact operation of the consumer code depends on the muon detector being calibrated. The calculated quantities are discussed in Section C.3.3. As of this writing, the consumer code for the CMU calibration is still under development.

C.2.4 DBANA

DBANA is a graphical interface to the data stored in the calibration database. It can be used to access calibration information about a number of detector systems at CDF. DBANA can be used to display the results of an individual calibration run. It can also compare the results from two or more runs. Users can track the calibration history of detectors using this code. In addition, DBANA provides for fast feedback on problems uncovered by detector calibrations. If the consumer code is run online from RUN CONTROL, DBANA can immediately view the calibration results and analyze them to find malfunctioning channels. Many CDF detector systems use

DBANA to detect and track bad channels. Although the general DBANA framework currently exists, the specific code required to display and check muon calibration results is still under development at this time.

C.3 Modes of Operation

The muon calibration hardware and software was intended to be as flexible as possible to allow different configurations for any imaginable test procedure or diagnostic. Nevertheless, it was designed with a few standard modes of operation in mind. The most common modes of operation for the muon calibration system are listed below. For each mode, one or more run configurations are stored in the run configuration database. Other configurations can quickly be built starting from a similar configuration in the database.

C.3.1 Continuous Pulsing Mode

This mode of operation, also known as scope mode, is intended for detector experts who wish to diagnose a problem with a muon detector by following the path of signals through the detector electronics with an oscilloscope. When operating in this mode, the muon calibration system sends a continuous, regular stream of pulses at a fixed period to the desired muon detectors. This stream of pulses continues until the user manually ends the run.

C.3.2 Single-Delay Mode

This mode of operation was used primarily to debug the calibration system during its initial commissioning. However, it is still useful in cases where a simple data file is desired for analyzing strange calibration results. In single-delay mode, the calibration system sends a set number of pulses, typically 1000, to the desired muon detectors all at the same delay value. The expected result is that the selected muon detectors will read out 1000 (or more if ringing or reflections are present) hits, all with the same time. This mode can be used to check that a channel is alive—in other words, that there is at least one hit for every pulse sent—and can also be used to deduce relative cable delays in systems where different lengths of cable were used for different parts of the detector. The spread in hit times received can be used to find channels with inconsistent timing. In addition, the average widths and the spread in widths give more information about the health of the detector in question.

C.3.3 Multiple-Delay Mode

In multiple-delay mode, the muon calibration system sends a set number of pulses at each of several delay values. For the BMU chambers, there is also the option for the first set of pulses to be ASD-only pulses instead of preamp pulses. Using the run configuration interface in RUN CONTROL, the user can specify the total number of pulses, number of pulses at each delay, the initial delay value, the amount to step the delay for subsequent sets of pulses, and whether the first set of pulses should be ASD-only for the BMU. A typical configuration would be 1000 pulses, 100 pulses at each delay value, an initial delay of 0 ns for the muon chambers and 250 ns for the scintillators, a delay step size of 100 ns for chambers and scintillators, and ASD-only pulses for the first set selected.

This mode of operation allows more checks than the single-delay pulsing mode. In addition to checking that a channel is alive by counting hits, one can fit the hit times versus the delay values to a straight line to make sure the relation ship between delay values and hits are linear. If this fit is successful, the *y*-intercept of the fit indicates the relative timing offset due to different signal path lengths. Channels that don't produce successful straight-line fits can be flagged for further investigation. Finally, all the same information on spread in hit times, average hit widths, and spread in hit widths are available in this pulsing mode as well. The muon calibration consumer, which will automatically calculate all the quantities discussed above for each channel in the muon detector and store the results in the calibration database, expects the multiple-delay mode of operation.

C.3.4 CMU Charge Division Calibration

In this mode of operation, the CMU chambers are pulsed with different amounts of injected charge at each of the four different charge injection points. A typical configuration in this mode might pulse the CMU chambers ten times each at up to 255 different amounts of charge injected for each of the four charge injection points. Data from this mode of operation can be used to calibrate the *z* position measurement in the CMU from charge division. The exact details of this calibration are still under development.

C.4 Performance

The plots in this section are taken from a typical muon calibration run for the CMP detector. The CMP is the detector for which the muon calibration system is most extensively tested. Other muon detectors have also been checked using the muon calibration system, but further work is required before the calibration results can be viewed as reliable or stable.

