Elliptic flow of Υ mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the CMS detector

By

Jared Stephens Jay

Dissertation

Submitted in partial satisfaction of the requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Office of Graduate Studies

of the

University of California

Davis

Approved:

Manuel Calderón de la Barca Sánchez, Chair

Daniel Cebra

Ramona Vogt

Committee in Charge

2021

To my wonderful wife Daiana,

I love you forever.

Contents

List of Figures

List of Tables

Abstract of the Dissertation

Elliptic flow of Υ mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the CMS detector

Quantum Chromodynamics (QCD) describes all the interactions of quarks and gluons, including the various phases of QCD matter such as the quark-gluon plasma (QGP) which is thought to exist at very high temperatures and baryon densities, and dominated the very early universe.

Heavy-ion collisions are one way in which the QGP can be created and studied in a laboratory environment. The Large Hadron Collider (LHC) is capable of colliding lead nuclei with a center-of-mass energy per nucleon pair of $\sqrt{s_{NN}} =$ 5.02 TeV. The CMS detector accurately and efficiently measures the energy and momenta of particles exiting these collisions, even in extremely high-occupancy environments, and provides a rich set of data which can be analyzed to learn more about the thermodynamic properties of the QGP, particularly through the measurement of muons.

Recent studies indicate that the QGP generated in heavy-ion collisions exhibits liquid-like collective behavior, expressed as an elliptic flow of the colliding medium generated by internal pressure gradients. The QGP flows with the lowest viscosity ever observed, making it a near perfect fluid. More studies are warranted to nail down the properties of this unusual form of matter.

The second-order Fourier coefficients (v_2) characterizing the elliptic flow of Υs in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV at CMS, are reported. The Υs are reconstructed through their dimuon decay channel in the rapidity range $|y| < 2.4$. The suppression of quarkonia in the presence of a quark-gluon plasma is expected to result in a very small v_2 signature. The v_2 is estimated using the event-plane method, validated by a closure test, and is found to be consistent with zero.

Acknowledgments

This thesis is the result of years of mentorship from experienced scientists, collaboration with colleagues, and support from family. I would like to extend my sincere gratitude to everyone who helped me to reach this important milestone.

First, I want to thank my main advisor and mentor, Manuel Calderón de la Barca Sánchez for inspiring me to enter this exciting field of study, for mentoring me every step of the way as a true expert in the field, for promoting a proper work-life balance, and making the overall experience of graduate school enlightening, empowering, and enjoyable.

I also thank my other thesis committee members. Daniel Cebra taught me much of what I know about nuclear physics, explaining complex concepts simply in a way that anyone could understand. We had many useful discussions about various topics that are described in this thesis. Ramona Vogt pushed me to a higher standard of scholarship and physics knowledge. Her many useful comments and insights improved the quality of this thesis immensely.

I want to give my thanks to my fellow students who served as mentors to me throughout this journey. To Ota, for teaching me the basics of tag-and-probe and providing mentorship throughout the entire analysis, and for always having ideas of how to improve any analysis that I presented to him. To Santona, now graduated, for her mentorship and leadership in my first analysis where I learned most of the tricks of the trade, and for providing useful and practical advice for my career. To Graham, who worked with me on many projects during this graduate experience.

I also want to extend my thanks to my parents who encouraged me and supported me throughout my life. My mother, Marsha, imparted much wisdom to me and encouraged me to do hard things. My Father, Rodney, was the first physicist in the family, a great example of analytical thinking, and probably the most brilliant man I've ever met.

I also extend my sincere gratitude to my beloved wife Daiana for being so

patient with me through all these years of graduate school, for always being there for me, for taking care of our three children while I worked, and for encouraging me to stay in graduate school and finish my PhD when I felt like giving up. I love you Daiana. Thank you.

Chapter 1

Theoretical Background

1.1 The Standard Model

The Standard Model is a Lagrangian theory of particle physics based on a small set of assumptions that provides a rich set of particles and interactions which, together with General Relativity, are sufficient to explain nearly everything we observe in the universe. The Standard Model assumes a universal gauge symmetry group $\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)$ and a Lorentz-invariant space-time with 3 spatial dimensions and time. All particles in the theory are treated as vectors or "states" embedded in the space-time. The states transform via the action of operators which are necessarily unitary in order to preserve probability. The unitarity of the operators allows for them to be written as $\vartheta = e^{iS}$ where S is a traceless Hermitian operator. In general, the S is a time-integral of the Lagrangian of the theory, which can also be written as an integral of the Lagrangian density over all of space-time. The Lagrangian density of the Standard Model is a sum over all possible traceless, Lorentz-invariant, Hermitian operators of weight 4 which can be built from the 17 fundamental particles that have been observed, keeping only "relevant" terms. See *e.g.* [1] and [2].

The Lagrangian density can be sub-divided into 3 parts: The electroweak sector, the quantum chromodynamics (QCD) sector, and the Higgs sector. The electroweak sector describes the leptons (neutrinos, electrons, muons, tauons) as well

as the photons and the W and Z bosons. The QCD sector describes the quarks and gluons. The Higgs sector describes the Higgs boson and its coupling to the other fundamental particles. This coupling provides the mechanism for electroweak symmetry breaking which generates the masses of the fundamental fermions, including the quarks.

1.2 Quantum Chromodynamics

Quantum Chromodynamics describes all of the interactions of quarks and gluons. The gauge-invariant Lagrangian density of QCD (after electroweak symmetry breaking) can be written as

$$\mathcal{L}_{\text{QCD}} = \bar{\psi}_i(i(\gamma^\mu D_\mu)_{ij} - m\delta_{ij})\psi_j - \frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a, \tag{1.1}$$

where $\psi_j(x)$ is the quark field in the fundamental representation of SU(3), indexed by $i, j \in \{1, 2, 3\}$, and a function of the space-time x, m is the quark mass, and $G^a_{\mu\nu}$ is the gluon field strength tensor. The operators D_μ and γ^μ are the gauge covariant derivative and the Dirac matrices, respectively. The indices μ and ν run over the four dimensions of space-time. The index a corresponds to the eight gluons of QCD. Einstein index notation is used in the equation. The object $G^a_{\mu\nu}$ is referred to as the "gluon field strength tensor" and is given by

$$G^a_{\mu\nu} = \partial_\mu \mathcal{A}^a_\nu - \partial_\nu \mathcal{A}^a_\mu + g f^{abc} \mathcal{A}^b_\mu \mathcal{A}^c_\nu, \tag{1.2}$$

where $\mathcal{A}^a_\mu(x)$ is the gluon field as a function of space-time x, g is the coupling constant of QCD, and f^{abc} are the structure constants of the gauge group SU(3). The gluon field strength tensor is analagous to the field strength tensor $F_{\mu\nu}$ of electromagnetism in which A_μ is the photon field, the difference being that photons are not self-interacting and thus the third term of Eq. (1.2) is not present. Symmetries of the gauge group give rise to the three "color" charges of QCD, which are analogous to the electric charge of SU(2). The eight linearly-independent generators of SU(3) correspond to the gluons which are the mediators of the color force, similar to the photons of quantum electrodynamics (QED). The self-interaction of

the colored gluons gives rise to the third term of Eq. (1.2) and is the reason why the strong force is strong.

The quarks are fundamental fermions which carry color and electric charge. They come in six flavors: up (u), down (d), charm (c), strange (s), top (t), and bottom (b). Three of these quarks (u, c, and t) carry a positive electric charge of $+2e/3$, and the other three (d, s, and b) carry a negative electric charge of $-e/3$, where e is the electric charge. The masses of the quarks vary widely, from as low as 1.7 MeV/c^2 for u to as high as 172 GeV/c^2 for t, extending over 5 orders of magnitude.

1.2.1 Hadrons and Color Confinement.

One of the distinguishing features of QCD, and a direct consequence of the third term of Eq. (1.2), is the fact that the coupling between colored objects increases with distance, whereas in QED, the coupling between electric charges tends to decrease with distance. The effect of this feature is that quarks, which carry a color charge, tend to stay bound together. If a quark q and an antiquark $\bar{q}$ are pulled apart from one another, the potential between them will continue to increase as the distance increases, until eventually it becomes energetically favorable to generate new $q\bar{q}$ pairs from the vacuum to bind with the original q and $\bar{q}$ rather than force them back together. For this reason, free color charge is never observed.

The particles that are made up of quarks are referred to as "hadrons," and they necessarily carry zero net color charge. Hadrons come in two forms: mesons and baryons. A meson consists of a quark and its antiquark bound together with opposite color charge (*e.g.* blue and anti-blue). Some examples of mesons are the pions, which are formed by combinations of u and d quarks and their antiquarks and are the mediators of the nuclear force. The most well-known participants of the nuclear force, protons and neutrons, are examples of baryons. A baryon consists of three quarks with three different color charges (*e.g.* red, green, and blue) which, together, are color neutral. The proton is uud, which results in an electric charge of $+1$, and the neutron is udd, which results in an electric charge

of 0.

There are many other species of hadrons that can be built by other combinations of the quarks. Two important mesons that will be mentioned repeatedly in thesis are the J/ψ ($c\bar{c}$) and the Υ ($b\bar{b}$).

1.2.2 The Quark-Gluon Plasma

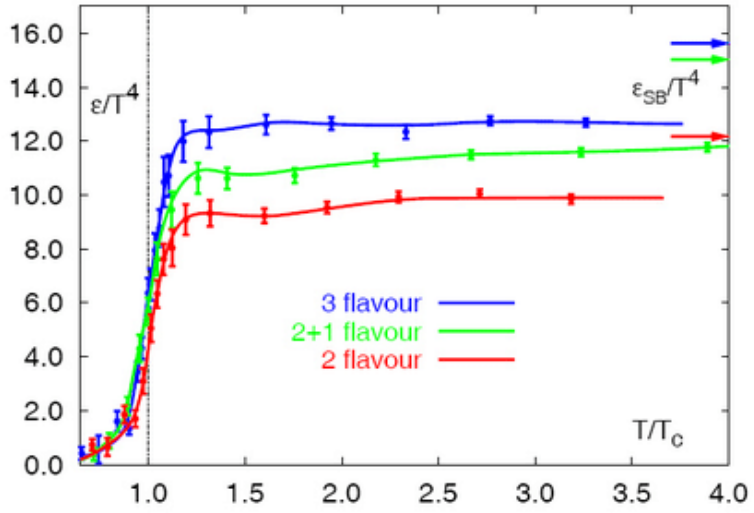

Figure 1.1. The energy density of QCD matter divided by temperature raised to the fourth power, calculated using lattice QCD. This quantity is proportional to the number of degrees of freedom of the system. The three lines correspond to different types of simulations. The red line assumes two massless quarks, the green line assumes two massless quarks and a massive strange quark, and the blue line assumes three massless quarks. A rapid increase is seen near the temperature T_c. The arrows at the upper right show the expected limit for an ideal gas of bosons. The figure is from Ref. [3].

At high temperatures and densities, hadronic matter is believed to cross over to a state of deconfined quarks and gluons called the quark-gluon plasma (QGP) [4, 5]. Calculations from lattice QCD [3, 6, 7] have predicted a jump in the degrees of freedom of QCD matter in the temperature range of 150–170 MeV. An example is shown in Fig. 1.1 [3]. The y-axis shows energy density over temperature to the fourth power, which is a number proportional to the degrees of freedom of the system. The x-axis shows the temperature of the medium over a specific temperature T_c. The three lines correspond to different types of simulations. The

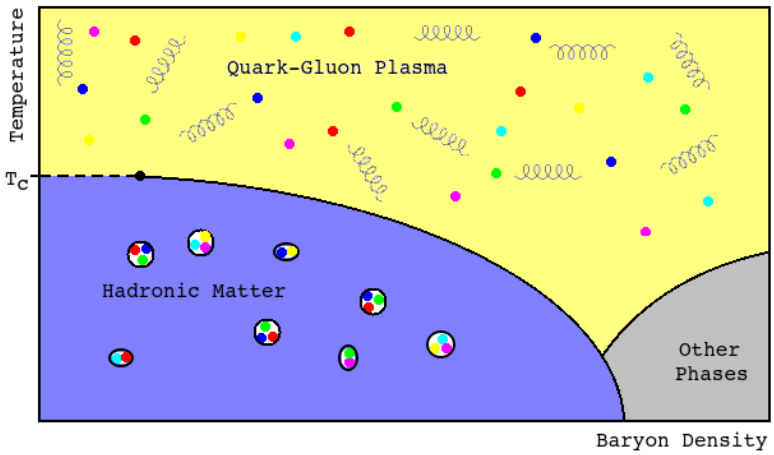

Figure 1.2. Schematic of the Phase Diagram of QCD Matter. The y-axis is temperature and the x-axis is net baryon density. At low temperatures and low (normal) densities, QCD matter is confined to hadrons. At low baryon density and high temperatures, there is a crossover (displayed as a dashed line) to the QGP phase. At low temperature and high net baryon density, there is a first-order phase transition, indicated by the solid black line. The point where the dashed line and the solid line meet is the critical point.

red line assumes two massless quarks, the green line assumes two massless quarks and a massive strange quark, and the blue line assumes three massless quarks. The arrows on the right, with the label $\varepsilon_{\mathrm{SB}}/T^4$, represent where the value should be if the QGP could be treated as an ideal Stefan-Boltzmann (SB) gas of quarks and gluons. The fact that it is different from the actual value indicates that the microscopic degrees of freedom are not simply free quarks and gluons. Some explanations for this discrepancy include possible colored bound states of quarks, or hadron states that are modified and become part of the gas [8].

The goal of our research is to narrow down the QCD equation of state and study the properties of the QGP phase, which include the initial-state temperature, viscosity, and density. A cartoon phase diagram of QCD matter is shown in Fig. 1.2 with temperature on the y-axis and net-baryon density on the x-axis. At low temperatures and densities, QCD matter is confined, forming the familiar hadrons. At high temperatures and densities, there is a phase transition to the QGP. There are also other phases at high density and low temperature that are not studied in this thesis. The lattice QCD calculation mentioned previously, where a rapid

crossover is found, requires zero net-baryon density [7]. The dashed line represents the crossover, and the solid line represents a first-order phase transition [9]. Since the line representing the first-order transistion cannot end at the temperature axis by virtue of the lattice QCD calculations, it must end at a critical point somewhere in the midst of the phase diagram [10]. The exact location is currently being investigated at RHIC in the Beam-Energy Scan project [11].

The critical temperature estimated from lattice QCD with two massless quark flavors is about 154 MeV [12, 13], which is equivalent to approximately 1.8×10^{12} K. This coincides with the temperature of the universe about 1 microsecond after the big bang, during the quark epoch [14]. For this reason, an understanding of the thermodynamic properties of the quark-gluon plasma allows for a better understanding of the very early universe.

1.3 Heavy-ion Collisions

The extreme conditions necessary to create a QGP make it relatively rare in nature. Current cosmological theories suggest that the QGP existed shortly after the Big Bang at high temperatures and low net baryon densities, corresponding to the upper left region of the QCD phase diagram. The QGP may also be present in the cores of neutron stars, at low temperature and very high net baryon density, corresponding to the lower right region of the QCD phase diagram. It is possible to create the QGP in a laboratory by colliding heavy ions at high energies in accelerators like the Large Hadron Collider (LHC) at CERN.

Our heavy ion of choice at the LHC is Lead-208 (^{208}Pb), because it is the heaviest stable nuclide. The reason for its stability lies in the fact that it has 82 protons and 126 neutrons, both of which are "magic" numbers in the Shell Model of nuclear physics. This also ensures that ^{208}Pb is a spherical nucleus, which greatly simplifies the geometry of the collisions. An example of what a collision between two Pb nuclei could look like is shown in Fig. 1.3. The image was generated from a simulation based on the Glauber Model (described in section 1.3.1). Notice that

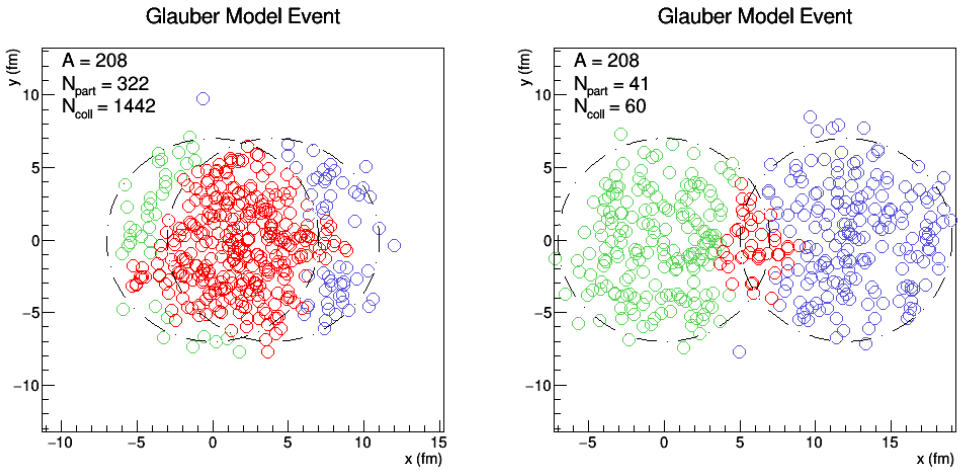

Figure 1.3. Heavy-ion collision events generated by a Glauber Model simulation. The impact parameter is 4 fm in the left plot and 12 fm in the right plot. The quantities A, N_{part}, and N_{coll}, are the number of nucleons in each nucleus, the number of nucleons participating in the collision (shown in red), and the total number of binary collisions that occured, respectively.

the amount of overlap between the two nuclei plays a major role in what happens in the collision.

The evolution of a heavy-ion collision is described schematically in Fig. 1.4. The two nuclei start colliding at time $t = 0$. The colliding matter forms a medium, possibly a QGP, which proceeds to expand under pressure, as shown in the figure on the left. On the right, a space-time diagram depicts the different stages of the expansion. A very short time after the collision, on the order of 1 fm/c ($\sim 3 \times 10^{-24}$ s), the system thermalizes. This is marked by the red hyperbola, labeled with a temperature of 230 MeV ($\sim 3 \times 10^{12}$ K), which is about 10^5 times hotter than estimated temperatures of the Sun's core. That is well above the crossover temperature, which means the system is in the deconfined QGP phase. The energy density of this medium is about ~ 3 GeV/fm^3, or about 20 times the density of normal nuclear matter. After time $t \sim 10$ fm/c, as the fireball continues to expand, it cools down below the crossover temperature and all the quarks become confined again. At about this same time, chemical freeze-out occurs, which means that all inelastic collisions have ceased and the particle yields are fixed, except for decays. After a few times 10 fm/c, thermal freeze-out occurs when all elastic collisions have

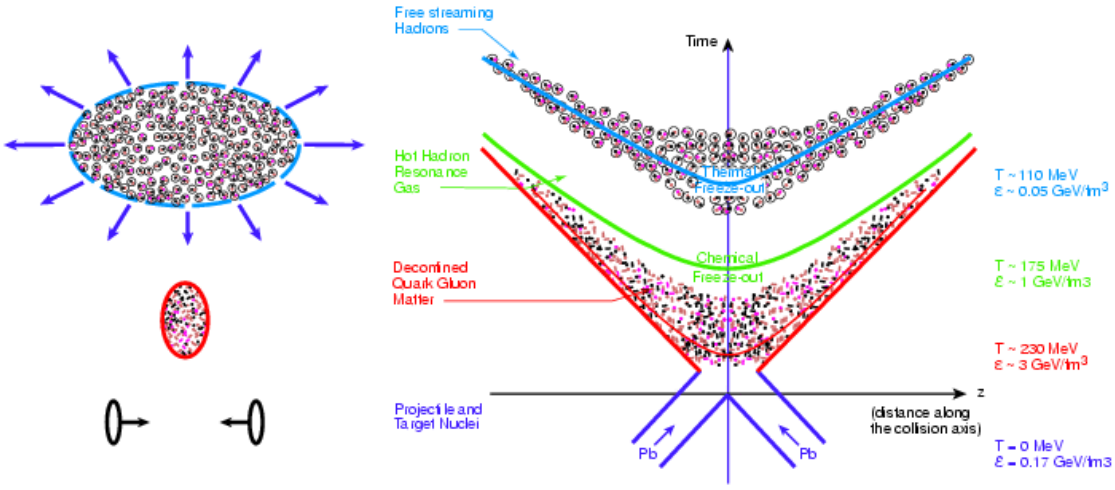

Figure 1.4. Schematic of the evolution of a heavy-ion collsion. The figure on the left displays a cartoon of the collision of two heavy ions and the expansion of the resulting medium, with time increasing up the page. The figure on the right is a space-time diagram displaying the different stages of the collision evolution, with time increasing up the page. The red, green, and blue lines denote specific times in the evolution, with the corresponding temperature and energy density displayed to the right. The red line marks the time of thermalization, where the medium is in the QGP state. The green marks the time of chemical freeze-out where the inelastic collisions cease. The blue line marks the time of thermal freeze-out where elastic collisions cease. Image from Ref. [15].

ceased and the particle momenta are now fixed until they reach the detectors. This is estimated to happen around $T \approx 110$ MeV, which corresponds to the cyan line in the figure. A more thorough study of the physics of heavy-ion collisions can be found in Ref. [16].

1.3.1 The Glauber Model and Centrality

The Glauber model is used to study the initial geometry of nucleus-nucleus collisions. In this model, the nucleus is treated as a collection of spherical nucleons. The structure of the generated nucleus is inspired by the Woods-Saxon potential

$$V(r) = -\frac{V_0}{1 + \exp[(r - R)/a]}, \tag{1.3}$$

where r represents the distance from the center of the nucleus and R is the nuclear radius. The parameter a is called the "skin depth" or "surface thickness," and is related to the distance over which the density function decreases from its maximum

value to zero at the edge of the nucleus. The depth of the potential well is denoted by the parameter V_0. This potential describes approximately the forces applied on each nucleon in the nuclear shell model.

The Glauber model assumes that the nuclear density profile matches the Woods-Saxon potential in Eq. (1.3), and can thus be described by the formula

$$\rho(r) = \frac{\rho_0}{1 + \exp[(r - R)/a]}. \tag{1.4}$$

The two nuclei are generated by randomly positioning 208 nucleons according to this density distribution.

In the case of two perfectly spherical nuclei, a collision will occur whenever the distance of closest approach between the centers of the two spheres is less than the sum of the two radii. The probability of a collision is the same for any size of sphere as long as the sum of the radii remains the same. Therefore, the probability of a collision between two spheres with equal radius R is equivalent to the probability of a point-sized projectile colliding with a sphere of radius $2R$. The pair of nuclei are thus assigned a random impact parameter b according to a linear probability distribution derived from the circular differential area $2\pi bdb$. The probability function is allowed to extend beyond $2R$ to be able to capture the most peripheral collisions which can occur when some nucleons are randomly placed very far from the center of the nucleus, due to the non-zero skin depth in the nuclear density distribution. In this simulation, a cutoff at $3R$ was sufficient to capture the full distribution of collision events.

The cross-sectional area occupied by the simulated nucleons is set such that the probability of a collision between two nucleons matches the true nucleon-nucleon inelastic cross section ($\sigma_{NN} = 67.6$ mb [17]) at the desired energy. Whenever a nucleon from one nucleus and a nucleon from the other nucleus have overlapping areas, it is counted as a collision, and the colliding nucleons are called participants. In practice, since the nucleons are modeled as hard spheres, this simply involves checking whether the nucleons are within 2 radii of each other, where the effective radius of each nucleon is taken from the nucleon-nucleon cross section as $r_N =$

$\sqrt{\sigma_{NN}/\pi}$. In the particular event shown in Fig. 1.3 on the left, there are 322 participants and 1442 binary nucleon-nucleon collisions.

The expected number of participants (N_{part}), and the number of binary nucleon-nucleon collisions (N_{coll}), can be calculated by use of the thickness function $T_A(s)$ which is defined as

$$T_A(s) = \int \rho(s,z)dz, \tag{1.5}$$

where ρ is again the nuclear density function defined in Eq. (1.4) normalized to the number of nucleons A in the nucleus, z is the longitudinal coordinate, and s is the radial distance from the z-axis. At a given impact parameter b, the nuclear overlap function of nuclei A and B is given by [16, 17]

$$T_{AB}(b) = \int T_A(s)T_B(|\vec{b}-\vec{s}|)d^2s. \tag{1.6}$$

Given these definitions, the expected values of N_{part} and N_{coll} are given by

$$\begin{aligned} N_{\text{part}}(b) = \int \Big[& T_A(s)\big(1-\exp[-\sigma_{NN}T_B(|\vec{b}-\vec{s}|)]\big) \\ & + T_B(|\vec{b}-\vec{s}|)\big(1-\exp[-\sigma_{NN}T_A(s)]\big)\Big]d^2s \end{aligned} \tag{1.7}$$

and

$$N_{\text{coll}}(b) = \sigma_{NN}T_{AB}(b). \tag{1.8}$$

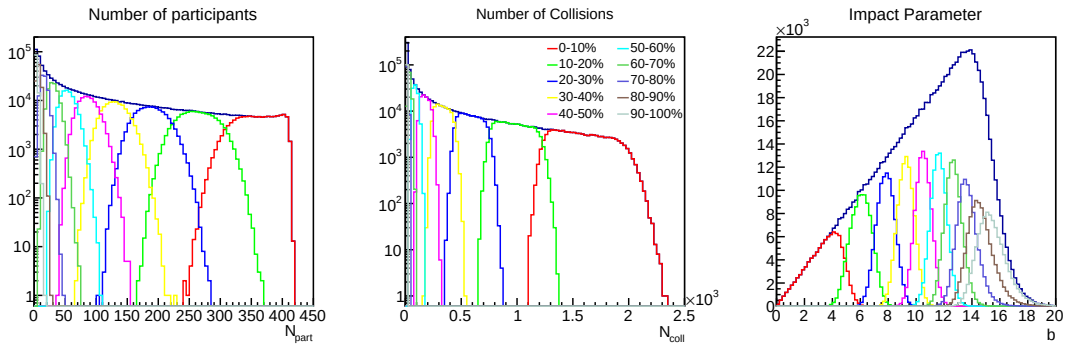

Figure 1.5. Distributions of the number of participating nucleons (left), number of collisions (middle) and the impact parameter (right) in a Glauber model simulation.

The distributions of N_{part}, N_{coll}, and b, after one million simulated events are shown in Fig. 1.5. The N_{coll} distribution, after simulating particle production, can

be used to fit the Glauber model to real data from the CMS detector in some activity-related variable such as the transverse energy deposited in the forward hadron calorimeters E_T^{HF}, as shown in Fig. 1.6. Particle production is simulated by convoluting N_{coll} with a specialized negative binomial distribution,

$$f_{\mathrm{NBD}}(x) = \frac{\Gamma(x+\mu)(k/\mu)^x}{\Gamma(x+1)\Gamma(\mu)(k/\mu+1)^{x+\mu}}, \tag{1.9}$$

where x symbolizes the amount of transverse energy deposited in the HF by the particles exiting a particular collision.

In this case, the fit was carried out by applying the Nelder-Mead algorithm [18] with a 3-point simplex in k-μ phase space. At each fitting step, the point of the trio with the largest χ^2 is moved to the opposite side and slightly closer to the other two points. At each step, the normalization parameter was also optimized to reduce the χ^2. Step by step, the 3 points converge to the point that minimizes χ^2. The process ends when the distance in phase space between the 3 points is smaller than a predefined threshold, which in this case was 0.00001. The progress of the three points and the path of the best fit are shown in Fig. 1.7. The best and worst χ^2 at each step are also shown in the bottom panel. The blue line represents the current best χ^2 at each step, which doesn't change until one of the other points manages to go lower. The red line represents the current worst, and since the worst point always changes, the red line changes every step.

The amount of activity in a collision event is usually quantified by a variable called "centrality," which is also related to the amount of nuclear overlap in the collisions, as we have seen. A head-on collision corresponds to about 0% centrality, and a grazing collision would be around 100% centrality. For example, the event shown in Fig. 1.4 on the left is in the 0-10% centrality range. The Glauber model is one way in which we can get estimates of centrality. The centrality bins are estimated by integrating E_T^{HF} starting from the right side. The top 10% most active events form the 0-10% centrality bin, the next 10% most active events form the 10-20% centrality bin, and so on up to 100%. The average values of N_{part}, N_{coll}, and the impact parameter b are displayed in Table 1.1 in bins of centrality,

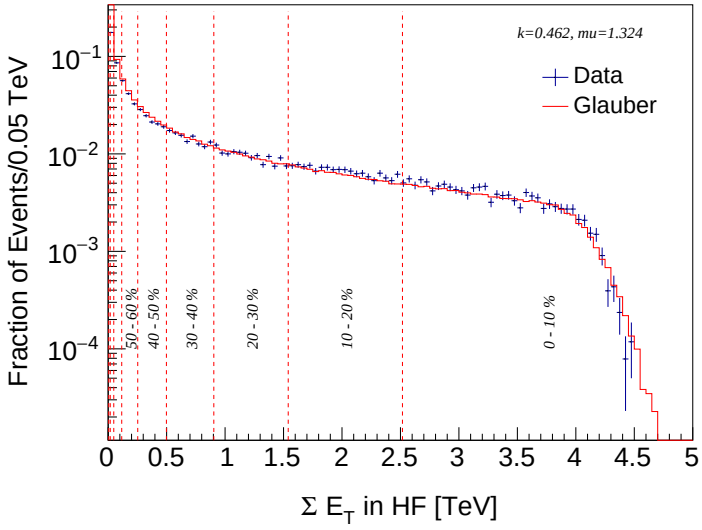

Figure 1.6. Centrality bins determined by a Glauber model fit to the E_T^{HF} distribution in minimum bias Pb + Pb data at $\sqrt{s_{NN}} = 5.02$ TeV.

as determined by this Glauber model. The documentation and results of a more sophisticated Glauber model can be found in Ref. [17].

Centrality is therefore more rigorously defined as the percentile of the total inelastic hadronic cross section. In CMS, it is found using the total energy in both HF calorimeters, where larger centrality percentiles correspond to smaller energy deposits in the calorimeters. Larger centrality percentiles are therefore associated with fewer tracks and fewer nucleon-nucleon collisions, as seen in Fig. 1.5 and Table 1.1, which is indicative of less overlap of the two colliding nuclei corresponding to larger impact parameters.

1.3.2 Experimental Observables

There are several thermodynamic properties of the QGP that can be measured with heavy-ion collisions, including the temperature, the density, and the viscosity. These quantities are measured indirectly via the study of particle distributions exiting heavy-ion collisions. For example, one way to probe the temperature of the QGP during the early stages of the collision is to measure the amount of suppression in the production of quarkonia, described in the next section. If there

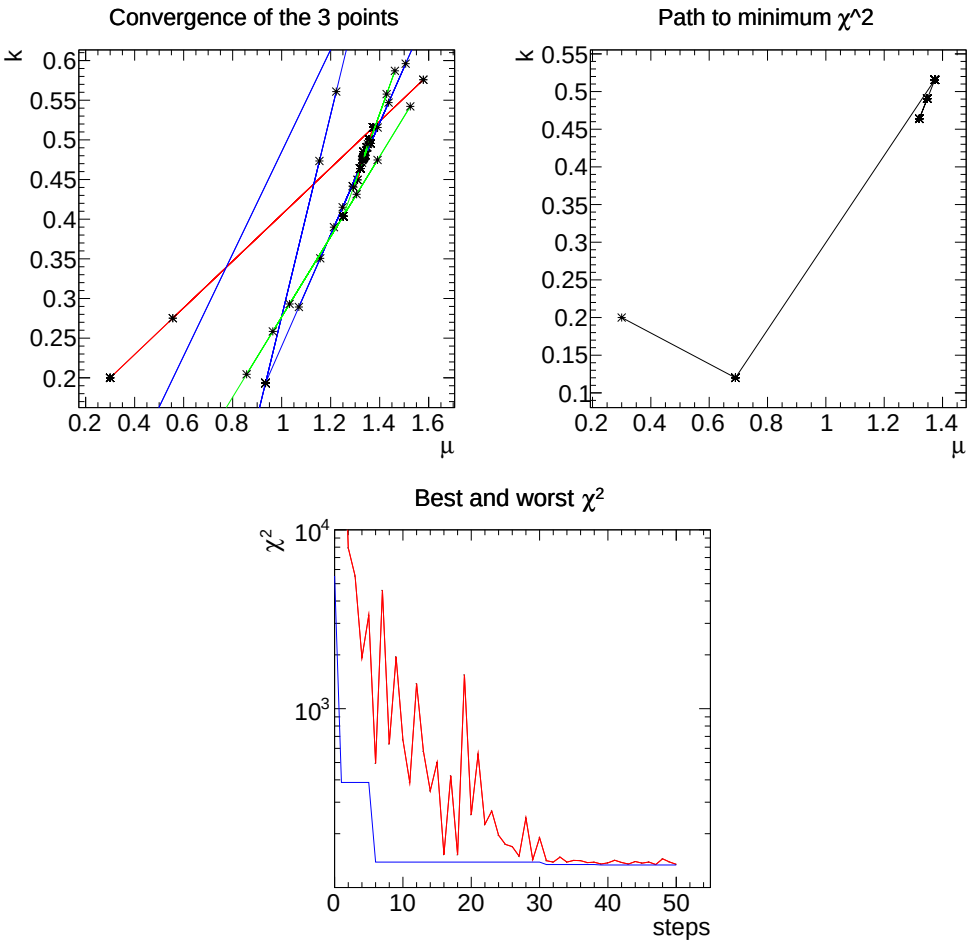

Figure 1.7. Progression of the k and μ parameters towards the minimum χ^2 in a Glauber model fit using a 3-point simplex. The upper-left plot displays the progress of all three points. The upper-right plot displays the progress of only the best fit at each step. The lower plot displays the best (blue) and worst (red) χ^2 for each step of the fitting process.

Centrality	$\langle N_{\text{part}} \rangle$	$\langle N_{\text{coll}} \rangle$	b
0 - 10%	357.116	1585.99	3.34351
10 - 20%	262.462	972.111	6.04368
20 - 30%	187.956	585.72	7.81843
30 - 40%	130.294	334.727	9.26266
40 - 50%	86.4019	178.383	10.5074
50 - 60%	53.6833	86.7814	11.6331
60 - 70%	30.4673	38.1272	12.6763
70 - 80%	15.4412	15.2242	13.6798
80 - 90%	6.9513	5.55382	14.6417
90 - 100%	3.02945	1.87529	15.432
0 - 20%	309.788	1279.05	4.69362
20 - 40%	159.124	460.22	8.54056
40 - 60%	70.0393	132.573	11.0704
60 - 80%	22.9719	26.7026	13.1769
80 - 100%	4.99007	3.71427	15.0369
0 - 50%	204.858	731.441	7.39484
50 - 100%	21.8918	29.4844	13.6151
0 - 100%	113.366	380.43	10.5053

Table 1.1. The average number of participating nucleons $\langle N_{\text{part}} \rangle$, average number of nucleon-nucleon collisions $\langle N_{\text{coll}} \rangle$, and the average impact parameter b for each centrality class, as estimated from a Glauber model simulation.

is *a priori* knowledge of the particular temperatures at which these quarkonia dissociate, then their suppression serves as a kind of thermometer for the medium.

The density and viscosity can be probed via multi-particle correlations and elliptic flow measurements, which can be extracted by measuring the azimuthal distribution of particles exiting the collisions, and comparing the results to hydrodynamic flow predictions which take density and viscosity as inputs. Recent comparisons between collision data and simulations of hydrodynamic flow have determined that the QGP exhibits liquid-like collective properties, and have estimated the QGP shear viscosity (a quantitative measure of the internal fluid friction) to be very near the theoretical lower limit [19, 20, 21, 22], indicating that the QGP is nearly a perfect fluid (*i.e.* a fluid that flows with zero viscosity).

1.4 Quarkonia Suppression

The term quarkonium refers to a flavorless meson composed of a quark and its antiquark. The term usually only refers to mesons composed of the heavy quarks c and b, such as charmonia (*i.e.* the J/ψ) and bottomonia (*i.e.* the Υ family shown in Fig. 1.8), and not the lighter mesons such as pions which are actually mixtures of $q\bar{q}$ states. The top quark is too heavy and short-lived to form such bound states.

Quarkonia are good probes of the early state of a heavy-ion collision. The high mass of the bottom and charm quarks ensures that they are unlikely to be generated thermally in the expanding fireball, which means their production is dominated by the initial hard scatterings [23]. This is especially true for the heavy Υs, which have masses displayed in Table 1.2. Quarkonia such as the J/ψ and the Υs also have a significant dimuon decay channel. The dimuon branching ratios are also listed in Table 1.2. Muons interact only weakly in matter and thus carry the information from the quarkonium state to the muon detectors with minimal modification.

Quarks bound in a meson state are not motionless in the rest frame of the meson. The lighter quarks such as u and d tend to move at relativistic speeds

Table 1.2. The masses of heavy quarkonia and their branching ratios to dimuons [24].

Quarkonium state	Mass (MeV/c^2)	Dimuon branching ratio
J/ψ(1S)	3096.900 ± 0.006	5.961 ± 0.033 %
Υ(1S)	9460.30 ± 0.26	2.48 ± 0.05 %
Υ(2S)	10023.26 ± 0.31	1.93 ± 0.17 %
Υ(3S)	10355.2± 0.5	2.18 ± 0.21 %

inside their respective mesons. In a heavy quarkonium, however, the speeds of the constituent quarks are relatively small, and the relativistic effects are significantly reduced. This allows for heavy quarkonia to be approximately described by non-relativistic QCD (NRQCD) [25].

Quarkonia are tightly bound. The size, estimated by the binding radius, is very small and the binding energy is very large, so that the bound state cannot be broken by interactions with normal hadrons. However, in the presence of a QGP, quarkonia can interact with and be modified by color fields in the QGP medium, and can thus be used as a testing ground for processes such as deconfinement. It has been predicted that QCD color screening in a QGP will result in a suppression of the production of quarkonia [26]. This suppression has been observed in experiments as modification of quarkonia yields. For example, Υs are suppressed in Pb + Pb collisions compared to p+p collisions at the same center-of-mass energy per nucleon pair [27].

The higher Υ states have successively smaller binding energies and larger size, which causes them to dissociate or "melt" at lower temperatures than the Υ(1S). This allows for the study of additional effects such as sequential suppression [27, 28, 29]. These excited Υ states, along with other bottomonia states such as the χ_b states, can decay to the lower states and contribute to the total yield [4, 16]. The decay modes of the bottomonia family are displayed in Fig. 1.8. This process is called "feed down," and it must be taken into account in theoretical predictions of Υ yields. The Υ(4S) and higher bottomonia states that lie above the $B\bar{B}$ threshold

decay strongly to $B\bar{B}$ and therefore do not contribute to feed down.

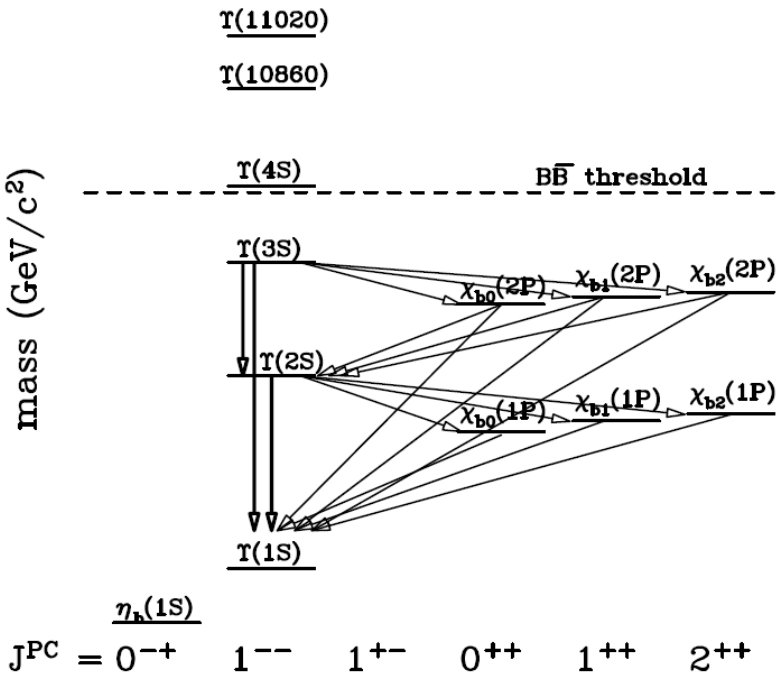

Figure 1.8. Illustration of the decay paths of the bottomonia family, with quantum numbers along the x-axis and mass on the y-axis. Thick arrows represent hadronic decays and thin arrows represent radiative decays. Figure taken from Ref. [16].

If the temperature of the medium is high enough, most of the members of the bottomia family will melt before decaying. The Υ(1S) is the most tightly bound of all the quarkonium states, and thus requires a very high temperature to melt. It follows that the Υ(1S) yields observed in extremely hot heavy-ion collisions at CMS are mostly from direct production ($\sim$67%) [30], because the melting of the excited states prevents the feed down decays from occurring.

The modification of Υ yields in Pb + Pb compared to $p + p$ can be seen in the plot on the right in Fig. 1.9, which shows a fit to the dimuon invariant mass spectrum from Pb + Pb collisions, and overlaid in red, what we would expect if Pb + Pb collisions were simply a superposition of many $p + p$ collisions. The ratio of the yield in Pb + Pb to the yield we would expect in this simpler scenario is

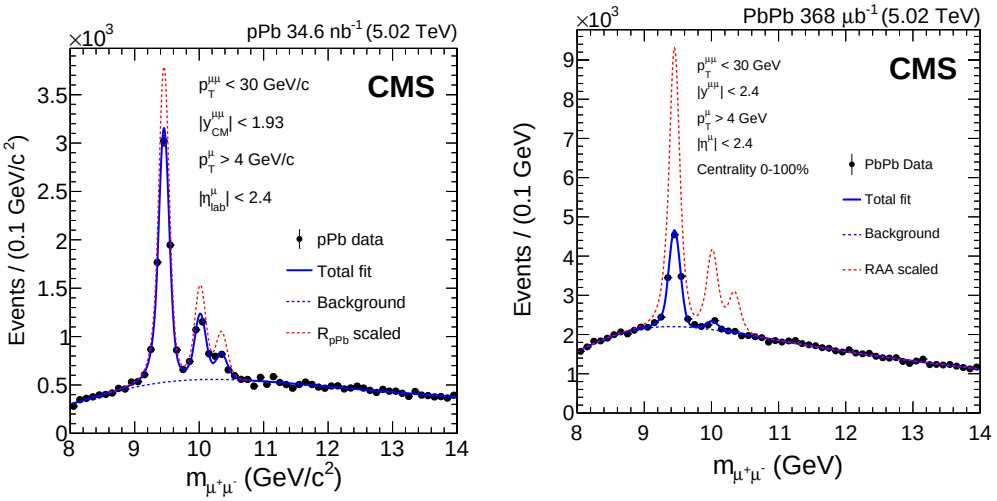

Figure 1.9. Plots of the invariant mass spectrum of Υ candidates in $p+\mathrm{Pb}$ [31] (left) and $\mathrm{Pb}+\mathrm{Pb}$ [27] (right). The yields that would be expected assuming no suppression are overlaid with the dashed red line.

called the nuclear modification factor, R_{AA}. The definition of this quantity is

$$R_{AA}^{\Upsilon}(p_T, y) = \frac{N_{AA}^{\Upsilon}(p_T^{\Upsilon}, y^{\Upsilon})}{\langle T_{AA}\rangle \sigma_{pp}^{\Upsilon}(p_T^{\Upsilon}, y^{\Upsilon})}. \tag{1.10}$$

The quantity N_{AA}^{Υ} is the Υ yield in $\mathrm{Pb}+\mathrm{Pb}$ and σ_{pp}^{Υ} is the Υ production cross section in $p+p$, given by

$$\sigma_{pp}^{\Upsilon}(p_T^{\Upsilon}, y^{\Upsilon}) = \frac{N_{pp}^{\Upsilon}/(a\cdot\varepsilon)}{\mathcal{L}_{int}\Delta p_T^{\Upsilon}\Delta y^{\Upsilon}}, \tag{1.11}$$

where a and ε are the acceptance and efficiency described in Ch. 5, N_{pp}^{Υ} is the Υ yield in $p+p$, and $\mathcal{L}_{int}$ is the integrated luminosity. The quantity T_{AA} is the nuclear overlap function, defined in Eq. (1.6), from which it can be shown that $\langle T_{AA}\rangle = \langle N_{\mathrm{coll}}\rangle/\sigma_{pp}^{\mathrm{in}}$ (see Eq. (1.8)), where $\sigma_{pp}^{\mathrm{in}}$ is the total inelastic $p+p$ cross section. In the case of $p+\mathrm{Pb}$, under the hypothesis of A-scaling, the nuclear modification factor is

$$R_{pA} = \frac{\sigma_{pA}^{\Upsilon}(p_T^{\Upsilon}, y^{\Upsilon})}{A\sigma_{pp}^{\Upsilon}(p_T^{\Upsilon}, y^{\Upsilon})}, \tag{1.12}$$

where σ_{pA}^{Υ} is the Υ production cross section in $p+\mathrm{Pb}$.

