Boston University

OpenBU **http://open.bu.edu**

Theses & Dissertations Boston University Theses & Dissertations

2016

A search for dark matter with bottom quarks

https://hdl.handle.net/2144/14522

Boston University

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

A SEARCH FOR DARK MATTER WITH BOTTOM QUARKS

by

MICHAEL EVANS KRUSKAL

B.S., Johns Hopkins University, 2009
M.S.E, Johns Hopkins University, 2009
M.A., Boston University, 2011

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2016

© Copyright by
MICHAEL EVANS KRUSKAL
2016

Approved by

First Reader _____

Steve Ahlen, Ph.D.
Professor of Physics

Second Reader _____

John Butler, Ph.D.
Professor of Physics

Acknowledgements

Although I am the sole author of this thesis, I could have never finished it without help from countless other people. First and foremost, I would like to thank my advisor Steve Ahlen for his guidance and support throughout this entire process, and giving me the opportunity to work on the ATLAS experiment. I would also like to express my gratitude to John Butler, who has provided guidance in many aspects of my research, and the rest of my thesis committee, Ed Kearns, Michael El-Batanouny, Claudio Chamon, and Ken Lane, for their time and efforts.

In addition to my thesis committee, I would like to acknowledge the rest of the ATLAS group at Boston University, with whom I've been working for the last five years. Specifically, I would like to thank Saul Youssef, Kevin Black, Jeremy Love, Zhen Yan, Louis Helary, Lidia Dell'Asta, Clare Bernard, and Alex Long for all their help.

This analysis would not have been possible without Bjoern Penning and Yangyang Cheng, who worked directly with me on the DM+b channel. I would also like to thank Steven Schramm, Emma Rosenfeld, Tim Tait, and Tongyan Lin for their help in various theoretical aspects of the analysis, and Henso Abreu, Eitan Gozani, Yoram Rozen, Sophio Pataraia, and Priscilla Pani for their work on the complementary DM+$t\bar{t}$ channels.

I want to express my sincere gratitude to my parents Joan and Vincent, for nurturing my curiosity and emphasizing the value of knowledge. I am also grateful to my brother Elliot, who first sparked my interest in physics with discussions about everything from black holes to quantum computing.

Finally, I would like to give a special thanks my girlfriend Sue Cano, for her all of her patience, love, and support throughout my entire graduate career. If it weren't for her encouragement, I would have never completed this thesis.

A SEARCH FOR DARK MATTER WITH BOTTOM QUARKS

MICHAEL EVANS KRUSKAL

Boston University, Graduate School of Arts and Sciences, 2016

Major Professor: Steve Ahlen, Professor of Physics

ABSTRACT

Despite making up over 80% of the matter in the universe, very little is known about dark matter. Its only well-established property is that it interacts gravitationally, but does not interact with ordinary matter through any of the other known forces. Specific details such as the number of dark matter particles, their quantum properties, and their interactions remain elusive and are only loosely constrained by experiments. In this dissertation I describe a novel search for a particular type of dark matter that couples preferentially to heavy quarks, using LHC proton-proton collisions at ATLAS. With a model-independent framework, comparisons are made to results obtained from other dark matter searches, and new limits are set on various interaction strengths.

Contents

List of Tables

List of Figures

List of Abbreviations

$b-$**FDM** bottom flavored dark matter

AF2 Atlfast-II

BAO baryon acoustic oscillation

BSM beyond the standard model

C charge reversal

CMB cosmic microwave background

CP charge-parity reversal

CPT charge-parity-time reversal

CR control region

CSC cathode strip chamber

DM dark matter

DQ data quality

ECAL electromagnetic calorimeter

EF event filter

EFE Einstein field equations

EFT effective field theory

eV electron volts

FLRW Friedmann–Lemaître–Robertson–Walker

FS frozen shower

FSR final state radiation

GRL good runs list

GUT grand unified theory

H_0 null hypothesis

H_a alternative hypothesis

HCAL hadronic calorimeter

HLT high-level trigger

ID inner detector

IP interaction point

IR infrared

ISM interstellar medium

ISR initial state radiation

JER jet energy resolution

JES jet energy scale

JVF jet vertex fraction

L1 level 1

L2 level 2

LAr liquid argon

$\Lambda - \mathbf{CDM}$ Lambda-cold dark matter

LCW local cluster weighting

LHC large hadron collider

linac linear accelerator

LLR log likelihood ratio

LO leading order

MACHO massive compact halo object

MC Monte Carlo

MCP muon combined performance

MDT monitored drift tube

ME matrix element

MFV minimal flavor violation

MOND modified Newtonian dynamics

MS muon spectrometer

MVA multivariate analysis

NLO next to leading order

NNLO next to next to leading order

NP new physics

P parity reversal

PDF parton density function

PS parton showering

QCD quantum chromodynamics

QFT quantum field theory

RF radio frequency

RoI region of interest

RPC resistive plate chamber

SCT semiconducting tracker

SD spin-dependent

SF scale factor

SI spin-independent

SM Standard Model

SR signal region

T time reversal

TF transfer factor

TGC thin gap chamber

TRT transition radiation tracker

UE underlying event

UV ultraviolet

VEV vacuum expectation value

VR validation region

WIMP weakly interacting massive particle

List of Variables

d_0	The distance from the perigee of a track to the z-axis.
$\Delta\phi_{min}$	The minimal azimuthal separation between $\vec{\not{p}}_T$ and any jet, $\min_i \Delta\phi(\vec{\not{p}}_T, \vec{p}_T(j_i))$. This variable is very efficient at picking out events where the $\not{E}_T$ doesn't come from primary neutrinos. Related variables are $\Delta\phi^{\mu}_{min}$ and $\Delta\phi^{\gamma}_{min}$, which use $\vec{\not{p}}^{\,\mu}_T$ and $\vec{\not{p}}^{\,\gamma}_T$ rather than the nominal $\vec{\not{p}}_T$.
$\Delta\phi^n_x$	The azimuthal separation between $\vec{\not{p}}_T$ and the nth object of class x, $\Delta\phi(\vec{\not{p}}_T, \vec{p}_T(x_n))$.
E_T	Transverse energy, defined as $E_T \equiv p_T^2 + m^2$ for a particle of mass m.
η	Pseudorapidity, defined as $-\ln\left(\tan\frac{\theta}{2}\right)$.
H^n_T	The scalar sum of each jet's p_T, $\sum_i p_T(j_i)$ starting with the nth jet.
$\not{E}_T$	Missing transverse energy calculated using the RefFinal algorithm. Related variables are $\not{E}^{\mu}_T$ and $\not{E}^{\gamma}_T$, which are recalculated without muon and photon contributions respectively.
$\not{E}^{sig}_T$	The $\not{E}_T$ significance, $\not{E}_T/\sqrt{H_T}$. This variable helps avoid flavor bias under the assumption that $\sqrt{H_T}$ is representative of statistical fluctuations in $\not{E}_T$.
$\vec{\not{p}}_T$	The 2-dimensional missing transverse momentum vector. Related variables are $\vec{\not{p}}^{\,\mu}_T$ and $\vec{\not{p}}^{\,\gamma}_T$, which are recalculated without muon and photon contributions respectively.
m^x_T	The transverse mass between an object x and the $\vec{\not{p}}_T$, defined as $m^x_T \equiv (\not{E}_T + E^x_T)^2 - (\vec{\not{p}}_T + \vec{p}_T)^2$. For low mass particles, this becomes $m^x_T \approx 2\not{E}_T E^x_T(1 - \cos\Delta\phi^x_{min})$.
$\langle\mu\rangle$	The average number of interactions per bunch crossing.
m_{xy}	The invariant mass between the highest p_T objects of type x and y.
n_x	The number of x particles reconstructed in an event.
ϕ	The azimuthal angle around the beamline, with respect to the x-axis.
ϕ_0	The azimuthal angle of a track around its perigee.
p_T	The component of momentum transverse to the beamline.
θ	The angle, in radians, with respect to the z-axis.
z_0	The distance along the z-axis between the perigee of a track and the interaction vertex.

Glossary

4-vector A vector in space-time coordinates 7

- **4-momentum** The combination of energy and momentum into a 4-vector. 14
- **4-velocity** The normalized tangent vector of a particle's worldline. 40

acoustic oscillations Sound waves produced in the early universe, which now show up as small anisotropies. 54

annihilation The interaction of a particle/anti-particle pair that destroys both via a bosonic field. 49

asymptotic freedom The effect of the strong force that results in bare color charges becoming confined within hadrons. 19

background Known processes that resemble the signal being searched for. 9

baryogenesis The period of time in the early universe in which an excess of baryons over anti-baryons was produced. 26

baryon Hadrons consisting of exactly three quarks, such as protons and neutrons. 48

baryonic matter Ordinary matter made up of protons and neutrons. 1

beam dump The method by which a collider ejects its beam in order to cease collisions. 71

beam injection The method by which a collider creates its beam in order to begin collisions. 71

Big Bang The singularity arrived at if one traces the evolution of our universe backwards in time. Typically defined as the beginning of the universe. 1

black hole A compact, inescapable, region of space-time described by General Relativity. 54

blind A technique for avoiding bias in an analysis. The final measurements are not observed until the details of the analysis are fixed. 87

blueshift A relativistic effect that increases the energy of a particle. 41

boson Particles with whole-integer spin, such as photons and α particles. Multiple bosons can exist in the same state. 13

- **gauge** Force-mediating particles corresponding to a Lie group symmetry. 21
- **Goldstone** The massless scalar field created by spontaneous symmetry breaking. 25
- **Higgs** A scalar field responsible for giving mass to all the known particles through spontaneous symmetry breaking of the electroweak force. 24
- **weak** Any of the heavy bosons that mediate weak interactions. 24
- $W^{\pm}$ ($W^{\pm}$) Heavy, electrically charged bosons that mediate weak interactions. 23
- Z^0 (Z^0) A heavy, electrically neutral boson that mediates weak interactions. 23

Bremmsstrahlung The emission of a photon from a charged particle. This is the dominant form of energy loss for electrons in high energy electromagnetic showers. 82

***b*-tag** A method of determining whether or not a jet is a b-jet. 115

bunch crossing The crossing of two proton bunches, potentially producing a collision. 109

calibration Adjustment of observations to obtain a better estimate of reality. 91

calorimeter A detector designed to absorb and measure incident energy. 81

charge The coupling strength to a gauge field. 11

- **bare** The divergent charge term one gets whenever particles are treated as points. This is typically removed by renormalization. 17
- **color** The charge associated with the $SU(3)$ force of the Standard Model. 23
- **electric** The charge associated with the electromagnetic force. 5
- **weak hypercharge** The charge associated with the unbroken the $U(1)$ force of the Standard Model. 23
- **weak isospin** The charge associated with the unbroken $SU(2)$ force of the Standard Model. 23

charged lepton (ℓ) The leptons with electric charge ± 1. 21

CKM matrix The matrix that governs the flavor mixing properties of the weak force. 114

confinement The effect of the strong force that results in bare color charges becoming confined within hadrons. 19

cosmic ray High energy particles that travel for millions of years throughout the galaxy before reaching Earth. Collisions between cosmic rays and atmospheric particles produce secondary particles that can be observed from Earth's surface. 77

cosmological constant A free parameter in the Einstein field equations that is able to explain the apparent acceleration in the expansion of the universe. 37

cosmological density The ratio of an energy density to the critical density of the universe. 41

cosmological principle The principle that on large enough scales, the universe should be homogeneous and isotropic. 39

coupling Any kind of interaction between two or more fields, usually associated with a specific Lagrangian term. 18

- **constant** A constant that represents the strength of a coupling. 14
- **electromagnetic** (α) The coupling associated with the electromagnetic force. 5
- **strong** (α_s) The coupling associated with the strong force. 180
- **universal** Identical couplings for up and down type quarks. 34
- **Yukawa** A coupling to the Higgs boson that creates an apparent mass. 24

cross-section A quantity with the units of area that determines the probability of an interaction occurring. This is a generalization of the classical notion of cross-section, which includes the scattering of quantum fields. 16

cut A requirement that a measured quantity fall within some range. 87

dark energy A scalar field that accelerates the expansion of the universe, and behaves like the cosmological constant on large scales. 19

dark matter An unknown form of matter responsible for a number of different cosmological phenomena. 19

- **cold dark matter** Heavy dark matter, which would be moving at non-relativistic speeds today. Planets are a known form of cold dark matter. 53

hot dark matter Light dark matter, which would be moving near the speed of light today. Neutrinos are a known form of hot dark matter. 54

data-driven A method of estimating a background's contribution using real data. This tends to reduce the systematic uncertainty present in simulations, but also increases statistical uncertainty. 90

delta ray The production of a high energy electron by a primary particle ionizing an atom. 124

direct detection An experiment designed to detect local dark matter interactions. 2

dust An idealized fluid composed of massive particles that do not interact with each other. 42

electromagnetic shower The process by which high energy electrons, positrons, and photons are absorbed by dense materials. 81

electron (e^-) The stable, lightest of the charged leptons. 16

electroweak symmetry The symmetry that unifies the electromagnetic and weak forces, broken by the Higgs mechanism. 23

epoch A period of time in the evolution of the universe since the Big Bang. 44

- **dark** The period immediately following recombination in which there were no appreciable thermal interactions, creating a very quiet universe. 54
- **electroweak** The period during which the electromagnetic and weak forces were still unified, and particles were massless. 46
- **grand unification** The period during which the strong and electroweak forces were unified, described by some theory not yet discovered. 44
- **hadron** The period after hadronization became possible, followed quickly by the freeze-out of hadrons. The end of the period coincides with the freeze-out of neutrinos, thereby freezing the proton to neutron ratio. 51
- **lepton** The period during which electrons and positrons dominated the energy density of the universe, kept in thermal equilibrium by annihilations and pair productions. 52
- **matter** The period in which dark matter dominated the energy density of the universe. 52
- **modern** The period in which dark energy dominated the energy density of the universe. This epoch began 9.8 Gyr after the Big Bang, and is currently theorized to continue indefinitely. 55
- **photon** The period during which photons dominated the energy density of the universe. A brief period of nucleosynthesis occurred early on in this epoch. 52
- **Planck** The earliest epoch that is strongly believed to have occurred, where the gravitational force was as strong as the three others. A full theory of quantum gravity is necessary to even begin discussing this time. 44
- **quark** The period during which quarks and gluons were unable to form bound states, making the universe a quark-gluon plasma. 51
- **reionization** The period after the dark epoch, in which newly formed stars and galaxies ionized the neutral hydrogen that had been created during recombination. 54

Euclidean space The vector space that describes the three spatial dimensions in the absence of gravity. 8

event A recorded collision. 80

factorization The process of removing IR divergences due to soft and collinear gluon radiation. 181

fermion (f) Particles with half-integer spin, such as electrons and protons. Fermions obey the Pauli exclusion principle, in which only a single particle can exist in a given state. 13

Feynman diagram The graphical representation of a term in the perturbative expansion of the probability amplitude of an interaction. 13

Feynman rule A rule used as part of the algorithmic computation of Feynman diagrams. 13

field Any function over space-time coordinates. 10

fine-tuning Any instance of a physical dimensionless constant that is orders of magnitude away from unity. 47

flatness problem An example of fine-tuning observed in the measured value of the cosmological constant. 47

flavor A quantum number that describes the fermionic content of a system. There are 12 flavors in the Standard Model, corresponding to the 6 quarks and 6 leptons. These are conserved by the electromagnetic and strong forces, but they are each explicitly broken by the weak force. Certain combinations of flavor such as baryon and lepton number *are* conserved by the weak force. 21

- **heavy** The group of heavy fermions that are all very unstable: tau leptons, and charm, bottom, and top quarks. 96
- **lepton** Six quantities corresponding to the electron, muon, tau, and neutrino flavors. 26
- **light** The group of light fermions that are relatively stable: electrons, muons, and up, down, and strange quarks. 174
- **quark** Six quantities corresponding to the up, down, strange, charm, bottom, and top flavors. 27

fluid Idealized, non-interacting matter fields. 40

- **perfect fluid** Idealized, non-interacting matter fields. 40

force Any interaction mediated by a gauge boson. 19

- **electromagnetic** The long-ranged force mediated by the massless photon. 18
- **electroweak** The $U(1) \bigotimes SU(2)$ gauge force of the Standard Model, which is broken spontaneously by the Higgs potential. 22
- **gravitational** The hypothetical quantum description of gravity, mediated by gravitons. 19
- **strong** The $SU(3)$ gauge force of the Standard Model, which confines quarks within hadrons. 18, 22
- **weak** The short-ranged force mediated by the massive bosons obtained after electroweak symmetry breaking. 25

freeze-out The process by which a class of particles exit thermal equilibrium with the rest of the universe, leaving a thermal relic. 50

galaxy Gravitationally bound structures made out of stars and dark matter, usually containing a supermassive black hole at their center. 37

- **cluster** A typical galactic structure, made out of galaxies, gas, and dark matter. 55
- **dwarf** Faint galaxies with unusually high masses, expected to be due to large dark matter concentrations. 37
- **halo** The spherical region surrounding a galaxy in which there is very little visible matter. 59
- **structure** Enormous structures composed of gravitationally bound galaxies. 54

gauge A condition placed to fix unphysical degrees of freedom in a theory. The choice of gauge will never affect any observable prediction. 22

symmetry Any local symmetry of a field that introduces unphysical degrees of freedom. 21

transformation A transformation corresponding to a gauge symmetry, which leaves the equations of motion unchanged. 22

unitary The gauge in which there is a single real scalar Higgs field. 25

General Relativity Our current theory of gravity, described by the Einstein field equations. 37

generation A group of fermions with identical quantum numbers, but different mass. 21

gluon A massless gauge boson that mediates strong interactions. 22

gravitational lensing An effect predicted by General Relativity in which photons are deflected around massive objects. 57

micro-lensing Lensing effect observed around a large planet or small star closely orbiting a brighter object. 58

strong-lensing The extreme lensing of light around a compact object, which distorts a background source into an arc. 57

weak-lensing Many small, but resolvable lensing effects around massive objects that can be combined statistically. 57

graviton A hypothetical massless gauge boson that mediates the gravitational force. 22

hadron A composite particle made of quarks, held together by the strong force. 27

hadronic shower The process by which hadrons are absorbed by dense materials. 82

hadronization The effect of the strong force that results in bare color charges becoming confined within hadrons. 19

hard interaction High energy interactions, where QCD can be treated perturbatively. 88

Higgs mechanism The mechanism by which the electroweak symmetry is broken, and fermions gain mass. 21

indirect detection An experiment designed to detect visible byproducts of distant dark matter annihilations. 2

inflation A proposed solution to the flatness and isotropy problems, where the universe briefly underwent very rapid expansion shortly after the Big Bang. 48

inflaton The hypothetical field responsible for inflation. 48

isotropy problem The isotropy observed across causally disconnected regions of space. 47

jet (j) A collimated stream of particles produced during the hadronization of a bare color charge. 74

b**-jet** (b) A jet originating from a bottom quark. 94

c**-jet** (c) A jet originating from a charm quark. 96

τ**-jet** (τ) A jet originating from a hadronically decaying τ lepton. 99

light A jet originating from a gluon or an up, down, or strange quark. 96

soft Low energy jets that are unavoidable backgrounds in hadronic collisions. 119

jet smearing A method for predicting sources of $\not{E}_T$ in multijet events. 153

Lagrangian A functional that encodes the physical laws governing a system. 10

leptogenesis The period of time in the early universe in which an excess of leptons of anti-leptons was produced. 48

lepton A fermion that does not interaction through the strong force. 21

Lie algebra A vector space, with a Lie bracket operator, that corresponds directly to a Lie group. These are typically represented as matrices over the real or complex numbers, and the commutation operation. 11

Lie group A mathematical group that is a smooth differentiable manifold. 7

lumiblock A set of events at fixed detector and beam conditions. 121

luminosity The number of interactions per unit time per unit area. This quantity is used to quantify the amount of data delivered to the detector, and when multiplied by an interaction cross-section gives the expected number of interactions per unit time. 71

- **effective** The luminosity with which a sample should be compared, corrected for various effects that can reduce efficiency. 86
- **integrated** The total luminosity integrated over some period of time. 70
- **peak** The highest luminosity obtained in a run. 71

meson Hadrons consisting of one quark and one antiquark, such as pions and kaons. 51

Minkowski space The vector space that describes the four space-time dimensions in the absence of gravity. 7

Monte Carlo generator A program that randomly samples possible outcomes of an experiment. 88

muon (μ) A fairly common particle that decays to electrons, via the weak force, very slowly. 51

neutrino (ν) The electrically neutral leptons, which have no mass under the Standard Model, but very small mass experimentally. 19

neutron A composite particle made of two down quarks and an up quark, held together by the strong force. 21

neutron star The densest celestial object that has been observed, next to black holes. When the gravitational collapse is strong enough to overcome the electron degeneracy pressure of a star, proton and electron pairs within become neutrons. 57

non-perturbative A theory whose predictions can not be calculated using perturbation theory. 88

non-relativistic The low velocity limit, $v \ll c$, in which an expansion around $v \sim 0$ is utilized. 37

non-renormalizable A theory or interaction whose divergences can't be handled by renormalization. 17

nucleosynthesis The period of time in the early universe in which light atomic nuclei were produced through nuclear fusion. 52

orthogonal group The Lie group associated with vector rotations, $O(N)$, composed of all N-dimensional orthogonal matrices. 11

orthogonality Any matrix, O, satisfying $O^T O = 1$. 12

pair production The production of a particle/anti-particle pair, and the time reversal of annihilation. 49

parton The quarks and gluons that can be found within a hadron. 72

perigee The point on a track that is closest to the z-axis. 77

perturbation A small addition to a simple theory that represents an approximation of a more complex theory. 88

perturbative A theory that can be expanded around a small value to obtain increasingly accurate predictions. 13

photon (γ) The massless gauge boson that mediates electromagnetic interactions. 41

physical density The cosmological density of some type of matter multiplied by the reduced Hubble constant squared. 41

pileup Additional collisions that occur nearby the interaction of interest. 77

positron (e^+) The antiparticle of the electron. The separate name is purely for historic reasons, and in LHC collisions they are usually included under a more general definition of "electron". 51

prescale A suppression factor applied to triggers, which prevent it from firing too often. 86

primary vertex The interaction vertex from which the particles triggered on originated. 80

principle of relativity The postulate that all inertial frames are governed by the same laws of physics, from which the theory of relativity can be inferred. 26

probability amplitude A complex quantity used by quantum theories, whose square represents the probability of some event occurring. 13

propagator An internal particle line in a Feynman diagram. 13

proton A composite particle made of two up quarks and a down quark, held together by the strong force. 21

proton bunch A group of protons that is accelerated and collided together. 71

pseudo-data A dataset that uses a combination of both real and simulated objects or events. 144

pseudoscalar field A field with spin 0, that is invariant under proper, but not improper, coordinate transformations. 20

pseudovector field A field with spin 1, that is invariant under proper, but not improper, coordinate transformations. 20

quark (q) Fermions that interact via the strong force. 21

- **down-type** (d) All quarks with electric charge $\pm 1/3$. 21
  - **bottom** (b) Third generation down-type quark. 34
  - **down** (d) First generation down-type quark. 21
  - **strange** (s) Second generation down-type quark. 27
- **up-type** (u) All quarks with electric charge $\pm 2/3$. 21
  - **charm** (c) Second generation up-type quark. 95
  - **top** (t) Third generation up-type quark. 34
  - **up** (u) First generation up-type quark. 21

quark-gluon plasma A phase of matter in which quarks and gluons are no longer confined to hadrons. 51

quasar Very luminous, distant, and newly formed galaxies around a supermassive black hole. 54

radiation Any form of energy that escapes a system. 54

recombination The time at which neutral hydrogen was able to form, quickly leaving the universe electrically neutral. This is very closely related to the decoupling of photons from the rest of the

universe. 52

redshift A relativistic effect that decreases the energy of a particle. 41

redshift parameter A measure of the redshift of a distant body. 41

region A collection of cuts defining some category of event. 29

relativistic An approximation in which it is assumed that all velocities $v \sim c$, which typically allows one to neglect particle masses. 43

relic A particle or object created shortly after the Big Bang. 50

relic abundance A population produced shortly after the Big Bang that remains relatively stable today. 1

renormalizable A theory or interaction whose divergences can be handled by renormalization. 17

renormalization The process of removing UV divergences due to divergent loop diagrams. 16

run A continuous period of stable beams where data was collected. 71

scalar field A field with spin 0, that is invariant under coordinate transformations. 13

scale Some approximate range of energies, which can also refer to distances or times by using Planck units. 5

- **cosmological** Distances at which the cumulative effect of gravity becomes strong. 37
- **electroweak** The energy scale of electroweak interactions, roughly equal to the Higgs VEV of 246 GeV. 1
- **factorization** The energy scale at which factorization is applied. Physical observables should not depend on this. 181
- **galactic** Typical galaxies are on the order of 10,000 parsecs in diameter. 55
- **Planck** The energy scale of gravitational interactions, defined by the Planck mass, 10^{19} GeV. 44
- **renormalization** The energy scale at which renormalization is applied. Physical observables should not depend on this. 180

sea particle Transient quarks and gluons found within hadrons that only appear at high energies. 72

secondary vertex A vertex created by particles created after, but still as a result of, the primary interaction. 80

shower A stream of particles created within the calorimeters around a single high energy incident particle. 77

signal The process or processes that the search is optimized to find. 9

smear Applying a random factor to the energy or momentum of a particle. This is typically done to correct for resolution differences between measurements and predictions. 91

soft interaction Low energy interactions that are usually background noise to hard ones. 88

space-time The four dimensional vector space that unifies Euclidean space with time. 2

space-time metric A tensor field that encapsulates the geometry of space-time. 7

speed of light The maximum velocity information can travel at, equivalent to the speed at which light propagates. 5

spin-dependent A direct detection experiment that looks for dark matter interactions coupling to the total spin of each nuclei. 2

spin-independent A direct detection experiment where only the density of nucleons in the target material determines sensitivity. 2

spinor A vector in spinor space, defined by the half integer Lorentz group representations. 33

spontaneous symmetry breaking An effect that occurs when a continuous symmetry of the physical laws is not obeyed by the lowest energy solution. 25

spurion field A theoretical tool for extending a violated symmetry to a new theory. 32

stress-energy tensor The generalization of energy and momentum to continuous objects, rather than point particles. 11

supernova A stellar explosion that can occur for a variety of reasons 65

- **type Ia** A specific type of supernova caused by the disintegration of a white dwarf, which produces a very consistent signature. 65

systematic The uncertainty in a measurement due to the nature of the measuring device. Unlike statistical uncertainties, these usually can't be reduced by simply collecting more data. 131

tau lepton (τ) The heaviest lepton in the Standard Model, which can decay leptonically *or* hadronically via the weak force. 93

tensor field An object with integer spin greater than 1, that transforms as a proper tensor under coordinate transformations. 13

thermally averaged cross-section The center of mass cross-section of an interaction multiplied by the velocity of the initial state particles, averaged over the allowed microstates corresponding to a large system. 49

topocluster A topological cluster of energy deposits within the calorimeters. 113

trigger An algorithm that decides whether or not to record an event. 85

ultra-relativistic fluid Relativistic matter than can be approximated as being composed of massless, non-interacting particles. 42

unblinding The process of making the final measurements after the analysis has been frozen. 88

unitary Any matrix, U, satisfying $U^\dagger U = 1$. 12

unitary group The Lie group associated with spinor rotations, $U(N)$, composed of all N-dimensional unitary matrices. 11

universality Equal couplings to up and down-type quarks, in the context of this analysis. 34

valence quark The bound quarks within a hadron that determine its quantum numbers. 72

vector field A field with spin 1, that transforms as a rank-1 tensor under coordinate transformations. 13

vector space Any space that defines a set of vectors with some fixed dimensionality, which are closed under linear transformations. 7

vertex A point at which particle lines in a Feynman diagram meet, directly corresponding to a term in the Lagrangian. 13

white dwarf A stellar remnant that is not massive enough to continue collapsing, due to its electron degeneracy pressure. 65

Yukawa matrix The matrices governing the coupling of fermions to the Higgs field, ultimately determining their masses. 24

Chapter 1

Introduction

Over the last 80 years, an abundance of cosmological evidence has been observed that strongly supports the existence of *dark matter* (DM), the strongest of which can be found in [1–6]. Despite all of these observations though, very little is known about what this invisible substance is made of. The only properties that have been gathered from the cosmological data are that it's stable, cold, abundant, and interacts very weakly with ordinary *baryonic matter*. So far, no non-gravitational interactions of DM have been observed.

A number of possible explanations for the DM phenomena have been offered. *Massive compact halo objects* (MACHOs) and neutrinos are two possibilities that already fit nicely into the existing models. While both contribute to the observed DM density, MACHOs are too sparse [7] and neutrinos are too warm to account for all of it. *Modified Newtonian dynamics* (MOND), which adds an acceleration-dependent component to the gravitational force, is another option [8], but it has difficulties explaining the observed spatial distributions of DM [3]. A more promising candidate is axions, which are very light particles that solve the strong CP problem [9]. The original motivation for proposing these particles had nothing to do with DM, but it was later found that they also provide a viable explanation.

Although none of these theories have been absolutely ruled out, the DM hypothesis that will be examined here is the currently favored *weakly interacting massive particle* (WIMP) model. The reason for its popularity stems from two observations. The first is that the *relic abundance* of any *electroweak scale* particle left over from the *Big Bang* is predicted to be incredibly close to the observed density of DM in the universe. The second is that most theories *beyond the standard model* (BSM) predict the existence of various WIMPs,

many of which are stable. These typically appear naturally, for unrelated reasons, but new WIMPs can also be very easily added to any theory. As long as they have the right mass and interaction strengths, they will provide a suitable description of DM. Neither of these observations *prove* the WIMP hypothesis, but they're suggestive coincidences that make it the target of many searches.

There are 3 different ways to search for WIMP DM, illustrated in Fig. 1.1, corresponding to the unique rotations in *space-time* of the interaction between a pair of WIMPs and baryonic matter. The first of these, known as *indirect detection*, looks for distant WIMP-WIMP interactions that produce unique signatures of visible *Standard Model* (SM) particles [10–13]. The second method, *direct detection*, looks for local DM particles recoiling off atomic nuclei. Direct detection searches are typically focused on either *spin-independent* interactions [14–18] or *spin-dependent* interactions [19–21], depending on the specific target nuclei used by the detector. The final search method, which will be discussed thoroughly here, is the collider search. Experiments such as ATLAS attempt to observe the production of WIMPs from the high energy collisions of SM particles.

This thesis is divided into three sections. In the first two chapters, I will describe the theoretical context in which this analysis is set. The next three chapters will describe the high-level details of the analysis itself. Finally, each of the appendices is a self-contained section that goes into deeper detail on studies conducted by me over the course of this analysis. The analysis performed here is based on, but not identical to, the published results in [22]. My personal contributions to the analysis, which are explained more in the appendices, include the following,

- Writing code to produce relevant histograms out of the raw ATLAS data, which were subsequently used to perform the analysis in [22].

- Designing the algorithm to optimize the selection cutflow, which was used to define two of the four signal regions in the published analysis, described in Appendix A. The signal in [22] was a combination of $\mathrm{DM} + b$ and $\mathrm{DM} + t\bar{t}$, rather than just the $\mathrm{DM} + b$

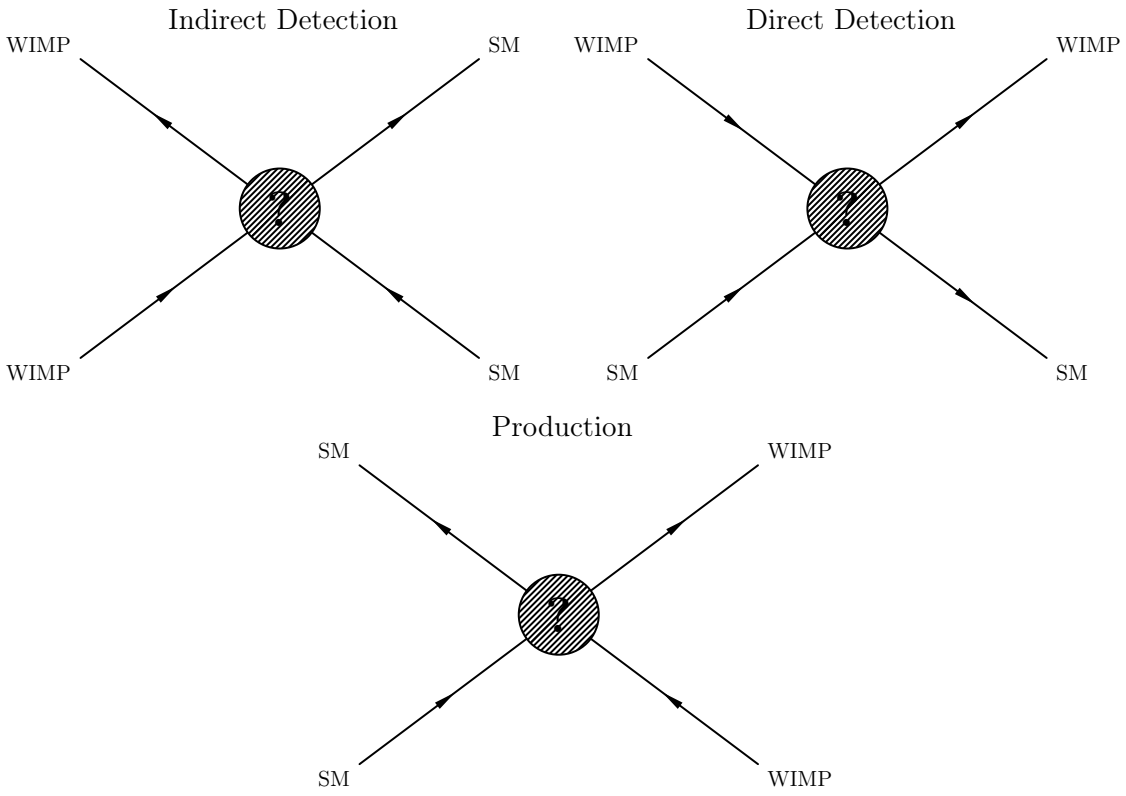

Figure 1.1: The Feynman diagrams representing the three types of WIMP search, with the time axis pointing to the right. Indirect detection involves WIMP-WIMP annihilation into visible SM particles, direct detection is the visible recoiling of a SM particle that has been hit by a WIMP, and production is the creating of WIMPs through the high energy collisions of SM particles. Each of these diagrams corresponds to the different rotations of the same process

signal focused on here.

- Implementing the hadronic and semileptonic DM + $t\bar{t}$ signal regions, SR3 and SR4. The results from these were not used here or by [22], but provided a useful cross-check of those reported by the DM + $t\bar{t}$ group.

- Implementing and validating the multijet estimate described in Appendix C, using the jet smearing method. A conservative 100% systematic uncertainty was used on this background by [22], as it had negligible contributions to the final results. More precise systematic uncertainties are calculated here, in Chapter 5.

- Implementing the data-driven $Z \to \nu\nu$ estimate using high energy γ-ray production, and expanding on work that had already been done using isolated $Z \to \mu\mu$ data. The latter method worked well at low energies, but suffered from very large statistical errors in the signal region. By combining the two methods we were able to get a very good estimate across a wide range of energy scales. This is explained further in Appendix B.

- Validating the $Z \to \nu\nu$ estimate, and calculating the flavor dependent transfer function. In the analysis, we used a phenomenologically derived function to fit the $\not{E}_T$-dependent transfer function for both the $Z \to \mu\mu$ and γ-ray methods. The fits were performed against data in a control region, meaning that the only systematic uncertainties originated from the fit parameters and the simulated backgrounds that were subtracted from the estimate. A much simpler, more robust method will be used here, where transfer factors are calculated using only *simulated* data, and are applied on a bin-by-bin basis for each histogram.

- Defining the control regions used to validate the background estimates. Three of these regions were used to fix the overall normalization of the top, $W \to \ell\nu$, and $Z \to \nu\nu$ backgrounds, which together made up over 95% of the background contributions in the signal region. I also defined control regions for the $Z \to \mu\mu$, γ + jets, and multijet

processes to calculate and validate the data-driven estimates used for the $Z \to \nu\nu$ and multijet backgrounds. In total, 8 of these control regions were used in [22].

- Computing the EFT validity regions to test the applicability of the results. This was done by collecting truth-level data on the simulated signal samples, and testing to see what fraction of events passed the validity requirements.

- Unless explicitly stated, any deviations from the analysis described in the public paper are a result of studies I personally performed.

Planck Units The Planck, or natural, unit system is used throughout this paper. It tends to be useful when dealing with any physics theory, as it eliminates the need for many of the physical constants that typically clutter equations. To illustrate the principle behind this system, a good example is the *speed of light*, c. From special relativity, we know that time and space are really just two manifestations of the same concept [23]. The fact that we measure time in seconds and space in meters, produces the arbitrary constant $c \approx 3 \times 10^8\,\mathrm{m\,s^{-1}}$. Analogously, if we were to measure along the x-axis in feet and the y-axis in meters, a new constant would be produced that relates the two. However, this constant would have no physical meaning and would simply represent our own ignorance. We can then remove these constants, by choosing a more natural system in which time and space are measured with the same units, and $c \to 1$. Similarly, we can also set Planck's constant $\hbar$, the Boltzmann constant k_B, and the Coulomb constant $(4\pi\epsilon_0)^{-1}$ to unity, eliminating units of length, temperature, and *electric charge* respectively, leaving only units of energy. Electric charge becomes a dimensionless quantity, with the fundamental electric charge e and the *fine structure constant*, α, related by $e^2 \equiv \alpha \approx 1/137$. In high energy physics, the most frequently used unit system is based on *electron volts* (eV), which are a measure of energy. When dealing with gravitational effects we can go one step further, by setting the gravitational constant $G \to 1$. This defines the *scale* of the *Planck energy* and makes all quantities dimensionless. All equations in this thesis assume Planck units, so the only

physical constants that will ever appear are the dimensionless ones that can't be eliminated from the theory. Explicit quantities will usually be expressed in eV for convenience, since the energy scale of this analysis is many orders of magnitude lower than the Planck energy. For reference, some common quantities are given in Tables 1.1 and 1.2, which can be used to convert to and from these unit systems.

Quantity	Planck	Formula	cgs
Planck length	L_P	$\sqrt{\frac{\hbar G}{c^3}}$	$1.6 \times 10^{-35}\,\mathrm{m}$
Planck time	t_P	$\sqrt{\frac{\hbar G}{c^5}}$	$5.4 \times 10^{-44}\,\mathrm{s}$
Planck energy	E_P	$\sqrt{\frac{\hbar c^5}{G}}$	$1.2 \times 10^{28}\,\mathrm{eV}$
Planck mass	M_P	$\sqrt{\frac{\hbar c}{G}}$	$2.2 \times 10^{-5}\,\mathrm{g}$
Planck temperature	T_P	$\sqrt{\frac{\hbar c^5}{G k_B^2}}$	$1.4 \times 10^{32}\,\mathrm{K}$
Planck charge	Q_P	$\sqrt{4\pi\epsilon_0 \hbar c}$	$11.7e$

Table 1.1: Definitions of Planck length, mass, and time.

Quantity	cgs	eV
Length	$1\,\mathrm{m}$	$5.1 \times 10^{6}\,\mathrm{eV}^{-1}$
Time	$1\,\mathrm{s}$	$1.5 \times 10^{15}\,\mathrm{eV}^{-1}$
Energy	$1\,\mathrm{J}$	$6.2 \times 10^{18}\,\mathrm{eV}$
Mass	$1\,\mathrm{kg}$	$5.6 \times 10^{35}\,\mathrm{eV}$
Temperature	$1\,\mathrm{K}$	$2.2 \times 10^{-5}\,\mathrm{eV}$
Charge	e	0.085

Table 1.2: Common conversions between cgs units and eV units. These can all by readily obtained from Table 1.1, by manipulating the appropriate Planck quantities.

Einstein Notation Einstein, or tensor, notation is typically used in any relativistic theory, as it makes complicated *tensor* equations significantly more readable. Tensors are

generalizations of vectors, where rank-0 and rank-1 tensors correspond to scalar values and vectors respectively. A rank-n tensor, T, can be represented by an n-dimensional matrix, and is labeled by n indices, $T^{i_1 i_2 \ldots i_n}$. As with matrices, every element in T is denoted by a unique combination of values for the indices $i_1 i_2 \ldots i_n$, which can take on a range of values either explicitly stated or inferred. Here, a convention will be used where tensor indices labeled by Roman letters $(i, j, k, \ldots)$ will indicate the spatial indices $1, 2, 3$ (i.e. x, y, z), and Greek letters $(\mu, \nu, \sigma, \ldots)$ will indicate space-time indices, $0, 1, 2, 3$ (i.e. t, x, y, z). Additionally, the Roman letters $(a, b, c, \ldots)$ will be reserved for *Lie group* indices, which will be discussed briefly in Chapter 2. Variables with superscript (subscript) indices, such as x^μ (x_μ), are known as *contravariant* (*covariant*) tensors. Tensors with Roman indices are called *3-tensors*, and those with Greek indices are called *4-tensors*, corresponding to the *vector spaces* they're defined on (space and space-time respectively). The key rule of tensor notation is an implicit summation over repeated indices. Unless explicitly stated, if an index is repeated over the product of a covariant and contravariant tensor, there is an implicit sum over that index. For example,

$$x^\mu y_\mu \rightarrow \sum_{\mu=0}^{3} x^\mu y_\mu. \tag{1.1}$$

More simply, an index is summed over if it appears as both a superscript and a subscript in a single product term, including the trivial case of a single high rank tensor (e.g. $T^\mu_\mu = \text{Trace}(T)$). Taking the tensor product of a *4-vector* with itself gives the *4-length*, denoted $x^2 = x^\mu x_\mu$. The covariant and contravariant *metric tensors*, $g_{\mu\nu}$ and $g^{\mu\nu}$, are used to "lower" and "raise" indices respectively, in the sense that $x_\mu \equiv g_{\mu\nu} x^\nu$. The metric tensor defines the geometry of space-time, and satisfies the identities $g^{\mu\nu} g_{\mu\sigma} = \delta^\nu_\sigma$ and $g_{\mu\nu} = g_{\nu\mu}$, where δ^ν_σ is the Kronecker delta function. In the absence of gravity, which distorts the metric into more complex forms, space-time is *flat*, meaning that it's described by a *Minkowski manifold* defined by the Minkowski metric $\eta_{\mu\nu} \equiv \text{diag}(1, -1, -1, -1)$. This can be compared to the metric of a flat *Euclidean* space, which would be $\text{diag}(1, 1, 1)$.

There is a slight complication to the tensor rule when it comes to derivatives. The derivative $\partial/\partial x^\mu$ is actually a covariant tensor, despite having an upper index on the x. A simple example to illustrate why is,

$$\frac{\partial}{\partial x^\mu} x^2 = \eta_{\sigma\rho} \frac{\partial}{\partial x^\mu} x^\sigma x^\rho = \eta_{\sigma\rho} \left(\delta^\sigma_\mu x^\rho + \delta^\rho_\mu x^\sigma \right) = 2\eta_{\mu\nu} x^\nu = 2x_\mu. \tag{1.2}$$

The term x^2 is a Lorentz scalar, so the derivative should be the same type of tensor as the right side of the equation, which is covariant. For this reason, and for simplicity, space-time derivatives are usually written as ∂_μ, shorthand for $\partial/\partial x^\mu$, so that the simple summation rule still works. More generally, a superscript in the denominator of any derivative should be treated as a subscript on the entire derivative, and vice versa.

Chapter 2

Theoretical Background

In any experimental search for *new physics* (NP), the hope is to observe a deviation from the expected outcome. In order to distinguish between the expectation and a deviation, a theoretical model of the experiment is required, with minimal uncertainties. Without one, it would be impossible to draw any conclusion from observations. In general we need a null hypothesis, or *background model*, and alternative hypotheses to give falsifiable predictions. The typical alternative hypotheses are either specific *signal models*, or generic statistically significant deviations from the null hypothesis, in the case of model-independent searches. Once the data is collected from the experiment, the hypotheses can be tested and conclusions can be drawn. If the data is inconsistent with the background model by a significant amount, a discovery has been made. Otherwise, constraints can be set on the alternative hypotheses, based on how consistent they are with the observed data.

2.1 Basic Concepts

Fully understanding the theory behind this search requires a lot of mathematical concepts that go well beyond the main focus of this thesis. This section contains a very brief summary of those that are most relevant to this analysis. For more in depth explanations or derivations, the reader is directed to [24–29], which cover these subjects in detail.

2.1.1 Lagrangian Formalism

The *Lagrangian* formalism is an incredibly powerful tool used to derive the equations of motion of complex systems evolving in time. The core concept behind the formalism is Hamilton's principle, which states that the *action*, S, of any system is *stationary*, meaning $\delta S = 0$. Classically, the action is the time integral of a functional with the form $L\{q_i(t), \dot{q}_i(t); t\}$, called the Lagrangian, which is defined over the dynamical parameters of the system, $\{q_i\}$, and their time derivatives, $\{\dot{q}_i\}$. When dealing with relativistic fields, the concept is extended to a Lagrangian density, $\mathscr{L}$, defined by,

$$S \equiv \int \mathscr{L}\{\phi_i(x), \partial_\mu \phi_i(x); x\} d^4x, \tag{2.1}$$

where each ϕ_i is a *field* over space-time coordinates. Applying Hamilton's principle to this action results in the Euler-Lagrange equations of motion for each field ϕ_i,

$$\frac{\partial \mathscr{L}}{\partial \phi_i} = \partial_\mu \left(\frac{\partial \mathscr{L}}{\partial(\partial_\mu \phi_i)} \right), \tag{2.2}$$

which define the dynamical properties of interacting fields.

2.1.2 Noether's Theorem

Since the earliest physical theories, symmetries and conservation laws have been core concepts in physics. Newton's first law (*a particle in motion stays in motion*) is a clear demonstration of the conservation of momentum. Newton's third law (*for every action there is an equal and opposite reaction*) points out a fundamental symmetry of nature, and Newton's law of gravitation takes full advantage of the spherical symmetry of celestial bodies. It wasn't until 1918 though, that the intimate connection between these two concepts was discovered. Noether's theorem [30] revealed that every *continuous* symmetry of a system is necessarily linked to a single conserved quantity. A symmetry of a system is formally defined as a transformation of either the dynamical or dependent variables of the system

that leaves the equations of motion unchanged. A very simple example of this is when the Lagrangian of a system does not depend *explicitly* on time, meaning that the transformation $t \to t + \delta t$ leaves it unchanged. It's straightforward to show that this symmetry leads directly to the definition of the Hamiltonian (or *stress-energy tensor* in the relativistic analog), which simply represents the conserved energy of the system.

In light of this, the underlying symmetries of any theory become very important for obtaining a deeper understanding, since they can be interpreted as an explanation for any conserved *charges*. For example, electric charge conservation is associated with gauge invariance, angular momentum conservation with rotational symmetry, and momentum conservation with translational symmetry. More generally, the fundamental quantities of any theory are those that are conserved, and each one arises from a symmetry within the theory. Note that symmetries don't need to be exact to make use of Noether's theorem. A symmetry that *almost* leaves the theory unchanged will produce a charge that is *nearly* conserved. Approximate symmetries are used wherever possible, because of all the advantages they provide.

2.1.3 Lie Groups

There is a very fundamental connection between physical symmetries and the abstract mathematical constructs known as Lie groups. Many of the symmetries observed in nature, especially those related to particle physics, obey the same relations as various Lie groups, and can be conveniently described by vector spaces known as *Lie algebras*. An algebra, in this case, is defined as a, frequently infinite, set of matrices that is closed under both matrix multiplication and summation. Lie algebras can be more compactly defined by some number, N, of *group generators* that create an homomorphic mapping between an N-dimensional vector space and the elements of the algebra. For every element in the algebra, M, there exists a unique vector $\boldsymbol{\alpha}$ such that, $M = e^{i\boldsymbol{\alpha}\cdot\boldsymbol{T}}$, where $\boldsymbol{T}$ is the set of generators.

The two most important families of Lie groups used in particle physics are the *orthogonal* ($O(n)$) and *unitary* ($U(n)$) groups, which each describe different types of rotations. $O(n)$ is

defined as the group represented by the $n \times n$ real matrices O that satisfy the *orthogonality* requirement, $O^T O = 1$, where O^T is the transpose of O. This group can be visualized as the group of rotational and inverse transformations in an n dimensional Euclidean space. The $SO(n)$ group is the subgroup of $O(n)$ whose matrices have determinant equal to 1, which is simply the set of proper rotations. A generalization of the orthogonal groups $O(n, m)$ can be made where the orthogonality condition becomes,

$$O^T O = \begin{pmatrix} \mathbf{1}_n & 0 \\ 0 & -\mathbf{1}_m \end{pmatrix}, \tag{2.3}$$

for the $i \times i$ unit matrices $\mathbf{1}_i$. One of the most frequently used orthogonal groups is $O(1, 3)$, which describes Lorentz transformations in 4-dimensional space-time.

The $U(n)$ groups are represented by the $n \times n$ complex matrices U that satisfy the *unitary* requirement, $U^\dagger U = 1$, where $U^\dagger$ is the Hermitian conjugate of U, $U^\dagger \equiv (U^T)^*$. The $SU(n)$ group is the subgroup of $U(n)$ composed of matrices with determinant equal to 1. These groups are closely related to the orthogonal groups, and can similarly be thought of as types of rotations and inversions in a *complex* vector space. Examples of the similarities are the homomorphisms, $\text{spin}(3) \equiv SU(2) \to SO(3)$ and $\text{spin}(4) \equiv SU(2) \otimes SU(2) \to SO(1, 3)$, where in both cases the spin groups are *dual covers* of the corresponding orthogonal groups, with two U matrices homomorphic to each O. The spin groups are frequently used in quantum mechanics since, as the name suggests, they describe the phenomenon of quantum spin. The relationship between $SU(2)$ and $SO(3)$ is analogous to the relationship between spin and angular momentum, which are very similar but not the same.

The $SU(n)$ groups are especially important in particle physics, because they describe most of the approximate and exact symmetries of the SM, which will be discussed later. Each group $SU(n)$ is generated by $n - 1$ matrices $\boldsymbol{T}$ satisfying,

$$T_a T_b = \frac{1}{2n}\delta_{ab} + \frac{1}{2}\left(i f_{abc} + d_{abc}\right) T^c. \tag{2.4}$$

Since $\det(M) = 1$ and $M^\dagger M = \mathbf{1}$ for each element M in the group, it's straightforward to show that each generator T_a must be an $n \times n$ traceless Hermitian matrix. The structure constants d_{abc} and f_{abc} are completely symmetric and antisymmetric respectively, over their indices, and are unique to each group.

2.1.4 Quantum Field Theories

The framework on top of which the SM, and most BSM theories, are built is known as *quantum field theory* (QFT) [26]. Any QFT, up to adjustable parameters not fixed by the theory, is uniquely determined by its Lagrangian density $\mathscr{L}$, which encodes all of the information necessary to make experimental predictions. Unlike classical fields, quantum fields are quantum *operators* that act on the vacuum to create or annihilate an integer number of field excitations, known as *quanta* or *particles*, and in general don't commute with each other. There are many different types of fields, distinguished by their associated Lorentz group representation, or *spin*. *Scalar*, *fermion*, *vector*, and *tensor* fields correspond to spin 0, 1/2, 1, and 2 particles respectively. Half-integer spin particles are known as fermions, and integer spin particles are known as *bosons*. The main difference between these two types of particles is their behavior under symmetry transformations.

Although most of the details behind QFT are well beyond the scope of this thesis, *Feynman diagrams* [31] provide a very useful tool for visualizing field interactions. Feynman diagrams and *Feynman rules* give an algorithmic solution to the very complex combinatoric problem of calculating the *probability amplitude* of an interaction. They produce an infinite series, but are very often *perturbative*, meaning that the contribution of each additional term falls off sufficiently fast that they can be truncated to arbitrary precision. Feynman diagrams allow us to iterate through all of the terms, and Feynman rules allow us to calculate the value of each diagram. The total probability of an interaction is given by the absolute square of the sum of all diagrams with identical initial and final states.

Feynman diagrams consist of *propagators*, interaction *vertices*, and external particles. External particle lines represent the initial and final state particles, with the time axis

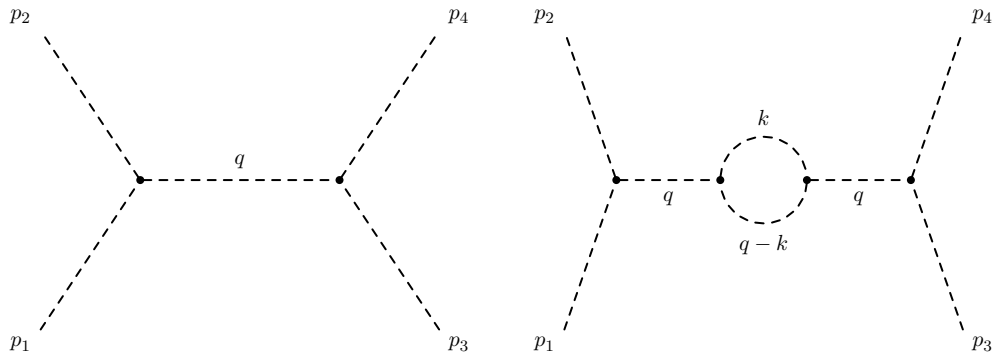

Figure 2.1: Two examples of Feynman diagrams. On the left is the leading order term of an interaction in the ϕ^3 theory, and on the right is a 2nd order term. In both cases $q = p_1 + p_2 = p_3 + p_4$ to conserve 4-momentum. In the right diagram though, k can take on any value and therefore must be integrated over. In general q^2 and k^2 don't have to be equal to the masses of their respective fields.

pointing to the right, so that incoming particles are on the left and outgoing particles are on the right. Internal lines are called propagators, which are commonly referred to as virtual, or intermediary particles. The ends of every line terminate at a vertex, except for the outer ends of the external lines. Arrows on lines are used to indicate the flow of charge, with forward and backward pointing arrows representing particles and anti-particles respectively. Both *4-momentum* and charge must be conserved at each vertex, and only vertices directly corresponding to interaction terms in the Lagrangian are allowed. For any given initial and final state, the amplitude is then just sum of all allowed Feynman diagrams.