In Figure C.3, the number of TDC hits per calibration pulse as well as the average width of the earliest hit on each channel are shown. As can be seen, because of reflections of the calibration signal from the unterminated end of the chamber wire and other locations, most channels register three to four hits per pulse. However, some channels register only one or two hits. These channels are also correlated with the channels that display unusually large hit widths. The deviation in number of hits and hit width in these channels is caused by the type of preamp installed on the chambers. A number of the CMP preamps have been modified to reduce their chance of oscillating. The modified preamps respond differently to the calibration pulse and are clearly recognizable in the plots.

Figure C.4 shows the distribution of time offsets obtained from the CMP calibration. Several clusters of offset times can be seen in this plot. Each represents a section of CMP chambers that were instrumented using cables of a different length. The observed delay values are consistent with expectations from cable lengths as well as observations from data. These results have been used to determine channel timing offsets for channels lacking sufficient statistics to determine the offsets from data. Figure C.4 also shows the timing slope information for CMP channels. Most channels report a slope near one, which is expected. Those channels that deviate significantly from a slope of one indicate possibly malfunctioning muon chambers.

Calibration results from the other muon systems are similar to those shown here, although there is a greater fraction of anomalous channels and the results lack stability. Development continues on the calibration hardware and software for these other systems.

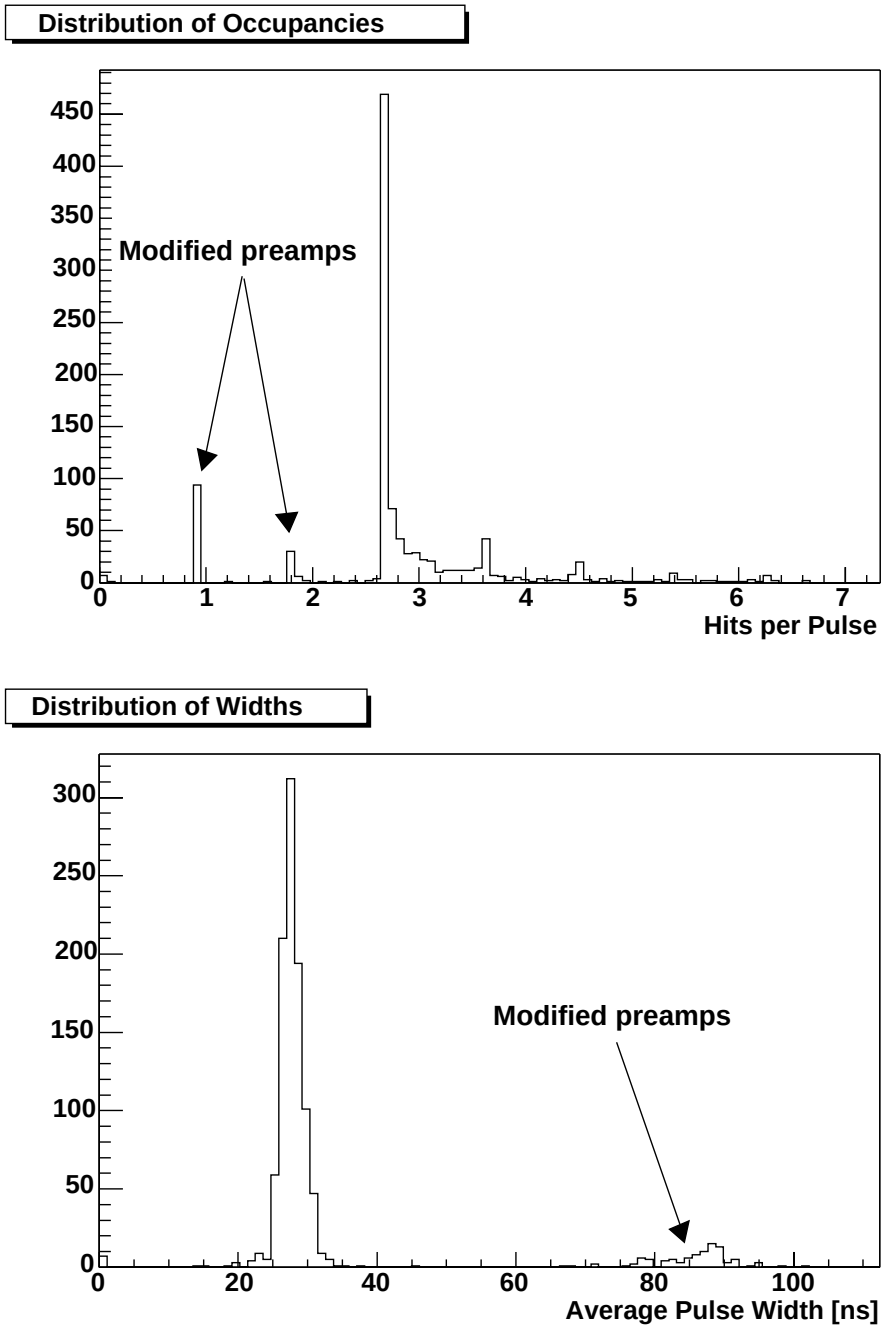

Figure C.3 The distribution of hits per calibration pulse and average pulse width for all CMP channels. Certain channels with modified preamps appear in these plots with fewer hits per pulse and larger pulse width.