There are a few effects due to the QGP that may cause this suppression. One major possibility is color screening, which is analogous to Debye screening in electromagnetism. When a $b\bar{b}$ pair is traveling through a QGP, the presence of other

color fields tends to screen the fields of the $b\bar{b}$ pair. The spatial scale at which the screening occurs is called the screening length μ, and it decreases with increasing temperature. The Υ(1S) is tightly bound with a very small radius, and thus μ must be very small (*i.e.* the temperature must be very high) to have any effect. If μ is small enough, then the screening quarks and antiquarks may be able to pass between the b and $\bar{b}$ and bind with them, causing the Υ state to dissociate. This would reduce the total yield of Υs in Pb + Pb collisions.

At temperature $T = 0$, the real part of the $q\bar{q}$ potential can be written approximately as [16]

$$Re\{V(r)\} = \sigma r - \frac{\alpha_c}{r}, \tag{1.13}$$

where r is the distance between the q and $\bar{q}$, α_c is the coupling of the Coulomb term, and σ describes the strength of the linear term. In a hot medium, both terms of Eq. (1.13) are modified. At finite temperatures, the real part of the in-medium $q\bar{q}$ potential can be expressed as [32]

$$Re\{V(r)\} = -\frac{\Gamma[\frac{1}{4}]}{2^{3/4}\sqrt{\pi}}\frac{\sigma}{\mu}D_{-\frac{1}{2}}(\sqrt{2}\mu r) + \frac{\Gamma[\frac{1}{4}]}{2\Gamma[\frac{3}{4}]}\frac{\sigma}{\mu}, \tag{1.14}$$

where Γ is the gamma function, $D_{-\frac{1}{2}}$ is a parabolic cylinder function, and μ is the screening length, which depends on the temperature.

The graphs in Fig. 1.10 are a depiction of how the in-medium $q\bar{q}$ potential V changes with temperature as a result of color screening [32]. The left side shows the real part of V as a function of distance, and the right side shows the imaginary part. The different colored lines represent the different temperatures, with purple being the coldest and red being the hottest, where T_c represents the crossover temperature. At low temperatures, the linear part of the potential tends to be dominant at large r and severely limits the distance that the bound q and $\bar{q}$ can move away from each other. At high temperatures, the linear part is seen to be depressed, and it becomes possible for the q and $\bar{q}$ to escape to large distances. The imaginary part of the potential, shown on the right, affects the width of the quarkonium state and thus its lifetime. It is related to the scattering (Landau

damping) and absorption (singlet-octet transition) of gluons from the medium which can lead to gluo-dissociation of the $q\bar{q}$ states (*i.e.* gluons break up the quarkonium state), and it is also seen to be modified by the temperature [32, 33].

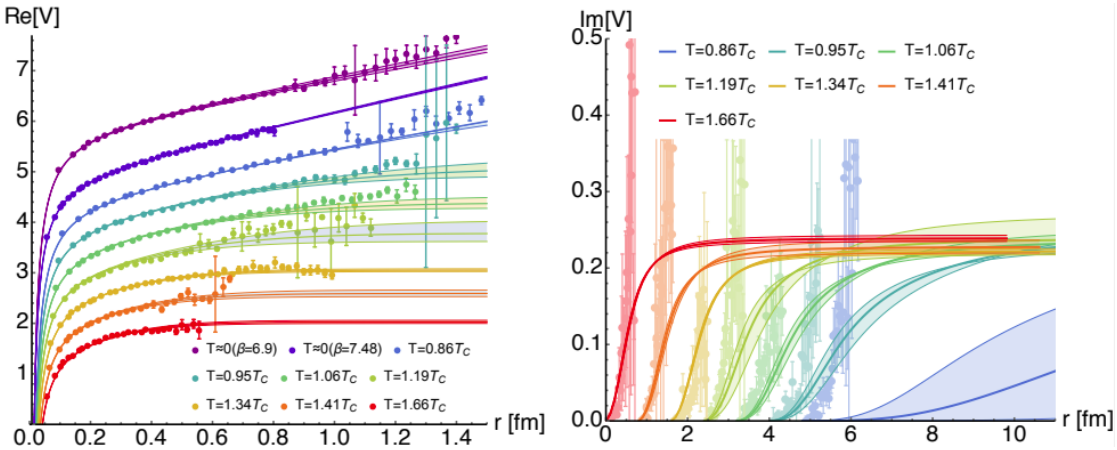

Figure 1.10. The real (left) and imaginary (right) parts of the in-medium $q\bar{q}$ potential V as a function of distance at various temperatures. Figure from Ref. [32].

Additional effects that modify the yields of quarkona include regeneration and recombination, which occur when an unbound $q\bar{q}$ pair that are close together in phase space bind together. They can be uncorrelated and moving freely in the QGP or they can be constituents of a melted quarkonium which remain close enough together in phase space that, at the time of hadronization, it is possible for the q and $\bar{q}$ to pair up and create or recreate the quarkonium. This tends to increase the quarkonium yields. This effect is fairly insignificant for Υs due to the fact that there are very few $b\bar{b}$ pairs, but it is important for the J/ψ [4, 5, 23, 34].

But the QGP may not be the only thing responsible for what the observed suppression. There are other effects due to the mere presence of a nucleus which have nothing to do with the QGP. We refer to these as cold nuclear matter effects [4]. The first are called shadowing and anti-shadowing, which are the names given to two modifications of the nuclear parton distribution function as shown in the example in Fig. 1.11. The parton distribution function (left) shows what fraction of the total momentum of the nucleon each parton has, as obtained from global analyses of data [35]. If a proton were comprised of only two up quarks and

a down quark, then we would expect each quark to peak around a value of 1/3, with the up peak twice as high as the down. In reality, the proton is filled with gluons and sea quarks generated by gluon splitting. The gluons and sea quarks tend to carry a small fraction of the total, and this tends to shift the x value of the u and d quarks down and create large tails in the distribution. The other types of sea quarks have only the large tails. The gluons follow the same pattern, but at a much larger scale. The gluon curve is scaled down by a factor of ten to be visible on the plot at a level similar to the quarks.

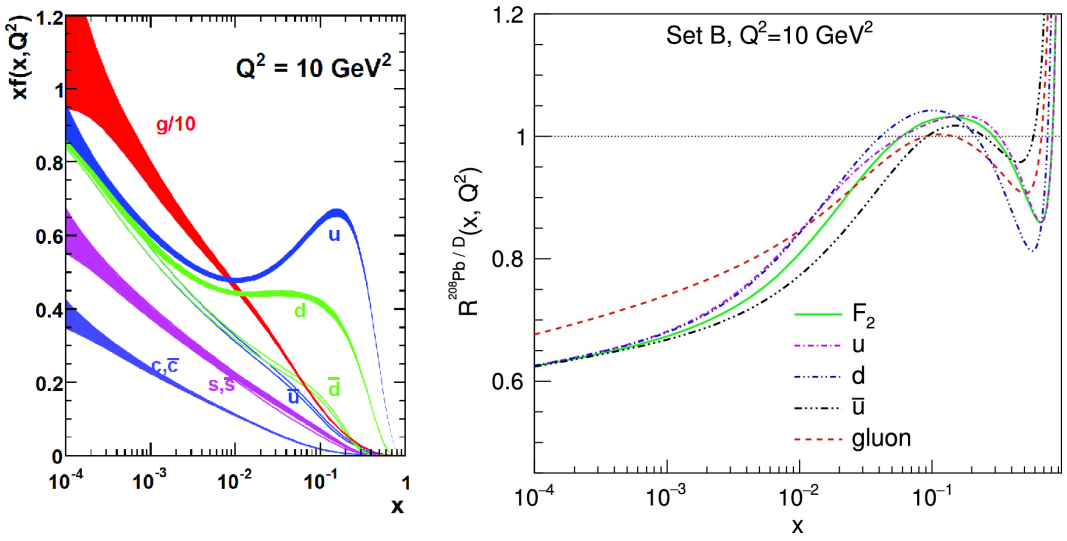

Figure 1.11. A proton parton distribution function (left) where the distributions of the various quarks and gluons are color-coded and labeled, taken from Ref. [35], and an example of the modification to the ^{208}Pb nuclear parton distribution function (right) compared to a deuteron, taken from Ref. [36].

The parton distribution function is modified in a nucleus compared to a nucleon. The modification is usually quantified as a ratio of the nuclear parton distribution function to that of a deuteron. An example of the modification for ^{208}Pb is shown in the figure on the right [36]. The parton distributions in the nucleus are lower than those in the deuteron in the low x region. This effect is referred to as shadowing. Around 0.1, there is an enhancement, which is referred to as anti-shadowing. At higher x, there is a local minimum referred to as the EMC effect because it was first discovered by the European Muon Collaboration (EMC) [37, 38, 39].

Some other non-QGP effects are energy loss as a particle propagates through the medium [40, 41], and the effects of comovers, which are particles in the same phase space which can interact with the quarkonium and dissociate it [42].

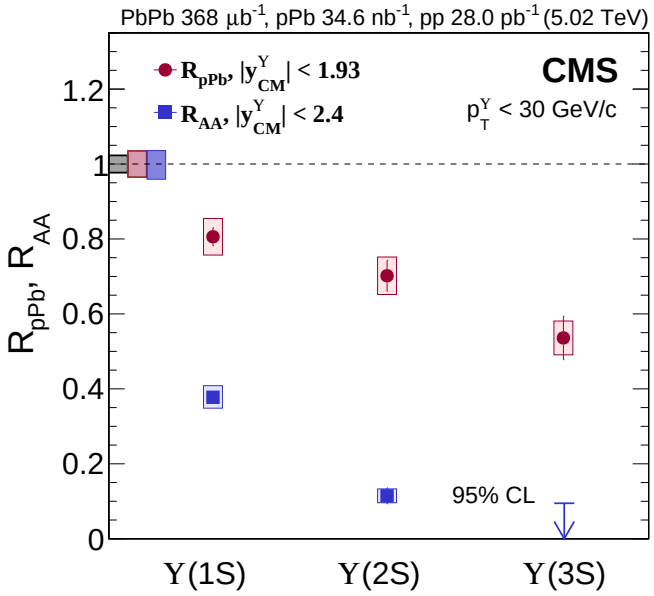

Figure 1.12. Comparison of the nuclear modification factors R_{pA} (red circles) and R_{AA} (blue squares) of $\Upsilon(n\mathrm{S})$ (n = 1,2,3) from p+Pb and Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. In the case of Υ(3S) in Pb+Pb, a 95% confidence interval is displayed. Figure from Ref. [31].

Cold nuclear matter effects can be separated from QGP effects by comparing $p+\mathrm{Pb}$ collisions with $\mathrm{Pb}+\mathrm{Pb}$ and $p+p$ collisions. A QGP is not expected to be created in typical p+Pb collisions, so any suppression of heavy-quarkonia in p+Pb is probably due to cold nuclear matter effects. The fit to the dimuon invariant mass spectrum in $p+\mathrm{Pb}$ is shown on the left in Fig. 1.9. The Υ yields are seen to be suppressed in $p+\mathrm{Pb}$, even without the presence of a QGP [31].

The nuclear modification factors R_{pA} and R_{AA} extracted from the fits are presented in the plot in Fig. 1.12 [31]. The level of suppression is much more significant in $\mathrm{Pb}+\mathrm{Pb}$ than in $p+\mathrm{Pb}$, due to the additional QGP effects. It can also be seen that the suppression is successively more significant for the higher Υ states, *i.e.* the Υ(3S) is more suppressed than the Υ(2S), which is more suppressed than the Υ(1S). This phenomenon is referred to as sequential suppression, and, as

discussed earlier, it is because the higher mass states are less tightly bound. The larger size and smaller binding energy make color screening more effective on the higher mass states, and thus they tend to dissociate at lower temperatures than the Υ(1S) [28, 29].

1.5 Elliptic Flow

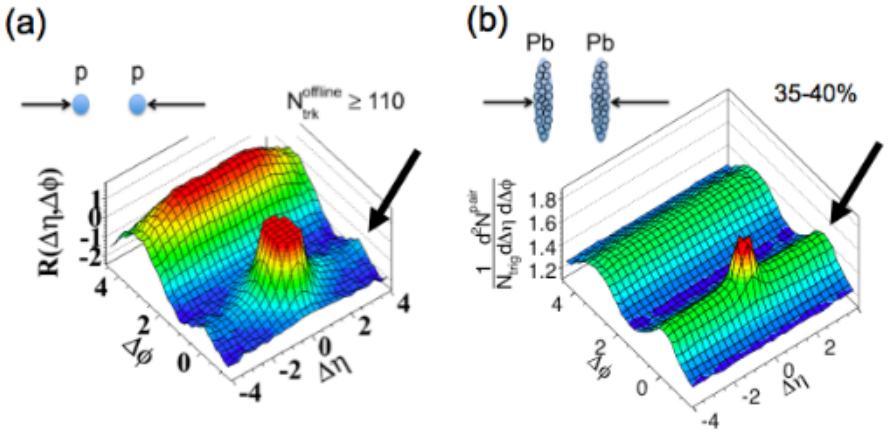

Figure 1.13. Two-particle correlations in high multiplicity events in $p+p$ at 7 TeV [43] (a) and in Pb + Pb at 2.76 TeV [44, 45] (b).

Apart from quarkonia suppression, another way of probing the QGP is through measurements of elliptic flow, which refers to a collective phenomenon seen in two-particle correlations in Pb + Pb that is not observed in $p+p$ [44, 45]. Such an effect can be seen in the plots in Fig. 1.13. Plot (a) shows a two-particle correlation function in p+p [43]. Given any particle in the event, this distribution shows where all the other particles are likely to be in pseudorapidity η and azimuthal angle ϕ relative to the selected particle. The two main features on the plot are the peak at $\Delta\phi = 0$ and $\Delta\eta = 0$ and the long ridge at $\Delta\phi = \pi$ which extends across all values of η. The peak corresponds to jets, which consist of many particles moving in approximately the same direction. The ridge on the other side describes particles that go in exactly the opposite azimuthal angle at many different values of pseudorapidity. This corresponds to the away-side jet. Both of these features are present in Pb + Pb, but there is one major difference which is indicated by

the black arrows in the figure. This ridge at $\Delta\phi = 0$ describes particles moving in the same azimuthal angle at all different rapidities. This is can only be explained by some sort of collective flow phenomenon. It is also present at $\Delta\phi = \pi$, but it is less obvious because the signal is mixed with the away-side jets. This pattern, which makes the shape of the function $\cos(2\Delta\phi)$, is referred to as elliptic flow.

In order to quantify the amount of elliptic flow present in a dataset, we parametrize the azimuthal distribution of particles as a Fourier series. Due to the symmetry of the collision region, as seen in Fig. 1.3, all the sine terms are zero and we get a series of cosines. If the system is properly centered, then the v_1 coefficient, which describes the horizontal offset, is also zero. Thus we obtain [46]

$$\frac{dN}{d\phi} \propto 1 + 2v_1\cos(\Delta\phi) + 2v_2\cos(2(\Delta\phi)) + 2v_3\cos(3(\Delta\phi)) + \cdots . \tag{1.15}$$

The elliptic flow is described by the coefficient of the second term, the v_2, which is given by

$$v_2 = \langle\cos(2(\Delta\phi))\rangle. \tag{1.16}$$

In fact, the terms "elliptic flow" and "v_2" are often used interchangeably.

Predictions of the magnitude of the elliptic flow are usually obtained via relativistic hydrodynamic simulations. The v_2 values measured from real data therefore tell us something about the inputs to those simulations, including viscosity, the equation of state, the speed of sound in the medium, and the level of thermalization achieved in the collision. For example, some estimates of v_2 from hydrodynamics are shown in Fig. 1.14 [22]. The y-axis variable is v_2 over eccentricity $e = \sqrt{1 - b^2/a^2}$ where a and b are the semi-major and semi-minor axes of the ellipse describing the shape of the collision region. The x-axis is total charged hadron multiplicity density per unit overlap area. The parameter being varied from line to line in these graphs is the specific shear viscosity (η/s). The data shown are estimates of the universal v_2 in gold-gold collisions at 200 GeV from RHIC [22]. The best match for the data seems to be with a specific shear viscosity of about 0.08. This is very near the theoretical lower limit of $(4\pi)^{-1} \approx 0.0796$ [19, 20, 21, 22],

which makes the QGP the closest to being a perfect fluid of any substance that has been observed. This discovery could not have been possible without v_2 measurements.

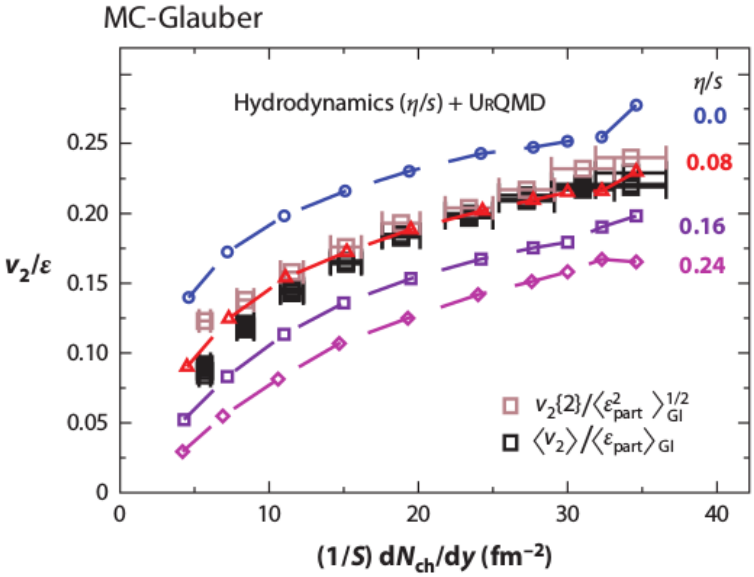

Figure 1.14. Eccentricity-scaled elliptic flow estimates from simulations using relativistic hydrodynamics in comparison with experimental data from gold-gold collisions at RHIC. The different colored lines correspond to different values of the specific shear viscosity (η/s). Figure from Ref. [22].

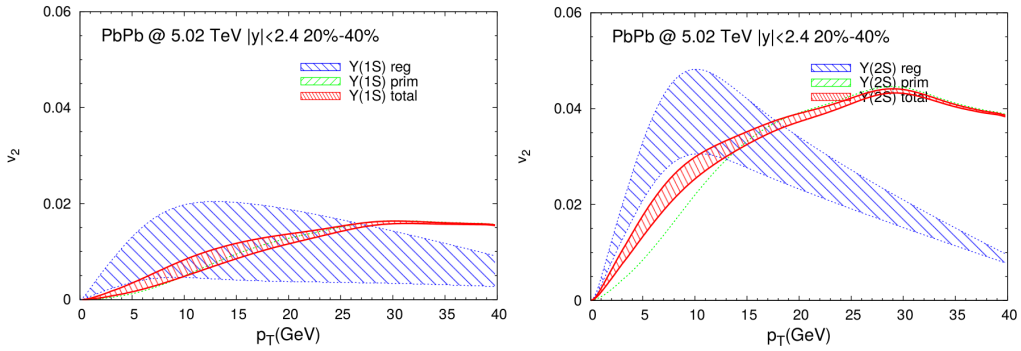

Figure 1.15. Elliptic flow predictions for Υ(1S) (left) and Υ(2S) (right) as a function of p_T in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. The red curve described the v_2 of the total signal, while the blue and green curves describe the v_2 of the regenerated and primordial components, respectively. Figure from Ref. [34].

Predictions have been made for the v_2 of Υs at $\sqrt{s_{NN}} = 5.02$ TeV [34]. Figure 1.15 shows predictions of the p_T dependence of the Υ(1S) v_2 on the left, and

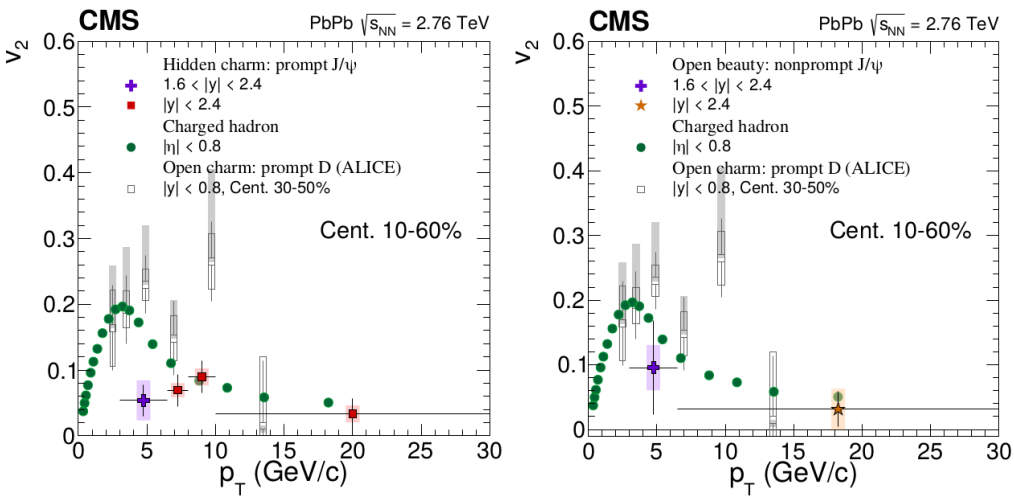

Figure 1.16. Elliptic flow of prompt (left) and nonprompt (right) J/ψ compared with the elliptic flow of charged hadrons and open charm as a function of p_T in Pb + Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. In both cases, the elliptic flow of J/ψ in forward rapidity ($1.6 < |y| < 2.4$) is displayed as a blue cross. The elliptic flow of J/ψ in full rapidity is shown as red squares on the left and a yellow star on the right. Figure from Ref. [47].

the Υ(2S) v_2 on the right. The v_2 of the regenerated and primordial components of the signal are shown by the blue and green curves, respectively, while the v_2 of the total signal is shown in red. For reference, the estimated v_2 of prompt and nonprompt J/ψs is shown in Fig. 1.16 along with the v_2 of charged hadrons and open charm [47]. The prefix "prompt" or "nonprompt" refers to whether the J/ψ was formed in the initial hard scattering or produced in the decay of b-hadrons [48]. As stated earlier, one of the benefits of studying the Υs is that the signal comes almost entirely from the inital hard scattering, which means the v_2 of the Υs can provide a clearer picture of the early stages of the collision.

The most obvious feature in Fig. 1.15 is that the v_2 of the Υ(1S) is predicted to be very small. This is due mainly to the high mass of the Υ(1S), which makes it more resistant to changes in momentum compared with the lighter particles (*e.g.* the charged hadrons), and results in comparatively little modification to the azimuthal distribution. The J/ψ also exhibits these characteristics, but to a lesser extent because it is not as massive as the Υ(1S), and with the added complication of the nonprompt contribution.

Another interesting feature is that the v_2 of the Υ(2S) is predicted to be about twice as large as that of the Υ(1S) [34, 49, 50]. The authors hypothesize that this is due to the fact that the Υ(1S) melts at a higher temperature than the Υ(2S), due to its smaller binding energy and larger width, and thus the suppression occurs earlier in the fireball evolution before it cools to lower temperatures. The suppression of the Υ(2S), on the other hand, is active to lower temperatures and thus to later times. The Υ(2S) therefore passes through more of the medium in a modified state and the resulting signal becomes more sensitive to the path length differences in the medium. Thus, the sequential suppression described in Sec. 1.4 may lead directly to the expected difference in elliptic flow signals of the Υ states.

The authors also claim that the total v_2 of the Υ(2S) is dominated by the primordial component, and that the regenerated component has very little effect [34]. However, Fig. 1.17 indicates that the primordial Υ(2S) yields are completely suppressed in the range of $N_{\rm part}$ that corresponds to centrality range (20-40%) in which their prediction was made. It can be seen on the left-hand side of Fig. 1.5 that the vast majority of events in that range have between 100 and 250 participants. The R_{AA} of the primordial component in Fig. 1.17 is near zero in that region.

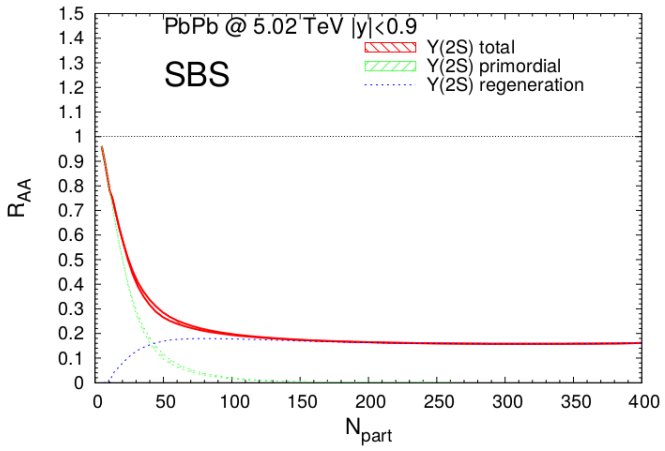

Figure 1.17. Nuclear modification of primordial (green), regenerated (blue), and total Υ(2S) yields as a function of $N_{\rm part}$ from simulation of Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. Figure from Ref. [34].

Another possible explanation for the factor of 2 difference is that the Υ(2S)

yield is much more influenced by regeneration than the Υ(1S) [34]. This means that the Υ(2S) mesons that exit the fireball were formed by b quarks that were unbound and interacting with the medium during the expansion, and therefore maintain a greater sensitivity to the properties of said expansion.

Chapter 2

Experimental Facilities

2.1 CERN LHC

The Large Hadron Collider (LHC), managed by the European Center for Nuclear Research (CERN), sits on the border between France and Switzerland. It is the largest particle accelerator in the world, with a circumference of 27 km. A schematic of the LHC accelerator complex is shown in Fig. 2.1. Protons and lead ions have different initial stages of acceleration. Protons are ionized from a hydrogen gas and then accelerated through the linear accelerator LINAC2. They then enter the Proton Synchrotron Booster (PSB) where they are accelerated up to speeds of $0.87c$ ($\sim$ 1 GeV of kinetic energy). They are then injected into the Proton Synchrotron (PS) where they are further accelerated. The lead ions also enter the PS after being initially accelerated through the LINAC3 and the Low Energy Ion Ring (LEIR). After reaching sufficiently high energies ($\sim$ 26 GeV for protons), the particles are sent into the Super Proton Synchrotron (SPS), which accelerates protons to $\sim$ 450 GeV. The final step is injection into the LHC. The two opposite-going beams cross one another at four points of intersection along the LHC ring where the detectors (ALICE, ATLAS, LHCb, and CMS) are placed.

The LHC can accelerate two beams of protons in opposite directions up to center-of-mass energies of 13 TeV per pair of colliding protons. The beams consist of bunches of protons, with about 1×10^{11} protons in each bunch. A pair of bunches

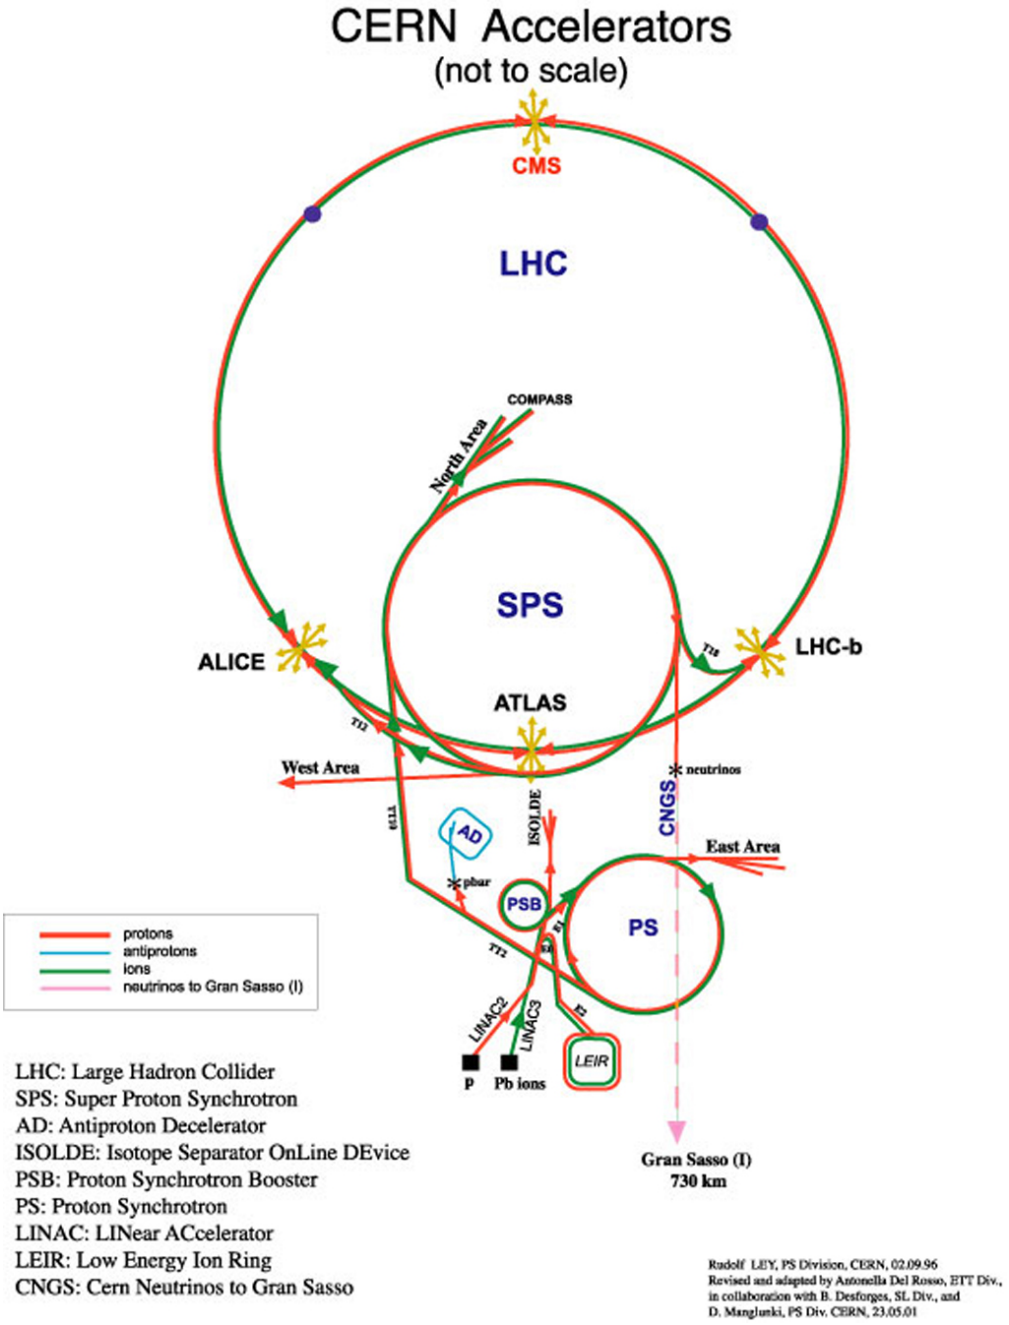

Figure 2.1. A schematic drawing of the CERN accelerators [51].

cross each other at an interaction point within a detector about every 25 ns. The rate of accelerated protons crossing the collision point is described by a quantity called luminosity [52]. The rate of collisions dR/dt is related to the luminosity through the total inelastic $p+p$ cross section σ_{pp}^{in} as follows:

$$dN/dt = \mathcal{L}\sigma_{pp}^{\text{in}} \tag{2.1}$$

where $\mathcal{L}$ is the luminosity, and is seen to have units of $\text{cm}^{-2}\text{s}^{-1}$. Integrating Eq. (2.1) over time gives the total number of collision events N in terms of the integrated luminosity $\mathcal{L}_{\text{int}}$:

$$N = \mathcal{L}_{\text{int}}\sigma_{pp}^{\text{in}} \tag{2.2}$$

where

$$\mathcal{L}_{\text{int}} = \int_{t_i}^{t_f} \mathcal{L}dt \tag{2.3}$$

The integrated luminosity is often used to characterize the size of a collision dataset. For any interaction of interest, the cross section of the interaction can be multiplied by $\mathcal{L}_{\text{int}}$ to give an estimate of the total number of events in which that interaction occurred. The integrated luminosity has units of inverse area, usually cm^{-2} or barns, b^{-1}, where $1\ \text{b} = 10^{-24}\ \text{cm}^2$.

Collision datasets are also characterized by the center-of-mass energy per nucleon pair $\sqrt{s_{NN}}$. With the same LHC configuration, the collision energy varies depending on the type of particle accelerated. In general, the center of mass energy per nucleon pair is given by

$$\sqrt{s_{NN}} = \sqrt{s_{pp}}\sqrt{\frac{Z_1 Z_2}{A_1 A_2}} \tag{2.4}$$

where $\sqrt{s_{pp}}$ is the equivalent proton-proton center-of-mass energy under the LHC settings, Z_1 and Z_2 are the number of protons in each nucleus, and A_1 and A_2 are the total number of nucleons in each nucleus.

The first heavy-ion beams were injected into the LHC in 2011, when the corresponding $\sqrt{s_{pp}}$ was 7 TeV. These were Pb + Pb collisions, so the center-of-mass energy per nucleon pair was $(7\ \text{TeV}) \times 82/208 \approx 2.76$ TeV. The next heavy-ion run

was a p+Pb run in 2013, when the corresponding $\sqrt{s_{pp}}$ had increased to 8 TeV. The center-of-mass energy per nucleon pair for this run was (8 TeV) $\times \sqrt{82/208} \approx 5.02$ TeV. The LHC energy capacity was greatly increased for Run 2, which began in 2015. The $\sqrt{s_{pp}}$ was increased to 13 TeV. With these settings, the 2015 Pb + Pb run was scheduled to be carried out with an energy of (13 TeV) $\times$ 82/208 $\approx$ 5.125 TeV, but the energy was decreased to match the 2013 p + Pb run. The year 2016 saw another set of p+Pb runs, one at the traditional 5.02 TeV, and another at the maximum energy of (13 TeV) $\times \sqrt{82/208} \approx 8.16$ TeV. All of the $p + p$ and heavy ion runs that have been carried out at the LHC for which CMS has collected data are summarized in Table 2.1.

The data analyzed in this thesis is from the most recent collisions of lead nuclei, collected in November and December 2018 with $\sqrt{s_{NN}} = 5.02$ TeV to match the previous Pb + Pb and p + Pb runs and the $p + p$ reference data collected in 2015 and 2017. The beam was ramped up to much higher luminosities in 2018, so that nearly 4 times as much Pb + Pb collision data was collected in 2018 compared to 2015.

2.2 The CMS Detector

The Compact Muon Solenoid (CMS) is one of the detectors at the Large Hadron Collider. A photograph of a cross section of the CMS detector is shown in Fig. 2.2. The detector weighs 14,000 metric tons and has a diameter of 15 meters and a length of 28.7 meters. It has the most powerful solenoidal magnet ever made. Though relatively large, the detector is compact compared to an alternative design proposal. From its inception, CMS had four main design goals [54]: 1) a very good muon system providing many possibilities for momentum measurement; 2) the best possible electromagnetic calorimeter consistent with the compactness of the detector; 3) high-quality central tracking to achieve both of the above requirements; and 4) an affordable detector.

A detailed description of the CMS detector can be found in Ref. [56]. A

Table 2.1. Summary of the datasets that have been collected by CMS at the LHC, displayed in order by year, with the collision system ($p+p$, $p+\mathrm{Pb}$, or $\mathrm{Pb}+\mathrm{Pb}$), the corresponding center-of-mass energy per nucleon pair $\sqrt{s_{NN}}$, and the minimum bias integrated luminosity $\mathcal{L}_{\mathrm{int}}$.

Run 1			
Dates	Collision system	$\sqrt{s_{NN}}$	$\mathcal{L}_{\mathrm{int}}$ (Min. Bias)
Mar-Oct 2010	$p+p$	7 TeV	41.5 pb^{-1}
Mar-Oct 2011	$p+p$	7 TeV	5.55 fb^{-1}
Nov-Dec 2011	$\mathrm{Pb}+\mathrm{Pb}$	2.76 TeV	174 $\mu\mathrm{b}^{-1}$
Apr-Dec 2012	$p+p$	8 TeV	21.8 fb^{-1}
Jan-Feb 2013	$p+\mathrm{Pb}$	5.02 TeV	35.5 nb^{-1}
Feb 2013	$p+p$	2.76 TeV	5.51 pb^{-1}
Run 2			
Dates	Collision system	$\sqrt{s_{NN}}$	$\mathcal{L}_{\mathrm{int}}$ (Min. Bias)
Jun-Nov 2015	$p+p$	13 TeV	3.86 fb^{-1}
Nov 2015	$p+p$	5.02 TeV	28.0 pb^{-1}
Nov-Dec 2015	$\mathrm{Pb}+\mathrm{Pb}$	5.02 TeV	464 $\mu\mathrm{b}^{-1}$
May-Oct 2016	$p+p$	13 TeV	38.3 fb^{-1}
Nov-Dec 2016	$p+\mathrm{Pb}$	5.02 TeV	509 $\mu\mathrm{b}^{-1}$
Nov-Dec 2016	$p+\mathrm{Pb}$	8.16 TeV	180 nb^{-1}
Jun-Nov 2017	$p+p$	13 TeV	45.0 fb^{-1}
Nov 2017	$p+p$	5.02 TeV	323 pb^{-1}
May-Oct 2018	$p+p$	13 TeV	63.7 fb^{-1}
Nov-Dec 2018	$\mathrm{Pb}+\mathrm{Pb}$	5.02 TeV	1.79 nb^{-1}

Figure 2.2. A photograph of a cross section of the CMS detector [53].

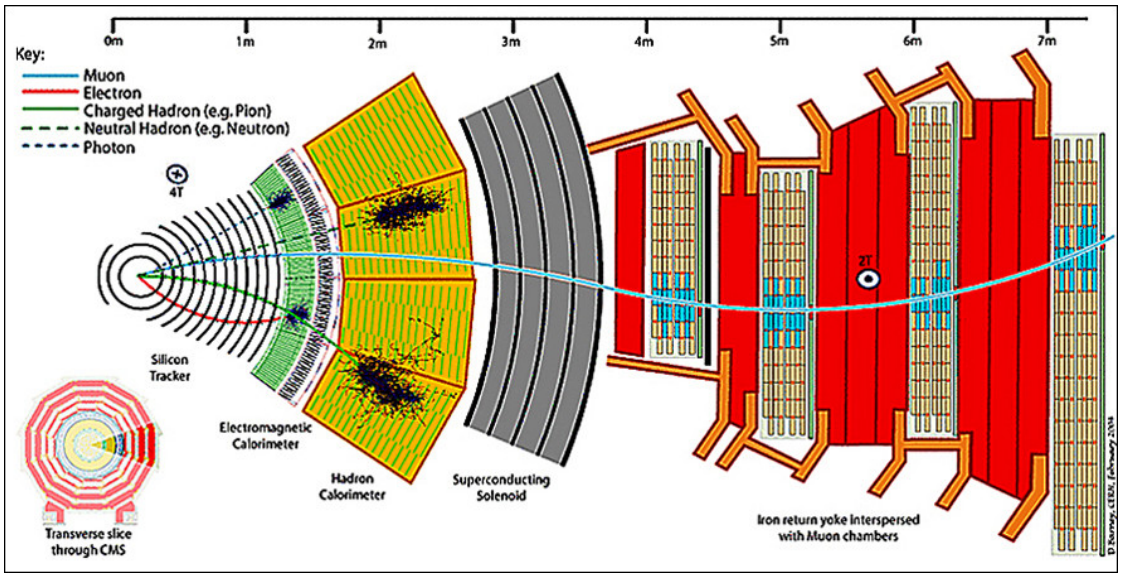

Figure 2.3. A schematic slice of the CMS detector [55].

schematic slice of the CMS detector is shown in Fig. 2.3. The beamline is directed into and out of the page in the center of the circle from which particles are seen to be emanating. The first detectors with which the particles interact are the silicon trackers, displayed as black circles near the collision point. The next detector is the electromagnetic calorimeter, shown in green, which detects and stops electrons and photons. Next, the hadron calorimeter, shown in yellow, detects and stops hadrons. Just outside the hadron calorimeter is the superconducting solenoid which provides an internal magnetic field of 3.8 T. The magnetic field bends the paths of the charged particles to allow for the determination of momentum. Outside the solenoid, muon detectors are interspersed with an iron yoke which serves to direct the magnetic field and to concentrate the field lines in the region outside the magnet.

The design of CMS makes it very adept at detecting muons. The calorimeters block all other particles so that only the muons escape and are easily identified in the other detectors, while the tracker and the solenoid combined provide excellent momentum measurements. A more detailed description of each detector component follows.

2.2.1 Coordinate System

The motion of particles in the detector requires knowledge of both position and momentum. Due to the cylindrical symmetry of the detector, the position of a particle can be conveniently described in cylindrical coordinates r, ϕ, and z, where r represents the radial distance from the beamline, ϕ represents the azimuthal angle, and z is distance in the direction of the beamline from the center of the detector.

The momentum can similarly be described with cylindrical coordinates, but certain modifications are applied for experimental convenience. The radial component of momentum in the transverse plane (perpendicular to the beam direction) is referred to as the transverse momentum p_T. It can be measured by analyzing the curvature of charged particle tracks in the magnetic field projected into the

transverse plane. In the longitudinal direction, a quantity called rapidity is often used. Rapidity is a representation of velocity that is additive under boosts,

$$y = \frac{1}{2} \ln \frac{E + p_z c}{E - p_z c}, \tag{2.5}$$

where E is the energy of the particle and p_z is the component of momentum along the z-axis. A rapidity of $y = 0$ corresponds to a velocity that is entirely perpendicular to the z-axis, while other values form an acute angle with respect to the z-axis.

One further related quantity depends on the angle θ between the particle trajectory and a line perpendicular to the z-axis. The pseudorapidity η is defined in a similar manner to the rapidity, except the energy E is replaced by the scalar momentum $|\mathbf{p}|$, as follows:

$$\eta = \frac{1}{2} \ln \frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} = \ln \left[\tan \left(\frac{\theta}{2} \right) \right]. \tag{2.6}$$

Since this quantity is directly related to the angle θ, regions of the CMS detector can be described in terms of η rather than θ. Another important advantage of pseudorapidity over rapidity is that it depends only on momentum and is independent of the mass of the particle, and thus it does not require particle identification. In the limit of $|\mathbf{p}| \gg mc$, the rapidity and pseudorapidity become identical.

2.2.2 Silicon Trackers

Reliable measurements of momentum are crucial for reconstructing a collision event. The detectors that measure momentum must be closer to the collision point than all the other detectors because the particles may lose momentum while passing through the many layers of material. The momentum is determined by measuring the curvature of charged particle paths in the presence of the magnetic field generated by the solenoid. The proper reconstruction of a curved track requires knowledge of the location of the particle at three points along the particle trajectory. When there are many particles, and particle decays that form secondary vertices, more detection points are needed in order to distinguish between

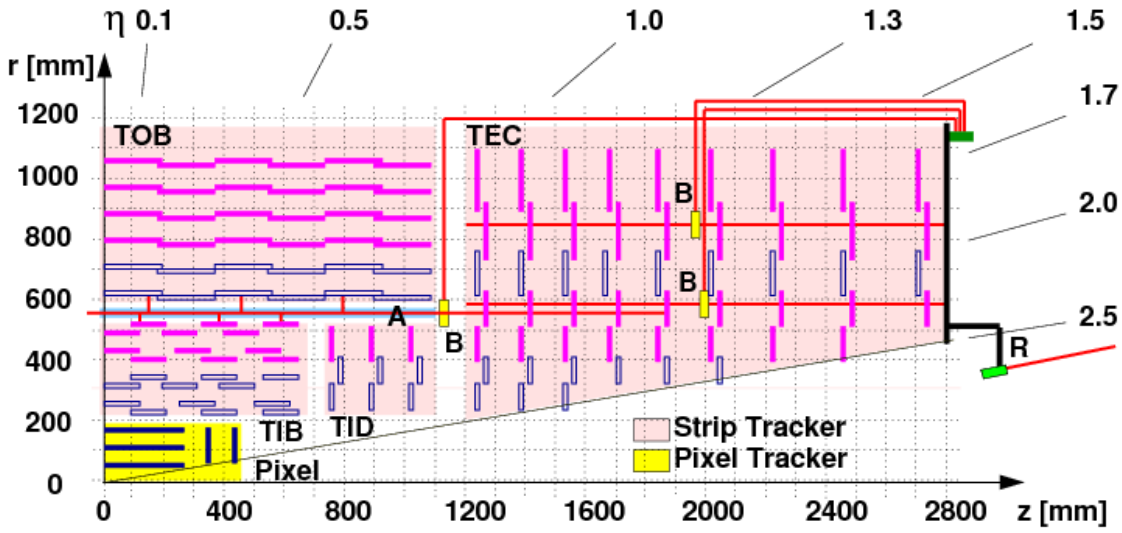

Figure 2.4. A schematic representation of one quadrant of the CMS tracking system, where the beamline extends left and right from the origin. The y-axis is the radial distance from the collision point, and the x-axis is the distance along the beamline. Pseudorapidity values are indicated at the top and right. Figure from Ref. [57].

the many different trajectories. It follows that the tracker must consist of several layers of detector material, and each layer must be thin enough to avoid causing significant change to the particle momentum, while at the same time be sensitive enough to accurately locate the particle. The tracker must also be able to withstand larger intensities of radiation than any other detector in CMS.