As an example, consider the very simple ϕ^3 QFT given by,

$$\mathscr{L} = -\frac{1}{2}\partial^\mu\phi\partial_\mu\phi - \frac{1}{2}m^2\phi^2 + \frac{1}{6}g\phi^3. \tag{2.5}$$

This is a theory with a single scalar particle with mass m and self *coupling constant* g. Figure 2.1 shows an example of two Feynman diagrams that would contribute to the scattering of two particles in this theory. Calculating the value of a Feynman diagram is done by multiplying the contributions of each line and vertex, using the associated Feynman rules.

The Feynman rules for the ϕ^3 QFT are defined as,

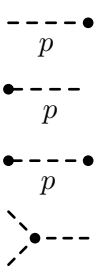

$$= 1,$$

$$= 1,$$

$$= \frac{-i}{p^2 - m^2},$$

$$= ig. \tag{2.6}$$

In this very simple theory, external particle lines contribute only a factor of 1 to the amplitude, although in less trivial theories these contributions can become quite complicated. The vertices and propagators usually take on fairly simple values whose form can be gathered from the Lagrangian. Each vertex introduces a factor of the relevant coupling constant, and each propagator with momentum p introduces a factor proportional to $(p^2 - m^2)^{-1}$ for bosons and $(\not{p} - m)^{-1}$ for fermions.

The total 4-momentum going in to any vertex in a Feynman diagram must sum to 0, in order to obey conservation laws. This regularly leads to propagators where $p^2 \neq m^2$, which are known as *off-shell*, or *virtual*, particles. This effect is loosely related to the Heisenberg uncertainty principle, where at short distances and times there is an intrinsic uncertainty in the 4-momentum of a particle. It creates a *resonant* behavior where if the total 4-momentum of the incoming particles is near the mass of the intermediate field, their interaction is greatly enhanced. Below the resonance, the amplitude is roughly proportional to m^{-2} or m^{-1} depending on the type of field. In contrast, external particles always have $p^2 = m^2$ and are known as *on-shell*, or *real*, particles. Complex diagrams can contain propagators that are under-constrained by 4-momentum conservation, and are known as *loop diagrams*. For each loop, the diagram is integrated over all allowed 4-momenta.

The important point to take away from this discussion is that each term in the Lagrangian of a QFT corresponds to a vertex or propagator with its own Feynman rule. The sum of all possible Feynman diagrams gives the amplitude for a specific interaction, whose

absolute square is the probability of that interaction occurring. By integrating over different possible final states, the *cross-section* of the collision between initial state particles can be found. This cross-section is one of the most frequently used observable quantities in high energy physics, and its dependence on particle masses and coupling constants can be easily obtained using these methods.

Mathematically speaking, Feynman diagrams represent terms in a perturbative series, and are simply a useful tool for keeping track of a complicated combinatoric problem. In order to extract useful quantities with this method, the series must decrease fast enough that the lower order terms are a good approximation. For interactions with coupling constant $g < 4\pi$, the amplitude can be perturbatively expanded in factors of $g/4\pi$. Likewise, for theories with very massive particles, expanding the amplitude by factors of the inverse mass M^{-1} produces a different perturbative series useful for *effective field theorys* (EFTs). Calculations that only involve terms up to $\mathcal{O}(g^2)$ are known as *leading order* (LO), and calculations that go to $\mathcal{O}(g^4)$ or $\mathcal{O}(g^6)$ are called *next to leading order* (NLO) and *next to next to leading order* (NNLO) respectively.

Renormalization The calculated amplitude of an interaction can frequently diverge, for a variety of reasons, and must be corrected using a process known as *renormalization*. To illustrate this, a simple example occurs in classical electromagnetism, where a charged particle's self-energy, E_{self}, diverges as its radius goes to 0. This would imply that all point particles have infinite mass, which is in direct conflict with experience. Although we can't demonstrate a particle has 0 radius, it has been shown that the *electron* has no discernable substructure past 3×10^{-19} m [32], which would give it a self-energy over 10,000 times its measured mass. One way to get around this is to treat the particle as a shell with a *bare mass* of $m - E_{\text{self}}$, so that the *total* mass of the particle remains m at any radius. This classical example is a somewhat awkward way to make the theory work, since what it's telling us is that charged point particles have a divergent negative mass, which is only canceled out by the divergent self-energy from its charge. However, it manages to keep the

theory self-consistent while still accurately describing observations.

In QFT, loop diagrams introduce integrals over arbitrarily high momenta. Just as in classical physics, problems occur at small distance or high energies. Divergences appear that need to be dealt with in order to calculate these higher order loop diagrams. Although at first glance it may look like an *ad hoc* procedure to fix theoretical shortcomings, renormalization in QFT has become a rich topic and is closely related to the concept of *scale invariance*. The process of renormalization mirrors the classical analogue by adding counter terms to the Lagrangian to balance out divergences. All of these terms are connected to parameters of the theory, such as the masses or charges, which generate divergences corresponding to bare masses and *bare charges*. If all of the divergences of a QFT can be renormalized using a finite number of additional terms, the theory is called *renormalizable*.

Interactions that require an infinite number of counter terms in the Lagrangian to remove all of their divergences are labeled *non-renormalizable*. A non-renormalizable QFT is one containing *any* non-renormalizable interaction. In general, any Lagrangian term with a coupling constant of negative energy dimension is non-renormalizable, and any term with a 0 or positive dimensional coupling constant is renormalizable. The action is dimensionless by definition, so that by basic dimensional analysis the Lagrangian density must have dimension d, where d is the number of space-time dimensions in the theory. The dimension of a given field can be determined by its kinetic term, which has no coupling constant by definition, using the fact that space-time derivatives have dimension 1. In 4-dimensional space-time, for example, bosonic fields have dimension 1 and fermionic fields have dimension $3/2$. Therefore any interaction between three or more fermions, or 5 or more bosons is non-renormalizable.

Non-renormalizability isn't necessarily a bad thing, and it will actually be taken advantage of later on. The general interpretation of a non-renormalizable term with coupling constant Λ^{-n} is that it's the nth-order term in an approximation of a higher energy theory parameterized by the energy scale Λ. This is because when a heavy particle of mass Λ is integrated out of a theory, it generates an infinite number of Lagrangian terms suppressed

by increasing powers of Λ^{-1}. With this interpretation, an alternative to renormalization is possible for dealing with divergent momenta integrals. Since the theory is viewed as a low energy approximation to some unknown higher energy theory, it can't be expected to work for energies above Λ. The integrals can be *cut-off* at Λ and calculations can be carried out by perturbatively expanding around Λ^{-1}. This produces an EFT, and can be very useful when the higher energy theory is unknown or difficult to perform calculations with. The SM is often viewed as an EFT for some unknown underlying theory, with an energy scale $\Lambda_{\text{NP}} = \mathcal{O}(1)$ TeV. Under this interpretation, the reason we haven't observed any fundamental non-renormalizable interactions is simply because we haven't reached high enough energies to see those higher order terms.

Running Couplings A concept very closely related to renormalization is the *running* of coupling constants in QFT. The process of renormalizing an interaction term introduces an energy scale dependence to its *coupling strength*. The beta function for an interaction,

$$\beta(g) \equiv \frac{\partial g}{\partial \log(\mu)}, \tag{2.7}$$

determines the dependence of its coupling strength, $g(\mu)$, on the energy scale μ. Essentially, the charge of a particle under any interaction changes as it's probed at different distances. The fine structure constant, for example, determines the strength of the *electromagnetic force* and the fundamental unit of electric charge. Although it approaches the asymptotic value of $\alpha \approx 1/137$ at low energies, it diverges at higher energies. A common method for visualizing this effect is to imagine virtual particle-antiparticle pairs being created and annihilating very close to an infinite bare charge. These create a screening effect where at large distances, we only see the asymptotic charges. As smaller distances are probed, the screening effect becomes less powerful and more of the bare charge is "visible". This effect has been observed in both the strong and electromagnetic forces, and is a well established phenomenon [33].

Unlike the electromagnetic force, strongly coupled interactions display a phenomenon known as *asymptotic freedom*. Their coupling strength disappears at high energies, but becomes divergent at low energies. The result of this is *confinement*, where isolated charges are not stable and quickly *hadronize* into chargeless states. As a concrete example, the strengths of the three SM *forces* are plotted in Fig. 2.2, showing the running of their couplings.

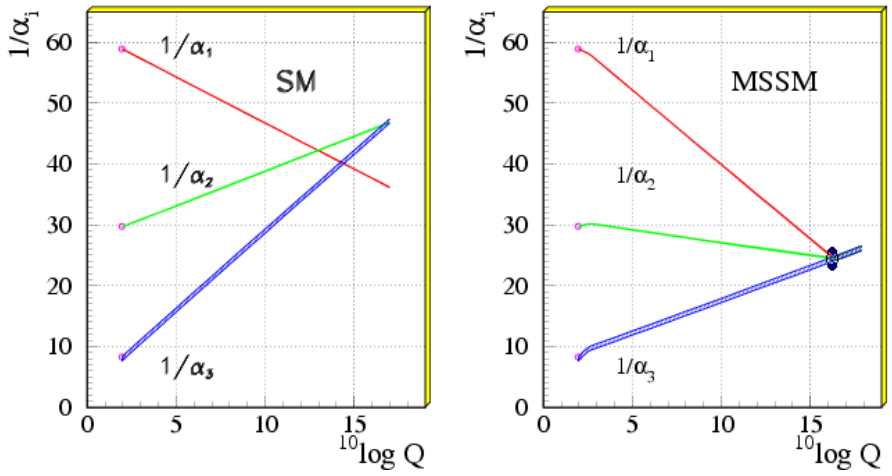

Figure 2.2: The running of the couplings of the three quantum forces. On the left are the SM couplings, and on the right are the same couplings under supersymmetry. The unification of all three forces at the GUT scale is one of many reasons why supersymmetry is popular. [34]

2.2 The Standard Model

The SM is currently the most accurate theory of subatomic physics we have, and has been probed to extraordinary precision. Despite this, there are a number of theoretical and experimental observations that suggest it's incomplete. The most notable of these are that the SM doesn't provide descriptions of the *gravitational force*, *dark matter*, *dark energy*, or *neutrino masses*. A very useful graphical representation of the SM is provided in Fig. 2.3, which compactly describes the different particles and interactions that will be discussed in this section.

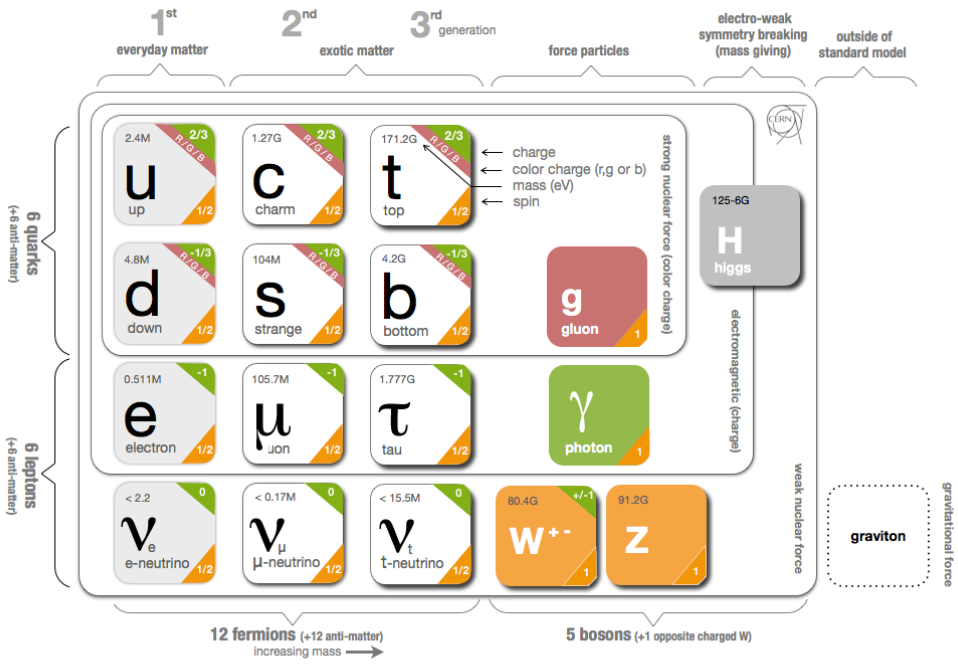

Figure 2.3: A graphical representation of the SM particles and interactions. [35]

2.2.1 Fields

The SM fields are ultimately what's observed in any experiment, as they make up all the matter that we interact with. There are 3 classes of fields in the SM, labeled scalar, fermion, and vector fields. These correspond to spin 1/2, 1, and 0 respectively, which are distinguished by how they behave under space-time rotations. Scalar fields take on a single value at every point in space-time, which stays constant under rotations. Vector and fermion fields, on the other hand, have 4 values at each point. While vector fields transform under $SO(3,1)$ as 4-vectors, fermion fields transform under the homomorphic spin group $SU(2) \bigotimes SU(2)$. The 4 components of a fermion field, in the chiral basis, correspond to the left/right parity components of a particle and its anti-particle. There are no tensor, *pseudoscalar*, or *pseudovector* fields in the SM, but they're frequently found in BSM theories.

Fermions There are 12 fermions in the SM, which can be organized by their charges into four families: *up-type quarks*, *down-type quarks*, *charged leptons*, and neutrinos. Within each family there are three fields with identical quantum numbers, distinguished by their *generation* and differing only in mass and *flavor*. The lightest fields in the two *quark* families are known as *up quarks* and *down quarks*, and are the constituents of *protons* and *neutrons*, which account for nearly all the visible matter in the universe. The lightest *leptonic fields* are the neutrinos, which are massless, and the electron, which is the first generation charged lepton. The heavier generations have higher masses and are all unstable, so they will eventually decay into combinations of these 6 stable fermions.

The kinematic Lagrangian term of a fermion field comes from the Dirac equation [29], and is given by,

$$\mathscr{L} \supseteq \bar{\psi}\left(i\not{\partial} - m\right)\psi, \tag{2.8}$$

where $\not{\partial}$ is Feynman slash notation for $\gamma^\mu\partial_\mu$ and $\bar{\psi} \equiv \psi^\dagger\gamma^0$. The four γ matrices are 4×4 complex matrices, with the additional $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$, defined by a number of identities derived from Dirac equation. They can take on different forms, and in the chiral basis,

$$\gamma^\mu \equiv \begin{pmatrix} 0 & \sigma^\mu \\ \bar{\sigma}^\mu & 0 \end{pmatrix}, \sigma^\mu \equiv (1, \sigma^i), \bar{\sigma}^\mu \equiv (1, -\sigma^i), \tag{2.9}$$

where the σ^i terms are the 2×2 Pauli matrices that generate the $SU(2)$ group. This will be the most useful here, as it splits the fields into left-handed and right-handed components. Note that in the SM, the mass terms of the Lagrangian are fixed to 0 because they explicitly break one of its symmetries. Instead, fermion masses are acquired dynamically through the *Higgs mechanism*, as discussed later.

Gauge Bosons *Gauge bosons* are quantum excitations of bosonic fields with an associated *gauge symmetry*, that mediate the fundamental forces. In the SM there are three of these,

corresponding to the *electroweak force* (B and W fields), and the *strong force* (*gluons*). The gravitational force is mediated by a tensor field that's invariant under differentiable coordinate transformations, so it's generally expected that this generates spin 2 particles called *gravitons*. However, gravity isn't included in the SM, and all of the gauge bosons considered here are spin 1 fields with $U(n)$ gauge symmetries.

Each gauge boson is represented by vector fields A^a_μ, where a runs over the associated Lie group's indices. There is 1 B field, 3 W fields, and 8 gluons corresponding to the number of generators for the groups $U(1)$, $SU(2)$, and $SU(3)$ respectively. Since the generators of these Lie algebras are orthogonal to each other under the trace operation, the compact notation,

$$\mathbf{A}_\mu \equiv A^a_\mu \mathbf{T}_a, \tag{2.10}$$

can be used to express each set of gauge fields as a single matrix. A *gauge transformation* is a local transformation of the form,

$$\boldsymbol{A}_\mu \rightarrow \boldsymbol{A}_\mu + \frac{i}{g} U \partial_\mu U^\dagger, \tag{2.11}$$

$$\psi \rightarrow U\psi, \tag{2.12}$$

where U is an element from the respective algebra, g is the charge operator, and ψ is a group of charged fermions. A *gauge-covariant* derivative and field tensor can be defined as,

$$\boldsymbol{D}_\mu \equiv \partial_\mu - ig\mathbf{A}_\mu, \tag{2.13}$$

$$\mathbf{F}_{\mu\nu} \equiv \frac{i}{g}\left[\boldsymbol{D}_\mu, \boldsymbol{D}_\nu\right], \tag{2.14}$$

which, under gauge transformations transform simply as,

$$\boldsymbol{D}_\mu \rightarrow U \boldsymbol{D}_\mu U^\dagger, \tag{2.15}$$

$$\boldsymbol{F}_{\mu\nu} \to U\boldsymbol{F}_{\mu\nu}U^{\dagger}. \tag{2.16}$$

By replacing all the partial derivatives of fermion fields with the covariant derivative, gauge invariance is automatically enforced. It's straightforward to show that any $SU(n)$ group, because its generators are non-commutative, can't support the covariant derivatives of fermion fields unless everything it couples to, including itself, has identical charge g. For the Abelian $U(1)$ group this restriction doesn't exist, which is why particles with different electric charge exist. The charges corresponding to the $U(1)$ and $SU(2)$ gauge symmetries are known as *weak hypercharge* (Y_W) and *weak isospin* ($\boldsymbol{T}$) respectively. Right-handed fermions are isospin singlets, meaning they don't couple to the W fields, while left-handed fermions form six isospin doublets corresponding to each lepton and quark generation. Neutrinos and down-type quarks have $T_3 = +1/2$, while charged leptons and up-type quarks have $T_3 = -1/2$. After *electroweak symmetry breaking*, right-handed fields gain electric charge equal to their weak hypercharge, while left-handed fields gain electric charge $Q = T_3 + Y_W/2$. Only quarks are charged under the $SU(3)$ gauge symmetry, which corresponds to their *color charge*.

The kinematic Lagrangian term for a gauge boson is,

$$\mathscr{L} \supseteq -\frac{1}{4}F^a_{\mu\nu}F_a^{\mu\nu} - \frac{1}{2}m^2A^a_\mu A^\mu_a, \tag{2.17}$$

where m is its mass. As with fermions, the mass term explicitly violates the gauge symmetry of the field, and is excluded from the SM. All of the SM gauge fields are massless and, like the fermions, they gain mass dynamically through the Higgs mechanism. This mixes the B and $\boldsymbol{W}$ fields to create the massless photon, the massive charged $W^{\pm}$ *bosons*, and the massive neutral Z^0 *boson*.

Higgs Boson The Higgs mechanism [36] adds two complex scalar fields to the SM Lagrangian, which are the only massive fields in the SM. The two fields form an $SU(2)$

doublet with weak hypercharge +1. Their kinematic Lagrangian terms, which come from the Klein-Gordon equation, are given by,

$$\mathscr{L} \supseteq \left(\boldsymbol{D}_\mu \phi\right)^\dagger \boldsymbol{D}^\mu \phi, \tag{2.18}$$

where ϕ is a vector in the $SU(2)$ space with elements corresponding to the two complex scalar fields, ϕ^+ and ϕ^0. This term defines the interactions between the *weak bosons* and the *Higgs*, through the covariant derivative operator. The key to the Higgs mechanism lies in the Higgs potential,

$$\mathscr{L} \supseteq -\frac{1}{4}\lambda^2 \left(|\phi|^2 - v^2\right)^2, \tag{2.19}$$

which defines its mass, and self-interaction. The parameters v and λ are known as the *vacuum expectation value* (VEV) and the self coupling constant respectively. The mass of the Higgs comes from the $|\phi|^2$ term, which has a coefficient corresponding to a mass of $m_H = v\lambda$. The fermion mass terms are replaced with the *Yukawa couplings*,

$$\mathscr{L} \supseteq -\frac{1}{v}(\bar{u}_L, \bar{d}_L)\phi \boldsymbol{M}^d d_R - \frac{1}{v}(-\bar{d}_L, \bar{u}_L)\phi^* \boldsymbol{M}^u u_R + \text{h.c.}, \tag{2.20}$$

where f_L/f_R are the left and right-handed components of a fermion field f, and u/d represent any two fermion fields that transform together as an $SU(2)$ doublet. Each matrix $\boldsymbol{M}^f$ is the *Yukawa matrix* that determines the masses of the fermion family f.

The two scalar fields can be viewed as a two dimensional vector with length equal to $|\phi|$, which transforms as

$$\phi \to U_Y U_W \phi \tag{2.21}$$

under the $U(1)$ and $SU(2)$ gauge symmetries. The $SU(2)$ transformation can then simply be viewed as a rotation of the ϕ vector around the origin. The only value of $|\phi|$ which is invariant under these kinds of rotations is $|\phi| = 0$, but the ground state corresponding to the minimum of the Higgs potential has length $|\phi| = v$. This means that any particular

ground state of the Higgs doublet is *not* invariant under electroweak gauge transformations, even though the Lagrangian is. This is directly related to the fact that the mass term of the ϕ doublet is positive, which makes them tachyonic fields of imaginary mass. This is a typical example of *spontaneous symmetry breaking*, which occurs whenever a continuous symmetry of the Lagrangian isn't obeyed by the ground state of the system. In general, this can be shown to produce one massless *Goldstone boson* for each broken symmetry [37]. Here, there are the three such symmetries, corresponding to the three $SU(2)$ generators.

We are always free to fix a specific gauge condition, as long as it can be satisfied for arbitrary ϕ under $SU(2)$ rotations. The *unitary gauge* fixes ϕ to the form,

$$\phi = \begin{pmatrix} v + H \\ \\ 0 \end{pmatrix}, \tag{2.22}$$

where H is now a *real* scalar field, with a vanishing ground state and a real mass $m_H = v\lambda$. In the low energy limit, we can expand around $H \sim 0$, which turns the lowest order terms of all the ϕ couplings into mass terms and interactions with the Higgs boson H. In this gauge, the three Goldstone bosons expected from spontaneous symmetry breaking are said to have been *eaten* by the three weak bosons, which all acquire masses. The high energy electroweak force decouples into the low energy interactions we call the electromagnetic and *weak forces*.

2.2.2 Symmetries

As mentioned earlier, symmetries play a huge role in high energy physics, and the SM contains a number of relevant ones. These can be classified by their symmetry group as either discrete or continuous, and can also be exactly or approximately obeyed.

Space-time Symmetries As with all relativistic theories, the SM is invariant under transformations described by the Poincare group. This is a Lie group consisting of the four

space-time translations, and the six Lorentz transformations of $SO(1,3)$. These symmetries are required by the *principle of relativity*, so they're almost universally assumed in any QFT.

Gauge Symmetries Gauge symmetries refer to continuous field transformations under which the Lagrangian is invariant. They are local, meaning that gauge transformations are generally functions of space-time, varying from point to point. Very often a specific gauge is chosen, specified by constraints placed on the fields that break the gauge symmetry of the theory. These symmetries point to redundancies in our description of the physical system, since any choice of gauge will produce the same observable results. The three forces of the SM each have conserved charges, which can be shown to follow from Noether's theorem and their *global* gauge symmetries.

Discrete Symmetries There are three important inversion symmetries in the SM, called *time reversal* (T), *charge reversal* (C), and *parity reversal* (P). The combination of all three, *charge-parity-time reversal* (CPT), is a consequence of Lorentz invariance and is expected to be a symmetry of any BSM theory. However, they are individually broken explicitly by the $SU(2)$ force, which only couples to certain combinations of charge and parity. The *charge-parity reversal* (CP), or equivalently T, symmetry is very nearly obeyed by the SM, but is expected to be violated by BSM theories. Although it's one of the necessary conditions for *baryogenesis* in the early universe [38], so far its breaking has only been observed in very rare circumstances involving the heaviest quarks [39–41].

Global Symmetries There are a number of accidental and approximate global symmetries in the SM, each with associated conservation laws. Accidental symmetries are those that were not explicitly added in the formulation of a theory. The SM has 4 of these which are all $U(1)$ global symmetries corresponding to the conservation of baryon number, and the three *lepton flavors*. The existence of massive neutrinos breaks the three lepton flavor symmetries, resulting in only a single $U(1)$ lepton number symmetry. Both baryon and lepton number symmetries are broken by quantum corrections to the classical field Lagrangian

(i.e. chiral anomalies). However, the combined $B - L$ symmetry is expected to be an exact symmetry of nature.

In the SM there are two notable approximate symmetries, corresponding to $SU(2)$ and $SU(3)$ rotations of the light *quark flavors*. The only terms in the SM that explicitly break the flavor symmetry between all six quarks are their Yukawa couplings to the Higgs field, which give them mass. Up, down, and *strange quarks* have such low masses compared to the *hadrons* they compose, that they can be treated as approximately massless. The approximately conserved current associated with these symmetries is called *isospin*, although it's usually only used to refer to the $SU(2)$ symmetry between up and down quarks. Strange quarks are much more massive than up or down quarks, so the $SU(3)$ isospin symmetry is less accurate.

2.3 Model Independent WIMP Searches

In general, it's very difficult to compare results from all of the different DM searches without choosing a model. Without a specific theory, it's impossible to predict how DM interactions would vary from search to search, due to the different approaches taken. Model-dependent analyses can be optimized to a specific theory, and tend to be more sensitive to *that* theory. However, their drawback is that there are a vast number of BSM models that have been formulated, each with many adjustable parameters. Optimizing a search to a single model biases the analysis against the other possibilities. The approach taken by most ATLAS DM searches is to work in a framework that makes three basic assumptions, in order to achieve near model-independence and test a wide range of different theories simultaneously.

The first assumption is the WIMP hypothesis. We assume that dark matter is made up of a *single* particle with mass and couplings at approximately the electroweak scale. By mass, there is about five times more DM than baryonic matter in the universe, so it may seem presumptuous to assume that the dark sector is occupied by a single particle.

However, even for baryonic matter this assumption isn't that far off, where free protons account for about 75% of the universe's baryonic matter, with the remaining 25% mostly made up of helium nuclei [42]. This is a simplifying assumption, not necessarily expected behavior, and is only used to turn a WIMP-production search into a DM search. Until we compare the results of the analysis to cosmological observations, this assumption isn't necessary.

2.3.1 Effective Field Theory

The second assumption is that the underlying theory can be approximated by an EFT. This implicitly requires that whatever particle mediates the interactions between the SM and DM is massive enough that it can be integrated out of the theory, resulting in EFT contact interactions. This assumption can be justified by looking at the typical DM candidate WIMPs. They need to be heavy, but they also need to be *dark* and *cold* in order to correctly account for DM. To remain dark and cold, a particle must be very stable. If DM could decay to SM particles with any significant frequency we would have seen it by now, and if it decays within the dark sector, then it would need an extremely long lifetime in order to have survived this long. Formulating a theory with a stable massive particle isn't trivial, and in fact, without adding some new symmetry, a sufficiently massive particle *can't* be stable. The typical way of accomplishing this is by adding some new conserved quantum number that differs between SM and DM particles. The lightest particle in the dark sector can then be labeled the DM candidate, since it's naturally stable. So in the typical BSM theory, the DM candidate is the lightest new particle, and any dark mediators should be more massive. It's still possible for the mediators to have a mass similar to or less than DM though, so this assumption should be kept in mind when interpreting the results.

Despite its restrictions, the EFT assumption is very powerful. By integrating out the unknown mediator, we can enumerate every possible contact interaction between an arbitrary WIMP and the SM particles. Since higher dimensional operators are suppressed by increasing powers of a new mass scale M_*, which is related to the mediator mass and

therefore large, we can limit ourselves to only the lowest dimensional operators. This leaves a finite number of possible DM-SM interactions, which are listed in Table 2.1 for quarks and gluons. *A priori*, each of the quark-DM interactions is a sum over 6 distinct interactions corresponding to the 6 quark flavors, each with its own M_* coupling.

Name	Operator	Coefficient	Interaction	Suppression
D1	$\bar{\chi}\chi\bar{q}q$	m_q/M_*^3	SI	-
D2	$\bar{\chi}\gamma^5\chi\bar{q}q$	im_q/M_*^3	SI	q^2
D3	$\bar{\chi}\chi\bar{q}\gamma^5 q$	im_q/M_*^3	SD	q^2
D4	$\bar{\chi}\gamma^5\chi\bar{q}\gamma^5 q$	m_q/M_*^3	SD	q^4
D5	$\bar{\chi}\gamma^\mu\chi\bar{q}\gamma_\mu q$	$1/M_*^2$	SI	-
D6	$\bar{\chi}\gamma^\mu\gamma^5\chi\bar{q}\gamma_\mu q$	$1/M_*^2$	SI	v^2, q^2
D7	$\bar{\chi}\gamma^\mu\chi\bar{q}\gamma_\mu\gamma^5 q$	$1/M_*^2$	SD	v^2, q^2
D8	$\bar{\chi}\gamma^\mu\gamma^5\chi\bar{q}\gamma_\mu\gamma^5 q$	$1/M_*^2$	SD	-
D9	$\bar{\chi}\sigma^{\mu\nu}\chi\bar{q}\sigma_{\mu\nu}q$	m_q/M_*^3	SD	-
D10	$\epsilon_{\mu\nu\sigma\rho}\bar{\chi}\sigma^{\mu\nu}\chi\bar{q}\sigma^{\sigma\rho}q$	im_q/M_*^3	SI	q^2
D11	$\bar{\chi}\chi G_{\mu\nu}G^{\mu\nu}$	$\alpha_s/4M_*^3$	SI	-
D12	$\bar{\chi}\gamma^5\chi G_{\mu\nu}G^{\mu\nu}$	$i\alpha_s/4M_*^3$	SD	q^2
C1	$\chi^\dagger\chi\bar{q}q$	m_q/M_*^2	SI	-
C2	$\chi^\dagger\chi\bar{q}\gamma^5 q$	im_q/M_*^2	SD	q^2
C3	$\chi^\dagger\partial_\mu\chi\bar{q}\gamma^\mu q$	$1/M_*^2$	SI	-
C4	$\chi^\dagger\partial_\mu\chi\bar{q}\gamma^\mu\gamma^5 q$	$1/M_*^2$	SD	v^2
C5	$\chi^\dagger\chi G_{\mu\nu}G^{\mu\nu}$	$\alpha_s/4M_*^2$	SI	-

Table 2.1: The lowest dimensional EFT operators for interactions between quarks and gluons with DM [43]. Dirac fermion DM is shown on top, and complex scalar DM is shown on the bottom. After applying MFV each quark operator picks up a sum over the three quark generations. In general, there are two different M_* values for each of these interactions corresponding to up and down type quarks. The quark fields can be switched out for leptons and the gluons for photons in order to get the other possible SM-DM interactions, but these are not relevant here. The dominant nuclear interactions for each operator are also listed, along with suppression factors in terms of the DM velocity v and total energy transfer q. These can be either spin-dependent (SD) or spin-independent (SI). [44]

2.3.2 EFT Validity

In order to evaluate the EFT approximation, its *region of validity* needs to be estimated. Generally, this is very difficult without knowledge of the underlying *ultraviolet*

(UV)-complete theory, and there is never any sharp cutoff. However, we can get a rough idea of where this validity region is by looking at a minimal model for each operator and making some conservative constraints, as done in [45]. Since, within the EFT validity region, the underlying theory is irrelevant, this provides a useful method of testing the bounds of the EFT assumption.

To construct a minimal model, we add a new WIMP of mass m_χ and a heavy mediator of mass M_ϕ to the SM, with the interactions,

$$\mathscr{L} \supset g_q \phi \bar{q} q + g_\chi \phi \bar{\chi} \chi \tag{2.23}$$

for fermionic WIMPs and,

$$\mathscr{L} \supset g_q \phi \bar{q} q + g_\chi \Lambda \phi \chi^\dagger \chi \tag{2.24}$$

for scalar WIMPs. The strength of the quark-mediator and WIMP-mediator couplings are given by the dimensionless constants g_q and g_χ respectively. Note that for scalar DM at leading order, there is an additional mass parameter Λ in the WIMP-mediator interaction that can be treated as the VEV of some additional massive field. In order for these interactions to be perturbative, the condition $\sqrt{g_q g_\chi} < 4\pi$ must hold. Although perturbativity isn't an absolute requirement, it's a reasonable assumption for electroweak scale interactions.

The final step is to place bounds on M_ϕ as a function of m_χ. The condition for EFT validity is that the momentum transferred to the WIMP pair in each collision, Q_{tr}, must be less than M_ϕ. This is directly related to the EFT assumption that $Q_{tr} \ll M_\phi$, and therefore gives a reasonable lower bound on M_ϕ. The simplest validity requirement is $M_\phi > 2m_\chi$, which follows from the observation that in any collision that produces a pair of WIMPs, $Q_{tr} \geq 2m_\chi$. By applying the validity condition on an event-by-event basis, a stronger bound can be set. The percentage of signal events passing the validity condition is then used as a metric for the validity of the final result.

In order to make use of these bounds, they need to be rephrased in terms of the model-

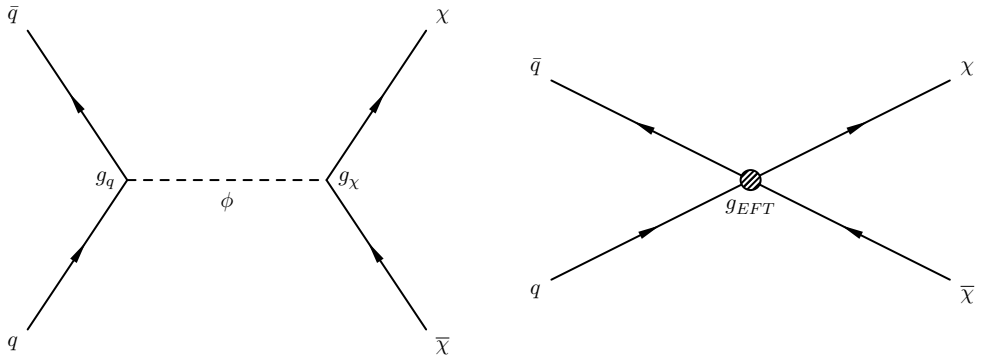

Figure 2.4: The Feynman diagrams used for EFT matching with the minimal UV-complete model.

Operator	Matching Condition	EFT Validity
D1	$\frac{g_q g_\chi}{M_\phi^2} = \frac{m_q}{M_*^3}$	$Q_{tr} \lesssim 4\pi\sqrt{\frac{M_*^3}{m_q}}$
D9	$\frac{g_q g_\chi}{M_\phi^2} = \frac{m_q}{M_*^3}$	$Q_{tr} \lesssim 4\pi\sqrt{\frac{M_*^3}{m_q}}$
C1	$\frac{g_q g_\chi}{M_\phi^2}\Lambda = \frac{m_q}{M_*^2}$	$Q_{tr} \lesssim 4\pi\frac{M_*^2}{m_q}$

Table 2.2: The EFT validity constraints for the operators of interest.

independent coupling M_*. This can be accomplished by performing a matching between the EFT operators and the minimal model chosen above. The interactions to be matched are shown in Fig. 2.4 at LO, and the resulting EFT validity conditions are shown in Table 2.2. Note that for C1, the assumption $\Lambda = M_\phi$ is made to remove the unknown mass scale Λ from consideration. This is a conservative measure, since M_ϕ is the maximum mass scale in this theory. If Λ were, for example, the Higgs VEV, the validity requirements would be significantly tighter than necessary. Applying these conditions allows us to estimate the EFT validity of the final results for each signal considered. Points outside the validity region may overestimate or underestimate the sensitivity of this analysis, depending on the details of the underlying UV-complete model.

2.3.3 Minimal Flavor Violation

The final assumption made in this search is known as *minimal flavor violation* (MFV), as defined in [46]. This is a purely phenomenologically motivated concept stemming from the observation that precision flavor experiments have set very strict limits on the energy scale of *generic* NP of roughly $\Lambda_{NP} > \mathcal{O}(10^4)$ TeV. This conflicts with the usual expectation of $\mathcal{O}(1)$ TeV scale physics beyond the SM, and is also much too high to be probed by any modern experiment. The only way we can hope to have any sensitivity to NP is by restricting its flavor structure at the TeV scale. MFV is a sufficient, but not unique, restriction that brings the experimental bounds down to $\mathcal{O}(1)$ TeV, as desired.

To apply MFV, we first note that the SM has the approximate global flavor symmetry G_F, described by

$$G_F \equiv U(3)_{q_L} \otimes U(3)_{u_R} \otimes U(3)_{d_R} \otimes U(3)_{\ell_L} \otimes U(3)_{e_R}, \tag{2.25}$$

which corresponds to 5 independent rotational symmetries among the three generations for left-handed quarks, right-handed up-type quarks, right-handed down-type quark, left-handed leptons, and right-handed charged leptons. Massive Dirac neutrinos would add a sixth $U(3)$ symmetry corresponding to right-handed neutrinos, but that is irrelevant to this search. The only terms in the SM that violate G_F are the Higgs Yukawa couplings that generate the fermion masses. Since the SM violates G_F explicitly, NP can't be expected to respect it. In fact, a generic BSM theory that *doesn't* violate G_F at all would still have very high bounds on its mass scale. MFV requires that any BSM theory break flavor symmetry *in the same way* as the SM, with the Yukawa matrices.

MFV is enforced by making use of *spurion fields*, which are a theoretical tool used to make an approximate symmetry exact until the theory is fully formulated. For each approximate symmetry, a very massive field is added to the theory with a VEV equal to the parameter that breaks the symmetry. Every instance of this parameter is then replaced with couplings to the new spurion field. The original theory can be recovered at any time

by taking the spurion mass to infinity and integrating it out, leaving only the VEV. In order to make G_F an exact symmetry, the Yukawa matrices, Y_f, are promoted to spurion fields that transform with the fermion fields as,

$$\begin{aligned} &Y_u \to V_{q_L} Y_u V_{u_R}^\dagger, \quad Y_d \to V_{q_L} Y_d V_{d_R}^\dagger, \quad Y_\ell \to V_{\ell_L} Y_\ell V_{e_R}^\dagger, \\ &Q_L \to V_{q_L} Q_L, \qquad D_R \to V_{d_R} D_R, \qquad U_R \to V_{u_R} U_R, \\ &E_R \to V_{e_R} E_R, \qquad L_L \to V_{\ell_L} L_L, \end{aligned} \tag{2.26}$$

under the 5 independent $U(3)$ flavor transformations V_{q_L}, V_{u_R}, V_{d_R}, V_{ℓ_L}, and V_{e_R}. Here Q_L, U_R, and D_R correspond to the three quark flavor triplets, and L_L and E_R correspond to the two lepton flavor triplets. Q_L and L_L are also $SU(2)$ doublets representing left handed quarks and leptons, while U_R, D_R, and E_R are $SU(2)$ singlets that represent right handed up-type quarks, down-type quarks, and electrons respectively. Once this substitution is made, every SM interaction becomes invariant under G_F. The formal definition of MFV is then simply that no interaction in the BSM theory of interest is allowed to violate G_F as long as the spurions are still in place. Once the full theory has been formulated, the spurions are demoted back to the Yukawa matrices, spontaneously breaking G_F and leaving behind a theory that satisfies MFV. Since this analysis is based on hadronic collisions, only interactions of WIMPs with quarks and gluons are important. The lowest order quark couplings allowed under MFV are $\overline{Q}_L Q_L$, $\overline{Q}_L Y_u U_R$, $\overline{Q}_L Y_d D_R$, $\overline{U}_R U_R$, $\overline{D}_R D_R$. Note that the two couplings with Yukawa matrix terms violate the $SU(2)$ gauge symmetry of the SM, and would need to couple to some additional $SU(2)$ doublet, such as the Higgs, to keep the symmetry intact. Theories obeying MFV can only couple to quarks via these terms, tightly constraining the flavor structure of all new interactions.

An arbitrary quark bilinear coupling to new fields will take the form $\bar{q}\Gamma q$, where Γ is any combination of gamma matrices, q is a *spinor* representing a quark field, and $\bar{q} \equiv q^\dagger \gamma^0$. Since each γ^μ matrix is anti-diagonal in the chiral basis and γ^5 is diagonal, any Γ with an even number of γ^μ matrices will put an anti-diagonal matrix between the two quark fields, q

and $q^{\dagger}$. This will mix the left and right handed components which, as shown above, requires the addition of a Yukawa matrix under MFV. After demoting the spurion fields and fixing the unitary gauge, every operator that mixes left and right handed quarks picks up a factor of that quark's mass. The allowed quark bilinears that remain then, are exactly the ones used in Table 2.1. The only difference is that there are now implicit sums over the quark generations, reducing the number of possible interactions by a factor of three and giving actual meaning to the m_q factors assigned to the scalar and tensor operators. Many analyses make an additional assumption of *universality*, where up-type and down-type quarks have equal M_* couplings. For this case, each entry in Table 2.1 represents a single interaction with a sum over all 6 quark flavors. However, there is no strong justification for assuming universality, and MFV does *not* require it. Therefore in this thesis, *universal couplings* will not be assumed in general, and each of the quark-DM operators listed represent two distinct interactions: one for down-type quarks and another for up-type quarks, each with their own mass scales M_*. Because it's rather simple to extend the analysis to include universal couplings, new operators C1u, D1u, and D9u with universal couplings will be considered at the end. It should be noted however, that the dominant signal process for these operators produces *top quarks* pairs, which this analysis isn't optimized to. The higher mass final states results in much tighter EFT validity conditions on operators with universal couplings.

2.3.4 Mono-b Search

Within this EFT/MFV framework, different types of WIMP searches can be easily compared in a model-independent way. The motivation for the search presented here comes from the quark mass dependence of the D1-4, C1-2, and D9-10 operators, where DM would couple much more strongly to *bottom quark* and top quarks than the lighter quarks. The down-type quark couplings to DM will be the main focus of this analysis, because there is no reason *a priori* to assume universal couplings. The search presented here is unique among ATLAS mono-X searches, in that the tagged particle recoiling off the WIMP is

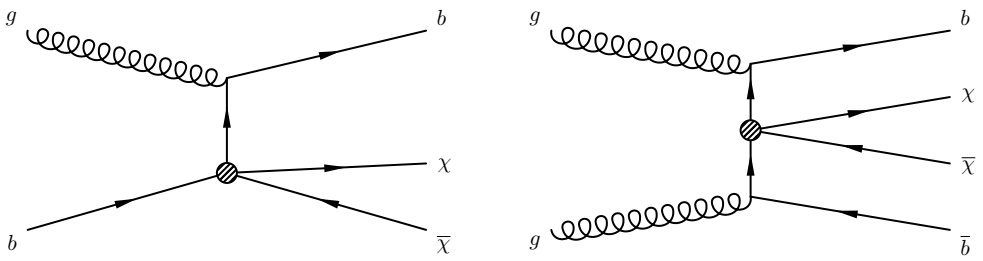

Figure 2.5: The dominant processes in the production of WIMPs for the interactions considered in this search.

crucial to the interaction, as shown in Fig. 2.5. The other Mono-X searches that have been done include mono-jet [47], mono-γ [48], mono-W [49] [50], and mono-Z [51] [50], which all rely on *initial state radiation* (ISR) to "tag" WIMP production, as illustrated in Fig. 2.6. This ISR has absolutely nothing to do with the WIMP production though, and its only purpose is to make it clear that something invisible was created. For the operators these searches are most sensitive to, the dominant production mechanism won't be triggered on, since nothing visible is produced.

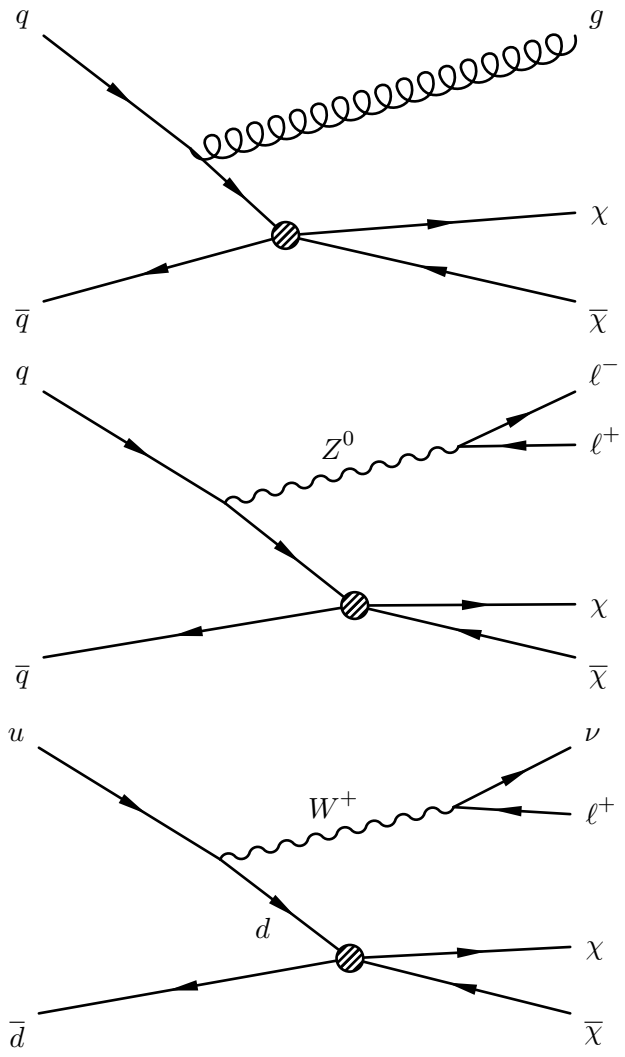

Figure 2.6: The dominant processes, from top to bottom, for Mono-jet, Mono-Z, and Mono-W.

Chapter 3

Cosmological Dark Matter

At this point in history, dark matter has only been observed at *cosmological scales*, with the smallest scale observations coming from *dwarf galaxies*. From these limited observations though, a great deal has been discovered about it. Although many details about the composition of dark matter are still unknown, others can be inferred. The fact that it's completely invisible, or *dark*, to us suggests that it does not interact through the electromagnetic force or the strong force, since any interactions of this nature would make it very easy to see. The observation that dark matter is gravitationally bound to *galaxies* suggests that it's *non-relativistic*, or *cold*. Simple calculations for the escape velocity of a typical galaxy shows that dark matter must be moving well under the speed of light, which also suggests that its constituents are relatively *massive*. Examining information such as this, obtained at the cosmological scale, can significantly narrow down searches for dark matter particles at the quantum scale.

3.1 $\Lambda - \mathrm{CDM}$

The "standard model" of cosmology, which has had extraordinary success, is known as the *Lambda-cold dark matter* ($\Lambda - \mathrm{CDM}$) model, where Λ corresponds to the *cosmological constant* term in the *Einstein field equations* (EFE). $\Lambda - \mathrm{CDM}$ is a shockingly simple model, considering the complexities of the universe, that provides a very clear picture of the universe's evolution since moments after the Big Bang. This is the *only* experimentally confirmed theory that combines quantum mechanics and thermodynamics with *General*

Relativity, and is crucial to the understanding of dark matter at the cosmological scale. Fig. 3.1 shows the timeline of our universe predicted by the model, which will be explained in more detail below. Note that in this section we will only discuss the *theory* behind $\Lambda - \text{CDM}$, which is capable of describing a wide range of universes, many of which are drastically different from ours. In the following sections, we will go over experiments that reinforce $\Lambda - \text{CDM}$ and constrain its parameters, thereby giving us information about which type of universe we live in. For a more in-depth look into these topics, which are far too rich to discuss exhaustively here, see [28, 34, 52].

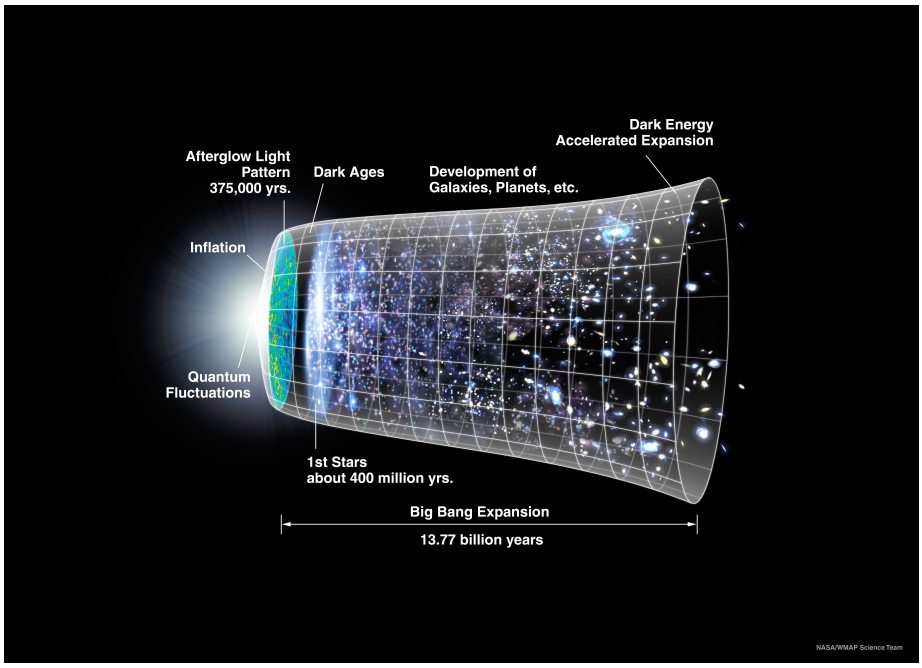

Figure 3.1: A visual timeline showing the evolution of the universe since the Big Bang. Immediately after the Big Bang the universe was in a very hot, very dense state where all of the fundamental forces may have been unified into a single force. As the universe expanded and cooled, the forces began to decouple. Around the end of the grand unification epoch, a rapid expansion of the universe occurred due to a hypothetical scalar field. This field then decayed into matter, from which baryogenesis occurred. The universe continued to cool, forming structures in a bottom-up pattern. First hadrons formed, then nucleons, atoms, stars, galaxies, galactic clusters, and finally the galactic super clusters we see today. [5]

3.1.1 FLRW Geometry

FLRW Metric At very large scales the universe appears to be extraordinarily homogeneous and isotropic, meaning that there are no preferred positions or directions respectively. This is closely related to the principle of relativity, and is known as the *cosmological principle*. A general space-time metric can be developed, under the assumption of perfect homogeneity and isotropy, that should be capable of describing our universe, at least on large scales. This is known as the *Friedmann–Lemaître–Robertson–Walker* (FLRW) metric, which is described by the infinitesimal line element,

$$ds^2 \equiv dx^\mu dx^\nu g_{\mu\nu} = dt^2 - a^2(t)\left[\frac{dr^2}{1-kr^2} + r^2\left(d\theta^2 + \sin^2\theta d\phi^2\right)\right]. \tag{3.1}$$

The only parameters to the model are the scale factor, $a(t)$, and the curvature constant, k. The curvature constant is a universal parameter with only three possible values $(-1, 0, 1)$ corresponding to open, flat, or closed universes respectively. In a closed universe, the total volume can be calculated as $2\pi^2 a^3$, which is exactly the surface "area" of a 3-sphere embedded in a 4-dimensional space. In this sense, the scale factor, $a(t)$, can roughly be thought of as the "radius" (or scale) of the universe, even in flat or open universes. Note that the coordinates t, r, θ, ϕ are *not* the usual Minkowski coordinates, as is typical in General Relativity. The labeling of these coordinates isn't arbitrary though, since if $k = 0$ and $a(t)$ is constant, the flat Minkowski metric is recovered by taking $ar \to r$.

The EFEs, which describe gravity in General Relativity, are given by

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R + g_{\mu\nu}\Lambda = 8\pi T_{\mu\nu}, \tag{3.2}$$

where $R_{\mu\nu}$ is known as the Ricci tensor and R is its trace R^μ_μ, both of which are functions of only the space-time metric. Λ is the cosmological constant, and $T_{\mu\nu}$ is the stress-energy tensor of the universe. The simplest model for describing large scale stress-energy tensors

is that of the *perfect fluid*,

$$T^{\mu\nu} = (\rho + p)\, v^{\mu} v^{\nu} - p g^{\mu\nu}, \tag{3.3}$$

where at any point in space-time, ρ is the energy density of the *fluid*, p is its pressure, and v^{μ} is its *4-velocity*. Note that the cosmological constant can be rephrased as a very exotic form of matter with an equation of state $p = -\rho$, giving a stress-energy tensor of $T^{\mu\nu} = -p g^{\mu\nu}$. It's easy to see that by adding matter of this form to the EFEs with $\Lambda = 0$, one recovers a new term identical to the original cosmological constant term if $\Lambda \to -8\pi p$. For this reason, Λ is more frequently called dark energy, to highlight the fact that it can be described as an exotic form of energy. Many extensions to $\Lambda -$ CDM add scalar fields that behave like the cosmological constant when averaged over large scales, but have dynamical effects at quantum scales.

Substituting a perfect fluid and the FLRW metric into the EFEs produces the three Friedmann-Lemaître equations,

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi}{3}\rho + \frac{1}{3}\Lambda - \frac{k}{a^2}, \tag{3.4}$$

$$\frac{\ddot{a}}{a} = \frac{1}{3}\Lambda - \frac{4\pi}{3}\left(\rho + 3p\right), \tag{3.5}$$

$$\dot{\rho} = -3\frac{\dot{a}}{a}\left(\rho + p\right), \tag{3.6}$$

which describe the evolution of a universe filled with a homogeneous, isotropic fluid of mass density ρ and pressure p.

FLRW Quantities A physically meaningful function, which will be discussed shortly, is the Hubble parameter $H(t)$ defined by,

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}. \tag{3.7}$$

The quantity t_0 is defined as the time here on earth, so $H_0 \equiv H(t_0)$ is defined as the Hubble parameter experienced in our local neighborhood. We can also define a *redshift parameter*,

$$z \equiv \frac{\nu_E - \nu_R}{\nu_R}, \tag{3.8}$$

corresponding to any frequency changes a *photon* might endure traveling from the emitter E at time t_E with frequency ν_E, to a receiver R, who measures frequency ν_R at time t_R. A value of $z < 0$ means that the photon was received at a higher frequency than it was emitted, called *blueshift*. Likewise, $z > 0$ means that the photon was received at a lower frequency than it was emitted, called *redshift*. In a FLRW geometry, this can be rephrased in terms of the scale factor $a(t)$,

$$z = \frac{a(t_R)}{a(t_E)} - 1. \tag{3.9}$$

Since $t_R > t_E$ by definition, in an expanding universe $z > 0$ and photons will always experience redshift. Two more important values are the critical density ρ_c and the reduced Hubble parameter h given by,

$$h \equiv \frac{H_0}{100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}}, \qquad \rho_c(t) \equiv \frac{3}{8\pi}H^2. \tag{3.10}$$

The critical density is the total energy density of a flat, $k = 0$, universe with $\Lambda = 0$, and can be easily derived from Eq. 3.4. The energy density of a particle species x, ρ_x, is typically divided by the critical density in order to give the more useful *cosmological density* Ω_x. From numerous observations we have determined that the total density of the universe is consistent with the critical density to within very small errors [53]. This is convenient, as it allows the parameter Ω_x to also be interpreted as the energy fraction of the universe taken up by x. Cosmological densities are also frequently expressed as *physical densities*, $\Omega_x h^2$, which reduces uncertainties from H_0.