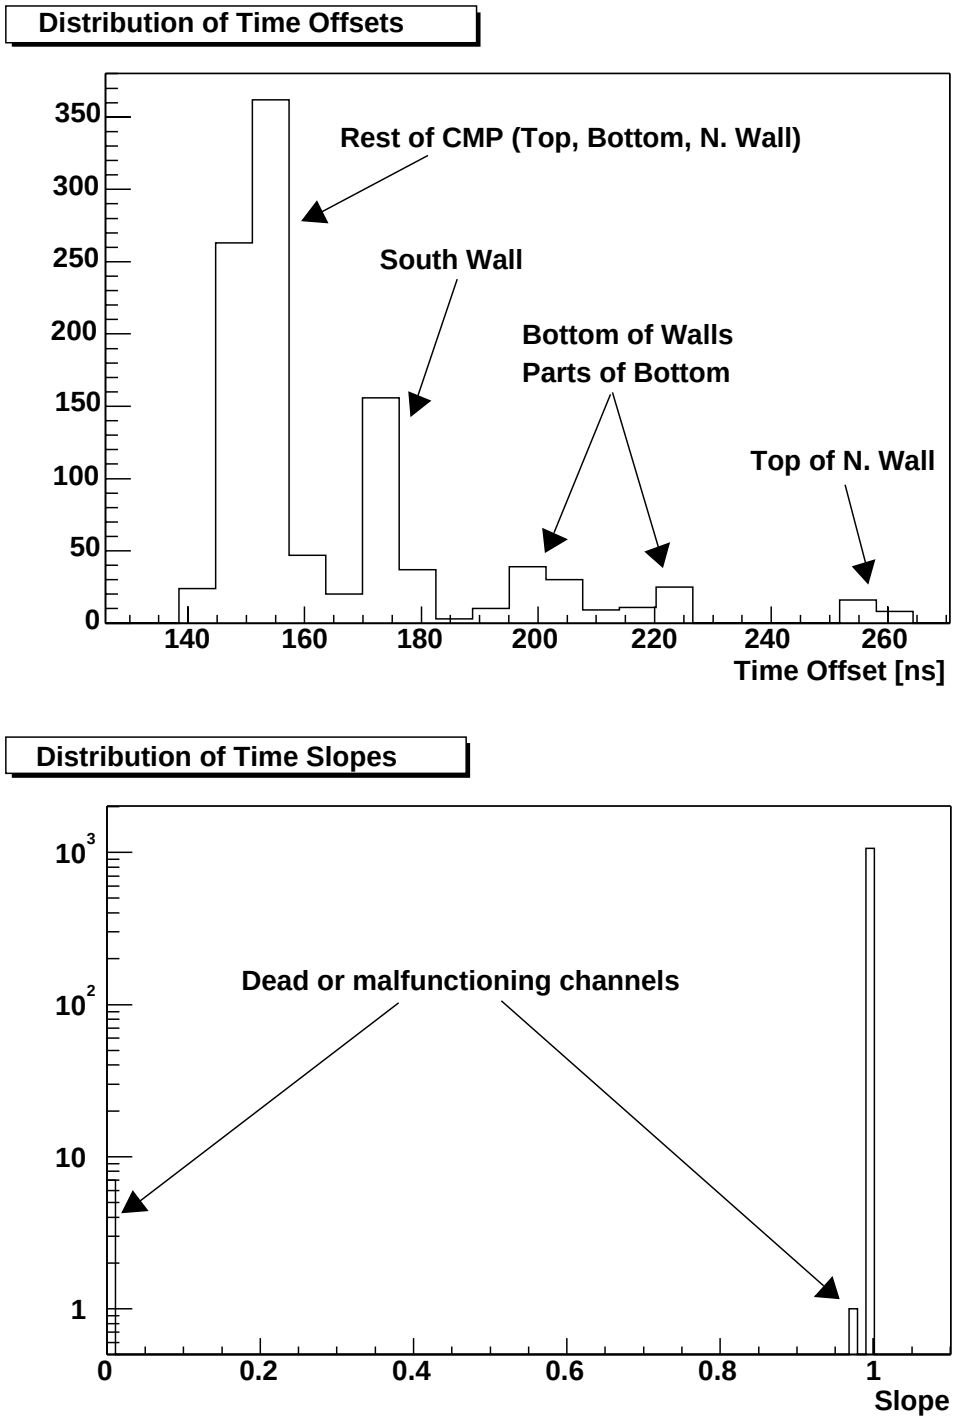

Figure C.4 The results of the straight-line fit to the CMP calibration data. The upper plot shows the distribution of timing offsets measured by the fit. There are several clusters of offset time that correspond to known sets of channels with different cable lengths. The lower plot shows the slope of the straight-line fit, with most channels having slopes nearly equal to one.

References

[1] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*. (Perseus Books, Reading, 1995).

[2] C. Quigg, *Guage Theories of the Strong, Weak, and Electromagnetic Interactions*. (Addison-Wesley, Reading 1997).

[3] R. P. Feynman, *Quantum Electrodynamics*. (Benjamin, New York, 1962).

[4] S. Willenbrock, "The Standard Model and the Top Quark," a lecture given at the 12th Advanced Study Institute on Techniques and Concepts of High-Energy Physics, St. Croix, U.S. Virgin Islands, 2002; M. K. Gaillard, P. D. Grannis, F. J. Sciulli. Rev. Mod. Phys. **71**, S96 (1999).

[5] S. L. Glashow, Nucl. Phys. **22**, 579 (1961); S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967); A. Salam, *Elementary Particle Theory: Relativistic Groups and Analyticity* (Nobel Symposium No. 8), edited by W. Svartholm (Almqvist and Wiksell, Stockholm 1968), 367.

[6] P. W. Higgs, Phys. Rev. Lett. **12**, 132 (1964); Phys. Rev. **145**, 1156 (1966); G. 't Hooft, Nucl. Phys. **B33**, 167 (1971); Nucl. Phys. **B35**, 167 (1971).