The material of choice was silicon. The silicon detectors work via ionization as the charged particles pass through. Electrons are excited and escape from the silicon atoms, creating free electron-hole pairs. An applied electric field causes the free electrons to move, generating a small electric current which is then amplified by an electronic silicon chip and carried off as an electrical signal.

A schematic depiction of the lateral view of one quadrant of the tracker is shown in Fig. 2.4. The y-axis is the radial distance from the collision point, and the x-axis is the distance along the beamline. The first 3 layers of the tracker consist of very fine silicon pixel detectors. There are 65 million pixels each with dimensions 100 μm by 150 μm. The next 4 layers are silicon strip detectors arranged in shells that form the Tracker Inner Barrel (TIB) and 2 sets of detectors in the inner end

caps called the Tracker Inner Disc (TID) made of silicon discs. Six more layers of silicon strip detectors in the outer barrel region make up the Tracker Outer Barrel (TOB), and two more endcap detectors make up the Tracker endcap (TEC). The tracking system is kept at $-20°$C by cooling tubes to minimize disorder and prevent damage to the detectors.

2.2.3 Calorimeters

The purpose of a calorimeter is to measure energy. In CMS, two different types of calorimeters, electromagnetic and hadronic, were developed to measure the energies of two different families of particles. In both cases, the calorimeters must completely stop the incoming particles in order to obtain a precise measurement of the total energy. Due to the compactness of CMS, the materials of choice had to be capable of stopping particles over very short distances. This limited the candidate materials to those with very short interaction lengths.

The electromagnetic calorimeter (ECAL) was designed to measure the energy of electrons and photons by inducing Bremsstrahlung radiation and pair production in a crystal scintillator. A scintillator is a type of material that absorbs energy from incoming ionizing particles and re-emits the energy as light. The resulting light signals could then travel to photodetectors at the end of the crystal and be converted into electrical signals. A high Z material was needed in order to increase the cross section for radiation and pair production. The crystals also needed to have a very fast scintillation decay time of the same order of magnitude as the bunch crossing time at the LHC. This would ensure that the signal from each event would be collected before the next event occurred. The main candidate materials for the electromagnetic calorimeter were cerium fluoride (CeF3), lead-tungstate crystals (PbWO4), and hafnium-fluoride glasses. Lead-tungstate was determined to be the optimal choice due its relatively low radiation length (0.9 cm), the availability of the raw materials to make it, and the fact that significant production capacity already existed [54]. Avalanche Photodiodes (APDs) detect the scintillating photons in the barrel region, magnify the signal with a gain of ~ 50, and transport the signal

away through fiber-optic cables. Vacuum phototriodes (VPTs), which have less gain (~ 10) but greater resistance to radiation, are used in the endcaps because the radiation in that region is too high for silicon photodiodes.

The hadronic calorimeter (HCAL) fills up the remaining space between the ECAL and the solenoid. The HCAL relies on nuclear interactions to measure the energy of hadrons. It consists of alternating layers of dense absorber and plastic scintillator. The absorber had to be a dense, rigid, high-Z material with a large nuclear interaction cross section that could last as long as 15 years in the CMS detector environment. The material chosen was brass, which is an alloy of copper and zinc. Most of the brass used in the HCAL was recycled from old Russian artillery shells that were melted and reshaped. The rest was made from copper from the United States. The interaction of hadrons with the nuclei in the brass absorber creates a hadronic shower of particles which emit flashes of light in the scintillator, where the light is detected and converted to an electrical signal similar to the process in the ECAL.

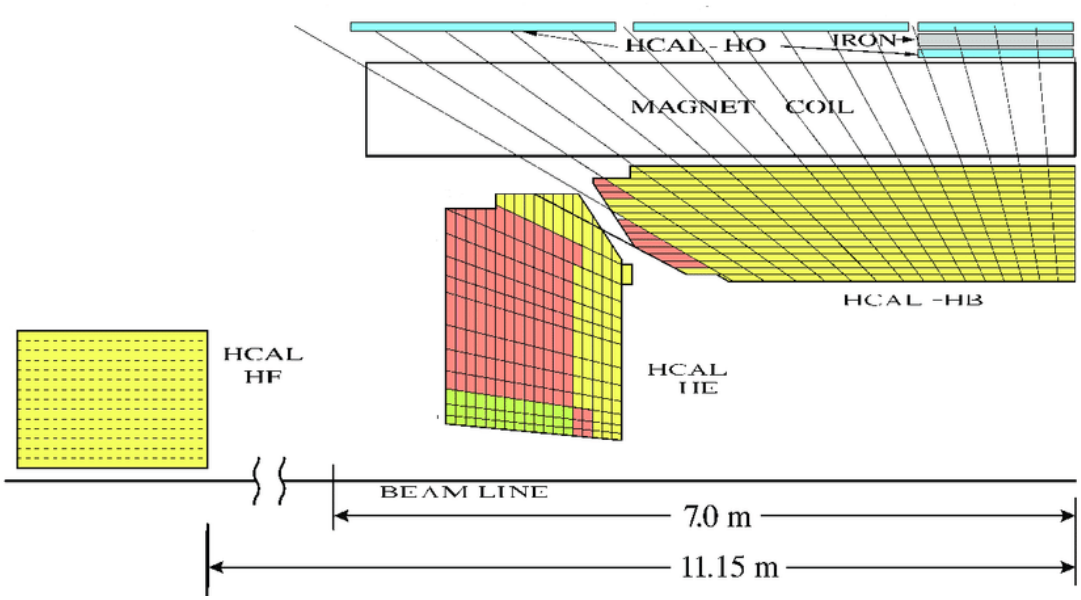

Figure 2.5. A schematic representation of one quadrant of the hadron calorimeter, where the beamline extends left and right from the origin. The y-axis is the radial distance from the collision point, and the x-axis is the distance along the beamline [58].

A schematic diagram of the HCAL is shown in Fig. 2.5 [58]. The HCAL was constructed as 36 large wedges, each weighing 26 metric tons, which were inserted into the barrel region forming the Barrel HCAL (HB). A few additional layers of HCAL were placed outside the solenoid, forming the Outer HCAL (HO) to ensure no energy escaped the HB undetected. Another 36 wedges form the HCAL in the endcap (HE). The HCAL forward calorimeters (HF) extend the pseudorapidity range of the HCAL to 5.0.

2.2.4 The Solenoid

The solenoid, for which CMS is named, is the central feature of the detector. It was designed to create the strongest magnetic field possible given the size and cost restrictions. The solenoid was built offsite and transported to CMS in the city of Cessy, France, by road, which limited the diameter of the solenoid to no more than 7 m. The length of the solenoid is 12.5 m and the internal diameter is 6 m, providing a snug fit for the tracker and calorimeters.

A strong magnetic field is necessary to obtain precise measurements of the charged particle momentum. In the presence of the magnetic field within the solenoid, the trajectory of a charged particle bends with a radius of curvature given by

$$R = \frac{p_T}{qB} \tag{2.7}$$

where p_T is the transverse momentum of the particle, q is the electric charge, and B is the magnetic field strength. Thus, with prior knowledge of the magnetic field strength, a good tracking system that accurately measures the curvature of the path, and by properly identifying the particle to obtain the charge, the momentum can then be determined. A larger magnetic field will result in a smaller radius of curvature (*i.e.* more bending), allowing for more precise measurements of p_T.

The solenoid consists of many coils of Rutherford cables, which are constructed from many flexible superconducting fibers wrapped together in a helix and flattened into a rectangular shape. The advantage of these cables is that they pack together very well, provide good control of dimensions, have good windability, and

the superconducting fibers are transposed. The cables in the solenoid consist of niobium titanium fibers co-extruded with aluminum and reinforced with aluminum alloy. The coils are cooled to a temperature of -268.5 °C to become superconducting, allowing electricity (18,000 amps) to flow without resistance and generating an interior magnetic field of 3.8 T. It is the largest superconducting magnet and the most powerful solenoidal magnet ever built.

Several layers of iron yoke extend out to a diameter of 14 m. The iron yoke serves to direct the magnetic field and to concentrate and contain the field lines in the region outside the magnet. The iron yoke, together with the solenoid, accounts for most of the mass ($\sim 14,000$ metric tons) and the structure of the CMS detector.

2.2.5 Muon Detectors

Outside the solenoid, interspersed between the layers of the iron yoke, are the muon detectors. Muons are the only particles that reach these detectors because all others are stopped in the calorimeters. Muons are very similar to electrons, but 200 times heavier, their greater mass prevents them from interacting with the ECAL. Electrons are detected in the ECAL via Bremsstrahlung radiation, where the power radiated is proportional to m^{-4}. This large dependence on mass is why electrons are stopped in the ECAL but muons are not. Charged hadrons such as protons also bypass the ECAL for the same reason, but they are stopped by nuclear interactions in the HCAL which the muons do not experience.

There are four layers of muon detectors that allow a muon track to be rectonstructed by fitting a curve to the four detection points, as seen in Fig. 2.3. The curve is bent by the magnetic field, which has a magnitude of about 2 T in this region, and provides information on the muon momentum. The muon track in the muon detectors is matched to a track in the inner tracker to complete the reconstruction of its trajectory.

The muon detectors consist of 250 drift tubes (DTs) in the barrel region and 540 cathode strip chambers (CSCs) in the endcap region, with 610 resistive plate chambers (RPCs) interspersed between each successive layer of DT or CSC.

A DT consists of a tube of gas (Ar and CO_2) with a stretched anode wire running through it. The wire is positively charged to create a magnetic field in the tube. When a muon passes through the tube, it ionizes the gas and the freed electrons then drift toward the wire because of the electric field. Proper timing allows this detector to determine how far away the electron was from the wire originally, and exactly where it hit the wire, providing information on two coordinates. Each DT chamber contains 12 layers of DTs, in three groups. The middle group is oriented to measure the coordinate along the beam axis and the other two groups are oriented to measure the perpendicular coordinates.

A CSC is made by crossing cathode strips with positively-charged anode wires within a gas volume. Similar to the DT, muons ionize the gas while passing through and the free electrons then move toward the anode wire. The ionized atoms, now positively charged, move toward the copper cathode strip and produce a signal. Since the cathode strips and the anode wires are perpendicular, the two signals pinpoint a location in two dimensions. The CSC modules are arranged like discs in the endcap region.

The RPCs are parallel-plate detectors, with a positive plate and a negative plate placed side-by-side with a thin gas volume in between. The plates are made of a very high-resistivity material called phenolic resin (bakelite) coated in conductive graphite. Passing muons ionize the gas particles in the same manner as in the DTs and CSCs. The electrons then move toward the positive plate while the positively-charged ions move toward the negative plate, inducing a signal in an array of external metallic strips. The RPCs are designed to give a very fast ($\sim$3 ns) measurement of muon momentum to serve as a basis for the muon triggers.

2.3 Triggering and Data Collection

With bunches crossing through the interaction point in CMS every 25 ns, the collision rate can be upwards of 40 MHz. It is not possible for CMS (or the world) to record and store every collision that occurs within CMS. Physicists are

insterested in studying particular events that meet certain pre-defined conditions, so even if it were possible to record every event, it would not be practical. The CMS detector has built-in trigger systems that fire whenever an event of interest occurs. These triggers decide which events will stored for analysis and which will be rejected.

Events of interest are selected using a two-tiered trigger system [59]. Custom hardware processors compose the first level (L1). These L1 processors use information from the calorimeters and muon detectors to select events at a rate of around 100 kHz for $p+p$ and $p+\mathrm{Pb}$ collisions and around 30 kHz for $\mathrm{Pb}+\mathrm{Pb}$ collisions within a fixed time interval of about 4 μs. The second level, known as the high-level trigger (HLT), consists of a farm of processors running a version of the full event reconstruction software optimized for fast processing. The HLT reduces the event rate to around 1 kHz for $p+p$, 20 kHz for $p+\mathrm{Pb}$, and 7 kHz for $\mathrm{Pb}+\mathrm{Pb}$ collisions before data storage.

The online muon reconstruction algorithms are identical for the various collision systems. However, heavy ion runs typically have a much lower luminosity than $p+p$, and hence the dataset occupies a much smaller space on tape. These runs usually happen at the end of the year, before the annual LHC shutdown, so offline processing constraints are relaxed. For these reasons, heavy ion runs use looser triggers than those used for $p+p$ runs.

An example of a heavy-ion event in CMS is shown in Fig. 2.6. Particle tracks reconstructed in the tracker are displayed as yellow lines. Energy deposits in the calorimeters are shown as blue and green towers. The muon tracks reconstructed in the muon chambers are shown as red curves. There tend to be many more tracks in $\mathrm{Pb}+\mathrm{Pb}$ collisions than in $p+p$ collisions, so the detectors experience very high occupancies that require excellent spatial resolution and track separation. However, while there is a large amount of hadronic activity in the inner detector, there are very few muons produced in a given event and their signal in the muon chambers is very clear even in this 2D projection of the event. The two

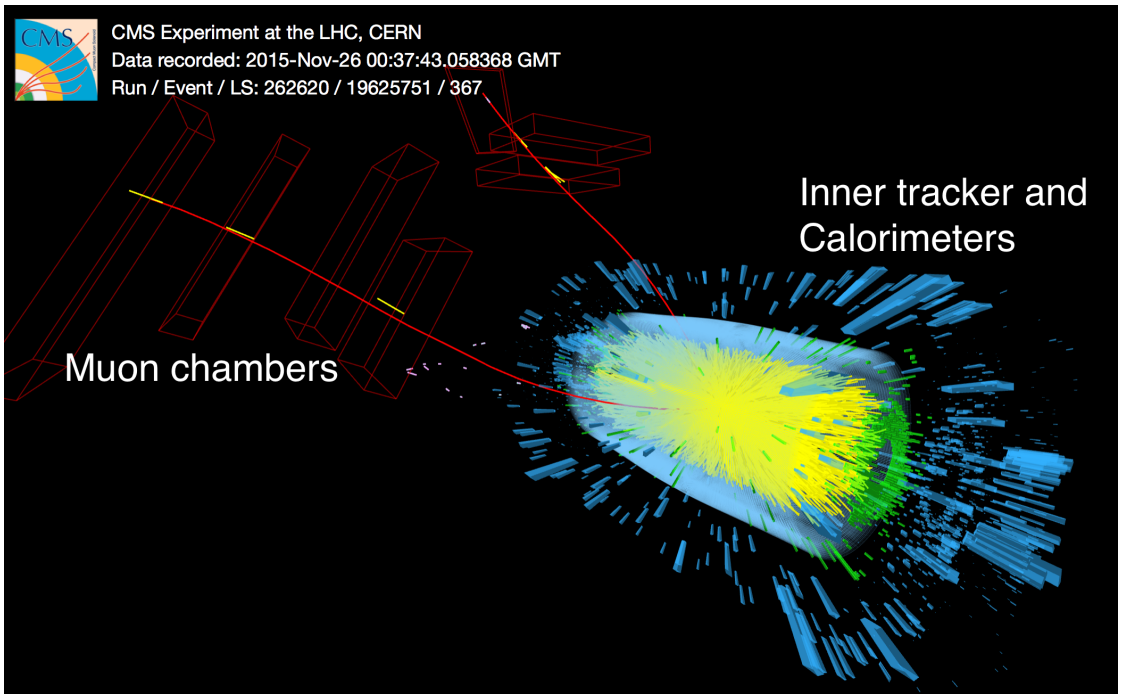

Figure 2.6. A dimuon event from a Pb + Pb collision in the CMS detector [60].

muons shown here are actually a candidate Upsilon. Their paths curve in opposite directions, which means they have opposite charge, and their combined invariant mass matches that of the $\Upsilon(1S)$. The superb muon-detection capabilities of CMS make it an excellent apparatus to study dimuon states, such as the Upsilon family, via $\Upsilon \to \mu^+\mu^-$.

Chapter 3

Detector Performance

3.1 Efficiency

To obtain an accurate yield of *e.g.* particle multiplicities in a given event, a determination of how well the detector can identify and reconstruct particle information is essential for understanding the reliability of the data and the measurement precision. The performance of the detector is estimated by the efficiency, defined as the ratio of the number of objects that were reconstructed or identified by the detector over the true number of objects that were present in the event. These estimates are typically made using Monte Carlo (MC) simulations in which the objects present in the event can be known a priori. However, estimates of efficiencies from MC assume a perfect description of a physical detector, which is typically not possible, and thus this method introduces large systematic uncertainties. To avoid this problem, data-driven approaches to estimating efficiencies have been developed [61, 62].

Performance studies of the CMS detector at the LHC at CERN have been carried out using the tag-and-probe method in the $p + p$ environment [61] and the heavy-ion environment since 2010 [62]. The tag-and-probe method involves selecting two of the objects of interest from real data and matching them to a resonance that decays to the objects of interest. One of the objects is labeled the "tag," and the other is the "probe." The tag is an object that passes a very strict

set of selection criteria. In this case, it is identified as a muon with a very high level of confidence. The probes have much looser selection criteria and represent the base sample from which the efficiencies of various selection criteria can be estimated. The tag-and-probe muon pairs are used to reconstruct the dimuon resonance from which they decayed. The resonance is then fit to extract the yield with and without certain selection criteria applied to the probes. The resonance yield reconstructed from "passing" probes divided by the total yield gives the efficiency of the selection criteria applied to the probe muons. Because the tag-and-probe method uses real data, and the true muons are not known, it cannot be used to measure absolute efficiencies where the denominator is "all true muons." Instead, it is used to estimate "relative" efficiencies where the denominator is "all objects passing certain criteria" and the numerator is "all of those objects passing tighter criteria."

The standard "candles" of the tag-and-probe analyses are the J/ψ and the Z. These two resonances lie at opposite extremes of the dimuon mass distribution (see Fig. 4.1), with world-average masses of 3.0969 GeV/c^2 and 91.1876 GeV/c^2, respectively [24]. The daughter muons born from the decays of the J/ψ and the Z cover the full p_T range of muons in the CMS detector, from $\sim$0.8 GeV/c up to $\sim$200 GeV/c. The efficiencies presented in Sec. 3.3 are estimated from these two resonances.

3.2 Muon Detection in the Heavy-ion Environment

The various collision systems of the LHC provide a wide range of environments for which the CMS detector has been prepared. The different collision systems all occur in the same detector and all muons are reconstructed using essentially the same algorithm. The main difference that is expected to affect the performance of the detector is the number of particles produced in the different systems, which is quantified employing the charged-particle multiplicity of a given event. The vari-

able used to characterize the total charged-particle multiplicity, which is measured in the pseudorapidity range $|\eta| < 2.4$, is labeled N_{tracks}.

The particle multiplicities in the Pb + Pb environment vary greatly due to the large number of particles involved and the different types of collisions that can occur (*e.g.* central versus peripheral collisions). In the most central Pb + Pb events, the multiplicities are typically thousands of times larger than the multiplicities observed in $p + p$ collisions. According to Table 1.1, the Pb + Pb events in 0-10% centrality have an average of ∼1600 binary collisions among ∼360 participating nucleons. This results in extremely high detector occupancies (the number of particles entering the detector) which may have an effect on the muon reconstruction and identification efficiencies. For this reason, the performance of the CMS detector was studied separately in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, p + Pb at $\sqrt{s_{NN}} = 8.16$ TeV, and $p + p$ at $\sqrt{s} = 5.02$ TeV (see Table 2.1) [60].

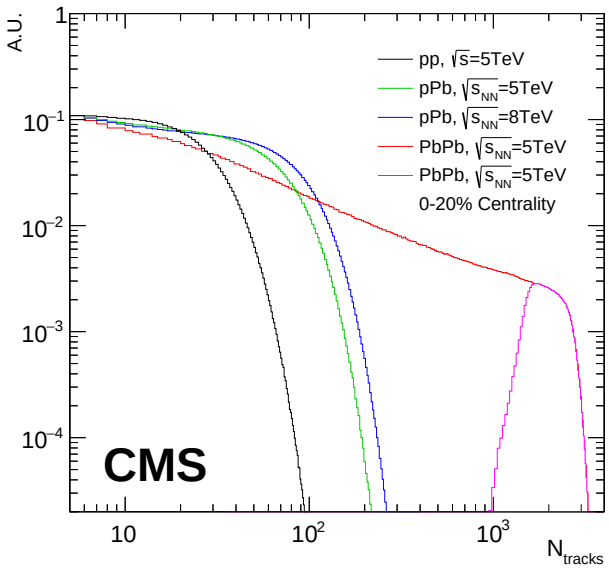

Figure 3.1. The distribution of N_{tracks} in Pb + Pb at $\sqrt{s_{\text{NN}}} = 5.02$ TeV (red), p + Pb at $\sqrt{s_{\text{NN}}} = 8.16$ TeV (blue), p + Pb at $\sqrt{s_{\text{NN}}} = 5.02$ TeV (green), and $p + p$ at $\sqrt{s} = 5.02$ TeV (black). The distribution of N_{tracks} in the most central collisions (0–20% centrality) is shown in magenta. The normalization of the curves is such that they align in the lowest N_{tracks} bin, hence the y-axis is in "arbitrary units" (A.U.) [60].

The distribution of N_{tracks} in $p+p$, $p+$Pb, and Pb+Pb events at $\sqrt{s_{\text{NN}}} = 5.02$

TeV is shown in Fig. 3.1 where p + Pb at $\sqrt{s_{NN}} = 8.16$ TeV is also included for comparison [60]. In the case of Pb + Pb, events with a large number of tracks correspond to the most central collisions (large nucleus-nucleus overlap and small impact parameter), while a small number of tracks indicates a peripheral collision. To illustrate the highest multiplicities achievable in Pb + Pb, the distribution of N_{tracks} in the most central Pb + Pb collisions (0–20% centrality) is displayed in magenta. The mean of the 0–20% centrality distribution is $\langle N_{\text{tracks}} \rangle_{0-20\%} = 2032$. Most Pb+Pb collisions are peripheral, which is why the distribution is heavily populated at lower values of N_{tracks} and less populated at the higher end values. The distribution quickly drops off when the number of participating nucleons reaches its limit.

The same effects are visible in the p + Pb system, but to a lesser extent. In p+Pb, the number of participating nucleons does not change as much with impact parameter as it does in Pb + Pb, so the N_{tracks} distribution decreases relatively slowly before dropping off. The total energy present in a p + Pb event is only about 1/3 of the total energy present in a Pb + Pb event at the same $\sqrt{s_{NN}}$. This is why the maximum N_{tracks} in p + Pb is only about 1/3 of the maximum value in Pb + Pb. The N_{tracks} distributions in two p + Pb systems with different energies demonstrate that a higher center-of-mass energy per nucleon results in more particle creation. In this case, an increase of 63% in the center-of-mass energy per nucleon pair resulted in an increase of 24% in the average value of N_{tracks}.

The $p + p$ collisions involve only two participating nucleons. However, at the luminosities achievable at the LHC, there can be more than one $p+p$ collision per bunch crossing. This is referred to as pileup, and it affects the shape and extent of the N_{tracks} distribution. Even with pileup, however, the $p+p$ collisions generate far fewer tracks than Pb + Pb collisions and even p + Pb collisions. The average number of tracks reconstructed in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV is shown in bins of centrality in Table 3.1. The average numbers of N_{tracks} present in p+Pb

Table 3.1. The average number of reconstructed tracks in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV in bins of centrality. The corresponding averages in $p + \text{Pb}$ and $p + p$ collisions are also given. The tabulated data are from Ref. [60].

PbPb Centrality range	Average N_{tracks}
0–10%	2345
10–20%	1721
20–30%	1224
30–40%	819
40–50%	503
50–60%	279
60–100%	33
pPb, $\sqrt{s_{NN}} = 5.02$ TeV	39
pPb, $\sqrt{s_{NN}} = 8.16$ TeV	48
pp, $\sqrt{s} = 5.02$ TeV	6
pp with pileup	14

collisions at two different energies and in $p + p$ collisions with and without pileup are also shown.

As seen in the table, the most central Pb+Pb collisions (0–10% centrality) have $\langle N_{\text{tracks}} \rangle_{0-10\%} = 2345$, which makes each event roughly equivalent to about 2345/6 $\approx$ 400 $p + p$ collisions occurring simultaneously. This is much less than the value of $\langle N_{\text{coll}} \rangle = 1586$ estimated in Table 1.1 because the nucleon-nucleon collisions in a Pb + Pb event re-use the same nucleons again and again. Energy is lost in each successive nucleon-nucleon collision so that the average center-of-mass energy per nucleon pair is much less than 5.02 TeV and generates fewer tracks. However, the difference in N_{tracks} between a central Pb + Pb event and a typical $p + p$ event at the same $\sqrt{s_{NN}}$ is still very large. This substantial increase in particle multiplicity warrants an investigation into the subsequent effects on the detector performance.

To convey an idea of these high-occupancy events, a typical dimuon event in the heavy-ion environment is displayed in Fig. 2.6. The high occupancy in the

inner detector is expected to impact the performance of the pixel detector and silicon strip tracker and reduce the muon tracking efficiency. The muon detectors, on the other hand, are expected to perform well even in the heavy-ion environment because muons are relatively rare and only the muons manage to traverse the entire inner detector and enter the muon chambers outside of the solenoidal magnet. Therefore, we do not expect a large change in muon identification efficiency, nor do we expect a large change in any low-level triggers which rely solely on information from the muon chambers. The efficiency of the high-level triggers (HLT), however, is expected to decrease since these triggers use tracking information from the inner detector. Many of the physics triggers (*e.g.* quarkonia, Z, etc.) use the HLT tracking information, and therefore are expected to be affected by the high occupancies in heavy-ion collisions.

3.3 Converting Centrality to Number of Tracks

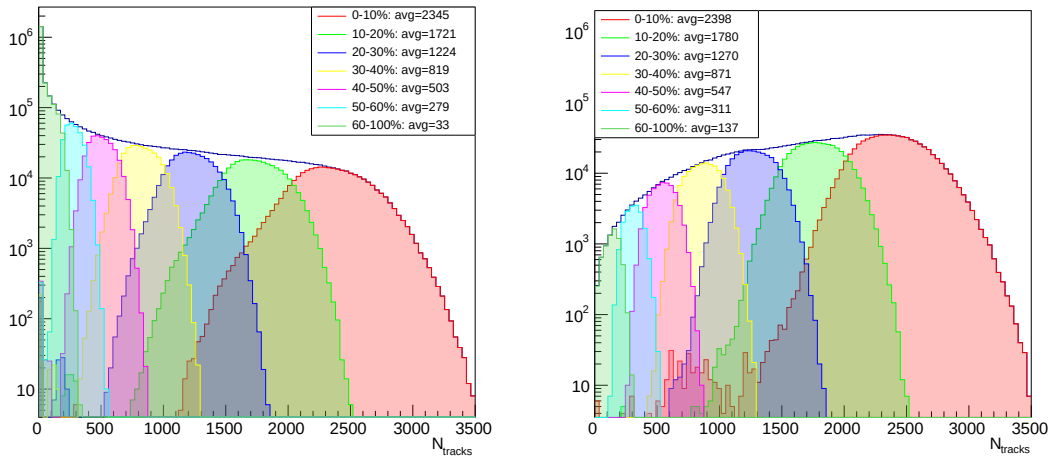

Figure 3.2. Distributions of N_{tracks} in the Pb + Pb dataset in bins of centrality in a minimum bias dataset (left) and a triggered dataset (right) [60].

It is common in analyses of Pb + Pb collisions to study the effects of event activity by binning in centrality, hence efficiencies in Pb + Pb were estimated in bins of centrality. Unfortunately, centrality is not well defined in $p+p$ and $p+$Pb systems and this limits our ability to make comparisons among the three systems.

Other variables such as the number of tracks (N_{tracks}) and the transverse energy deposited in the forward hadron calorimeters (E_T^{HF}) are available to describe the event activity in all three systems. In order to improve the quality and readability of the comparison between efficiencies in Pb + Pb and p + Pb as a function of event activity, it was necessary to convert the centrality bins of the Pb + Pb efficiencies to N_{tracks}.

The conversion was done by plotting the distribution of N_{tracks} in the Pb + Pb dataset in the various centrality bins and taking the mean of each distribution as the value of N_{tracks} for that data point of efficiency. Examples of this are shown in Fig. 3.2 [60]. The figure on the left was generated from a 'minimum bias' dataset, which means no selections were applied to the collision events other than what is necessary to ensure that a hadronic collision actually occurred. The minimum bias trigger requires at least two energy deposits in the calorimeters in the range $|\eta| < 4.7$ and a signal above noise level in the very forward pseudorapidity region $3.2 < |\eta| < 4.7$ [63, 64]. The figure on the right is from a 'triggered' dataset which required events to have at least two muons passing various selection criteria common to dimuon analyses. The extra requirements tend to bias the average N_{tracks} toward higher values. This is because the high-mass dimuons are more likely to occur in central events with large N_{coll}, and hence large N_{tracks}.

The figures showing efficiency as a function of centrality were translated to N_{tracks} by plotting the efficiency in each centrality bin at the mean value of N_{tracks} in that bin. The resulting efficiency as a function of N_{tracks} in Pb + Pb extracted from the Z resonance were merged with the corresponding results from p + Pb and $p + p$ and are shown in Figures 3.3, 3.4, and 3.5 [60].

Comparisons of efficiencies across the various collision systems at the J/ψ resonance was complicated by the different approaches used in the tag-and-probe analyses. At the level of reconstruction, the Pb + Pb analysis probed the inner tracking efficiency of "stand-alone muons" (muons with tracks in the muon detectors), while the p + Pb analysis probed the efficiency of "tracker muons" (muons

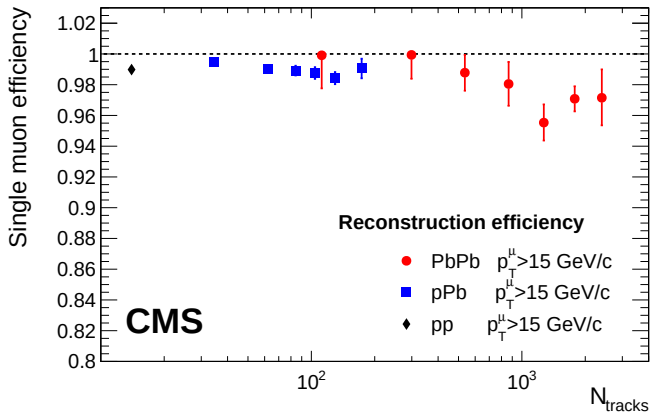

Figure 3.3. Muon reconstruction efficiency as a function of N_{tracks} estimated from $p+p$ at $\sqrt{s} = 5.02$ TeV (black diamond), $p+\text{Pb}$ at $\sqrt{s_{NN}} = 8.16$ TeV (blue squares), and $\text{Pb}+\text{Pb}$ at $\sqrt{s_{NN}} = 5.02$ TeV (red circles) collision data at the Z resonance [60].

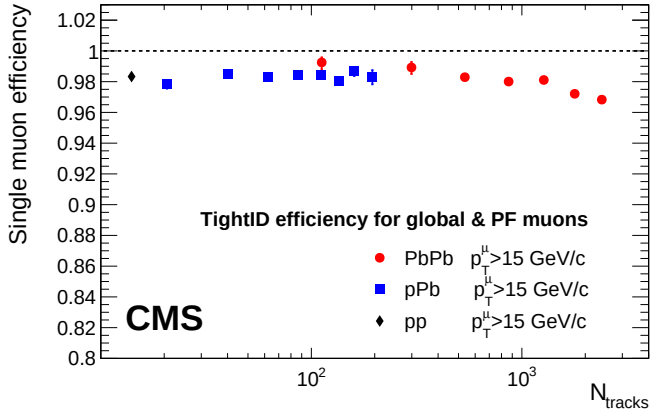

Figure 3.4. Muon identification efficiency as a function of N_{tracks} estimated from $p+p$ at $\sqrt{s} = 5.02$ TeV (black diamond), $p+\text{Pb}$ at $\sqrt{s_{NN}} = 8.16$ TeV (blue squares), and $\text{Pb}+\text{Pb}$ at $\sqrt{s_{NN}} = 5.02$ TeV (red circles) collision data at the Z resonance [60].

with tracks in the inner tracker matched with tracks in the muon chambers), and the $p+p$ analysis probed the efficiency of "global muons" (muon tracks in the muon chambers matched to tracks in the inner tracker). The different methods made it difficult to draw meaningful comparisons between systems. Therefore only the $\text{Pb}+\text{Pb}$ tracking efficiency at the J/ψ resonance is shown in Fig. 3.6.

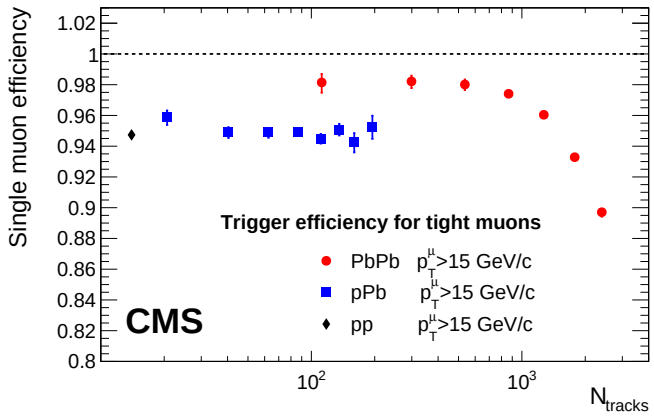

Figure 3.5. Muon trigger efficiency as a function of N_{tracks} estimated from $p+p$ at $\sqrt{s} = 5.02$ TeV (black diamond), $p+\text{Pb}$ at $\sqrt{s_{NN}} = 8.16$ TeV (blue squares), and $\text{Pb}+\text{Pb}$ at $\sqrt{s_{NN}} = 5.02$ TeV (red circles) collision data at the Z resonance [60].

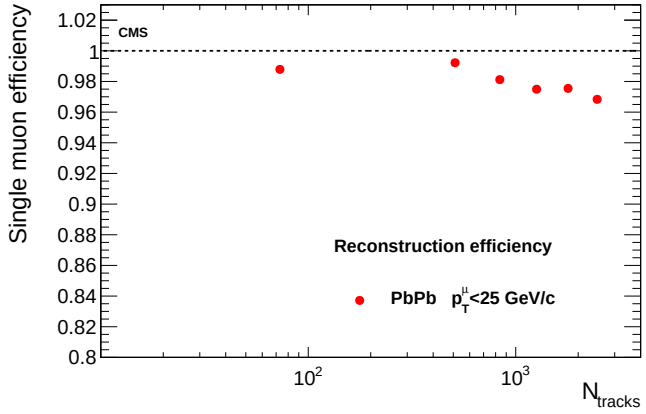

Figure 3.6. Muon tracking efficiency as a function of N_{tracks} estimated from $\text{Pb}+\text{Pb}$ at $\sqrt{s_{NN}} = 5.02$ TeV (red circles) collision data at the J/ψ resonance. Other collision systems are not included due to a lack of commensurability of the probed efficiencies among the different analyses.

3.4 Mass Scale and Resolution

In addition to the muon reconstruction efficiency, it is possible that the high occupancy environment in central $\text{Pb}+\text{Pb}$ collisions affects the precision and accuracy of the reconstructed muon momentum. Since the mass of the J/ψ and Z are well known, and the reconstructed mass is derived from the muon momentum, the dimuon mass resolution provides a suitable window into the precision of the momentum measurement.

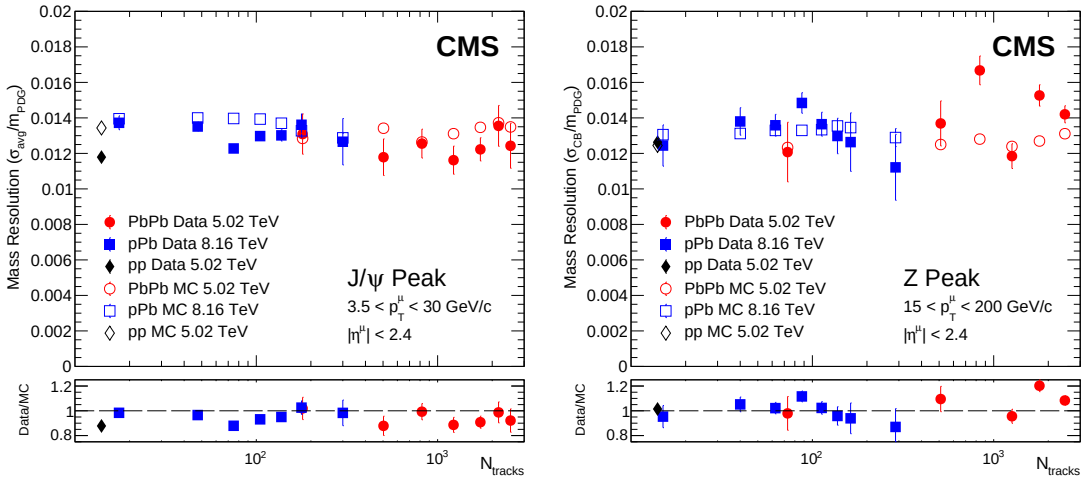

Figure 3.7. Mass resolution as a function of N_{tracks}, estimated as the width of the signal peak (σ_{avg} for J/ψ and σ_{CB} for Z) in fits to $p+p$ (black diamonds), $p+\text{Pb}$ (blue squares) and $\text{Pb}+\text{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances using the tag-and-probe method [60].

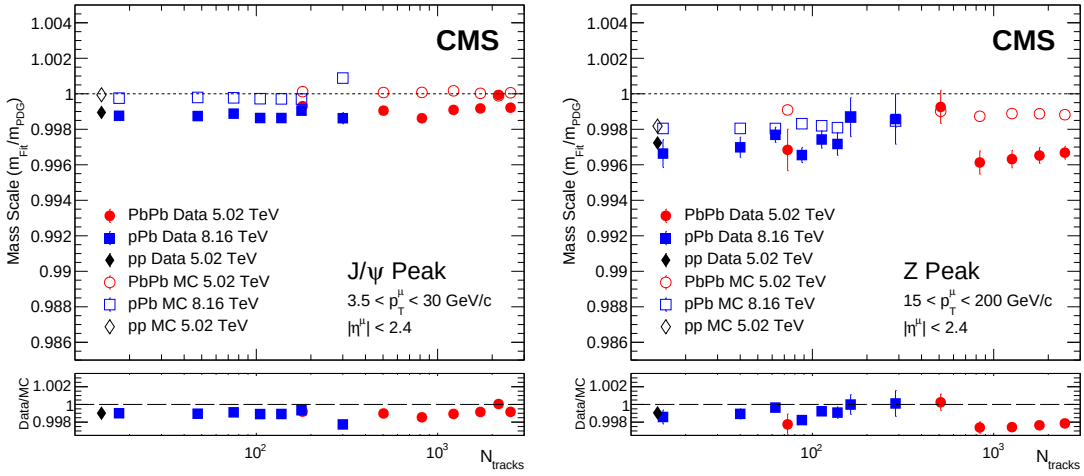

Figure 3.8. Mass scaling as a function of N_{tracks} estimated from $p+p$ (black diamond), $p+\text{Pb}$ (blue squares), and $\text{Pb}+\text{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances [60].

The mass resolution is obtained from the Gaussian width of the dimuon resonance after all selections are applied, taking into account the natural width of the resonance, divided by the world-average mass m_{PDG}, as reported by the Particle Data Group (PDG) [24]. The quantity shown therefore represents the unitless

relative deviation in the reconstructed mass compared to m_{PDG}.

In the case of the Z, the signal peak is described by a Breit-Wigner (BW) function convoluted with a Crystal Ball (CB) function. The width of the BW is fixed to the natural width of the Z ($\sim$2.5 GeV/c^2), so the width of the CB function, σ_{CB}, which is a free parameter in the fit, contains all the information on the resolution of the detector. In the case of the J/ψ, the natural width is negligible, and the signal is described by a CB plus a Gaussian with different widths to account for varying resolution in different pseudorapidity regions of the detector. In this case, the width used to estimate the mass resolution is a weighted average of the CB and Gaussian widths, weighted by the amount of signal contained in each function. The estimated mass resolution of the J/ψ and the Z are shown as a function of N_{tracks} in Fig. 3.7. The resolution is consistent with a constant value, which indicates that it is not affected by the level of tracker occupancy.

The mass scaling measures how accurately the detector reconstructs the known mass of the resonance. It is obtained by dividing the reconstructed mass m_{fit} by the PDG mass m_{PDG}. It is worth noting that mass measurements are rare using heavy-ion collisions because more more data and cleaner signals are available in *e.g.* $p+p$ collisions. It is also not expected that a small bias on the momentum scale will have a significant impact on the mass measurement because the natural widths of the J/ψ, Υ, and Z are very small compared to the detector resolution. A bias in the mass scaling is therefore not a major issue for most heavy-ion analyses. However, a degraded resolution could have more visible effects, such as altering the signal to background ratio for resonances, for instance.

The estimated mass scaling is shown as a function of N_{tracks} in Fig. 3.8. The reconstructed mass is lower than the true mass, but this scaling is consistent across the different collision systems and for all values of N_{tracks}, indicating that the momentum scale is unaffected by the level of occupancy in the tracker or the data taking year, and is accurate to approximately $\sim$0.3% in the Z region and about $\sim$0.1% in the J/ψ region.

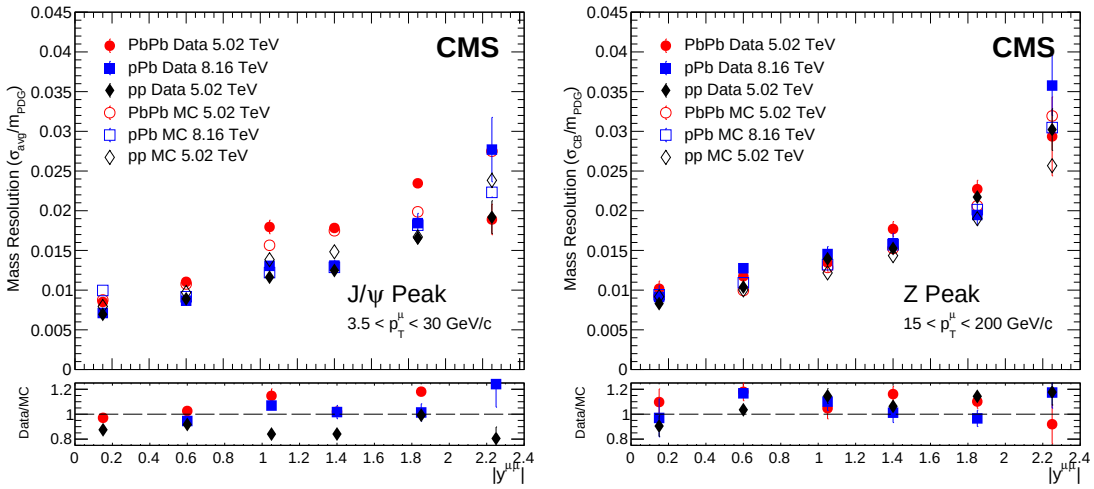

Figure 3.9. Mass resolution as a function of dimuon rapidity $|y^{\mu\mu}|$ estimated from $p+p$ (black diamond), $p+\mathrm{Pb}$ (blue squares), and $\mathrm{Pb}+\mathrm{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances [60].