FLRW Solutions Using Eq. 3.6, the mass density ρ of a fluid filling the universe can be solved for analytically, as a function of $a(t)$, under the assumption that it has an equation of state $p = w\rho$, for some adjustable parameter w. This gives $\rho_w = Na^{-3(1+w)}$, with an arbitrary normalization N. There are three perfect fluid models typically used to describe the different populations in the universe, for which w can be readily obtained. An ideal *dust*, which approximates weakly interacting non-relativistic matter, is defined as a non-interacting fluid with vanishing pressure, so that $w = 0$. *Ultra-relativistic fluids*, which are similar to photon gases, approximate any sufficiently hot gas, and are defined as non-interacting massless particles moving at the speed of light. These types of fluids exhibit radiation pressure equal to $1/3\rho_r$, so $w = 1/3$. The third type of fluid is used exclusively to describe the cosmological constant contributions to the energy density. As mentioned before, the cosmological constant can be treated as an exotic dark energy fluid with equation of state $p_\Lambda = -\rho_\Lambda$, and therefore $w = -1$. Generally, cold matter is approximated as dust, while hot matter, neutrinos, and photons are approximated as ultra-relativistic fluids.

Although it's clearly an over-simplification, it can be useful to imagine a universe containing only one of these three fluid. Plugging the solution for ρ_w into Eq. 3.4,

$$\dot{a}^2 = \frac{8\pi}{3} N a^{-3w-1} - k, \tag{3.11}$$

where for very small $a(t)$, the k term can be neglected as long as $w > -1/3$. Since we live in an expanding universe, this approximation can be used to describe the early universe when matter and radiation dominated its energy density. For very large $a(t)$, the k term can also be neglected if $w < -1/3$. This would correspond to a very large universe, where dark energy dominates. Isolating each of the three types of energy gives the solutions,

$$\begin{aligned} a_m(t) &\propto t^{2/3}, & H_m(t) &= \frac{2}{3t}, \\ a_r(t) &\propto t^{1/2}, & H_r(t) &= \frac{1}{2t}, \end{aligned} \tag{3.12}$$

$$a_\Lambda(t) \propto e^{\sqrt{1/3\Lambda}t}, \qquad H_\Lambda(t) = \sqrt{\frac{\Lambda}{3}},$$

where the subscripts m, r, and Λ refer to universes in which matter, radiation, and dark energy dominate respectively. From these, it's clear that the expansion rate, $\dot{a}$, decreases in matter and radiation-dominated universes, and increases exponentially in a universe dominated by dark energy. Different forms of energy have dominated throughout our universe's evolution, giving it a rich history that will be discussed further in the following sections.

Thermodynamics The very early universe was so hot that it can be approximated as a radiation-dominated universe where all particles are moving relativistically. An ideal photon gas obeys the thermodynamic relations [54],

$$\rho_\gamma = \frac{\pi^2}{15}T^4, \quad p_\gamma = \frac{1}{3}\rho_\gamma, \quad n_\gamma = \frac{30\zeta(3)\rho_\gamma}{\pi^4 T}, \quad s_\gamma = \frac{4\rho_\gamma}{3T}, \tag{3.13}$$

where T is its temperature, n_γ is its number density, and s_γ is its entropy density. From the Friedmann equations, we know that $\rho \propto a^{-4}$ for radiation-dominated universes, meaning that $T \propto a^{-1}$ and $s \propto a^{-3}$. For the more general case of a *relativistic* gas of more than one type of particle, a function $g_*(T)$ is included, where the total density is now related to the photon density by $\rho \equiv g_*(T)\rho_\gamma/2$. For a thermal bath composed of both bosons and fermions, g_* is given by,

$$g_*(T) = \sum_{\text{bosons}} g_i \left(\frac{T_i}{T}\right)^4 + \frac{7}{8}\sum_{\text{fermions}} g_i \left(\frac{T_i}{T}\right)^4, \tag{3.14}$$

where for each i, T_i and g_i are the temperature and internal degrees of freedom for particle species i. Combining this with Eqns. 3.12 and 3.4, we find that the temperature of a

radiation-dominated universe is related to the age of the universe by

$$t = \frac{1}{4\pi T^2}\sqrt{\frac{45}{g_*\pi}}. \tag{3.15}$$

The function $g_*(T)$, plotted in Fig. 3.2, represents the relativistic degrees of freedom in the expanding universe, which falls as the universe cools and heavier particles fall out of thermal equilibrium. The heaviest known particles are top quarks, with a mass of 173 GeV and a corresponding temperature scale of 2×10^{15} K. At this temperature $g_* \approx 106.75$, and above it g_* depends strongly on whether or not more massive particles exist. As the temperature drops, so does g_* until it reaches the electron temperature scale, 6×10^9 K, at which point g_* stays a constant 7.25 until the relativistic approximation falls apart (assuming massless neutrinos).

3.1.2 The Big Bang

The Big Bang, occurring at $t = 0$, was quite literally the beginning of time. Immediately after its formation, the universe was very dense, very hot, and rapidly expanding. As it expanded it underwent a number of phase transitions occurring at the boundaries of *epochs* [56]. The first moments after the Big Bang are known as the *Planck epoch*, but without a full theory of quantum gravitation, it's pointless to speculate about what may have happened during this time. We can really only begin to talk about the Big Bang after it had cooled to below T_P and the gravitational force had become sufficiently weaker than the other known forces. Assuming the four fundamental forces are unified at the *Planck scale*, the end of the Planck epoch corresponds to the point at which the temperature dropped low enough for gravity to decouple from the other three forces. This would presumably be described by some sort of spontaneous symmetry breaking phase transition in a complete theory of quantum gravity. Using Eq. 3.15, we can calculate that the Planck epoch ended at approximately $t \approx 10^{-45}$ s.

After the Planck epoch came the *grand unification epoch*, where the strong and elec-

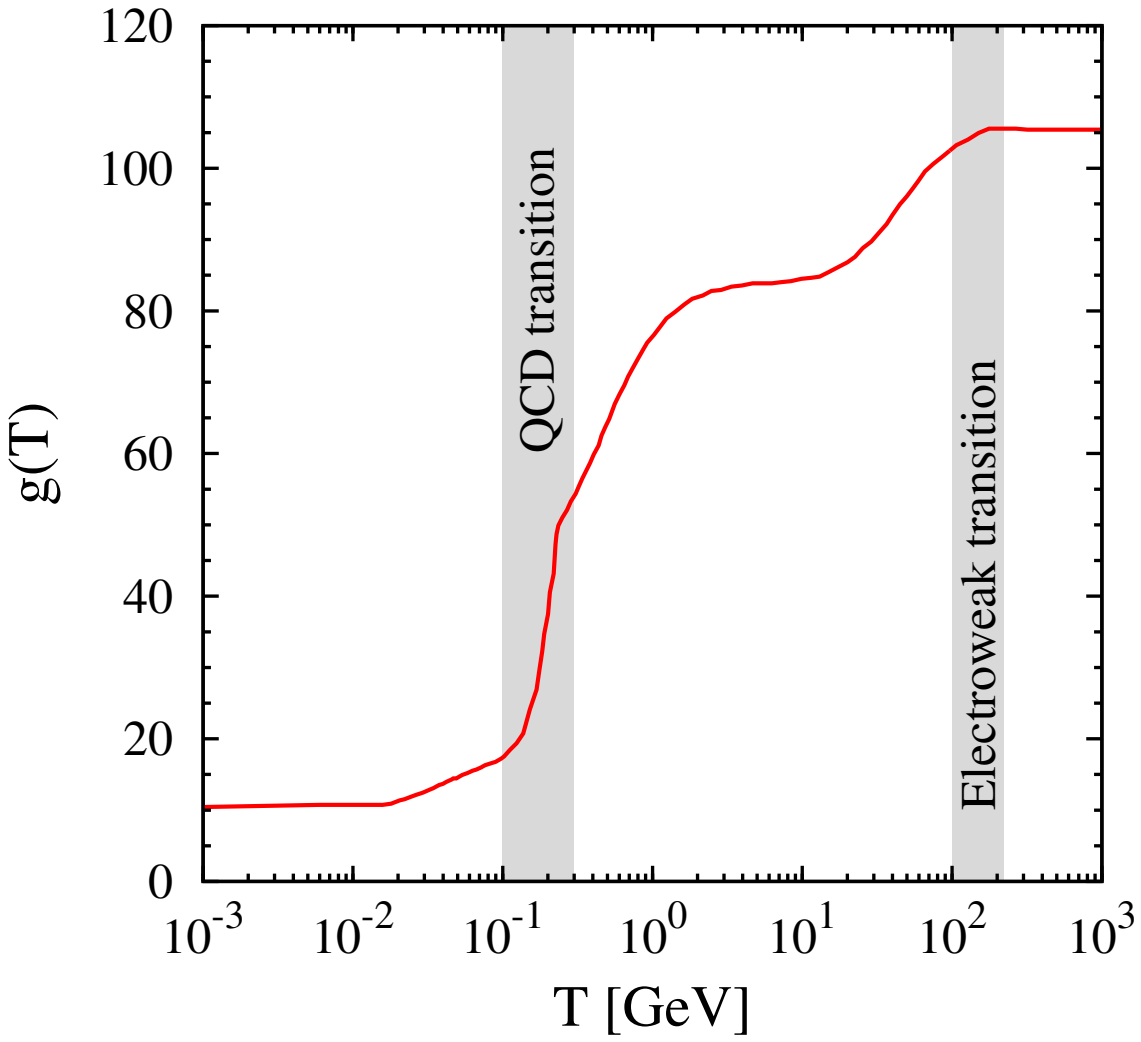

Figure 3.2: The relativistic degrees of freedom, $g_*(T)$, at various temperatures. This is based off *known* particles, so could vary quite significantly at higher temperatures. It also assumes neutrinos are massless, which we now know to be false. [55]

troweak forces are believed to have been unified. As with quantum gravity, we haven't confirmed the existence of any *grand unified theory* (GUT) that unifies these forces, but there are good, albeit controversial, reasons to believe one may exist. The GUT scale is determined by extrapolating the strengths of the three known quantum forces to 10^{16} GeV, at which point they appear to *nearly* converge to a single value, and can be made to exactly converge under various BSM theories. The GUT energy scale corresponds to a temperature of 10^{29} K, which means that the grand unification epoch ended at around $t \approx 10^{-39}$ s.

After the strong force decoupled, the universe entered into the *electroweak epoch*. Although the early stages of this epoch are probably not modeled well by the SM, the later stages should be. In this epoch all known particles were created in abundance and, with the exception of the Higgs fields, were massless. This epoch continued until around 10^{-12} s when the temperature of the universe fell below the electroweak scale, of about 10^{15} K. At this point, the electroweak symmetry was spontaneously broken through the Higgs mechanism, which was described in Chapter 2. Quarks, leptons, and the weak bosons all gained mass, and the Higgs fields condensed into a single massive boson. The end of this epoch is when the universe began to resemble itself today, at least in terms of the underlying physics.

3.1.3 Inflation

There are two major problems with the conventional Big Bang theory when it's compared to our observations. The more significant of these is that the total density of the universe is precisely the critical density, meaning $\Omega = 1$, to within very small errors. Multiplying Eq. 3.4 by a^2, neglecting Λ in the early universe, and pulling out ρ gives

$$\left(\frac{1}{\Omega} - 1\right) \rho a^2 = -\frac{3}{8\pi} k, \tag{3.16}$$

where the right hand side is trivially constant in time. As shown earlier, $\rho \sim R^{-3(1+w)}$, leading to

$$\left(\frac{1}{\Omega} - 1\right) \propto a^{1+3w}, \tag{3.17}$$

which describes the deviation of Ω from unity as a function of time. Since the early universe is dominated by matter and radiation, any deviation from $\Omega = 1$ at the origin of the universe would grow with time no slower than $\sim t^{2/3}$. Given that the measured age of the universe is $\sim 10^{61} t_p$ [53], $(\Omega^{-1} - 1)$ would be over 10^{40} times larger today than it was initially. This is known as the *flatness problem*, and is an example of *fine-tuning*. There is no known law forcing Ω to be exactly 1, but it would need to have been extremely close to 1 at the beginning of the universe in order to have its observed value today. Fine-tuning is usually treated as a sign that something is missing from a theory that would explain the specific value of the parameter we observe. Note, that in a dark energy-dominated universe,

$$\left(\frac{1}{\Omega} - 1\right) \propto e^{-2t\sqrt{\Lambda/3}}, \tag{3.18}$$

which exponentially *decreases* with time. However, the measured value $\Lambda \approx 10^{-37}\,\mathrm{s}^{-2}$ [53] causes a decrease that occurs on the timescale of around 50 Gyr, which is too long to drive Ω so close to 1. The age of the universe is a only about 13.8 Gyr, and dark energy only began to dominate 4 Gyr ago. The dark energy content of a much older universe would drive Ω to unity, but it can't explain the value we see today in *our* universe.

The second problem, known as the *isotropy problem*, is related to the very isotropy and homogeneity that allowed us to model the universe so simply. All around the edge of our visible horizon we can see regions of space that are completely causally disconnected from each other [57]. The diameter of our observable universe is about 90 Gyr, while its age is only 13.8 Gyr, creating a space-like separation between its edges. The FLRW metric predicts that in a decelerating universe, as ours was early on, these regions *were* never causally connected. Further, in the accelerating universe we find ourselves in today, these

regions *will* never become causally connected. This raises the question of how these disjoint regions could be *so* similar, with nearly identical temperatures and matter distributions, if no information ever passed between them.

Although somewhat controversial, the most successful theory that's been proposed to solve these problems, and a few other incidental ones, is the theory of *inflation*. The basic concept of inflation is that somewhere around the electroweak epoch, the universe underwent incredibly rapid expansion before slowing down to the FLRW geometry. Regions that were causally connected initially, were ripped apart at faster than the speed of light, permanently disconnecting them. As long as the universe was able to thermalize before inflation, the isotropy problem is solved.

Inflation is usually described using some unknown scalar *inflaton* field. This field has a potential with some local minimum, and a much lower energy global minimum. Inflation occurs from quantum thermal fluctuations, which push the field to fall to its global minimum, releasing a huge amount of energy in the process. Similar to dark energy, which is also a scalar field, the inflaton field can drive the energy density of the universe exponentially close to its critical density. No matter what the initial Ω was, inflation would very rapidly drive it arbitrarily close to 1, thus solving the flatness problem as well.

At the end of inflation, a huge amount of energy was released into the universe, creating an abundance of all of the particles we've discovered, and probably many we haven't. Baryogenesis and *leptogenesis*, are the processes that occurred during this time to produce a small excess of matter over anti-matter. The exact mechanisms are unknown, but whatever they are, they must meet the Sakharov conditions [38] that require explicit breaking of *baryon*, C, and CP symmetries. Although the SM can accommodate all these conditions, it can't account for anywhere close to the amount of matter we observe today, and this remains an open problem in physics.

3.1.4 Freezeout

In the early universe, thermal equilibrium for any given particle species was kept through cycles of *pair production*, *annihilation*, and other, more complicated, interactions. However, as the universe expanded, these thermal processes became weaker and each species of particle eventually left thermal equilibrium with the rest of the universe. The evolution of this transition out of thermal equilibrium is described by the Boltzmann transport equation [58],

$$\dot{n} = \langle \sigma v \rangle \left(n_{eq}^2 - n^2 \right) - 3Hn, \tag{3.19}$$

where n_{eq} is the number density at thermal equilibrium for the particle species of interest, n is the *actual* number density, and $\langle \sigma v \rangle$ is the *thermally averaged cross-section* for any thermalizing interactions. It's useful to define the dimensionless quantities,

$$x \equiv \frac{m}{T}, \quad y \equiv \frac{n}{s}, \tag{3.20}$$

where m is the particle mass, T is the temperature, and s is the entropy density as defined earlier in Eq. 3.13. Remembering that for an ultra-relativistic fluid, $s \propto a^{-3}$,

$$\dot{y} = \frac{\dot{n}}{s} - \frac{n\dot{s}}{s^2} = \frac{1}{s}\left(\dot{n} + 3Hn \right), \tag{3.21}$$

and using the fact that $T \propto a^{-1}$,

$$dx = -\frac{m}{T}\frac{dT}{T} = x\frac{da}{a} = xHdt. \tag{3.22}$$

Plugging these back into the Boltzmann equation and defining the quantities,

$$\lambda \equiv \frac{sx}{H}, \quad \Gamma \equiv \frac{\langle \sigma v \rangle}{x^2}, \tag{3.23}$$

gives,

$$\frac{dy}{dx} = -\lambda\Gamma\left(y^2 - y_{eq}^2\right). \tag{3.24}$$

This equation is difficult to solve analytically, but its behavior is fairly simple to understand qualitatively, as illustrated in Fig. 3.3. For relativistic particles, λ is constant in time and $\Gamma \propto \langle\sigma v\rangle T^2$. At high temperatures where $\Gamma \gg 1$, the particle density y tightly follows the equilibrium density y_{eq}. As the universe expands and cools, both v and T drop, causing Γ to drop as well. When $\Gamma \sim 1$, the density begins to level out and asymptotically approaches a constant value, which is strongly dependent on $\langle\sigma v\rangle$. This is known as *freeze-out*, because the particles become too cold to interact at any significant rate, and decouple from the rest of the universe. The remaining population is known as a thermal *relic*. Throughout the evolution of the universe, each type of particle eventually experienced freeze-out, leaving behind the relics we observe today.

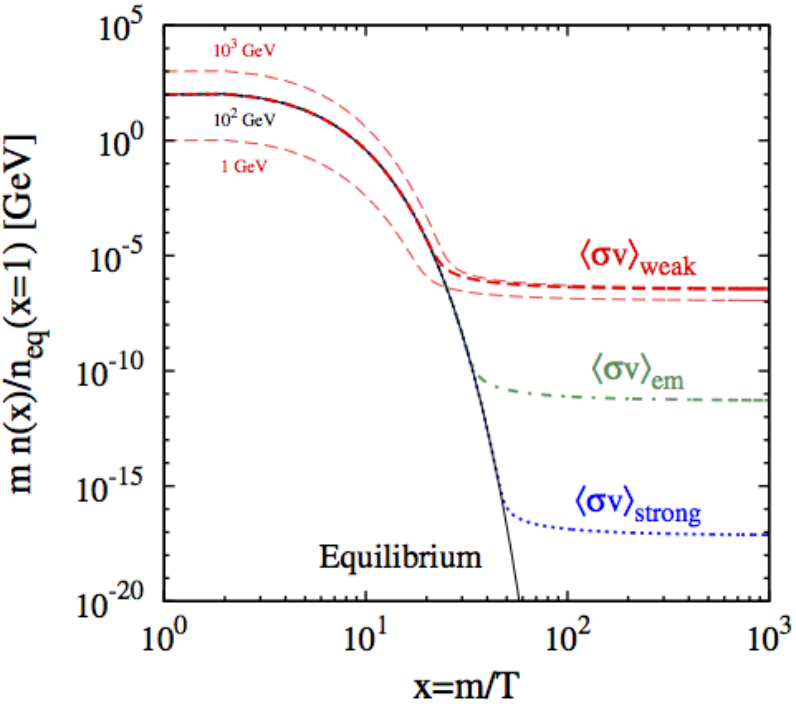

Figure 3.3: Examples of the number density evolution with time for varying cross-sections and masses. The relic abundance and freeze-out temperature primarily depend on the value of $\langle\sigma v\rangle$, with only weak dependence on the particle's mass. [55]

3.1.5 Early Universe

The end of the electroweak epoch marks the point at which we can reliably discuss the universe in terms of our current theories, which are experimentally confirmed at these cooler temperatures. WIMPs, which are the focus of this analysis, are expected to have frozen out shortly after the electroweak epoch. A WIMP is defined as having mass and interaction strengths near the electroweak scale, which would tend to give them non-relativistic velocities at the time of freeze-out. This has some ramifications on the treatment of the freeze-out calculation, described further in Appendix E, but qualitatively the behavior is very similar.

Above the *quantum chromodynamics* (QCD) scale of approximately 217 MeV, quarks are unable form bound states within hadrons, labeling the following period the *quark epoch*. The entire universe was filled with a *quark-gluon plasma*, until cooling to around 10^{12} K at 10^{-5} s. At this point, the *hadron epoch* began as the strong force overtook thermal excitations and the universe hadronized.

A number of important events occurred in the hadron epoch that are worth mentioning. First, baryons were too heavy at this point to thermalize via pair production. Therefore, quickly after this epoch began, baryon-antibaryon pairs began annihilating, leaving behind the small excess of protons and neutrons created during baryogenesis. Next, at around 3×10^{-5} s, the remaining *mesons* became too massive for pair production, and also froze out. At about 5×10^{-5} s after the Big Bang, *muons* froze out, leaving only neutrinos, electrons, *positrons*, and photons in abundance.

The hadron epoch ended approximately 1 s after the Big Bang, and is marked by the freeze-out of neutrinos, protons, and neutrons. The hadron pair production processes had already frozen out, but there were still thermalizing processes between neutrons and protons, such as $n \rightarrow p\overline{\nu}_e e^-$, $n\nu_e \rightarrow pe^-$ and their reverse interactions. The energy released in neutron decay is about 800 keV [59], corresponding to a temperature of 10^{10} K, which was reached at $t \approx 1$ s.

After the freeze-out of neutrinos, the *lepton epoch* began, where electrons and positrons dominated the energy density of the universe. Pair production and annihilation via high energy photons kept thermal equilibrium until about $t \approx 4\,\mathrm{s}$. At this point the temperature fell too low to continue pair production, and the majority of the remaining leptons annihilated into photons. This left behind a relic population of excess electrons created during leptogenesis.

The lepton freeze-out led into the *photon epoch*, where almost all of the energy in the universe was in the form of photons at a temperature of about $10^9\,\mathrm{K}$. About 3 min into this epoch, the universe entered into a period of nuclear fusion known as *nucleosynthesis*, where the photons had cooled down enough to allow protons and neutrons to form into bound nuclei. During nucleosynthesis, deuterium, tritium, helium-3, helium-4, lithium-7, and beryllium-7 nuclei were produced through the fusion reactions of relic protons and neutrons. The very limited timespan of this period prevented any additional fusion interactions from occurring, which would have required more time and higher energies. The tritium population all decayed to helium-3 before the end of the epoch, and the beryllium-7 decayed to lithium-7 some time later via electron capture. After about 17 minutes, the coulomb repulsion of these nuclei overpowered the thermal fusion, and a relic population of light atomic nuclei was left behind. The interesting thing about this period is that the relic abundances of all of those nuclei relative to hydrogen is determined by only a single parameter: the baryon to photon ratio at the time of nucleosynthesis. Fig. 3.4 shows predictions of the relative abundance of each type of nuclei, as a function of this ratio, along with two experimental measurements that were made using different techniques. The incredible agreement between the predictions and observations of these light nuclei abundances was the first major success of the Big Bang model.

The photon epoch lasted about 70 kyr, at which point the matter content of the universe, which can't be arbitrarily redshifted during expansion, overtook it. This lead into the *matter epoch*, which lasted until about 380 kyr after the Big Bang when the universe had cooled to 4,000 K and stable atoms were finally able to form. This process, known as *recombination*,

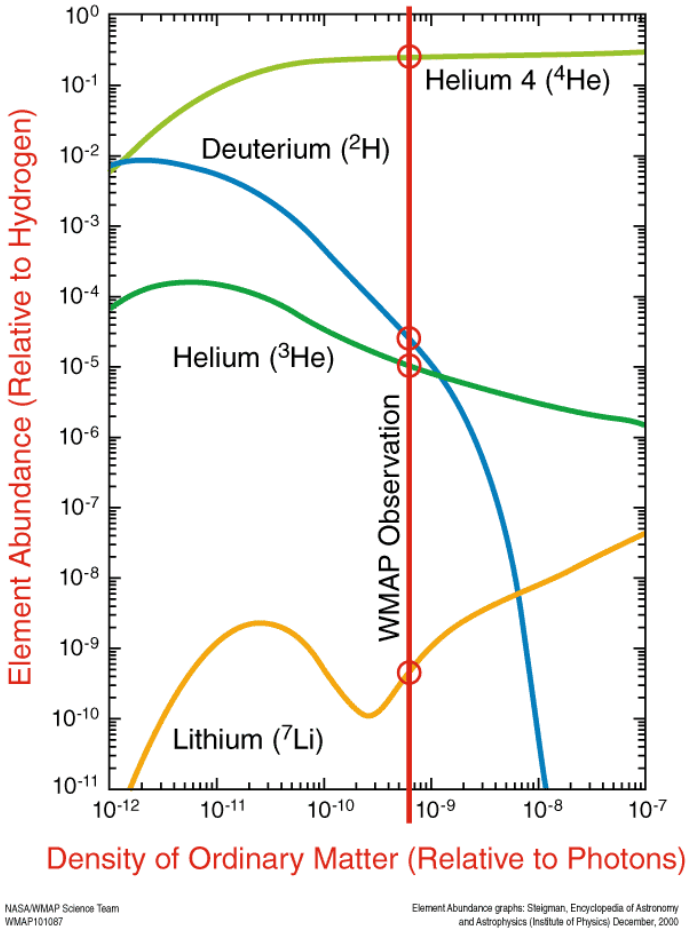

Figure 3.4: An illustration of one the major successes of the Big Bang theory. The four curves are the abundances of 4 isotopes predicted from the theory, as a function of the total baryonic matter density. The four circles are the observed abundances, and the red line is the observed baryon density from WMAP. Not only do the relic abundances line up perfectly with the predictions, but a completely separate measurement of the baryon density agrees extraordinarily well. [5]

is when the universe became electrically neutral, and therefore transparent to photons. This freeze-out left behind a relic population of photons that's known today as the *cosmic microwave background* (CMB).

The CMB is the oldest observable relic population that is still intact today, giving us a snapshot of the universe at the time of recombination. The importance of this is that just before recombination, there were three dominant forms of energy that all had comparable densities. *Cold dark matter*, the most abundant of the three, only interacts gravitationally with the rest of the universe. Baryonic matter, on the other hand, interacts

with itself and photons through the much stronger electromagnetic force. Photons, in turn, only interact with charged baryonic matter. This gives three distinct populations with very unique behaviors. The radiation pressure of the photons and the self-interaction of baryonic matter tended to smooth out any inhomogeneities in the universe, while at the same time, gravitational collapse of dark matter and baryonic matter amplified minute inhomogeneities left over from early quantum fluctuations. This led to *acoustic oscillations* that became frozen into the CMB at recombination. Because of this, measurements of the tiny anisotropies observed in the CMB today set very tight constraints on the nature of our universe.

3.1.6 Structure Formation

The 150 Myr period after recombination is known as the *dark epoch*, because nearly all of the thermal processes had frozen out and on large scales, the composition of the universe didn't change. The universe was a very homogeneous mixture of light atoms, dark matter, and photons, none of which interacted very strongly with each other. However, there were tiny inhomogeneities left over from earlier quantum fluctuations, and 150 Myr years of gravitational amplification finally allowed the formation of the first stars and super massive *black holes* destined to become the first galaxies, or *quasars*.

Over the next billion years the (mostly hydrogen) gas that made up the universe began collapsing into stars and quasars, and young galaxies began to form. As they evolved, they started producing a significant amount of *radiation*. This radiation was strong enough to ionize most of the neutral gas that had formed during recombination, naming this period the *reionization epoch*. At this point, the universe had expanded and matter had condensed quite significantly, so that any baryonic matter was sparsely distributed enough for the universe to remain mostly transparent. By the end of this epoch, the free baryonic matter content of the universe was almost entirely ionized.

One of the key predictions for distinguishing models with cold dark matter from those with *hot dark matter* is in large scale *galactic structure* formation after the reionization

epoch. Dark matter, making up the majority of the matter in the universe, plays a significant role in the formation of any gravitationally bound structure on *galactic scales* or larger. Hot dark matter moves relativistically, making the formation of small structures much more difficult than larger ones. In a universe dominated by hot dark matter, one would expect to see structure formation in a top-down fashion, where *galaxy clusters* would have formed before the galaxies that compose them. Cold dark matter, on the other hand, predicts a bottom-up structure formation where galaxy formation occurs first, and is then followed by increasingly large galactic structures.

Around 9.8 Gyr after the Big Bang, shortly after the formation of our solar system, the dark energy density of the universe overtook the matter and radiation densities. This marks the start of the *modern epoch*, which extends to today at 13.8 Gyr after the Big Bang, where dark energy fuels an exponential expansion of our universe.

3.2 Rotational Curves

The first experimental hints of dark matter came from Zwicky's analysis of the Coma galaxy cluster in 1933 [1]. Using the observed redshift of the galaxies within the cluster, he inferred that the dispersion of their velocities was over 10 times the maximum velocity allowed to keep the galaxies bound within the cluster, based on the amount of visible matter. He concluded that the invisible matter within the Coma cluster must be a few hundred times more abundant than its visible matter. Although there was additional evidence of dark matter in the following decades, the concept wasn't taken seriously until 1970 when Rubin and Ford made detailed measurements of the stellar velocities within the nearby Andromeda galaxy [2]. The distribution of visible matter within a galaxy is highly concentrated towards its center, known as the bulk. Using Newtonian gravity, which is a valid approximation at these scales, the gravitational acceleration within the bulk, at a distance r from its center, should be $MR^{-3}r$ for a galaxy with bulk mass M and radius R. Outside of the bulk the gravitational acceleration would simply be Mr^{-2}, just as

with any other spherical mass. The velocity of any object in a stable orbit of radius r with centripetal acceleration a is $\sqrt{ar}$, from which it's straightforward to show that matter inside and outside the bulk should move with a velocity proportional to r and $r^{-1/2}$ respectively. However, what was observed was a nearly constant velocity throughout the outer regions of the galaxy, as shown in Fig. 3.5.

The only way to explain this, without modifying our theory of gravity, is to conclude that a large portion of this galaxy is completely invisible to us. The only model that retains the cylindrical symmetry observed in galaxies, is one where the cumulative mass of the galaxy increases linearly with radius, corresponding to a spherically symmetric mass density falling off like $\propto r^{-2}$. The acceleration would then be Mr^{-1}, and the velocity of stable orbits would be roughly proportional to $\sqrt{M}$. By 1978, not only had Rubin's results had been independently confirmed by other groups, but she had observed the same phenomenon in 10 other spiral galaxies [60]. Since then, similar discrepancies in rotational velocities have been found within most galaxies and galaxy clusters.

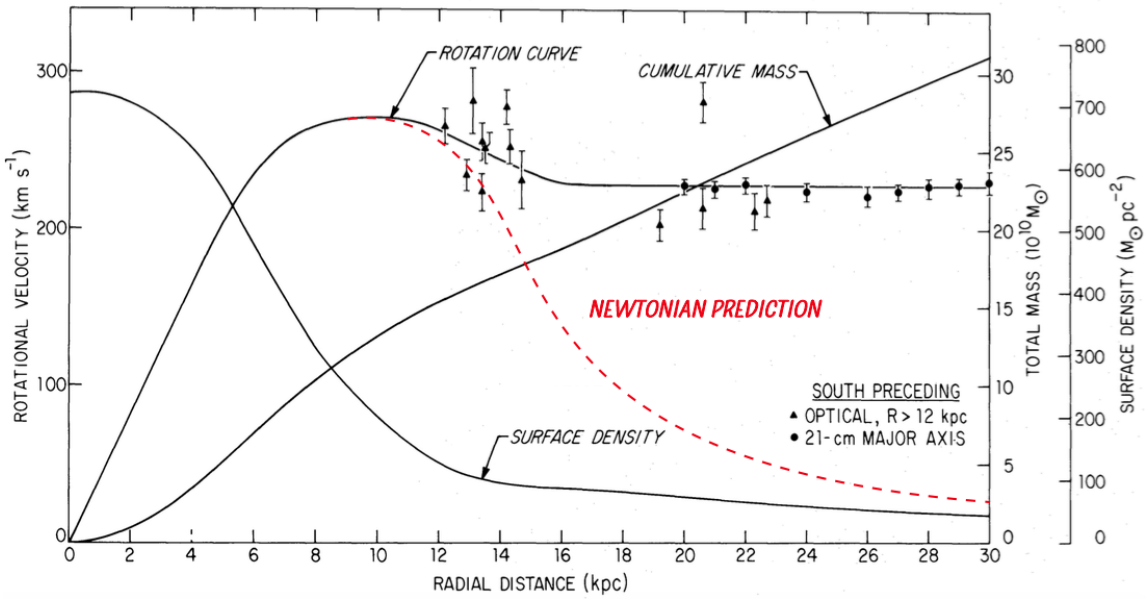

Figure 3.5: The observed rotational velocity distribution of gas and stars within the Andromeda galaxy. The inferred mass distribution is also plotted, showing a roughly linear trend with respect to radius. [61]

3.3 Gravitational Lensing

Rotation curves alone don't prove the existence of dark matter, since MOND theories predict the same phenomenon. The strongest support for the dark matter hypothesis over MOND comes from *gravitational lensing* experiments. In General Relativity, photons follows null geodesics that are predicted to, and have been experimentally confirmed to, bend around massive objects that distort the structure of space-time. This bending can produce a lensing effect where light from a point source can take multiple paths around a massive object to reach the same destination. The light's angle of deflection is proportional to the lensing object's mass and inversely proportional to the distance of closest approach. Sufficiently compact stellar objects can have such a strong effect that the light reaching us from more distance sources is visibly distorted. Since the lensing effect depends only on the gravitational mass of the matter between the observer and the source, the degree to which the light is bent can be used to make a mass measurement of that matter, even if it's invisible to us.

Strong-lensing is a phenomenon that occurs around only the most compact objects, such as black holes and *neutron stars*, where background sources are visibly distorted into arcs, rings, or multiple images of the same source. This provides a very accurate mass measurement, but can only be used under very specific conditions, and is therefore not suited for more general purposes.

A more commonly used approach, called *weak-lensing*, which looks at the minute distortions of a large sample of background sources behind the region of interest. Although any single background source isn't distorted enough to make a measurement, they can all be combined to give a statistical estimate of the mass contained within a target region. This allows the matter content of the nearby universe to be mapped out using the light from the distant universe. Images from the Hubble telescope of both strong and weak-lensing are shown in Fig. 3.6.

Both strong and weak-lensing have provided evidence supporting the dark matter hy-

Figure 3.6: The strong lensing effect on the left [a], and weak lensing on the right [b].

[a] NASA, ESA, J. Rigby (NASA Goddard Space Flight Center), K. Sharon (Kavli Institute for Cosmological Physics, University of Chicago), and M. Gladders and E. Wuyts (University of Chicago)

[b] NASA, ESA, M. Postman (STScI) and the CLASH Team

pothesis, but the strongest example that sets it it apart from the alternatives is the bullet cluster, shown in Fig. 3.7. This galaxy cluster was formed by the merging of two smaller clusters, as is evident from the X-ray images, where the shock-waves from the collision can still be seen in the cluster's gas distribution. During the collision of galaxy clusters, the galaxies within each cluster are very sparse and mostly pass through each other with little interference. The intergalactic gas of each cluster, on the other hand, is very dense and interacts quite strongly in the collision, which is why the gas lags behind. Weak-lensing around the bullet cluster shows that the center of mass is significantly displaced from center of *visible* mass, which tends to follow the gas that makes up approximately 90% of the baryonic matter in a galaxy cluster. This observation is very difficult to reconcile with MOND theories, which can't avoid the prediction that the center of mass would follow the baryonic matter.

A third type of lensing is known as *micro-lensing*, and occurs when a dim star or large planet briefly passes in front of a nearby light source. These objects are not compact enough to create any resolvable spatial displacement, but as they pass by they amplify the light

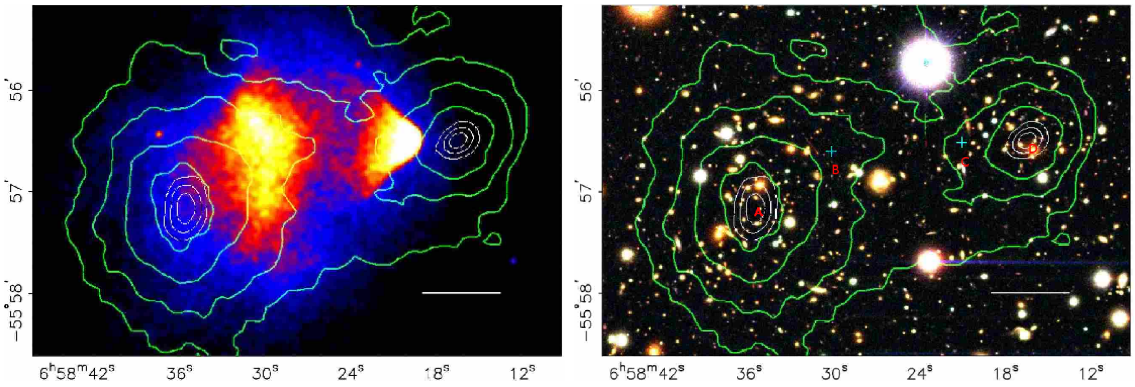

Figure 3.7: The bullet cluster with mass contours superimposed, derived from weak lensing. On the right is the visible light coming from the galaxies within the cluster, and on the left is an X-ray image showing the intergalactic gas distribution. Most of the cluster's visible mass comes from the gas, but the mass contours coincide with the galactic distributions. [3]

source's luminosity in a unique way. Measuring this brightening as a function of time can provide information about the object's mass. Micro-lensing surveys [7, 62, 63] search for baryonic MACHOs in the form of small stars and planets, which are expected to make up the majority of the MACHO density. They have all placed very strong limits on the MACHO content within our *galactic halo*, ruling out the possibility of dark matter being composed *entirely* of baryonic matter. The exact percentage they account for is debated, and very model-dependent, but going by micro-lensing data alone, they are very unlikely to make up 100% [62] and the latest estimates put them at under 8% [7].

3.4 Large Scale Structure

As mentioned earlier, the best way to differentiate between cold dark matter and hot dark matter is to observe how galactic structures formed during the evolution of the universe. Although at the largest scales the universe appears homogeneous, its matter distribution remains "clumpy" well beyond the galactic scale. It's well known that gas collapses into stars, and stars form into galaxies, but galaxies also form into increasingly larger populations known as groups, clusters, and superclusters. These large scale galactic structures form a web-like pattern with gigantic regions of space in between that are almost entirely

void of matter. Being the dominant form of matter in the universe, dark matter plays a crucial role in the formation of these structures, so that their evolution over time can give us insight into properties of dark matter.

To illustrate the difference between cold dark matter and hot dark matter, let us consider a homogeneous ideal gas distribution at thermal equilibrium, with radius R, mass density ρ, and temperature T. The average velocity of each particle is related to the temperature by $3T = mv^2$, where m is the particle mass. The escape velocity at radius $r \leq R$ is easily shown to be,

$$v_e = r\sqrt{\frac{4}{3}\pi\rho}. \tag{3.25}$$

Requiring that the $v < v_e$, for the gas to be gravitationally bound, gives,

$$r > \frac{3}{2}\sqrt{\frac{T}{\pi\rho m}}. \tag{3.26}$$

The matter at any radius beyond this bound will experience gravitational collapse, while the matter within it will be dominated by thermal forces. This is an idealized scenario, but it shows the temperature and mass dependence of gravitational collapse.

For hot dark matter, such as neutrinos, the dark matter particles are moving relativistically and any small-scale inhomogeneities in the early universe would have been quickly washed out. Only very large structures would be able to form at first, followed by smaller sub-structures as the universe cooled. On the other hand, in a cold dark matter model, where the dark matter particles were moving non-relativistically even in the early universe, quantum fluctuations would be amplified by gravitational collapse into small-scale structures in an otherwise homogeneous universe. These would then continue to grow into increasingly larger structures. Hot dark matter can be associated with top-down, and cold dark matter with bottom-up structure formation. The Millennium Simulation Project was a simulation of the $\Lambda - \mathrm{CDM}$ model performed with over 10 billion "particles" at the Max Planck Institute [64]. Fig. 3.8 shows the final result of the simulation, and Fig. 3.9 shows

the evolution of the matter distribution over time. These results support the idea that cold dark matter produces bottom-up structure formation, and results in large scale web-like galactic structures.

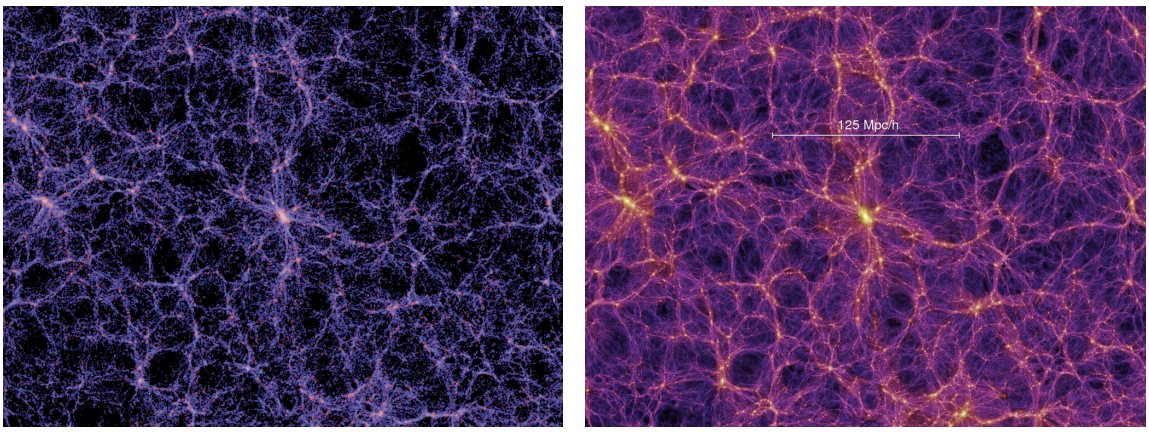

(a) Baryonic Matter **(b)** Dark Matter

Figure 3.8: The results of the Millennium Simulation for the matter distributions of the universe today, as predicted by $\Lambda - \mathrm{CDM}$. [64]

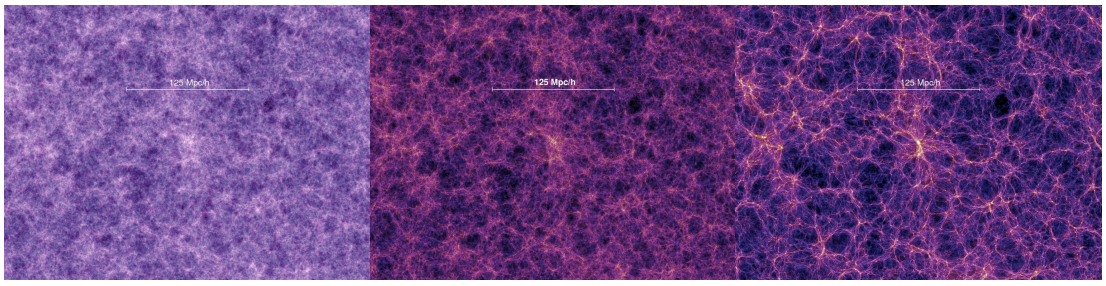

(a) 210 Myr **(b)** 1 Gyr **(c)** 4.7 Gyr

Figure 3.9: The results of the Millennium Simulation for the dark matter distribution of the universe at three different times after the Big Bang, as predicted by $\Lambda - \mathrm{CDM}$. [64]

We can compare these predictions to a three dimensional map of our universe, obtained by measuring the redshift of distant galaxies. Since in our expanding universe, redshift determines both time and distance, the third dimension of this map can be treated as both a spatial axis and a timeline of the universe. The 2dF Galaxy Redshift Survey [65] surveyed the redshift of known galaxies in order to calculate their distance, producing the results shown in Fig. 3.10. The formation of web-like structures similar to those observed in the

Millennium simulation can be clearly seen in these results. Above the redshift of $z > 0.2$, or 2 Gyr ago, these structures disappear, and the universe appears to be a nearly homogeneous collection of galaxies .

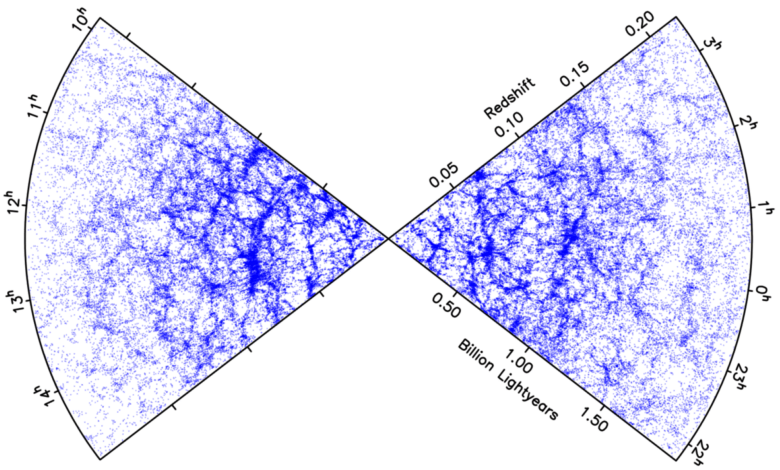

Figure 3.10: The results of the 2dF galaxy Redshift Survey, showing the evolution of galactic structure over the last 2 Gyr. This result supports the $\Lambda-$CDM model, where galaxies form in a roughly homogeneous pattern and then group into increasingly large structures. [65]

A more rigorous approach can be taken with respect to the large scale structures, by quantifying the observed anisotropies. These measurements show *baryon acoustic oscillations* (BAOs), which are related to oscillatory interactions between baryons and other forms of energy in the universe [66]. This is discussed further in the next section, with the analogous CMB anisotropies.

3.5 Cosmic Microwave Background

The relic abundance of photons that make up the CMB provide us with one of the earliest glimpses into the Big Bang. The freeze-out occurred at about 380 kyr, when the universe had a temperature of about 10^3 K. In theory, there should also be a cosmic neutrino background, created 1 second after the Big Bang, but only indirect evidence for its existence

has been found so far. This population today consists of incredibly low energy neutrinos, which are nearly impossible to detect. The CMB, on the other hand, interacts much more strongly and is easy to detect. Today, it has a temperature of approximately 2.7 K [67] and is made up of an almost perfectly homogeneous and isotropic black body spectrum. From its interactions with the rest of the universe, it has acquired anisotropies on top of any initial ones present at the time of recombination. These minute deviations from perfect isotropy can provide a wealth of information and tightly constrain the parameters of $\Lambda - \mathrm{CDM}$.

The CMB spectrum is measured by sampling photons coming from different directions in the sky. This produces a projection of a 3-dimensional density onto a 2-dimensional spherical shell. The measured quantity, $\mathcal{O}$, which is typically either temperature or polarization, can be parameterized as a function of the spherical coordinates θ and ϕ. The spectrum can then be decomposed into spherical harmonics using the equation [28],

$$\Delta\mathcal{O}(\theta,\phi) \equiv \sum_{\ell}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta,\phi), \tag{3.27}$$

where $Y_{\ell m}$ are the spherical harmonic functions defined by,

$$Y_{\ell m} \propto e^{im\phi} P_{\ell m}\left(\cos\theta\right), \tag{3.28}$$

and $P_{\ell m}$, the associated Legendre polynomials, satisfy the differential equations

$$\frac{d}{dx}\left[\left(1-x^2\right)\frac{d}{dx}P_{\ell m}(x)\right] + \left[\ell(\ell+1) - \frac{m}{1-x^2}\right] P_{\ell m}(x) = 0. \tag{3.29}$$

The functions $Y_{\ell m}$ satisfy normalization and orthogonality requirements,

$$\int_0^{2\pi}\int_{-1}^{1} d(\cos\theta)d\phi Y_{\ell m}(\theta,\phi)Y^*_{\ell' m'}(\theta,\phi) = \delta_{\ell\ell'}\delta_{mm'}, \tag{3.30}$$

so that they can be used to decompose any angular function into the coefficients $a_{\ell m}$. The

CMB *power spectrum* is then given by,

$$C_\ell \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2 \,, \tag{3.31}$$

which is simply a function of ℓ, the multipole moment. The constant term C_0 is just the isotropic average, which gives no information about anisotropies, and the dipole term C_1 is dominated by relativistic Doppler effects, due to our relative motion with respect to the CMB rest frame. Therefore, when looking at the CMB anisotropies these two terms are typically excluded and only $\ell \geq 2$ are considered.

Fig. 3.11 shows the measured CMB power spectrum, along with the $\Lambda-$CDM prediction fit to data. The various peaks in the spectrum are called *acoustic peaks*, and their location and size provide important information about the cosmological parameters. The precise relationship between the predicted power spectrum and the parameters is fairly complex, and model-dependent, but fitting them to data sets very tight constraints. Multiple parameter sets have been shown to describe identical CMB power spectra [68] though, so results from other experiments are typically used to distinguish between degeneracies in the fit.

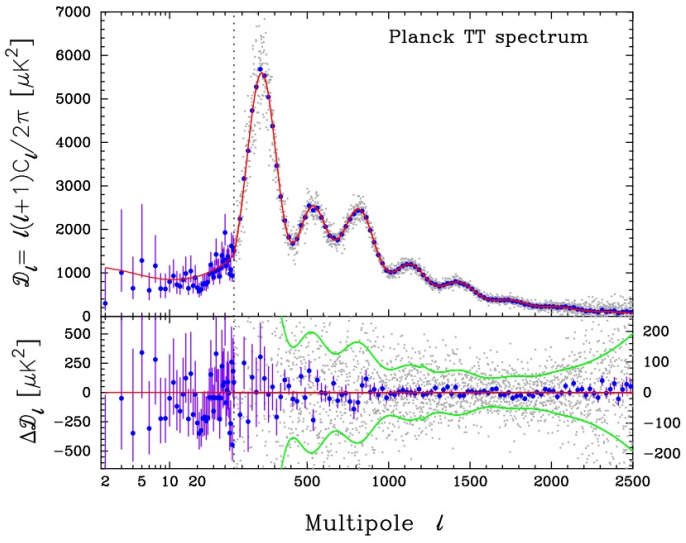

Figure 3.11: Results from the PLANCK experiment showing the CMB anisotropy power spectrum. The best fit $\Lambda-$CDM model corresponding to the parameters in Table 3.1 are superimposed on top of the data. [53]

3.6 Supernovae

As mentioned above, the CMB measurements alone can't constrain a unique set of $\Lambda-$CDM parameters. This is illustrated in Fig. 3.12, where the constraints set by three very different observations are plotted. The CMB measurements provide very tight constraints on the curvature of the universe, showing that it's nearly flat, but they say very little about the ratio between the dark energy and matter densities. BAO measurements, on the other hand, constrain the matter density of the universe, but give little information about the dark energy content or the curvature of the universe. A third independent measurement comes from observations of *type Ia supernovae*, which provide constraints on the expansion of the universe that complement the other two very well [6].

Unlike most *supernovae*, which result from the gravitational collapse of a single massive star, type Ia supernovae are believed to be caused by a thermonuclear explosion of low mass stars in binary systems, that consequently rip them apart [69]. *White dwarfs* are stable, low mass stellar remnants whose internal electron degeneracy pressure prevents any further collapse. The evolution of a star into a white dwarf leaves a dim remnant primarily composed of carbon and oxygen, without enough energy to undergo heavier fusion reactions. The *Chandrasekhar limit* [70] of about 1.4 solar masses represents the mass at which the gravitational force overcomes this degeneracy pressure and the star is able to continue its collapse. In a binary system where a white dwarf is able to gain mass from its partner star, its mass can gradually reach the Chandrasekhar limit. At this point the internal temperature rises to a level where carbon fusion suddenly becomes possible, and the star undergoes a runaway fusion reaction. In a matter of seconds enough energy is built up to completely disintegrate the star, and a type Ia supernovae occurs. A number of different binary systems can produce these supernovae, but since the mechanism behind the explosion is always the same, the characteristics of the supernovae create a very consistent signature.

Because of the consistency of these supernovae, they are typically referred to as *standard candles*. They all should have roughly the same total brightness, so the distance to a given

supernova can be calculated from the *observed* brightness simply by taking advantage of the inverse square law. This allows us to obtain very accurate distance measurements for these events, independent of any assumptions about the nature of our universe. The elemental composition of these supernovae is very well known, since the initial composition is primarily carbon and oxygen. After the carbon fusion, oxygen is the lowest mass element and elements up to iron can be produced through fusion reactions. Matching the spectral lines from these elements provides an accurate measurement of the redshift parameter, z. The emission time of the light received from the supernova is simply $t_E = d$, where d is the distance to the supernovae, and using Eq. 3.9 to obtain the expansion rate of the universe at time t_E,

$$a(t_E) = \frac{a(t_0) - 1}{z}. \tag{3.32}$$

By repeating this calculation for supernovae at various distances, the Hubble parameter $H(t)$ can be estimated and the relative densities of radiation, matter, and dark energy can be fit using Eqns. 3.12. Combining all of these observations results in the very precise fit of the $\Lambda - \mathrm{CDM}$ parameters shown in Table 3.1.