[7] M. Kobayashi and K. Maskawa, Progr. Theor. Phys. **49**, 652 (1973).

[8] W. J. Marciano, Nucl. Phys. Proc. Suppl. **B59**, 339 (1997)

[9] K. Hagiwara *et al*., Phys. Rev. **D66**, 010001-1 (2002).

[10] H. Fritzsch, M. Gell-Mann, and H. Leutwyler, Phys. Lett. **B47**, 365 (1973).

[11] V. W. Hughes and T. Kinoshita, Rev. Mod. Phys. **71**, S133 (1999).

[12] K. Wilson, Phys. Rev. **D10**, 2445 (1974); M. Creutz, *Quarks, Gluons, and Lattices*. (Cambridge University Press, Cambridge, 1983).

[13] S. Frixione *et al*., Adv. Ser. Direct High Energy Phys. **15**, 69 (1998).

[14] C. Peterson *et al*., Phys. Rev. **D27**, 105 (1983).

[15] G. Abbiendi *et al*, hep-ex/0210031 submitted to Eur.Phys.J.**C**.

[16] M Cacciari and P. Nason, Phys. Rev. Lett. **89**, 122003 (2002).

[17] J. Wess and J. Bagger, *Supersymmetry and Supergravity*. (Princeton University Press, 1983).

[18] E. Berger, Int. J. Mod. Phys. **A18**, 1263 (2003).

[19] P. Nason, S. Dawson, and R. K. Ellis, Nucl. Phys. **B327**, 49 (1989); *erratum*, Nucl. Phys. **B335**, 260 (1990); Nucl. Phys. **B303**, 607 (1988).

[20] R. D. Field, Phys Rev. **D65**, 094006 (2002).

[21] C. Albaja *et al*., Z. Phys. **C61**, 41 (1994).

[22] M. Mangano, P. Nason, and G. Ridolfi, Nucl. Phys. **B373**, 295 (1992).

[23] H. Baer, F.E. Paige, S. D. Protopopescu and X. Tata, arXiv:hep-ph/0001086; F. Paige and S. Protopopescu, BNL Report BNL38034, 1986 (unpublished), version 7.32.

[24] F. Abe *et al*., Phys. Rev. **D55**, 2546 (1997).