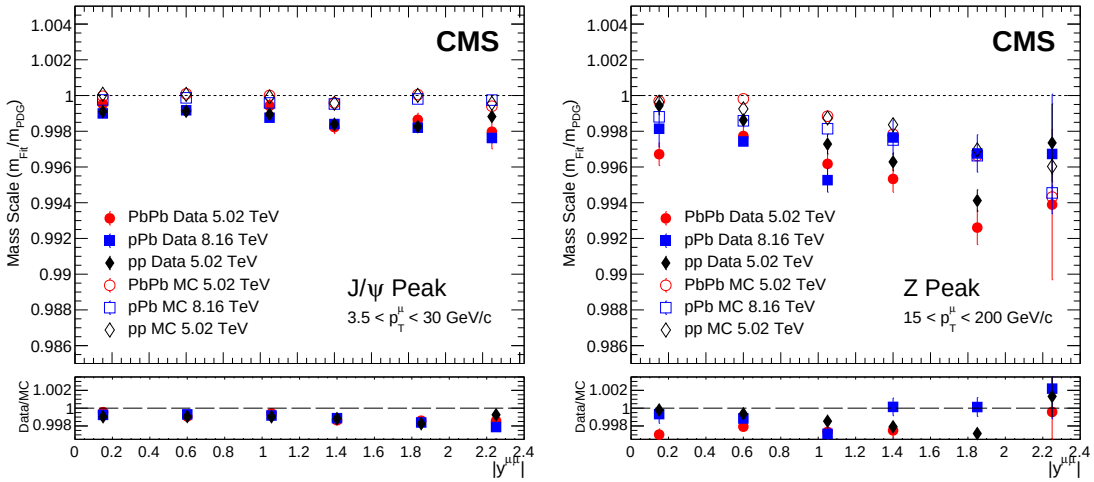

Figure 3.10. Mass scaling as a function of dimuon rapidity $|y^{\mu\mu}|$ estimated from $p+p$ (black diamond), $p+\mathrm{Pb}$ (blue squares), and $\mathrm{Pb}+\mathrm{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances [60].

The dependence of the mass resolution and the mass scaling on dimuon rapidity $y^{\mu\mu}$ was also studied. The mass resolution as a function of $|y^{\mu\mu}|$ is shown in Fig. 3.9 at the J/ψ (left) and Z (right) resonances. The mass scaling is similarly shown in Fig. 3.10 [60]. The resolution varies significantly with rapidity, due mainly to

the non-uniformites in the magnetic field in the forward and backward rapidity regions of the detector which tend to smear the resolution. The different detector technologies in the barrel and endcap also complicate the comparison between the regions.

The mass scaling decreases slightly with rapidity at the J/ψ peak in data but not in MC, presumably due to the said non-uniformities. At the Z peak, the trend is more prominent and is present in both data and MC. In almost all the rapidity bins, the mass scale tends to be lower in data compared to MC. The reason for the differences between data and MC indicates some imperfection in the simulation that warrants further study.

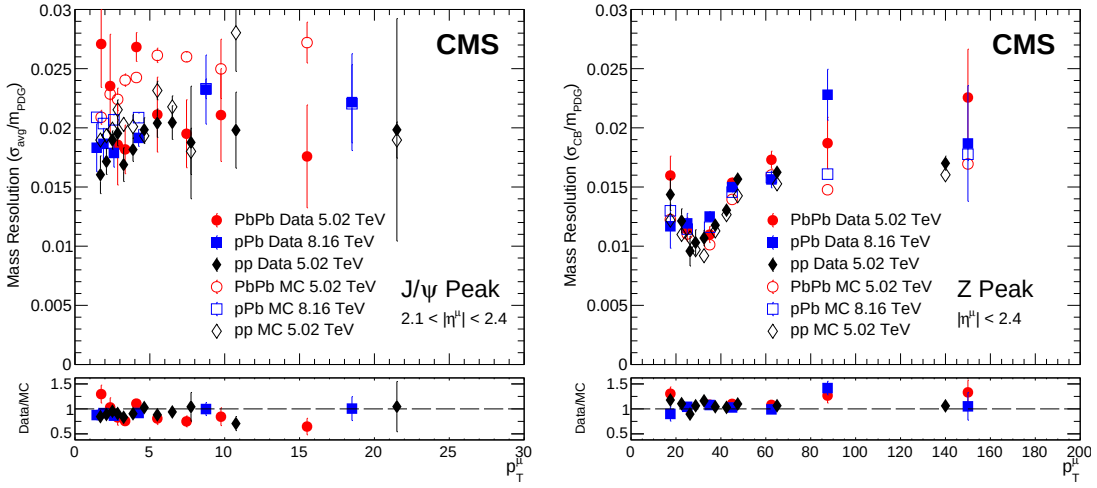

Figure 3.11. Mass resolution as a function of p_T estimated from $p+p$ (black diamond), $p+\mathrm{Pb}$ (blue squares), and $\mathrm{Pb}+\mathrm{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances. The data in the J/ψ region are restricted to $2.1 < |\eta^{\mu}| < 2.4$.

The mass resolution and scaling are also presented as functions of p_T in Fig. 3.11 and Fig. 3.12 at the Z (right) and J/ψ (left) peaks, respectively. In the case of the J/ψ, the pseudorapidity range is restricted to the forward region $2.1 < |\eta| < 2.4$, because that region allows the largest range of p_T. These figures demonstrate that the detector maintains satisfactory performance for reconstructing muons for a large range of momenta, as high as 200 GeV/c and as low as 2 GeV/c, in all three collision systems covering a wide range of N_{tracks}.

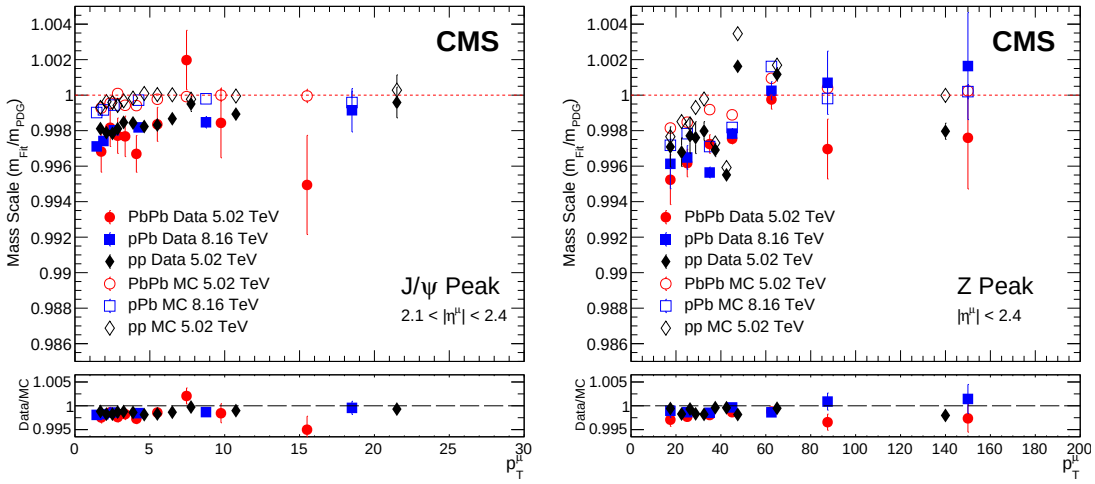

Figure 3.12. Mass scaling as a function of p_T estimated from $p+p$ (black diamond), $p+\mathrm{Pb}$ (blue squares), and $\mathrm{Pb}+\mathrm{Pb}$ (red circles) in real data (filled markers) and MC (open markers) at the J/ψ (left) and Z (right) resonances. The data in the J/ψ region are restricted to $2.1 < |\eta^{\mu}| < 2.4$.

A significant difference in performance between collision systems is not observed as a function of $|y^{\mu\mu}|$ or p_T. However, future studies may restrict $\mathrm{Pb}+\mathrm{Pb}$ to the lower-efficiency region at high N_{tracks} to investigate how the degradation varies with rapidity and p_T.

Estimates of mass resolution and mass scaling were carried out previously on a $p+p$ dataset at a center-of-mass energy $\sqrt{s} = 7$ TeV in 2010 [65]. A comparison of our results versus probe muon pseudorapidity η^{μ} to the previous results is shown in Fig. 3.13, where the division by the PDG mass m_{PDG} has been omitted for ease of comparison. The results of the previous analysis were averaged over the forward and backward pseudorapidity bins in order to show them as a function of $|\eta^{\mu}|$. Despite differences in the collision energy, the data-taking year, the reconstruction algorithms, and the analysis strategies, there is good agreement between the results of the two analyses.

3.5 Performance Summary

Performance studies were carried out using $\mathrm{Pb}+\mathrm{Pb}$ data at $\sqrt{s_{NN}} = 5.02$ TeV and $p+\mathrm{Pb}$ data at $\sqrt{s_{NN}} = 8.16$ TeV. The efficiencies were also measured in

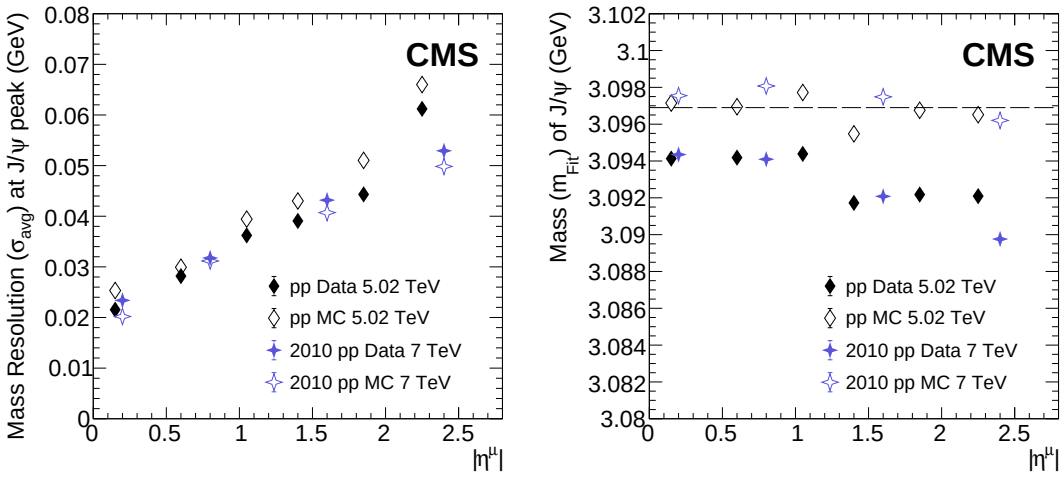

Figure 3.13. Comparison of mass resolution (left) and mass scaling (right) as a function of muon pseudorapidity at the J/ψ peak in $p+p$ collisions from 2017 (black diamonds), without division by the PDG mass m_{PDG}, with estimates from a previous analysis (four-pointed stars). Filled markers represent real data and open markers represent MC [60].

$p+p$ data at $\sqrt{s} = 5.02$ TeV for comparison across all systems as a function of charged-particle multiplicity. Efficiencies of muon reconstruction, identification, and triggering, as well as dimuon mass scales and resolutions were estimated using the data-driven tag-and-probe technique. Previous studies in $p+p$ have now been extended to the heavy-ion environment [60].

We observe a reduction in the high-level trigger efficiency in Pb + Pb at high values of N_{tracks} that are not attainable in the p + Pb or $p+p$ environments. The relative drop in efficiency in Pb + Pb between the lowest and highest N_{tracks} bins is about ∼8%. This reduction reflects the extremely high multiplicities and high occupancies in the inner detector in the most central Pb + Pb collisions. A similar trend is seen in the estimated muon identification efficiencies but to a much lesser extent. The drop in efficiency is on the order of ∼1%. In both cases, a decrease in efficiency with activity was expected due to the fact that both the muon identification and high-level trigger algorithms rely on information from the inner tracker, which is heavily occupied in the most central Pb + Pb collisions.

In all cases, the efficiencies are high (typically above 90%), even with extremely

high occupancies, and the dimuon mass scale and resolution are observed to be unaffected by the high-occupancy environment. This excellent performance of the CMS detector has made possible a robust muon and dimuon program in the Pb+Pb environment, leading to many muon-based measurements, such as the J/ψ the Υ, the Z, the top quark, and Z + jets.

Chapter 4

Data Selection and Simulation

4.1 Dataset

The data analyzed in this thesis are from the most recent Pb+Pb run at the LHC, collected in November and December 2018 with $\sqrt{s_{NN}} = 5.02$ TeV using a dimuon trigger in CMS, and correspond to an integrated luminosity of 1.6 nb^{-1}. The data are stored in a format called analysis object data (AOD), which contains all of the information about the event. The prompt AOD datasets are listed in Table 4.2.

The data were reduced to an `Onia` (from quark-"onia") tree which contains information on all the dimuons in the dataset along with general event information such as centrality, N_{tracks}, E_T^{HF}, etc. The `Onia` tree used in this thesis is stored in CERN's open source storage (EOS) at the following location:

`/eos/cms/store/group/phys_heavyions/dileptons/Data2018/PbPb502TeV/`

`TTrees/ReReco/AOD/DoubleMuon/ReReco_Oniatree_addvn_part`$*$.root.

The `Onia` tree includes information on all the triggers satisfied by each dimuon. These are stored in a binary format where each trigger is represented by one bit. In each bit, a 1 means that the trigger was satisfied and a 0 means that it was not. The high-level triggers corresponding to each bit are listed in Table 4.1. The information in the table is called the "HLT menu" for the 2018 Pb + Pb run.

Each trigger has a descriptive name beginning with `HLT` which stands for "high-level trigger," followed by `HI` which stands for "heavy ion" because these triggers

are designed for heavy-ion events. The level of the trigger is denoted by `L1`, `L2`, or `L3`, and is related to how loose or tight the trigger criteria are. Some triggers look for single muons with p_T above a certain threshold, such as `Mu12` which triggers on a single muon with $p_T > 12$ GeV/c. Other triggers look for dimuons, hence the `DoubleMu` in the name. Cuts can be applied to centrality (`Centrality`), mass (`M`), number of quality hits in the tracker (`NHitQ`), and other variables. Triggers can be updated from time to time, and the version number is displayed at the end of the name, usually as `v1`.

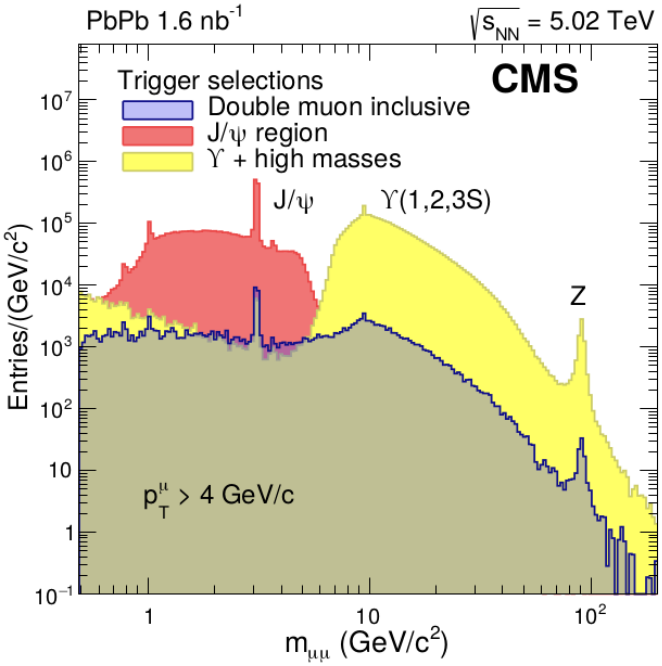

Figure 4.1. The full dimuon invariant mass distribution for dimuons in the 2018 Pb + Pb dataset satsifying `HLT_HIL1DoubleMuOpen_v1` (Double muon inclusive) in blue, `HLT_HIL3Mu0NHitQ10_L2Mu0_MAXdR3p5_M1to5_v1` (J/ψ region) in red, and `HLT_HIL3Mu2p5NHitQ10_L2Mu2_M7toinf_v1` (Υ + high masses) in yellow [66].

Figure 4.1 demonstrates how the HLTs are used to select events of interest. The figure shows the dimuon invariant mass distribution for dimuons passing three distinct triggers:

- `HLT_HIL1DoubleMuOpen_v1` (Double muon inclusive),
- `HLT_HIL3Mu0NHitQ10_L2Mu0_MAXdR3p5_M1to5_v1` (J/ψ region),

Table 4.1. HLT menu for 2018 Pb+Pb data. The prescales and rates are typical values and are collected from LHC run 327524 on December 1st, 2018.

Bit	Trigger Name	Prescale	Unprescaled Rate (Hz)
1	HLT_HIL1DoubleMuOpen_v1	1223	2225.86
2	HLT_HIL1DoubleMuOpen_OS_Centrality_40_100_v1	1	105.75
3	HLT_HIL1DoubleMuOpen_Centrality_50_100_v1	1	129.35
4	HLT_HIL1DoubleMu10_v1	1	2.30
5	HLT_HIL2_L1DoubleMu10_v1	1	0.75
6	HLT_HIL3_L1DoubleMu10_v1	1	0.46
7	HLT_HIL2DoubleMuOpen_v1	953	1143.60
8	HLT_HIL3DoubleMuOpen_v1	1	53.22
9	HLT_HIL3DoubleMuOpen_M60120_v1	1	0.40
10	HLT_HIL3DoubleMuOpen_JpsiPsi_v1	1	8.02
11	HLT_HIL3DoubleMuOpen_Upsi_v1	1	5.75
12	HLT_HIL3Mu0_L2Mu0_v1	269	333.56
13	HLT_HIL3Mu0NHitQ10_L2Mu0_MAXdR3p5_M1to5_v1	1	65.49
14	HLT_HIL3Mu2p5NHitQ10_L2Mu2_M7toinf_v1	1	59.53
15	HLT_HIL3Mu3_L1TripleMuOpen_v1	0	0.00
16	HLT_HIL1MuOpen_Centrality_70_100_v1	0	0.00
17	HLT_HIL1MuOpen_Centrality_80_100_v1	1	37.58
18	HLT_HIL2Mu3_NHitQ15_v1	0	0.00
19	HLT_HIL2Mu5_NHitQ15_v1	0	0.00
20	HLT_HIL2Mu7_NHitQ15_v1	0	0.00
21	HLT_HIL3Mu3_NHitQ10_v1	1	148.45
22	HLT_HIL3Mu5_NHitQ10_v1	1	43.12
23	HLT_HIL3Mu7_NHitQ10_v1	1	12.64
24	HLT_HIL3Mu12_v1	1	12.39
25	HLT_HIL3Mu15_v1	1	6.78
26	HLT_HIL3Mu20_v1	1	3.69

- and HLT_HIL3Mu2p5NHitQ10_L2Mu2_M7toinf_v1 (Υ + high masses),

corresponding to bits 1, 13, and 14, respectively. The J/ψ trigger focuses on the J/ψ region and cuts out much of the higher mass dimuons that would not be needed in a J/ψ analysis. The Υ trigger does the same for the Υ, cutting out much of the lower mass dimuons and focusing on the Υ peak. The double muon inclusive trigger selects all dimuons, but does not necessarily contain the other two datasets. Similar plots have been made for every LHC run, in which the dimuon mass distribution has always been drawn in yellow, and hence this type of plot has received the nickname "Yellow Plot."

A large prescale ($\sim$1000) was applied to the double muon inclusive trigger during the run, so that only about 1 out of every 1000 events that could satisfy this trigger were actually recorded. Hence the prescale effectively reduces the sampled luminosity of the triggered events by the prescale factor. The prescale was put in place because of the high rate of events that could satisfy the trigger. The rate, which can reach above 3 kHz, had to be slowed down to allow the proper amount of time for the trigger to analyze and store each event. This large prescale can be contrasted with the Υ trigger, which had a prescale of 1. If it were possible to give the double muon inclusive trigger a prescale of 1, then the normalization would effectively be increased by a factor of 1000, and the histogram would contain far more events than the other triggers.

Upsilon candidate events for this thesis were selected using the specialized trigger

HLT_HIL3Mu2p5NHitQ10_L2Mu2_M7toinf_v1,

corresponding to bit 14, which requires a pair of muons with dimuon invariant mass at least 7 GeV/c^2. One of the muons is required to pass an L2 trigger, reconstructed with track fits in the outer muon spectrometer, with $p_T > 2$ GeV/c. The other muon is required to pass an L3 trigger, reconstructed from an L2 muon combined with inner tracker information, with $p_T > 2.5$ GeV/c. The L3 muon is also required to have at least 10 high-quality hits (NHitQ10). The efficiency of this

trigger will be discussed and efficiency corrections will be calculated in Sec. 5.3.

4.2 Monte Carlo Simulation

Table 4.2. Monte Carlo samples for 2018 5.02 TeV Pb + Pb collision

Type	Dataset	Events
Double Muon triggered data	/HIDoubleMuon/HIRun2018A−04Apr2019−v1/AOD	69507589
Minbias MC for embedding	/MinBias_Hydjet_Drum5F_2018_5p02TeV/ HINPbPbAutumn18GS−103X_upgrade2018 _realistic_HI_v11−v1/GEN−SIM	999111
Embedded MC	/Upsilon1S_pThat−2_TuneCP5_HydjetDrumMB _5p02TeV_Pythia8/HINPbPbAutumn18DR−mva98 _103X_upgrade2018_realistic_HI_v11_ext1−v1/AODSIM	4987183
Embedded MC	/Upsilon1S_pThat−2_TuneCP5_HydjetDrumMB _5p02TeV_Pythia8/HINPbPbAutumn18DR−mva98 _103X_upgrade2018_realistic_HI_v11−v1/AODSIM	4749176

Monte Carlo simulations of the collisions and detector response are useful for testing whether or not current physics models accurately reproduce the real data that has been collected, and also serve as a testing ground for determining how well the detectors reconstruct the data and how accurately the analysis methods extract quantities of interest.

The MC used in the v_2 analysis contains Υs embedded in Pb + Pb collision events reconstructed by a virtual CMS detector. This MC is used in calculations of acceptance and efficiency, and in a closure test described in Sec. 6.6. Signal events are generated using the high-energy physics event generator PYTHIA8.212 [67] with a set of parameters ("tunes") called Tune CP5 [68]. These tunes are based on the NNPDF3.1 set of parton distribution functions [69] at leading (LO), next-to-leading (NLO), or next-to-next-to-leading order (NNLO) calculations in perturba-

tive quantum chromodynamics and the strong coupling evolution at LO or NLO at 5.02 TeV. The cross sections of Υ mesons were calculated from ${}^3S_1$,${}^3P_J$,${}^3D_J$ bottomonium states (see Fig. 1.8) via the colour-singlet [70, 71] and colour-octet [72] mechanisms. Events are required to include at least one Υ(1S) with $p^{\Upsilon} > 2.5$ GeV/c that decays to two daughter muons satisfying $p_T^{\mu} > 0$ and $|\eta^{\mu}| < 2.5$. All the generated Υ events are embedded into a Pb + Pb background event simulated by a heavy-ion event generator called HYDJET ("HYDrodynamics plus JETs") version 1.8, tune Drum5F [73]. The MC events are stored in AOD format to match the real data. The MC datasets are shown in Table 4.2.

4.3 Acceptance Cuts

The CMS detector covers a wide kinematic range, but it does not cover the entire available phase space. For example, many muons never reach the muon chambers becase their transverse momentum p_T is too low. In order for a muon to reach the muon detectors in the barrel region ($|\eta| < 1$) outside the solenoid (about 4 m from the interaction point), the radius of curvature must satisfy $R \geq 2$ m. This means the minimum p_T required for a muon to be detected in that region is, using Eq. (2.7),

$$p_{T,\text{min}} \approx (1.609 \times 10^{-19}\ \text{C})(3.8\ \text{T})(2\ \text{m}) \approx 1.2 \times 10^{-18}\ \frac{\text{kg} \cdot \text{m}}{\text{s}} \approx 2.3\ \text{GeV}/c.$$

The muon detectors in the endcap extend the pseudorapidity coverage to $|\eta| < 2.4$. The p_T threshold for muon detection varies with pseudorapidity, since the endcap detectors are closer to the beamline. A plot of the distribution of muons produced from Υ decays in the MC dataset is shown in Fig. 4.2 as a function of p_T and $|\eta|$. An "acceptance cut" is applied to the data, indicated by the red curve. Specifically, we accept only muons into the dataset with $p_T > 3.5$ GeV/c and within the pseudorapidity range $|\eta| < 2.4$.

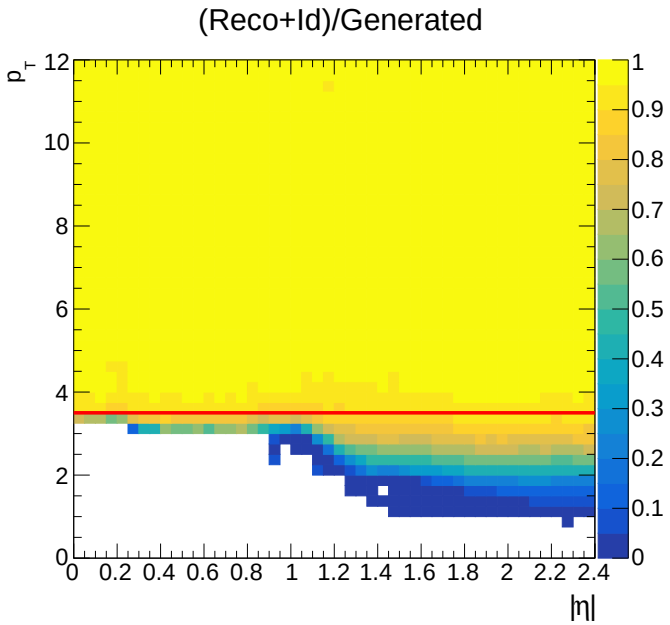

Figure 4.2. The distribution of muons produced from Υ decays in MC that pass reconstruction and hybrid-soft muon identification criteria versus p_T and $|\eta|$. The red line shows the acceptance cut.

4.4 Additional Cuts

Additional cuts are applied to the dimuon dataset in order to reduce background and focus on the Υ signal. Only pairs of muons with opposite sign are accepted, because the two daughter muons resulting from the decay of the Υ always have opposite sign. Each candidate muon is required to pass an idenfication requirement called `HybridSoftId`, which imposes the following requirements on the muon:

- at least 6 hits in the silicon strip layers,
- at least one hit in the pixel detectors,
- a distance of closest approach between the track and the event vertex less than 0.3 cm in the transverse plane and less than 20 cm in the longitudinal direction,
- the track should be matched to a track in the muon detectors.

Additionally, the reconstructed dimuon vertex probability is required to be greater than 1%.

Finally, the dimuon invariant mass is restricted to the region

$$8 \text{ GeV}/c^2 \leq m_{\mu^+\mu^-} \leq 14 \text{ GeV}/c^2.$$

This mass window is typically used in Υ analyses, as it provides a wide range for extraction of the Υ(1S), Υ(2S), and Υ(3S) signals. Since there is no interference from other known dimuon resonances, this mass region provides a good anchor for the background description.

Chapter 5

Analysis

5.1 Signal Extraction

The raw yields of the Υ states are obtained via unbinned maximum-likelihood fits to the invariant mass spectra. Quarkonia signals are often fit with a Crystal Ball (CB) function, first implemented in Ref. [74]. The CB function consists of a Gaussian peak with a power-law tail grafted onto the lower-valued side such that the function and its derivative are continuous. It is given by,

$$\mathrm{CB}(x;\bar{x},n,\alpha,\sigma) = N \cdot \begin{cases} \exp(-\frac{(x-\bar{x})^2}{2\sigma^2}) & \text{for } \frac{x-\bar{x}}{\sigma} > -\alpha \\ A \cdot (B - \frac{x-\bar{x}}{\sigma})^{-n} & \text{for } \frac{x-\bar{x}}{\sigma} \leq -\alpha\,, \end{cases} \tag{5.1}$$

where

$$\begin{aligned} A &= \left(\frac{n}{|\alpha|}\right)^n \cdot \exp\left(-\frac{|\alpha|^2}{2}\right)\,, \\ B &= \frac{n}{|\alpha|} - |\alpha|\,. \end{aligned} \tag{5.2}$$

The advantage of the CB over a simple Gaussian is that its power-law tail can account for quarkonia which are reconstructed at lower masses due to energy loss of the decay products, usually via bremsstrahlung radiation of the muons in the detector material.

Due to the varying resolution in the detector in the forward and midrapidity regions (see Fig. 3.9), the Υ signals at CMS are fit with a sum of two functions with different widths. Two possible combinations were tried, a CB plus a Gaussian, and

a sum of two CBs called a "Double Crystal Ball" (DCB). After some goodness-of-fit tests using Monte Carlo data, it was decided that the DCB performed better and it was chosen as the nominal fit function for the Υ resonances [27, 31]. The double Crystal Ball signal function is defined by

$$\begin{aligned}\Sigma_{\mathrm{1S}}(m_{\mu\mu}; m_{\mathrm{1S}}, n_1, n_2, &\alpha_1, \alpha_2, \sigma_1, f, x) \\ &= f \cdot \mathrm{CB}_1\left(m_{\mu\mu}; m_{\mathrm{1S}}, n_1, \alpha_1, \sigma_1\right) \\ &\quad + (1-f) \cdot \mathrm{CB}_2\left(m_{\mu\mu}; m_{\mathrm{1S}}, n_2, \alpha_2, x \cdot \sigma_1\right), \end{aligned} \tag{5.3}$$

and normalized by the Υ(1S) yield $\mathcal{N}_{\mathrm{1S}}$. The parameters α and n are the tail parameters, where α is related to the location where the power law tail is grafted onto the Gaussian and n is the exponent. The parameter f is the ratio of the amount of signal contained in the first CB over the amount of signal in the second CB. The parameter σ represents the Gaussian width of the first CB, and x represents the width of the second CB over σ. The CB plus Gaussian is defined similarly and was kept as a reasonable alternative fit function to be used for estimating systematic uncertainties.

The background function is an error function multiplied by a decaying exponential, defined by

$$\mathcal{B}\left(m_{\mu\mu}; \mu_{\mathrm{err}}, \sigma_{\mathrm{err}}, \lambda\right) = \exp\left(-\frac{m_{\mu\mu}}{\lambda}\right) \cdot \frac{1 + \mathrm{Erf}\left(\frac{m_{\mu\mu}-\mu_{\mathrm{err}}}{\sqrt{2}\sigma_{\mathrm{err}}}\right)}{2}. \tag{5.4}$$

The decaying exponential models the dominant combinatorial background. The error function is there to model the turn-on shape in the dimuon invariant mass spectrum around ~ 8 GeV/c^2 which arises from the single-muon requirement $p_T^{\mu} > 4$ GeV/c. In the high-p_T^{Υ} bins ($p_T^{\Upsilon} > 6$ GeV/c), the turn-on becomes negligible and the background can be modeled simply by an exponential.

Due to the large number of parameters in the three DCB functions (nine for each DCB), several physics-motivated constraints were applied. First, the Υ(2S) and Υ(3S) masses and widths were constrained such that their ratio to the mass and width of the Υ(1S) should be the same as the ratio of the world-average masses

$m^{n\mathrm{S}}_{\mathrm{PDG}}$ as defined by the PDG [24] as follows,

$$m_{n\mathrm{S}} = m_{1\mathrm{S}} \cdot \frac{m^{n\mathrm{S}}_{\mathrm{PDG}}}{m^{1\mathrm{S}}_{\mathrm{PDG}}}, \tag{5.5}$$

$$\sigma_{n\mathrm{S}} = \sigma_{1\mathrm{S}} \cdot \frac{m^{n\mathrm{S}}_{\mathrm{PDG}}}{m^{1\mathrm{S}}_{\mathrm{PDG}}}, \tag{5.6}$$

while the mass and width of the $\Upsilon(1\mathrm{S})$ were left as free parameters. The world-average masses $m^{n\mathrm{S}}_{\mathrm{PDG}}$ are tabulated in Table 1.2.

Within each DCB, the radiative tail parameters of the two CB functions were constrained to be the same (*i.e.* $\alpha_1 = \alpha_2$ and $n_1 = n_2$). The three Υ states were also constrained to have the same tail parameters and the same fraction of yield in one CB over the other, since the physics of the radiative tail is the same for all three, and the effects of the kinematic differences will be much smaller than the uncertainties on the tail parameters. With these constraints, the DCB is now more simply defined as

$$\begin{aligned}&\Sigma_{1\mathrm{S}}(m_{\mu\mu};\, m_{1\mathrm{S}}, n, \alpha, \sigma, f, x) \\ &\qquad = f \cdot \mathrm{CB}_1\left(m_{\mu\mu}; m_{1\mathrm{S}}, n, \alpha, \sigma\right) + (1-f) \cdot \mathrm{CB}_2\left(m_{\mu\mu}; m_{1\mathrm{S}}, n, \alpha, x\cdot\sigma\right),\end{aligned} \tag{5.7}$$

and the functions for the higher states are given by

$$\Sigma_{n\mathrm{S}}\left(m_{\mu\mu}\right) = \Sigma_{1\mathrm{S}}\left(m_{\mu\mu}; m_{1\mathrm{S}} \cdot \frac{m^{n\mathrm{S}}_{\mathrm{PDG}}}{m^{1\mathrm{S}}_{\mathrm{PDG}}}, n, \alpha, \sigma \cdot \frac{m^{n\mathrm{S}}_{\mathrm{PDG}}}{m^{1\mathrm{S}}_{\mathrm{PDG}}}, f, x\right). \tag{5.8}$$

The total signal is the sum of the three DCBs scaled by the yields $\mathcal{N}_{n\mathrm{S}}$,

$$\begin{aligned}&\mathcal{S}(m_{\mu\mu}; \mathcal{N}_{1\mathrm{S}}, \mathcal{N}_{2\mathrm{S}}, \mathcal{N}_{3\mathrm{S}}, m_{1\mathrm{S}}, n, \alpha, \sigma, f, x) \\ &\qquad = \mathcal{N}_{1\mathrm{S}} \cdot \Sigma_{1\mathrm{S}}\left(m_{\mu\mu}\right) + \mathcal{N}_{2\mathrm{S}} \cdot \Sigma_{2\mathrm{S}}\left(m_{\mu\mu}\right) + \mathcal{N}_{3\mathrm{S}} \cdot \Sigma_{3\mathrm{S}}\left(m_{\mu\mu}\right).\end{aligned} \tag{5.9}$$

After these constraints, and putting all the functions together, there are 9 total signal parameters (α, f, $m_{1\mathrm{S}}$, n, σ, x, and the 3 yields) and 4 background parameters (λ, μ_{err}, σ_{err}, and the background yield) which are allowed to vary in

the fits. The final fit function can be summarized as

$$\begin{aligned}\mathcal{F}(m_{\mu\mu};\, &\mathcal{N}_{1S}, \mathcal{N}_{2S}, \mathcal{N}_{3S}, \mathcal{N}_{\text{Bkg}}, m_{1S}, n, \alpha, \sigma, f, x, \mu_{\text{err}}, \sigma_{\text{err}}, \lambda) \\ &= \mathcal{S}(m_{\mu\mu}; \mathcal{N}_{1S}, \mathcal{N}_{2S}, \mathcal{N}_{3S}, m_{1S}, n, \alpha, \sigma, f, x) + \mathcal{N}_{\text{Bkg}} \cdot \mathcal{B}(m_{\mu\mu}; \mu_{\text{err}}, \sigma_{\text{err}}, \lambda).\end{aligned} \tag{5.10}$$

In order to narrow down the relevant phase space for the signal parameters, an iterative fitting approach was taken. In the first round of fitting, the five signal shape parameters (α, f, n, σ, and x) were all free. The fits were required to pass a set of quality checks: First, the chi-squared per degree of freedom (χ^2/dof) must be less than 4.0 and greater than 0.5, and second, the fit parameters α and n must not hit their limits. Specifically, the distance of the fitted value from each limit must be greater than 3% of the total available range of the parameter. The parameter limits were the same as those used in other recent Υ analyses by CMS [27, 31].

In each of the subsequent rounds of fitting, certain signal parameters were fixed in order to decrease the available phase space and determine what values could be considered physically reasonable for the other free parameters. The information obtained from these iterative fits was used to construct Gaussian constaints on the signal parameters for the nominal fits. In this way, the parameters were still mostly free, but constrained to physically reasonable regions of the parameter space. The process is outlined below:

- Round R0: All parameters free.
- Round R1: α fixed to the average value from R0. The average value of n and its RMS from this round are used to construct the Gaussian constraint on n.
- Round R2: n fixed to the average value from R1. The average value of α and its RMS from this round are used to construct the Gaussian constraint on α.

- Round R3: n fixed to the average value from R1 and α fixed to the average value from R2. The average value of x and its RMS from this round are used to construct the Gaussian constraint on x.

- Round R4: (Nominal fits) Gaussian constraints are applied to n, α, and x as determined from their average values from the previous fits. The parameter f is also constrained bin by bin using the fitted value and error from each fit in round R3.

- Alternative constraints for estimating systematic uncertainty are obtained by first fixing n in round R1 instead of α.

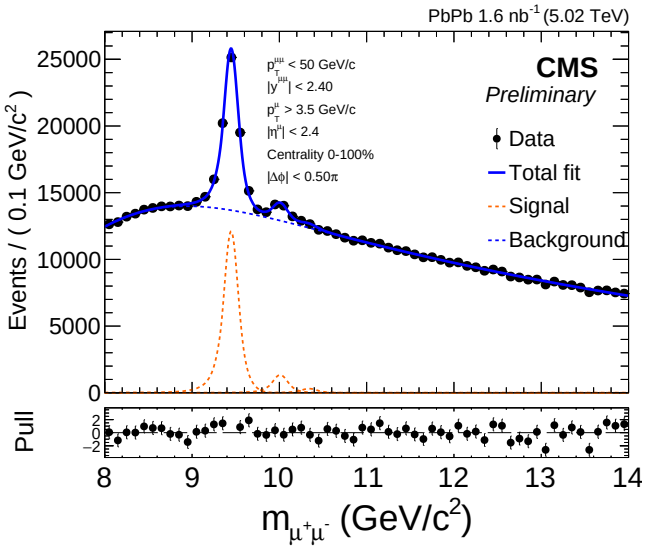

Figure 5.1. An unbinned maximum-likelihood fit to the invariant mass spectra in Pb + Pb at $\sqrt{s_{NN}} = 5.02$ TeV in the integrated bin. The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

An example fit is shown in Fig. 5.1, where the data are shown as black circles and the fit function is the blue curve. The separate signal and background functions are displayed as red and blue dashed curves, respectively. The pull distribution in the bottom panel plots the difference between the data and the fitted curve,

divided by the uncertainty in the data. The total signal extracted from this fit to uncorrected data in the integrated bin is displayed in Table 5.1.

Table 5.1. The total yields of Υ(nS) in the uncorrected 2018 Pb + Pb dataset, obtained from an unbinned maximum-likelihood fit in the integrated bin.

Quarkonium state	Total yield
Υ(1S)	29381.5 $\pm$ 588.139
Υ(2S)	3438.51 $\pm$ 289.264
Υ(3S)	817.114 $\pm$ 274.433

The final nominal fits for the v_2 analysis are carried out simultaneously in four bins of $\Delta\phi$, the azimuthal angle made by the Υ relative to the event plane, as described in Sec. 5.6. These fits can be seen in Figs. 5.8 through 5.17.

5.2 Acceptance Corrections

While the acceptance cut on the muon kinematics described in Sec. 4.3 ensures that we have good data, it results in an underestimate of the number of Υs in the events. Not all of the Υs in the kinematic region of interest ($|\eta| < 2.4$) will be reconstructed from the accepted muons. Some of these Υs decay into two muons of which one may fall outside the acceptance. Because of the cut, this perfectly good Υ will be left out of the analysis. To account for these missing Υs, we rescale the extracted Υ yields by factors called "acceptance corrections" which we estimate from Monte Carlo simulations.

We estimate acceptance corrections using GEN-level Monte Carlo (before the simulation of the detector response and reconstruction), so that we have *a priori* knowledge of all the Υs decaying to two muons. We check how many of the Υs have both muons passing the acceptance cut. The ratio of passing Υs over total Υs is the acceptance correction factor. The correction factors are estimated as a function of p_T^{Υ} and then applied to the data by weighting each Υ by the inverse of the correction factor.

It can be seen in Fig. 5.2 that the acceptance corrections show a strong depen-

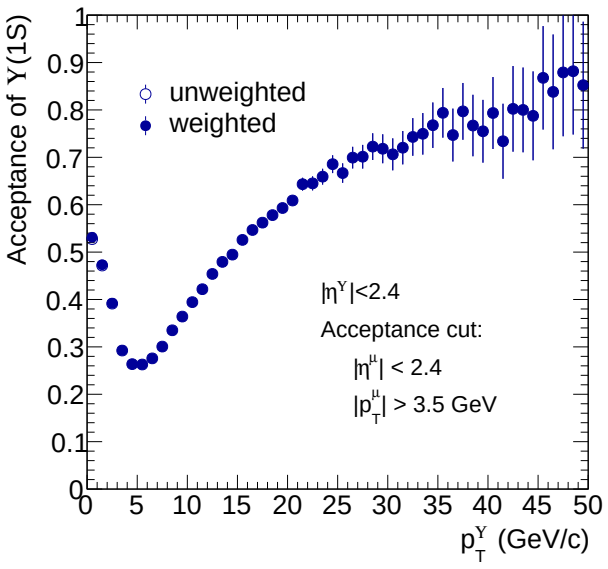

Figure 5.2. Acceptance corrections for Υ(1S) as a function of p_T^{Υ} with (closed circles) and without (open circles) p_T^{μ} weighting. The effect of the weighting is negligible.

dence on p_T^{Υ}. A large number of Υs are missed at low p_T^{Υ} with a minimum around ~ 5 GeV/c, but a progressively larger portion of the Υs are properly captured with increasing p_T^{Υ}. The dependence can be understood from the decay kinematics. The Υs with high p_T^{Υ} are more likely to have their momentum oriented toward the barrel region of the detector, so the decaying muons are more likely to be detected in that region as well. The angle between the two daughter muons decreases with increasing p_T^{Υ}, resulting in an even greater likelihood of each muon being detected in the barrel region. The low-p_T^{Υ} Υs, on the other hand, decay with a very wide angle, making it more likely that one of the muons will lie outside the barrel region and miss the acceptance cut. However, at very low p_T^{Υ} (< 5 GeV/c), the Υs are likely headed toward the endcaps, and a significant number of these will decay to two muons with a large enough opening angle that they will both be detected in the barrel region.

To improve the accuracy of the MC simulation, the p_T^{μ} distribuion of single muons is weighted to match the corresponding p_T^{μ} distribution in data. The accep-

tance corrections as a function of p_T^{Υ} are shown in Fig. 5.2 with and without the p_T^{μ} weights. The effect of the weighting is negligible.

5.3 Efficiency Corrections

Similar to the acceptance cut, the steps taken to reconstruct, identify, and trigger on Υs are not 100% efficient and therefore underestimate the true Υ yield. Hence, the yields must be scaled up by "efficiency corrections" in the same way that the acceptance corrections were applied. We estimate the efficiencies of the selection criteria using generated (GEN) and reconstructed (RECO) MC. The GEN MC allows us to know exactly how many Υs were generated in each simulated event. The RECO MC has passed through a simulated detector response, and thus allows us to estimate the fraction of Υs that passed all the selection criteria. All dimuons in RECO and GEN MC must have already passed the acceptance cut.

As noted in Sec. 3.1, the MC is not a perfect description of the real data, and these differences can affect the efficiencies that are estimated from MC. It is therefore necessary to rescale the efficiencies by scale factors derived from comparisons of data to MC. These scale factors are calculated from single-muon efficiencies using the tag-and-probe method, as discussed in Sec. 3.1, where the scale factor is the ratio of the single-muon efficiency estimated from data over the single-muon efficiency from MC, as follows,

$$\mathrm{SF}(p_T^{\mu}, \eta^{\mu}) = \frac{\varepsilon^{\mu}_{\mathrm{data}}(p_T^{\mu}, \eta^{\mu})}{\varepsilon^{\mu}_{\mathrm{MC}}(p_T^{\mu}, \eta^{\mu})}. \tag{5.11}$$

The Υ efficiency is then calculated as

$$\varepsilon^{\Upsilon} = \frac{N^{\mu^+\mu^-}_{\mathrm{RECO}}[\mathrm{acc.}, \mathrm{muid.}, \mathrm{trig.}]}{N^{\Upsilon}_{\mathrm{GEN}}[\mathrm{acc.}]} \otimes [\mathrm{SF}(p_T^{\mu-}, \eta^{\mu-}) \times \mathrm{SF}(p_T^{\mu+}, \eta^{\mu+})], \tag{5.12}$$

where the $\otimes$ indicates that, for each dimuon, the appropriate scale factor on the right is applied independently to each of the two muons. In the case of the Υ trigger, the scale factor applied to each muon depends on whether the muon passed the `L2` or `L3` single-muon trigger (see the Υ trigger description in Sec. 4.1). In the cases where both muons pass the `L3` trigger, one of them is randomly chosen to act as the

L2 muon. The numerator $N^{\mu^+\mu^-}_{\rm RECO}[\text{acc.}, \text{muid.}, \text{trig.}]$ is the number of reconstructed opposite-sign dimuons that pass the acceptance cut, muon identification, and the Υ trigger. The denominator $N^{\Upsilon}_{\rm GEN}[\text{acc.}]$ is the number of generated Υs that pass the acceptance cut.