Parameter	Symbol	Value
Universe age	t_0	$(13.799 \pm 0.021)\,\mathrm{Gyr}$
Physical baryon density	$\Omega_b h^2$	0.02230 ± 0.00014
Physical dark matter density	$\Omega_c h^2$	0.1188 ± 0.0010
Scalar spectral index	n_s	0.9667 ± 0.0040
Reionization optical depth	τ	0.066 ± 0.012
Density fluctuations at $8h^{-1}$Mpc	σ_8	0.8159 ± 0.0086
Dark energy density	Ω_Λ	0.6911 ± 0.0062
Matter density	Ω_m	0.3089 ± 0.0062
Age at recombination	t_*	$(377.7 \pm 3.2)\,\mathrm{kyr}$
Hubble constant	H_0	$(67.74 \pm 0.46)\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$

Table 3.1: The results reported by PLANCK that very precisely measure the $\Lambda - \mathrm{CDM}$ parameters. In addition to the CMB results, BAO, lensing, and type Ia supernovae observations were used to fit these parameters. The lower half of the table are derived values of interest. [53]

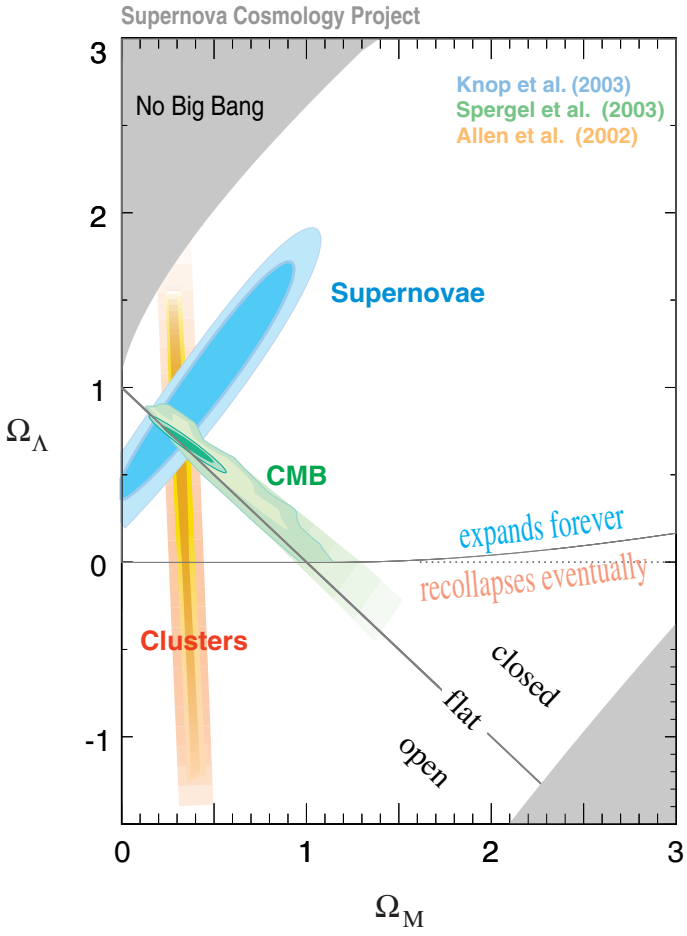

Figure 3.12: The combined results from supernovae [6], CMB [5], and galaxy clustering [66] observations. The x-axis is the total matter density, and the y-axis is the total dark energy density. The three experiments overlap in a region that is nearly flat, originated in a Big Bang, and will expand forever. [71]

Chapter 4

The ATLAS Experiment

4.1 The Large Hadron Collider

At everyday low energies, the electromagnetic and weak forces look vastly different. The electromagnetic force is mediated by a massless, chargeless, gauge boson, which allows interactions over very large distances. The weak force, on the other hand, is a short-ranged force mediated by three massive, electrically charged bosons, the $W^{\pm}$ and Z^0 fields. Their high mass makes the strength of this force fall off exponentially, and also explicitly breaks any possible gauge symmetry. However, the low energy theory governing the weak force breaks down at TeV scale energies. The Higgs boson was originally introduced to solve this problem, and unifies the two forces into one electroweak force mediated by four massless gauge bosons. It's only at low energies that the weak bosons *appear* to be massive due to the VEV of the Higgs field. Although there are alternative theories, they all share one thing in common: something new must show up at the TeV scale to break the electroweak symmetry. The *large hadron collider* (LHC) was built in order to probe this scale through high energy collisions, with the goal of getting a better understanding of the electroweak force [72].

There are many different options to consider when designing a high energy particle collider, and each has its own benefits and drawbacks. Both the particles chosen, and the method used to accelerate them have a huge impact on the cost of running the collider, the energy of the collisions, and their frequency. One of the earliest particle accelerators was the Cockroft-Walton generator, which in 1932 was used to disintegrate atomic nuclei

for the first time [73]. This device used an electrostatic field to accelerate protons to 0.7 MeV and collide them into a stationary lithium target. Electrostatic accelerators are very simple, but also very limited due to the difficulties of maintaining a high DC current. The method of acceleration used by all modern high energy accelerators is known as *radio frequency* (RF) acceleration. In this method, RF cavities are built along the beam line that produce oscillating electromagnetic fields, roughly in the RF range. The frequency of these oscillations are tuned to the frequency of incoming particles, and are fixed at resonant frequencies within the cavity to create standing waves. As a group of particles passes through the cavity, they can be accelerated, decelerated, compressed, or expanded just by adjusting the phase of the waves.

Linear accelerators (linacs) only have a single pass to accelerate particles, so in order to achieve high energies they need to be very long and very strong. For reference, SLAC currently has the highest energy linac, which accelerates electrons and positrons up to 50 GeV over a distance of 3 km. For a general purpose probe into the high energy realm, a circular accelerator is more appropriate, as it can be built more compactly and offers a wider range of collision energies. Circular accelerators can accelerate particles through multiple passes of the ring, allowing for very high energies to be achieved in relatively small rings.

The largest obstacle for any particle accelerator is the radiation reaction of electrically charged particles under the influence of an external force. Electrically charged particles radiate off energy as they're accelerated, making it increasingly difficult to achieve higher energy collisions. The relativistic Larmor formula [25] is,

$$\frac{dE}{dt} = \frac{2}{3}\alpha\gamma^6 \left(\left| \frac{d\vec{v}}{dt} \right|^2 - \left(\frac{d\vec{v}}{dt} \times \vec{v} \right)^2 \right), \tag{4.1}$$

which gives the radiated power of a particle in terms of its velocity, $\vec{v}$, and Lorentz factor, γ. Given that $\vec{p} \equiv \gamma m \vec{v}$ and $E \equiv \gamma m$, it's straightforward to show that for a relativistic

particle in circular motion, where $\dot{\vec{v}} \cdot \vec{v} = 0$,

$$\frac{dE}{dt} \approx \frac{2}{3}\alpha \frac{E^2}{m^4} \left| \frac{d\vec{p}}{dt} \right|^2, \tag{4.2}$$

and in linear motion where $\dot{\vec{v}} \times \vec{v} = 0$,

$$\frac{dE}{dt} \approx \frac{2}{3}\alpha \frac{1}{m^2} \left| \frac{d\vec{p}}{dt} \right|^2. \tag{4.3}$$

Each of these equations describe the power lost to radiation during circular and linear acceleration respectively, as a function of the applied force. In a collider experiment, this radiation is viewed as wasted energy and can severely limit its reach. One of the advantages to linacs is that they experience less radiation, by a factor of $\gamma^2 = E^2/m^2$. For a fixed target energy, this effect will be dramatically higher for lighter particles.

Any stable, electrically charged particle can be accelerated and collided, such as electrons, positrons, charged hadrons, and atomic nuclei. The important factors to consider in choosing a particle is abundance, mass, and composition. Anti-matter, which needs to be manufactured in separate collisions, is the least abundant choice and severely limits the rate of collisions. As mentioned above, electrons produce much more radiation than the heavier options, which limits the maximum energy of a detector on a fixed budget. Composite particles such as protons and atomic nuclei produce very complex collisions, making analyzing the data significantly more difficult. Collisions occur between constituents which only contain a fraction of the total energy, and glancing collisions frequently occur which fragment particles and create a very noisy background. All of the known composite particles are made of quarks and gluons, whose color charge creates a number of additional difficulties in reconstructing collisions.

LEP [74] was an electron-positron collider, which ran for 11 years in the tunnel now housing the LHC. It was unable to reach the energy scale of electroweak symmetry breaking, with a maximum collision energy of 209 GeV and a total *integrated luminosity* under $1\,\mathrm{fb}^{-1}$. The Tevatron [75] was a proton-antiproton collider in the US that was the highest

energy collider until the LHC, with a 1.96 TeV center of mass collision energy. However, the machine's *luminosity* was limited by the need for antiprotons, which are difficult to produce. The *peak luminosity* of the Tevatron was about $300\,\mu\text{b}^{-1}\text{s}^{-1}$, with a total integrated luminosity of about $10\,\text{fb}^{-1}$ over 10 years. The LHC, in contrast to both of these, was built to provide maximal energy collisions at the highest luminosity possible by using proton-proton collisions. Although the collisions pose a number of technical difficulties, the energies and luminosities achieved are significantly higher than any collider in the past. The LHC was designed to collide protons at an energy of 14 TeV with a peak luminosity of $10\,\text{nb}^{-1}\text{s}^{-1}$. In the 2012 *run*, collisions were performed at 8 TeV and delivered about $20\,\text{fb}^{-1}$ worth of usable data to each of its detectors, almost double what the Tevatron collected in its 28 year lifetime.

The protons used for LHC collisions are collected and accelerated using the complex *injection* chain illustrated in Figure 4.1. Hydrogen is first ionized using a *duoplasmatron*, which strips the electrons from hydrogen atoms using a strong electric field. The protons isolated using this method are then accelerated to 50 MeV using the linac2 accelerator. From there, the protons are injected into the Proton Synchrotron Booster, which accelerates them to 1.4 GeV, followed by the Proton Synchrotron, accelerating them further to 25 GeV. Before they can be injected into the LHC, they are accelerated by the Super Proton Synchrotron to an energy of 450 GeV. Finally, the protons are put into the two LHC beams, which have a nominal energy of 7 TeV each. The protons are grouped into *bunches* of about 100 billion protons each that circle the 27 km ring at nearly the speed of light. The bunches are nominally spaced 25 ns apart with gaps left for beam injection and *dumping*, leaving room for a maximum of 2808 bunches in each beam. The two proton beams intersect at 4 points along the ring, corresponding to the 4 experiments that analyze the collisions: ATLAS, CMS [76], ALICE [77], and LHCb [78]. ATLAS and CMS are general purpose detectors designed to search for any NP signatures. LHCb is a specialized detector designed to measure CP violation in bottom decays, and ALICE is used in heavy ion collisions to probe the strong force.

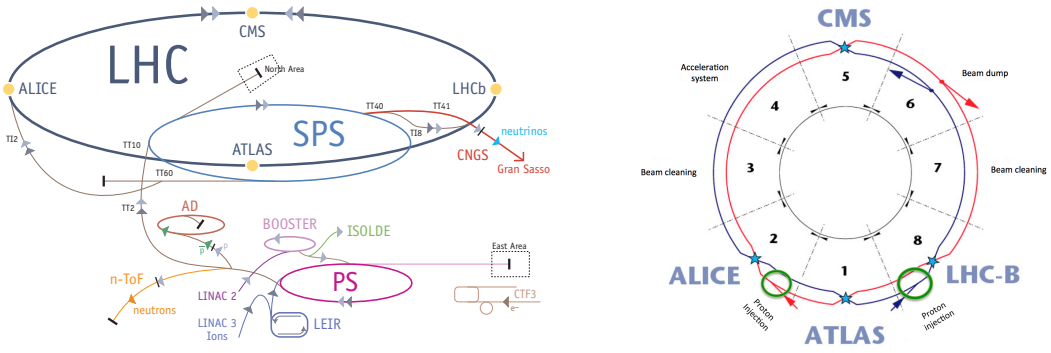

Figure 4.1: Schematics for the LHC operation. On the left is the injections chain for the protons, and on the right is the layout of the two beams and where they cross. [72]

4.1.1 Collider Physics

The physics of hadron collisions can be very complicated due to the nature of the strong interactions that dominate. To start with, at high enough energies, collisions will occur between *partons* within the protons rather than between the protons as a whole. Although hadrons are typically defined by their *valence quarks*, transient *sea* quark/anti-quark pairs and gluons are frequently produced at small distances. At LHC energies, the proton is no longer accurately described as simply a down and two up quarks, and the virtual sea particles become important, as illustrated in Fig. 4.2. Partons consist of gluons and the five lighter quarks, with each owning some fraction, x, of the total 4-momentum, p, of the proton. The presence of a parton can be parameterized by using a *parton density function* (PDF), $f_i(x, Q^2)$, for each parton species i and energy scale Q^2, which represents the probability of finding particle i with 4-momentum xp in a collision with an energy transfer of Q. Fig. 4.3 shows examples of the predicted proton PDFs at two different energy scales.

Although the most common type of collision at the LHC is low energy proton-proton scattering, it's the high energy parton-parton scattering events that are of interest. We know that each proton has a fixed energy, of 4 TeV in 2012, but the parton energies can range anywhere from 0 to about 1/3 the proton energy. Without knowing the 4-momentum of the

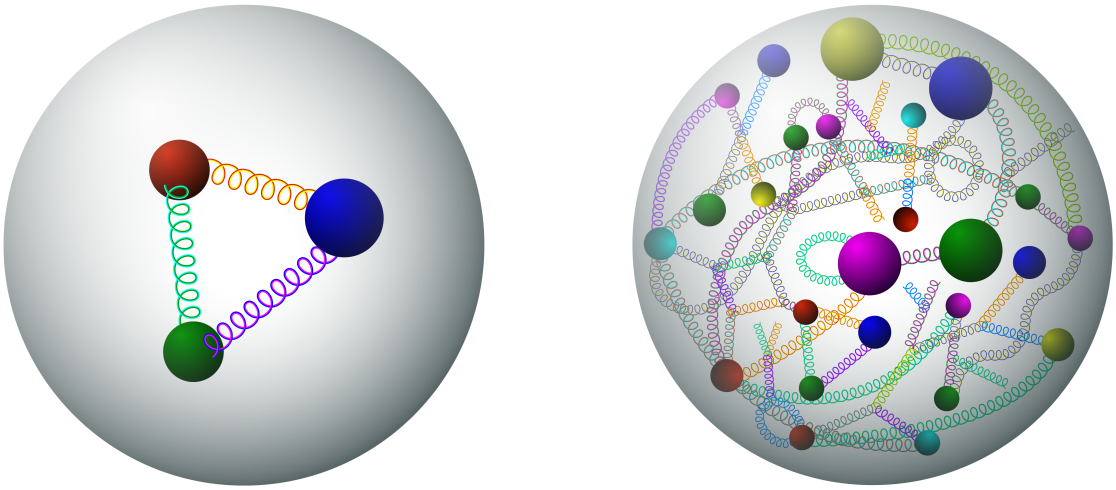

(a) Low Energy **(b)** High Energy

Figure 4.2: Illustration of proton substructure [79] showing how at higher energies, sea quarks and gluons become relevant. On the left is the usual description of a proton at low energies, which can be treated as three bound quarks. In LHC collisions though, protons more closely resemble the image on the right, where the valence quarks become less important. Strange, charm, and bottom quarks can all be found within these protons in virtual quark/anti-quark pairs, giving a wide array of possible initial states.

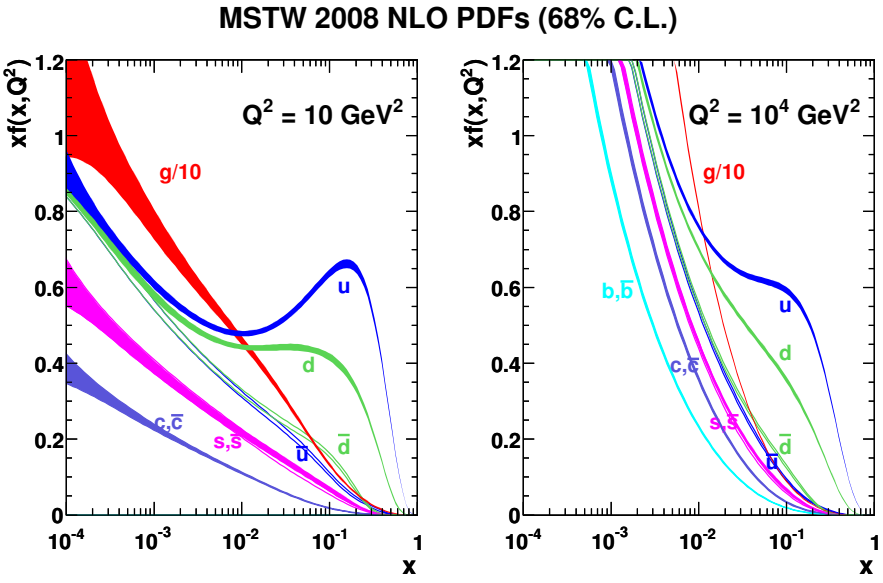

Figure 4.3: The proton PDFs at two different energy scales [80]. The characteristic enhancement of up and down quarks is due to the valence quarks, which become less prominent at higher energy scales.

initial state, reconstructing the final state can be more difficult. Fortunately, because each parton's momentum is $\vec{p_i} \sim x\vec{p}$, the total momentum *transverse* to the beam-line vanishes in any collision. This fact allows for the partial reconstruction of invisible particles, such as neutrinos.

In addition to the initial state complexities, the final states of hadronic collisions can be very difficult to work with, as explained by Fig. 4.4. Protons are color neutral, but their partons are not. When high energy interactions occur between partons, they're ejected from their parent proton, temporarily leaving bare color charges. As they separate from the rest of the proton, the strength of the strong force increases rapidly. At some point, it becomes energetically favorable for a particle/anti-particle pair to be produced in the space between, neutralizing and slowing down the charges. This process continues until all of the color charges are bound within hadrons, producing a stream of new particles that is known as a *jet*. Heavier hadrons produced in jets can decay quickly, creating even more particles. Jets are the most common, and the most complicated objects produced in LHC collisions.

4.2 The ATLAS Detector

The ATLAS detector [82] is located at Point 1 along the LHC ring, and is designed as a general purpose particle detector. It's composed of four subdetectors, organized into concentric cylindrical structures, that can identify and make precision measurements of nearly every type of particle produced in collisions. Figure 4.5 shows a cutout of the detector, which may help in the visualization of the following sections. ATLAS and each of its subdetectors are split into a cylindrical barrel region around the beampipe, and two circular endcaps closing it off on either end.

4.3 Coordinate Systems

There are a few different coordinate systems used in ATLAS. The most basic is a local, right-handed Cartesian system where the z-axis points along the beampipe towards

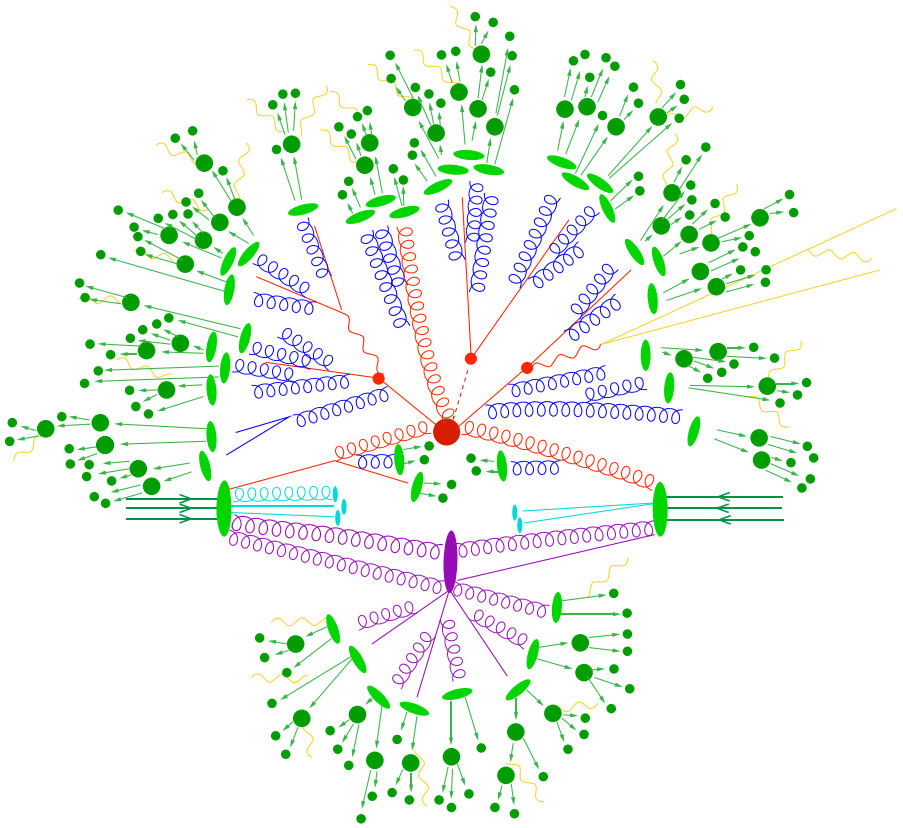

Figure 4.4: An illustration of the complexity of hadronic collisions. The hard interaction of interest to physics analyses is in red, obscured by all the other interactions. Parton showering (blue) and hadronization (light green) are soft QCD processes originating from the primary quarks and gluons, producing jets. Heavier particles decay (dark green) before reaching the detector, and interactions with the detector material can produce additional particles (orange). Secondary interactions (purple) can occur between additional partons within the primary protons, and the non-interacting partons (teal) left behind produce soft backgrounds that are correlated with the hard collision. There can also be uncorrelated soft jets from pileup collisions between additional protons in the same or nearby bunches. External radiation, such as cosmic rays or radioactive decays within detector material, create additional backgrounds completely uncorrelated with the collision. [81]

Point 8, and the x-axis points towards the center of the LHC ring. This can be useful for describing the detector, but it does a poor job of dealing with collision products. Instead, a coordinate system based on cylindrical coordinates is preferred, which takes advantage of the symmetrical nature of collisions. The z-axis is still directed along the beampipe, and the ϕ-axis is the ordinary azimuthal angle. However, instead of the usual ρ-axis, a second angular coordinate, η, is used. This coordinate is known as *pseudo-rapidity*, and is defined

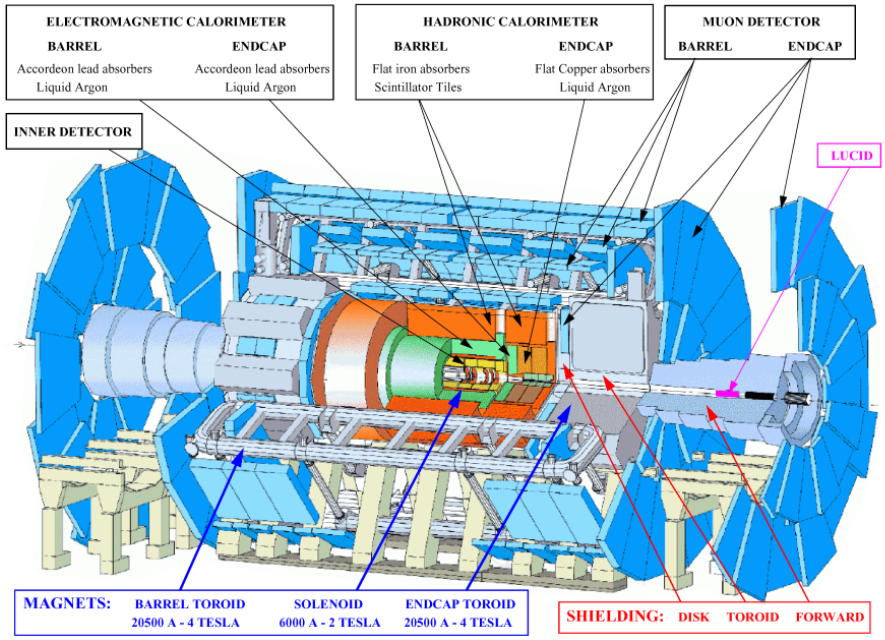

Figure 4.5: A cutout of the ATLAS detector showing its layout and all of its subdetectors [82]

in terms of the angle with respect to the z-axis, θ, as,

$$\eta \equiv -\ln\left(\tan\frac{\theta}{2}\right). \tag{4.4}$$

There are two benefits to using this coordinate system to describe the typical collision. The first is that the probability distribution of collision products, to first order, is flat as a function of both η and ϕ. More importantly though, both angular coordinates transform linearly under Lorentz boosts along the z-axis. This means that the *difference* between any two η or ϕ coordinates is invariant under such boosts, which removes much of the dependence on the unknown initial state longitudinal momentum. A related variable that is frequently used is ΔR, which defines an angular cone whose boundary is a circle in η/ϕ coordinates with $\Delta R^2 = \Delta\eta^2 + \Delta\phi^2$. This variable is usually used for parameterizing distances between particle tracks, and is also invariant under Lorentz boosts.

In addition to the Cartesian and cylindrical coordinate systems defined above, there is a third system used specifically for reconstructed particle tracks. There are five coordinates

for any given track, $(d_0, z_0, \phi_0, \theta_0, q/p)$, which encode all of the information necessary to reconstruct a single particle track going through the detector. The parameters d_0 and z_0 describe the location of the *perigee*, where d_0 is the distance between the perigee and the z-axis, and z_0 is the z distance of the perigee from the *interaction point* (IP). Particles not originating from the primary interaction vertex, such as *cosmic rays* or *pileup*, typically have large d_0 or z_0 values, giving these powerful discriminating power. Emerging from the perigee, ϕ_0 and θ_0 represent the direction of the track. For straight tracks, these four coordinates alone are sufficient to describe the entire track. However, charged tracks curve through the magnetic fields within ATLAS so the q/p parameter is necessary to determine the degree of curvature. Knowledge of the magnetic field distribution within the detector allows these coordinates to uniquely describe any track.

4.4 Operation

All particle detectors are based on three simple principles. The first is that high energy charged particles ionize atoms through inelastic collisions as they pass by. If the energy loss is relatively small, this effect can be used to track particles with minimal interference by measuring the current induced by the ionized electrons in a nearby circuit. The second principle is that charged particles curve in magnetic fields as a function of their momentum transverse to the field. By placing a magnetic field around a tracking detector, a transverse momentum measurement can be made. The last principle is that particles incident on dense materials, such as steel or lead, will generate particle *showers*. The constituents of these showers will have lower energy than the original particle, and the charged ones can be detected by making use of the first two principles. Neutral products will either continue to shower or go undetected. The size and shape of these showers are determined by the material used, the energy of the original particle and its type. By combining these three principles, complex machines such as ATLAS can be designed to precisely identify and measure incident particles. Figure 4.6 shows a schematic view of the unique signatures left

by the various SM particles produced in collisions.

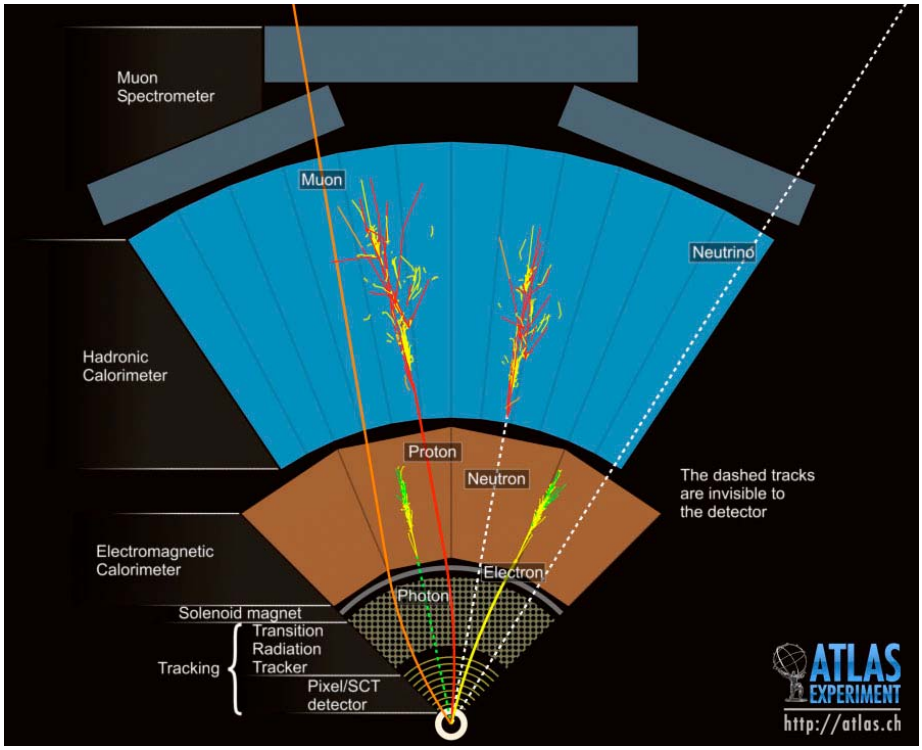

Figure 4.6: A schematic diagram of the signatures left in the ATLAS subdetectors by a representative sample of all the known particles that reach the detector. Dashed lines represent regions where a particle leaves no measurable track, and solid lines represent regions where particles can be tracked through energy deposits. Typical particles produced in collisions leave a unique signature across all the subdetectors, allowing for both identification and reconstruction. [82]

The primary goal of a general purpose particle detector is to be able to identify and reconstruct the 4-momentum of every incoming particle. If this were done perfectly, all accessible information about a collision could be reconstructed from this information. No detector is perfect of course, but ATLAS was built with this goal in mind and performs extraordinarily well. Because any particle with a mass comparable to the collision energies would decay well before reaching the detector, a zero mass assumption can be made during reconstruction. Typically the only particles stable enough to reach the detector are photons, electrons, muons, and light hadrons, the heaviest of which are about 1 GeV. Of course, if a particle of known mass can be identified, corrections can be included, but the effect is usually negligible considering the high energy scale of collisions. Since $m^2 = E^2 - p^2$, only three

of the 4-momentum components need to be measured after making this assumption. The three components commonly measured in reconstruction are the 2-dimensional transverse momentum vector ($\vec{p}_T$) and pseudo-rapidity (η).

The exception to the discussion above is neutrinos, which are impossible to detect with ATLAS because their interaction cross-section is far too low. In this case, which also applies to WIMPs, the concepts of missing transverse momentum and energy ($\vec{\not{p}}_T$ and $\not{E}_T$) are introduced. Because the longitudinal momenta of the two partons involved in the collision are unknown, and any number of high energy particles can escape if they're at a small enough angle relative to the beam pipe, momentum conservation can't be applied to the detected collision products. However, the component of momentum *transverse* to the beam vanishes for all of the initial-state partons. The closer to the beam line a collision product gets, the lower its p_T is relative to its energy. With a large enough η coverage, the p_T of these escaping products becomes very small and the total p_T of the *measured* products should nearly vanish in the absence of neutrinos. For example, a particle with the near-maximal energy of 2.5 TeV and an η of 5, which is the maximal η coverage of ATLAS, would only have a p_T of 34 GeV. Neutrino production is inferred by taking the vector sum of every measured particle's p_T, $\vec{\not{p}}_T \equiv -\sum \vec{p}_T$, and comparing it to the expected value of 0. If this number is above some threshold, it's likely that there were neutrinos produced, since other phenomenon that contribute significantly to $\vec{\not{p}}_T$ are very rare. Although it's not the full 4-momentum, $\vec{\not{p}}_T$ provides two components that can be very useful in analyses. $\not{E}_T$ is the vector norm of $\vec{\not{p}}_T$, and is much more commonly used due to the axial symmetry of collisions and the detector.

4.5 Inner Detector

The *inner detector* (ID), shown in detail in Fig. 4.7, is responsible for tracking any electrically charged particles emerging from a collision. A solenoidal magnet around the ID provides a longitudinal magnetic field that bends charged particles in the transverse plane.

By taking advantage of the ionizing behavior of charged particles, the semiconducting pixel detectors, the *semiconducting tracker* (SCT), and the *transition radiation tracker* (TRT) can be utilized to track their paths as they escape the IP.

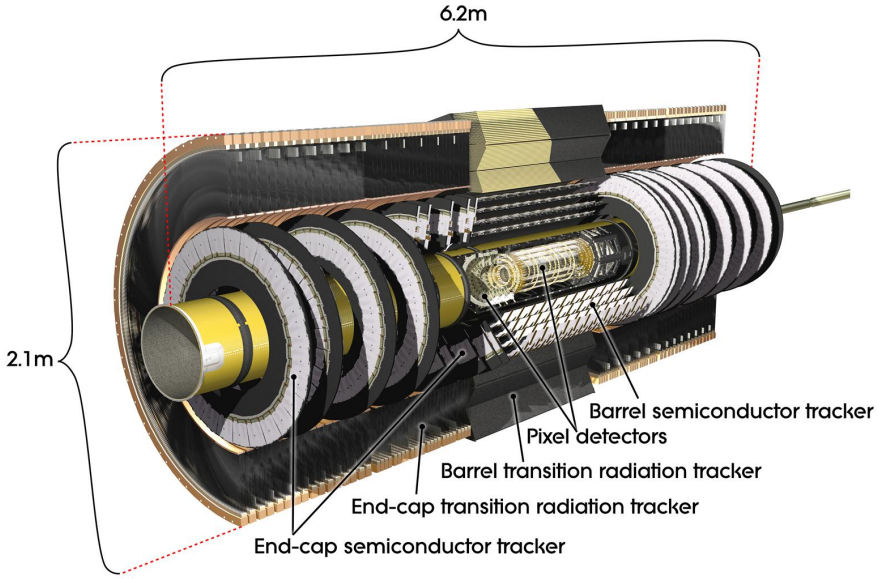

Figure 4.7: A scale cutout of the inner detector in ATLAS, with each of the subdetector systems visible. [82]

The precise tracking measurements made by the ID are necessary in reconstructing the higher level *event* details. Combining the measured track with a map of the solenoidal field provides a p_T measurement, and can be traced back to the beam line to determine the interaction vertex from which it originated. This can be used to locate both collisional (*primary*) and decay (*secondary*) vertices, as well as eliminate pileup and cosmic ray tracks. With the ID measurements alone, there is usually enough information to reconstruct the location of the interaction vertices, and the p_T of all electrically charged particles.

The ID has a barrel region that covers $|\eta| < 1$, and two endcaps extending out to $|\eta| = 2.5$. The semiconducting pixel system is the inner most part of ATLAS, and therefore gets the highest density particle flux. These detectors have very high granularity and precision, and are organized into three barrel layers and three endcap disks on each side. The SCT system is just outside the pixel system, and consists of eight layers of semiconducting

strip chambers in both the barrel and endcaps. The TRT system is the outermost part of the ID, and is made out of straw-like drift tubes filled with a xenon gas mixture. The barrel is composed of 73 layers, and has a total of about 50,000 straws oriented parallel to the beam pipe. Each endcap has 18 adjacent wheels with about 320,000 straws oriented radially. Immediately outside of the barrel TRTs is the superconducting solenoidal magnet providing a 2 T field oriented along the z-axis throughout the ID.

4.6 Calorimeters

The *calorimeters*, shown in Fig. 4.8, are located just outside the solenoid, and are designed to fully stop particles produced in collisions. In doing so, showers of secondary particles are produced. These showers can be reconstructed to provide an estimate of the original particle's energy and direction. For charged particles, the shower can be matched to an ID track and the 4-momentum can be reconstructed. Neutral particles don't have corresponding ID tracks, but also don't curve as they pass through the magnetic field. Therefore, the shower direction and energy provide the three components needed to reconstruct their 4-momentum. The only particles that can easily escape the calorimeters are muons and neutrinos. Although particles can make it through small cracks between the barrel and endcaps, and very energetic showers can extend beyond the calorimeters, these are both rare events and in general electrons, photons, and hadrons can all be accurately reconstructed with the ID and calorimeter measurements.

The innermost calorimeter is the *electromagnetic calorimeter* (ECAL), which is optimized for the detection of electrons and photons. It uses layers of lead to produce *electromagnetic showers* from the incident particles, and *liquid argon* (LAr) chambers to sample them. In order to distinguish direct photons and electrons from decays such as $\pi^0 \to \gamma\gamma$, this detector must have very fine angular and energy resolution. The η and ϕ resolutions are about 0.003 and 0.025 rad respectively, and the energy resolution σ_E/E is better than 1% at high energies. With over 26 radiation lengths of material, this detector stops nearly

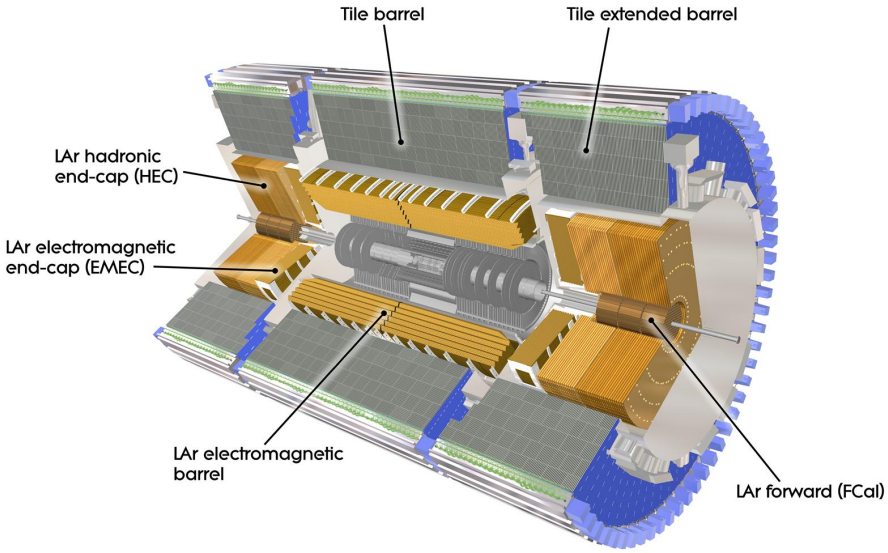

Figure 4.8: A scale cutout of the ATLAS calorimeters, with each of the subdetector systems visible. [82]

all photons and electrons within its eta coverage of $|\eta| < 3.2$. Hadrons *will* produce showers in the ECAL, but won't typically be fully absorbed.

The outer calorimeter is the *hadronic calorimeter* (HCAL), which is designed to stop the hadrons that escape the ECAL. Unlike electromagnetic showers, which are narrow jets of *Bremmsstrahlung* photons and electron pair production, *hadronic showers* can become very complicated. They are typically much larger, and produce a wide array of particles through inelastic collisions. Muons, neutrinos, and neutrons can all be produced in these showers, which then escape ATLAS entirely. These escaping particles can be corrected for, but degrade the resolution of the HCAL. The HCAL makes use of two technologies to create and absorb these showers. In the endcaps from $1.5 < |\eta| < 3.2$, LAr detectors similar to the ECAL are used, but with copper absorbers rather than lead. In the barrel region $|\eta| < 1.7$ is a tile calorimeter that uses steel absorbers and scintillating tiles. The η and ϕ resolutions are about 0.1 and 0.1 rad respectively, and the energy resolution σ_E/E approaches about 3% at high energies.

In addition to the ECAL and HCAL, there is also a forward calorimeter that covers the

region $3.1 < |\eta| < 4.9$. This detector is primarily used for improving $\not{E}_T$ measurements, and uses LAr samplers. The absorbing materials are a combination of copper and tungsten, and the η, ϕ, and E resolutions are similar to that of the HCAL. Overall, the calorimeter system covers the wide range of $|\eta| < 4.9$, with only a small gap around $|\eta| \approx 1.4$ where the barrel and endcaps meet. The calorimeter system is capable of absorbing almost all particles emerging from the collision, with the primary exceptions being muons and neutrinos.

4.7 Muon Spectrometer

Unlike neutrinos, which we can't detect, muons have electric charge and therefore *do* interact electromagnetically with the detector. They're very difficult to stop though, since they're about 200 times more massive than electrons. When a high energy muon collides with an atomic electron, the electron is ionized but due to the mass difference, the muon doesn't lose any significant amount of energy. Unless the muon hits an atomic nucleus, which is very unlikely, it will be able to pass through a large amount of matter before stopping. The *muon spectrometer* (MS), shown in Fig. 4.9, is the largest subdetector in ATLAS, and is a standalone muon detection system designed to measure muon tracks with very high precision. Integrated into the structure is a magnet system that produces a toroidal field around the detector of about 4 T. Because muons can't reliably be stopped, the goal is to make a *second* orthogonal momentum measurement in order to reconstruct the 4-momentum. The η component of the track is measured with very high precision, and the ϕ component, or *second coordinate* is measured a second time with lower precision. The MS track is matched to the ID track in order to reconstruct muons with remarkable accuracy, which is necessary because of the important role muons play in NP searches. In addition to having virtually no *non-muon* background, muons are only produced in the weak decays of massive particles, so their detection is a clear signal that an interesting electroweak interaction occurred. The momentum resolution of the MS is designed to be 3% at 100 GeV and 10% at 1 TeV.

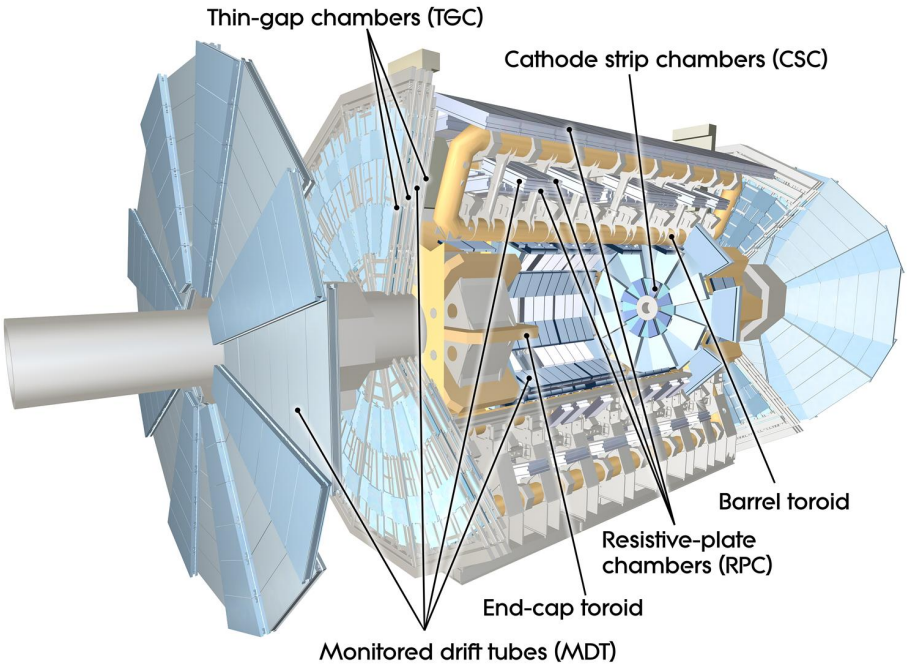

Figure 4.9: A scale cutout of the muon spectrometer in ATLAS, with each of the subdetector systems visible. [82]

The MS is composed of four different detector technologies. For precision tracking, *monitored drift tubes* (MDTs) and *cathode strip chambers* (CSCs) are used. The MDTs are aluminum tubes 30 mm in diameter, with a central wire of 50 µm and filled with a pressurized argon gas mixture. In order to meet the momentum resolution requirements, each MDT must have an average spatial resolution of 80 µm. There are a total of about 372,000 MDTs in the MS, organized into three concentric layers of over 1,000 total chambers. Except for the very forward region of the inner layer, the MS precision tracking system is made entirely of MDTs and covers the region $|\eta| < 2.7$. In the region $2 < |\eta| < 2.7$, on the inner endcap wheel, the particle flux is too high for the MDT technology, so CSCs are used. The CSCs are fast multiwire chambers containing two orthogonal cathode strips that can make both precision and second coordinate measurements of muons comparable to the other technologies.

The remaining two technologies are the *resistive plate chambers* (RPCs) and the *thin*

gap chambers (TGCs). Both serve the same purpose, with the RPCs in the barrel and the TGCs in the endcaps. They use the same basic principles as the CSCs and MDTs, except that precision is traded for speed and each chamber only makes a single measurement of each track. Due to their speed, they're mainly used for triggering on muon events, but also complement the MDT system by providing the second-coordinate of muon tracks.

4.8 Trigger System

With around 10^8 collisions per second, there isn't enough storage space or processing power in the entire world to record every collision at ATLAS. However, the overwhelming majority of collisions are soft QCD interactions that have no relevance to physics searches. They *are* used for calibration and validation of the detector and background studies, but there is no reason to keep anything more than a small sample of these events. A method of quickly filtering out these events with high efficiency is needed to keep the collected data stream at a manageable size. On the other hand, we *never* want to lose events relevant to NP searches. In order to meet both demands, a tiered system of *triggers* is used to incrementally discard events that are not considered interesting. Triggers are organized hierarchically into chains, so that the worst events are discarded more quickly and eat up fewer resources. The *level 1* (L1) triggers have about 2.5 µs to make a decision, and only have access to very limited parts of the detector around a *region of interest* (RoI). The L1 triggers are hardware-based to provide maximal speed, and use rough calculations to veto events that are unambiguously boring. Next are the *level 2* (L2) triggers, which only execute if the parent L1 trigger on their chain has fired. The L2 triggers have access to *all* of the detector data within the RoI, and can make a much more educated decision. The L2 trigger decisions take an average of 50 ms to decide on each event. The final stage of the trigger chain is the *event filter* (EF) trigger. The EF triggers have access to all of the detector data, and can take up to 180 s to make a decision on any given event. After passing the EF stage, an event will be recorded to tape. L2 and EF triggers are part of

the *high-level trigger* (HLT) system, which is implemented in software. For a more detailed explanation of the trigger system, see [83].

Although the above system works well for processes that have a slow enough rate, it fails if we want to analyze processes with very large cross-sections. In order to handle these high rates, while still keeping the trigger rate low, trigger *prescales* are introduced. The prescale of any given trigger is a value representing the reduction in rate that is artificially applied. For example, a trigger with a prescale of 20 will only trigger for 1 in 20 events that would otherwise pass, on average. Prescales at L1 are implemented with a simple deterministic counter that accepts every Nth trigger, while the HLT trigger prescales use a pseudo-random generator that accepts events with a probability of N^{-1}. By using prescales, processes that are generally considered uninteresting, such as those described in Appendix C, can still be studied at a reduced *effective luminosity*.

Chapter 5

The Mono-b Search

The analysis presented here is named the mono-b search, or DM + b [22], and looks for generic NP processes that can be described by the scalar and tensor EFT operators coupling down-type quarks to WIMPs. This is a cut-based analysis, meaning that simple bounds, or cuts, are placed on event variables, trying to reveal the presence NP processes. Each cut is a single value, corresponding to either a maximum or minimum value for an individual variable. Any set of n event variables form an n-dimensional vector space, so that any collection of cuts carves out a volume, or region, within this space. Associated with any given region in this space is the number of observed events, and the number of expected events from each signal and background process. The regions designed to maximize the sensitivity to certain signals over the known backgrounds are called *signal regions* (SRs). In general, an analysis can contain any number of disjoint or overlapping SRs. Each individual region forms a hyper-rectangle in the variable space, so combining multiple regions is a way of building more complex volumes while still relying on simple cuts. *Multivariate analysis* (MVA) algorithms, such as neural networks and boosted decision trees, perform decisions on multiple variables simultaneously, and are able to define much more generic regions in the variable space, which can greatly improve sensitivity. However, the invisible signal of this search leaves very few discriminating variables, limiting the gains of a multivariate approach over a traditional cut-based analysis.

While optimizing the analysis, the observed number of events in any region with high signal sensitivity is always kept hidden. This is known as *blinding*, and is done to avoid potential bias. Looking at the SR while freely modifying analysis details can easily lead to

the creation of an artificial signal or to the obscuring of a real signal. Since no changes should be made after *unblinding*, it's important to define *control regions* (CRs) and *validation regions* (VRs) in order to make sure the background estimates will be accurate in the SR. A CR is a region dominated by a specific background process, whose purpose is to correct any mismodeling in that background. These are typically very removed from the SR to avoid any signal contamination, and have looser cuts to decrease statistical uncertainties. A VR is a region closer to the SR that can be used to evaluate the background estimates without biasing the analysis. Usually the VRs are made as close to the SR as possible, avoiding regions with *significant* signal contamination. In practice these two concepts are often convoluted, where distributions in a CR can be used for validation and a VR can illuminate problems in the background estimate that need to be corrected.

5.1 Data Samples

The most commonly used technique to estimate background and signal processes at ATLAS is known as *Monte Carlo* (MC) simulation. MC samples are generated by estimating the probabilities of the various possible outcomes of a collision, which are then sampled using pseudo-random number generators. Simulations are split up into multiple stages, chaining together the results of different *generators*. Because QCD becomes *non-perturbative* in the low energy limit, *perturbation theory* can only be used to estimate the high energy, or *hard*, processes that take place between the partons within colliding protons. Simulating the low energy effects requires other methods, and the problem is separated into different phenomena that occur in *soft* interactions.

In order to properly simulate the initial state of any collision, the distribution of quarks and gluons within each proton must be modeled. There are a wide number of PDF sets available, but in generating the MC samples used here only the two plotted in Fig. 5.1 were used. CTEQ6L1 is an old PDF set calculated at LO [84], and was used in a large number of the early ATLAS simulations. This set was calculated in 2001, well before the

LHC began running, and has a variety of known issues at higher energies. A more recent PDF, CT10, was produced at NLO and was used for most of the nominal background MC samples in this analysis.

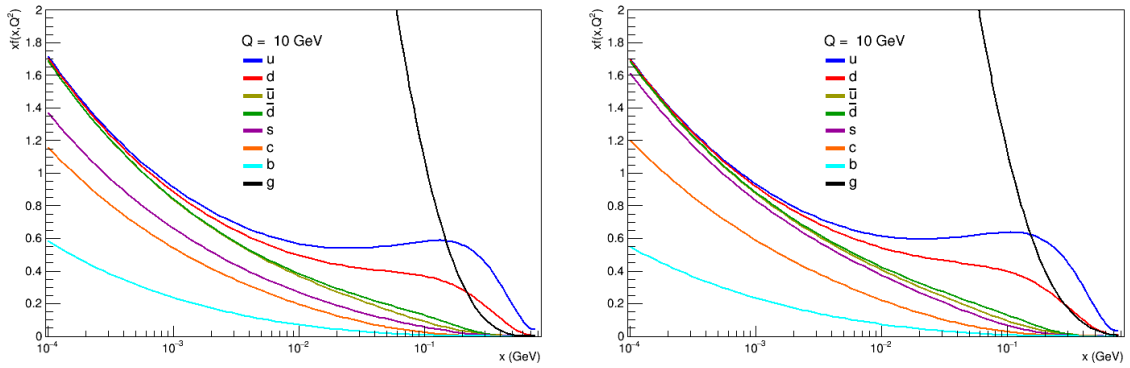

(a) CTEQ6L1 **(b)** CT10

Figure 5.1: The two parton distribution functions used to generate the MC samples, evaluated at a momentum transfer of $Q^2 = 100\,\text{GeV}$. On the left are the PDFs for the CTEQ6L1 set, and CT10 on the right [84]. Both PDF sets are from the same family, but CT10 is more recent and performs much better at higher energies.

The first stage of any simulation involves calculating the *matrix element* (ME) of the primary hard interaction up to LO or NLO. For a given final state, ME generators use the proton PDFs to select two initial-state partons, and calculate the corresponding Feynman diagrams. This is a pretty straightforward procedure, but it only works at high energies where QCD becomes perturbative. For any given hard interaction, there are an infinite number of nearly identical interactions containing very low energy particles radiated from either the initial partons, ISR, or the final-state particles, *final state radiation* (FSR). Quarks are most likely to radiate gluons which, because they're also charged, can continue to radiate. These soft higher order contributions to the hard process is known as *parton showering* (PS), and isn't considered in the initial ME calculation. The ME results are passed to a PS generator, which can approximate the calculation of PS contributions. Because soft and hard are poorly defined terms, there is a region in between that is covered by both the ME and PS generators. Matching algorithms, such as MLM [85] and CKKW [86], are needed to remove this overlap and combine the results without double counting.

After the parton-level processes are completed, the simulation still contains bare color charges, which should hadronize into color-neutral states. After the hadronization step, there can be any number of heavy hadrons produced, which will decay before reaching the detector. Typical PS generators handle both of these steps, and output a collection of final-state particles that are stable enough to reach the detector.

In addition to the primary collision of interest, many other low energy interactions typically occur in each event. These are known as the *underlying event* (UE) and include sources such as the primary protons' remnants, secondary parton interactions, and pileup collisions. Although *data-driven methods* exist, the UE is usually simulated by PS generators, whose results are inserted into every event to account for the soft backgrounds of the typical hadron collision.

The final step in MC generation is to simulate the detector itself. The ME, PS, and UE results are put through an ATLAS simulator, which takes into account the efficiencies and noise rates of each subdetector, the magnetic fields, and interactions with the materials composing the detector. The results of the detector simulation are then passed through the exact same digitization and reconstruction algorithms as real data, producing the final MC sample to be used for physics analyses.

Pythia [87] is the most frequently used PS generator in ATLAS MC production, and is also used for the UE simulation. For this analysis, at LO, the MadGraph [88], AcerMC [89], and Alpgen [90] generators are used for the ME, while at NLO, Powheg [91] and MC@NLO [92] are used. Sherpa [93–96] and Herwig [97] are used as standalone generators, capable of both ME and PS calculations and ME/PS matching at LO. All of the samples produced with Sherpa made use of the NLO PDF set CT10, while the other generators used CTEQ6L1.

Detector simulations are done primarily with the Geant4 [98] program. FullSim, the slowest but most accurate method, uses a detailed, complete detector model in Geant4. The problem with FullSim is that it's slow, and over 90% of the CPU power is spent simulating calorimeter showers. The *frozen shower* (FS) method [99] is the easiest way to increase

speed, and uses a library of pre-made showers rather than simulating new ones for each event. An even faster option is *Atlfast-II* (AF2), which uses the FastCaloSim program with a simplified geometry model to simulate the calorimeters and Geant4 only for the ID and MS. As always, increased speed is coupled with decreased accuracy. Depending on the needs of the analysis, FS is usually used for high accuracy and AF2 for high statistic MC samples.

One issue that needs to be addressed is that these simulations aren't exact. They were all produced well in advance of the 2012 run, using versions of various programs that are now outdated. Even at the time of production, approximations were used at every stage in order to decrease computational demands. The final samples can have any number of issues that need to be corrected before the analysis is performed. There are three main types of corrections that we will use: *calibration*, *smearing*, and *scale factors* (SFs). Calibrations are the adjustment of event variables to better reflect reality. For example, the jet measurements written to disk are really just a first pass. Pileup, *jet energy scale* (JES), and origin corrections can be applied to both MC and data to increase the accuracy and resolution of the energy measurement. Smearing, on the other hand, is an artificial modification of event variables in order to *worsen* resolution. This is usually applied to MC simulations that overestimate the resolution of a detector, to achieve better agreement with data. The third type of correction is SFs, where the weight of each event is modified. Every MC event has a number of weights associated with it, which are multiplied to obtain its total weight. These weights can be positive or negative, and are summed for the calculation of the yield estimate in any region. The following SFs are applied to every MC event,

Generator Certain ME generators assign weights to each event, in order to sample the phase space in a more flexible way. Rare regions can be reached by assigning a small weight, and other regions can be subtracted by assigning a negative weight.