[25] B. Abbott *et al*., Phys. Lett. **B487**, 264 (2000).

[26] F. Abe *et al*., Phys. Rev. **D53**, 1051 (1996).

[27] D. Buskulic *et al*., Phys. Lett. **B313**, 535 (1993).

[28] F. Abe *et al*., Phys. Rev. **D61**, 032001 (2000).

[29] F. Abe *et al*., Phys. Rev. **D50**, 2966 (1994).

[30] J.C. Collins, D. E. Soper, and G. Sterman, Nucl. Phys. **B263**, 37 (1986).

[31] S. Frixione and B. R. Webber, JHEP **06**, 029 (2002); S. Frixione, P. Nason, and B. R. Webber, hep-ph/0305252.

[32] P. Nason, S. Dawson, and R.K. Ellis, Nucl. Phys. **B303**, 607 (1988); P. Nason, S. Dawson, and R.K. Ellis, Nucl. Phys. **B327**, 49 (1988).

[33] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, Computer Phys. Commun. **135**, 238 (2001); T. Sjostrand, Phy. Lett. **157B**, 321 (1985).

[34] G. Corcella, I.G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour and B.R. Webber, JHEP **0101**, 010 (2001) [hep-ph/0011363];hep-ph/0201201; G. Marchesini, B.R. Webber, G. Abbiendi, I.G. Knowles, M.H. Seymour and L. Stanco, Computer Phys. Commun. **67**, 465 (1992).

[35] A.H. Mueller, Phys. Lett. **104B**, 161 (1981); B.I. Ermolaev, V.S. Fadin, JETP Lett. **33**, 269 (1981).

[36] H. L. Lai *et al*., Eur. Phys. J. **C12**, 375 (2000).

[37] A. D. Martin *et al.*, Eur. Phys. J. **C4**, 463 (1998).

[38] P. Nason and C. Oleari, Nucl. Phys. **B565**, 245 (2000).

[39] M.G. Bowler, Z. Phys. **C11**, 169 (1981).

[40] H. Edwards. Ann. Rev. Nucl. Part. Sci. **35**, 605 (1985).

[41] J. Thompson, Fermilab Technical Memo, FERMILAB-TM-1909, http://library.fnal.gov/archive/test-tm/1000/fermilab-tm-1909.shtml

[42] C. W. Schmidt and C. D. Curtis, "Operation of the Fermilab H^- Magnetron Source." Proc. 4^{th} Int. Symp. on the Production and Neutralization of Negative Ions and Beams, Brookhaven, US, AIP Conf. Proc. No 158, 425 (1986).

[43] K. Junck *et al*., "Commissioning and First Operational Experience of the 400-MeV Linac at Fermilab." FERMILAB-Conf-94-181 (1994).

[44] M. D. Church and J. P. Marriner. Ann. Rev. Nucl. Part. Sci. **43**, 253 (1993).

[45] A. Ruggiero. IEEE Trans. Nucl. Sci. **30**, 2478 (1983).

[46] B. Austin *et al*., IEEE Trans. Nucl. Sci. **30,** 2593 (1983).

[47] A. Ando *et al*., IEEE Trans. Nucl. Sci. **30,** 2031 (1983).

[48] F. Abe *et al*., Nucl. Instrum. Methods Phys. Res. **A271**, 387 (1988); F Abe *et al*., Phys. Rev. Lett. **74**, 2626 (1995); F. Abe *et al*., Phys. Rev. **D50**, 2966 (1994).

[49] C. Newman-Holmes, E. E. Schmidt, and R. Yamada, Nucl. Instrum. Meth. **A274**, 443 (1989).

[50] P. Azzi *et al*., Nucl. Instrum. Meth. **A360**, 137 (1995).

[51] The vertex time projection chamber used for this analysis is similar to an earlier chamber described in F. Snider *et al*., Nucl. Instr. Methods **A268**, 75 (1988).

[52] F. Bedeschi *et al*., Nucl. Instrum. Mech. **A271**, 387 (1988).

[53] L. Balka *et al*., Nucl. Instrum. Meth. **A267**, 272 (1988).

[54] S. Bertolucci *et al*., Nucl. Instrum. Meth. **A267**, 301 (1988).

[55] G. Ascoli *et al*, Nucl. Instrum. Meth. **A268** (1988).

[56] D. Amidei *et al*., Nucl. Instrum. Methods Phys. Res. **A269**, 68 (1988).

[57] G. Foster *et al*., Nucl. Instrum. Methods Phys. Res. **A269**, 93 (1988).