The final estimated efficiency corrections are shown in Fig. 5.3 as a function of p_T in bins of centrality. The efficiencies obtained without the p^{μ}_T weighting are also displayed. The effect of the weighting is observed to cause an increase in the efficiency at low p_T and a reduction at high p_T. The inverse of the efficiencies obtained with the p^{μ}_T weighting were applied as weights to each Υ candidate in the data before fitting.

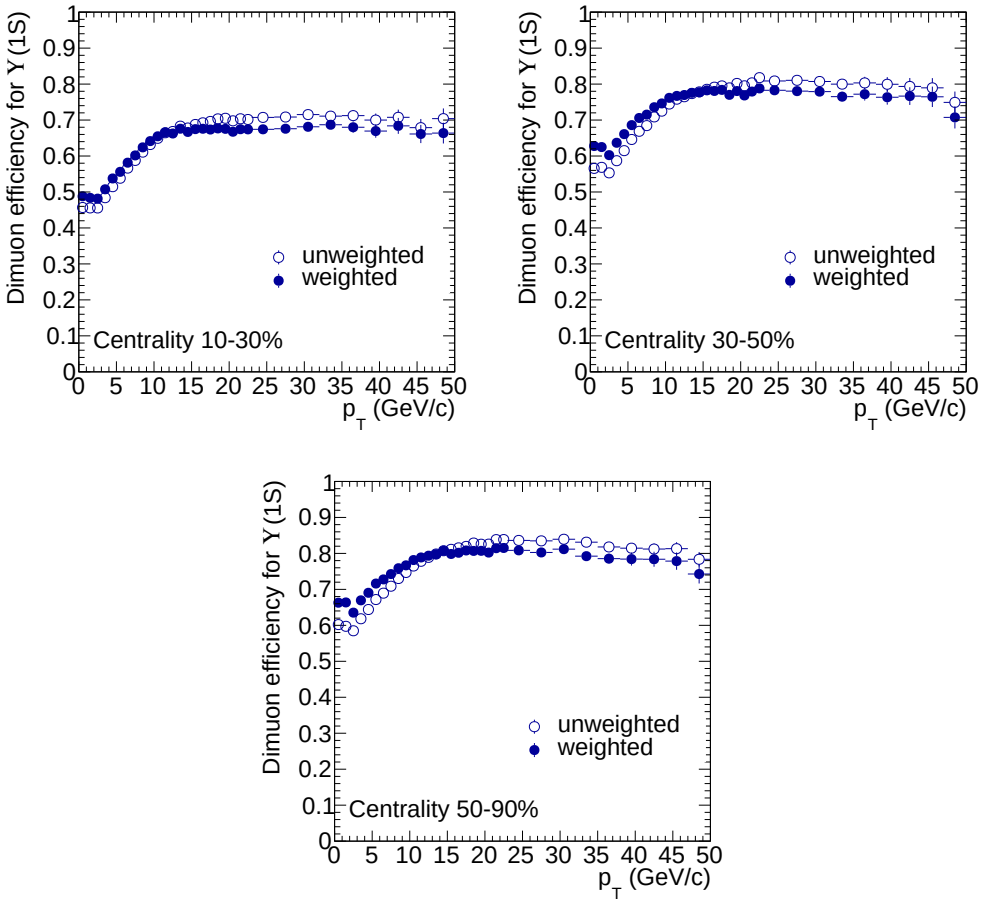

Figure 5.3. Dimuon efficiency corrections for Υ(1S) as a function of $p^{\mu\mu}_T$ in centrality ranges 0-30% (upper left), 30-50% (upper right), and 50-100% (lower middle), with (closed circles) and without (open circles) p^{μ}_T weighting.

5.4 Event Plane Estimate

In this analysis, the elliptic flow coefficient v_2 is estimated using the event plane method. This method has been used previously to estimate the elliptic flow of charmonium [47]. The event plane is defined by the beam direction and the line connecting the centers of the two nuclei. It is customary to use the term "reaction plane" when referring to the true plane which cannot be known in real data, and to use the term "event plane" when referring to our estimate of where the plane is, although they are often used interchangeably. The angle ψ is the angle between the event plane (or reaction plane) and the x-axis, as seen in Fig. 5.4. The azimuthal angles in a v_2 analysis must all be measured with respect to this plane, which is why the azimuthal plots and equations in this thesis are in terms of $\Delta\phi = \phi - \psi$ rather than ϕ.

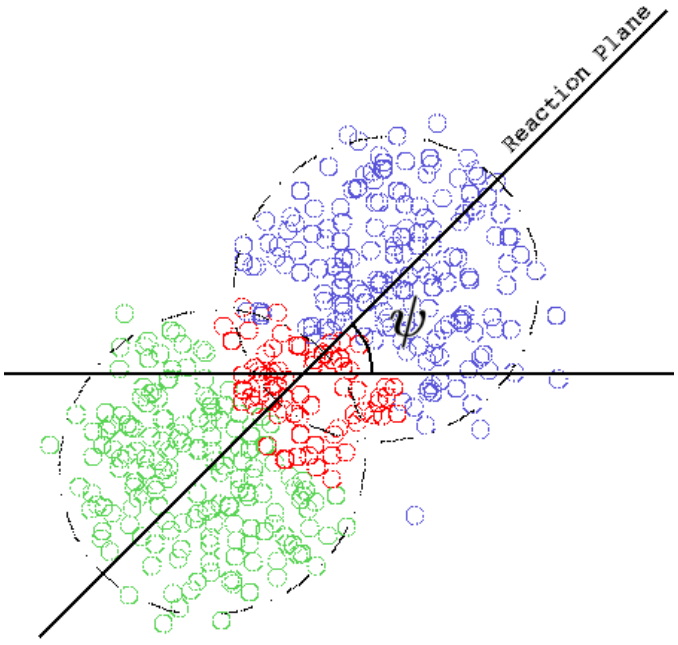

Figure 5.4. Depiction of a heavy-ion collision event, generated by the Glauber model simulation described in Sec. 1.3.1. The beamline extends into and out of the page. Nucleons are represented by the blue, green, and red circles. The participating nucleons are marked in red. The reaction plane angle is indicated by the angle ψ.

The event plane is estimated using energy deposited in a region of the HF

spanning $3 < |\eta| < 5$. The decision to use this region of the detector avoids autocorrelations by avoiding overlap between the data used to find the event plane and the data used to calculate the v_2, given the event plane. We define a vector

$$Q_n = (Q_{nx}, Q_{ny}) \tag{5.13}$$

pointing along the event plane angle where

$$Q_{nx} = \sum_i w_i \cos(n\phi_i) \tag{5.14}$$

and

$$Q_{ny} = \sum_i w_i \sin(n\phi_i), \tag{5.15}$$

where the sums are over all the particles, and the weights, w_i, are the energies of the particles. The subscript n represents the order of the event plane. Measurements of elliptic flow make use of the second-order event plane, $n = 2$. The event plane angle is then given by

$$\psi_n = \frac{1}{n}\tan^{-1}\left(\frac{Q_{ny}}{Q_{nx}}\right). \tag{5.16}$$

Given the randomness of the event plane angle, we expect the average values of Q_{nx} and Q_{ny} over many events to be zero. If they are not, then it is likely that the beam line does not pass through the exact center of the CMS detector and it is necessary to recenter the distributions as follows:

$$Q'_{nx} = Q_{nx} - \langle Q_{nx} \rangle, \tag{5.17}$$

$$Q'_{ny} = Q_{ny} - \langle Q_{ny} \rangle. \tag{5.18}$$

Then the recentered reaction plane angle is given by

$$\psi'_n = \frac{1}{n}\tan^{-1}\left(\frac{Q'_{ny}}{Q'_{nx}}\right). \tag{5.19}$$

The components of the Q_n vector before and after recentering are plotted in Fig. 5.5. There is a slight shift between the raw and recentered distributions, but it is very small compared to the width of the distributions. The effect on the event plane distribution is therefore small.

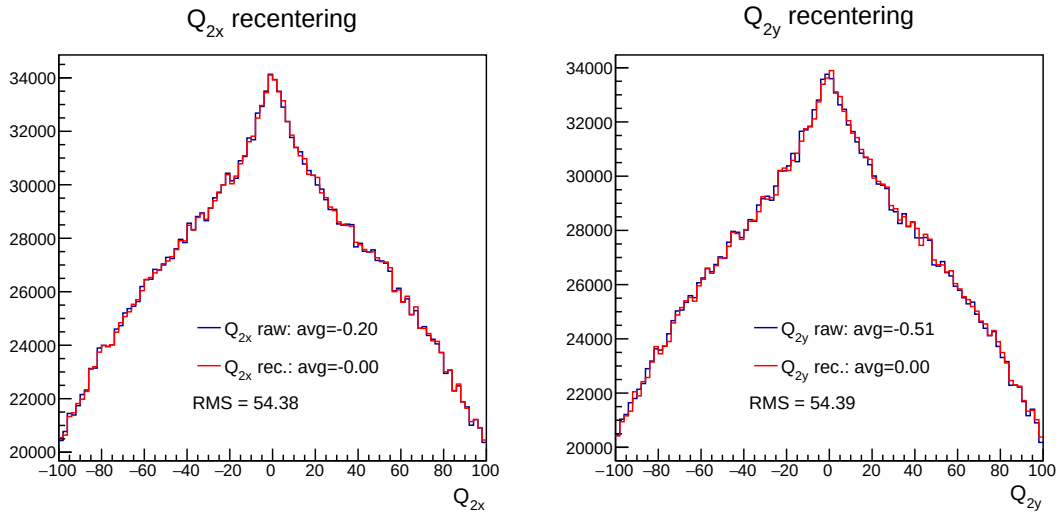

Figure 5.5. Recentering of the Q_2 vector components. The raw distribution is shown in blue and the recentered distribution is shown in red.

Because of its randomness, we expect the recentered event plane angle to have an isotropic distribution. However, asymmetries in the detector acceptance and other laboratory effects could affect the distribution and the subsequent analysis. It is therefore necessary to flatten the event plane distribution via some transformation.

We let ψ' represent the raw event plane angle and $\rho(\psi')$ represent the distribution of the event plane angle over all Υ events. In order to flatten the distribution, we apply a transformation $\psi' \to \psi$ such that $\rho(\psi) = \text{constant}$. If the range of the event plane angle is restricted to $[-\pi/2, \pi/2]$, then we find specifically that $\rho(\psi) = N/\pi$, where N is the total number of Υ events. Assuming a continuous transformation, we observe that

$$\rho(\psi)d\psi = \rho(\psi')d\psi'. \tag{5.20}$$

We then integrate both sides of the equation to find

$$\psi = \frac{\pi}{N}\int_{-\pi/2}^{\pi/2} \rho(\psi')d\psi'. \tag{5.21}$$

We then expand $\rho(\psi')$ as a Fourier series,

$$\rho(\psi') = \sum_{j=0}^{j_{\max}} [A_j \cos(2j\psi') + B_j \sin(2j\psi')], \tag{5.22}$$

where only even frequencies have been kept because the period is π rather than 2π. The following formulae can be derived from the Fourier series:

$$A_0 = \frac{N}{\pi}, \tag{5.23}$$

$$B_0 = 0, \tag{5.24}$$

$$A_j = \frac{2N}{\pi}\langle\cos(2j\psi')\rangle \text{ for } j > 0, \tag{5.25}$$

$$B_j = \frac{2N}{\pi}\langle\sin(2j\psi')\rangle \text{ for } j > 0. \tag{5.26}$$

The $\langle\rangle$ represents the averages, which are taken over all events of interest. Inserting the above formulae into Eq. 5.21 and carrying out the integral, setting arbitrary constants to zero, we find

$$\psi = \psi' - \sum_{j=1}^{j_{\max}} \frac{1}{j}\Big(\langle\sin(2j\psi')\rangle\cos(2j\psi') - \langle\cos(2j\psi')\rangle\sin(2j\psi')\Big) \tag{5.27}$$

where ψ' is the event plane angle before flattening and ψ is the event plane angle after flattening. The quantity $j_{\max}$ is the order of the flattening correction. In theory it can be infinite, but in practice a value of 21 was more than sufficient. In this analysis, the flattening procedure was carried out in three centrality bins: 10–30%, 30–50%, and 50–90%. The flattened nominal event plane distribution as a function of ψ is shown in Fig. 5.6 (upper left) with the raw and recentered distributions shown for comparison.

5.5 Event Plane Resolution Correction

As discussed in the previous section, the event plane angle is estimated from an analysis of the momenta of particles emanating from the collision. But because the number of particles in the detector is finite and the momentum resolution of the CMS detector is imperfect (see Sec. 3.4), the resolution of the event plane angle estimate is limited. This lack of resolution in the event plane angle tends to smear the distribution of Υ yields as a function of $\Delta\phi$, decreasing the size of the oscillations (see Figs. 5.18 through 5.20), and leading to an underestimate of the v_2. To correct for this lack of resolution, the raw extracted v_2 must be

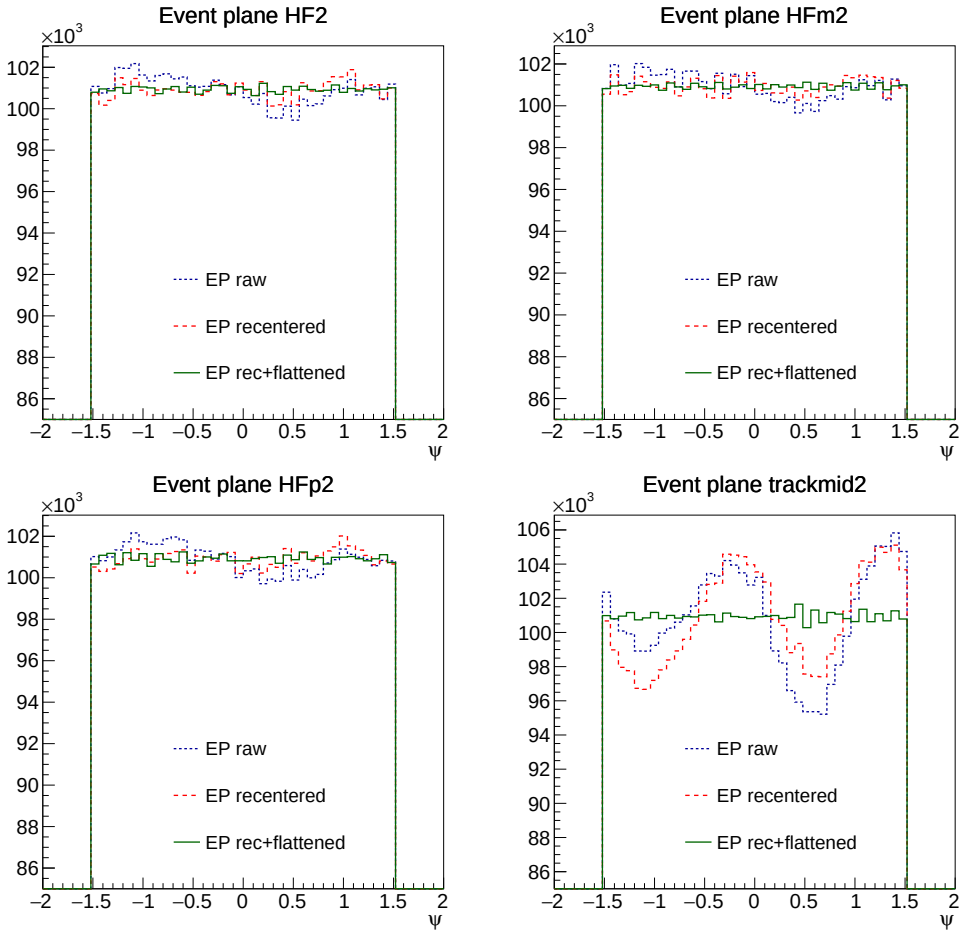

Figure 5.6. Recentering and flattening of the event plane (ψ) distributions as estimated from the forward and backward HF (upper left), the backward HF (upper right), the forward HF (lower left), and the tracker at midrapidity (lower right). The raw distribution is shown as a blue dotted curve, the recentered distribution is shown as a red dashed curve, and the flattened distribution is shown as a green solid curve. The flattened distribution is observed to be uniform in all cases. The event planes plotted are from candidate Υ events in the centrality range 10–90%.

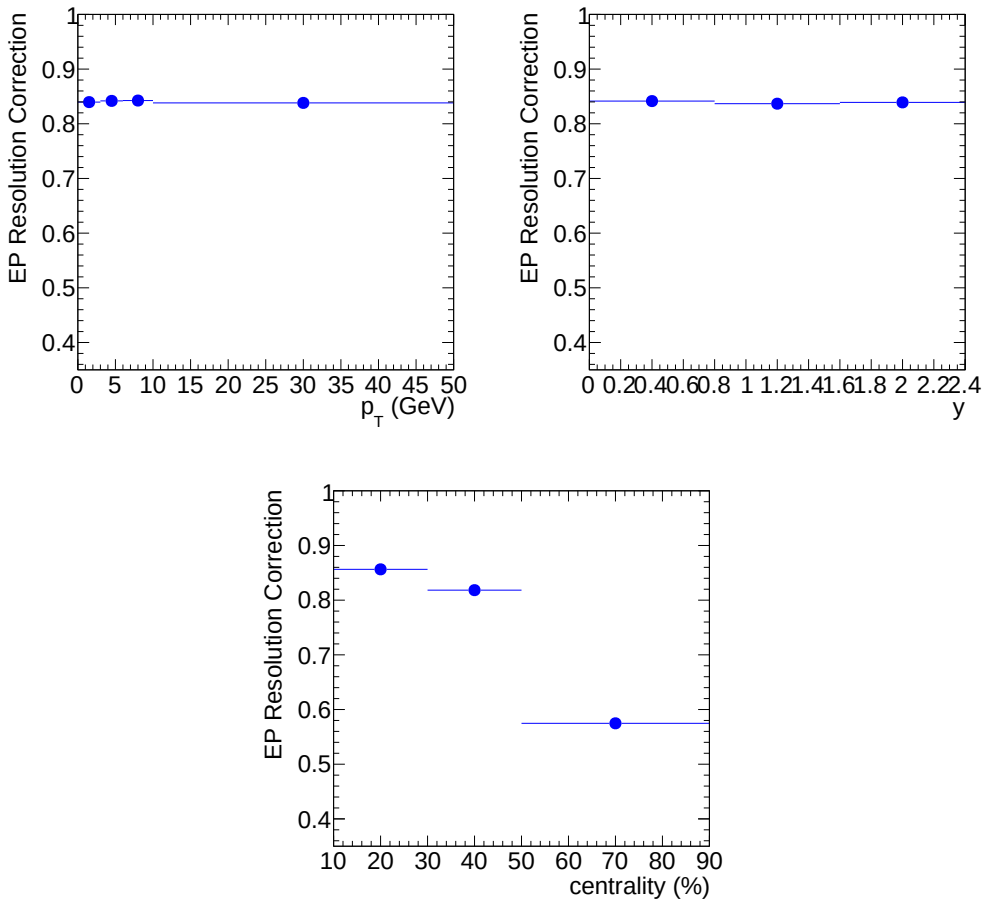

Figure 5.7. The resolution correction factor as a function of p_T (upper left), rapidity (upper right), and centrality (lower middle).

divided by an event plane resolution correction factor. This is implemented using the three-sub-event technique [75] which was also used in the v_2 analysis of the J/ψ [47].

The three-sub-event technique relies on three independent pseudorapidity regions to construct three independent event plane estimates which are then used to estimate the event plane resolution. The three regions used are $-5 < \eta < -3$, corresponding to the HF at negative pseudorapidity, $3 < \eta < 5$, corresponding to the HF at positive pseudorapidity, and $-0.8 < \eta < 0.8$, corresponding to the tracker in the barrel region. The three regions are kept at least 2 units of rapidity apart from one another to minimize correlations. Each of the event planes is also

recentered and flattened through the process described in Sec. 5.4. Plots of each of the event plane distributions estimated from the different regions are shown in Fig. 5.6. The raw distributions are observed to be nonisotropic, especially the distribution of event planes estimated from the tracker, indicating asymmetries in the detector. However, all the distributions look uniform after the recentering and flattening procedure. The event-by-event values are then used to obtain the resolution correction.

For each dimuon, whichever event plane angle was estimated from the region of the detector farthest away from the dimuon in pseudorapidity is given the label ψ^A, the event plane angle estimated from the tracker in the barrel region is given the label ψ^C, and the the other event plane angle is given the label ψ^B. The resolution correction R_A is then given by:

$$R_A = \sqrt{\frac{\langle \cos(2(\psi^A - \psi^B)) \rangle \langle \cos(2(\psi^A - \psi^C)) \rangle}{\langle \cos(2(\psi^B - \psi^C)) \rangle}}, \tag{5.28}$$

where the $\langle\rangle$ indicates an average taken over all events used in the analysis.

The resolution correction factor estimated in each analysis bin is plotted in Fig. 5.7. The raw v_2 values extracted from the fits to the $\Delta\phi$ distribution are divided by these resolution correction factors, bin by bin, to obtain the corrected v_2 results.

5.6 Fitting the v_2

With our estimate of the event plane, we proceed to estimate the value of the azimuthal angle of the Υ relative to the event plane, $\Delta\phi = \phi - \psi$, for each Υ candidate in the data. We first fit the dimuon invariant mass spectrum in bins of p_T, η, and centrality. Each of these bins is then split into four bins of $|\Delta\phi|$ which are fit simultaneously to extract the azimuthal distribution of yields in each kinematic region. Because of the azimuthal symmetry of the CMS detector, all of the parameters of the fit except for the Υ yields and the number of background events are constrained to be the same across the four bins. These fits can be seen in Figs. 5.8 through 5.17.

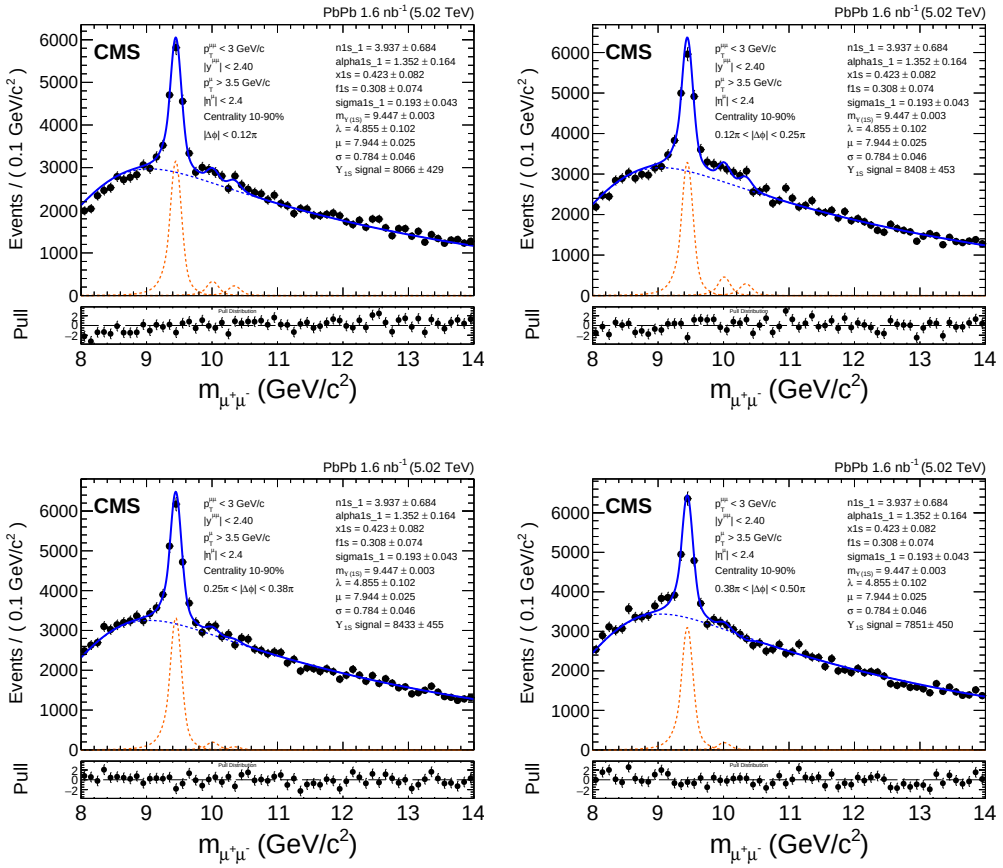

Figure 5.8. Nominal fits to data with $0 < p_T^{\mu\mu} < 3$ GeV/c and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

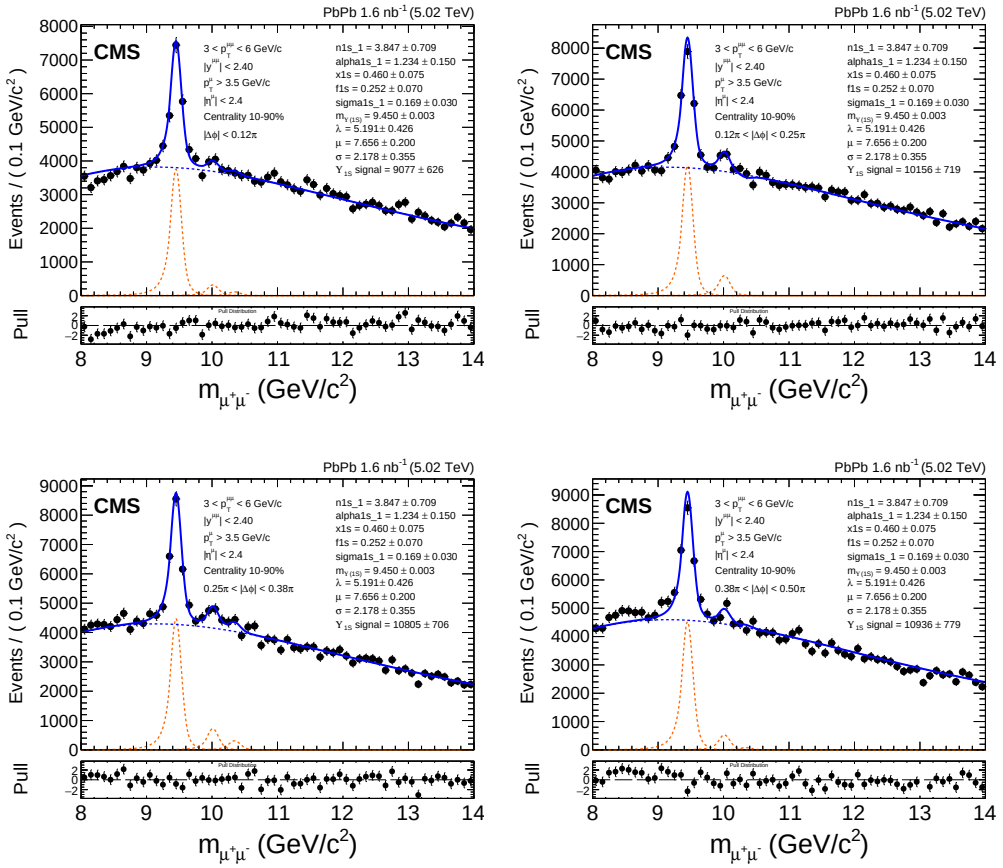

Figure 5.9. Nominal fits to data with $3 < p_T^{\mu\mu} < 6$ GeV/c and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

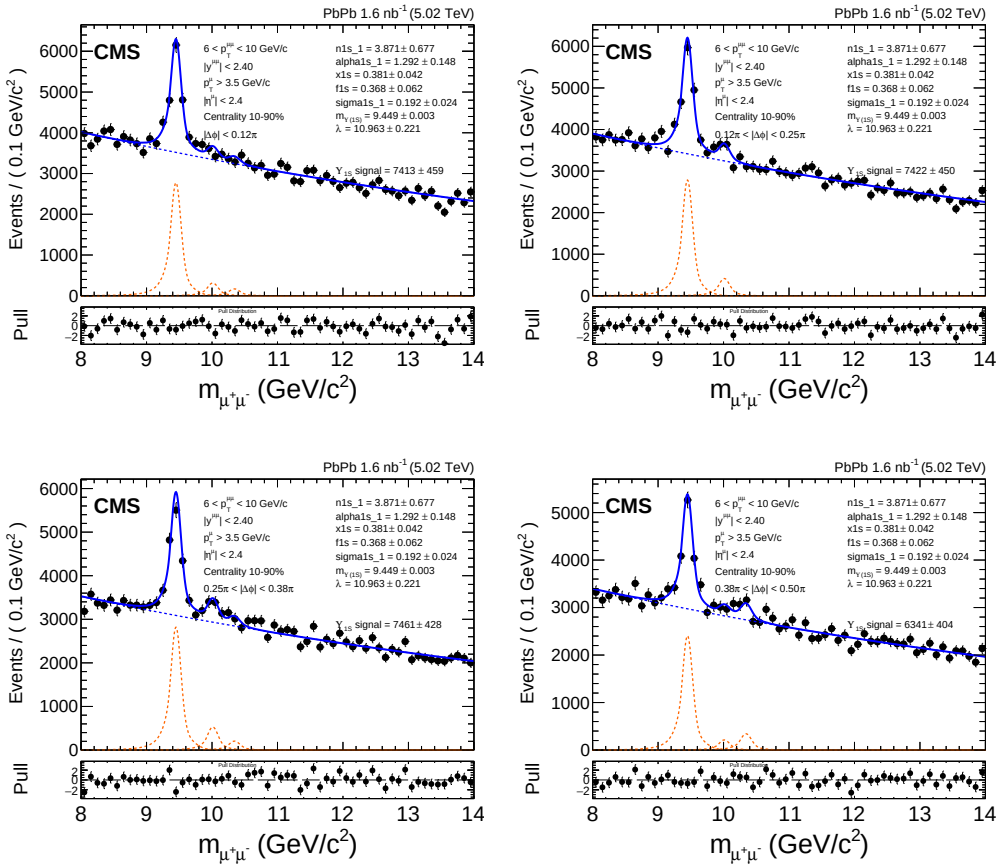

Figure 5.10. Nominal fits to data with $6 < p_T^{\mu\mu} < 10$ GeV/c and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

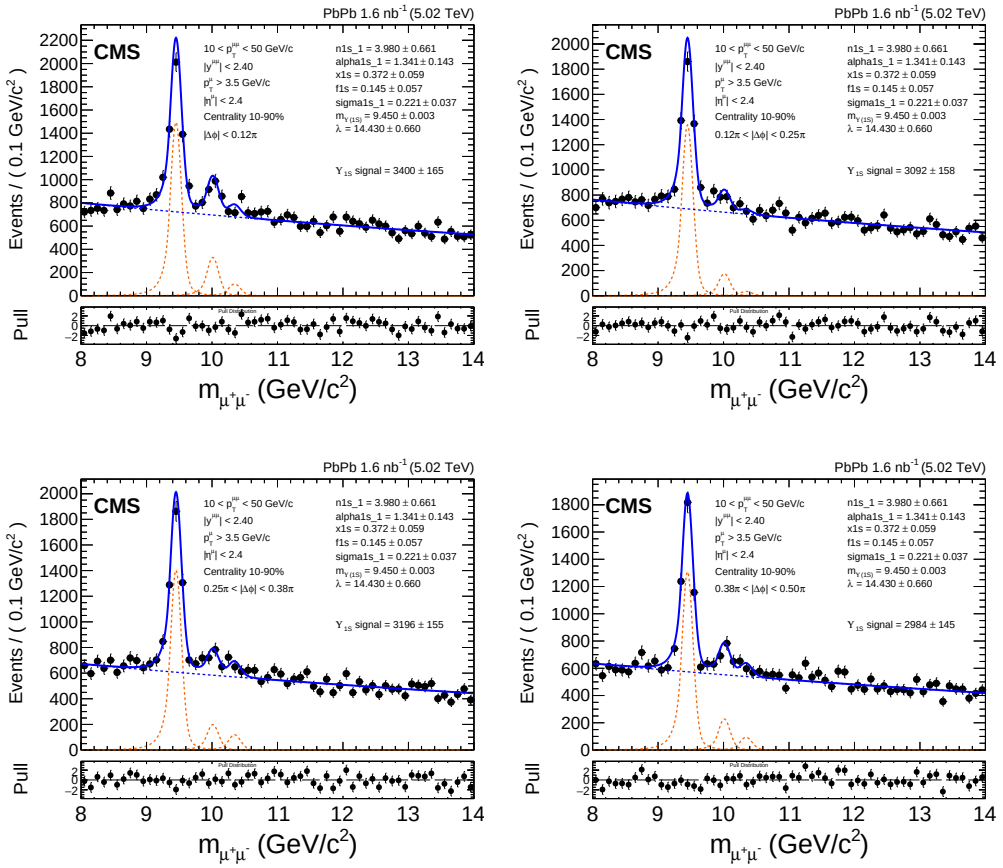

Figure 5.11. Nominal fits to data with $10 < p_T^{\mu\mu} < 50$ GeV/c and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

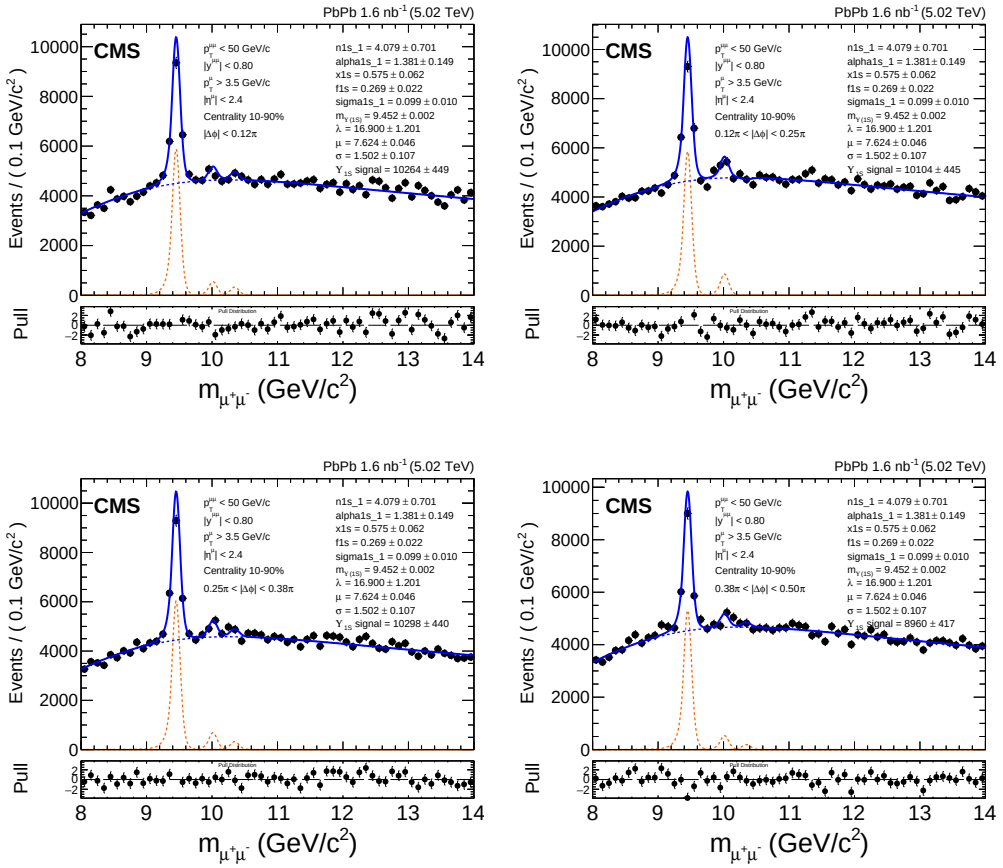

Figure 5.12. Nominal fits to data with $0.0 < y^{\mu\mu} < 0.8$ and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

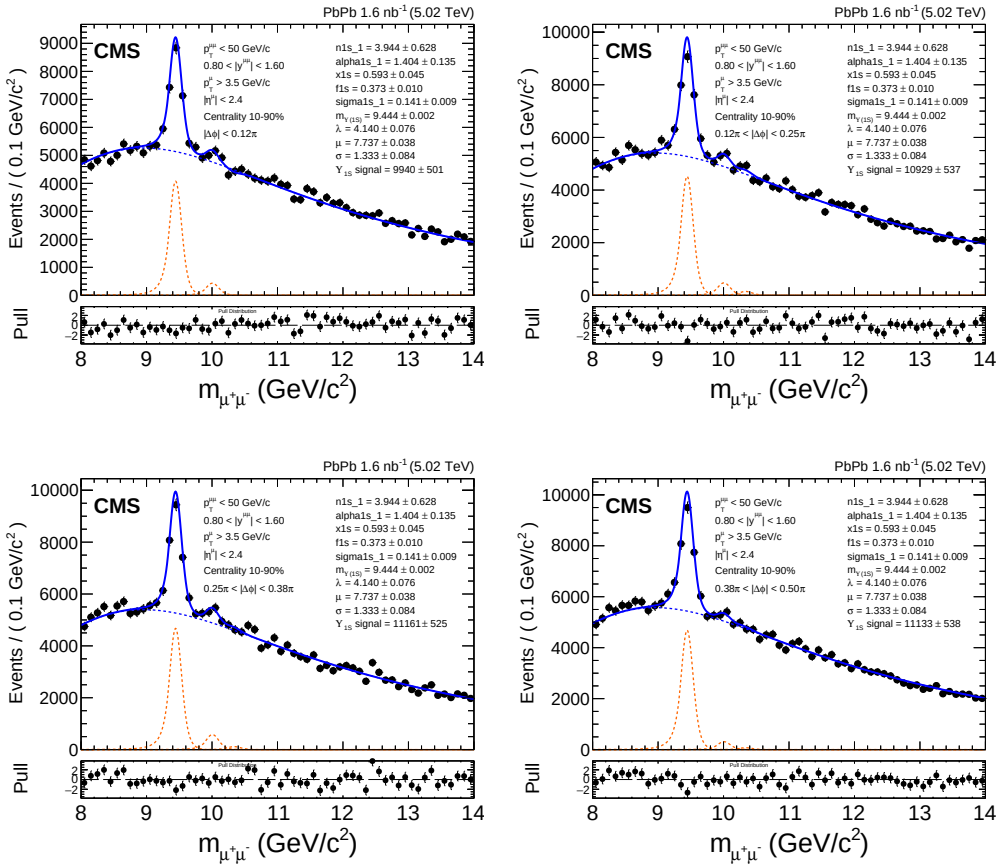

Figure 5.13. Nominal fits to data with $0.8 < y^{\mu\mu} < 1.6$ and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

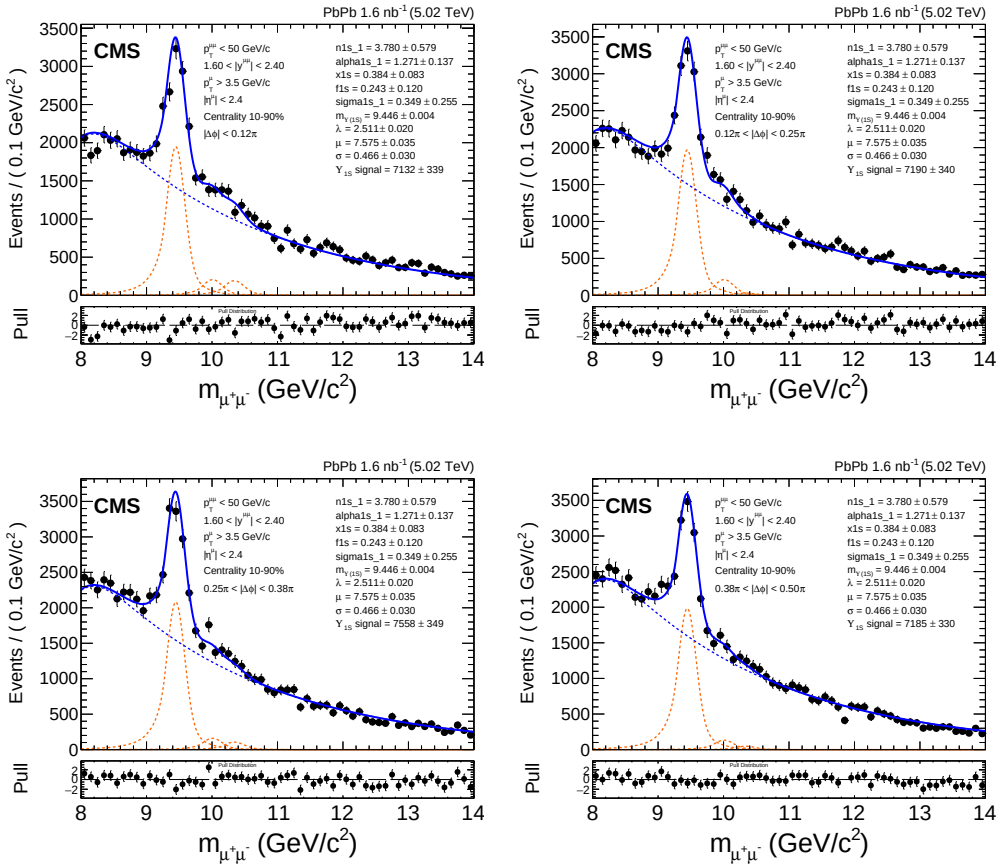

Figure 5.14. Nominal fits to data with $1.6 < y^{\mu\mu} < 2.4$ and centrality 10–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

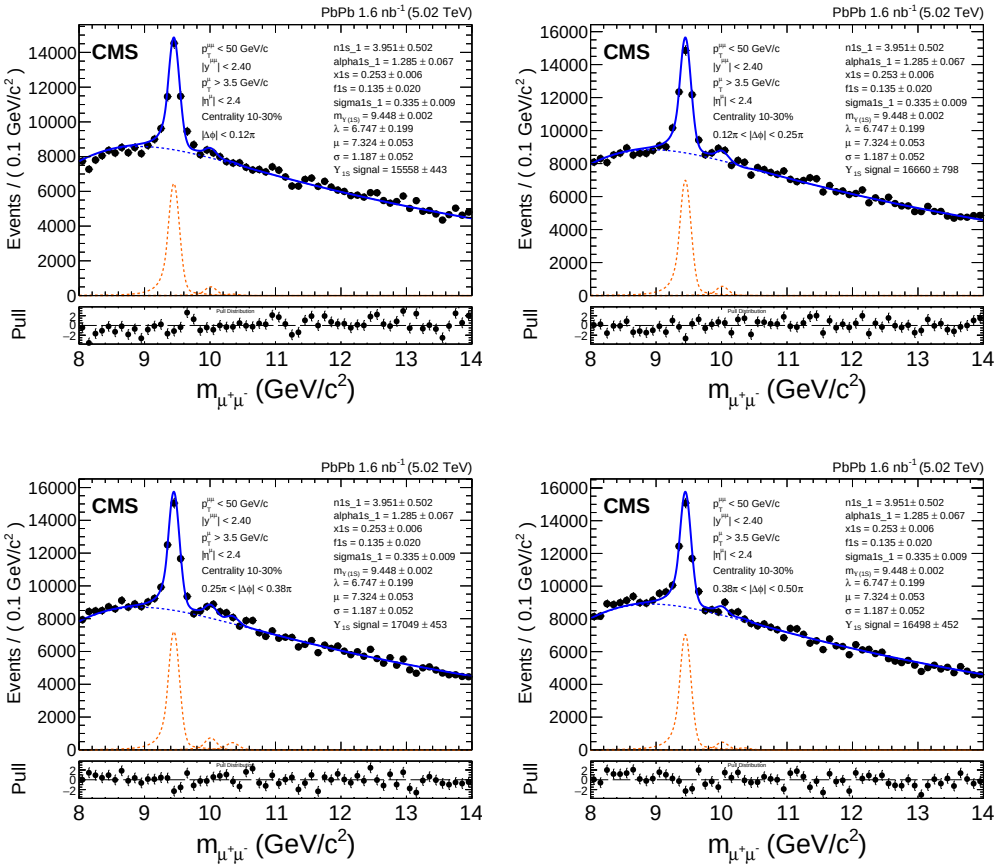

Figure 5.15. Nominal fits to data with centrality 10–30% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

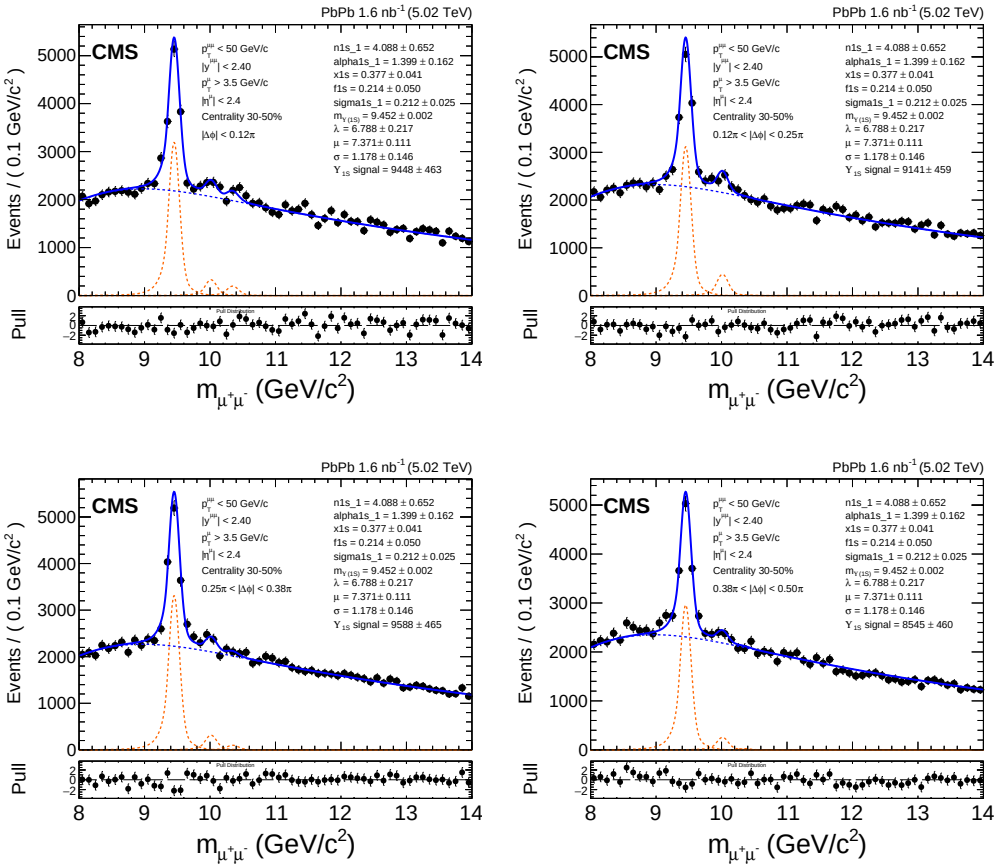

Figure 5.16. Nominal fits to data with centrality 30–50% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

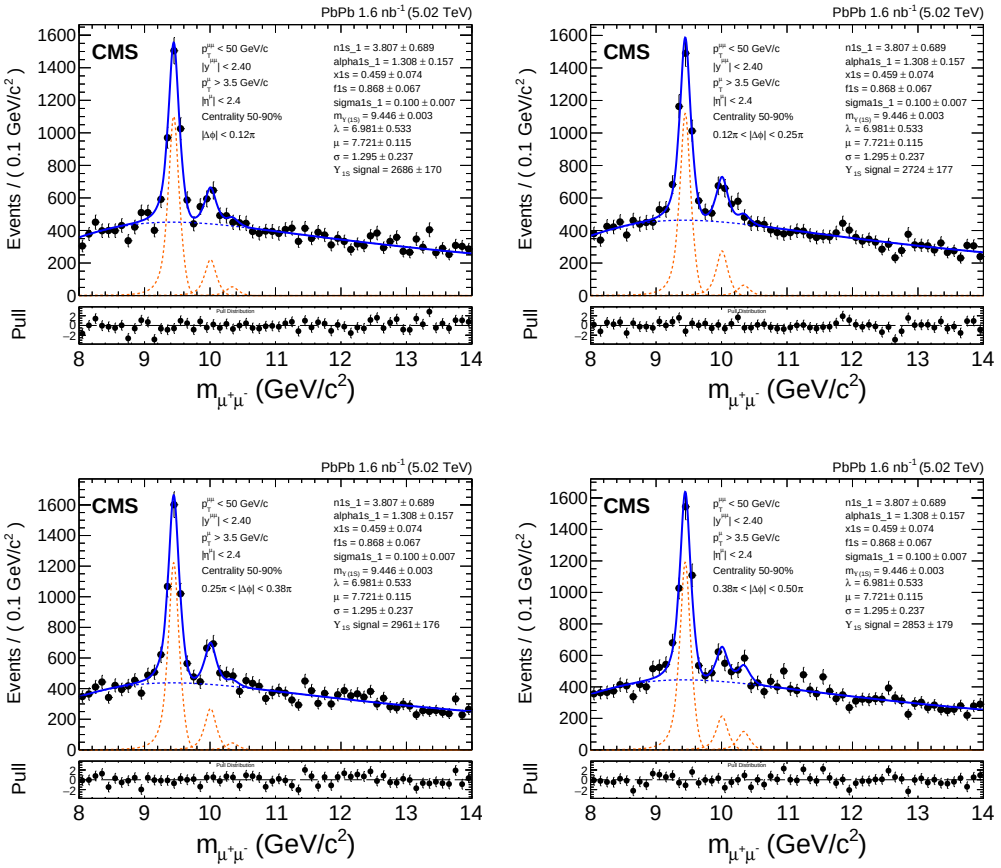

Figure 5.17. Nominal fits to data with centrality 50–90% in the four $|\Delta\phi|$ bins: $[0, \pi/8)$ (upper left), $[\pi/8, \pi/4)$ (upper right), $[\pi/4, 3\pi/8)$ (lower left), and $[3\pi/8, \pi/2)$ (lower right). The data are plotted as black circles, and are fit with the blue line. The signal and background components of the fit function are plotted as red and blue dashed lines, respectively. The pull distribution in the bottom panel shows the difference between each data point and the fit function.