Sample Size Every sample is normalized to the sum of generator weights, to make its predictions independent of the number of events generated.

Cross-section Once the sample has been normalized to unity, it's multiplied by the total calculated cross-section of the simulated process.

k-factor If NLO or NNLO calculations for the total cross-section are available, the ratio to the LO result can be used to provide increased accuracy. These factors are applied uniformly to an entire sample, and don't correct for higher order kinematic deviations.

Filter efficiency Various filters can be applied at the generator-level, to single out more important regions of the phase space. This shrinks the effective cross-section, and the sample weight needs to be corrected accordingly.

Luminosity The final theoretical SF is the total integrated luminosity of the data sample. After scaling each sample to its effective cross-section using the above weights, the luminosity determines the expected number of events.

Trigger Different effects in the detector simulation can result in mismodeling of the trigger efficiencies. Weighting each event by the ratio in trigger efficiency between data and MC is an effective method of correcting this. This is discussed further in Section 5.2, and is applied whenever a trigger requirement is placed.

Object Mismodeling in the identification and reconstruction of the simulated physics objects are also corrected using event weights. Each reconstructed object in an event contributes a weight, and is discussed further in Section 5.3.

Pileup The pileup environment was poorly modeled in the MC simulations produced before data collection, and event weights are used to correct this. This is discussed further in Section 5.4

5.1.1 Collision Data

The three data streams from the 2012 collisions used in this analysis are named Egamma, Muon, and JetTauEtMiss. Every trigger chain in ATLAS has associated streams that it

will write an event to, when fired. If any of the electron or photon triggers fired, the event will be written to the Egamma stream. Likewise, events in which muon triggers fired will be written to the Muon stream, and events with any jets, *tau leptons*, or $\not{E}_T$ triggers will be written to the JetTauEtmiss stream. There are a large number of "coincidental" overlaps between these streams, as well as multi-object triggers that will *always* place an event in multiple streams when fired. This is frequently taken advantage of, as in Appendix D, to measure trigger efficiencies. The total integrated luminosity recorded during the 2012 run, shown in Fig. 5.2, was estimated to be $(20.3 \pm 0.6)\,\mathrm{fb}^{-1}$ [100, 101].

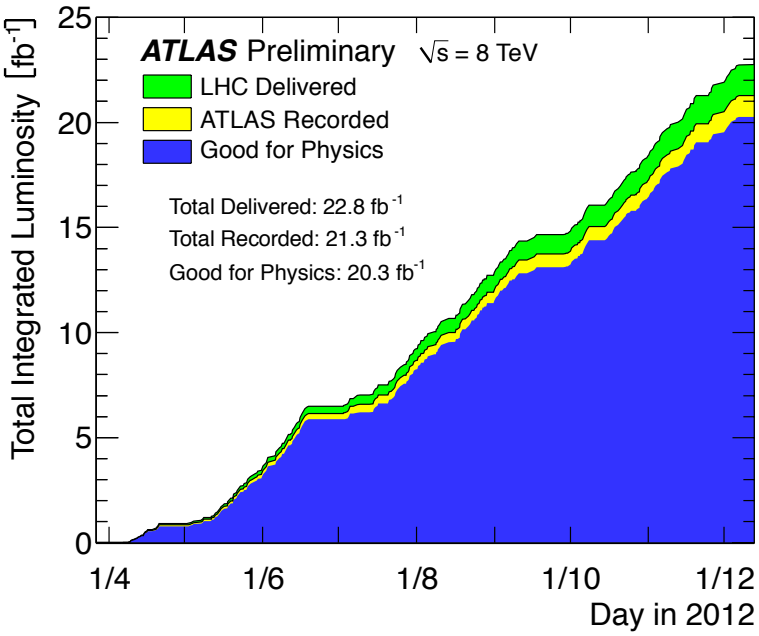

Figure 5.2: The total integrated luminosity over the entire 2012 run period. The LHC provided a total of $22.8\,\mathrm{fb}^{-1}$, of which about 89% was usable for analyses. The missing 11% is due to brief periods where ATLAS was not recording data or there was some known problem in the detector.

5.1.2 Signals

The focus of this search is the EFT interactions coupling down-type quarks to WIMPs that exhibit a quark mass dependency under MFV. In principle, this would require a search for 10 distinct signals out of the 24 interactions listed in Table 2.1. However, ignoring the different placements of the γ^5 matrix, it's important to note that as far as this analysis is

concerned, there are only 3 unique operators. With direct or indirect detection experiments, the γ^5 placement in these 10 operators can produce vastly different results. In collider searches though, many of the operators are nearly identical up to constant factors in the cross-section. The three groupings of (D1, D2, D3, D4), (D9,D10), and (C1,C2) are all expected to behave nearly identically, demonstrated in Appendix F, allowing the analysis to be reduced to a search on just three operators: D1, D9, and C1.

24 signal samples, summarized in Table 5.1, were produced for this analysis, broken up into 8 mass points for each of the three operators. WIMP masses of 10, 50, 100, 200, 400, 700, 1000, and 1,300 GeV were simulated, with the EFT coupling constant M_* set to 1 TeV. For 8 TeV collisions, there is no sensitivity to WIMP masses above 1,300 GeV, and below 10 GeV the kinematics become independent of mass, so the results can be easily extrapolated by comparing cross-sections. All of the signal samples were generated using MadGraph for the ME calculations and Pythia6 for the PS. There was a small complication with the D9 samples, due to the fact that they were not produced following the strict MFV guidelines. The quark mass dependence was not included in these samples, the motivation for which was never made clear. However, a filter on bottom quarks *was* used in the simulation so the *kinematics* are expected to be very similar to the MFV variant. This deviation can mostly be corrected for by simply scaling each D9 sample by m_b^2/M_*^2 to account for the change in cross-section.

5.1.3 Backgrounds

There are only a handful of different types of SM background processes that can mimic NP. These primarily consist of processes that produce some number of electroweak bosons, top quarks, or high p_T jets. The differences from search to search are in the decays of the heavier particles, which can result in a large variety of final-state signatures. Usually only final states with high energy photons, leptons or $\not{E}_T$ are considered, since purely hadronic events are overwhelmed by the soft QCD background. For this analysis, the final state signature of the signal is one or two bottom quarks recoiling off $\not{E}_T$. The largest

Operator	M_{DM} [GeV]	Generator	σ [pb]	N_{sim}
D1	10	Madgraph+Pythia6 AF2	$3.8 \cdot 10^{-8}$	40,000
D1	50	Madgraph+Pythia6 AF2	$3.1 \cdot 10^{-8}$	40,000
D1	100	Madgraph+Pythia6 AF2	$2.2 \cdot 10^{-8}$	40,000
D1	200	Madgraph+Pythia6 AF2	$9.2 \cdot 10^{-9}$	40,000
D1	400	Madgraph+Pythia6 AF2	$1.6 \cdot 10^{-9}$	40,000
D1	700	Madgraph+Pythia6 AF2	$1.4 \cdot 10^{-10}$	40,000
D1	1,000	Madgraph+Pythia6 AF2	$1.3 \cdot 10^{-11}$	40,000
D1	1,300	Madgraph+Pythia6 AF2	$1.2 \cdot 10^{-12}$	40,000
D9	10	Madgraph+Pythia6 AF2	$5.6 \cdot 10^{-8}$	20,000
D9	50	Madgraph+Pythia6 AF2	$4.5 \cdot 10^{-8}$	20,000
D9	100	Madgraph+Pythia6 AF2	$3.1 \cdot 10^{-8}$	20,000
D9	200	Madgraph+Pythia6 AF2	$1.4 \cdot 10^{-8}$	20,000
D9	400	Madgraph+Pythia6 AF2	$3.5 \cdot 10^{-8}$	20,000
D9	700	Madgraph+Pythia6 AF2	$6.5 \cdot 10^{-8}$	20,000
D9	1,000	Madgraph+Pythia6 AF2	$1.4 \cdot 10^{-8}$	20,000
D9	1,300	Madgraph+Pythia6 AF2	$2.6 \cdot 10^{-10}$	20,000
C1	10	Madgraph+Pythia6 AF2	$3.4 \cdot 10^{-7}$	40,000
C1	50	Madgraph+Pythia6 AF2	$2.8 \cdot 10^{-7}$	40,000
C1	100	Madgraph+Pythia6 AF2	$7.6 \cdot 10^{-8}$	40,000
C1	200	Madgraph+Pythia6 AF2	$1.7 \cdot 10^{-8}$	40,000
C1	400	Madgraph+Pythia6 AF2	$1.4 \cdot 10^{-9}$	40,000
C1	700	Madgraph+Pythia6 AF2	$6.7 \cdot 10^{-11}$	40,000
C1	1,000	Madgraph+Pythia6 AF2	$4.2 \cdot 10^{-12}$	40,000
C1	1,300	Madgraph+Pythia6 AF2	$3.0 \cdot 10^{-13}$	40,000

Table 5.1: Signal samples for D1, D9, and C1 operators. k-factors and filter efficiencies were all 1, and an M_* value of 1 TeV was used.

background to this is the irreducible process $Z \to \nu\nu$ produced in association with *b-jets*. The second largest is $W \to \ell\nu$, where the charged lepton either escapes detection or is a hadronically decaying tau lepton. Top production is the smallest of the three major backgrounds, consisting mostly of events with semi-leptonically decaying top pairs. Since the top decays to a $W^{\pm}$ boson and a bottom with a branching ratio of nearly 100% [34], this background is similar to $W \to \ell\nu$, except with more jets, and two guaranteed b-jets. The remaining four backgrounds considered are $Z \to \ell\ell$, γ + jets, diboson and multijet production, none of which contribute significantly in the SR.

V + jets The V + jets processes represent the production of single γ-ray, or a leptonically decaying weak boson in association with any numbers of jets. For all of these, the Sherpa generator was used to produce samples, taking the masses of both *charm* and bottom quarks into account. For each boson, the simulations were split up into 3 flavor and 5 boson p_T slices, in order to get adequate statistics in the rarer parts of the parameter space. The $W^{\pm}$

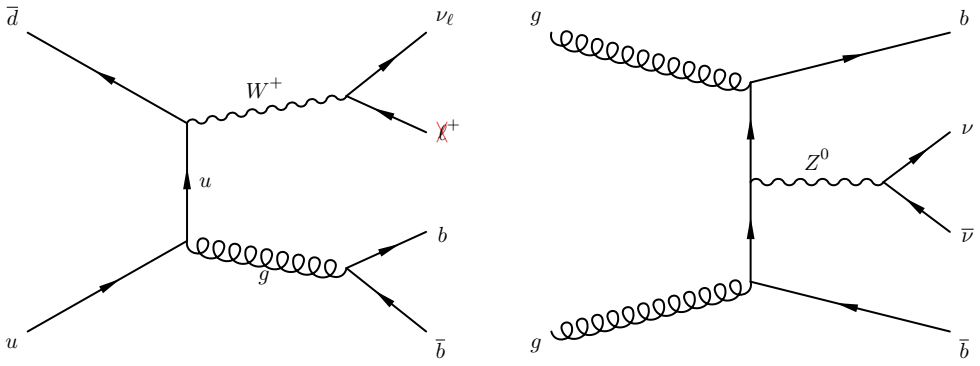

Figure 5.3: Example $W^{\pm} \to \ell^{\pm}\nu$ and $Z^0 \to \nu\bar{\nu}$ background processes.

and Z^0 simulations were additionally split by their leptonic decay modes. This resulted in the 15 γ, 45 $W \to \ell\nu$, 45 $Z \to \ell\ell$, and 15 $Z \to \nu\nu$ samples detailed in Tables 5.2, 5.3, 5.4, and 5.5 respectively. The flavor bins of each process were defined by filtering on or vetoing D and B hadrons in the final state to get orthogonal *light jet*, *c-jet*, and *b*-jet samples. The p_T slices were defined by using the truth value of the boson's p_T, defined for the ranges 0 GeV–70 GeV, 70 GeV–140 GeV, 140 GeV–280 GeV, 280 GeV–500 GeV, and an inclusive bin for $p_T > 500$ GeV.

Although the Sherpa samples were used in the final analysis, simulations produced with Alpgen and Pythia were also used as cross-checks. As with the Sherpa samples, samples were produced for each of the leptonic decay modes for each boson. However, rather than organizing them by the boson p_T, they were split according to the additional final state partons produced in the hard interaction. Samples were produced in 5 exclusive bins corresponding to 0–4 additional partons, and an inclusive bin containing events with five or more. Special samples were also generated for *heavy flavor* jets produced in association with gauge bosons. The Zbb, Zcc, Wbb, and Wcc processes were produced with three exclusive bins, and one inclusive bin for events with three or more additional partons, while the Wc process was produced with four exclusive bins and one inclusive. The Alpgen samples for γ + jets and $Z \to \nu\nu$ were treated slightly differently. For the Alpgen γ + jets samples, additional slicing was performed on the truth p_T for 35 GeV–70 GeV, 70 GeV–140 GeV, 140 GeV–280 GeV, 280 GeV–500 GeV, and an inclusive bin for $p_T > 500$ GeV.

For the flavor-blind $Z \to \nu\nu$ samples, a filter of at least 1 jet and a binning of 0 GeV–70 GeV, 70 GeV–140 GeV, 140 GeV–280 GeV, 280 GeV–500 GeV, and $p_T > 500$ GeV was used. Overlap between the flavor-blind samples and the heavy flavor samples occurs in certain parts of the phase space, and was appropriately removed. In total there were 25 γ, 32 $W \to \ell\nu$, 42 $Z \to \ell\ell$, and 34 $Z \to \nu\nu$ Alpgen samples, which are detailed in Tables 5.6, 5.7, 5.8, and 5.9 respectively.

These three electroweak interactions are by far the most important to this analysis. The $Z \to \nu\nu$ and $W \to \ell\nu$ processes, examples of which are shown in Fig. 5.3, are the two largest backgrounds to the signal, contributing over 75% of the expected SM background. The $Z \to \mu\mu$ and γ processes, although minor backgrounds to the signal, are crucial in the data-driven estimate of $Z \to \nu\nu$, detailed in Appendix B. Because of their importance, a lot of effort was put into making sure these backgrounds were well estimated, as will be shown later.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
SinglePhoton MassiveCB Pt100-140 CVetoBVeto	Sherpa AF2	430.8	1.00	0.50	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt100-140 CFilterBVeto	Sherpa AF2	430.4	1.00	0.42	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt100-140 BFilter	Sherpa AF2	428.8	1.00	0.08	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt140-280 CVetoBVeto	Sherpa AF2	138.0	1.00	0.49	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt140-280 CFilterBVeto	Sherpa AF2	137.8	1.00	0.42	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt140-280 BFilter	Sherpa AF2	137.3	1.00	0.08	$2.0 \cdot 10^6$
SinglePhoton MassiveCB Pt280-500 CVetoBVeto	Sherpa FS	5.9	1.00	0.48	$7.5 \cdot 10^5$
SinglePhoton MassiveCB Pt280-500 BFilter	Sherpa FS	5.9	1.00	0.10	$2.5 \cdot 10^5$
SinglePhoton MassiveCB Pt280-500 CFilterBVeto	Sherpa FS	5.9	1.00	0.42	$7.5 \cdot 10^5$
SinglePhoton MassiveCB Pt500 CVetoBVeto	Sherpa FS	0.3	1.00	0.47	$2.0 \cdot 10^5$
SinglePhoton MassiveCB Pt500 CFilterBVeto	Sherpa FS	0.3	1.00	0.42	$2.0 \cdot 10^5$
SinglePhoton MassiveCB Pt500 BFilter	Sherpa FS	0.3	1.00	0.11	$1.0 \cdot 10^5$

Table 5.2: γ + jets Sherpa samples used in the analysis.

Diboson The diboson background represents the simultaneous production of any two electroweak bosons. For WW, WZ, and ZZ production, inclusive samples were produced using the Herwig generator, while Sherpa was used for $W\gamma$ and $Z\gamma$. Details of these can be found in Table 5.10. The diboson background processes can easily mimic the signal, an example of which is shown in Figure 5.4. However, their cross-sections are over 137 times smaller than single boson production, because of the additional electroweak vertex,

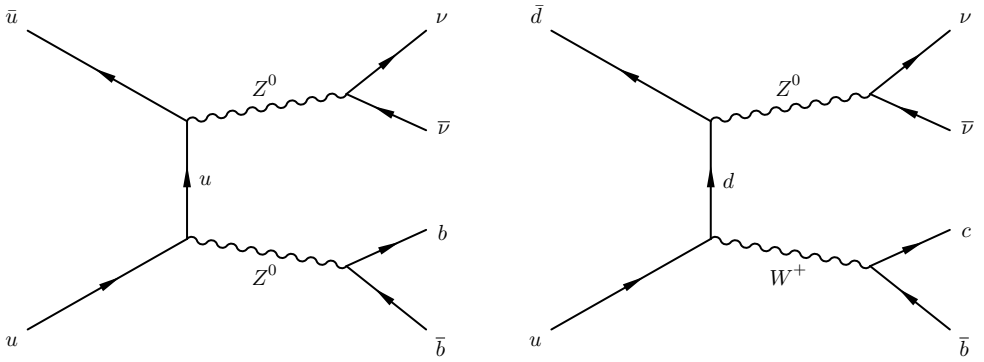

Figure 5.4: Example diboson background processes.

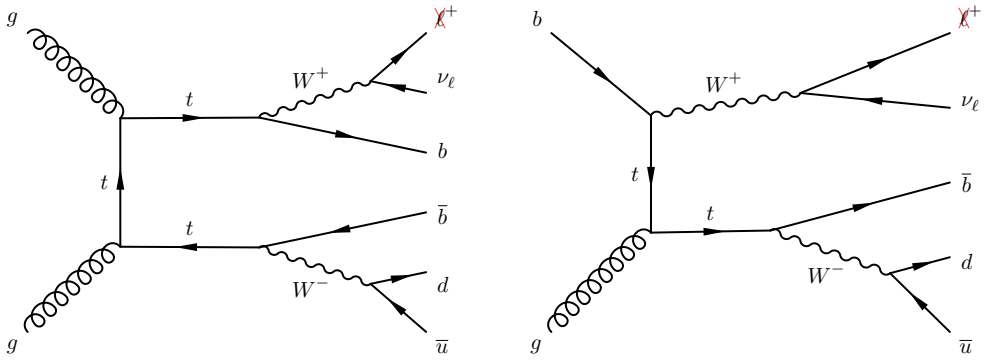

Figure 5.5: Example $t\bar{t}$ and single top background processes where a lepton is not identified.

and therefore contribute very little to the final background estimate. Because of this, and because mismodeling will be partially compensated for by corrections to the V + jets estimates, these estimates were not directly validated. Sanity checks were performed though, to make sure there was no major mismodeling, and large systematic uncertainties were assigned to reflect this.

Top Production Top quarks can either be produced alone, or more commonly as a quark/anti-quark pair. The $t\bar{t}$ process is dominated by QCD interactions, while single top quarks are produced via $W^{\pm}$ boson exchange. To generate single top samples, AcerMC with Pythia6 was used for t-channel production, which corresponds to a flavor changing $W^{\pm}$ exchange between two light quarks that turns one of them into a top. MC@NLO with Herwig++ was used for s-channel production, in which a virtual $W^{\pm}$ boson produces a top and down-type quark. MC@NLO and Herwig++ were also used to simulate Wt-channel

production, in which both a $W^{\pm}$ boson and a top are produced. For top pair production, a single inclusive sample was simulated using Powheg and Pythia6. Details of these samples can be found in Table 5.11. Note that simulations of purely hadronically decaying top quarks were not used, and are considered to be negligible contributions to the $W \to \ell\nu$ and multijet backgrounds.

Top quarks decay to a $W^{\pm}$ and a bottom nearly 100% of the time, which makes them a significant background to the signal. Leptonic top decays with a misidentified lepton or τ-*jet* can very easily mimic WIMP production, since there will always be a real neutrino. Examples of the dominant single top and $t\bar{t}$ processes that contribute to the expected SM background are shown in Fig. 5.5.

Multijets The multijet background is one of the most difficult SM backgrounds to deal with for two reasons. The first is that in hadron collisions, the cross-section for multijet events are many orders of magnitude higher than any other SM background ($\mathcal{O}(10^6)$ higher than $W \to \ell\nu$). On top of that, the generation of high $\not{E}_T$ in multijet events is incredibly rare. In our SR this translates to a tiny fraction of a gigantic number, which is problematic for MC simulation which don't produce enough events in the SR to be reliable. The data-driven estimate of this background is described in Appendix C, which is used as the multijet estimate throughout the analysis. In order to validate this estimate in higher statistic regions, Pythia8 simulations of the multijet processes were used for 8 different jet p_T slices, as listed in Table 5.12.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$W \to e\nu$ MassiveCB Pt0 BFilter	Sherpa AF2	10973.0	1.11	0.01	$4.4 \cdot 10^6$
$W \to e\nu$ MassiveCB Pt0 CJetFilterBVeto	Sherpa AF2	10971.0	1.11	0.05	$3.0 \cdot 10^6$
$W \to e\nu$ MassiveCB Pt0 CJetVetoBVeto	Sherpa AF2	10987.0	1.11	0.94	$2.8 \cdot 10^7$
$W \to \mu\nu$ MassiveCB Pt0 BFilter	Sherpa AF2	10973.0	1.11	0.01	$4.4 \cdot 10^6$
$W \to \mu\nu$ MassiveCB Pt0 CJetFilterBVeto	Sherpa AF2	10970.0	1.11	0.04	$2.9 \cdot 10^6$
$W \to \mu\nu$ MassiveCB Pt0 CJetVetoBVeto	Sherpa AF2	10981.0	1.11	0.94	$2.8 \cdot 10^7$
$W \to \tau\nu$ MassiveCB Pt0 BFilter	Sherpa AF2	10974.0	1.11	0.01	$4.3 \cdot 10^6$
$W \to \tau\nu$ MassiveCB Pt0 CJetFilterBVeto	Sherpa AF2	10971.0	1.11	0.05	$3.0 \cdot 10^6$
$W \to \tau\nu$ MassiveCB Pt0 CJetVetoBVeto	Sherpa AF2	10969.0	1.11	0.94	$2.8 \cdot 10^7$
$W \to e\nu$ MassiveCB Pt70-140 BFilter	Sherpa AF2	250.6	1.11	0.05	$4.3 \cdot 10^5$
$W \to e\nu$ MassiveCB Pt70-140 CJetFilterBVeto	Sherpa AF2	250.7	1.11	0.20	$6.9 \cdot 10^5$
$W \to e\nu$ MassiveCB Pt70-140 CJetVetoBVeto	Sherpa AF2	250.4	1.11	0.75	$3.8 \cdot 10^6$
$W \to \mu\nu$ MassiveCB Pt70-140 BFilter	Sherpa AF2	250.6	1.11	0.05	$4.3 \cdot 10^5$
$W \to \mu\nu$ MassiveCB Pt70-140 CJetFilterBVeto	Sherpa AF2	250.6	1.11	0.20	$6.9 \cdot 10^5$
$W \to \mu\nu$ MassiveCB Pt70-140 CJetVetoBVeto	Sherpa AF2	250.8	1.11	0.76	$3.8 \cdot 10^6$
$W \to e\nu$ MassiveCB Pt140-280 BFilter	Sherpa AF2	31.2	1.11	0.06	$1.0 \cdot 10^6$
$W \to e\nu$ MassiveCB Pt140-280 CJetFilterBVeto	Sherpa AF2	31.2	1.11	0.22	$4.3 \cdot 10^5$
$W \to e\nu$ MassiveCB Pt140-280 CJetVetoBVeto	Sherpa AF2	31.1	1.11	0.71	$4.7 \cdot 10^5$
$W \to \mu\nu$ MassiveCB Pt140-280 BFilter	Sherpa AF2	31.2	1.11	0.06	$1.0 \cdot 10^6$
$W \to \mu\nu$ MassiveCB Pt140-280 CJetFilterBVeto	Sherpa AF2	31.2	1.11	0.22	$4.3 \cdot 10^5$
$W \to \mu\nu$ MassiveCB Pt140-280 CJetVetoBVeto	Sherpa AF2	31.2	1.11	0.72	$4.6 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt70-140 BFilter	Sherpa FS	250.6	1.11	0.05	$4.3 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt70-140 CJetFilterBVeto	Sherpa FS	250.6	1.11	0.20	$6.9 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt70-140 CJetVetoBVeto	Sherpa FS	250.6	1.11	0.75	$3.8 \cdot 10^6$
$W \to \tau\nu$ MassiveCB Pt140-280 BFilter	Sherpa FS	31.2	1.11	0.06	$2.1 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt140-280 CJetFilterBVeto	Sherpa FS	31.2	1.11	0.22	$4.3 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt140-280 CJetVetoBVeto	Sherpa FS	31.2	1.11	0.72	$4.7 \cdot 10^5$
$W \to e\nu$ MassiveCB Pt280-500 BFilter	Sherpa FS	1.8	1.11	0.08	$2.0 \cdot 10^4$
$W \to e\nu$ MassiveCB Pt280-500 CJetFilterBVeto	Sherpa FS	1.8	1.11	0.23	$4.1 \cdot 10^4$
$W \to e\nu$ MassiveCB Pt280-500 CJetVetoBVeto	Sherpa FS	1.8	1.11	0.68	$1.1 \cdot 10^5$
$W \to \mu\nu$ MassiveCB Pt280-500 BFilter	Sherpa FS	1.8	1.11	0.08	$2.0 \cdot 10^4$
$W \to \mu\nu$ MassiveCB Pt280-500 CJetFilterBVeto	Sherpa FS	1.8	1.11	0.23	$4.1 \cdot 10^4$
$W \to \mu\nu$ MassiveCB Pt280-500 CJetVetoBVeto	Sherpa FS	1.8	1.11	0.69	$1.1 \cdot 10^5$
$W \to \tau\nu$ MassiveCB Pt280-500 BFilter	Sherpa FS	1.8	1.11	0.08	$2.0 \cdot 10^4$
$W \to \tau\nu$ MassiveCB Pt280-500 CJetFilterBVeto	Sherpa FS	1.8	1.11	0.23	$4.1 \cdot 10^4$
$W \to \tau\nu$ MassiveCB Pt280-500 CJetVetoBVeto	Sherpa FS	1.8	1.11	0.68	$1.1 \cdot 10^5$
$W \to e\nu$ MassiveCB Pt500 BFilter	Sherpa FS	0.1	1.11	0.10	$2.0 \cdot 10^3$
$W \to e\nu$ MassiveCB Pt500 CJetFilterBVeto	Sherpa FS	0.1	1.11	0.24	$2.0 \cdot 10^3$
$W \to e\nu$ MassiveCB Pt500 CJetVetoBVeto	Sherpa FS	0.1	1.11	0.66	$2.1 \cdot 10^3$
$W \to \mu\nu$ MassiveCB Pt500 BFilter	Sherpa FS	0.1	1.11	0.10	$2.0 \cdot 10^3$
$W \to \mu\nu$ MassiveCB Pt500 CJetFilterBVeto	Sherpa FS	0.1	1.11	0.24	$2.0 \cdot 10^3$
$W \to \mu\nu$ MassiveCB Pt500 CJetVetoBVeto	Sherpa FS	0.1	1.11	0.66	$1.1 \cdot 10^4$
$W \to \tau\nu$ MassiveCB Pt500 BFilter	Sherpa FS	0.1	1.11	0.10	$2.0 \cdot 10^3$
$W \to \tau\nu$ MassiveCB Pt500 CJetFilterBVeto	Sherpa FS	0.1	1.11	0.24	$2.0 \cdot 10^3$
$W \to \tau\nu$ MassiveCB Pt500 CJetVetoBVeto	Sherpa FS	0.1	1.11	0.66	$1.1 \cdot 10^4$

Table 5.3: $W \to \ell\nu$ Sherpa samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$Z \to ee$ MassiveCB Pt0 BFilter	Sherpa AF2	1110.7	1.12	0.03	$1.1 \cdot 10^6$
$Z \to ee$ MassiveCB Pt0 CFilterBVeto	Sherpa AF2	1109.6	1.12	0.28	$1.4 \cdot 10^6$
$Z \to ee$ MassiveCB Pt0 CVetoBVeto	Sherpa AF2	1107.1	1.12	0.69	$2.8 \cdot 10^6$
$Z \to \mu\mu$ MassiveCB Pt0 BFilter	Sherpa AF2	1109.8	1.12	0.03	$1.1 \cdot 10^6$
$Z \to \mu\mu$ MassiveCB Pt0 CFilterBVeto	Sherpa AF2	1112.0	1.12	0.28	$1.4 \cdot 10^6$
$Z \to \mu\mu$ MassiveCB Pt0 CVetoBVeto	Sherpa AF2	1108.7	1.12	0.69	$2.8 \cdot 10^6$
$Z \to \tau\tau$ MassiveCB Pt0 BFilter	Sherpa AF2	1109.1	1.12	0.03	$1.1 \cdot 10^6$
$Z \to \tau\tau$ MassiveCB Pt0 CFilterBVeto	Sherpa AF2	1110.2	1.12	0.28	$1.4 \cdot 10^6$
$Z \to \tau\tau$ MassiveCB Pt0 CVetoBVeto	Sherpa AF2	1112.1	1.12	0.69	$2.8 \cdot 10^6$
$Z \to ee$ MassiveCB Pt70-140 BFilter	Sherpa AF2	29.5	1.12	0.08	$3.0 \cdot 10^5$
$Z \to ee$ MassiveCB Pt70-140 CFilterBVeto	Sherpa AF2	29.5	1.12	0.35	$2.4 \cdot 10^5$
$Z \to ee$ MassiveCB Pt70-140 CVetoBVeto	Sherpa AF2	29.5	1.12	0.56	$5.1 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt70-140 BFilter	Sherpa AF2	29.5	1.12	0.08	$3.0 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt70-140 CFilterBVeto	Sherpa AF2	29.4	1.12	0.35	$2.4 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt70-140 CVetoBVeto	Sherpa AF2	29.5	1.12	0.56	$5.1 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt70-140 BFilter	Sherpa FS	29.5	1.12	0.08	$3.0 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt70-140 CFilterBVeto	Sherpa FS	29.5	1.12	0.36	$2.4 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt70-140 CVetoBVeto	Sherpa FS	29.5	1.12	0.56	$5.1 \cdot 10^5$
$Z \to ee$ MassiveCB Pt140-280 BFilter	Sherpa AF2	4.0	1.12	0.10	$2.1 \cdot 10^5$
$Z \to ee$ MassiveCB Pt140-280 CFilterBVeto	Sherpa AF2	4.0	1.12	0.37	$8.9 \cdot 10^4$
$Z \to ee$ MassiveCB Pt140-280 CVetoBVeto	Sherpa AF2	4.0	1.12	0.53	$1.4 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt140-280 BFilter	Sherpa AF2	4.0	1.12	0.10	$2.1 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt140-280 CFilterBVeto	Sherpa AF2	4.0	1.12	0.37	$8.9 \cdot 10^4$
$Z \to \mu\mu$ MassiveCB Pt140-280 CVetoBVeto	Sherpa AF2	4.0	1.12	0.53	$1.4 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt140-280 BFilter	Sherpa FS	4.0	1.12	0.10	$4.1 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt140-280 CFilterBVeto	Sherpa FS	4.0	1.12	0.37	$8.9 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt140-280 CVetoBVeto	Sherpa FS	4.0	1.12	0.53	$1.4 \cdot 10^5$
$Z \to ee$ MassiveCB Pt280-500 BFilter	Sherpa FS	0.2	1.12	0.11	$5.7 \cdot 10^4$
$Z \to ee$ MassiveCB Pt280-500 CFilterBVeto	Sherpa FS	0.2	1.12	0.39	$1.4 \cdot 10^5$
$Z \to ee$ MassiveCB Pt280-500 CVetoBVeto	Sherpa FS	0.2	1.12	0.51	$1.5 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt280-500 BFilter	Sherpa FS	0.2	1.12	0.11	$5.7 \cdot 10^4$
$Z \to \mu\mu$ MassiveCB Pt280-500 CFilterBVeto	Sherpa FS	0.2	1.12	0.39	$1.4 \cdot 10^5$
$Z \to \mu\mu$ MassiveCB Pt280-500 CVetoBVeto	Sherpa FS	0.2	1.12	0.51	$1.5 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt280-500 BFilter	Sherpa FS	0.2	1.12	0.11	$4.5 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt280-500 CFilterBVeto	Sherpa FS	0.2	1.12	0.38	$1.2 \cdot 10^5$
$Z \to \tau\tau$ MassiveCB Pt280-500 CVetoBVeto	Sherpa FS	0.2	1.12	0.51	$1.2 \cdot 10^5$
$Z \to ee$ MassiveCB Pt500 BFilter	Sherpa FS	0.0	1.12	0.12	$1.8 \cdot 10^4$
$Z \to ee$ MassiveCB Pt500 CFilterBVeto	Sherpa FS	0.0	1.12	0.40	$1.9 \cdot 10^4$
$Z \to ee$ MassiveCB Pt500 CVetoBVeto	Sherpa FS	0.0	1.12	0.48	$6.2 \cdot 10^4$
$Z \to \mu\mu$ MassiveCB Pt500 BFilter	Sherpa FS	0.0	1.12	0.11	$2.0 \cdot 10^3$
$Z \to \mu\mu$ MassiveCB Pt500 CFilterBVeto	Sherpa FS	0.0	1.12	0.40	$2.1 \cdot 10^3$
$Z \to \mu\mu$ MassiveCB Pt500 CVetoBVeto	Sherpa FS	0.0	1.12	0.49	$5.9 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt500 BFilter	Sherpa FS	0.0	1.12	0.12	$1.2 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt500 CFilterBVeto	Sherpa FS	0.0	1.12	0.39	$1.2 \cdot 10^4$
$Z \to \tau\tau$ MassiveCB Pt500 CVetoBVeto	Sherpa FS	0.0	1.12	0.49	$3.2 \cdot 10^4$

Table 5.4: $Z \to \ell\ell$ Sherpa samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$Z \to \nu\nu$ MassiveCB Pt0 BFilter	Sherpa AF2	5990.8	1.12	0.03	$6.8 \cdot 10^6$
$Z \to \nu\nu$ MassiveCB Pt0 CFilterBVeto	Sherpa AF2	5988.3	1.12	0.28	$9.5 \cdot 10^6$
$Z \to \nu\nu$ MassiveCB Pt0 CVetoBVeto	Sherpa AF2	5987.5	1.12	0.69	$1.4 \cdot 10^7$
$Z \to \nu\nu$ MassiveCB Pt70-140 BFilter	Sherpa AF2	166.6	1.12	0.08	$1.3 \cdot 10^6$
$Z \to \nu\nu$ MassiveCB Pt70-140 CFilterBVeto	Sherpa AF2	166.6	1.12	0.35	$7.2 \cdot 10^5$
$Z \to \nu\nu$ MassiveCB Pt70-140 CVetoBVeto	Sherpa AF2	166.6	1.12	0.56	$1.3 \cdot 10^6$
$Z \to \nu\nu$ MassiveCB Pt140-280 BFilter	Sherpa AF2	22.5	1.12	0.10	$1.0 \cdot 10^6$
$Z \to \nu\nu$ MassiveCB Pt140-280 CFilterBVeto	Sherpa AF2	22.5	1.12	0.37	$4.4 \cdot 10^5$
$Z \to \nu\nu$ MassiveCB Pt140-280 CVetoBVeto	Sherpa AF2	22.5	1.12	0.53	$7.0 \cdot 10^5$
$Z \to \nu\nu$ MassiveCB Pt280-500 BFilter	Sherpa FS	1.4	1.12	0.11	$4.1 \cdot 10^4$
$Z \to \nu\nu$ MassiveCB Pt280-500 CFilterBVeto	Sherpa FS	1.4	1.12	0.38	$5.3 \cdot 10^4$
$Z \to \nu\nu$ MassiveCB Pt280-500 CVetoBVeto	Sherpa FS	1.4	1.12	0.51	$2.2 \cdot 10^5$
$Z \to \nu\nu$ MassiveCB Pt500 BFilter	Sherpa FS	0.1	1.12	0.12	$1.0 \cdot 10^4$
$Z \to \nu\nu$ MassiveCB Pt500 CFilterBVeto	Sherpa FS	0.1	1.12	0.40	$1.0 \cdot 10^4$
$Z \to \nu\nu$ MassiveCB Pt500 CVetoBVeto	Sherpa FS	0.1	1.12	0.48	$4.2 \cdot 10^4$

Table 5.5: $Z \to \nu\nu$ Sherpa samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
SinglePhoton Np1 Pt35	Alpgen+Jimmy FS	9554	1.0	1.0	$8.5 \cdot 10^4$
SinglePhoton Np2 Pt35	Alpgen+Jimmy FS	4515	1.0	1.0	$3.0 \cdot 10^4$
SinglePhoton Np3 Pt35	Alpgen+Jimmy FS	1717	1.0	1.0	$1.5 \cdot 10^4$
SinglePhoton Np4 Pt35	Alpgen+Jimmy FS	513.9	1.0	1.0	$4.0 \cdot 10^3$
SinglePhoton Np5 Pt35	Alpgen+Jimmy FS	163.8	1.0	1.0	$7.5 \cdot 10^3$
SinglePhoton Np1 Pt70	Alpgen+Jimmy FS	577.5	1.0	1.0	$8.5 \cdot 10^4$
SinglePhoton Np2 Pt70	Alpgen+Jimmy FS	572	1.0	1.0	$1.0 \cdot 10^5$
SinglePhoton Np3 Pt70	Alpgen+Jimmy FS	306	1.0	1.0	$1.5 \cdot 10^4$
SinglePhoton Np4 Pt70	Alpgen+Jimmy FS	115.90	1.0	1.0	$2.5 \cdot 10^4$
SinglePhoton Np5 Pt70	Alpgen+Jimmy FS	44.3	1.0	1.0	$4.3 \cdot 10^4$
SinglePhoton Np1 Pt140	Alpgen+Jimmy FS	26.2	1.0	1.0	$9.5 \cdot 10^4$
SinglePhoton Np2 Pt140	Alpgen+Jimmy FS	38.7	1.0	1.0	$1.5 \cdot 10^5$
SinglePhoton Np3 Pt140	Alpgen+Jimmy FS	28.6	1.0	1.0	$1.5 \cdot 10^5$
SinglePhoton Np4 Pt140	Alpgen+Jimmy FS	14.2	1.0	1.0	$1.6 \cdot 10^5$
SinglePhoton Np5 Pt140	Alpgen+Jimmy FS	7.0	1.0	1.0	$1.3 \cdot 10^5$
SinglePhoton Np1 Pt280	Alpgen+Jimmy FS	0.83	1.0	1.0	$2.5 \cdot 10^4$
SinglePhoton Np2 Pt280	Alpgen+Jimmy FS	1.7	1.0	1.0	$5.0 \cdot 10^5$
SinglePhoton Np3 Pt280	Alpgen+Jimmy FS	1.5	1.0	1.0	$5.0 \cdot 10^4$
SinglePhoton Np4 Pt280	Alpgen+Jimmy FS	0.92	1.0	1.0	$4.5 \cdot 10^4$
SinglePhoton Np5 Pt280	Alpgen+Jimmy FS	0.54	1.0	1.0	$5.0 \cdot 10^4$
SinglePhoton Np1 Pt500	Alpgen+Jimmy FS	0.03	1.0	1.0	$1.0 \cdot 10^4$
SinglePhoton Np2 Pt500	Alpgen+Jimmy FS	0.076	1.0	1.0	$2.0 \cdot 10^4$
SinglePhoton Np3 Pt500	Alpgen+Jimmy FS	0.077	1.0	1.0	$6.0 \cdot 10^3$
SinglePhoton Np4 Pt500	Alpgen+Jimmy FS	0.051	1.0	1.0	$8.0 \cdot 10^3$
SinglePhoton Np5 Pt500	Alpgen+Jimmy FS	0.033	1.0	1.0	$1.5 \cdot 10^4$

Table 5.6: $\gamma +$ jets Alpgen samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$(W \to \ell\nu)bb$ Np0	Alpgen+Pythia FS	52.3	1.14	1.00	$4.8 \cdot 10^5$
$(W \to \ell\nu)bb$ Np1	Alpgen+Pythia FS	45.5	1.14	1.00	$3.6 \cdot 10^5$
$(W \to \ell\nu)bb$ Np2	Alpgen+Pythia FS	23.7	1.14	1.00	$1.7 \cdot 10^5$
$(W \to \ell\nu)bb$ Np3	Alpgen+Pythia FS	12.5	1.14	1.00	$5.0 \cdot 10^4$
$W \to e\nu$ Np0	Alpgen+Pythia FS	8136.8	1.14	1.00	$3.5 \cdot 10^6$
$W \to e\nu$ Np1	Alpgen+Pythia FS	1791.5	1.14	1.00	$2.5 \cdot 10^6$
$W \to e\nu$ Np2	Alpgen+Pythia FS	541.6	1.14	1.00	$3.8 \cdot 10^6$
$W \to e\nu$ Np3	Alpgen+Pythia FS	146.7	1.14	1.00	$1.0 \cdot 10^6$
$W \to e\nu$ Np4	Alpgen+Pythia FS	37.3	1.14	1.00	$2.5 \cdot 10^5$
$W \to e\nu$ Np5	Alpgen+Pythia FS	11.4	1.14	1.00	$7.0 \cdot 10^4$
$W \to \mu\nu$ Np0	Alpgen+Pythia FS	8133.4	1.14	1.00	$3.5 \cdot 10^6$
$W \to \mu\nu$ Np1	Alpgen+Pythia FS	1792.7	1.14	1.00	$2.5 \cdot 10^6$
$W \to \mu\nu$ Np2	Alpgen+Pythia FS	541.3	1.14	1.00	$3.8 \cdot 10^6$
$W \to \mu\nu$ Np3	Alpgen+Pythia FS	146.5	1.14	1.00	$1.0 \cdot 10^6$
$W \to \mu\nu$ Np4	Alpgen+Pythia FS	37.3	1.14	1.00	$2.6 \cdot 10^5$
$W \to \mu\nu$ Np5	Alpgen+Pythia FS	11.4	1.14	1.00	$6.5 \cdot 10^4$
$W \to \tau\nu$ Np0	Alpgen+Pythia FS	8135.7	1.14	1.00	$3.4 \cdot 10^6$
$W \to \tau\nu$ Np1	Alpgen+Pythia FS	1793.7	1.14	1.00	$2.5 \cdot 10^6$
$W \to \tau\nu$ Np2	Alpgen+Pythia FS	541.2	1.14	1.00	$3.8 \cdot 10^6$
$W \to \tau\nu$ Np3	Alpgen+Pythia FS	146.5	1.14	1.00	$1.0 \cdot 10^6$
$W \to \tau\nu$ Np4	Alpgen+Pythia FS	37.3	1.14	1.00	$2.5 \cdot 10^5$
$W \to \tau\nu$ Np5	Alpgen+Pythia FS	11.5	1.14	1.00	$6.5 \cdot 10^4$
$(W \to \ell\nu)c$ Np0	Alpgen+Pythia FS	758.9	1.14	1.00	$6.5 \cdot 10^6$
$(W \to \ell\nu)c$ Np1	Alpgen+Pythia FS	274.2	1.14	1.00	$2.1 \cdot 10^6$
$(W \to \ell\nu)c$ Np2	Alpgen+Pythia FS	71.6	1.14	1.00	$5.2 \cdot 10^5$
$(W \to \ell\nu)c$ Np3	Alpgen+Pythia FS	16.4	1.14	1.00	$1.1 \cdot 10^5$
$(W \to \ell\nu)c$ Np4	Alpgen+Pythia FS	4.7	1.14	1.00	$2.0 \cdot 10^4$
$(W \to \ell\nu)cc$ Np0	Alpgen+Pythia FS	143.1	1.14	1.00	$1.3 \cdot 10^6$
$(W \to \ell\nu)cc$ Np1	Alpgen+Pythia FS	143.7	1.14	1.00	$1.0 \cdot 10^6$
$(W \to \ell\nu)cc$ Np2	Alpgen+Pythia FS	80.8	1.14	1.00	$5.2 \cdot 10^5$
$(W \to \ell\nu)cc$ Np3	Alpgen+Pythia FS	35.9	1.14	1.00	$1.7 \cdot 10^5$

Table 5.7: $W \to \ell\nu$ Alpgen samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$(Z \to ee)cc$ Np0	Alpgen+Pythia FS	15.1	1.18	1.00	$6.0 \cdot 10^5$
$(Z \to ee)cc$ Np1	Alpgen+Pythia FS	7.2	1.18	1.00	$2.6 \cdot 10^5$
$(Z \to ee)cc$ Np2	Alpgen+Pythia FS	3.0	1.18	1.00	$1.1 \cdot 10^5$
$(Z \to ee)cc$ Np3	Alpgen+Pythia FS	1.2	1.18	1.00	$4.0 \cdot 10^4$
$(Z \to \mu\mu)cc$ Np0	Alpgen+Pythia FS	15.1	1.18	1.00	$6.0 \cdot 10^5$
$(Z \to \mu\mu)cc$ Np1	Alpgen+Pythia FS	7.2	1.18	1.00	$2.7 \cdot 10^5$
$(Z \to \mu\mu)cc$ Np2	Alpgen+Pythia FS	3.0	1.18	1.00	$1.2 \cdot 10^5$
$(Z \to \mu\mu)cc$ Np3	Alpgen+Pythia FS	1.2	1.18	1.00	$4.0 \cdot 10^4$
$(Z \to \tau\tau)cc$ Np0	Alpgen+Pythia FS	15.1	1.18	1.00	$6.0 \cdot 10^5$
$(Z \to \tau\tau)cc$ Np1	Alpgen+Pythia FS	7.2	1.18	1.00	$2.7 \cdot 10^5$
$(Z \to \tau\tau)cc$ Np2	Alpgen+Pythia FS	3.0	1.18	1.00	$1.2 \cdot 10^5$
$(Z \to \tau\tau)cc$ Np3	Alpgen+Pythia FS	1.2	1.18	1.00	$4.0 \cdot 10^4$
$(Z \to ee)bb$ Np0	Alpgen+Pythia FS	8.0	1.18	1.00	$1.5 \cdot 10^5$
$(Z \to ee)bb$ Np1	Alpgen+Pythia FS	3.2	1.18	1.00	$8.0 \cdot 10^4$
$(Z \to ee)bb$ Np2	Alpgen+Pythia FS	1.1	1.18	1.00	$4.5 \cdot 10^4$
$(Z \to ee)bb$ Np3	Alpgen+Pythia FS	0.5	1.18	1.00	$4.5 \cdot 10^3$
$(Z \to \mu\mu)bb$ Np0	Alpgen+Pythia FS	8.0	1.18	1.00	$1.5 \cdot 10^5$
$(Z \to \mu\mu)bb$ Np1	Alpgen+Pythia FS	3.2	1.18	1.00	$8.0 \cdot 10^4$
$(Z \to \mu\mu)bb$ Np2	Alpgen+Pythia FS	1.1	1.18	1.00	$4.5 \cdot 10^4$
$(Z \to \mu\mu)bb$ Np3	Alpgen+Pythia FS	0.5	1.18	1.00	$5.0 \cdot 10^3$
$(Z \to \tau\tau)bb$ Np0	Alpgen+Pythia FS	8.0	1.18	1.00	$1.5 \cdot 10^5$
$(Z \to \tau\tau)bb$ Np1	Alpgen+Pythia FS	3.2	1.18	1.00	$8.0 \cdot 10^4$
$(Z \to \tau\tau)bb$ Np2	Alpgen+Pythia FS	1.1	1.18	1.00	$4.5 \cdot 10^4$
$(Z \to \tau\tau)bb$ Np3	Alpgen+Pythia FS	0.5	1.18	1.00	$5.0 \cdot 10^3$
$Z \to ee$ Np0	Alpgen+Pythia FS	718.9	1.18	1.00	$6.6 \cdot 10^6$
$Z \to ee$ Np1	Alpgen+Pythia FS	175.6	1.18	1.00	$1.3 \cdot 10^6$
$Z \to ee$ Np2	Alpgen+Pythia FS	58.8	1.18	1.00	$4.0 \cdot 10^5$
$Z \to ee$ Np3	Alpgen+Pythia FS	15.6	1.18	1.00	$1.1 \cdot 10^5$
$Z \to ee$ Np4	Alpgen+Pythia FS	3.9	1.18	1.00	$3.0 \cdot 10^4$
$Z \to ee$ Np5	Alpgen+Pythia FS	1.2	1.18	1.00	$1.0 \cdot 10^4$
$Z \to \mu\mu$ Np0	Alpgen+Pythia FS	718.9	1.18	1.00	$6.6 \cdot 10^6$
$Z \to \mu\mu$ Np1	Alpgen+Pythia FS	175.8	1.18	1.00	$1.3 \cdot 10^6$
$Z \to \mu\mu$ Np2	Alpgen+Pythia FS	58.8	1.18	1.00	$4.0 \cdot 10^5$
$Z \to \mu\mu$ Np3	Alpgen+Pythia FS	15.6	1.18	1.00	$1.1 \cdot 10^5$
$Z \to \mu\mu$ Np4	Alpgen+Pythia FS	3.9	1.18	1.00	$3.0 \cdot 10^4$
$Z \to \mu\mu$ Np5	Alpgen+Pythia FS	1.2	1.18	1.00	$1.0 \cdot 10^4$
$Z \to \tau\tau$ Np0	Alpgen+Pythia FS	718.9	1.18	1.00	$6.6 \cdot 10^6$
$Z \to \tau\tau$ Np1	Alpgen+Pythia FS	175.8	1.18	1.00	$1.3 \cdot 10^6$
$Z \to \tau\tau$ Np2	Alpgen+Pythia FS	58.6	1.18	1.00	$4.1 \cdot 10^5$
$Z \to \tau\tau$ Np3	Alpgen+Pythia FS	15.5	1.18	1.00	$1.1 \cdot 10^5$
$Z \to \tau\tau$ Np4	Alpgen+Pythia FS	4.0	1.18	1.00	$3.0 \cdot 10^4$
$Z \to \tau\tau$ Np5	Alpgen+Pythia FS	1.2	1.18	1.00	$1.0 \cdot 10^4$

Table 5.8: $Z \to \ell\ell$ Alpgen samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$Z \to \nu\nu$ Np0 1JetFilter	Alpgen+Jimmy FS	4152	1.232	0.006	$5.0 \cdot 10^3$
$Z \to \nu\nu$ Np1 1JetFilter	Alpgen+Jimmy FS	892.4	1.232	0.46	$8.5 \cdot 10^4$
$Z \to \nu\nu$ Np2 1JetFilter	Alpgen+Jimmy FS	282.1	1.232	0.76	$4.0 \cdot 10^4$
$Z \to \nu\nu$ Np3 1JetFilter	Alpgen+Jimmy FS	82.0	1.232	0.91	$1.5 \cdot 10^4$
$Z \to \nu\nu$ Np4 1JetFilter	Alpgen+Jimmy FS	21.6	1.232	0.96	$4.0 \cdot 10^3$
$Z \to \nu\nu$ Np5 1JetFilter	Alpgen+Jimmy FS	6.6	1.232	0.99	$8.0 \cdot 10^3$
$Z \to \nu\nu$ Np1 1JetFilter Pt70	Alpgen+Jimmy FS	69.5	1.232	1.0	$7.5 \cdot 10^4$
$Z \to \nu\nu$ Np2 1JetFilter Pt70	Alpgen+Jimmy FS	61.0	1.232	0.99	$1.0 \cdot 10^5$
$Z \to \nu\nu$ Np3 1JetFilter Pt70	Alpgen+Jimmy FS	27.5	1.232	0.99	$7.5 \cdot 10^4$
$Z \to \nu\nu$ Np4 1JetFilter Pt70	Alpgen+Jimmy FS	9.4	1.232	0.99	$2.5 \cdot 10^4$
$Z \to \nu\nu$ Np5 1JetFilter Pt70	Alpgen+Jimmy FS	3.4	1.232	1.0	$1.9 \cdot 10^4$
$Z \to \nu\nu$ Np1 1JetFilter Pt140	Alpgen+Jimmy FS	6.0	1.232	1.0	$9.5 \cdot 10^4$
$Z \to \nu\nu$ Np2 1JetFilter Pt140	Alpgen+Jimmy FS	8.0	1.232	1.0	$1.5 \cdot 10^5$
$Z \to \nu\nu$ Np3 1JetFilter Pt140	Alpgen+Jimmy FS	5.2	1.232	1.0	$1.1 \cdot 10^5$
$Z \to \nu\nu$ Np4 1JetFilter Pt140	Alpgen+Jimmy FS	2.3	1.232	1.0	$1.2 \cdot 10^5$
$Z \to \nu\nu$ Np5 1JetFilter Pt140	Alpgen+Jimmy FS	1.0	1.232	1.0	$3.0 \cdot 10^4$
$Z \to \nu\nu$ Np1 1JetFilter Pt280	Alpgen+Jimmy FS	0.24	1.232	1.0	$2.5 \cdot 10^4$
$Z \to \nu\nu$ Np2 1JetFilter Pt280	Alpgen+Jimmy FS	0.46	1.232	1.0	$5.0 \cdot 10^4$
$Z \to \nu\nu$ Np3 1JetFilter Pt280	Alpgen+Jimmy FS	0.39	1.232	1.0	$1.0 \cdot 10^4$
$Z \to \nu\nu$ Np4 1JetFilter Pt280	Alpgen+Jimmy FS	0.22	1.232	1.0	$4.5 \cdot 10^4$
$Z \to \nu\nu$ Np5 1JetFilter Pt280	Alpgen+Jimmy FS	0.12	1.232	1.0	$1.7 \cdot 10^4$
$Z \to \nu\nu$ Np1 1JetFilter Pt500	Alpgen+Jimmy FS	0.001	1.232	1.0	$5.0 \cdot 10^3$
$Z \to \nu\nu$ Np2 1JetFilter Pt500	Alpgen+Jimmy FS	0.023	1.232	1.0	$1.3 \cdot 10^4$
$Z \to \nu\nu$ Np3 1JetFilter Pt500	Alpgen+Jimmy FS	0.023	1.232	1.0	$1.3 \cdot 10^4$
$Z \to \nu\nu$ Np4 1JetFilter Pt500	Alpgen+Jimmy FS	0.014	1.232	1.0	$1.0 \cdot 10^4$
$Z \to \nu\nu$ Np5 1JetFilter Pt500	Alpgen+Jimmy FS	0.009	1.232	1.0	$4.0 \cdot 10^3$
$(Z \to \nu\nu)cc$ Np0	Alpgen+Jimmy AF2	70.8	1.232	1.0	$8.6 \cdot 10^6$
$(Z \to \nu\nu)cc$ Np1	Alpgen+Jimmy AF2	33.8	1.232	1.0	$4.0 \cdot 10^6$
$(Z \to \nu\nu)cc$ Np2	Alpgen+Jimmy AF2	15.0	1.232	1.0	$1.8 \cdot 10^6$
$(Z \to \nu\nu)cc$ Np3	Alpgen+Jimmy AF2	7.2	1.232	1.0	$6.4 \cdot 10^5$
$(Z \to \nu\nu)bb$ Np0	Alpgen+Jimmy AF2	44.1	1.232	1.0	$5.0 \cdot 10^6$
$(Z \to \nu\nu)bb$ Np1	Alpgen+Jimmy AF2	17.5	1.232	1.0	$2.3 \cdot 10^6$
$(Z \to \nu\nu)bb$ Np2	Alpgen+Jimmy AF2	6.4	1.232	1.0	$1.0 \cdot 10^6$
$(Z \to \nu\nu)bb$ Np3	Alpgen+Jimmy AF2	2.3	1.232	1.0	$1.7 \cdot 10^5$

Table 5.9: $Z \to \nu\nu$ Alpgen samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
WW Inclusive	Herwig AF2	32.5	1.68	1.00	$1.0 \cdot 10^7$
WZ Inclusive	Herwig AF2	12.0	1.90	1.00	$2.0 \cdot 10^7$
ZZ Inclusive	Herwig AF2	4.7	1.55	1.00	$2.0 \cdot 10^6$
$Z\gamma \to ee\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$8.8 \cdot 10^6$
$Z\gamma \to \mu\mu\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$9.2 \cdot 10^6$
$Z\gamma \to \tau\tau\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$4.0 \cdot 10^6$
$Z\gamma \to \nu\nu\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$5.5 \cdot 10^6$
$W\gamma \to e\nu\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$1.2 \cdot 10^7$
$W\gamma \to \mu\nu\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$1.2 \cdot 10^7$
$W\gamma \to \tau\nu\gamma$ MassiveCB	Sherpa FS	0.76	1.00	1.00	$6.5 \cdot 10^6$

Table 5.10: Diboson samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
$t\bar{t}$	Powheg+Pythia AF2	253.0	1.00	0.54	$1.0 \cdot 10^8$
$t\bar{t}$	Powheg+Pythia FS	253.0	1.00	0.54	$5.0 \cdot 10^7$
$t \to (W \to e\nu)b$ (s-channel)	McAtNlo+Jimmy AF2	0.6	1.07	1.00	$8.5 \cdot 10^5$
$t \to (W \to \mu\nu)b$ (s-channel)	McAtNlo+Jimmy AF2	0.6	1.07	1.00	$8.4 \cdot 10^5$
$t \to (W \to \tau\nu)b$ (s-channel)	McAtNlo+Jimmy AF2	0.6	1.07	1.00	$8.5 \cdot 10^5$
$(W \to \ell\nu)t$	McAtNlo+Jimmy AF2	20.7	1.08	1.00	$4.4 \cdot 10^6$
$t \to (W \to e\nu)b$ (t-channel)	AcerMC+Pythia AF2	8.6	1.10	1.00	$2.0 \cdot 10^6$
$t \to (W \to \mu\nu)b$ (t-channel)	AcerMC+Pythia AF2	8.6	1.10	1.00	$2.0 \cdot 10^6$
$t \to (W \to \tau\nu)b$ (t-channel)	AcerMC+Pythia AF2	8.6	1.10	1.00	$2.0 \cdot 10^6$

Table 5.11: Top quark samples used in the analysis.