[58] F. Abe *et al*., Phys. Rev. **D60**, 072003 (1999).

[59] S. Dell'Angello, Ph.D. Thesis, University of Pisa (1994).

[60] F. Abe *et al*., Phys. Rev. **D57**, 5382 (1998).

[61] K. Hagiwara *et al*. (Particle Data Group), Phys. Rev. **D66**, 010001 (2002) and 2003 partial update for edition 2004 (http://pdg.lbl.gov)

[62] O. Long, Ph.D. thesis, University of Pennsylvania (1998).

[63] M. Peters, Ph.D. thesis, University of California, Berkeley (1997).

[64] R. Field, private communication.

[65] E. Norrbin and T. Sjöstrand, Phy. Lett. **B442**, 407 (1998); E. Norrbin and T. Sjöstrand, Eur. Phy. J. **C17**, 137 (2000).

[66] R. D. Field, http://www.phys.ufl.edu/~rfield/cdf/FNAL_Workshop_10-4-02.pdf.

[67] 'QQ - The CLEO Event Generator,' http://www.lns.cornell.edu/public/CLEO/soft/QQ (unpublished).

[68] W. Taylor, Ph.D. Thesis, University of Toronto (1999).

[69] S. Pappas, Ph.D. Thesis, Yale University (1999).

[70] J. P. Done, Ph.D. Thesis, Texas A & M University (1999).

[71] S. Pappas, "DIMUTG: Dimuon Trigger Simulation for Run 1B," CDF Note 3537; private version containing single muon efficiencies obtained from J. D. Lewis, private communication.

[72] W. Ashmanskas *et al*., Nucl. Instrum. Meth. **A447**, 218 (2000).

[73] I. Furic, http://www-cdf.fnal.gov/physics/talks_transp/2003/lathuile_charm_furic.ps.

[74] A. D. Sakharov, J. Exp. Theor. Phys. Lett. **5**, 24 (1967).

[75] H. Quinn, Phys. Today **56N2**, 30 (2003).

[76] S. Stone, "Prospects for *B* Physics in the Next Decade," in *Techniques and Concepts of High-Energy Physics IX: Proceedings*, *NATO ASI Series*, edited by T. Ferbel (Plenum, New York, 1997), p. 465.

[77] F. Abe *et al*., Phys. Rev. Lett. **80**, 2057 (1998); F. Abe *et al*., Phys. Rev. **D57**, 032001 (1999).

[78] F. Maltoni, Z. Sullivan, and S. Willenbrock, Phys. Rev. **D67**, 093005 (2003); J. Campbell *et al*., Phys. Rev. **D67**, 095002 (2003).

[79] R. Demina *et al*., Phys. Rev. **D62**, 035011 (2000).

[80] G. Ascoli *et al*., Nucl. Instrum. Meth. **A268**, 33 (1988).

[81] K. Byrum *et al*., Nucl. Instrum. Meth. **A268**, 46 (1988).

[82] W. Badgett *et al*., "CDF Run II Run Control and Online Monitor," presented at CHEP'01: Computing in High-Energy Physics and Nuclear, Beijing, China, 2001, http://www.ihep.ac.cn/~chep01/paper/9-048.pdf.

Vita

Kevin Patrick Lannon [redacted]. Growing up, he possessed a great curiosity about the world around him, although he displayed a far greater talent for taking things apart than for putting them back together. After his high-school dreams of becoming a professional runner and breaking the world record in the 800m failed to materialize, Kevin went off to St. Norbert College, in De Pere, Wisconsin. There he studied English, physics, math, and religion—and spent a fair amount of time running, both across fields and in circles. At the end of four years, Kevin received his B.A. from St. Norbert with majors in English, math, and physics. Uncertain how to proceed, he resolved to attend a graduate school of some kind. Based on three happy summers of research as an undergraduate with Bob Wagner at Argonne National Laboratory and the conviction that there was more money in physics than in English, Kevin settled on the graduate program in physics at the University of Illinois. After approximately three years of classes, in which he finally learned what a Lagrangian was, Kevin left Champaign-Urbana in order to do research with the CDF collaboration at Fermi National Accelerator Laboratory. During this time, he learned to drive lifts around the assembly area and collision hall, discovered the joys of crawling under the detector and climbing over it, wrote programs in an alphabet soup of different languages, and studied the bottom quark. Finally in 2003, he finished his thesis, completed his Ph.D., and moved on to a new job at The Ohio State University, working with fancy trigger electronics and studying the top quark.