The Υ yields are then extracted from the four fits in each kinematic bin or centrality bin and plotted together as a function of $\Delta\phi$. The total Υ yields extracted in each kinematic bin in the full $\Delta\phi$ range are tabulated in Table 5.2.

Table 5.2. The yields of Υ(nS) in the 2018 Pb + Pb dataset in all of the analysis bins in the full $\Delta\phi$ range. The yields were obtained from fits to data after the application of acceptance and efficiency corrections.

Bin	Υ(1S) yield	Υ(2S) yield	Υ(3S) yield
$0 \leq p_T < 3$ GeV	32758.2 ± 601.2	2863.54 ± 298.12	1336.28 ± 286.75
$3 \leq p_T < 6$ GeV	40974.1 ± 1283.0	5741.13 ± 540.91	1332.92 ± 417.94
$6 \leq p_T < 10$ GeV	28637.2 ± 557.8	3466.18 ± 291.08	2346.88 ± 287.79
$10 \leq p_T < 50$ GeV	12672 ± 292	2117 ± 130	763.98 ± 117.24
$0 \leq y < 0.8$	39626 ± 1447	5061.92 ± 445.55	1750.35 ± 335.31
$0.8 \leq y < 1.6$	43163.3 ± 665.1	4341.29 ± 359.60	229.69 ± 104.86
$1.6 \leq y < 2.4$	29064 ± 468	3022.54 ± 274.86	1823.49 ± 209.21
Centrality 10-30%	65765.7 ± 1240.2	6509.02 ± 500.99	2007.29 ± 460.44
Centrality 30-50%	36722 ± 614	4426.82 ± 274.29	2009.46 ± 242.91
Centrality 50-90%	11224.1 ± 279.1	2387.93 ± 139.27	776.40 ± 109.92
Centrality 10-90%	106135 ± 1216	12147.7 ± 538.1	3480.36 ± 491.23

The distribution of yields as a function of $\Delta\phi$ is then normalized and fit using the function

$$\frac{1}{N_{\text{bins}}}(1 + 2v_2 \cos(2|\Delta\phi|)), \tag{5.29}$$

where N_{bins} is the number of bins of $|\Delta\phi|$, which was 4 in this analysis. The fitted v_2 parameter is the raw v_2 value for the particular analysis bin. The fits to the nominally extracted Υ(1S) normalized yields as a function of $|\Delta\phi|$ are shown in Figs. 5.18 through 5.20. The corresponding fits to the Υ(2S) normalized yields are shown in Figs. 5.21 through 5.23.

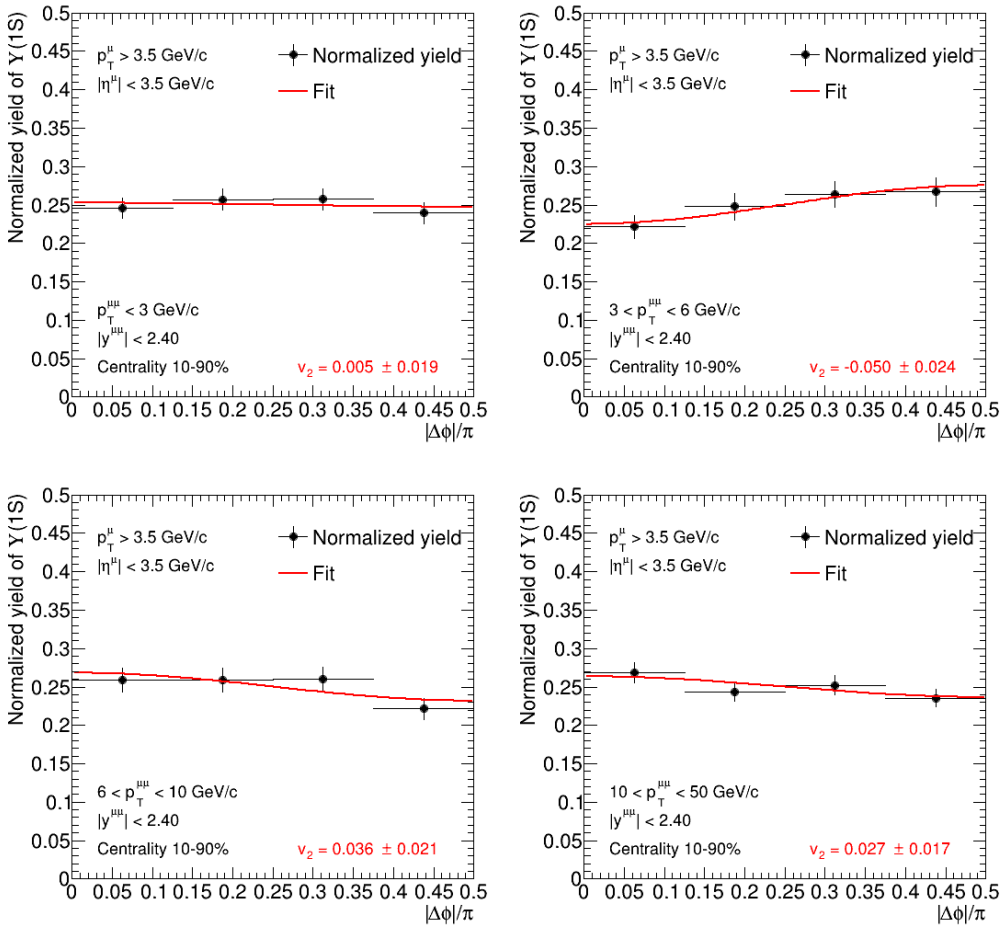

Figure 5.18. The nominally extracted Υ(1S) yields as a function of $\Delta\phi$ in the transverse momentum bins $0 < p_T < 3$ GeV/c (upper left), $3 < p_T < 6$ GeV/c (upper right), $6 < p_T < 10$ GeV/c (lower left), and $10 < p_T < 50$ GeV/c (lower right), fit with the function in Eq. 5.29 to extract v_2.

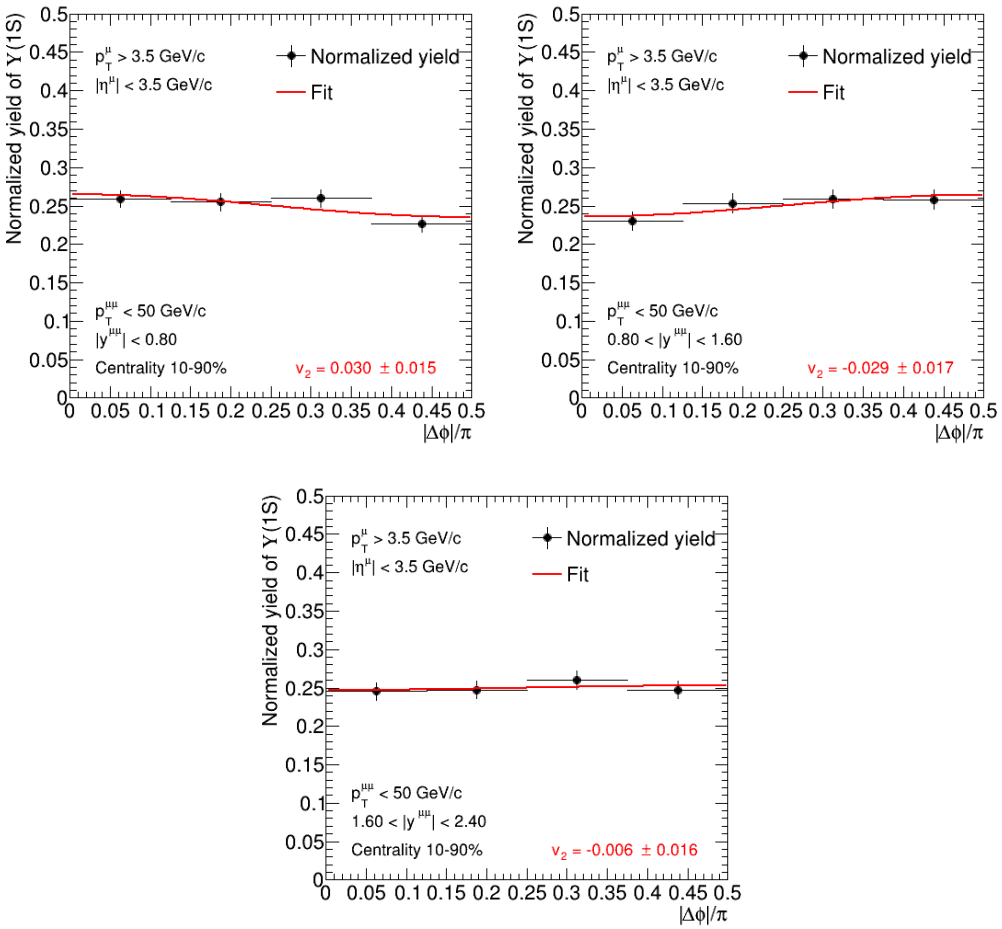

Figure 5.19. The nominally extracted Υ(1S) yields as a function of $\Delta\phi$ in the rapidity bins $0.0 < y < 0.8$ (upper left), $0.8 < y < 1.6$ (upper right), and $1.6 < y < 2.4$ (lower middle), fit with the function in Eq. 5.29 to extract v_2.

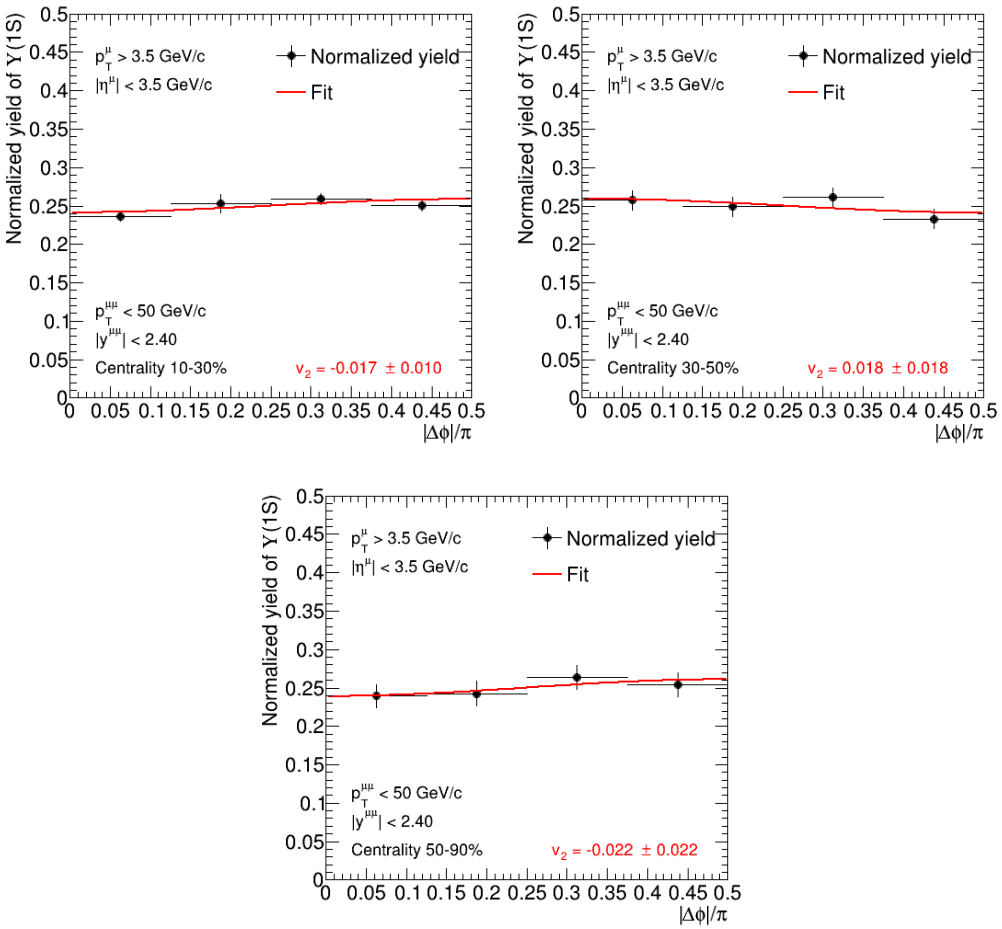

Figure 5.20. The nominally extracted Υ(1S) yields as a function of $\Delta\phi$ in the centrality regions 10–30% (upper left), 30–50% (upper right), and 50–90% (lower middle), fit with the function in Eq. 5.29 to extract v_2.

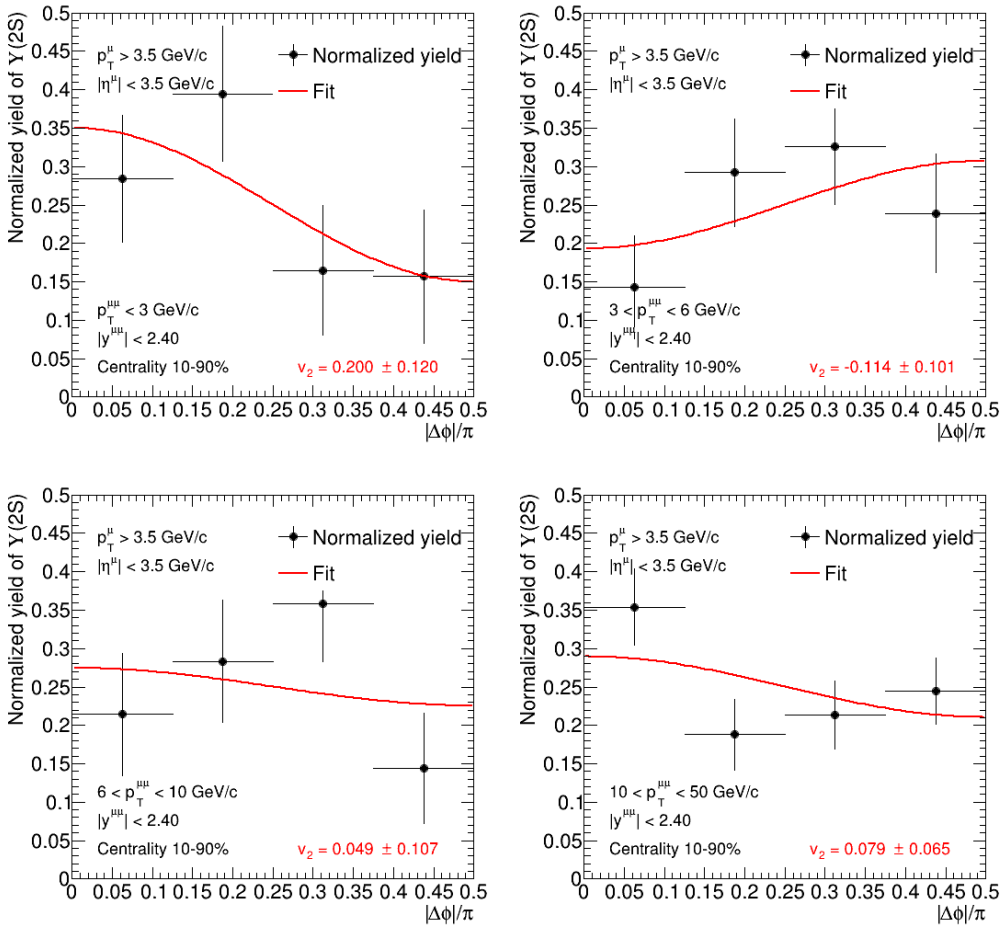

Figure 5.21. The nominally extracted Υ(2S) yields as a function of $\Delta\phi$ in the transverse momentum bins $0 < p_T < 3$ GeV/c (upper left), $3 < p_T < 6$ GeV/c (upper right), $6 < p_T < 10$ GeV/c (lower left), and $10 < p_T < 50$ GeV/c (lower right), fit with the function in Eq. 5.29 to extract v_2.

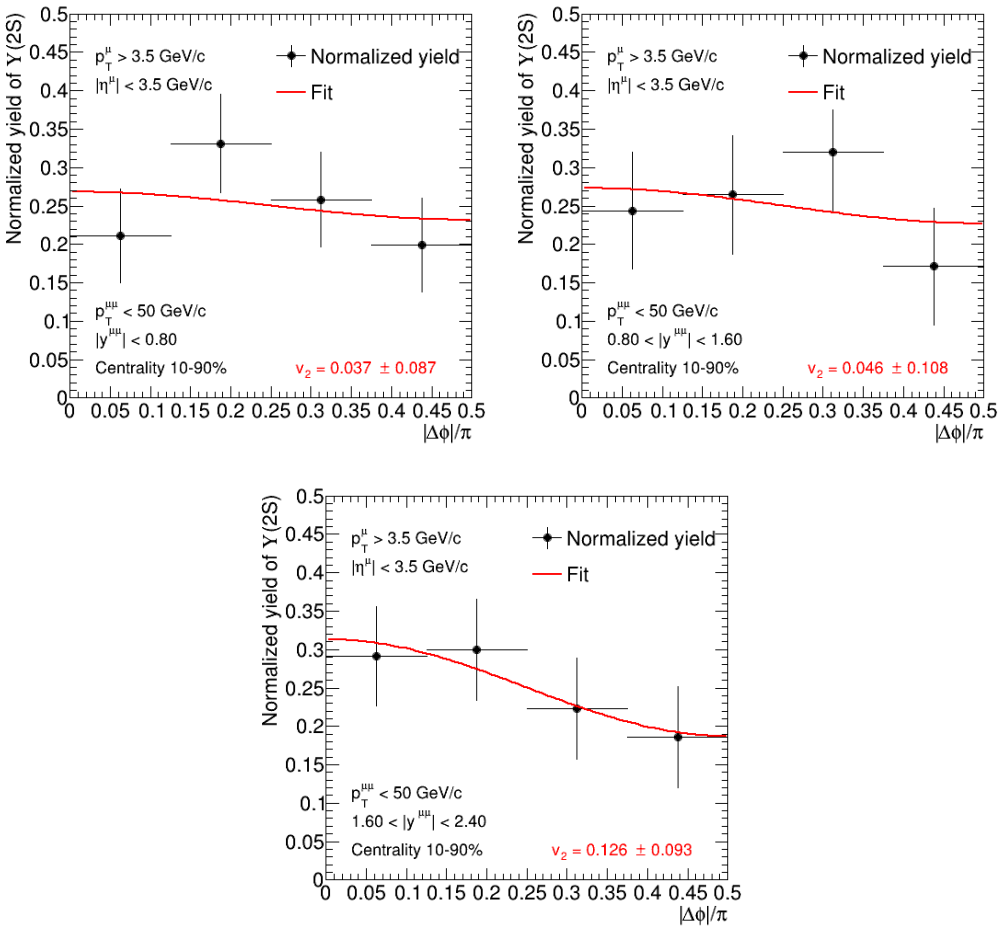

Figure 5.22. The nominally extracted Υ(2S) yields as a function of $\Delta\phi$ in the rapidity bins $0.0 < y < 0.8$ (upper left), $0.8 < y < 1.6$ (upper right), and $1.6 < y < 2.4$ (lower middle), fit with the function in Eq. 5.29 to extract v_2.

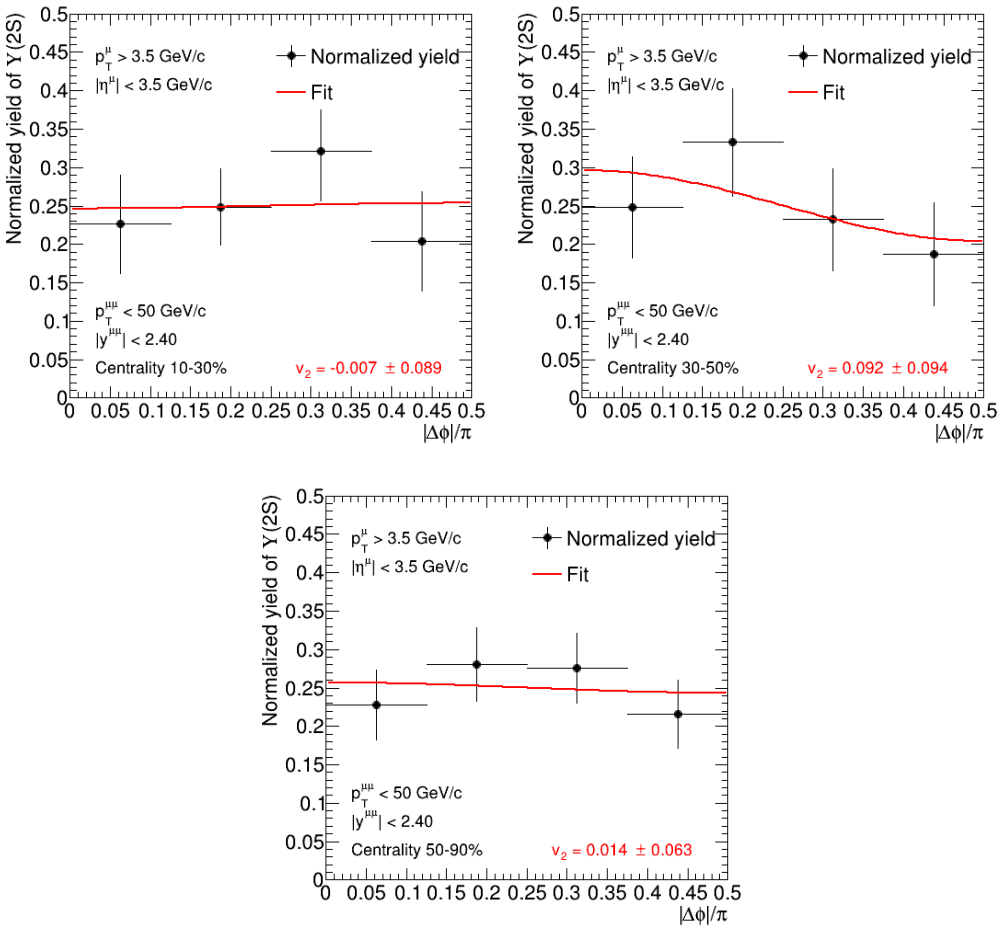

Figure 5.23. The nominally extracted Υ(2S) yields as a function of $\Delta\phi$ in the centrality regions 10–30% (upper left), 30–50% (upper right), and 50–90% (lower middle), fit with the function in Eq. 5.29 to extract v_2.

5.7 Feasibility Test

To test the feasibility of a v_2 analysis of Υ's before the 2018 dataset became available, a simplified version of the analysis was carried out using the Pb + Pb data from 2015 (see Table 2.1). This feasibility test was done to estimate how large the uncertainties on the $\Upsilon(1S)$ v_2 signal would be. It was therefore not necessary to precisely estimate the reaction plane angle since only the statistical uncertainty was needed. Systematic uncertainty calculations were also omitted, and instead the average systematic uncerainties from a previous analysis of the v_2 of the J/ψ [47] was used. The results are shown in Fig. 5.24. The only quantities of interest in these figures are the statistical uncertainties. We expected about a factor of 4 increase in data with the new dataset collected in November 2018, which means these uncertainties will be reduced by 50% once the larger dataset is used. Also, if necessary, we could use larger bins to further reduce the uncertainties. The two most important results we are looking for are the presence of a nonzero v_2, and a larger v_2 for the $\Upsilon(2S)$ compared to the $\Upsilon(1S)$. The desired results do not require fine binning, and thus the analysis seemed to be feasible.

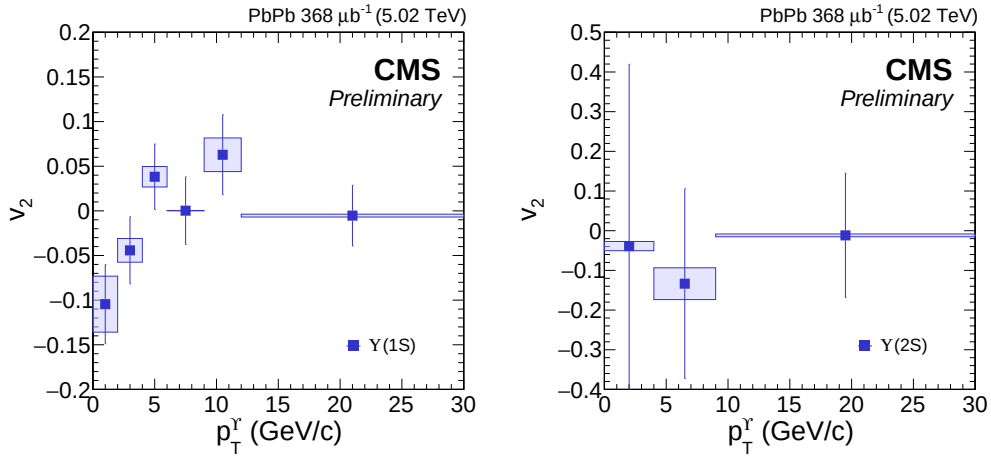

Figure 5.24. Results of a simplified v_2 analysis using 2015 Pb + Pb data as a feasibility test. The only quantities of interest in the plots are the statistical uncertainties.

Another feasibility test was a simple comparison with the J/ψ analysis which had already been done. The analysis extracted a total prompt J/ψ yield of ap-

proximately ∼10,000. In a certain bin, they extracted a yield of 2086 prompt J/ψs, and determined a v_2 of 0.034 ± 0.021 (stat) ± 0.007 (syst) [47]. The 2015 Pb + Pb data yielded 5,868 Υ(1S) [27] without applying acceptance and efficiency correction factors. With a factor of 4 increase, our statistics would be comparable to those of the J/ψ analysis.

5.8 The Scalar-Product Method

In parallel with this analysis, another group of researchers investigated the Υ v_2 using the scalar-product (SP) method [76]. This method makes use of the Q vectors described in Sec. 5.4 [77, 78]. The Q vectors are converted to complex numbers, q, as follows:

$$q_n = \frac{\sum_{i=1}^{M} w_i e^{in\phi_i}}{\sum_{i=1}^{M} w_i} = \frac{Q_{nx} + iQ_{ny}}{W}, \tag{5.30}$$

where the weights, w_i, are the energies of the particles, as in Sec. 5.4, and W is the sum of the weights in a given event. The scalar-product coefficients are then given by

$$v_n\,\{\mathrm{SP}\} \equiv \frac{\langle q_n q_{nA}^* \rangle}{\sqrt{\frac{\langle q_{nA} q_{nB}^* \rangle \langle q_{nA} q_{nC}^* \rangle}{\langle q_{nB} q_{nC}^* \rangle}}}, \tag{5.31}$$

where A and B refer to event planes derived from the forward and backward calorimeter regions, and C refers to event planes derived from the tracker in midrapidity. The averages $\langle q_n q_{nA}^* \rangle$ in Eq. 5.31 are weighted averages taken in this manner:

$$\langle q_n q_{nA}^* \rangle = \mathrm{Re}\left[\frac{\sum_{i=1}^{N_{evt}} W_{Ai} W_{Bi} q_{nAi} q_{nBi}^*}{\sum_{i=1}^{N_{evt}} W_{Ai} W_{Bi}}\right]. \tag{5.32}$$

The W_{Ai} and W_{Bi} are the sums of of weights used for event planes A and B, respectively, for each event i. The averages $\langle q_{nA} q_{nB}^* \rangle$, $\langle q_{nA} q_{nC}^* \rangle$ and $\langle q_{nB} q_{nC}^* \rangle$ are taken over all events, while the average $\langle q_n q_{nA}^* \rangle$ is over all Υ(1S) candidates in all events. For a measurement of elliptic flow, v_2, we use $n = 2$.

The event plane method is known to give a v_2 somewhere between the mean, $\langle v_2 \rangle$, and RMS, $\sqrt{\langle v_2^2 \rangle}$, of the true v_2, while the scalar-product method removes the ambiguity and gives the RMS [78]. Thus the two methods are not guaranteed

to give the same result, but if there is good agreement on the measured v_2 between the two analyses, then we will have greater confidence in the accuracy of the measurement.

Chapter 6

Systematic Uncertainties

6.1 Choice of Signal Function

In order to estimate the systematic uncertainty due to the choice of signal function, we tested how much the resulting v_2 values changed if we used a reasonable alternative function, which in this case was a CB plus a Gaussian. In order to isolate the systematic uncertainty from any statistical uncertainties, we performed 100 "pseudoexperiments" in each analysis bin and compared the results obtained using the two functions. Each pseudoexperiment is a repetition of the analysis described in Ch. 5 using a set of generated data dubbed "pseudodata." The procedure is as follows:

1. We generate pseudodata using the nominal fit to the data. The generated Υ yields, and thus the v_2, are fixed to the nominally extracted values from real data.

2. We fit the pseudodata with both the nominal fit function and the alternative function in each $|\Delta\phi|$ bin.

3. We fit the extracted yields as a function of $|\Delta\phi|$ to extract the nominal and alternative v_2 values. The difference between the alternative and nominal values is recorded.

4. We repeat steps 1 through 3 one hundred times to get a distribution of differences.

5. We take the larger value of the mean or the standard deviation of the

distribution of the differences and divide it by the value of v_2 nominally extracted from real data to get the relative systematic uncertainty in that bin. In almost all cases, the standard deviation was the larger of the two values, while the mean, which represents a systematic bias, was usually negligible.

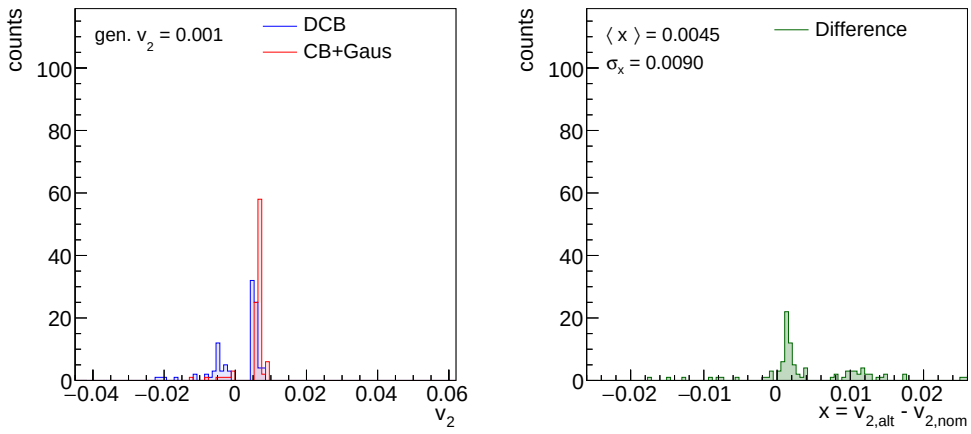

Figure 6.1. Results of pseudoexperiments in the kinematic bin $0 < p_T < 3$ GeV/c. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian. The distribution of v_2 extracted from 100 generated sets of pseudodata using the nominal fit function is shown on the left in blue. The distribution of v_2 extracted from the same 100 sets of pseudodata using the alternative fit function is also shown on the left in red. The generated v_2 is noted in the top left corner. The plot on the right shows the distribution of the differences in the extracted v_2 between the two fitting methods. The mean ($\langle x \rangle$) and standard deviation (σ_x) are noted in the figure.

The plots in Fig. 6.1 show an example of this in one bin. On the left are the nominal v_2 values (blue) and the alternative v_2 values (red) from 100 pseudoexperiments. On the right is the distribution of differences in the estimated v_2. The results of the pseudoexperiments in all the analysis bins can be found in Appendix A.1.

6.2 Constraining Parameters

The systematic uncertainty due to the choice of parameter constraints was estimated by carrying out the analysis with a different set of constraints. These alternative constraints were obtained following the same procedure outlined in section 5.1, except the first parameter fixed was n rather than α. In order to isolate

the systematic uncertainty from the statistical uncertainty, 100 pseudoexperiments were performed in the same manner as was done for the signal systematics, as described in section 6.1. For each pseudoexperiment, the nominal fit function was used to generate the pseudodata, which was then fit by both the nominal fit function and the alternative fit function. In this case, the alternative signal function had the same functional form as the nominal signal function (DCB), the only difference being the set of constraints applied to the signal parameters. The results of the pseudoexperiments in all the analysis bins can be found in Appendix A.2.

6.3 Choice of Background Function

The systematic uncertainty due to the choice of background function was estimated by fitting the data with an alternative background function. In the high p_T bins ($p_T > 6$ GeV/c), the alternative background function was a power law rather than the nominal exponential function. In the low p_T and integrated p_T bins, the alternative background function was a fourth-order Chebychev polynomial instead of the nominal error function times an exponential. In order to isolate the systematic uncertainty from the statistical uncertainty, 100 pseudoexperiments were performed in the same manner as was done for the signal systematics, as described in section 6.1. The results of the pseudoexperiments in all the analysis bins can be found in Appendix A.3.

6.4 Application of Correction Factors

6.4.1 Acceptance Corrections

The process of obtaining the acceptance corrections is described in Sec. 5.2. The systematic uncertainty due to the p_T-weighting of the acceptance corrections was obtained by repeating the analysis with a set of unweighted acceptance corrections. The relative difference in the resulting v_2 value from the nominal value in each bin was taken as the systematic uncertainty.

6.4.2 Efficiency Corrections

The process of obtaining the efficiency corrections is described in Sec. 5.3. The systematic uncertainty due to the p_T-weighting of the efficiency corrections was obtained by repeating the analysis with a set of unweighted efficiency corrections. The relative difference in the resulting v_2 value from the nominal value in each bin was taken as the systematic uncertainty.

6.5 Summary of Systematic Uncertainties

The total systematic uncertainty is obtained by summing in quadrature each of the systematic uncertainties from the choice of signal function, the choice of background function, the parameter constraints, the acceptance corrections, and the efficiency corrections. The total systematic uncertainty in each bin, along with the relative contributions from each of the sources of systematic uncertainty, can be seen in Fig. 6.2.

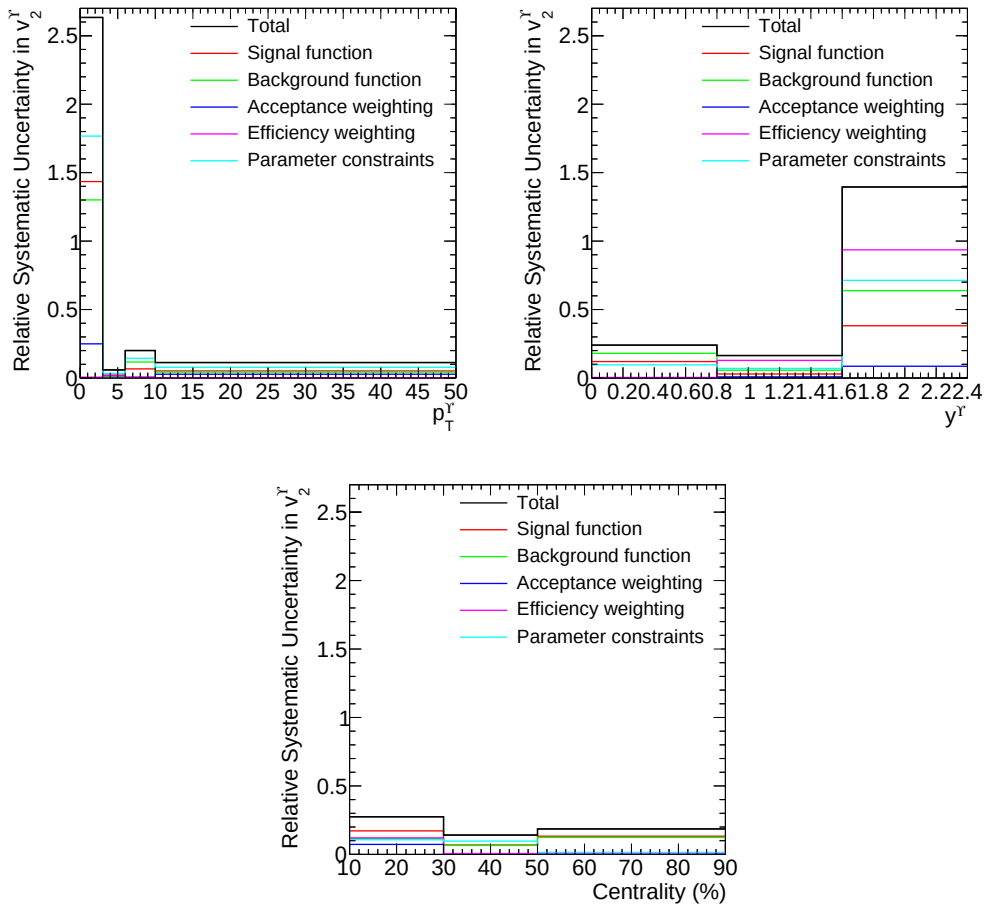

Figure 6.2. Systematic uncertainties as a function of p_T^{Υ} (upper left), y^{Υ} (upper right), and centrality (bottom middle). The total systematic uncertainty, shown in black, is the sum in quadrature of the uncertainty due to the choice of signal function (red), the choice of background function (green), the acceptance weighting (blue), the efficiency weighting (violet), and the choice of signal parameter constraints (cyan).

6.6 Monte Carlo Closure Test

In order to further solidify our confidence in the accuracy of our results and the robustness of the event-plane method in general, the analysis was carried out on a simulated dataset with a value of elliptic flow that was input by hand. If the method can extract a value of v_2 that matches the input value, then it is proven to be a robust and accurate method.

The analysis was carried out on the GEN-level MC in the same way that it was carried out on the data, with a few adjustments. Since the Υ yields are already known *a priori*, there is no need for acceptance and efficiency corrections. Also, without the simulated detector response, the resolution correction is unnecessary. Without a simulated background, the fitting of the Υ signal is also simplified, because the yields are known perfectly in each bin and can be counted directly. After counting, in order to provide comparable statistics for the v_2 extraction between MC and real data, the uncertainties on the Υ(1S) yields were scaled to mimic the average uncertainties seen in fits to real data ($\sim 5\%$).

The MC closure test was carried out using various different input v_2 values. The v_2 weight was applied to the $\Delta\phi$ distribution of the MC Υs as shown in Fig. 6.3. The plot on the left is the invariant mass distribution of the dimuons that are to be weighted, demonstrating that they are indeed Υs, and it shows the various selections that were applied to this particular bin. The plot on the right displays the $\Delta\phi$ distribution for these MC Υs before and after the v_2 weight is applied. The v_2-weighted dataset was then analyzed via the event-plane method, as described in Ch. 5, and the extracted v_2 for inputs 0.5, 0.2, 0.1, and 0.05 are shown in Figures 6.4, 6.5, 6.6, and 6.7, respectively. The event-plane method successfully extracts the v_2 in every case to within $\sim 3\%$. This is well within the statistical uncertainty for most values of v_2, especially for $v_2 = 0.05$ (see Fig. 6.7), which is most comparable to the true v_2 of the Υ(1S).