Name	Generator	σ [pb]	k-factor	ϵ_{filter}	N_{sim}
jetjet JZ0W	Pythia8 FS	72850000000.0	1.00	0.99	$6.0 \cdot 10^6$
jetjet JZ1W	Pythia8 FS	72850000000.0	1.00	0.0001	$5.8 \cdot 10^6$
jetjet JZ2W	Pythia8 FS	26359000.0	1.00	0.004	$6.0 \cdot 10^6$
jetjet JZ3W	Pythia8 FS	544190.0	1.00	0.001	$6.0 \cdot 10^6$
jetjet JZ4W	Pythia8 FS	6445.3	1.00	0.001	$6.0 \cdot 10^6$
jetjet JZ5W	Pythia8 FS	39.7	1.00	0.002	$3.0 \cdot 10^6$
jetjet JZ6W	Pythia8 FS	0.4	1.00	0.005	$3.0 \cdot 10^6$
jetjet JZ7W	Pythia8 FS	0.04	1.00	0.02	$3.0 \cdot 10^6$

Table 5.12: Multijet samples used in the analysis.

5.2 Triggers

Because this analysis deals with a few different triggers, it's useful to understand their naming scheme. Trigger names are of the form L_NPX_A, and contain up to five separate elements, corresponding to each letter. The first element L, is the trigger level, which, as mentioned in Chapter 4, can be either L1, L2, or EF. Unless otherwise specified, any level L requires that all of the previous triggers had fired (e.g. EF_g120_loose corresponds to the full chain rather than just the EF level trigger). The element P in a trigger name corresponds to the type of object that is triggered on, and X corresponds to the trigger thresholds applied to it. Triggers can have multiple PX terms corresponding to multi-particle triggers, and the N element is used for identical copies of any PX. The final element A contains any additional information about the algorithms used by the trigger. Some examples of common trigger name elements are listed in Tables 5.13 and 5.14. Typically X is some energy threshold, in GeV, and A is used for isolation requirements.

P	Definition
mu$X$	muon with $p_T > X$ GeV
mu$X i$	isolated muon with $p_T > X$ GeV
gX	photon with $E_{\mathrm{T}} > X$ GeV
jX	central jet with $E_{\mathrm{T}} > X$ GeV and $\eta < 3.2$
fj$X$	forward jet with $E_{\mathrm{T}} > X$ GeV and $\eta > 3.2$
bX	b-jet with $E_{\mathrm{T}} > X$ GeV
xe$X$	Calorimeter $\not{E}_T > X$ GeV
xe$X$T	Calorimeter $\not{E}_T > X$ GeV using BGRP7 at L1

Table 5.13: Some common objects that are selected by triggers in the P term.

Due to mismodeling in the detector simulation, trigger efficiencies can differ between data and MC by a significant amount. To correct this effect, MC events are each given trigger SFs. These SFs are generally functions of some variable related to the trigger, and vary from event to event. They're calculated by measuring the trigger efficiency in both

A	Definition
loose	loose isolation requirements
tight	tight isolation requirements
tclcw	uses topoclusters at EF level
tclcw_loose	tclcw with lower L1/L2 thresholds
a4tchad	topoclusters, anti-kT $\Delta R = 0.4$, EM+JES scale

Table 5.14: Some common A elements in trigger names.

data and MC, through various techniques, and taking the ratio. In principle these SFs should be used every time a trigger is required in an event, although in practice, there are regions where triggers have nearly 100% efficiency in both data and MC and no SF is needed.

For any analysis, it's important to choose the right triggers in order to maximize sensitivity. As mentioned in Chapter 4, prescales are assigned to high rate triggers, lowering the effective luminosity and degrading any results obtained. On the other hand, triggers that aren't prescaled throw away a lot of potentially important data because of their stricter requirements. In physics analyses, prescaled triggers tend to be a last resort and the most commonly used triggers are known as the *lowest prescaled triggers*. The tighter triggers are all subsets of these, making them a good starting point for an analysis of the full $20\,\text{fb}^{-1}$ of 8 TeV data. For reference, the six trigger chains discussed below are listed in Table 5.15.

EF	L2	L1
EF_xe80T_tclcw_loose	L2_xe45T	L1_XE40_BGRP7
EF_xe80_tclcw_loose	L2_xe45	L1_XE40
EF_xe80_tclcw	L2_xe55	L1_XE50
EF_mu24i_tight	L2_mu24_tight	L1_MU15
EF_mu36_tight	L2_mu36_tight	L1_MU15
EF_g120_loose	L2_g120_loose	L1_EM30

Table 5.15: The trigger chains considered in the analysis.

$\not{E}_T$ Triggers To get the highest acceptance, there were three $\not{E}_T$ triggers considered (EF_xe80_tclcw, EF_xe80_tclcw_loose, and EF_xe80T_tclcw_loose), which above 160 GeV $\not{E}_T$, all plateau with an efficiency of over 99% in both data and MC. EF_xe80_tclcw_loose wasn't turned on until later in the 2012 run though, and the BGRP7 L1 trigger ignores certain *bunch crossings*, reducing the effective luminosity by 10% [102]. The third trigger has tighter criteria, but was enabled throughout the entire data collecting period. Therefore, since the region of interest here is well above the efficiency plateau, only EF_xe80_tclcw was used.

Although the offline $\not{E}_T$ calculation takes into account all of the subdetectors, the online $\not{E}_T$ triggers only use calorimeter data. Since muons leave very little trace in the calorimeters, the $\not{E}_T$ triggers will fire in events with high p_T muons and no actual $\not{E}_T$. This is perfect for performing efficiency studies, as is done in Appendix D. The results of this study are used to calculate the trigger SF for events with $\not{E}_T < 160$ GeV. Above 160 GeV, no SFs are applied.

Photon Trigger The single photon trigger EF_g120_loose is used to select high p_T prompt photon events. This trigger plateaus at 125 GeV with over 99% efficiency [103], as shown in Fig. 5.6. By requiring that the reconstructed photon have $p_T > 125$ GeV we avoid the need for any trigger SFs. This trigger is only used as part of the $Z \to \nu\nu$ estimate, detailed in Appendix B, and does not directly enter into the analysis.

Muon Triggers Muon events in this analysis are selected using the logical OR of the EF_mu24i_tight and EF_mu36_tight triggers, which are designed to trigger on events with at least one well measured muon. The lower threshold 24 GeV trigger requires the muon to be isolated, while the 36 GeV does not. As shown in Fig. 5.7, the efficiency of this trigger combination sharply plateaus at just under 25 GeV, after which they have constant $\sim 86\%$ and $\sim 70\%$ efficiencies with respect to p_T in the endcap and barrel detectors respectively [105]. SFs are applied, as functions of η and p_T, to correct $\mathcal{O}(1\%)$ efficiency

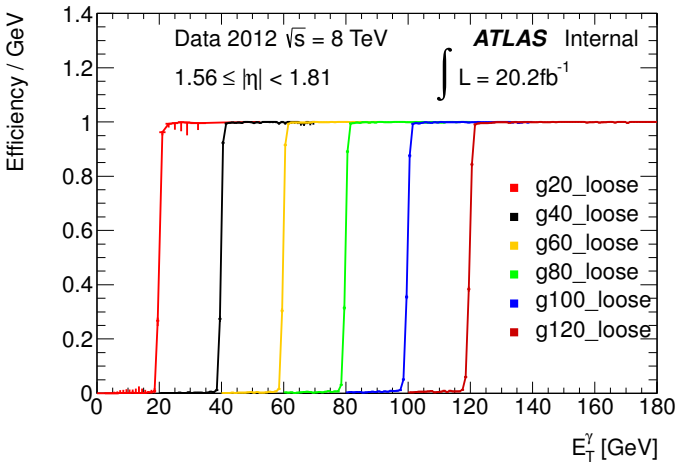

Figure 5.6: Efficiency curves for various photon triggers, including the EF_g120_loose trigger used in this analysis. [104]

differences between data and MC.

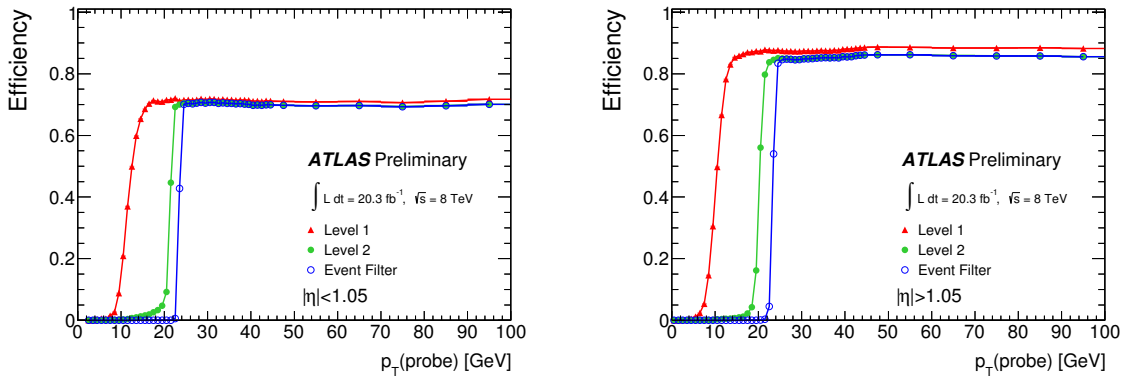

(a) Barrel **(b)** Endcap

Figure 5.7: Efficiency curves for the combination of single muon triggers. [105]

Jet Triggers A number of single jet triggers are used to select the multijet seed sample, and are discussed more in Appendix C. All but one of these are necessarily prescaled, because of the high rate of soft multijet events. To maximize acceptance, the trigger requirement for any given event is the highest threshold trigger that has reached its 99% efficiency threshold on the leading jet's p_T. Because events are only selected in this high efficiency region, no trigger SFs are applied. However, to correct for the luminosity differ-

ences between the different triggers, weights are applied to each event corresponding to the prescale value of the trigger used.

5.3 Object Selection

Physics analyses at ATLAS are done using high-level *objects*, reconstructed from raw data, that represent physical particles that were likely to have been produced in a collision. There are six classes of objects used in this analysis: muons, electrons, photons, jets, b-jets and $\vec{\not{p}}_T$. Given the similarities between this analysis and the sbottom search [102], their object definitions were a used as a starting point, and minor adjustments were made where needed.

In general, each object in an event has associated identification and reconstruction efficiencies. Identification efficiency refers to the probability that a particle is reconstructed as the correct *type* of object, while reconstruction efficiency refers to the probability of it being reconstructed *at all*. Both of these can differ between data and MC, so SFs are applied as corrections.

5.3.1 Primary Vertex

Although not a physical object, the primary vertex plays a crucial role during reconstruction, and is defined as the location at which the interaction that was triggered on, known as the primary interaction, took place. Primary particles are particles produced in the primary interaction, and therefore originate from the primary vertex. Secondary particles, on the other hand, are particles produced after, but as a result of, the collision and originate from vertices displaced from the beam-line. Being able to tell the difference between these two classes of particles is very important, especially when dealing with heavy quark decays, since charm and bottom quarks typically decay at a resolvable distance from the primary vertex. In addition to primary and secondary particles, particles can also be classified as pileup or non-collisional. Pileup particles originate from a separate collision,

and can be traced back to a pileup vertex along the beam-line. Non-collisional particles are those that aren't associated with any vertex, such as cosmic rays.

5.3.2 Muons

Muons are by far the simplest particles to reconstruct offline, partly because they have an entire subdetector dedicated to tracking them. They are reconstructed using the STACO algorithm, which matches ID and MS tracks, and MuTag, which tags ID tracks as muons by using partial MS tracks [106]. All muon tracks are subject to the quality cuts recommended by the *muon combined performance* (MCP) group, that ensure that they're well-measured by the ID. Each muon is likewise required to have a corresponding ID track, and at least a partial MS track within the acceptance region. Every muon that satisfies the requirements listed in Cutflow 5.16 is labeled a *loose* muon.

Name	Cut	Notes
Pixel Hits	$N_{\text{pixel}}^{\text{hit}} + N_{\text{pixel}}^{\text{dead}} \geq 1$	At least one pixel sensor crossed by the track
SCT Hits	$N_{\text{SCT}}^{\text{hit}} + N_{\text{SCT}}^{\text{dead}} \geq 5$	At least five SCT sensors crossed by the track
ID Holes	$N_{\text{pixel}}^{\text{holes}} + N_{\text{SCT}}^{\text{holes}} \leq 2$	No more than 2 missing sensors along the track
TRT Extension	If $0.1 < \lvert\eta_{\text{track}}\rvert < 0.9$, $N_{\text{TRT}} > 5$ and $N_{\text{TRT}}^{\text{outliers}} < 0.9 \cdot N_{\text{TRT}}$	Require a successful TRT extension if the track lies within the TRT coverage
Minimum p_T	$p_T > 6\,\text{GeV}$	Reconstruction efficiency plummets below 6 GeV
ID+MS Coverage	$\lvert\eta_{\text{track}}\rvert < 2.5$	Track lies in region covered by the ID and MS

Cutflow 5.16: Loose muon selection cuts

Although the loose selection has a very high efficiency, it has a relatively low purity of primary muons. Many of the particles reconstructed as loose muons are actually secondary muons produced after the primary interaction, and some of them aren't even muons. When we need a high purity sample of primary muons in the analysis, a stricter selection must be applied. One method of doing this is to add isolation cuts that remove muons in the vicinity of other high energy particles, since they're unlikely to be primary. Isolation cuts set an upper limit on the additional energy measured in a $\Delta R < 0.3$ cone around the

muon's track, for either ID tracks ($p_T\text{Cone}_{30}$), or calorimeter clusters ($E_T\text{Cone}_{30}$). Another method of singling out primary muons is to extrapolate their tracks back to the IP. The track coordinates d_0 and z_0 will be very small for primary muons, but can be quite large for cosmic or secondary muons. *Tight* muons are defined by Cutflow 5.17, which has an extraordinarily high rejection rate for non-primary muons, but a much lower efficiency than the loose selection.

Name	Cut	Notes
Loose	Cutflow 5.16	Loose criteria are applied
High p_T	$p_T > 20\,\text{GeV}$	Reduces the number of secondary muons significantly
Isolation	$p_T\text{Cone}_{30} < 0.12 \cdot p_T$	Track-based isolation
	$E_T\text{Cone}_{30} < 0.12 \cdot p_T$	Calorimeter-based isolated
Primary Vertex	$\lvert z_0 \sin\theta_{\text{track}}\rvert < 0.4\,\text{mm}$	Reject muons originating far from the IP
	$\lvert d_0/\sigma_{d_0}\rvert < 3$	Track's origin must have been resolved well

Cutflow 5.17: Tight muon selection cuts

Simulated muons have their p_T smeared to correct for resolution discrepancies, which is done *before* the selection cuts. SFs are assigned to each *tight* muon in order to correct for the differences in identification and reconstruction efficiencies between data and MC. [106]

5.3.3 Jets

Jets are the most complex objects produced in hadron collisions, but also the most common. Identifying jets begins in the ECAL, by grouping the data into three-dimensional *topoclusters* [107]. Each cluster is constructed from a seed calorimeter cell with a significant energy deposit $|E_{\text{cell}}| > 4\sigma$, where σ is the RMS of the background noise in the cell. Neighboring cells are iteratively added to the cluster if they have $|E_{\text{cell}}| > 2\sigma$, after which the outer layer of surrounding cells is also included.

Once the topoclusters have been collected, they're classified as either hadronic or electromagnetic, depending on their shape. They are then individually calibrated using the *local cluster weighting* (LCW) technique [108], to get a better estimate of the energy of

the actual particles that produced them. First, the HCAL energy deposits are added to hadronic topoclusters and calibrated to account for invisible contributions from muons, neutrinos, or non-ionizing interactions, and lost contributions from low energy deposits not included in the cluster. Both types of clusters are then corrected for any regions of dead material, which can't measure energy deposits. These calibrated topoclusters are passed to the anti-k_t algorithm [109] with $R = 0.4$, which groups them into the jet objects used by the analysis. At the analysis-level, JES corrections are applied to obtain an estimate of the energy of the original parton from the measured energy deposits within each jet.

Jets are grouped into two different categories in this analysis, based on where they're detected. Central jets are ones found within the central region, $|\eta| < 2.4$, and any within the forward region, $2.4 < |\eta| < 4.5$, are forward jets. The p_T cuts of 25 GeV and 30 GeV are applied to central and forward jets respectively. The *jet vertex fraction* (JVF), as defined in Ref. [110], is a quantity used to suppress various sources of non-primary jets and is calculated as the scalar sum of the track p_T originating from the primary vertex divided by the sum of all track p_T associated with the jet. Because of the limited ID coverage, this can only be calculated for central jets which are required to have a JVF greater than 0.5 *if and only if* they have $p_T < 50$. Non-primary jets tend to have a much lower p_T, so for $p_T > 50$ a JVF cut becomes detrimental to the selection.

5.3.4 *b*-jets

Physical *b*-jets are jets produced during the hadronization of a bottom quark, which generally have a very unique signature. With a mass of about 4 GeV, bottom quarks hadronize into B hadrons, which have a typical mass of 5 GeV–6 GeV [34]. As with most heavy fermions, they can usually only decay via weak force interactions, which can result in the production of leptons. Their couplings to up and charm quarks are heavily suppressed by the *CKM matrix*, which strongly favors the top couplings, giving them a characteristically long lifetime of about 10^{-12} s. A B hadron with energy E, mass M, and rest lifetime of τ, will then have an average *observed* lifetime of about $\tau E/M$. Assuming $M \approx 5$ GeV,

$\tau \approx 1 \times 10^{-12}$ s, and a speed of $v \approx c$, the mean distance it would travel before decaying is about 6×10^{-5} m GeV. This translates to about 1 mm–6 mm for the typical b-jet of interest, in the range 20 GeV–100 GeV. This distance scale is large enough to reconstruct a secondary vertex, displaced from the beamline, corresponding to the point at which the B hadron decayed. If the B hadron decays into a D hadron, which contain charm quarks, then a third vertex can sometimes be reconstructed as well. Because they're so much heavier than their final decay products, which are all under 1 GeV in mass, they also tend to produce wider, higher multiplicity jets. These unique properties lead to the concept of b-tagging, where jets can be classified according to how likely they are to be b-jets.

The b-tagging algorithm used by this analysis is called MV1, which is a neural network trained on the results of the three other b-tagging algorithms used at ATLAS (SV1, IP3D, and JetFitter) [111], for which example distributions are plotted in Fig. 5.8. The MV1 output is a number between 0 and 1, where 1 represents the ideal b-jet and 0 represents the ideal light jet. A number of different *working points* have been recommended by the flavor tagging group that correspond to the efficiency of MV1 on b-jets. The working point used in this analysis is 0.9827, which corresponds to a b-tagging efficiency of 60%. This working point provides a purity of about 95%, with rejection factors of 8, 24, and 585 for c-jets, τ-jets, and light jets respectively [112].

b-jets are reconstructed in exactly the same way as light jets, before calculating their MV1 value. The main selection difference is that the p_T requirement is loosened, and the η requirement is tightened. The η cut comes from the coverage of the ID, which is crucial to b-tagging, and the lower p_T cut is to increase acceptance, since fake jets are much less likely to pass the b-tag requirement. The selection cuts for b-jets are summarized in Cutflow 5.18.

SFs are applied to *all* simulated jets to correct for mismodeling of the b-jet identification efficiency. Using the truth information from the MC generator to determine the flavor of each jet, separate sets of SFs are calculated for each, as functions of p_T and η. Due to limited statistics, the b-tag calibration SFs are only calculated explicitly up to 300 GeV for c-jets and b-jets and 750 GeV for light jets [113]. Beyond this limit, extrapolations are used

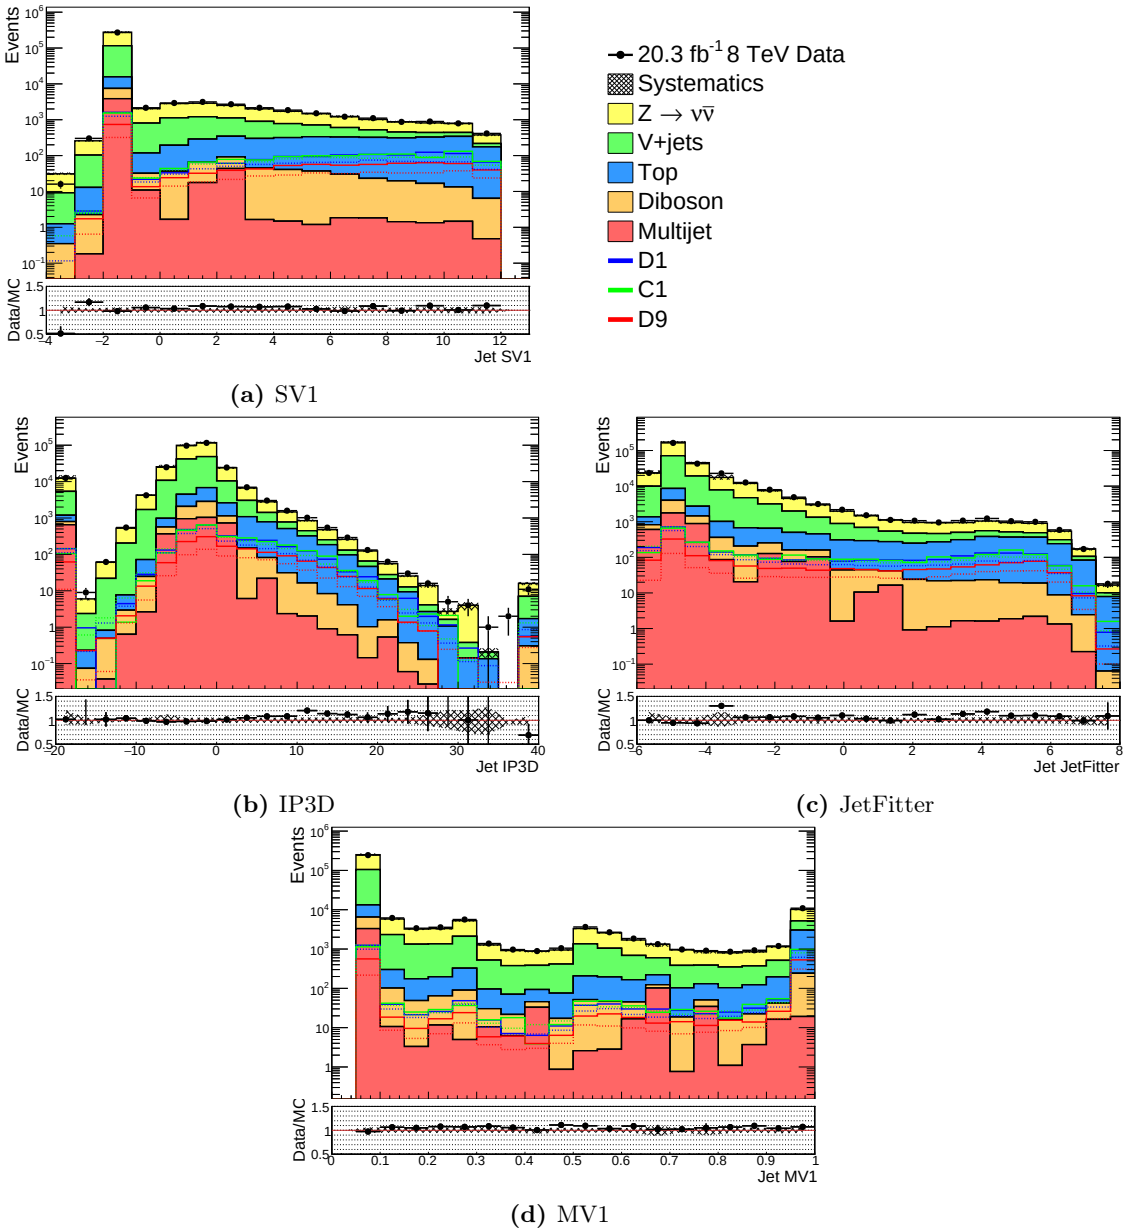

(a) SV1

(b) IP3D

(c) JetFitter

(d) MV1

Figure 5.8: Distributions for the four b-tagging algorithms used. The outputs of SV1, IP3D, and JetFitter are fed into MV1 to get the final discriminant.

Name	Cut	Notes
High p_T	$p_T > 20\,\text{GeV}$	Displaced vertices only at high p_T
JVF	$p_T > 50\,\text{GeV}$ or JVF> 0.5	Suppress pileup contributions to low p_T b-jets
ID Coverage	$\lvert\eta\rvert < 2.5$	ID tracks are necessary for b-tagging
MV1	MV1 > 0.9827	b-tag corresponding to 60% efficiency

Cutflow 5.18: b-jet selection cuts

to estimate the effect.

5.3.5 Electrons

Electrons are detected by ATLAS in a similar fashion as hadrons, although they're significantly simpler. Electrons are absorbed by the ECAL through the generation of electromagnetic showers, which consist primarily of Bremmsstrahlung and pair production at high energies. These processes are very well understood, and contain many fewer degrees of freedom than hadronic showers. Additionally, electrons only have a single associated ID track leading into the cluster, and don't typically reach the HCAL. So while similar algorithms can be used to reconstruct electrons, the background rejection rate and energy resolution are much better than for jets.

Analysis-level electrons are selected by applying Cutflow 5.19. Electrons are only used as an event veto, so the selection is fairly loose to keep the efficiency high. Each reconstructed electron has an object quality flag from the ECAL, which flags electrons near problematic cells. Electrons are required to pass this flag, as well as the medium++ criteria set by the eGamma group [114]. The only other cuts applied to the electron selection are on the acceptance region and a minimal p_T cut to avoid vetoing secondary electrons, which are common in heavy flavor decays.

Name	Cut	Notes
Minimum p_T	$p_T > 7\,\text{GeV}$	Reconstruction efficiency begins to plummet at around 7 GeV
ID Coverage	$\lvert\eta_{\text{cluster}}\rvert < 2.47$	Must be within the acceptance region of the ID
Quality	medium++	Defined by the eGamma group [114]

Cutflow 5.19: Electron selection cuts

As with jets, electrons are calibrated to account for energy scale variations in the ECAL. The momenta of simulated electrons are also smeared due to mismodeling of the ECAL in MC [115]. SFs are typically used for the reconstruction and identification efficiencies [116], but since no electrons are actually used in the analysis these are irrelevant.

5.3.6 Photons

Photon detection is very similar to electron detection, producing the same type of electromagnetic showers in the ECAL. The two are distinguished from each other by using the ID data, for which electrons have one corresponding track and *unconverted* photons have none. In the interaction with the ID material, *conversion* of a photon to an electron-positron pair is common, and results in *two* tracks originating from a displaced vertex. Converted and unconverted photons create different signatures, but they're both easily distinguishable from electrons. The main background to photon reconstruction is from primary π^0 mesons, which decay to two photon almost instantly and can be mistaken for a single photon if they're highly collimated. Calorimeter isolation is used to reduce this background, since π^0 production is usually accompanied by other hadronic processes and the showers from two photons are typically wider than from one.

Loose photons are selected by applying Cutflow 5.20. As with electrons, the object quality flag of each photon must show that the photon cluster does not coincide with any problematic cells and its location must fall within the acceptance region. The tight requirements from the eGamma working group are placed on the electromagnetic shower shape [117], and photons with p_T under 10 GeV are ignored to meet the recommendations for the $\not{E}_T$ calculation. *Tight* photons must satisfy the additional p_T and isolation requirements in Cutflow 5.21 [118].

Name	Cut	Notes
Quality	tight	Defined by the eGamma group Ref. [117]
Minimum p_T	$p_T > 10\,\text{GeV}$	Reconstruction efficiency is very low below 10 GeV
ID Coverage	$\lvert\eta_{\text{cluster}}\rvert < 2.37$	Must be within the acceptance region of the ID
Crack Veto	$1.37 \leq \lvert\eta\rvert$ or $\lvert\eta\rvert \geq 1.52$	Region around the calorimeter crack is unreliable

Cutflow 5.20: Loose photon selection cuts

Photons, being very similar to electrons in the calorimeter, have the same type of calibrations, smearing, and SFs applied [115,119]. One additional correction is also used, which corrects an issue with the FS samples. "Fudge Factors" are applied to the electromagnetic

Name	Cut	Notes
Loose	Cutflow 5.20	Loose criteria are applied
High p_T	$p_T > 125\,\text{GeV}$	Puts the photon trigger at 100% efficiency
Isolation	$E_T\text{Cone}_{40} < 5\,\text{GeV}$	ECAL isolation to remove large π^0 background

Cutflow 5.21: Tight photon selection cuts

shower shape to bring the simulation into agreement with data.

5.3.7 Missing E_T

By using the conservation of transverse momentum, $\vec{\not{p}}_T$ can be measured by summing all the p_T vectors associated with a particular collision. Since these should sum to 0 in principle, any non-zero value of $\not{E}_T = |\vec{\not{p}}_T|$ suggests that energy escaped the detector in one form or another. After calibrating and smearing all of the physics objects, the $\not{E}_T$ will be altered and needs to be recalculated. In this analysis we use a variant of the algorithm known as 'RefFinal', which calculates the total $\not{E}_T$ of an event, and accounts for pileup [120]. Here, since we don't reconstruct hadronically decaying τ-jets as separate objects, any tau lepton contributions are included in the jet term. There are five components to the $\not{E}_T$ calculation,

$$\vec{\not{p}}_T + \vec{p}_T^{\,e} + \vec{p}_T^{\,\gamma} + \vec{p}_T^{\,\mu} + \vec{p}_T^{\,j} + \vec{p}_T^{\,\text{soft}} = 0, \tag{5.1}$$

which correspond to the electron, photon, muon, jet, and soft radiation contributions to the total $\not{E}_T$ respectively. The $\vec{p}_T^{\,\text{soft}}$ term includes all topoclusters not included in reconstructed objects, and all *soft jets* with $p_T < 20\,\text{GeV}$. The other terms are calculated before overlap removal using,

$\vec{p}_T^{\,j}$ *All* reconstructed jets in the event with $p_T > 20\,\text{GeV}$

$\vec{p}_T^{\,\mu}$ Loose muons with $p_T > 10\,\text{GeV}$

$\vec{p}_T^{\,e}$ Loose electrons with $p_T > 10\,\text{GeV}$

$\vec{p}_T^{\,\gamma}$ Loose photons with $p_T > 10\,\text{GeV}$

Each object is assigned a weight on E_T, p_x, and p_y to correct for deviations between the cluster variables, which determine $\not{E}_T$, and the track variables, which determine the 4-momenta. These weights are propagated to the recalculation, where they're applied to the calibrated 4-momenta to obtain the final $\not{E}_T$. Overlap is removed on a cell-by-cell basis, where no calorimeter cell can be used more than once in the calculation.

In addition to the nominal $\not{E}_T$, there are two other variations used in the analysis. The first, denoted $\not{E}_T^{\mu}$, is the calorimeter $\not{E}_T$ obtained by recalculating the $\not{E}_T$ without the muon term. The second is $\not{E}_T^{\gamma}$, which is calculated without photon contributions. These are used in various CRs to simulate the effect of particles escaping detection.

5.3.8 Overlap Removal

Using only the object selections defined above, one physical object can easily be double-counted and reconstructed as multiple objects. This can be caused by both real effects, where the two objects correspond to two real particles, and fake effects, where one particle is reconstructed twice. For example, a secondary muon emerging from a b-jet is considered part of that b-jet, rather than a separate muon object at the analysis level, and should be removed. On the other hand, electrons are frequently reconstructed as *both* an electron and a jet, in which case we want the jet removed. In addition to these two examples, there are a number of other overlap effects that need to be addressed after object selection. The procedure is as follows:

- For any two electrons within $\Delta R < 0.05$ of each other, remove the lower p_T one.
- Remove any electrons within $\Delta R < 0.2$ of a loose photon.
- Remove any light jets within $\Delta R < 0.2$ of a loose electron or photon. This step is not taken for b-jets (defined by the looser 80% working point), since it drastically reduces their acceptance.
- Remove any electrons, muons, or photons within $\Delta R < 0.4$ of a jet. This must be

done *after* the previous step.

5.4 Event Preselection

There are two stages to the preselection of events before the analysis-level cuts are applied. The first, is a set of cleaning cuts applied to every event, which are designed to avoid problematic events and select a sample of generally well measured ones. Any event passing these is then put through a *topological* preselection. There are five different event topologies used in this analysis, each with its own set of preselection cuts. These are applied in order to ensure that the relevant objects in each event are accurate representations of the underlying physical process. These topological preselections are all orthogonal to each other, so that any single event will pass, at most, one of them.

5.4.1 Common Preselection

The following are the standard preselection cuts applied by most of ATLAS in order to ensure high quality data. These cuts are designed to rule out various sources known to result in irreparable collisional data. Many of them are applied to simulation as well, either because the detector simulation models the effect or because the cut affects relevant processes. The acceptances for each cut on the three data streams are listed in Table 5.22.

Cut	JetTauEtMiss	Muons	Egamma
GRL	94	92	96
LAr	94	92	95
Tile Error	94	92	95
Tile Trip	94	92	95
Core Flags	94	92	95
Bad Jets	90	90	93
Dead Tiles	89	90	93
Primary Vertex	89	89	93
Bad Muons	89	89	93
Lepton Overlap	89	89	93
Cosmic Muons	88	81	92

Table 5.22: Percentage acceptance of the common preselection on the data streams.

Good Runs List Every event recorded at ATLAS contains *data quality* (DQ) flags that indicate the status of all the subdetectors. The DQ group, using these flags, has compiled *good runs lists* (GRLs) that contain recommended *lumiblocks* to use in physics analyses. The integrated luminosity of the 2012 data is computed using these lumiblocks, and therefore any event not in the GRL should be excluded. Simulated events don't have DQ flags, so no GRL cut is placed on them.

Data Quality In addition to GRL requirements, there are a few additional flags that need to be checked to ensure quality data. LAr noise bursts can occur in the calorimeters, and certain tiles in the HCAL can become corrupted. There can also be incomplete events where only a fraction of the detector is functional, due to the detector restarting in the middle of a run. Data events with flags indicating any of these issues are vetoed and, as with the GRL cut, these aren't required in MC.

Jet Cleaning After reconstructing the jets in an event, certain checks need to be made to avoid *bad* or *ugly* jets. Bad jets are reconstructed jet objects that don't correspond to any physical jet coming from a collision. Various noise effects in the calorimeters, and non-collisional backgrounds can both mimic real jets. Bad jets are selected with the *looser* criteria recommended by the JetEtMiss group, and any event with bad jets after overlap removal is vetoed in both data and MC.

Ugly jets are formed from real jets falling near broken regions of the HCAL. Throughout the 2012 data taking period, small regions of the HCAL were turned off for a number of reasons. This was far too common, and not serious enough to justify vetoing every event it occurred in. Instead, the broken modules were added to the detector simulation and for most events, they don't cause any problems. If a jet is produced close to one of these dead regions though, it can be very badly reconstructed and affect both the jet p_T and the total $\not{E}_T$. If the energy correction applied to account for these dead regions is too large, the jet is considered ugly. Events with such jets, in both data and MC, are vetoed using the *medium*

criteria recommended by the JetEtMiss group.

Beam Background An additional jet cleaning cut was recommended by the mono-jet group [47] to remove beam backgrounds. The jet charge fraction, which is the ratio of the jet's total track p_T to its calorimeter E_T, is a measure of how much of the jet is charged and passes through the ID. This is required to be at least 10% of the maximum energy fraction in any calorimeter layer, for the leading jet in the event. The effect of this cut is minimal on any good event, but it drastically reduces beam backgrounds. To check this, the ϕ distribution of the leading jet can be examined, as is done in Fig. 5.9. Beam backgrounds are highly asymmetric in ϕ, causing structure on top of the flat expectation.

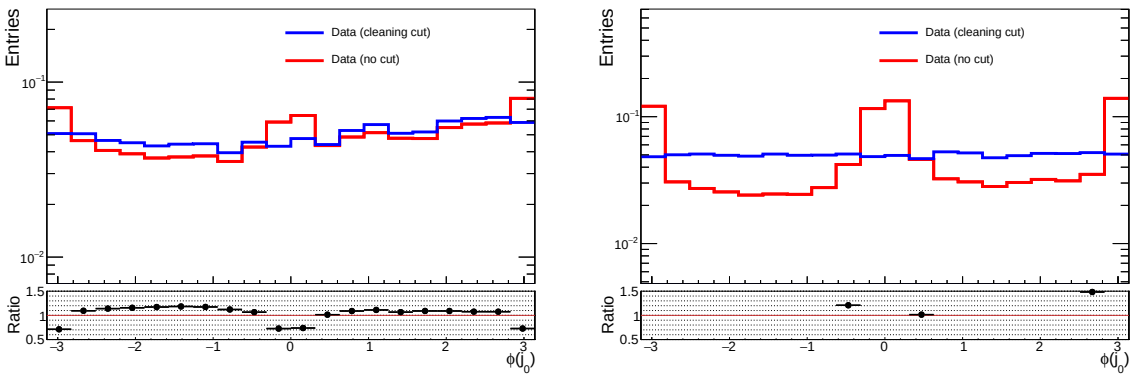

(a) Preselection **(b)** $\not{E}_T > 200\,\text{GeV}$

Figure 5.9: The normalized ϕ distribution of the leading jet, with and without the beam background rejection cut at two different selection stages. In the high $\not{E}_T$ region there is significant structure, suggesting the presence of beam backgrounds.

Primary Vertex If the primary vertex has fewer than 4 tracks with a p_T of at least 500 MeV, the event is excluded from the analysis. The resolution of the reconstructed interaction vertex is heavily dependent on the number of tracks emerging from it, and without a well measured primary vertex, the event becomes very difficult to distinguish from pileup, degrading the $\not{E}_T$ measurement. Because soft QCD backgrounds are so abundant in hadron collisions, this cut has very little effect on any legitimate events of interest.

Bad Muon Veto A bad muon is defined as one that satisfies $\sigma(q/p)/|q/p| > 0.2$, for charge q and momentum p, meaning that the momentum measurement has a significant uncertainty. This will interfere with the $\not{E}_T$ calculation, so any event containing a bad muon is vetoed. Note that this cut is applied *before* overlap removal, so it includes muons originating from b-jets.

Lepton Overlap Veto If a baseline electron and muon are reconstructed within $\Delta R < 0.01$, the most likely scenario is that a *real* muon transferred a significant amount of energy to an atomic electron, creating a *delta ray*. In this case neither is likely to be well measured and the event is removed. In principle, the electron momentum could be added to the muon, but it has little effect on any analysis with a lepton veto.

Cosmic Muon Veto Muons are one of the few particles able to easily pass through the earth and reach ATLAS. They are also produced abundantly in the atmosphere from the decays of pions created in cosmic ray collisions. If a muon track is sufficiently displaced from the primary vertex, it's most likely to be a cosmic ray and the entire event should be excluded. Therefore, events with loose muons that have $|z_0| > 1\,\text{mm}$ or $|d_0| > 0.2\,\text{mm}$, relative to the primary vertex, are vetoed.

Pileup Reweighting The MC12a simulations, which include all of the samples used here, were produced well before the 2012 data was collected. Since the pileup profile would have been impossible to predict with any accuracy, a very generic one was chosen. In addition to the trigger and object SFs used for MC events, there are also corrections that need to be applied to take the mismodeling of these pileup effects into account [121]. Fig. 5.10 shows the distribution of $\langle\mu\rangle$, the average number of interactions per bunch crossing, that was used for this round of MC production, compared with the observed distribution in the 8 TeV data.

Although a lot of work has been put into removing it, pileup can affect an analysis drastically by introducing particles that didn't actually come from the primary vertex. In

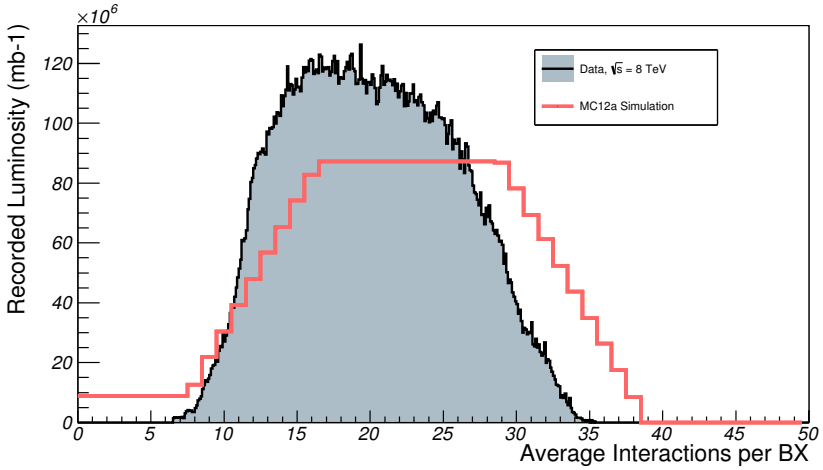

Figure 5.10: Comparison of observed $\langle\mu\rangle$ with the distribution used in early MC simulations. The MC distribution is scaled to have the same normalization as the data.

principle, the results of an analysis should be independent of $\langle\mu\rangle$, but this can never be achieved perfectly. To account for this effect, it's important to have $\langle\mu\rangle$ well-modeled by MC, which can be done by reweighting events on $\langle\mu\rangle$. Each MC event is randomly assigned some continuous sequence of data run numbers, from which the luminosity-averaged $\langle\mu\rangle$ distribution is used to calculate a SF.

An additional complication comes from the fact that the ID visible cross-section, defined as the $p-p$ inelastic cross-section times the efficiency of detection by the ID, is overestimated by Pythia8. This creates a bias towards larger $\langle\mu\rangle$ values in MC, which is well modeled by a simple scaling of the $\langle\mu\rangle$ value for either MC or data, before reweighting. Since the simulated $\langle\mu\rangle$ only takes on integer values, it makes more sense to scale the measured $\langle\mu\rangle$ from data, which is a continuous variable averaged over multiple collisions. Dividing $\langle\mu\rangle$ in each data event by 1.09 provides excellent agreement after applying the pileup reweighting.

5.4.2 Topological Preselections

There are five distinct sets of preselection cuts used by this analysis, corresponding to the five event topologies of interest,

Zero Lepton Hadronic events with high $\not{E}_T$ and no leptons, defined by Cutflow 5.23

Single Muon Events with a single well measured muon, defined by Cutflow 5.24

Di-muon Events with two well measured muons that are likely to originate from the decay of a Z^0 boson, defined by Cutflow 5.25

Prompt Photon Events with a single high p_T photon and no leptons, defined by Cutflow 5.26

Multijet Events with low $\not{E}_T$, no leptons, and at least one high p_T jet, defined by Cutflow 5.27

These preselection cuts are made orthogonal by adding loose selection cuts where necessary, which splits all of the used data into five distinct regions. The effect of the topological preselection cuts on data and MC can be seen in Tables 5.28-5.31.

Name	Cut	Notes
Trigger	EF_xe80_tclcw	See Appendix D
Lepton Veto	$n_\ell^{\text{loose}} = 0$	No loose leptons
Photon Veto	$n_\gamma^{\text{tight}} = 0$	Orthogonality requirement
Jet Requirement	$n_j > 0$	Must recoil against something
$\not{E}_T$	$\not{E}_T > 100\,\text{GeV}$	Reject low $\not{E}_T$ events

Cutflow 5.23: $\not{E}_T$ preselection cuts.

Name	Cut	Notes
Quality	$n_\mu^{\text{tight}} > 0$	There must be a tight muon
Trigger	EF_mu24i_tight or EF_mu36_tight	Single muon triggers
	Trigger Matching	It must be matched to the trigger
Lepton Veto	$n_\ell^{\text{loose}} = 1$	No additional loose leptons
$\not{E}_T^\mu$	$\not{E}_T^\mu > 100\,\text{GeV}$	Reject low $\not{E}_T^\mu$ events

Cutflow 5.24: Single muon preselection cuts.

Name	Cut	Notes
Quality	$n_\mu^{\text{tight}} \geq 2$	Two tight muons
Trigger	EF_mu24i_tight or EF_mu36_tight	Single muon triggers
	Trigger Matching	One of them must match a trigger
Z^0 Decay	$\mid m_{\mu\mu} - M_Z \mid < 20\,\text{GeV}$	Must be consistent with a Z^0 decay
	Opposite Sign	
Lepton Veto	$n_\ell^{\text{loose}} = 2$	No additional loose leptons
$\not{E}_T^\mu$	$\not{E}_T^\mu > 100\,\text{GeV}$	Reject low $\not{E}_T^\mu$ events

Cutflow 5.25: Di-muon preselection cuts.

Name	Cut	Notes
Trigger	EF_g120_loose	Single photon trigger
Quality	$n_\gamma^{\text{tight}} \geq 1$	One tight photon
Lepton Veto	$n_\ell^{\text{loose}} = 0$	No loose leptons
$\not{E}_T^\gamma$	$\not{E}_T^\gamma > 100\,\text{GeV}$	Reject low $\not{E}_T^\gamma$ events

Cutflow 5.26: Single photon preselection cuts

Name	Cut	Notes
Trigger	EF_jX_a4tchad	See Appendix C
Lepton Veto	$n_\ell^{\text{loose}} = 0$	No loose leptons
Photon Veto	$n_\gamma^{\text{tight}} = 0$	Orthogonality requirement
$\not{E}_T$ Significance	$\not{E}_T^{sig} < 0.7$	Low $\not{E}_T$
Leading Jet p_T	$p_T(j_0) > 130\,\text{GeV}$	At least one high p_T jet

Cutflow 5.27: Multijet preselection cuts.

Cut	JetTauEtMiss	$Z \to \nu\nu$	$W \to \ell\nu$	$t\bar{t}$
Event Preselection	81	26	34	93
Lepton Veto	77	25	25	35
Photon Veto	77	25	25	35
Trigger	2.9	2.3	0.50	7.2

Table 5.28: Percentage acceptance of the $\not{E}_T$ preselection on data and MC.

Cut	Muons	$Z \to \ell\ell$	$W \to \ell\nu$	$t\bar{t}$
Event Preselection	46	45	34	93
Trigger	15	6.4	3.2	22
Tight	4.3	3.3	2.8	17
Muon p_T	3.8	3.1	2.7	16
Trigger Matching	3.8	3.0	2.7	16
Electron Veto	3.8	3.0	2.7	14
Muon Veto	3.7	2.3	2.7	14

Table 5.29: Percentage acceptance of the single muon preselection on data and MC.

Cut	Muons	$Z \to \ell\ell$	$W \to \ell\nu$	$t\bar{t}$
Event Preselection	46	45	34	93
Trigger	15	6.4	3.3	22
Tight	0.26	2.7	0.00	0.93
Trigger Matching	0.26	2.7	0.00	0.93
Muon p_T	0.22	2.3	0.00	0.78
Opposite Sign	0.22	2.3	0.00	0.78
Z Mass	0.21	2.2	0.00	0.23
Electron Veto	0.21	2.1	0.00	0.23
Muon Veto	0.21	2.1	0.00	0.22

Table 5.30: Percentage acceptance of the di-muon preselection on data and MC.

Cut	Egamma	$\gamma + \text{jets}$	$V\gamma$	$t\bar{t}$
Event Preselection	69	97	42	93
Trigger	4.2	42	0.28	1.8
Tight	0.46	26	0.11	0.05
p_T	0.46	26	0.11	0.05
Lepton Veto	0.45	26	0.05	0.04

Table 5.31: Percentage acceptance of the single photon preselection on data and MC.

5.5 Signal Selection

To carry out this search, SRs must be chosen that maximize the sensitivity of the analysis. As mentioned earlier, it's important to remain blinded during this process, by not looking at the data in any region predicted to contain a significant amount of signal over the SM backgrounds. So before we select the SRs, we need estimates for each signal and background process. The signal samples are relatively straightforward, but the background processes pose a problem. We need to have a signal selection set before choosing appropriate CRs and VRs, but without these regions we can't say with any certainty that the background estimates are good. In practice, this circular problem is solved iteratively with a method that basically consists of educated guessing and checking. To avoid that confusion here, I will just present and discuss the final SRs that were chosen. Section 5.6 will go through the CRs and VRs, and Appendix A discusses how the SRs were chosen and optimized.

Name	Cut	Notes
Preselection	Event Preselection 0 Lepton Preselection	
Jet Multiplicity	$0 < n_j < 5$	DM + $t\bar{t}$ orthogonality
$\Delta\phi_{min}$	$\Delta\phi_{min} > 1$	Multijet removal
Jet p_T	$p_T(j_0) > 100\,\text{GeV}$	Redundant cleaning cut
Base $\not{E}_T$	$\not{E}_T > 200\,\text{GeV}$	Redundant cleaning cut
b-jet Multiplicity	$n_b > 0$	60% working point
b-jet p_T	$p_T(b_0) > 100\,\text{GeV}$	Recoiling b-jet
SR300	$\not{E}_T > 300\,\text{GeV}$	
SR350	$\not{E}_T > 350\,\text{GeV}$	
SR400	$\not{E}_T > 400\,\text{GeV}$	Signal Regions
SR450	$\not{E}_T > 450\,\text{GeV}$	
SR500	$\not{E}_T > 500\,\text{GeV}$	

Cutflow 5.32: The nominal SR selection cuts. Only a single region is used for any specific signal process.

In addition to just $\not{E}_T$, there are a number of kinematic variables we will be looking at here, many of which are not directly related to the analysis. Because the selection

optimization and background estimates all use a pretty limited set of variables, it can be useful to look at unrelated variables for validation. A background estimate is more reliable if it can accurately predict distributions that are uncorrelated with the ones it was fit to. The different types of variables that will be used throughout the remainder of this analysis are described below.

n_x This represents the multiplicity of some type of reconstructed object x. Unless otherwise specified, this corresponds to the tightest selection cuts available for x. The variables n_ℓ, n_μ, n_γ, n_j, n_b correspond to the number of tight leptons, muons, photons, jets, and b-jets respectively. The topological preselections cut heavily on these, so they are of limited use in the analysis.

$p_T(x_n)$ This corresponds to the p_T of the nth x object, in order of decreasing p_T. For example, $p_T(j_0)$, $p_T(b_0)$, $p_T(\mu_0)$ and $p_T(\gamma_0)$ correspond to the p_T of leading jet, b-jet, muon, and photon in the event respectively. Since no process considered here has more than two objects defined by the event topology, n will usually be 0 or 1.

$\Delta\phi_{min}$ This is the minimal separation in the ϕ coordinate between the $\vec{\not{p}}_T$ and any reconstructed jet in the event. A common way to get high $\not{E}_T$ backgrounds is from events where a jet contains real or fake $\not{E}_T$ contributions. Fake $\not{E}_T$ can occur when the object's momentum isn't measured properly, and real $\not{E}_T$ is when particles escape detection. In both of these situations, if the jet is still able to be reconstructed, the $\not{E}_T$ it produces will be collinear with it's $\vec{p}_T$, giving a very low value for $\Delta\phi_{min}$. On the other hand, the signal process involves WIMPs recoiling *against* b-jets, so $\Delta\phi_{min}$ will typically be very large. This variable provides excellent separation between the signal and backgrounds that have no primary leptons or neutrinos. There are a number of closely related variables that are also used, where the definition of $\vec{\not{p}}_T$ is modified. Examples are $\Delta\phi^{\mu}_{min}$ and $\Delta\phi^{\gamma}_{min}$, which use a $\vec{\not{p}}_T$ recalculated without the muon and photon contributions respectively.

m_{xx} The invariant mass between the two leading x objects, or 0 if they don't exist. In the CRs, $m_{\mu\mu}$ is directly cut on to isolate $Z \to \mu\mu$ events. The quantities m_{jj} and m_{bb} are also useful for validation since they have varying shapes for different background processes.

H_T^n This is the scalar sum of each jet object's p_T, $\sum_i p_T(j_i)$, for $i \geq n$. H_T^0 is often abbreviated H_T and sets the overall energy scale of hadronic events. H_T^1 and H_T^2 can also be useful since they contain information about the soft jets, which aren't necessarily related to the $\not{E}_T$.

m_T^x The transverse mass between an object x and the $\vec{\not{p}}_T$, defined as,

$$\sqrt{2E_T(x)\not{E}_T(1-\cos\Delta\phi)} \tag{5.2}$$

in the limit of $E_T(x)\not{E}_T \gg m_x m_\chi$, where $E_T(x)$ is the transverse energy of x, and $\Delta\phi$ is the azimuthal angular separation between $\vec{p}_T(x)$ and $\vec{\not{p}}_T$. Here, m_T^ℓ and m_T^b will be useful variables for validation. In backgrounds with leptonic $W^\pm$ decays, such as $W \to \ell\nu$ and $t\bar{t}$, m_T^ℓ is directly related to the mass of the $W^\pm$ boson. The m_T^b variable is a very good signal discriminant, as discussed in Appendix A, because it depends on both the momentum transferred in the collision *and* the separation between $\vec{\not{p}}_T$ and $\vec{p}_T(b_0)$. Requiring a high m_T^b can simultaneously reduce *all* of the SM backgrounds with respect to the signal.

Figs. 5.11-5.19 show distributions for a selection of these variables at each stage of the selection. Although there is some noticeable mismodeling in earlier stages, it's all well covered by the *systematic uncertainties*, discussed in Section 5.7, and it becomes indistinguishable over statistical fluctuations in the later stages. During the development of this analysis, any event with a b-jet *and* $\not{E}_T > 300\,\text{GeV}$ was kept blinded to avoid bias. After the analysis was frozen, the data was unblinded and the resulting yields are listed in Table 5.33.

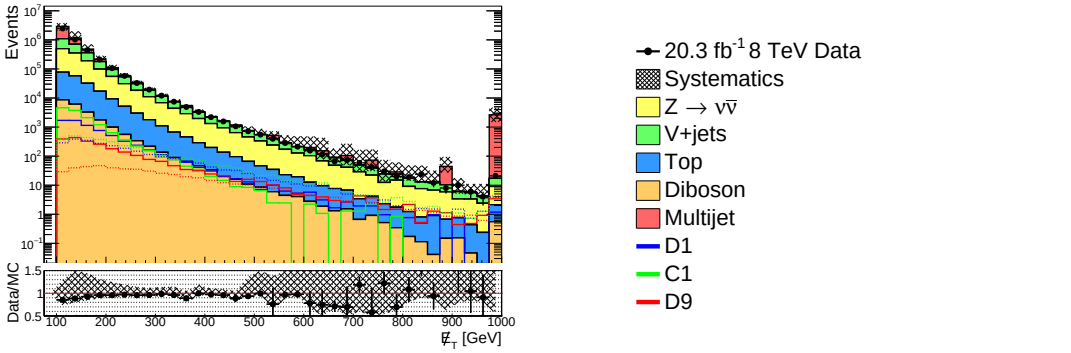

(a) Preselection

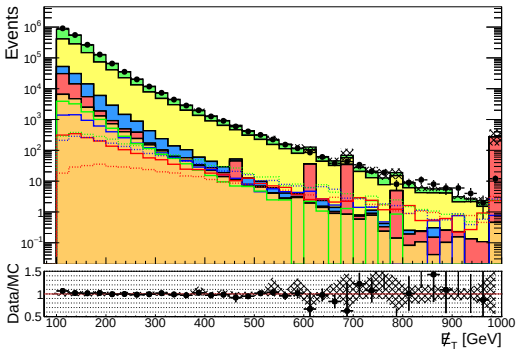

(b) $\Delta\phi_{min}$

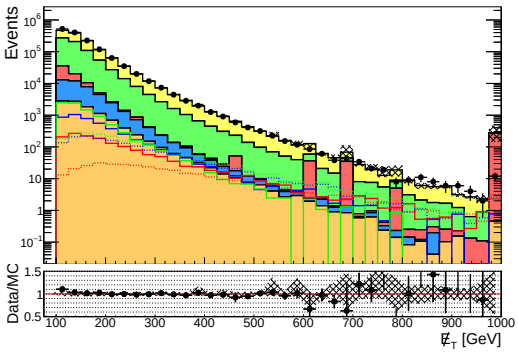

(c) Jet p_T

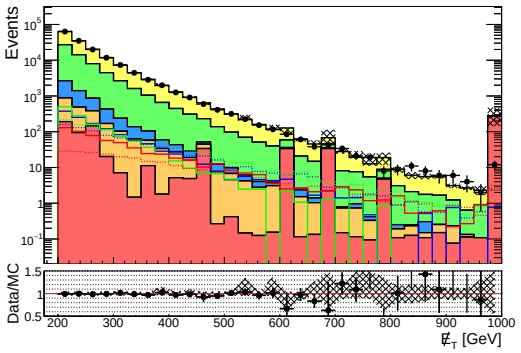

(d) Base $\not{E}_T$

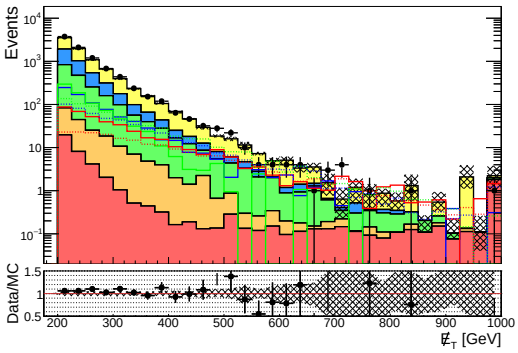

(e) b-jet Multiplicity

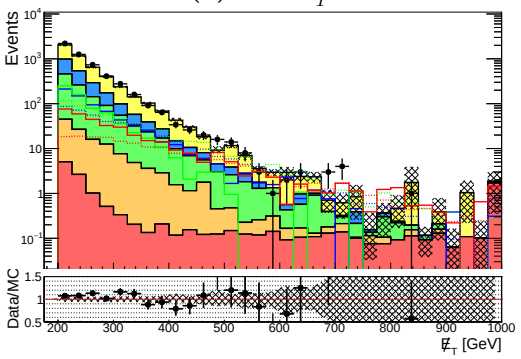

(f) b-jet p_T

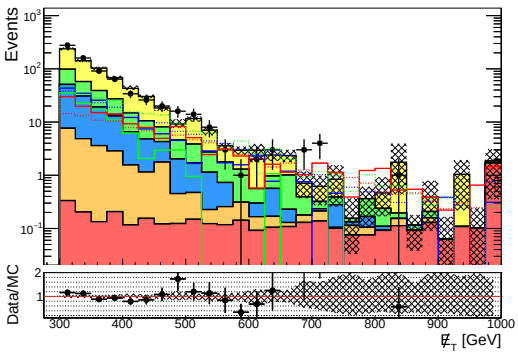

(g) SR300

Figure 5.11: $\not{E}_T$ distributions after each of the selection cuts.

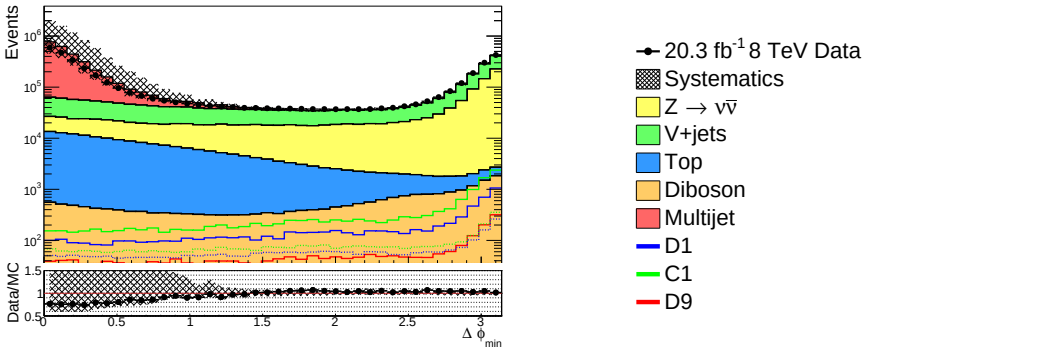

(a) Preselection

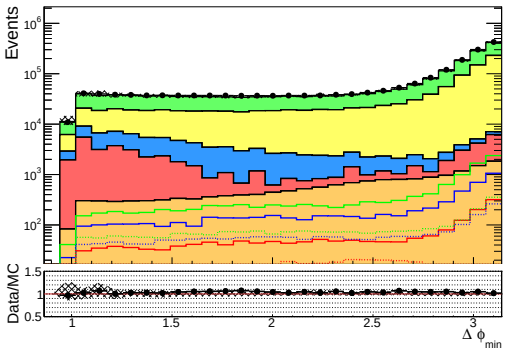

(b) $\Delta\phi_{min}$

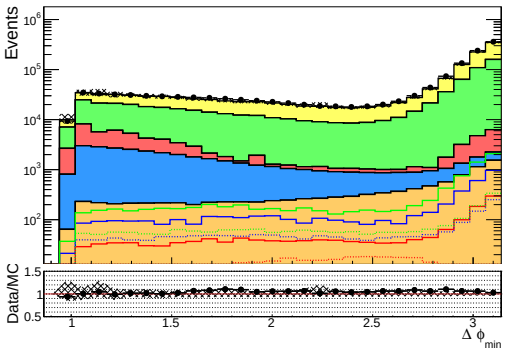

(c) Jet p_T

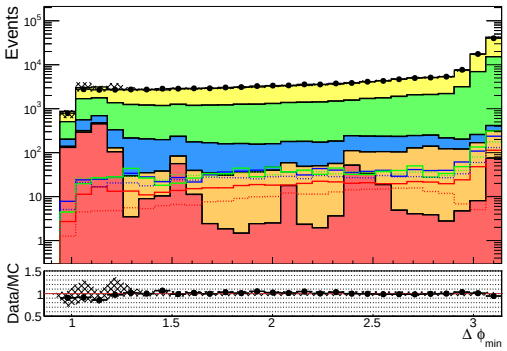

(d) Base $\not{E}_T$

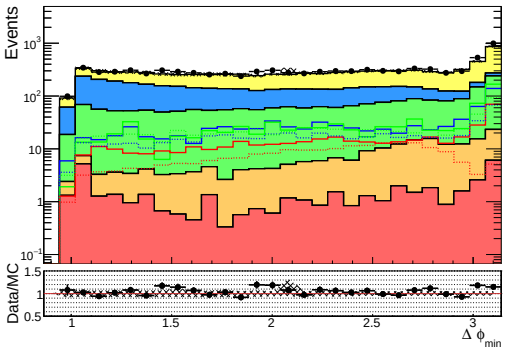

(e) b-jet Multiplicity

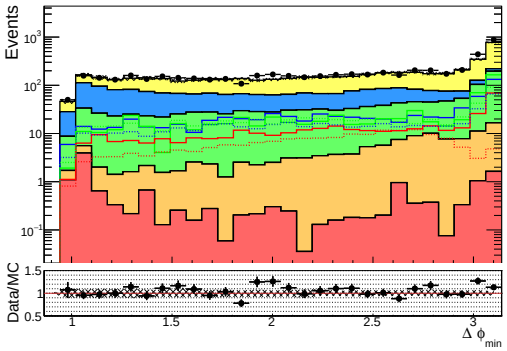

(f) b-jet p_T

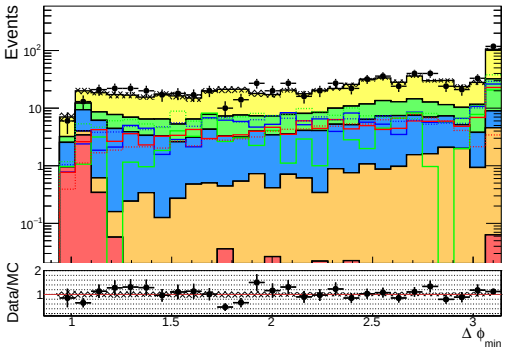

(g) SR300

Figure 5.12: $\Delta\phi_{min}$ distributions after each of the selection cuts.

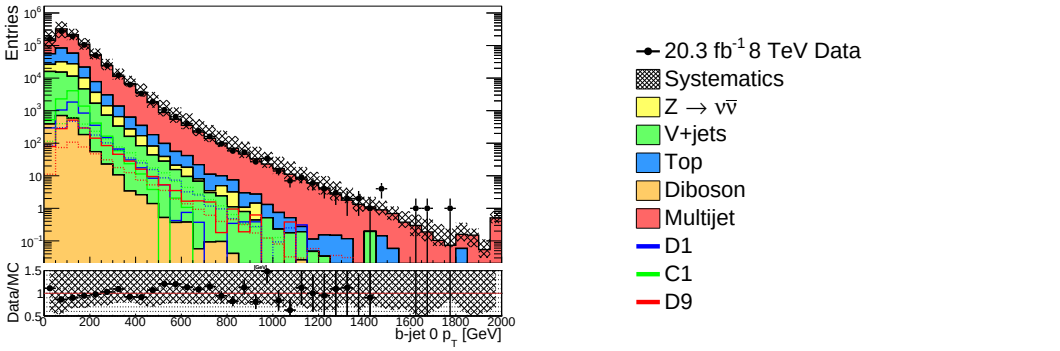

(a) Preselection

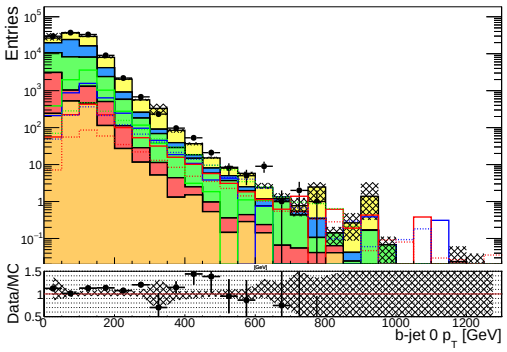

(b) $\Delta\phi_{min}$

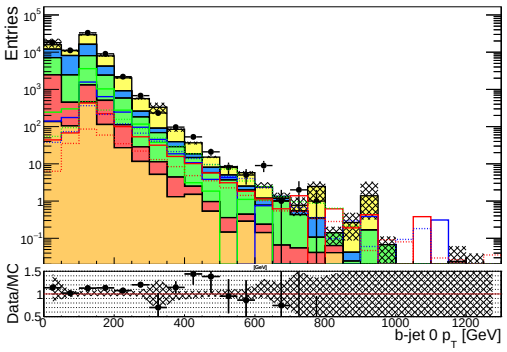

(c) Jet p_T

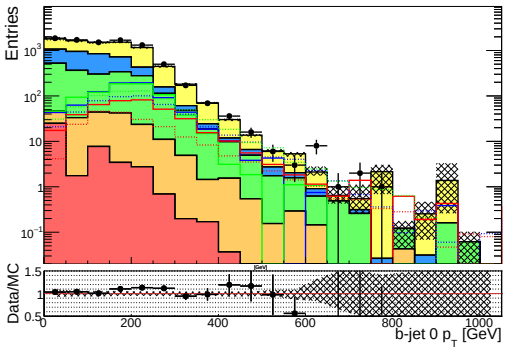

(d) Base $\not{E}_T$

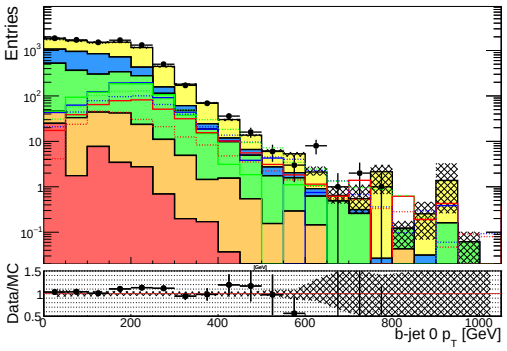

(e) b-jet Multiplicity

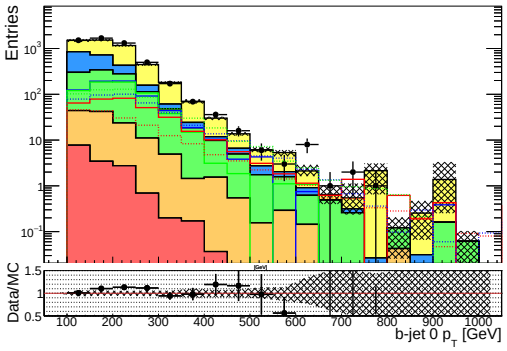

(f) b-jet p_T

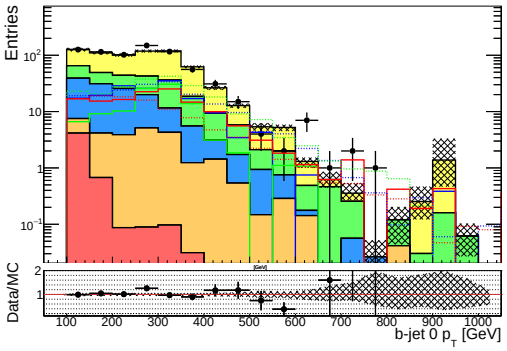

(g) SR300

Figure 5.13: $p_T(b_0)$ distributions after each of the selection cuts.

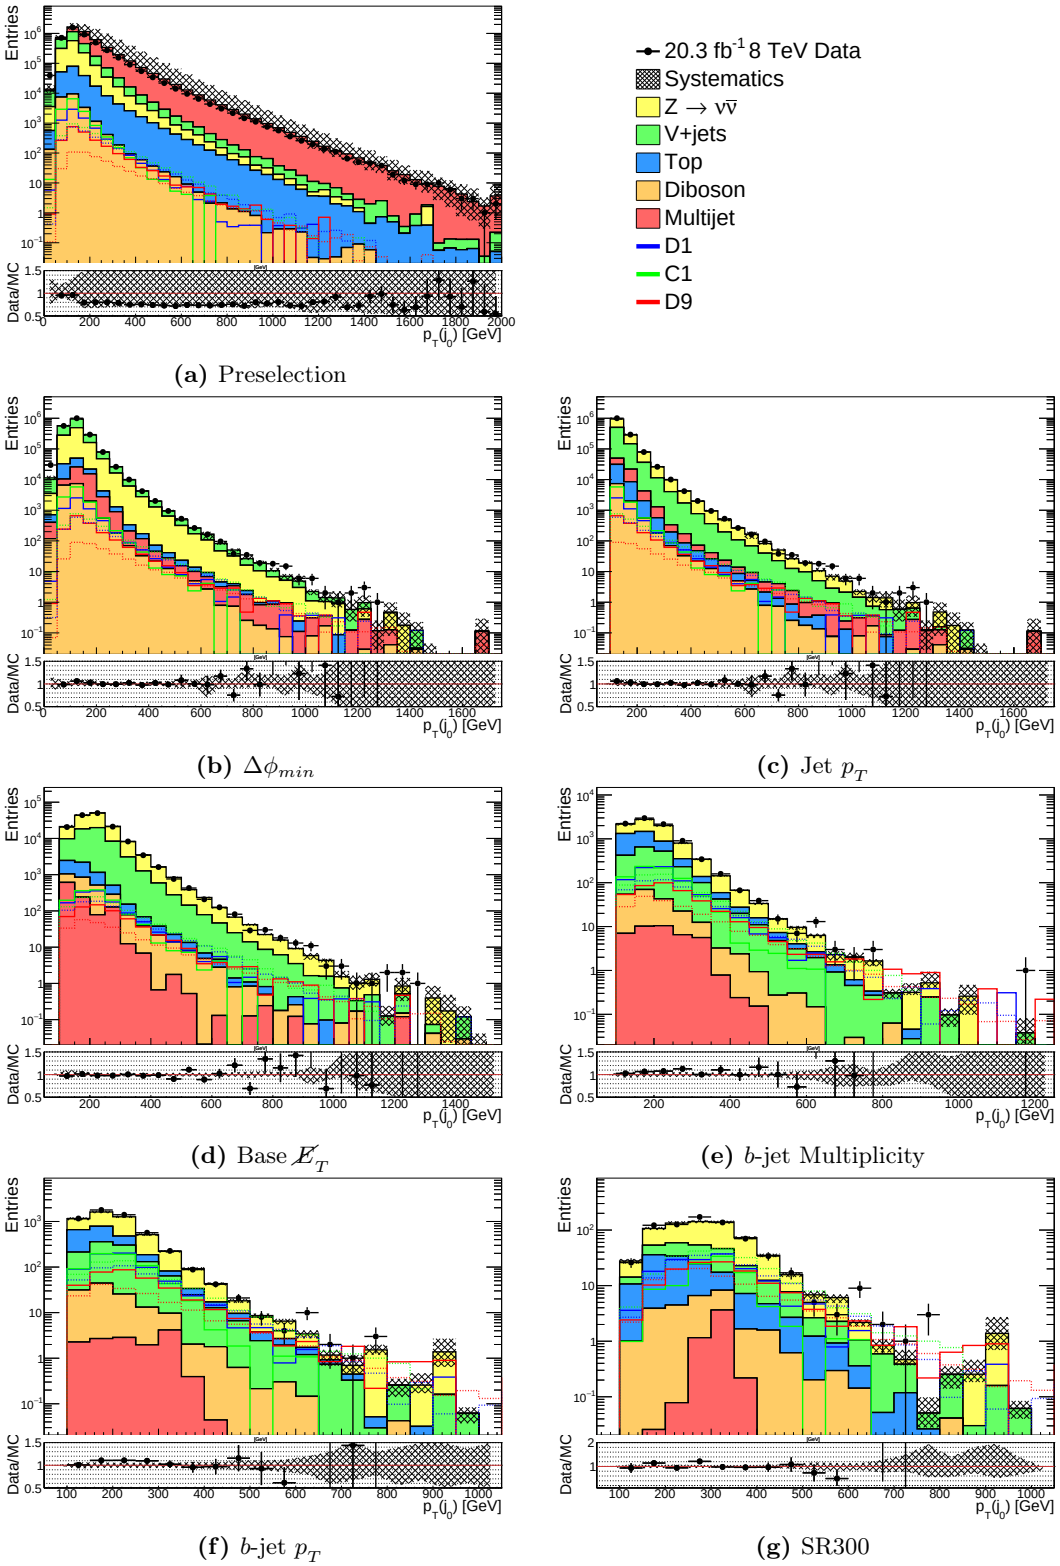

(a) Preselection

(b) $\Delta\phi_{min}$

(c) Jet p_T

(d) Base $\not{E}_T$

(e) b-jet Multiplicity

(f) b-jet p_T

(g) SR300

Figure 5.14: $p_T(j_0)$ distributions after each of the selection cuts.

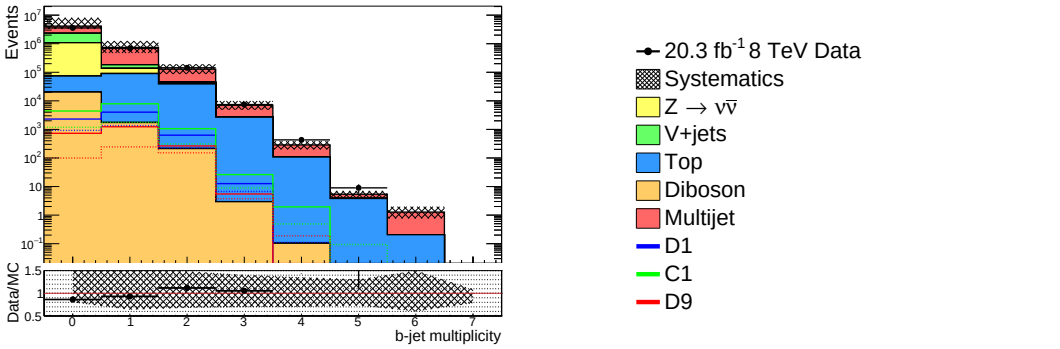

(a) Preselection

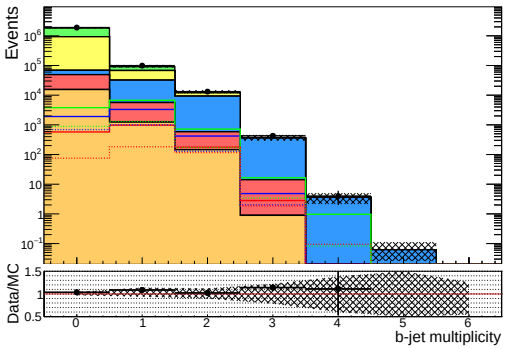

(b) $\Delta\phi_{min}$

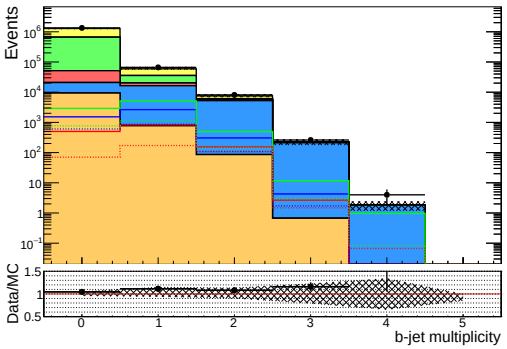

(c) Jet p_T

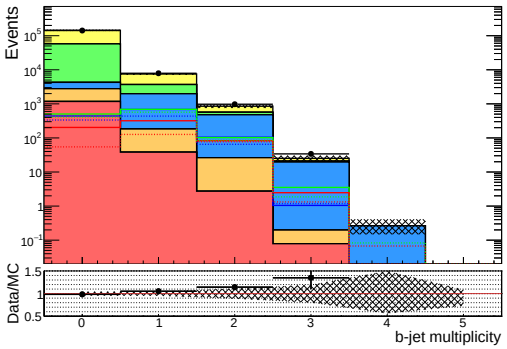

(d) Base $\not{E}_T$

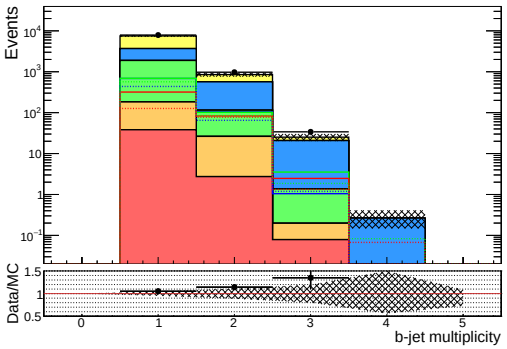

(e) b-jet Multiplicity

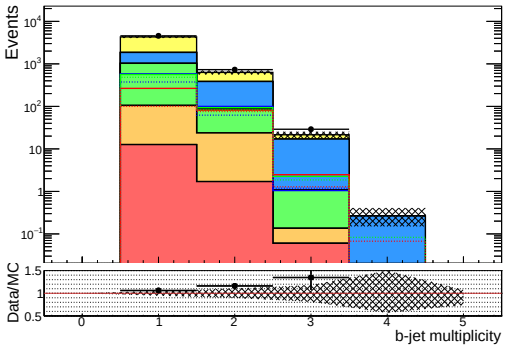

(f) b-jet p_T

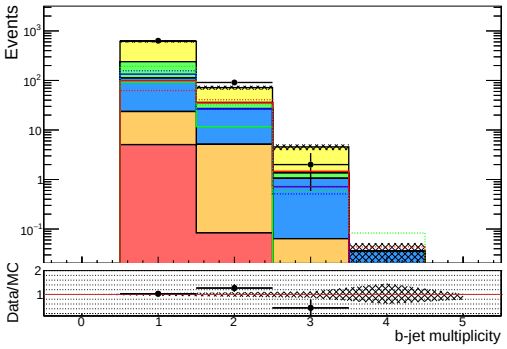

(g) SR300

Figure 5.15: n_b distributions after each of the selection cuts.

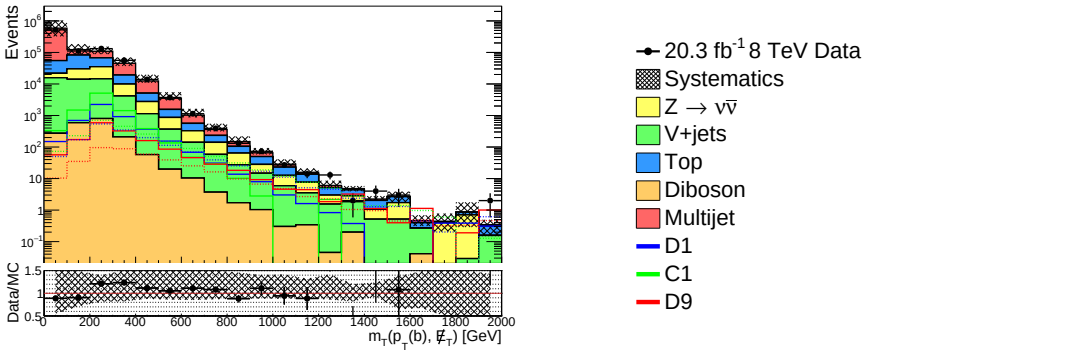

(a) Preselection

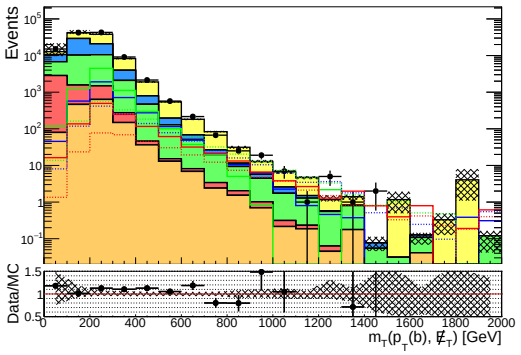

(b) $\Delta\phi_{min}$

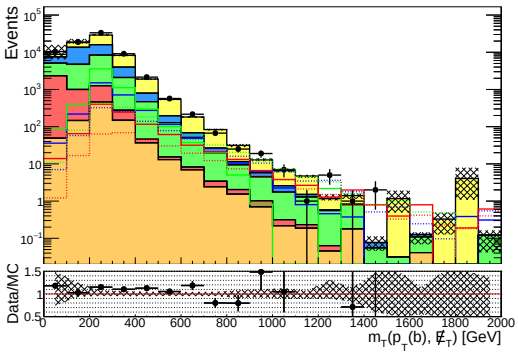

(c) Jet p_T

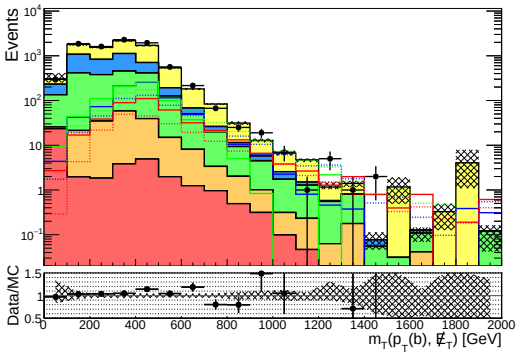

(d) Base $\not{E}_T$

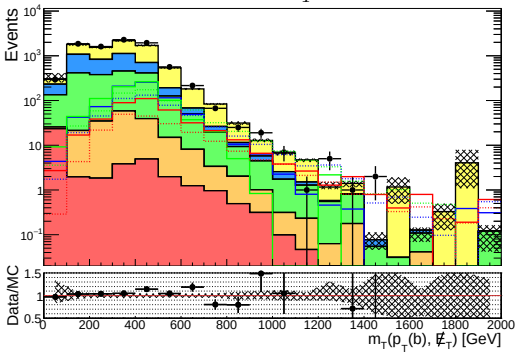

(e) b-jet Multiplicity

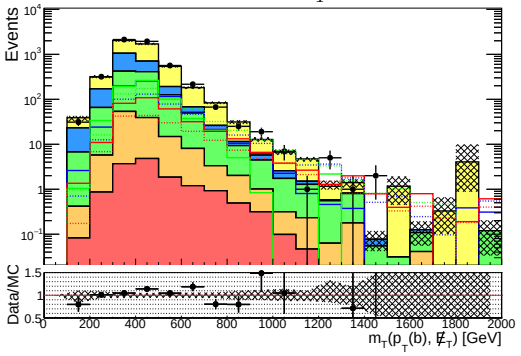

(f) b-jet p_T

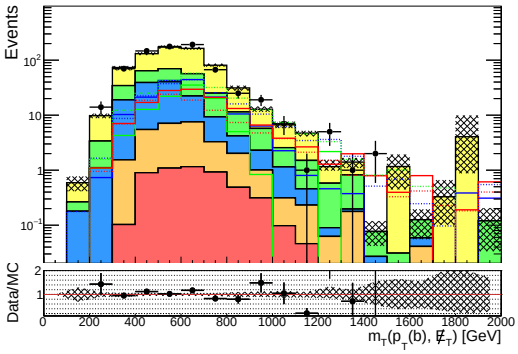

(g) SR300

Figure 5.16: m_T^b distributions after each of the selection cuts.

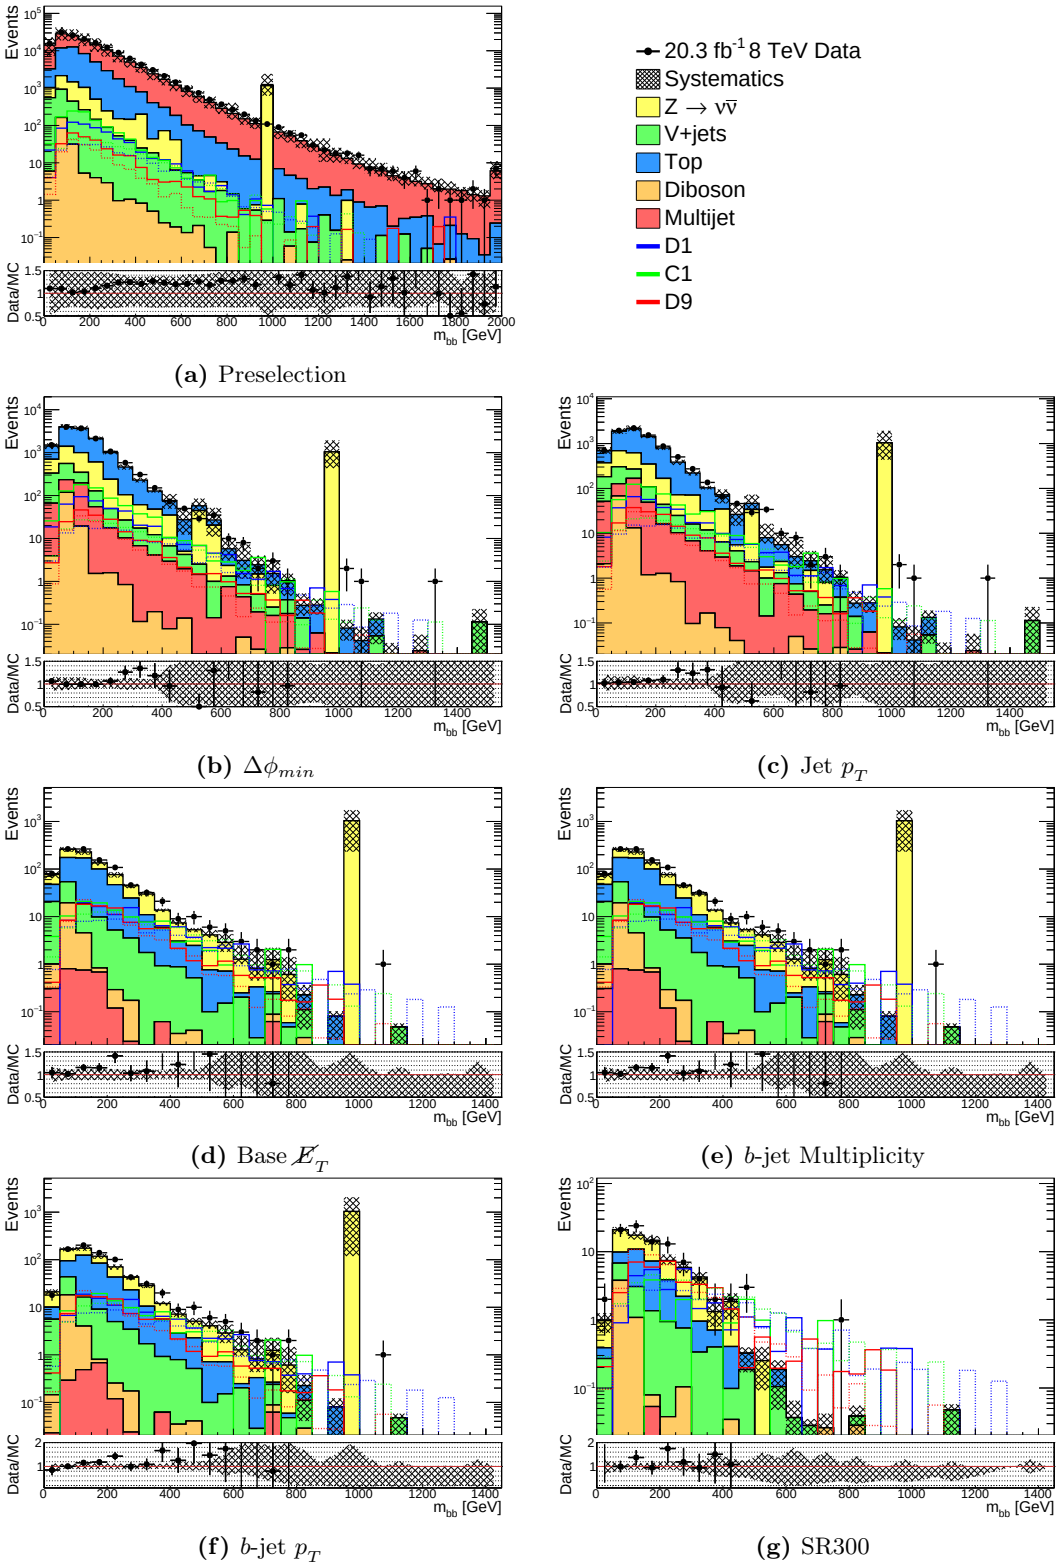

(a) Preselection

(b) $\Delta\phi_{min}$

(c) Jet p_T

(d) Base $\not{E}_T$

(e) b-jet Multiplicity

(f) b-jet p_T

(g) SR300

Figure 5.17: m_{bb} distributions after each of the selection cuts. The spike in the $Z \to \nu\nu$ estimate is actually due to a statistical downward fluctuation in the γ + jets sample. After dividing by this sample, the estimate receives a very large upward fluctuation. However, the statistical uncertainties on this bin are very large, and no statistically significant deviations occur.

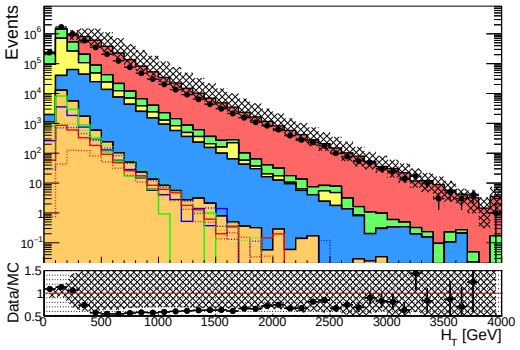

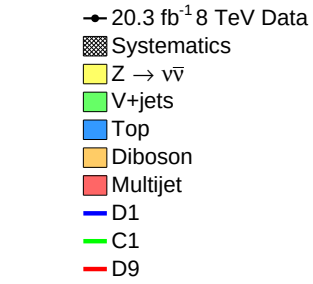

(a) Preselection

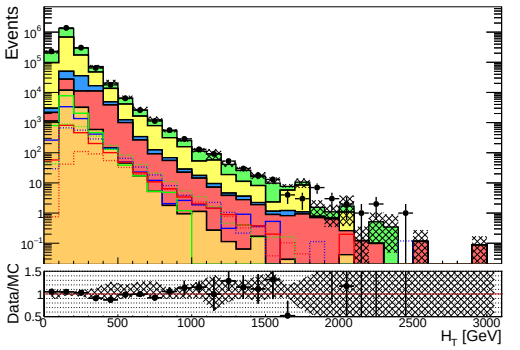

(b) $\Delta\phi_{min}$

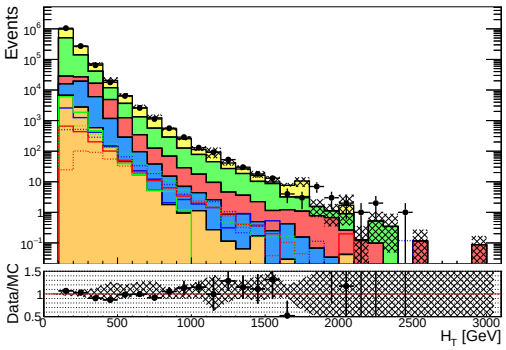

(c) Jet p_T

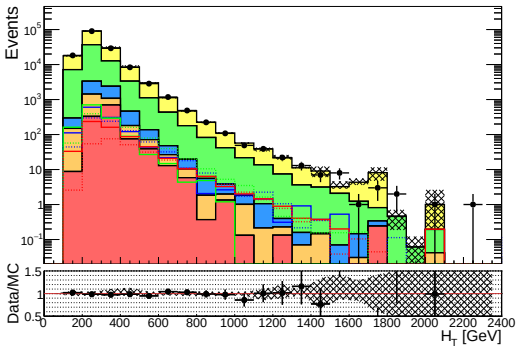

(d) Base $\not{E}_T$

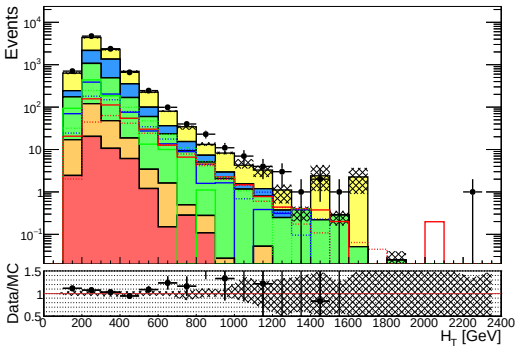

(e) b-jet Multiplicity

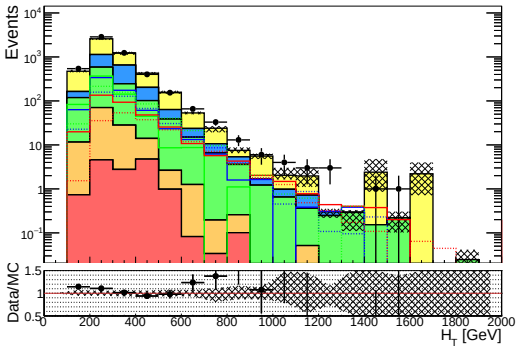

(f) b-jet p_T

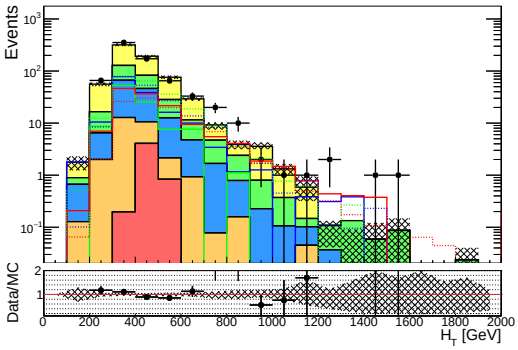

(g) SR300

Figure 5.18: H_T distributions after each of the selection cuts.

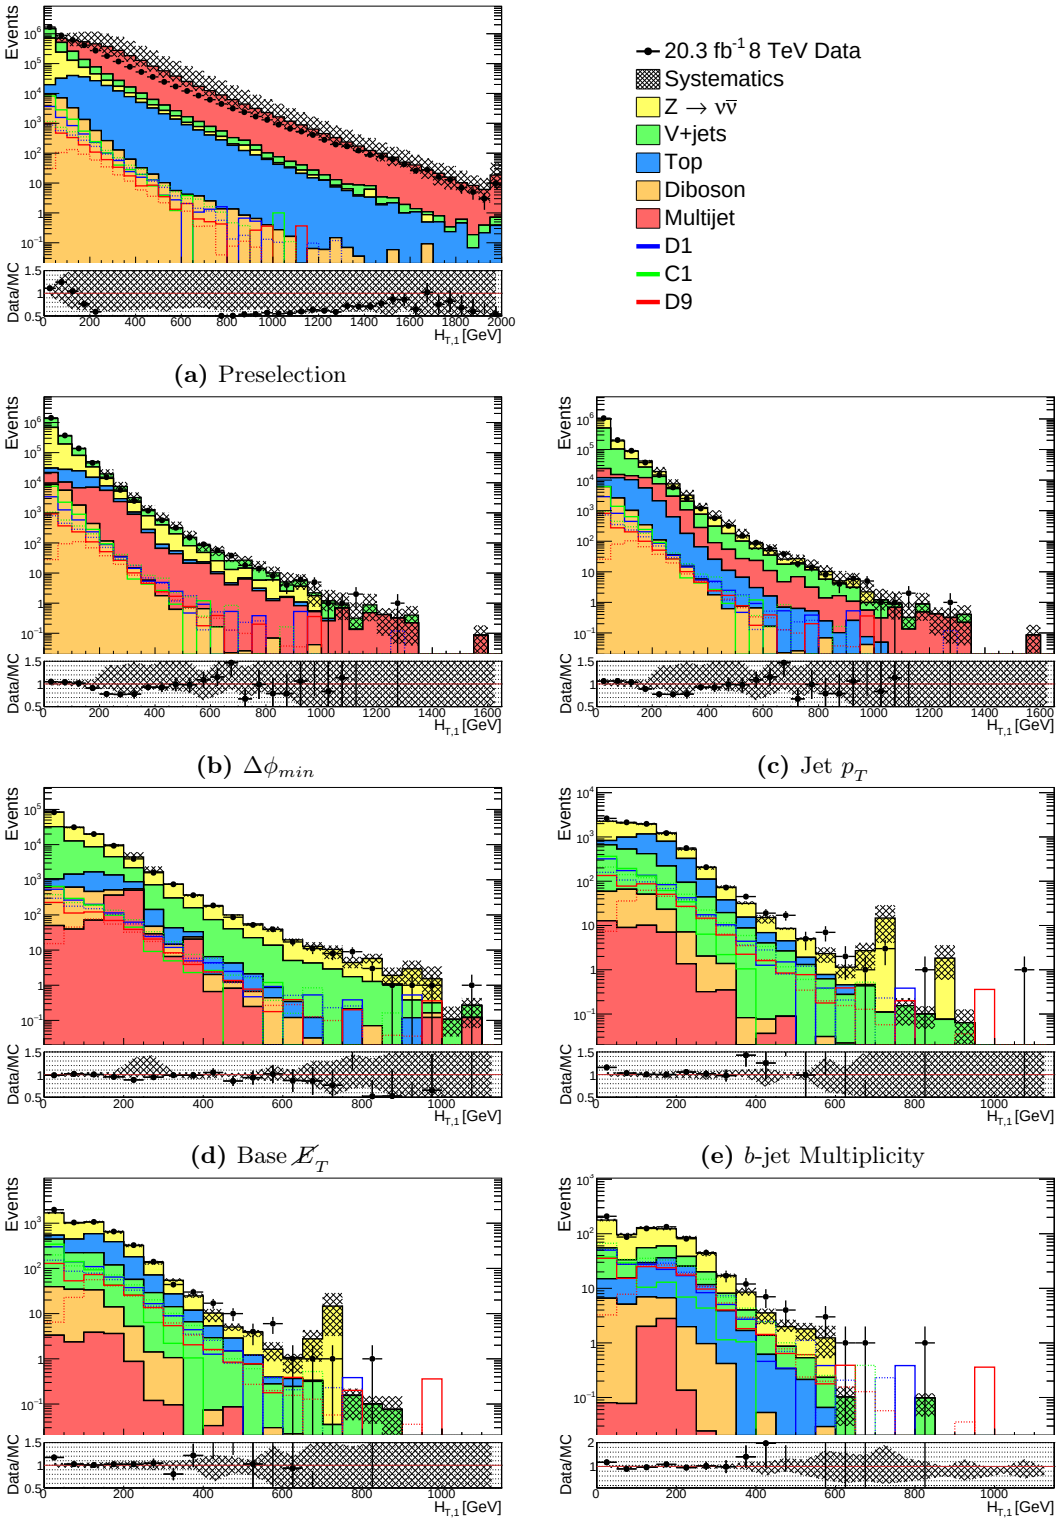

(a) Preselection

(b) $\Delta\phi_{min}$

(c) Jet p_T

(d) Base $\not{E}_T$

(e) b-jet Multiplicity

(f) b-jet p_T

(g) SR300

Figure 5.19: H_T^1 distributions after each of the selection cuts.

Cut	Multijet	VV	Top	V + jets	$Z \to \nu\nu$	Total SM	Data
Preselection	8,490,000 ± 13,000,000	28,000 ± 2,400	240,000 ± 17,000	1,760,000 ± 30,000	1,340,000 ± 93,000	11,900,000 ± 13,000,000	13,920,151 ± 3,731
$\not{E}_T > 100$ GeV	2,470,000 ± 4,500,000	22,300 ± 2,300	184,000 ± 16,000	1,290,000 ± 81,000	1,060,000 ± 66,000	5,030,000 ± 4,600,000	4,409,360 ± 2,100
Jet Multiplicity	1,510,000 ± 2,700,000	21,500 ± 2,100	106,000 ± 7,700	1,240,000 ± 73,000	1,040,000 ± 64,000	3,920,000 ± 2,900,000	3,990,230 ± 1,998
$\Delta\phi_{min}$	38,600 ± 82,000	16,900 ± 1,600	55,000 ± 4,100	927,000 ± 50,000	903,000 ± 58,000	1,940,000 ± 180,000	2,012,811 ± 1,419
Jet p_T	34,800 ± 68,000	10,200 ± 1,100	32,500 ± 2,400	626,000 ± 45,000	650,000 ± 31,000	1,350,000 ± 140,000	1,418,943 ± 1,191
Base $\not{E}_T$	1,220 ± 2,900	1,780 ± 210	3,760 ± 370	55,000 ± 4,800	90,900 ± 4,200	153,000 ± 12,000	150,804 ± 388
b-jet Multiplicity	41.0 ± 230	168 ± 25	2,270 ± 260	1,800 ± 210	4,110 ± 220	8,390 ± 880	8,906 ± 94
b-jet p_T	14.4 ± 49	115 ± 18	1,140 ± 130	988 ± 120	2,700 ± 150	4,960 ± 450	5,328 ± 73
SR300	5.14 ± 2.7	23.6 ± 4.3	111 ± 15	137 ± 19	414 ± 26	691 ± 62	727 ± 27
SR350	4.60 ± 0.68	11.7 ± 2.2	41.9 ± 6.7	62.6 ± 8.1	190 ± 14	311 ± 28	290 ± 17
SR400	4.26 ± 0.38	5.57 ± 1.3	16.2 ± 2.9	28.6 ± 4.0	85.0 ± 7.5	140 ± 14	135 ± 12
SR450	3.99 ± 0.44	3.13 ± 0.89	7.60 ± 1.7	14.6 ± 2.1	37.0 ± 4.4	66.3 ± 7.7	75 ± 9
SR500	3.74 ± 0.39	1.16 ± 0.36	3.51 ± 0.90	8.22 ± 1.7	22.3 ± 3.2	39.0 ± 5.0	39 ± 6

Table 5.33: Yields in the SR. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

5.6 Background Estimates

In addition to the SRs defined in the previous section, various CRs and VRs are necessary. These serve to constrain free parameters of the background estimates, and also provide a signal-free region with which to validate the estimates while keeping the analysis blinded. There are 6 CRs used by this analysis to control dominant background processes, defined in Table 5.34, and 2 VRs to validate them, defined in Table 5.35. While the kinematics of each process are well modeled, there are large uncertainties on their overall normalizations. Leaving these as free parameters that are fit in CRs can remove a significant source of systematic uncertainty. Below we will briefly discuss the strategies used to estimate each background and validate them using a variety of methods. The expected and observed yields in each of these regions are listed after every cut in Tables 5.38-5.44.

The major SM backgrounds, with the exception of $Z \to \nu\nu$, can only contribute to the SR through the misidentification of a leptonic $W^{\pm}$ decay. Electrons and muons can escape detection, and tau leptons can decay hadronically and be misidentified as b-jets. Focusing on the former, we can inspect these backgrounds in a region where identified muons are removed from the event. Recalculating the $\not{E}_T$ without the muon contributions, denoted $\not{E}_T^{\mu}$, results in a signal-free region that is *very* close to the SR for such processes. The regions CR**wjets**, CR**top**, CR**zmm**, and VR**lep** all require at least one well measured muon, and unless explicitly stated, the $\vec{\not{p}}_T$ of these events will be replaced with $\vec{\not{p}}_T^{\,\mu}$. A similar procedure is used in CR**gamma**, where we remove the contribution of a high p_T photon to the $\not{E}_T$ of the event. This new quantity is labeled $\vec{\not{p}}_T^{\,\gamma}$, and is used instead of $\vec{\not{p}}_T$ in this region.