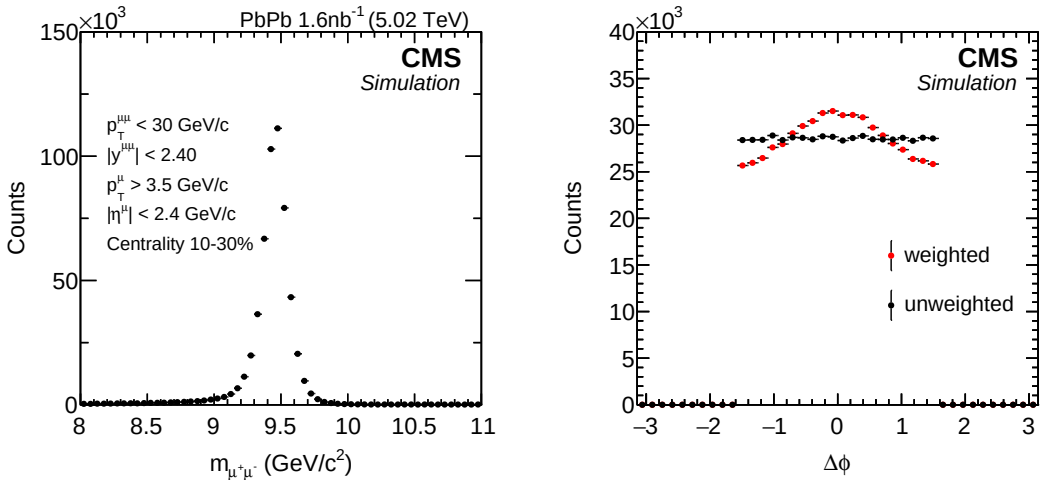

Figure 6.3. Weighting of the $\Delta\phi$ distribution in one of the analysis bins. The left plot shows the invariant mass distribution of the dimuons to be weighted, demonstrating that they are Υs. The right plot shows the $\Delta\phi$ distribution of the Υs before and after applying a v_2 weight of 0.05.

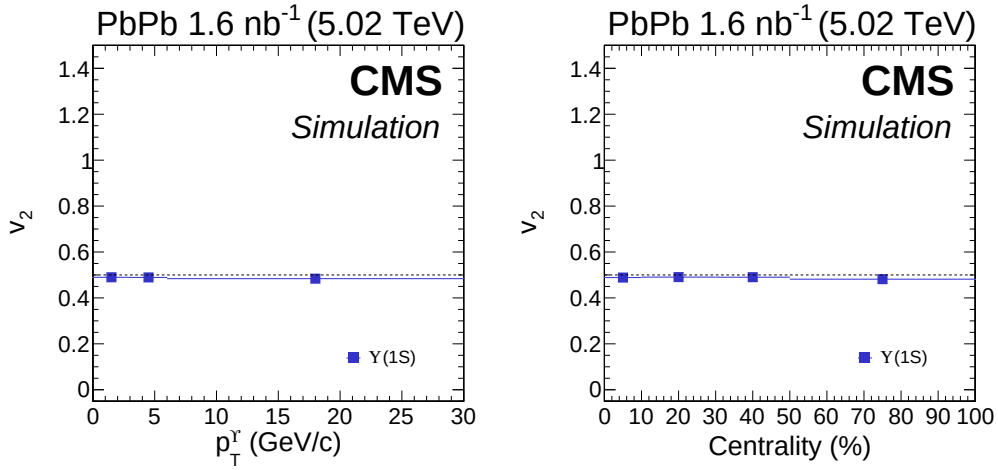

Figure 6.4. Estimated v_2 as a function of p_T and centrality from MC simulations using an input $v_2 = 0.5$.

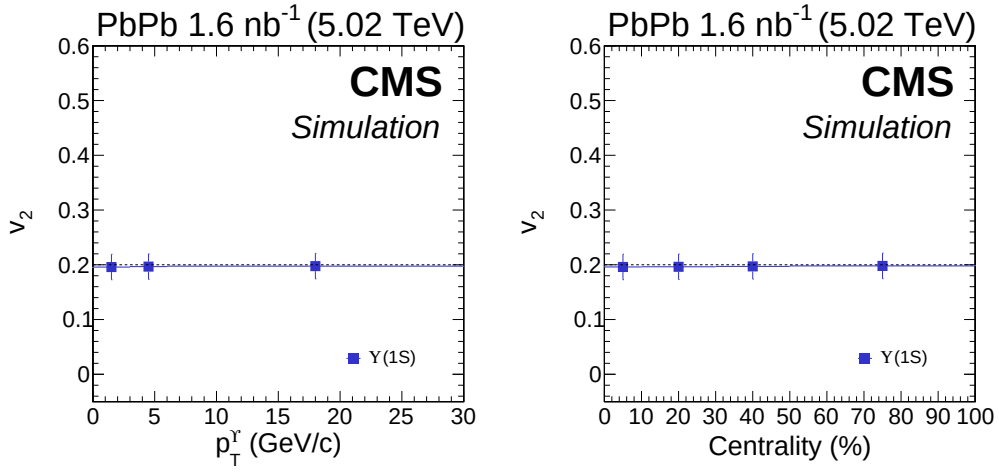

Figure 6.5. Estimated v_2 as a function of p_T and centrality from MC simulations using an input $v_2 = 0.2$.

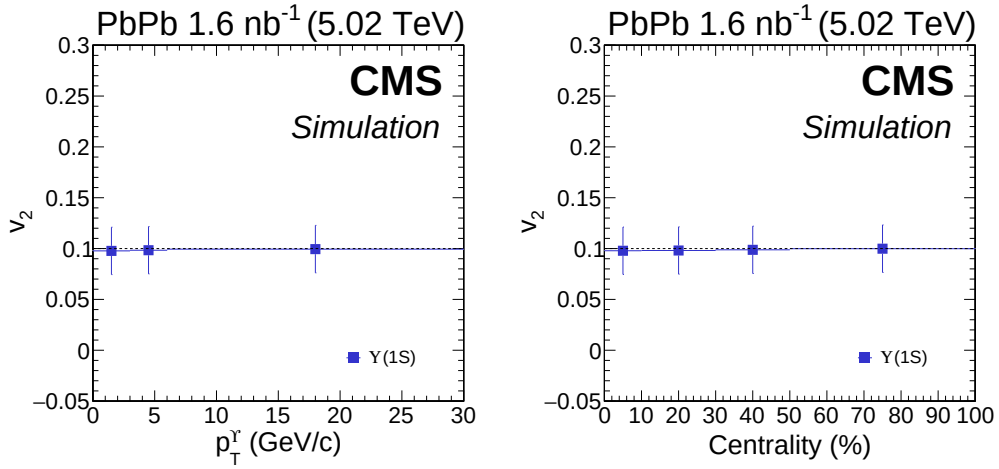

Figure 6.6. Estimated v_2 as a function of p_T and centrality from MC simulations using an input $v_2 = 0.1$.

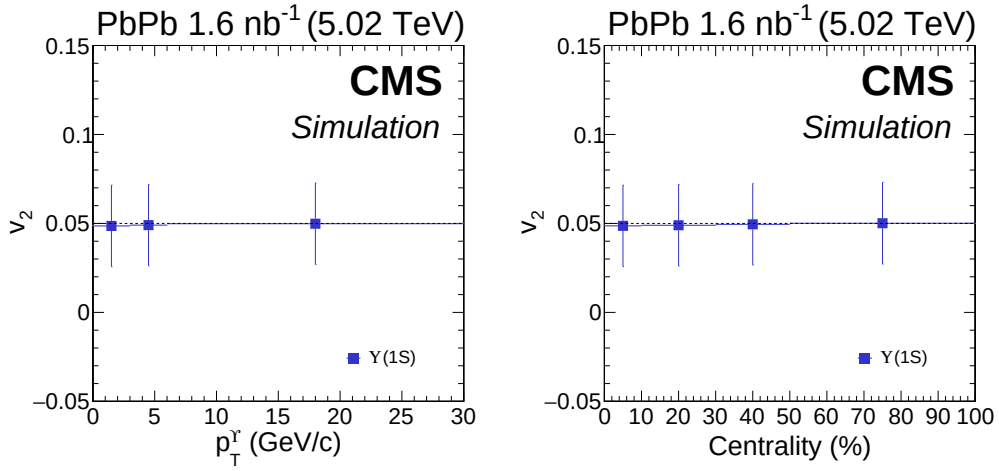

Figure 6.7. Estimated v_2 as a function of p_T and centrality from MC simulations using an input $v_2 = 0.05$.

Chapter 7

Results

The second-order Fourier coefficient, v_2, of the azimuthal distribution relative to the event plane, known as the elliptic flow, of the Υ(1S) is presented. The signal was extracted from Pb + Pb collisions at CMS with $\sqrt{s_{NN}} = 5.02$ TeV. The v_2 is shown as a function of transverse momentum, p_T, in Fig. 7.1; rapidity in Fig. 7.2; and collision centrality in Fig. 7.3. Due to symmetry, the forward and backward rapidity regions were combined. In all cases, the measured v_2 fluctuates around zero with no systematic behavior.

In parallel with this analysis, another group of researchers carried out a measurement of the Υ v_2 using the scalar-product method [76], described in Sec. 5.8. A comparison of the results of the two analyses is shown in Fig. 7.4. The two methods are not expected to yield identical results [78], as discussed in Sec. 5.8. However, they should be similar. Given the statistical uncertainties in the measurements, we cannot conclude with confidence that the measured v_2 values are different.

A comparison of the Υ(1S) v_2 measurements with current theoretical predictions is shown in Fig. 7.5. The blue dashed line (Yao *et al.* [79, 80, 81]) is computed via simulations of heavy quarks in a QCD medium described by 2+1D viscous hydrodynamics, where the transport of the quarks is described by coupled Boltzmann transport equations. The purple dot-dashed line (Hong, Lee [82, 83]) is a kinetic model in which heavy quarks scatter with partons in the medium using matrix ele-

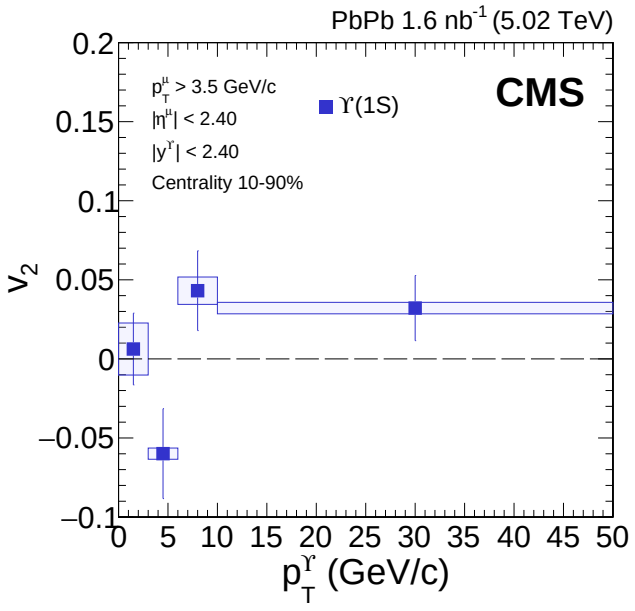

Figure 7.1. The measured v_2 of Υ(1S) (blue squares) as a function of p_T. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

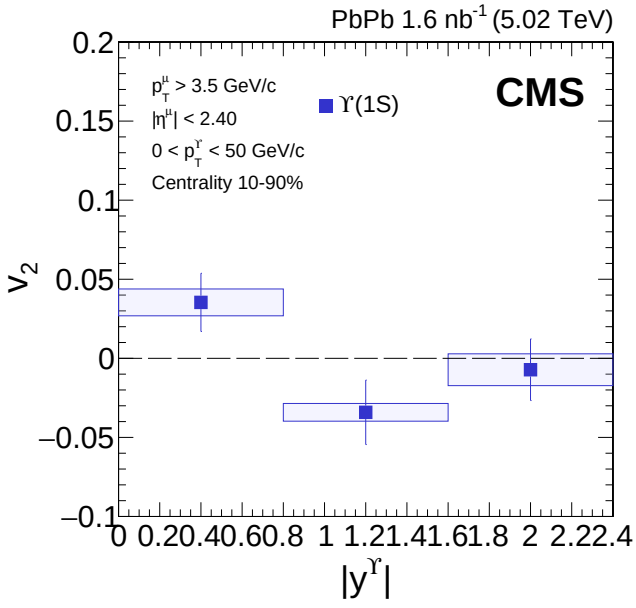

Figure 7.2. The measured v_2 of Υ(1S) (blue squares) as a function of the absolute value of rapidity. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

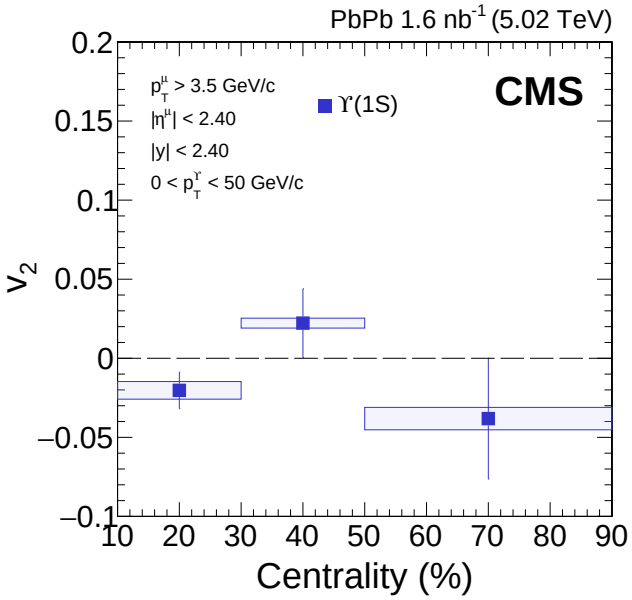

Figure 7.3. The measured v_2 of Υ(1S) (blue squares) as a function of centrality. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

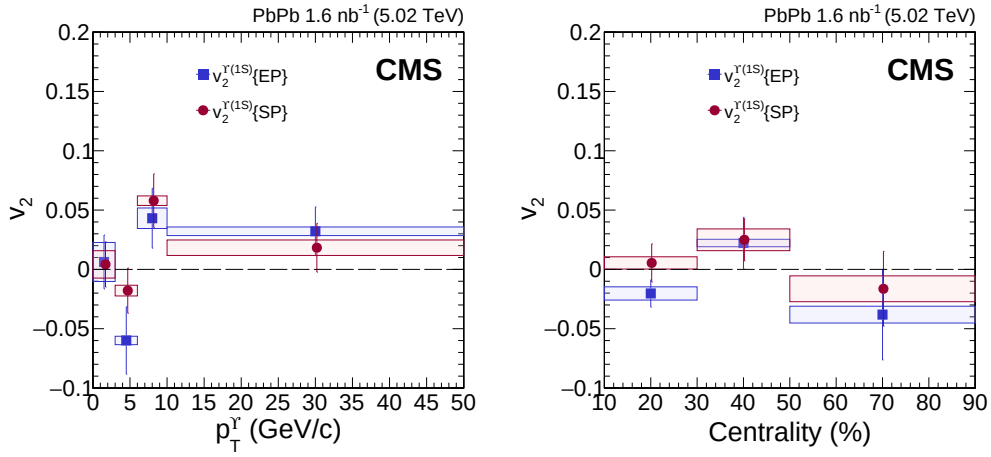

Figure 7.4. Comparison of the v_2 of Υ(1S) measured via the event-plane method (blue squares) and the scalar-product method (red circles) as a function of p_T (left) and centrality (right). The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty. The red circles are shifted slightly to the right to aid in distinguishing the statistical uncertainties.

ments calculated in perturbative QCD, while the non-perturbative soft part of the interaction between these discrete hard scatterings is modeled using a Langevin equation with transport coefficients chosen empirically. The QCD medium is again modeled by 2+1D viscous hydrodynamics. The green shaded area (Du, Rapp [34]) is the prediction referenced in Sec. 1.5 of this thesis, which models the time evolution of bottom quarks using a kinetic-rate equation in a medium modeled by a lattice-QCD based equation of state. It includes medium effects such as regeneration with temperature dependent binding energies and reaction rates, and B-meson resonance states. The red dashed line (Bhaduri *et al.* [84]) is computed from simulations of various bottomonium states with temperature-dependent decay widths in the 3+1D quasiparticle anisotropic hydrodynamic model. The details of the models can be found in the references.

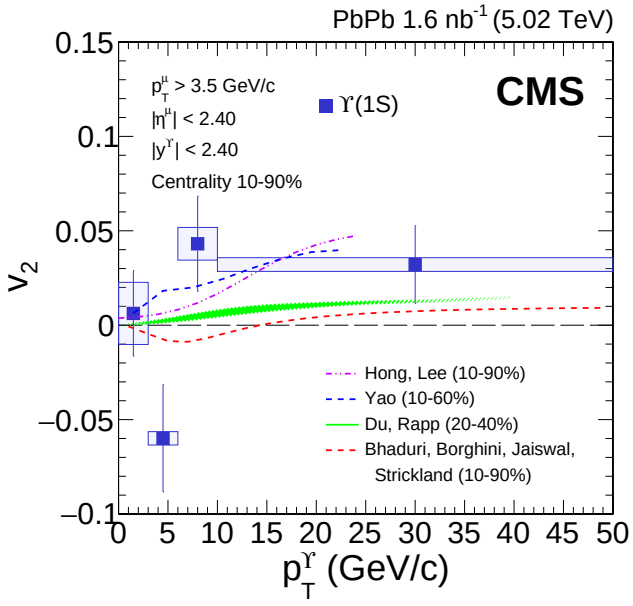

Figure 7.5. The v_2 of Υ(1S) measured via the event-plane method (blue squares) as a function of p_T compared model calculations from Du and Rapp [34], Yao [79, 80, 81], Hong and Lee [82, 83], and Bhaduri *et al.* [84].

All these models, despite their differences, predict a very small v_2 for the Υ(1S). The measured v_2 presented here and in Ref. [76] represent the first measurements of Υ(1S) v_2 using CMS data. The ALICE Collaboration has also measured the

v_2 for Υ(1S) mesons at forward rapidity ($2.5 < y^{\Upsilon} < 4$) in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV [85]. All these measurements are consistent with zero and are thus in agreement with the model predictions, given the statistics. The small v_2 is indicative of the high mass of the bottom quark and its consequent resistance to modifications in its momentum from medium effects. All the models predict a rise in the v_2 above a p_T of about 10 GeV/c, but the change is much greater for two of the models (Yao and Hong, Lee) compared to the other two, so that discrimination between these models may be possible with more data at high p_T.

The v_2 of the Υ(2S) was also extracted in this analysis in the same manner as the Υ(1S), and the results are shown as a function of transverse momentum p_T in Fig. 7.6; rapidity in Fig. 7.7; and collision centrality in Fig. 7.8. There are fewer Υ(2S) than Υ(1S), which results in larger uncertainties compared to the Υ(1S). The measured v_2 of the Υ(2S) is also consistent with zero.

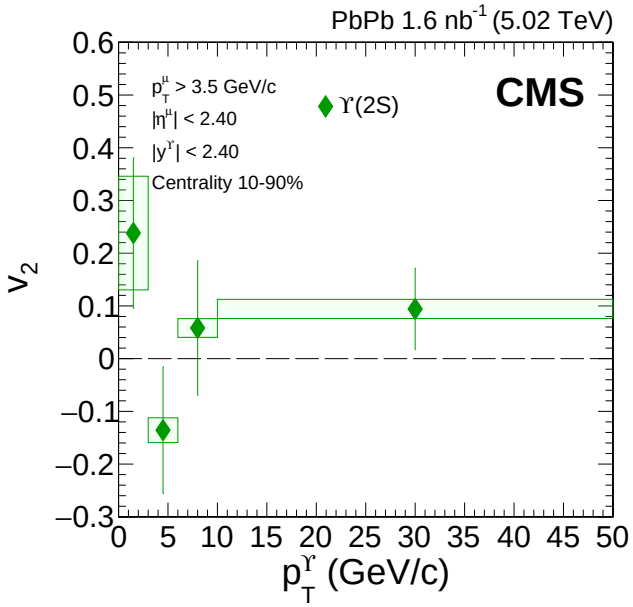

Figure 7.6. The measured v_2 of Υ(2S) (green diamonds) as a function of p_T. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

An interesting prediction from Refs. [34, 49, 50] was that the v_2 of the Υ(2S) was expected to be approximately twice as large as the v_2 of the Υ(1S) as a function

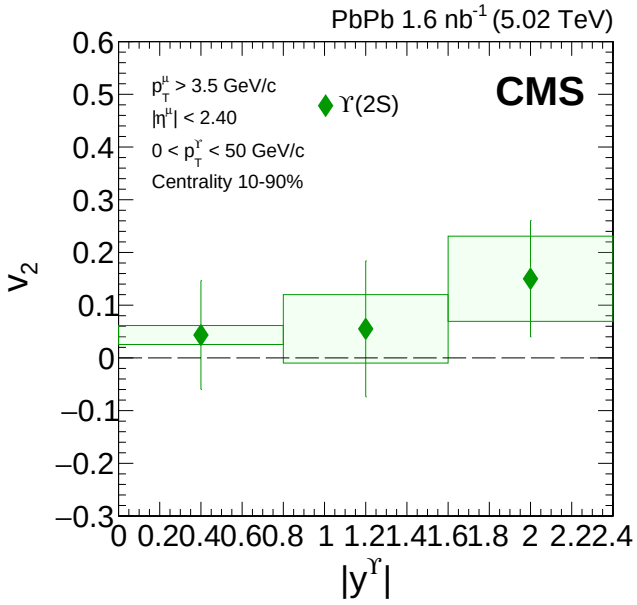

Figure 7.7. The measured v_2 of Υ(2S) (green diamonds) as a function of the absolute value of rapidity. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

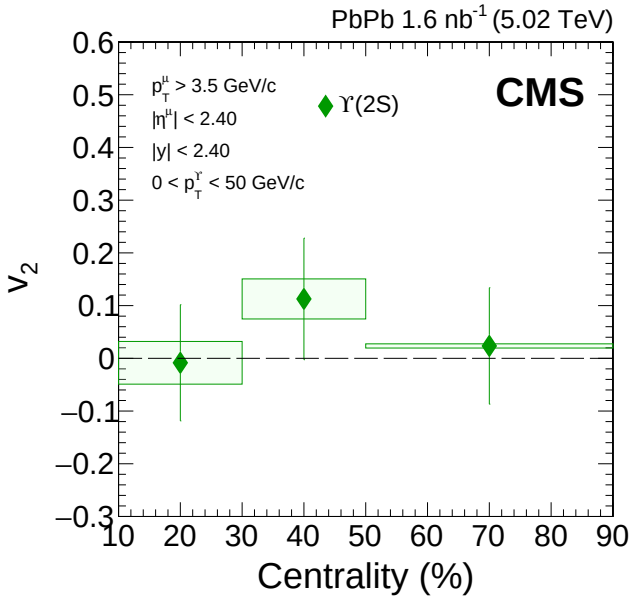

Figure 7.8. The measured v_2 of Υ(2S) (green diamonds) as a function of centrality. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

of p_T. To test this prediction, the ratio of the v_2 of the Υ(2S) and the v_2 of the Υ(1S) was calculated and the result is shown in Fig. 7.9 in three p_T bins. The lowest p_T bin is omitted because the denominator of the ratio is very close to zero and results in extremely large uncertainties in that bin. The ratio is centered around a value of 2, but given the size of the vertical bars, we cannot conclude with confidence that the measured v_2 of the Υ(2S) is larger than that of the Υ(1S). A future repetition of this study with more data can yield more robust results for the v_2 ratio.

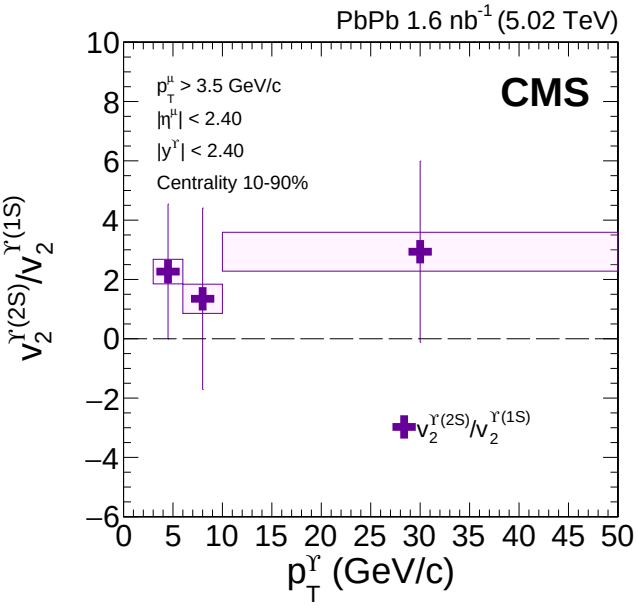

Figure 7.9. The ratio of the v_2 of Υ(2S) over the v_2 of Υ(1S) (purple crosses) as a function of p_T. The vertical bars represent the statistical uncertainty, and the shaded boxes represent the systematic uncertainty.

Chapter 8

Conclusion

The high-energy heavy-ion collisions at the LHC make possible the study of the QGP in a laboratory environment. Heavy quarkonia, including Υs, are useful probes of the QGP medium that is believed to form in these collisions. The thoughtfully designed and highly effective CMS detector is our microscope which allows us to see the intricate details of these tiny energy-dense collisions. The clearly reconstructed muon tracks provide the window into the quarkonia signals. The reconstruction, muon identification, and trigger efficiencies of CMS are very high (typically above $\sim$90%), even with the extremely high multiplicities observed in the most central Pb + Pb events.

The elliptic flow coefficient v_2 of the Υ(1S) and Υ(2S) was measured using the event-plane method with Pb + Pb data collected by the CMS Collaboration in 2018. The measurement of elliptic flow using the event-plane method is in good agreement with a parallel measurement carried out using the scalar-product method. The validity of the analysis was tested using a closure test, in which a simulated v_2 was accurately extracted using the event-plane method. Systematic uncertainties were estimated and reported with the measurements.

The v_2 of the Υ(1S) and Υ(2S) are observed to fluctuate around zero with no systematic behavior, given the current statistics. All Υ(1S) results lie within 2σ of zero, which means a 5σ nonzero measurement will require a reduction in the uncertainties of at least a factor of $\sim$3, or in other words a nine-fold increase in

data. Measurements with similar statistics for the Υ(2S) will require much more data.

The measurements are in agreement with several theoretical models which predict a very small v_2 due to the relatively high mass of the bottom quarks compared to other particles composed of lighter quarks. Discrimination among the various models is not possible with the current data. The deviation in the models is greatest at high p_T, but much more data is need to facilitate more robust estimates of the v_2 at high p_T and the ratio of the v_2 of the Υ(2S) and Υ(1S).

Appendix A

Results of Pseudoexperiments

A.1 Alternative Signal Function

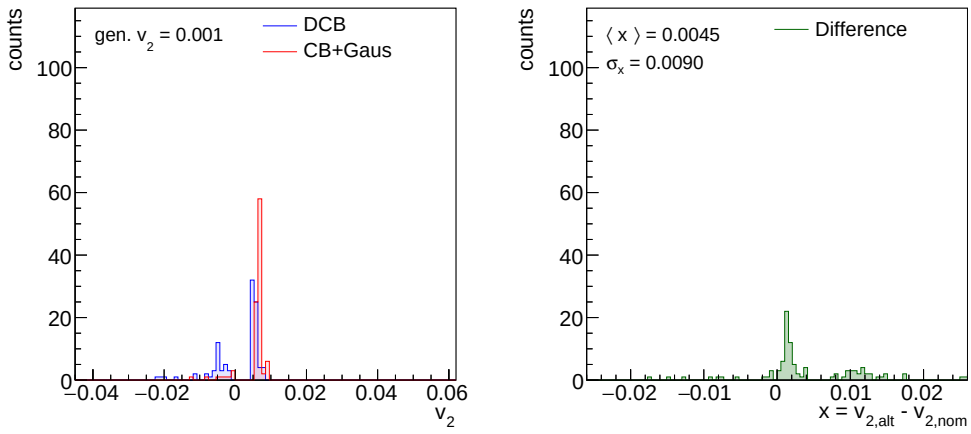

Figure A.1. Results of pseudoexperiments in the kinematic bin $0 \leq p_T < 3$ GeV/c. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1. The distribution of v_2 extracted from 100 generated sets of pseudodata using the nominal fit function is shown on the left in blue. The distribution of v_2 extracted from the same 100 sets of pseudodata using the alternative fit function is also shown on the left in red. The generated v_2 is noted in the top left corner. The plot on the right shows the distribution of the differences in the extracted v_2 between the two fitting methods. The mean ($\langle x \rangle$) and standard deviation (σ_x) are noted in the figure.

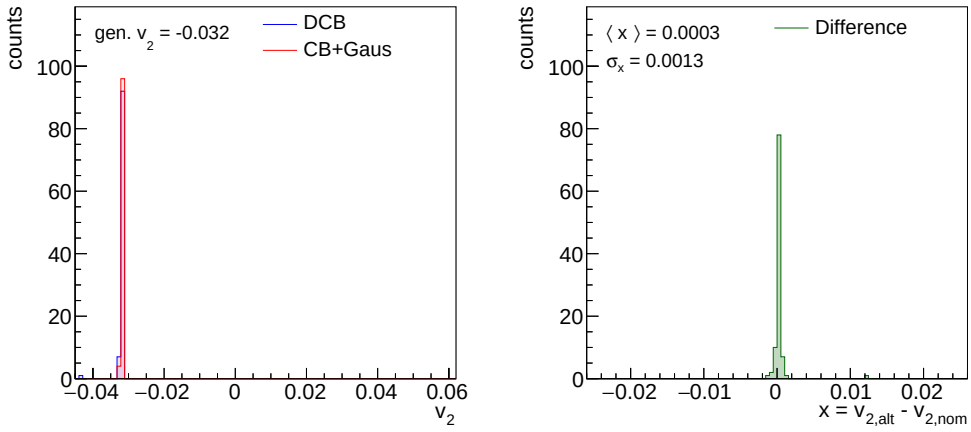

Figure A.2. Results of pseudoexperiments in the kinematic bin $3 \leq p_T < 6$ GeV/c. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

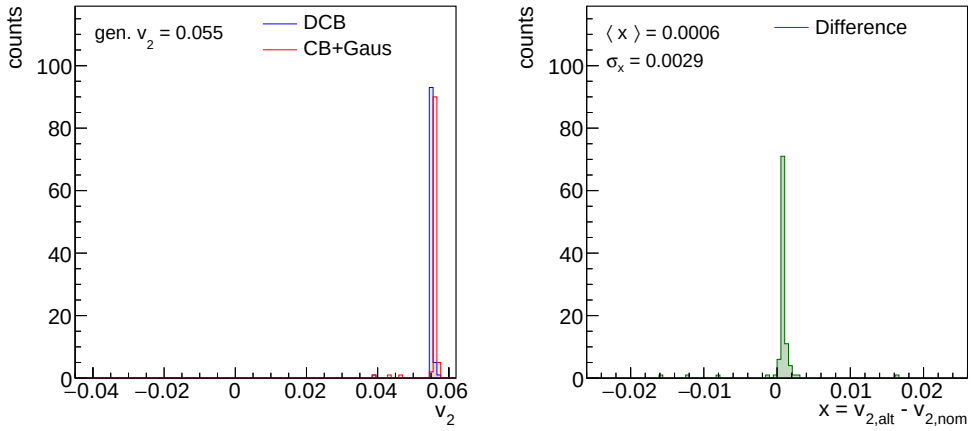

Figure A.3. Results of pseudoexperiments in the kinematic bin $6 \leq p_T < 10$ GeV/c. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

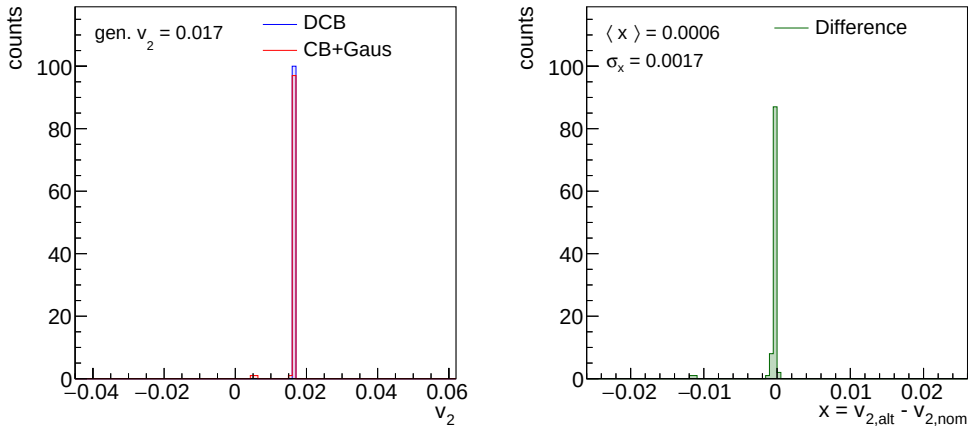

Figure A.4. Results of pseudoexperiments in the kinematic bin $10 \leq p_T < 50$ GeV/c. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

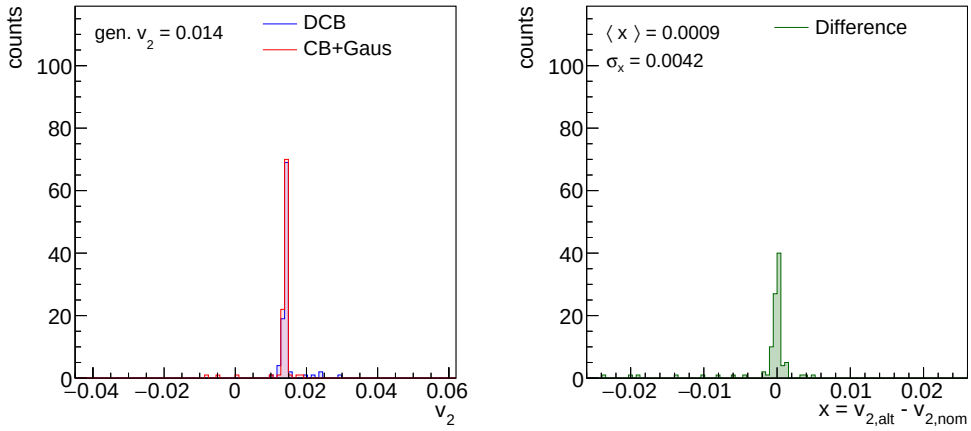

Figure A.5. Results of pseudoexperiments in the kinematic bin $0 < y < 0.8$. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

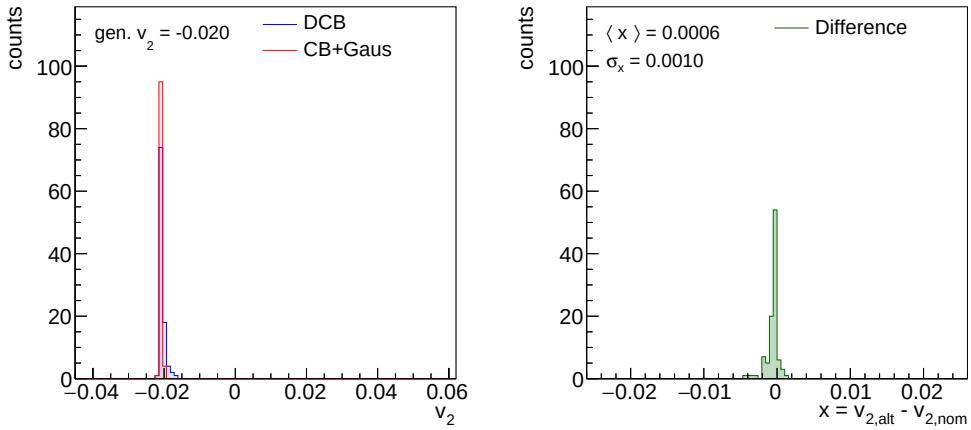

Figure A.6. Results of pseudoexperiments in the kinematic bin $0.8 < y < 1.6$. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

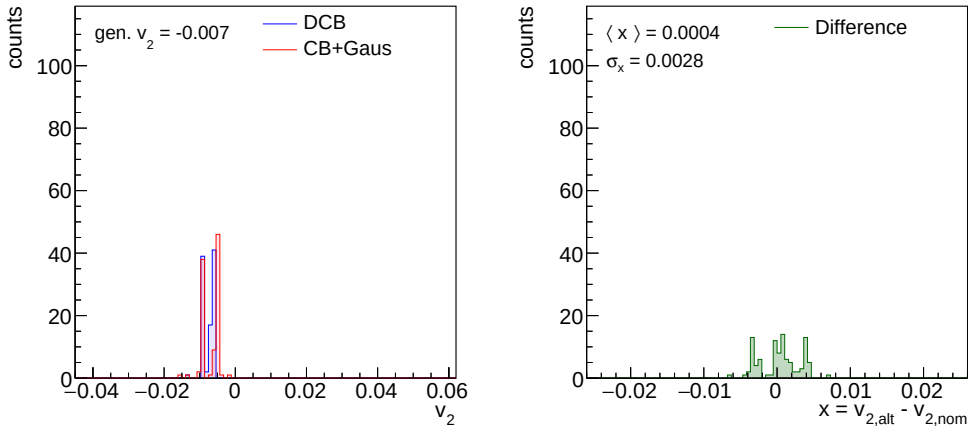

Figure A.7. Results of pseudoexperiments in the kinematic bin $1.6 < y < 2.4$. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

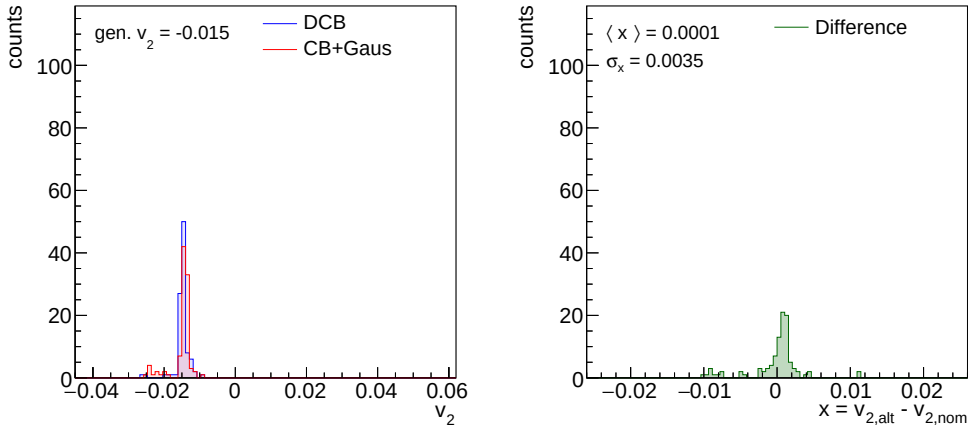

Figure A.8. Results of pseudoexperiments in the 10–30% centrality bin. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

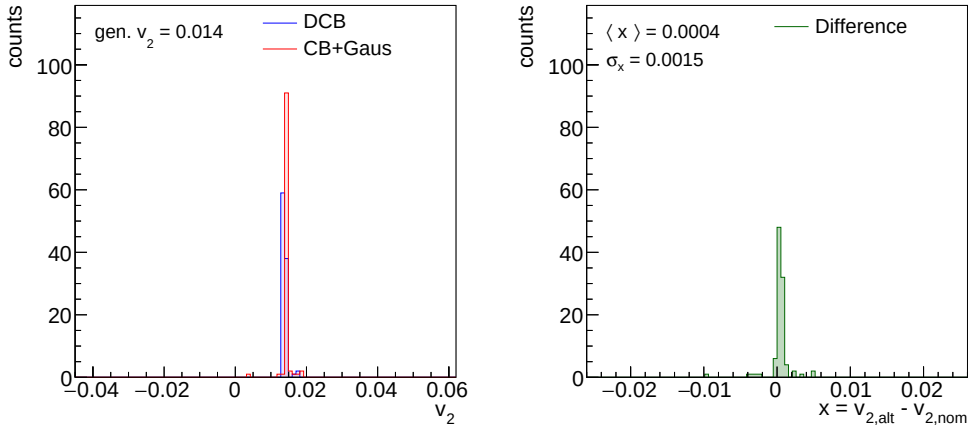

Figure A.9. Results of pseudoexperiments in the 30–50% centrality bin. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

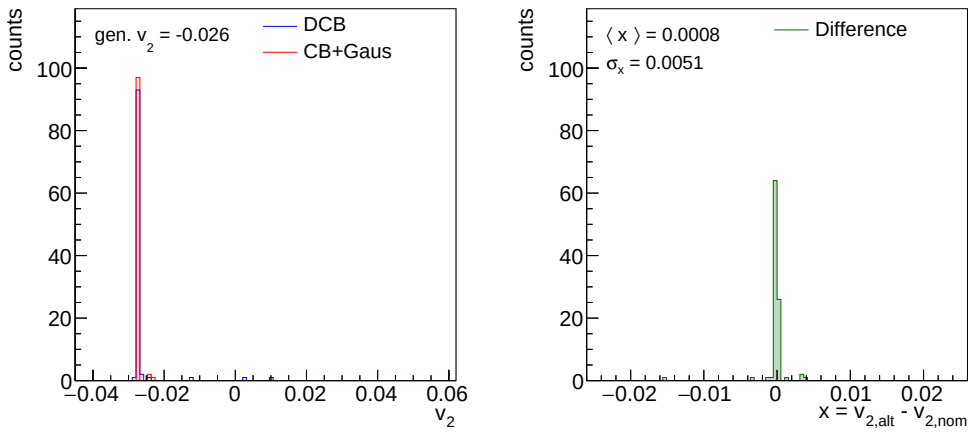

Figure A.10. Results of pseudoexperiments in the 50–90% centrality bin. The nominal signal function is a DCB (see Sec. 5.1) and the alternative signal function is a CB plus a Gaussian as described in Sec. 6.1.

A.2 Alternative Constraints

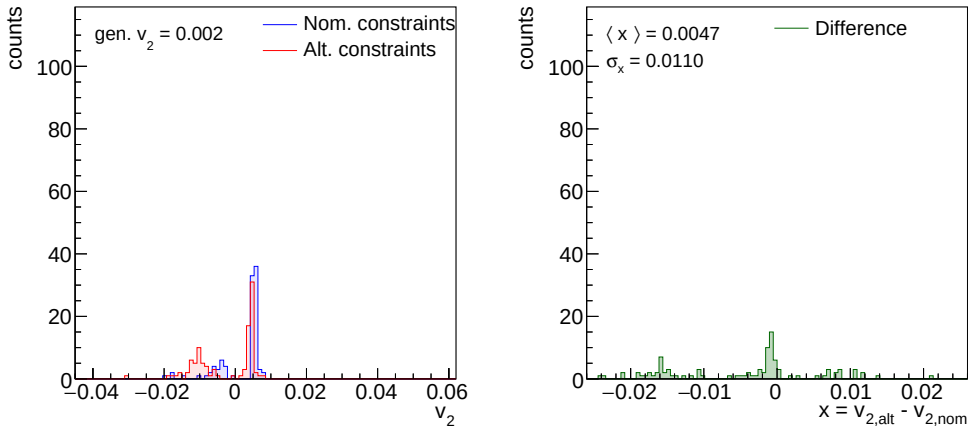

Figure A.11. Results of pseudoexperiments in the kinematic bin $0 \leq p_T < 3$ GeV/c. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2. The distribution of v_2 extracted from 100 generated sets of pseudodata using the nominal fit function is shown on the left in blue. The distribution of v_2 extracted from the same 100 sets of pseudodata using the alternative fit function is also shown on the left in red. The generated v_2 is noted in the top left corner. The plot on the right shows the distribution of the differences in the extracted v_2 between the two fitting methods. The mean ($\langle x \rangle$) and standard deviation (σ_x) are noted in the figure.

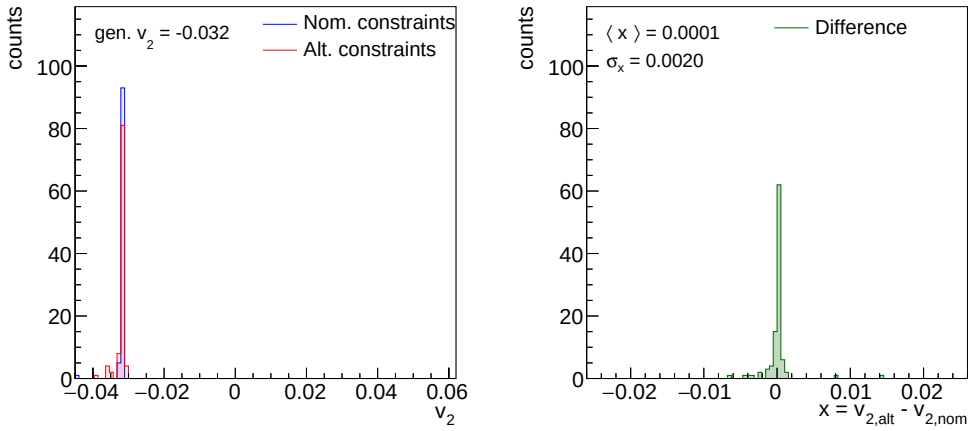

Figure A.12. Results of pseudoexperiments in the kinematic bin $3 \leq p_T < 6$ GeV/c. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

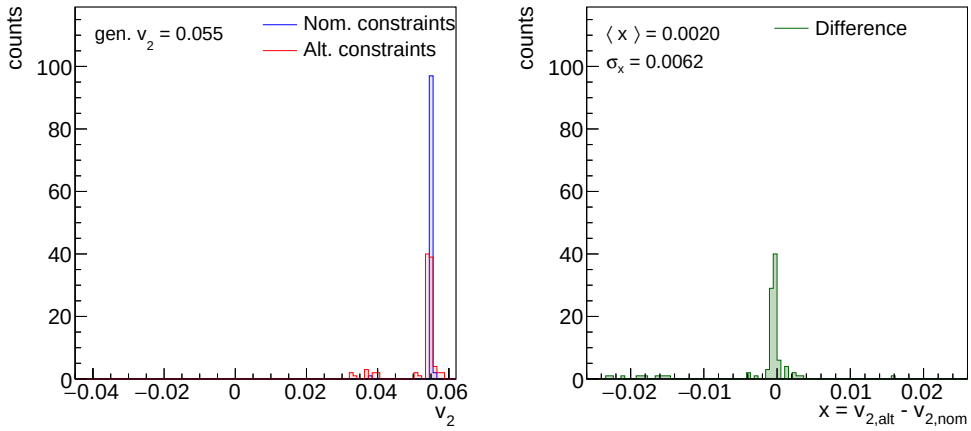

Figure A.13. Results of pseudoexperiments in the kinematic bin $6 \leq p_T < 10$ GeV/c. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

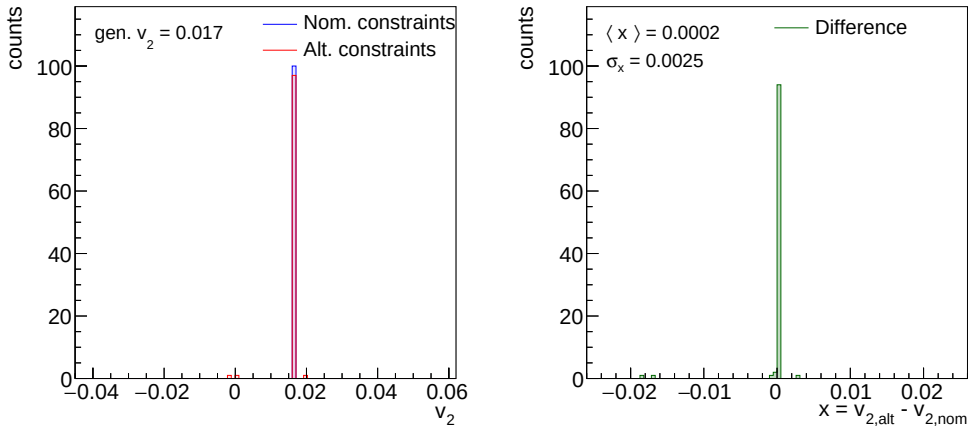

Figure A.14. Results of pseudoexperiments in the kinematic bin $10 \leq p_T < 50$ GeV/c. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

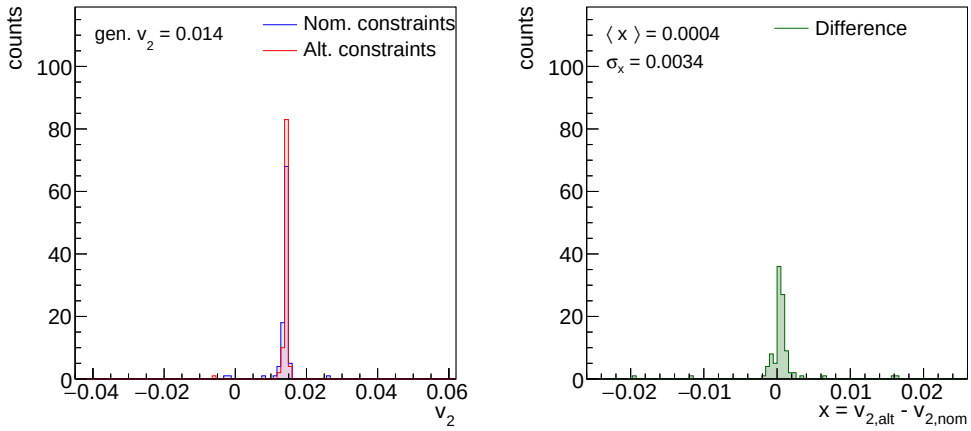

Figure A.15. Results of pseudoexperiments in the kinematic bin $0 < y < 0.8$. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

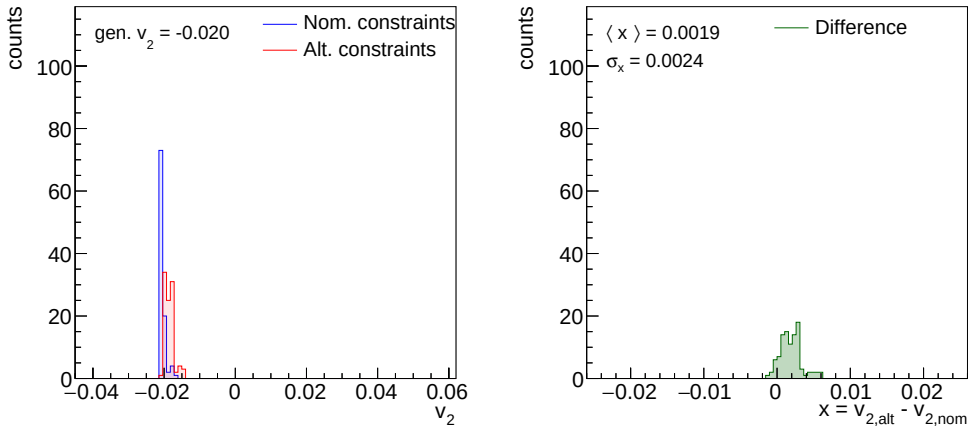

Figure A.16. Results of pseudoexperiments in the kinematic bin $0.8 < y < 1.6$. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

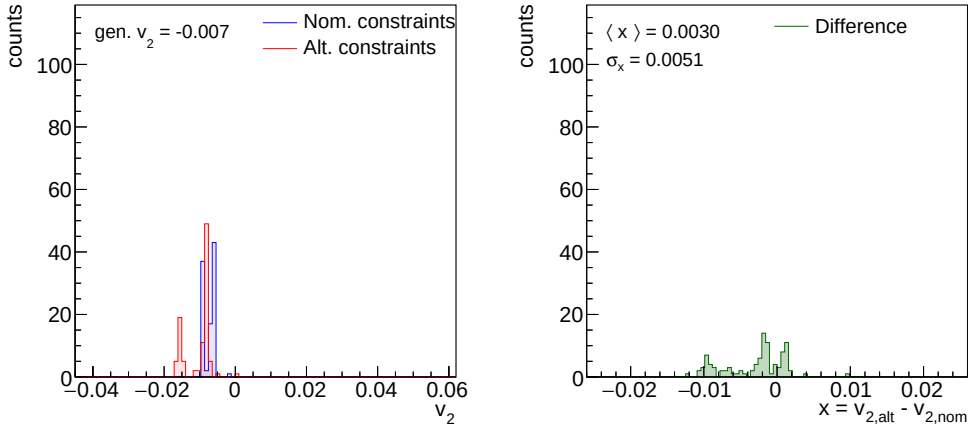

Figure A.17. Results of pseudoexperiments in the kinematic bin $1.6 < y < 2.4$. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

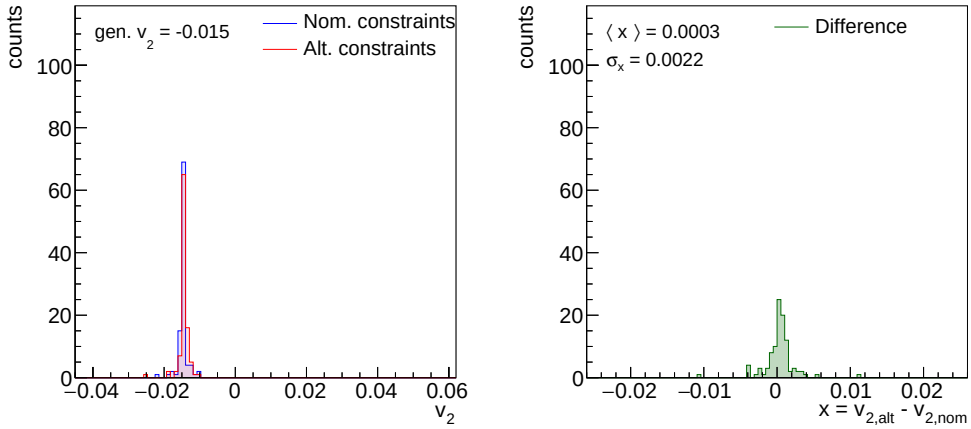

Figure A.18. Results of pseudoexperiments in the 10–30% centrality bin. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

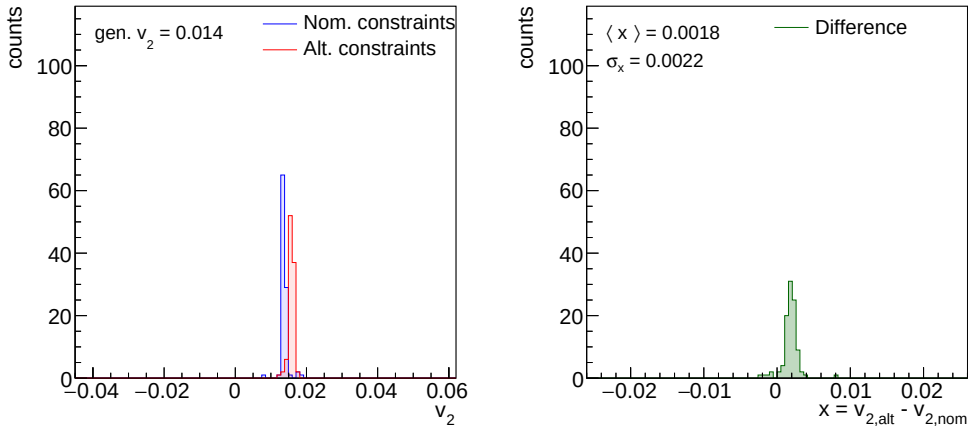

Figure A.19. Results of pseudoexperiments in the 30–50% centrality bin. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

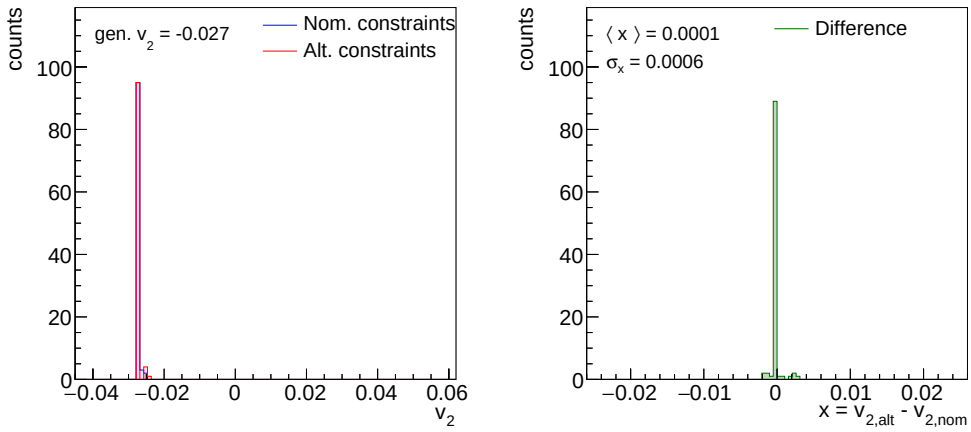

Figure A.20. Results of pseudoexperiments in the 50–90% centrality bin. The nominal constraints on the signal parameters are described in Sec. 5.1. The alternative constraints are described in Sec. 6.2.

A.3 Alternative Background Function

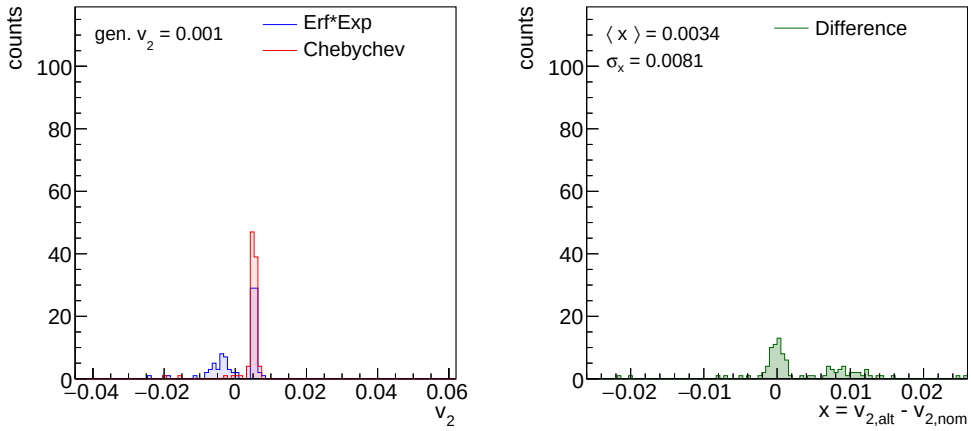

Figure A.21. Results of pseudoexperiments in the kinematic bin $0 \leq p_T < 3$ GeV/c. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3. The distribution of v_2 extracted from 100 generated sets of pseudodata using the nominal fit function is shown on the left in blue. The distribution of v_2 extracted from the same 100 sets of pseudodata using the alternative fit function is also shown on the left in red. The generated v_2 is noted in the top left corner. The plot on the right shows the distribution of the differences in the extracted v_2 between the two fitting methods. The mean ($\langle x \rangle$) and standard deviation (σ_x) are noted in the figure.

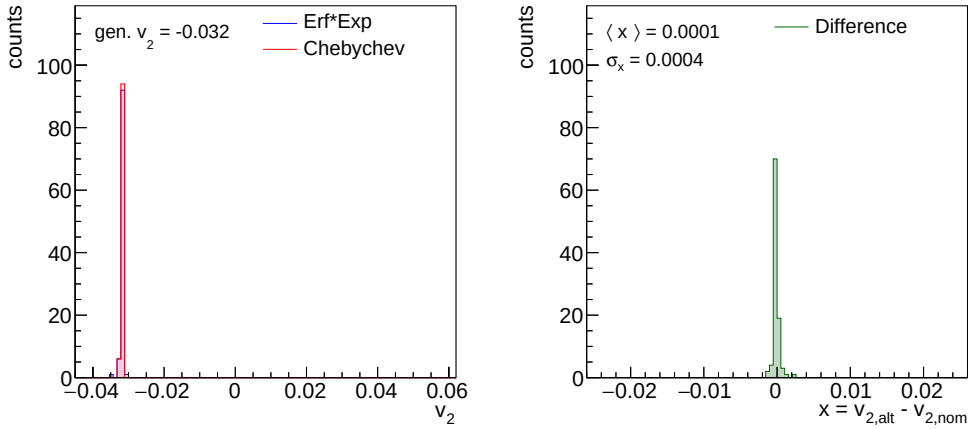

Figure A.22. Results of pseudoexperiments in the kinematic bin $3 \leq p_T < 6$ GeV/c. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

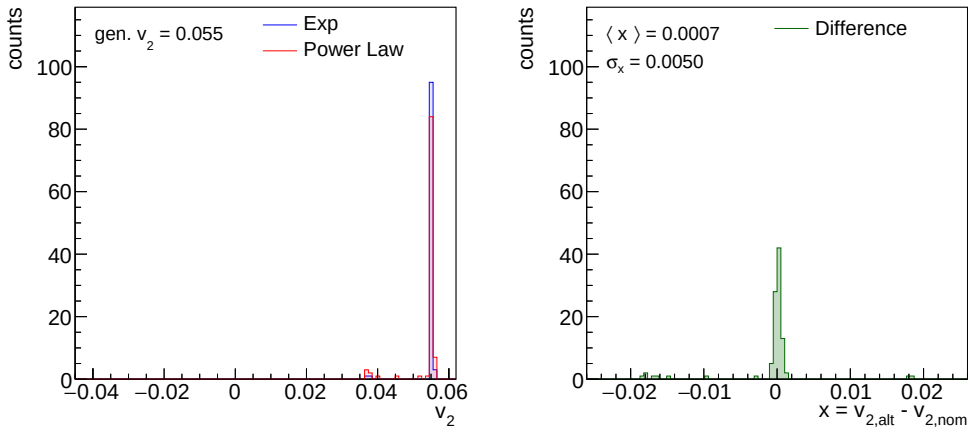

Figure A.23. Results of pseudoexperiments in the kinematic bin $6 \leq p_T < 10$ GeV/c. The nominal background function is an exponential and the alternative background function is a power law as described in Sec. 6.3.

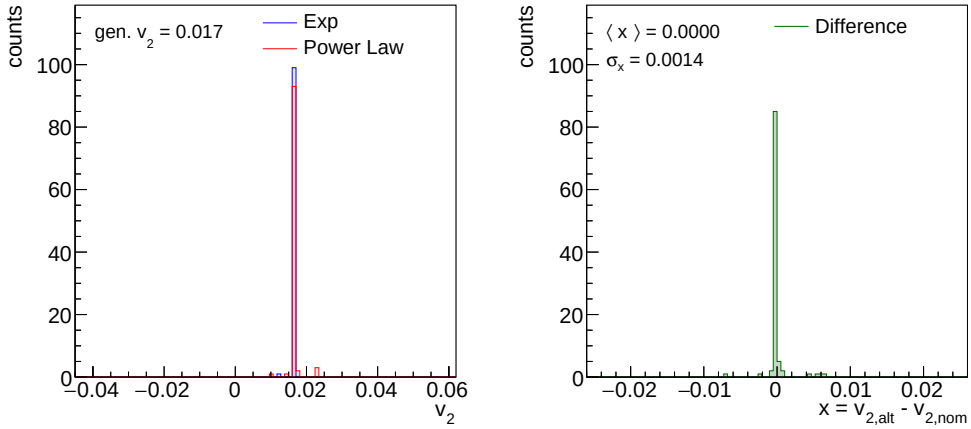

Figure A.24. Results of pseudoexperiments in the kinematic bin $10 \leq p_T < 50$ GeV/c. The nominal background function is an exponential and the alternative background function is a power law as described in Sec. 6.3.

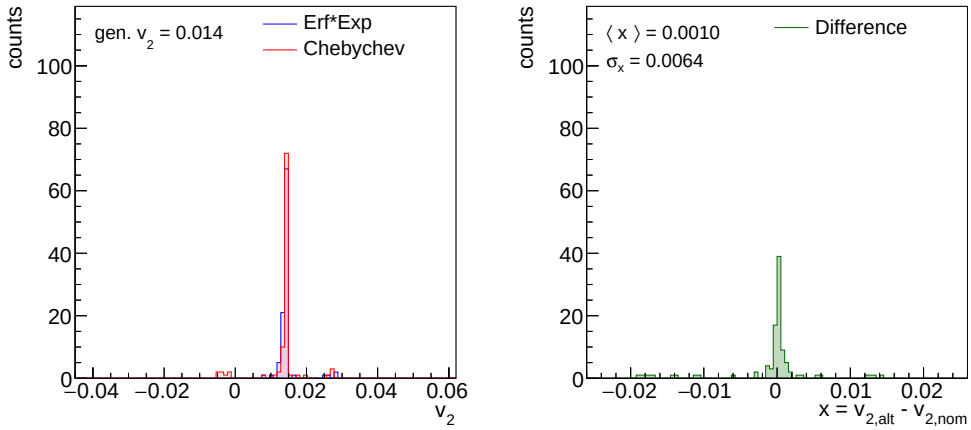

Figure A.25. Results of pseudoexperiments in the kinematic bin $0 < y < 0.8$. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

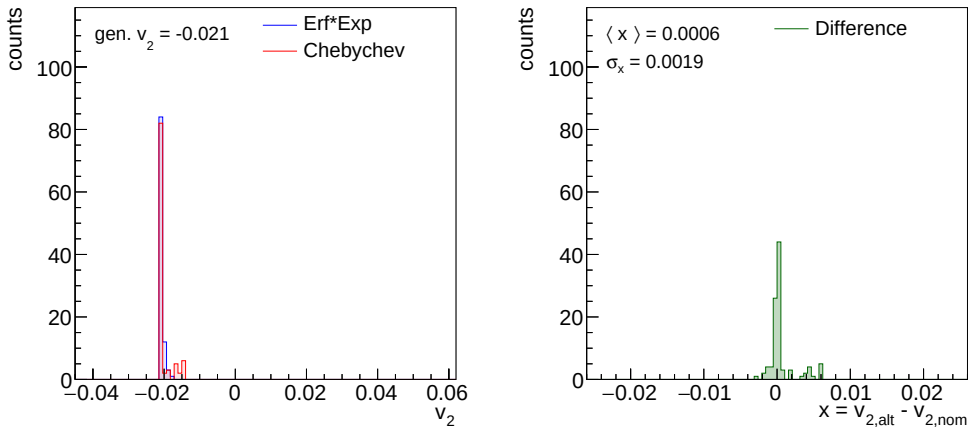

Figure A.26. Results of pseudoexperiments in the kinematic bin $0.8 < y < 1.6$. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

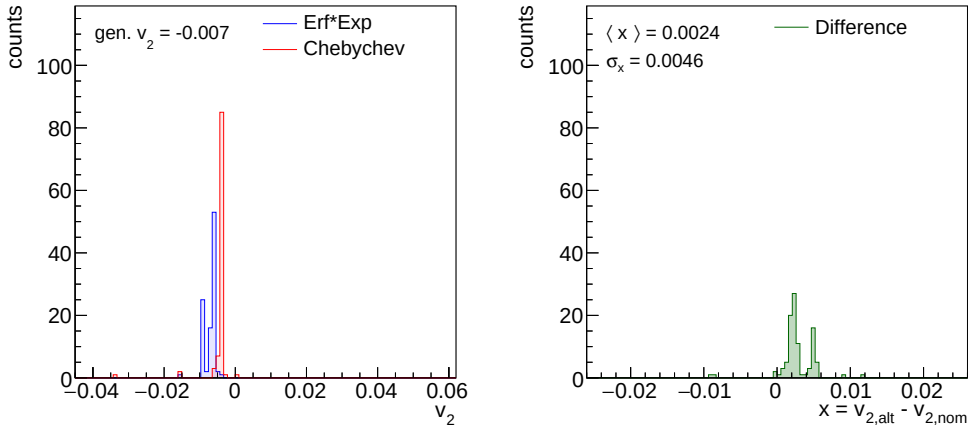

Figure A.27. Results of pseudoexperiments in the kinematic bin $1.6 < y < 2.4$. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

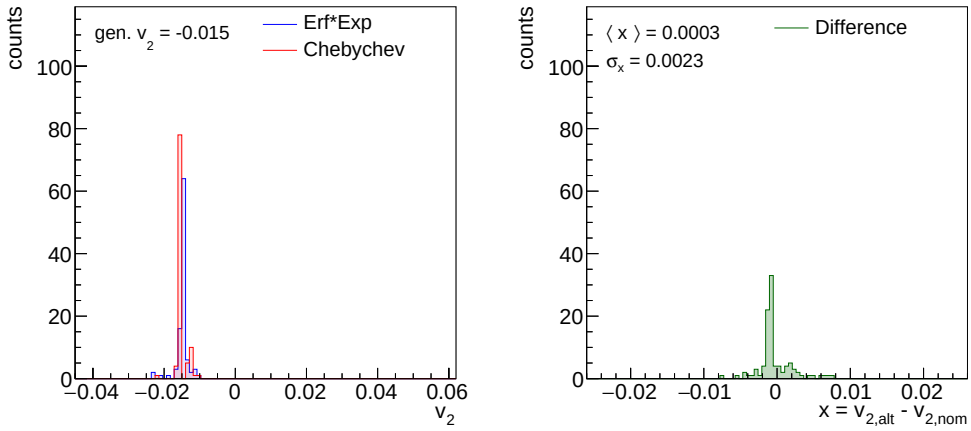

Figure A.28. Results of pseudoexperiments in the 10–30% centrality bin. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

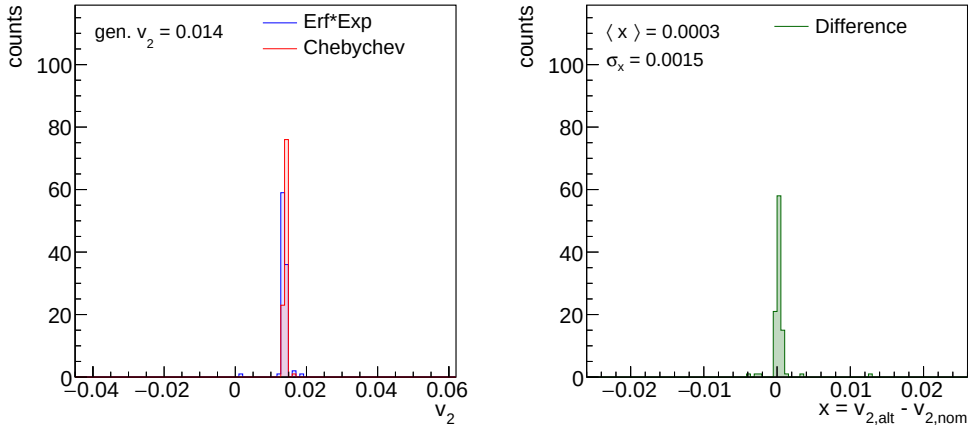

Figure A.29. Results of pseudoexperiments in the 30–50% centrality bin. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

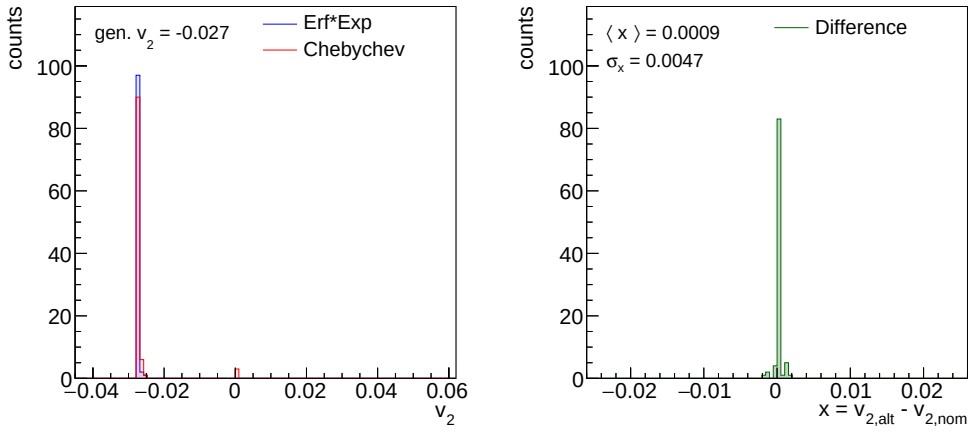

Figure A.30. Results of pseudoexperiments in the 50–90% centrality bin. The nominal background function is an error function times an exponential and the alternative background function is a fourth-order Chebychev polynomial as described in Sec. 6.3.

References

[1] P. Langacker, *The standard model and beyond; 2nd ed.* High energy physics, cosmology and gravitation, CRC Press, 2017.

[2] J. F. Donoghue, E. Golowich, and B. R. Holstein, *Dynamics of the Standard Model*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology, Cambridge University Press, 2 ed., 2014.

[3] G. Altarelli, "Collider Physics within the Standard Model: a Primer," 2013.

[4] N. Brambilla, S. Eidelman, B. K. Heltsley, R. Vogt, G. T. Bodwin, E. Eichten, A. D. Frawley, A. B. Meyer, R. E. Mitchell, V. Papadimitriou, extitet al., "Heavy quarkonium: progress, puzzles, and opportunities," *European Physical Journal C*, vol. 71, no. 2, 2011.

[5] R. Averbeck, "Heavy-flavor production in heavy-ion collisions and implications for the properties of hot QCD matter," *Progress in Particle and Nuclear Physics*, vol. 70, p. 159209, 2013.

[6] G. Martinez, "Advances in Quark Gluon Plasma," 2013.

[7] F. R. Brown, F. P. Butler, H. Chen, N. H. Christ, Z. Dong, W. Schaffer, L. I. Unger, and A. Vaccarino, "On the existence of a phase transition for QCD with three light quarks," *Physical Review Letters*, vol. 65, pp. 2491–2494, 1990.

[8] L. Turko, D. Blaschke, D. Prorok, and J. Berdermann, "Effective degrees of freedom in QCD thermodynamics," *EPJ Web of Conferences*, vol. 71, p. 00134, 2014.

[9] M. Asakawa and K. Yazaki, "Chiral Restoration at Finite Density and Temperature," *Nuclear Physics A*, vol. 504, pp. 668–684, 1989.

[10] M. A. Stephanov, "QCD phase diagram and the critical point," *International Journal of Modern Physics A*, vol. 20, no. 19, p. 43874392, 2005.

[11] D. Tlusty, "The RHIC Beam Energy Scan Phase II: Physics and Upgrades," in *13th Conference on the Intersections of Particle and Nuclear Physics*, 2018.

[12] A. Bazavov, T. Bhattacharya, M. Cheng, C. DeTar, H.-T. Ding, S. Gottlieb, R. Gupta, P. Hegde, U. M. Heller, F. Karsch, extitet al., "Chiral and deconfinement aspects of the QCD transition," *Physical Review D*, vol. 85, no. 5, 2012.

[13] F. Karsch, "Lattice QCD results on cumulant ratios at freeze-out," *Journal of Physics: Conference Series*, vol. 779, p. 012015, 2017.

[14] D. Schwarz, "The first second of the Universe," *Annalen der Physik*, vol. 12, no. 4, p. 220270, 2003.

[15] NA49 Collaboration, "Our current understanding on the evolution of heavy ion collisions." `https://na49info.web.cern.ch/Public/Press/findings.html`.

[16] R. Vogt, *Ultrarelativistic Heavy-Ion Collisions*. Elsevier Science, 2007.

[17] C. Loizides, J. Kamin, and D. d'Enterria, "Improved Monte Carlo Glauber predictions at present and future nuclear colliders," *Physical Review C*, vol. 97, p. 054910, 2018.

[18] J. A. Nelder and R. Mead, "A Simplex Method for Function Minimization," *Comput. J.*, vol. 7, pp. 308–313, 1965.

[19] P. Romatschke and U. Romatschke, "Viscosity Information from Relativistic Nuclear Collisions: How Perfect is the Fluid Observed at RHIC?," *Physical Review Letters*, vol. 99, p. 172301, 2007.

[20] J. I. Kapusta, "Viscous Properties of Strongly Interacting Matter at High Temperature," *Landolt-Bornstein*, vol. 23, pp. 563–580, 2010.

[21] X. Li, "Creation of quarkgluon plasma droplets with three distinct geometries," *Nature Physics*, vol. 15, 2019.

[22] H. Song, S. A. Bass, U. Heinz, T. Hirano, and C. Shen, "200 A GeV Au + Au Collisions Serve a Nearly Perfect Quark-Gluon Liquid," *Physical Review Letters*, vol. 106, p. 192301, 2011.

[23] M. Strickland, "Thermal Bottomonium Suppression," *AIP Conf. Proc.*, vol. 1520, no. 1, pp. 179–184, 2013.

[24] Particle Data Group, "Review of Particle Physics," *PTEP*, vol. 2020, no. 8, p. 083C01, 2020.

[25] G. Bodwin, E. Braaten, and G. Lepage, "Rigorous QCD analysis of inclusive annihilation and production of heavy quarkonium," *Physical Review D, Particles and fields*, vol. 51, no. 3, p. 11251171, 1995.

[26] T. Matsui and H. Satz, "J/ψ Suppression by Quark-Gluon Plasma Formation," *Physics Letters B*, vol. 178, pp. 416–422, 1986.

[27] CMS Collaboration, "Measurement of nuclear modification factors of Υ(1S), Υ(2S), and Υ(3S) mesons in PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV," *Physics Letters B*, vol. 790, pp. 270–293, 2019.

[28] A. Andronic, F. Arleo, R. Arnaldi, A. Beraudo, E. Bruna, D. Caffarri, Z. C. del Valle, J. G. Contreras, T. Dahms, and A. e. Dainese, "Heavy-flavour and quarkonium production in the LHC era: from protonproton to heavy-ion collisions," *European Physical Journal C*, vol. 76, no. 3, 2016.

[29] S. Digal, P. Petreczky, and H. Satz, "Sequential quarkonium suppression," in *International Workshop on the Physics of the Quark Gluon Plasma*, 2001.

[30] B. Krouppa, A. Rothkopf, and M. Strickland, "Bottomonium suppression using a lattice QCD vetted potential," *Physical Review D*, vol. 97, no. 1, 2018.

[31] CMS Collaboration, "Nuclear modification of Υ states in pPb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV," Tech. Rep. CMS-PAS-HIN-18-005, CERN, 2019.

[32] Y. Burnier, O. Kaczmarek, and A. Rothkopf, "Quarkonium at finite temperature: towards realistic phenomenology from first principles," *Journal of High Energy Physics*, vol. 2015, no. 12, p. 134, 2015.

[33] Y. Burnier, O. Kaczmarek, and A. Rothkopf, "Static Quark-Antiquark Potential in the Quark-Gluon Plasma from Lattice QCD," *Physical Review Letters*, vol. 114, no. 8, 2015.

[34] X. Du, M. He, and R. Rapp, "Color screening and regeneration of bottomonia in high-energy heavy-ion collisions," *Physical Review C*, vol. 96, no. 5, 2017.

[35] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, "Parton distributions for the LHC," *European Physical Journal C*, vol. 63, no. 2, p. 189285, 2009.

[36] R. Wang, X. Chen, and Q. Fu, "Global study of nuclear modifications on parton distribution functions," *Nuclear Physics B*, vol. 920, p. 119, 2017.

[37] European Muon Collaboration, "The ratio of the nucleon structure functions F_2^n for iron and deuterium," *Physics Letters B*, vol. 123, pp. 275–278, 1983.

[38] J. Gomez, R. G. Arnold, P. E. Bosted, C. C. Chang, A. T. Katramatou, G. G. Petratos, A. A. Rahbar, S. E. Rock, A. F. Sill, Z. M. Szalata, A. Bodek, N. Giokaris, D. J. Sherden, B. A. Mecking, and R. M. Lombard-Nelsen, "Measurement of the A dependence of deep-inelastic electron scattering," *Physical Review D*, vol. 49, pp. 4348–4372, 1994.

[39] European Muon Collaboration, "New Measurements of European Muon Collaboration Effect in Very Light Nuclei," *Physical Review Letters*, vol. 103, p. 202301, 2009.

[40] F. Arleo, C.-J. Nam, and S. Platchkov, "Initial-state energy loss in cold QCD matter and the Drell-Yan process," *Journal of High Energy Physics*, vol. 2019, no. 1, 2019.

[41] N. Armesto and E. Scomparin, "Heavy-ion collisions at the Large Hadron Collider: A review of the results from Run 1," *European Physical Journal Plus*, vol. 131, no. 3, 2016.

[42] E. G. Ferreiro, "Excited charmonium suppression in protonnucleus collisions as a consequence of comovers," *Physics Letters B*, vol. 749, pp. 98–103, 2015.

[43] CMS Collaboration, "Observation of Long-Range Near-Side Angular Correlations in Proton-Proton Collisions at the LHC," *Journal of High Energy Physics*, vol. 09, p. 091, 2010.

[44] CMS Collaboration, "Long-range and short-range dihadron angular correlations in central PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV," *Journal of High Energy Physics*, vol. 2011, no. 7, 2011.

[45] CMS Collaboration, "Centrality dependence of dihadron correlations and azimuthal anisotropy harmonics in PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV," *European Physical Journal C*, vol. 72, no. 5, 2012.

[46] R. Snellings, "Elliptic flow: a brief review," *New Journal of Physics*, vol. 13, no. 5, p. 055008, 2011.

[47] CMS Collaboration, "Suppression and azimuthal anisotropy of prompt and nonprompt J/ψ production in PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV," *European Physical Journal C*, vol. 77, no. 4, 2017.

[48] CMS Collaboration, "Prompt and non-prompt J/ψ production in pp collisions at $\sqrt{s} = 7$ TeV," *European Physical Journal C*, vol. 71, no. 3, 2011.

[49] A. Islam and M. Strickland, "Bottomonium suppression and elliptic flow from real-time quantum evolution," *Physics Letters B*, vol. 811, p. 135949, 2020.

[50] N. Brambilla, M. A. Escobedo, M. Strickland, A. Vairo, P. V. Griend, and J. H. Weber, "Bottomonium production in heavy-ion collisions using quantum trajectories: Differential observables and momentum anisotropy," 2021.

[51] J.-L. Caron, "Accelerator complex of CERN: an overview of all accelerators of CERN. Vue d'ensemble de tous les acclrateurs du CERN." 2001.

[52] W. Herr and B. Muratori, "Concept of luminosity," in *CERN Accelerator School and DESY Zeuthen: Accelerator Physics*, 2003.

[53] S. Boreham, M. Brice, P. Ginter, C. Marcelloni, and M. Hoch, "Photos from the CMS Photo Book." CMS Collection., 2008.

[54] C. Sutton, "CMS: A study in compactness (archive)," *CERN Courier*, 2008. https://cerncourier.com/a/cms-a-study-in-compactness-archive/.

[55] D. d'Enterria, M. Ballintijn, M. Bedjidian, D. Hoffman, O. Kodolova, C. Loizides, I. P. Lokhtin, C. Lourenco, C. Mironov, S. V. Petrushanko, C. Roland, G. Roland, F. Sikler, and G. Veres, "CMS Physics: Technical Design Report v.2: Addendum on High Density QCD with Heavy Ions," *Journal of Physics G*, vol. 34, pp. 2307–2455, 2007.

[56] CMS Collaboration, "The CMS experiment at the CERN LHC," *Journal of Instrumentation*, vol. 3, no. 08, pp. S08004–S08004, 2008.

[57] CMS Collaboration, "Alignment of the CMS silicon tracker during commissioning with cosmic rays," *Journal of Instrumentation*, vol. 5, no. 03, p. T03009T03009, 2010.

[58] CMS Collaboration, "Identification and filtering of uncharacteristic noise in the CMS hadron calorimeter," *Journal of Instrumentation*, vol. 5, no. 03, p. T03014T03014, 2010.

[59] CMS Collaboration, "The CMS trigger system," *Journal of Instrumentation*, vol. 12, no. 01, p. P01020, 2017.

[60] CMS Collaboration, "Performance of CMS muon reconstruction in heavy-ion collisions," CMS Note 2021/001, CERN, 2021.

[61] CMS Collaboration, "Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV," *Journal of Instrumentation*, vol. 7, no. 10, p. P10002P10002, 2012.

[62] CMS Collaboration, "Suppression of non-prompt J/ψ, prompt J/ψ, and $\Upsilon(1S)$ in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV," *Journal of High Energy Physics*, vol. 05, no. 5, 2012.

[63] D. Dutta, "Minimum Bias, MPI and DPS, Diffractive and Exclusive measurements at CMS," *Nuclear and Particle Physics Proceedings*, vol. 273-275, pp. 1986–1990, 2016.

[64] CMS Collaboration, "Zero bias and HF-based minimum bias triggering for pp collisions at 14 TeV in CMS," Tech. Rep. CMS-PAS-QCD-07-002, CERN, 2007.

[65] CMS Collaboration, "Measurement of Momentum Scale and Resolution of the CMS Detector using Low-mass Resonances and Cosmic Ray Muons," tech. rep., CERN, Geneva, 2010.

[66] CMS Collaboration, "Measurement of the anisotropic flow of upsilon mesons in PbPb collisions at 5.02 TeV," CMS Note 2019/001, CERN, 2019.

[67] T. Sjstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, "An introduction to PYTHIA 8.2," *Computer Physics Communications*, vol. 191, p. 159177, 2015.

[68] CMS Collaboration, "Extraction and validation of a new set of CMS pythia8 tunes from underlying-event measurements," *European Physical Journal C*, vol. 80, no. 1, 2020.

[69] R. D. Ball, V. Bertone, S. Carrazza, L. D. Debbio, S. Forte, P. Groth-Merrild, A. Guffanti, N. P. Hartland, Z. Kassabov, J. I. Latorre, and extitet al., "Parton distributions from high-precision collider data," *European Physical Journal C*, vol. 77, no. 10, 2017.

[70] R. Baier and R. Ruckl, "Hadronic Collisions: A Quarkonium Factory," *Z. Phys. C*, vol. 19, p. 251, 1983.

[71] R. Gastmans, W. Troost, and T. T. Wu, "Cross sections for gluon + gluon $\rightarrow$ heavy quarkonium + gluon," *Physics Letters B*, vol. 184, no. 2, pp. 257–260, 1987.

[72] P. Cho and A. K. Leibovich, "Color-octet quarkonia production. II," *Physical Review D*, vol. 53, no. 11, p. 62036217, 1996.

[73] I. P. Lokhtin, L. V. Malinina, S. V. Petrushanko, A. M. Snigirev, I. Arsene, and K. Tywoniuk, "Heavy ion event generator HYDJET++ (HYDrodynamics plus JETs)," *Comput. Phys. Commun.*, vol. 180, pp. 779–799, 2009.

[74] J. E. Gaiser, "Charmonium Spectroscopy From Radiative Decays of the J/ψ and ψ'," other thesis, Stanford University, 1982.

[75] A. M. Poskanzer and S. A. Voloshin, "Methods for analyzing anisotropic flow in relativistic nuclear collisions," *Physical Review C*, vol. 58, no. 3, p. 16711678, 1998.

[76] CMS Collaboration, "Measurement of the azimuthal anisotropy of Υ(1S) and Υ(2S) mesons in PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV," *Physics Letters B*, vol. 819, p. 136385, 2021.

[77] STAR Collaboration, "Elliptic flow from two- and four-particle correlations in Au+Au collisions at $\sqrt{s_{NN}} = 130$ GeV," *Physical Review C*, vol. 66, no. 3, 2002.

[78] M. Luzum and J.-Y. Ollitrault, "Eliminating experimental bias in anisotropic-flow measurements of high-energy nuclear collisions," *Physical Review C*, vol. 87, no. 4, 2013.

[79] X. Yao, W. Ke, Y. Xu, S. Bass, and B. Müller, "Quarkonium production in heavy ion collisions: coupled Boltzmann transport equations," *Nuclear Physics A*, vol. 982, p. 755758, 2019.

[80] X. Yao and B. Müller, "Quarkonium inside the quark-gluon plasma: Diffusion, dissociation, recombination, and energy loss," *Physical Review D*, vol. 100, no. 1, 2019.

[81] X. Yao, W. Ke, Y. Xu, S. A. Bass, and B. Müller, "Coupled Boltzmann transport equations of heavy quarks and quarkonia in quark-gluon plasma," *Journal of High Energy Physics*, vol. 2021, no. 1, 2021.

[82] J. Hong and S. H. Lee, "Υ(1S) transverse momentum spectra through dissociation and regeneration in heavy-ion collisions," *Physics Letters B*, vol. 801, p. 135147, 2020.

[83] W. Ke, Y. Xu, and S. A. Bass, "Linearized Boltzmann-Langevin model for heavy quark transport in hot and dense QCD matter," *Physical Review C*, vol. 98, no. 6, 2018.

[84] P. P. Bhaduri, M. Alqahtani, N. Borghini, A. Jaiswal, and M. Strickland, "Fireball tomography from bottomonia elliptic flow in relativistic heavy-ion collisions," *European Physical Journal C*, vol. 81, no. 7, 2021.

[85] ALICE Collaboration, "Measurement of Υ(1S) Elliptic Flow at Forward Rapidity in Pb-Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV," *Physical Review Letters*, vol. 123, no. 19, 2019.