Before we begin comparing the background estimates to data, they need to be fit in their respective CRs. This is done using the HistFitter program [122] to perform a simultaneous fit of all the SM backgrounds in each CR. HistFitter adjusts the normalization of each background to maximize a joint likelihood function that takes into account all the various uncertainties. The results of this fit are given in Table 5.36, for both the nominal and auxiliary samples. In all of the plots shown from this point forward, these SFs will

Cut	CR**znn**	CR**wjets**	CR**top**	CR**zmm**	CR**gamma**	CR**multijet**
Preselection	0 lepton	1 muon	1 muon	2 muon	1 photon	0 lepton
Orthogonality	$n_b = 0$	$n_b = 0$	$n_b = 2$	$n_\mu = 2$	$n_\gamma = 1$	$\Delta\phi_{min} < 1$
Jet Multiplicity	$0 < n_j < 5$	$0 < n_j < 5$	$0 < n_j < 5$	$0 < n_j < 5$	$0 < n_j < 5$	$0 < n_j < 5$
$\Delta\phi_{min}$	$\Delta\phi_{min} > 1$	$\Delta\phi^{\mu}_{min} > 1$	$\Delta\phi^{\mu}_{min} > 1$	$\Delta\phi^{\mu}_{min} > 1$	$\Delta\phi^{\gamma}_{min} > 1$	
Jet p_T (GeV)	$p_T(j_0) > 100$	$p_T(j_0) > 100$	$p_T(j_0) > 100$	$p_T(j_0) > 100$	$p_T(j_0) > 100$	$p_T(j_0) > 100$
Base $\not{E}_T$ (GeV)	$\not{E}_T > 200$	$\not{E}^{\mu}_T > 200$	$\not{E}^{\mu}_T > 200$	$\not{E}^{\mu}_T > 200$	$\not{E}^{\gamma}_T > 200$	$\not{E}_T > 200$
b-jet Multiplicity				$n_b > 0$	$n_b > 0$	$n_b > 0$
b-jet p_T (GeV)			$p_T(b_0) > 100$	$p_T(b_0) > 100$	$p_T(b_0) > 100$	$p_T(b_0) > 100$
Final $\not{E}_T$ (GeV)	$\not{E}^{\mu}_T > 300$	$\not{E}^{\mu}_T > 300$			$\not{E}^{\gamma}_T > 300$	

Cutflow 5.34: Control region selection cuts.

Cut	VR**lep**	VR**met**
Preselection	1 muon	0 lepton
Orthogonality	$n_\mu = 1$	$\not{E}_T < 300$
Jet Multiplicity	$0 < n_j < 5$	$0 < n_j < 5$
$\Delta\phi_{min}$	$\Delta\phi^{\mu}_{min} > 1$	$\Delta\phi_{min} > 1$
Jet p_T (GeV)	$p_T(j_0) > 100$	$p_T(j_0) > 100$
Base $\not{E}_T$ (GeV)	$\not{E}^{\mu}_T > 200$	$\not{E}_T > 200$
b-jet Multiplicity	$n_b = 1$	$n_b > 1$
b-jet p_T (GeV)	$p_T(b_0) > 100$	$p_T(b_0) > 100$
Final $\not{E}_T$ (GeV)	$\not{E}^{\mu}_T > 300$	

Cutflow 5.35: Validation region selection cuts.

be automatically applied to the appropriate backgrounds. The single top, diboson, and signal processes are normalized to their theoretical cross-section, and are assigned a larger systematic uncertainty, discussed further in Section 5.7.

Background	Scale Factor
$Z^0 \to \mu^{\pm}\mu^{\mp}$	0.94 ± 0.06
$Z^0 \to \nu\bar{\nu}$	0.88 ± 0.02
$W^{\pm} \to \ell^{\pm}\nu$	0.81 ± 0.01
$\gamma +$ jets	0.94 ± 0.03
Top	1.13 ± 0.08
Multijet	1.20 ± 0.04
$Z^0 \to \nu\bar{\nu}$ (Sherpa)	0.80 ± 0.02
Multijet (Pythia)	0.79 ± 0.11
$Z^0 \to \nu\bar{\nu}$ (Alpgen)	0.70 ± 0.07
$W^{\pm} \to \ell^{\pm}\nu$ (Alpgen)	0.84 ± 0.03
$Z^0 \to \mu^{\pm}\mu^{\mp}$ (Alpgen)	1.30 ± 0.17
$\gamma +$ jets (Alpgen)	1.59 ± 0.08

Table 5.36: The results of a background-only fit in HistFitter, taking both statistical and systematic uncertainties into account in each CR.

5.6.1 $Z^0 \to \nu\bar{\nu}$

In the mono-b search, $Z \to \nu\nu$ is the largest background process, which is why data-driven techniques were chosen to estimate it more precisely. The details of how this estimate was created are discussed in Appendix B, which resulted in a pseudo-data sample simulating the $Z \to \nu\nu$ process. To study this background, a region CR**znn** is defined in which the b-jet requirements of the nominal selection are reversed by adding a veto on b-jets to the preselection cuts, providing orthogonality to the SR, and the nominal b-jet cuts are removed. This leaves a region with very little signal and a $Z \to \nu\nu$ purity of about 65%. Figs. 5.20–5.23 show relevant kinematic variables plotted at difference stages in the CR**znn** selection, and Table 5.37 lists the total yields after each cut.

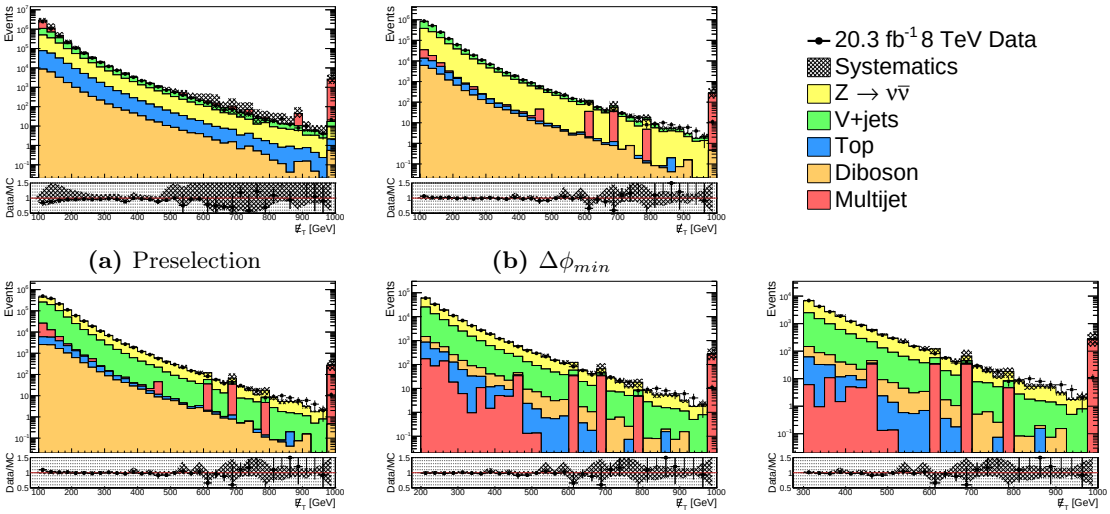

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.20: $\not{E}_T$ distributions after each of the CR**znn** cuts.

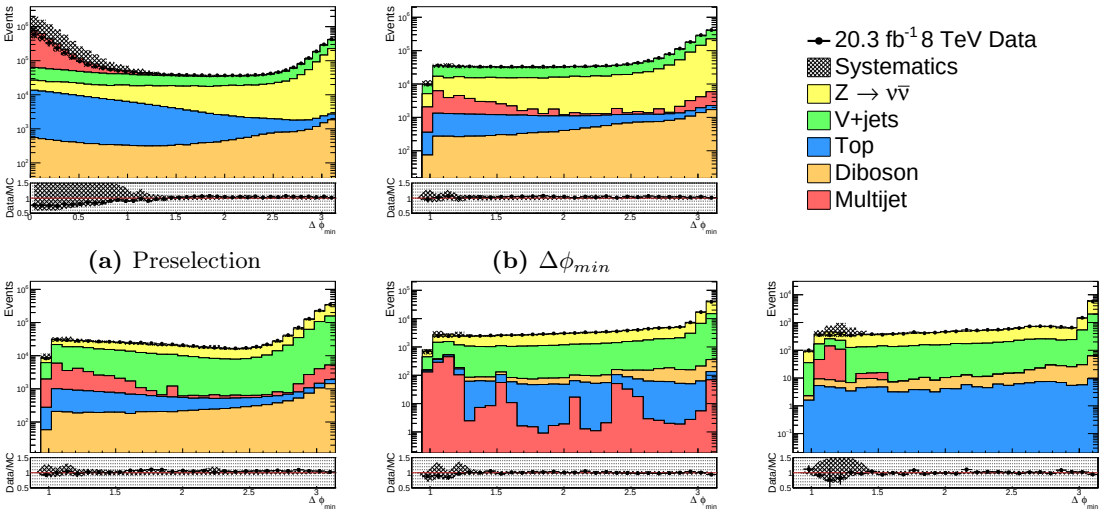

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.21: $\Delta\phi_{min}$ distributions after each of the CR**znn** cuts.

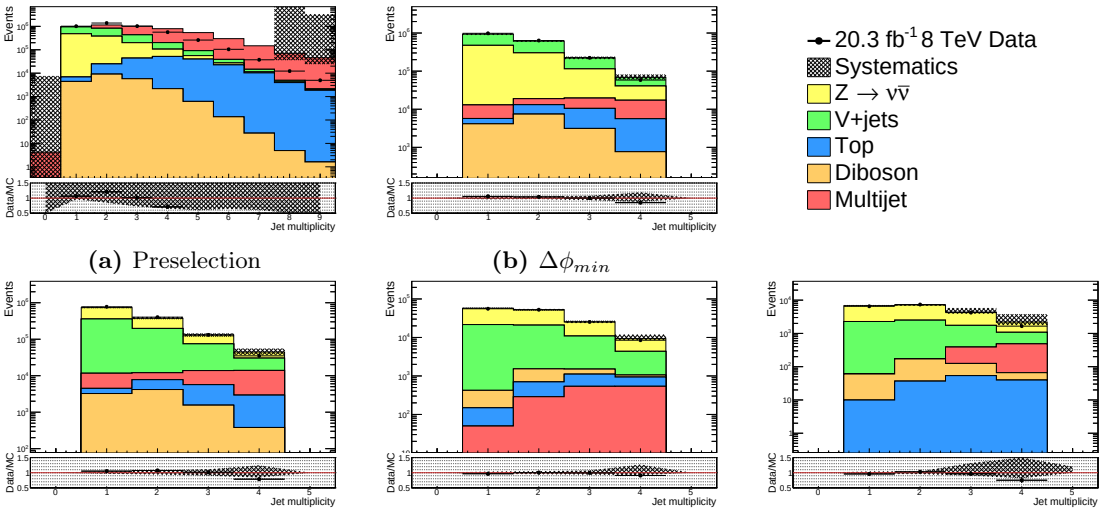

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.22: n_j distributions after each of the CR**znn** cuts.

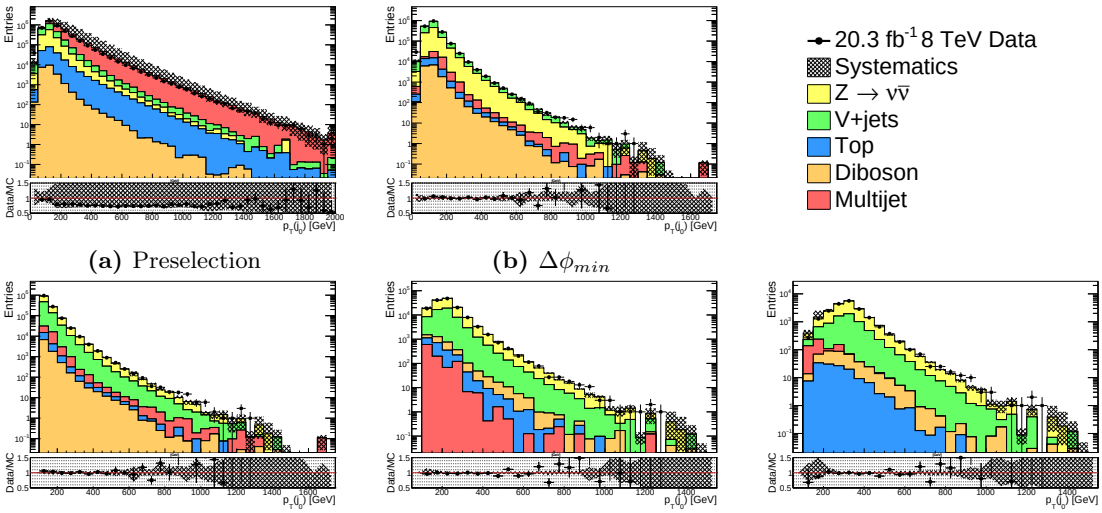

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.23: $p_T(j_0)$ distributions after each of the CR**znn** cuts.

Cut	Multijet	VV	Top	V + jets	$Z \to \nu\nu$	Total SM	Data
Preselection	8,490,000 ± 13,000,000	28,000 ± 2,400	240,000 ± 17,000	1,760,000 ± 30,000	1,340,000 ± 93,000	11,900,000 ± 13,000,000	13,920,151 ± 3,731
$\not{E}_T > 100$ GeV	2,470,000 ± 4,500,000	22,300 ± 2,300	184,000 ± 16,000	1,290,000 ± 81,000	1,060,000 ± 66,000	5,030,000 ± 4,600,000	4,409,360 ± 2,100
b-jet Multiplicity	1,810,000 ± 4,000,000	20,300 ± 2,100	53,700 ± 7,800	1,240,000 ± 78,000	1,010,000 ± 62,000	4,140,000 ± 4,100,000	3,568,046 ± 1,889
Jet Multiplicity	1,100,000 ± 2,400,000	19,600 ± 1,900	35,600 ± 4,600	1,200,000 ± 71,000	993,000 ± 61,000	3,340,000 ± 2,600,000	3,322,544 ± 1,823
$\Delta\phi_{min}$	33,700 ± 69,000	15,500 ± 1,500	19,500 ± 2,400	903,000 ± 49,000	864,000 ± 55,000	1,840,000 ± 160,000	1,900,341 ± 1,379
Jet p_T	30,600 ± 59,000	9,360 ± 1,000	11,500 ± 1,400	610,000 ± 43,000	624,000 ± 29,000	1,290,000 ± 130,000	1,344,237 ± 1,159
Base $\not{E}_T$	1,170 ± 2,700	1,610 ± 190	1,490 ± 210	53,200 ± 4,700	86,800 ± 4,000	144,000 ± 12,000	141,898 ± 377
SR300	530 ± 2,800	282 ± 34	140 ± 20	6,490 ± 490	12,900 ± 590	20,400 ± 3,900	19,780 ± 141

Table 5.37: Yields for CRznn region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature

Cut	VV	Top	$Z \to \ell\ell$	$W \to \ell\nu$	Total SM	Data
Preselection	80,300 ± 7,200	507,000 ± 35,000	1,290,000 ± 90,000	13,600,000 ± 560,000	15,500,000 ± 690,000	22,800,383 ± 4,775
$\not{E}_T > 100$ GeV	10,400 ± 1,200	45,400 ± 7,000	58,400 ± 8,000	885,000 ± 84,000	999,000 ± 99,000	1,134,846 ± 1,065
b-jet Multiplicity	75,400 ± 6,700	146,000 ± 20,000	1,260,000 ± 87,000	13,300,000 ± 560,000	14,800,000 ± 680,000	21,651,161 ± 4,653
Jet Multiplicity	10,200 ± 1,200	35,100 ± 4,800	56,700 ± 7,700	867,000 ± 79,000	969,000 ± 92,000	1,107,677 ± 1,052
$\Delta\phi_{min}$	9,090 ± 1,000	24,600 ± 3,200	46,000 ± 6,200	771,000 ± 69,000	850,000 ± 78,000	969,747 ± 985
Jet p_T	4,840 ± 550	12,900 ± 1,700	25,500 ± 3,300	457,000 ± 35,000	500,000 ± 40,000	565,226 ± 752
Base $\not{E}_T$	922 ± 110	2,640 ± 380	2,250 ± 350	58,900 ± 5,000	64,700 ± 5,800	71,502 ± 267
SR300	183 ± 22	358 ± 55	288 ± 49	9,940 ± 790	10,800 ± 900	10,782 ± 104

Table 5.38: Yields for CRwjets region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

Cut	VV	Single t	$t\bar{t}$	V + jets	Total SM	Data
Preselection	80,300 ± 7,200	96,700 ± 12,000	418,000 ± 29,000	15,200,000 ± 620,000	15,800,000 ± 670,000	22,800,383 ± 4,775
$\not{E}_T > 100\,\text{GeV}$	82.4 ± 14	2,000 ± 380	33,500 ± 5,200	1,830 ± 290	37,400 ± 5,800	34,298 ± 185
b-jet Multiplicity	391 ± 78	8,230 ± 1,600	106,000 ± 15,000	7,880 ± 1,400	123,000 ± 18,000	123,538 ± 351
Jet Multiplicity	78.0 ± 13	1,530 ± 290	20,000 ± 3,000	1,530 ± 240	23,200 ± 3,500	22,127 ± 149
$\Delta\phi_{min}$	66.5 ± 11	914 ± 170	11,600 ± 1,800	1,100 ± 170	13,700 ± 2,100	13,440 ± 116
Jet p_T	31.5 ± 5.6	490 ± 92	6,000 ± 930	551 ± 87	7,080 ± 1,100	6,829 ± 83
Base $\not{E}_T$	9.15 ± 1.6	94.0 ± 19	756 ± 130	135 ± 21	994 ± 170	986 ± 31
b-jet p_T	8.61 ± 1.5	73.7 ± 14	496 ± 86	98.3 ± 15	677 ± 120	682 ± 26

Table 5.39: Yields for CRtop region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature

Cut	VV	Top	$Z \to \ell\ell$	$W \to \ell\nu$	Total SM	Data
Preselection	5,510 ± 510	7,110 ± 500	1,160,000 ± 99,000	3.86 ± 4.8	1,170,000 ± 100,000	1,285,016 ± 1,134
$\not{E}_T > 100\,\text{GeV}$	809 ± 92	2,100 ± 190	85,300 ± 9,200	0.201 ± 1.2	88,200 ± 9,500	90,385 ± 301
Jet Multiplicity	795 ± 88	1,890 ± 160	83,500 ± 8,800	0.201 ± 0.20	86,200 ± 9,000	88,574 ± 298
$\Delta\phi_{min}$	714 ± 76	1,490 ± 130	74,500 ± 7,700	0.0564 ± 0.23	76,700 ± 7,900	78,437 ± 280
Jet p_T	405 ± 44	693 ± 64	46,500 ± 4,400	0.0564 ± 0.23	47,600 ± 4,500	48,246 ± 220
Base $\not{E}_T$	85.4 ± 11	68.4 ± 7.1	6,680 ± 700	0.0564 ± 0.057	6,840 ± 710	6,922 ± 83
b-jet Multiplicity	8.28 ± 1.7	33.5 ± 4.0	286 ± 33	0. ± 0.	327 ± 38	328 ± 18
b-jet p_T	6.16 ± 1.5	11.8 ± 1.9	180 ± 23	0. ± 0.	198 ± 26	209 ± 14

Table 5.40: Yields for CRzmm region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

Cut	γ + jets	VV	Top	V + jets	Total SM	Data
Preselection	2,710,000 ± 110,000	4,100 ± 370	1,350 ± 100	11,700 ± 430	2,720,000 ± 110,000	3,131,293 ± 1,770
$\not{E}_T > 100$ GeV	2,650,000 ± 110,000	2,970 ± 270	1,250 ± 94	11,100 ± 420	2,670,000 ± 110,000	3,065,591 ± 1,751
Jet Multiplicity	2,560,000 ± 110,000	2,830 ± 260	827 ± 69	10,700 ± 400	2,580,000 ± 110,000	2,979,623 ± 1,726
$\Delta\phi_{min}$	2,170,000 ± 94,000	2,280 ± 210	542 ± 47	9,390 ± 360	2,180,000 ± 94,000	2,554,231 ± 1,598
Jet p_T	1,700,000 ± 68,000	1,700 ± 160	395 ± 30	7,880 ± 280	1,710,000 ± 68,000	1,951,725 ± 1,397
Base $\not{E}_T$	404,000 ± 35,000	690 ± 76	186 ± 15	2,600 ± 190	408,000 ± 36,000	432,954 ± 658
b-jet Multiplicity	17,800 ± 2,000	70.2 ± 11	112 ± 11	105 ± 13	18,000 ± 2,000	18,911 ± 138
b-jet p_T	10,900 ± 1,300	41.7 ± 6.6	61.5 ± 6.2	66.6 ± 9.2	11,100 ± 1,300	11,928 ± 109
SR300	1,660 ± 190	12.3 ± 2.8	12.7 ± 1.8	13.2 ± 2.1	1,700 ± 200	1,701 ± 41

Table 5.41: Yields for CRgamma region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature

Cut	Multijet	VV	Top	V + jets	$Z \to \nu\nu$	Total SM	Data
Preselection	8,490,000 ± 13,000,000	28,000 ± 2,400	240,000 ± 17,000	1,760,000 ± 30,000	1,340,000 ± 93,000	11,900,000 ± 13,000,000	13,920,151 ± 3,731
$\not{E}_T > 100$ GeV	2,470,000 ± 4,500,000	22,300 ± 2,300	184,000 ± 16,000	1,290,000 ± 81,000	1,060,000 ± 66,000	5,030,000 ± 4,600,000	4,409,360 ± 2,100
Jet Multiplicity	1,510,000 ± 2,700,000	21,500 ± 2,100	106,000 ± 7,700	1,240,000 ± 73,000	1,040,000 ± 64,000	3,920,000 ± 2,900,000	3,990,230 ± 1,998
$\Delta\phi_{min}$	1,420,000 ± 2,700,000	3,930 ± 460	43,700 ± 3,100	288,000 ± 21,000	129,000 ± 6,400	1,880,000 ± 2,700,000	1,858,402 ± 1,363
Jet p_T	1,450,000 ± 2,700,000	14,200 ± 1,600	76,200 ± 5,500	913,000 ± 65,000	778,000 ± 37,000	3,240,000 ± 2,800,000	3,277,345 ± 1,810
Base $\not{E}_T$	18,100 ± 29,000	34.4 ± 78	2,150 ± 430	887 ± 2,700	16,100 ± 1,500	37,300 ± 34,000	7,920 ± 89
b-jet Multiplicity	5,740 ± 7,100	56.8 ± 10	3,630 ± 320	2,010 ± 280	1,160 ± 66	12,600 ± 7,700	12,150 ± 110
b-jet p_T	4,240 ± 4,000	34.4 ± 6.3	2,150 ± 210	887 ± 120	632 ± 39	7,950 ± 4,400	7,920 ± 89

Table 5.42: Yields for CRmultijet region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

In addition to the nominal data-driven estimate, four alternatives were used for validation. The Sherpa and Alpgen MC samples, the $Z \to \mu\mu$ based estimate, and an estimate using the Alpgen samples to calculate the *transfer factor* (TF) were all used. The Sherpa MC samples are the same ones used to derive the nominal TF, so any comparison to them is more of a closure test. The other three estimates have very low statistics, and are not reliable near the SR. However, in earlier selection stages the agreement between all four estimates is very good, as shown in Fig. 5.24.

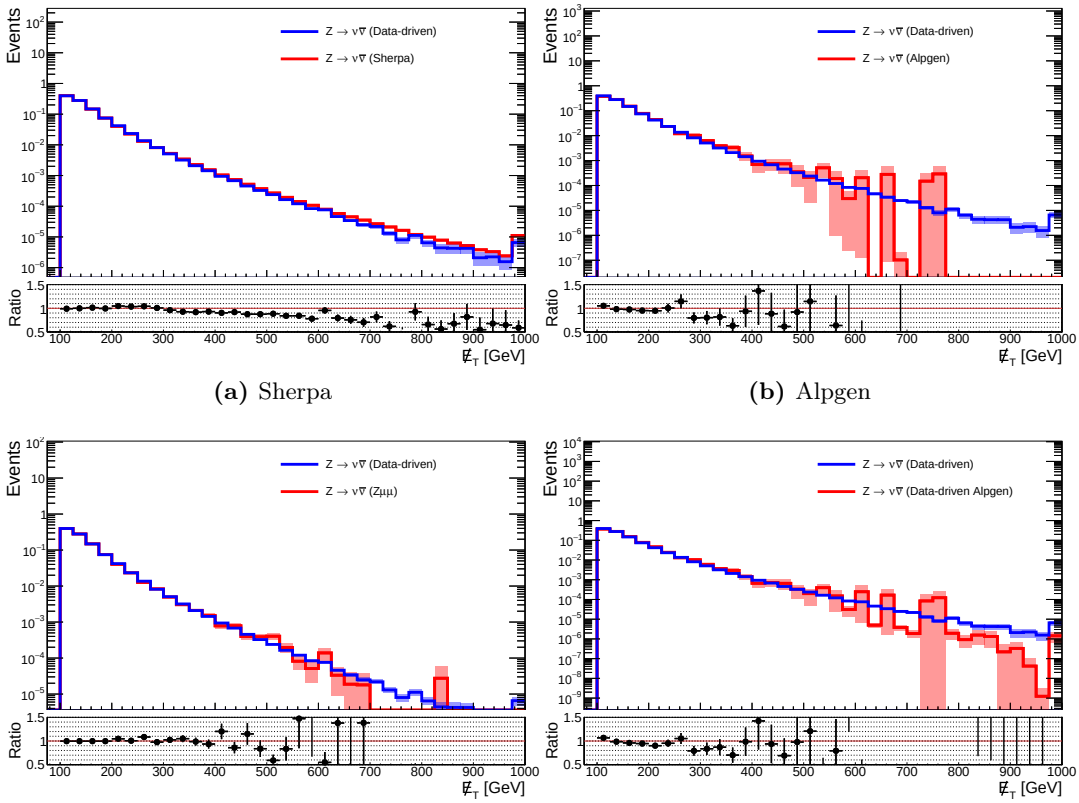

(a) Sherpa **(b)** Alpgen

(c) $Z \to \mu\mu$ **(d)** Alpgen TF

Figure 5.24: Comparison of $\not{E}_T$ between different estimates of the $Z \to \nu\nu$ background after the $\Delta\phi_{min}$ cut.

5.6.2 $W^{\pm} \to \ell^{\pm}\nu$

The second largest background in the SR is the $W \to \ell\nu$ process, which is isolated by requiring 1 tight muon. Additionally, a b-veto is placed to differentiate it from the top backgrounds. Figs. 5.25–5.28 show relevant kinematic variables plotted at difference stages in the CR**wjets** selection, and Table 5.38 lists the total yields after each cut. There is known mismodeling of the boson p_T in all the Sherpa V + jets samples, which can be seen clearly in the earlier stages of the $\not{E}_T^{\mu}$ distribution, and results in an underestimation compared to data. Ideally we would reweight each of these samples on the gauge boson p_T, as is done by [123]. However, this effect is relatively small in the SR, is well covered by systematic uncertainties, and appears to level out in the high $\not{E}_T^{\mu}$ region. To further validate these samples, the nominal Sherpa samples are compared to alternate Alpgen samples at early cut stages in Fig. 5.29.

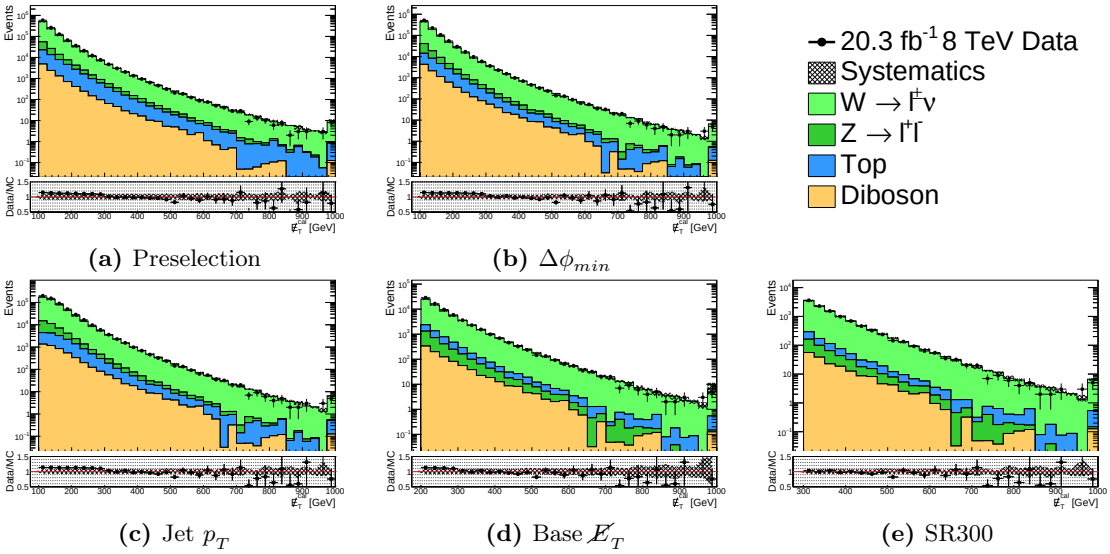

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.25: $\not{E}_T^{cal}$ distributions after each of the CR**wjets** cuts.

5.6.3 $t\bar{t}$

The $t\bar{t}$ process is the third largest background in the SR, and is validated in the single muon region with the requirement of 2 b-jets. Figs. 5.30–5.33 show relevant kinematic

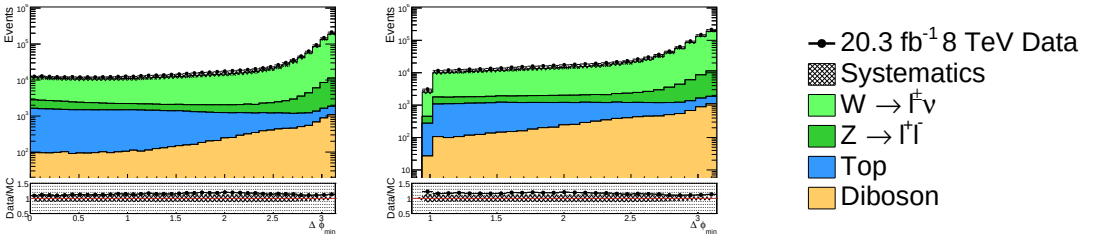

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.26: $\Delta\phi_{min}^{cal}$ distributions after each of the CR**wjets** cuts.

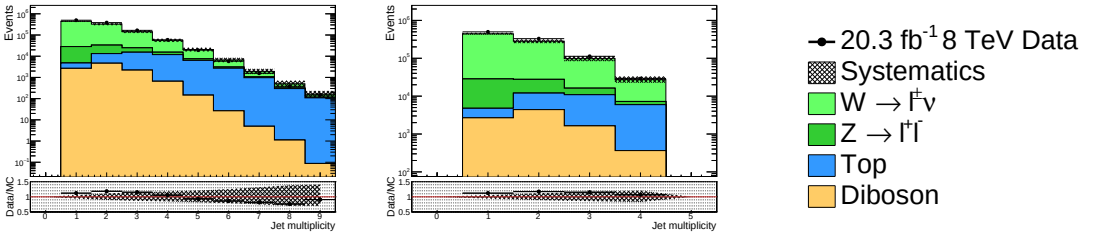

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.27: n_j distributions after each of the CR**wjets** cuts.

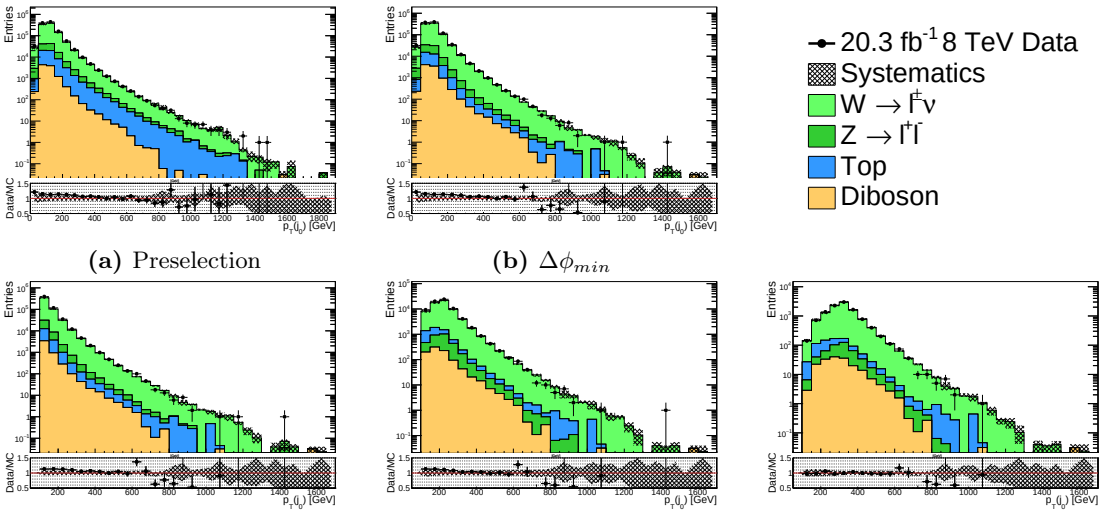

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** SR300

Figure 5.28: $p_T(j_0)$ distributions after each of the CR**wjets** cuts.

variables plotted at difference stages in the CR**top** selection, and Table 5.39 lists the total yields after each cut. Only the Powheg samples were available for this analysis, so no comparisons between estimates could be made. However, this background contributes significantly to CR**top**, CR**multijet**, *and* VR**lep**, tightly constraining it. The excellent agreement in shape and normalization across a wide variety of variables, in both hadronic and leptonic regions, shows the accuracy of this estimate.

5.6.4 Multijet

Although it has a very small contribution to the SR, the multijet background is the most difficult to estimate and typically has the highest uncertainties of any SM background. This can potentially introduce large uncertainties into the expected SR yields, and needs to be taken into account appropriately. The *jet smearing* method, discussed more in Appendix C, is used to generate pseudo-data that simulates the multijet process. This background can be isolated by reversing the $\Delta\phi_{min}$ cut of the nominal selection, which was primarily introduced just to suppress it. Figs. 5.34-5.37 show relevant kinematic variables plotted at difference stages in the CR**multijet** selection, and Table 5.42 lists the total yields after

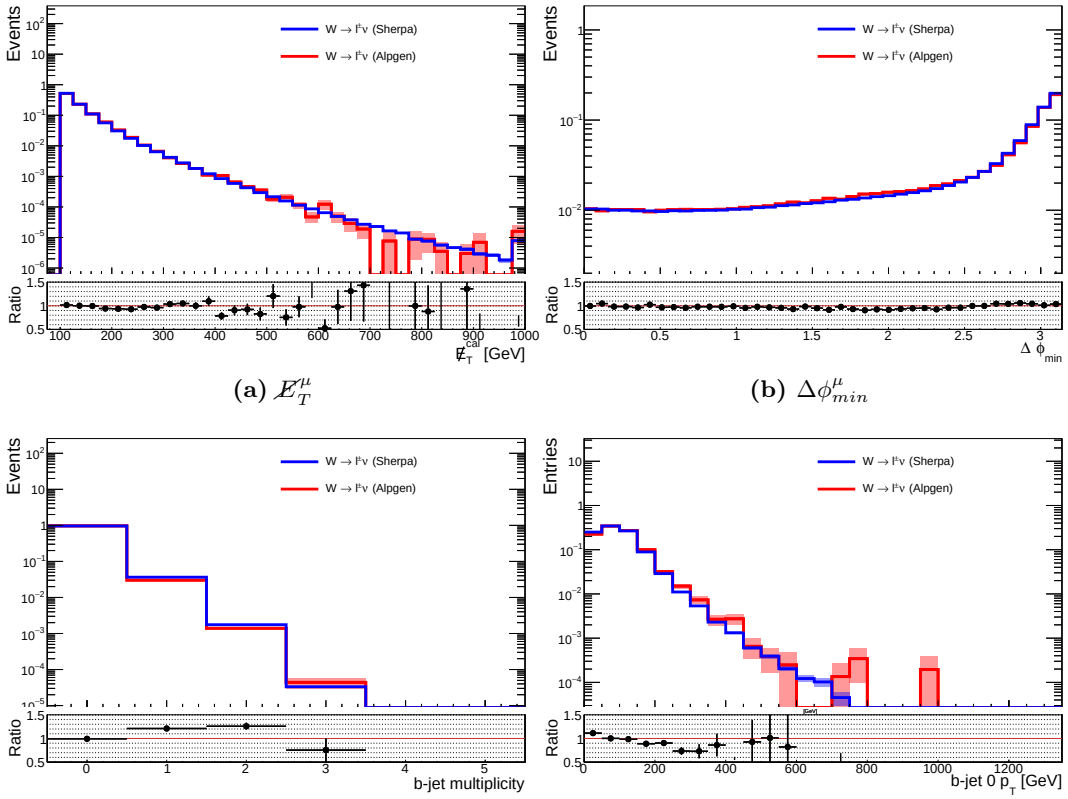

(a) $\not{E}_T^{\mu}$ **(b)** $\Delta\phi_{min}^{\mu}$

(c) n_b **(d)** $p_T(b_0)$

Figure 5.29: Comparison of kinematic variables between different estimates of the $W \to \ell\nu$ background at early stages of the single muon region

each cut.

In addition to the jet smearing estimate, Pythia samples of the multijet processes were available. As is illustrated in Fig. 5.38, above $\Delta\phi_{min} > 1$ the statistics are very poor for these samples and can't be used reliably. However, focusing on the CR**multijet** region allows for validation of the jet smearing method. Fig. 5.39 shows comparisons of relevant variables between the two estimates after the n_b cut of the CR**multijet** selection. The agreement is excellent, especially considering that the systematic uncertainties are $\mathcal{O}(100)\%$ on both estimates.

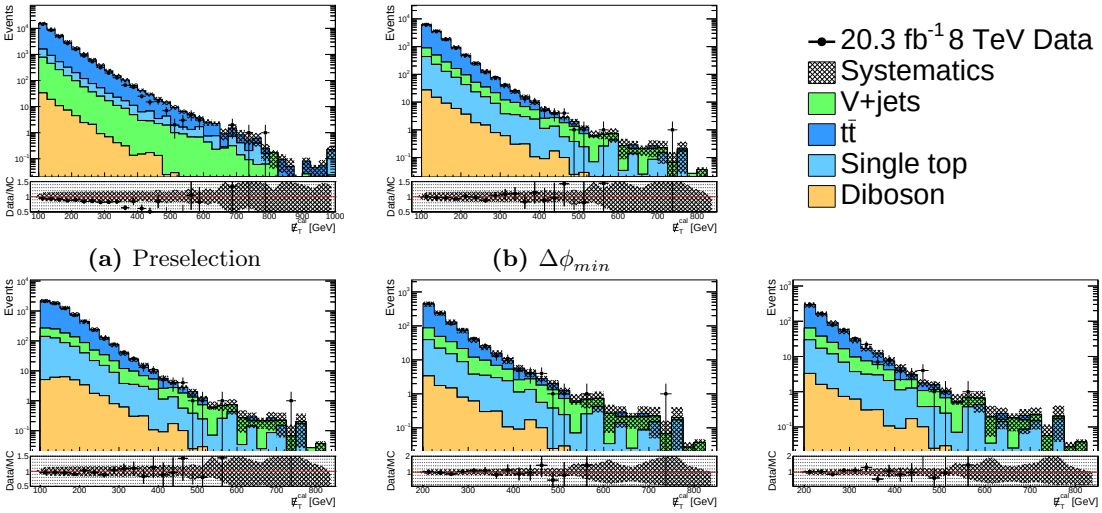

(a) Preselection **(b)** $\Delta\phi_{min}$ **(c)** Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.30: $\not{E}_T^{cal}$ distributions after each of the CR**top** cuts.

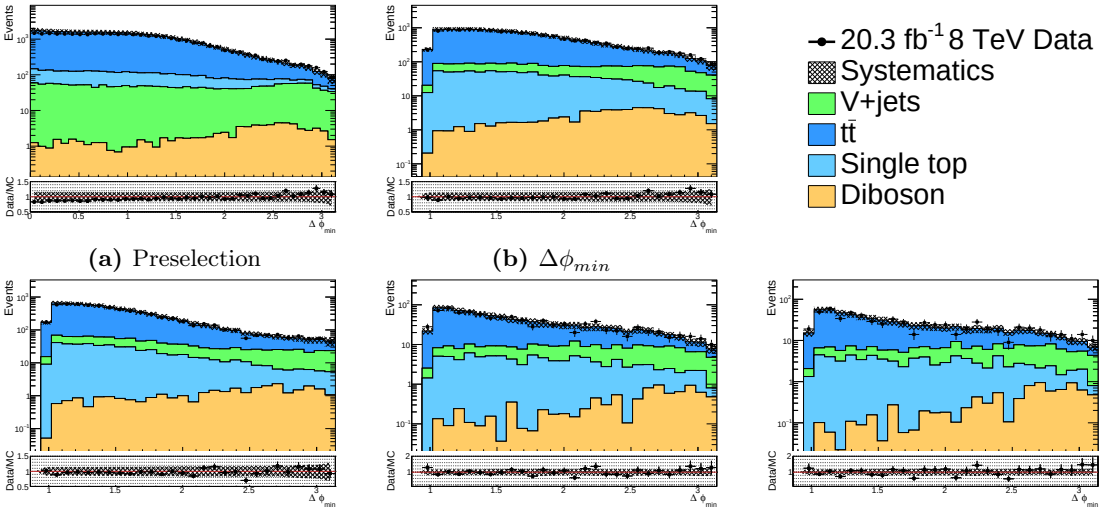

(a) Preselection **(b)** $\Delta\phi_{min}$ **(c)** Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.31: $\Delta\phi_{min}^{cal}$ distributions after each of the CR**top** cuts.

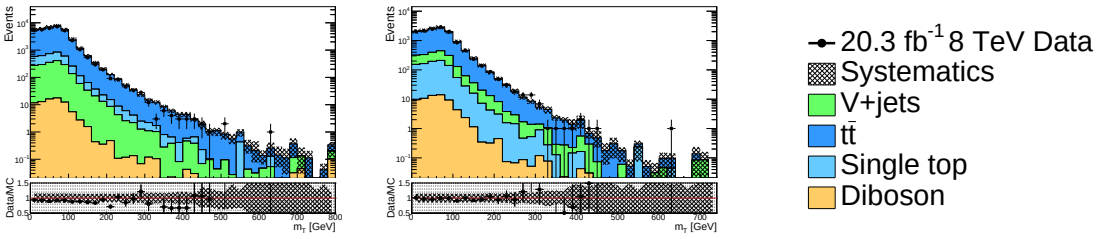

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.32: m_T^{μ} distributions after each of the CR**top** cuts.

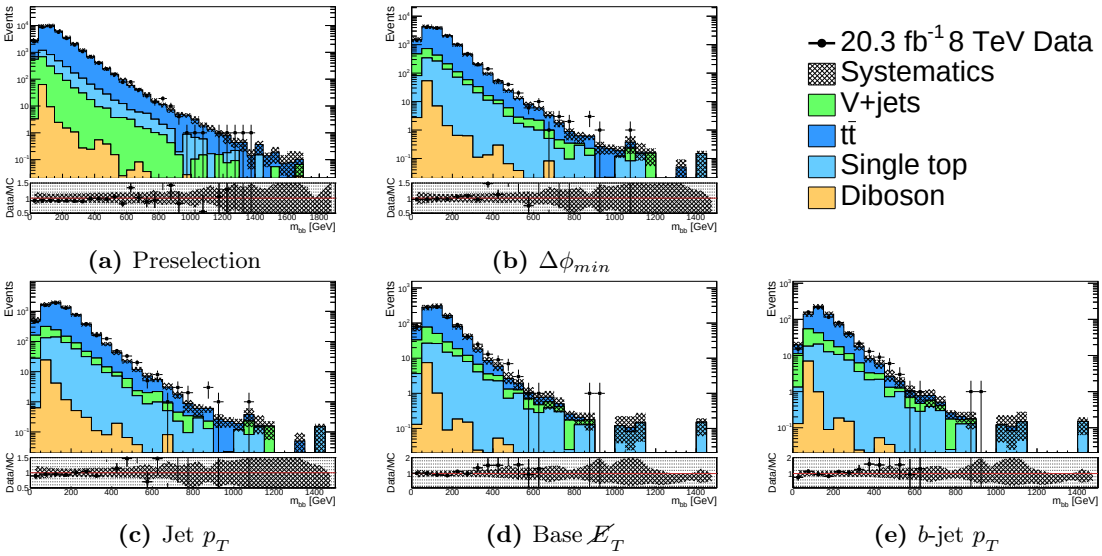

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.33: m_{bb} distributions after each of the CR**top** cuts.

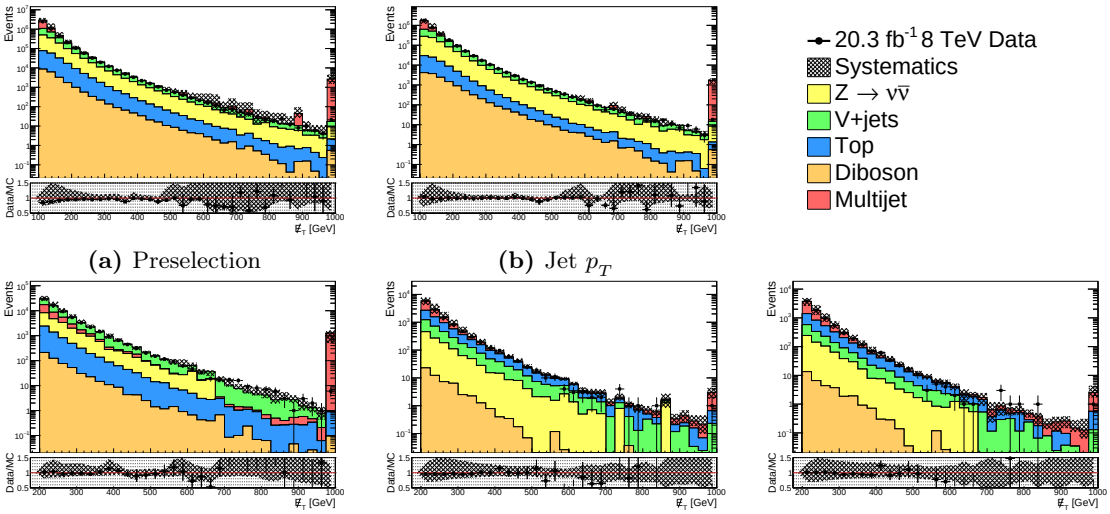

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.34: $\not{E}_T$ distributions after each of the CR**multijet** cuts.

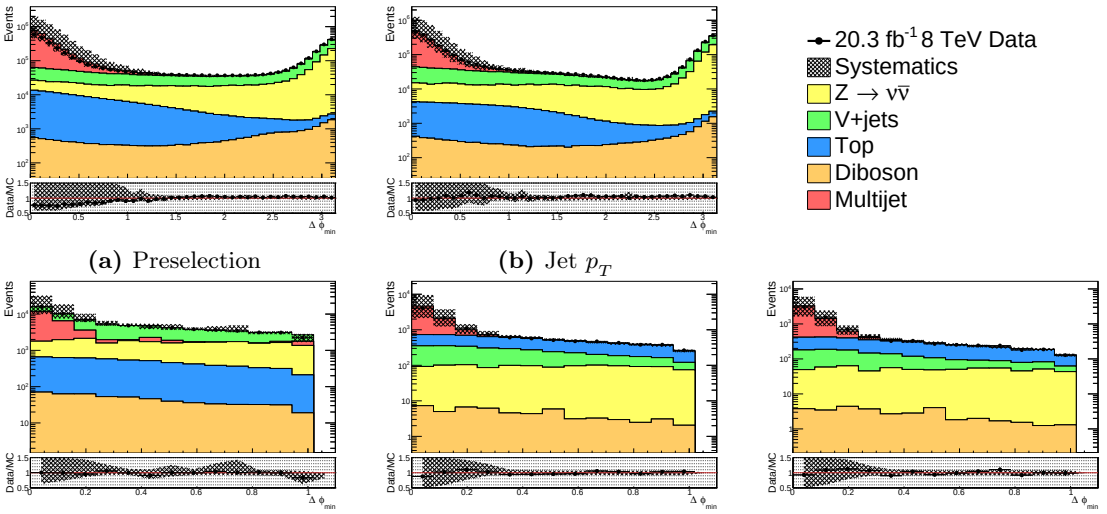

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.35: $\Delta\phi_{min}$ distributions after each of the CR**multijet** cuts.

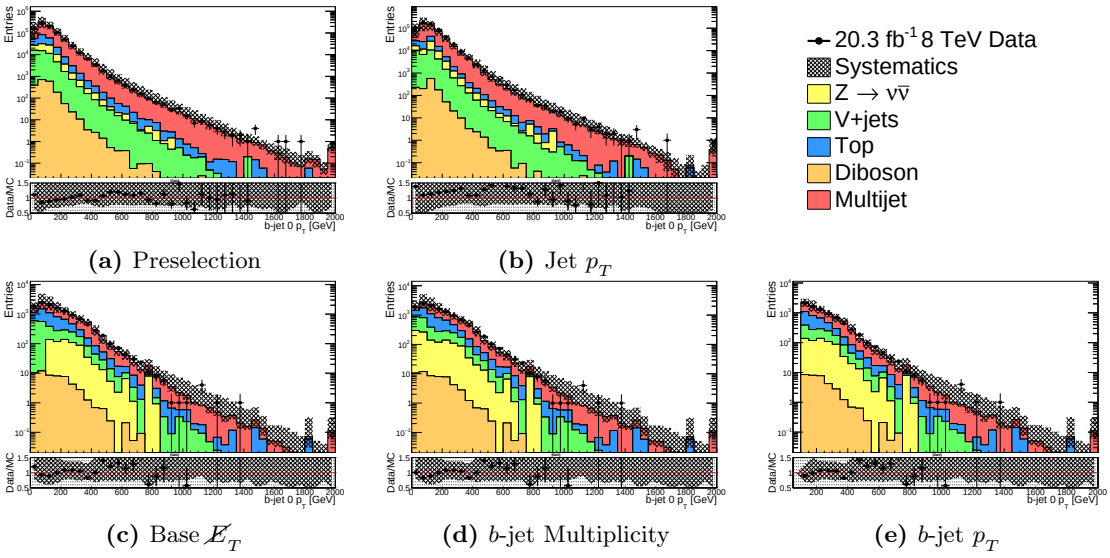

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.36: $p_T(b_0)$ distributions after each of the CR**multijet** cuts.

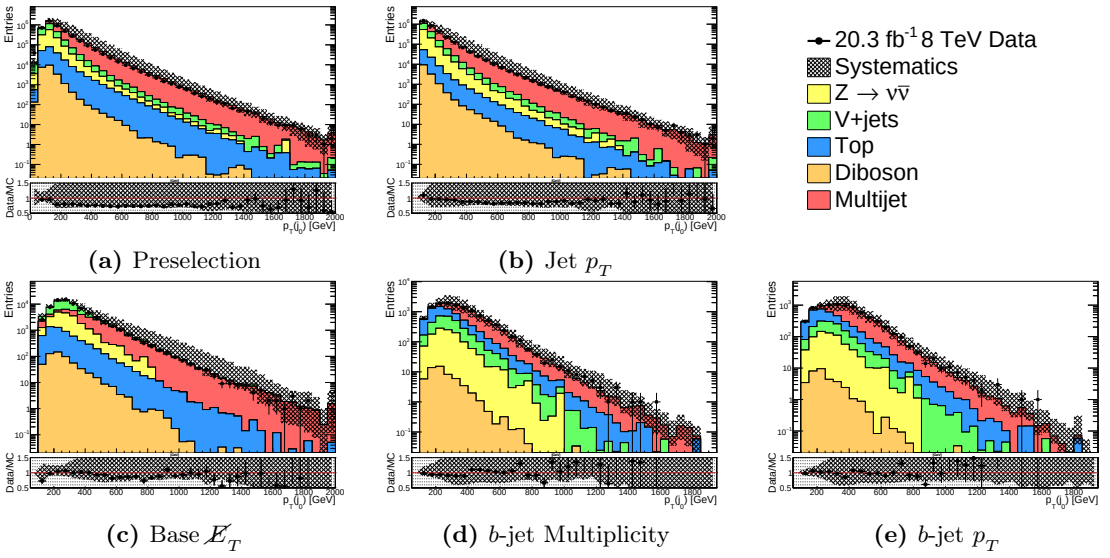

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.37: $p_T(j_0)$ distributions after each of the CR**multijet** cuts.

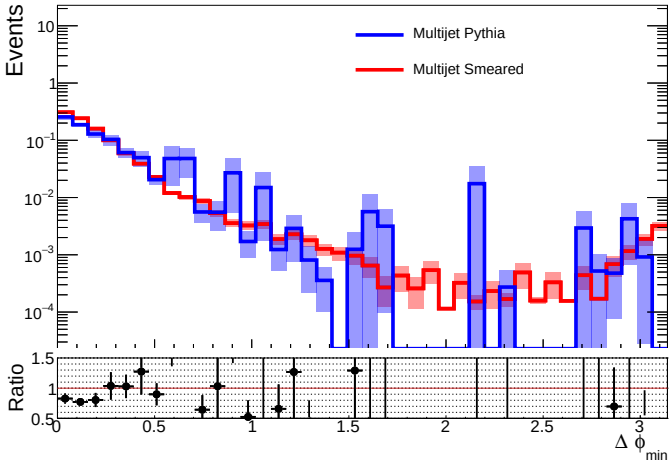

Figure 5.38: Comparison of $\Delta\phi_{min}$ between the estimates of the multijet background at preselection

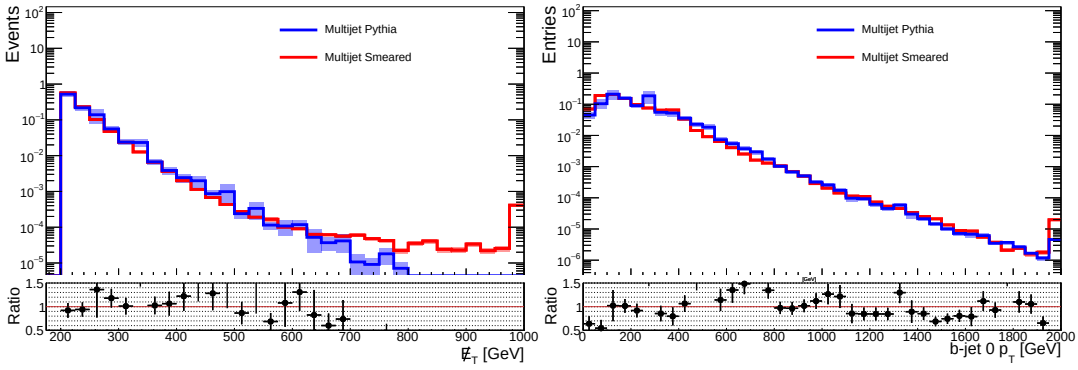

(a) $\not{E}_T$

(b) $p_T(b_0)$

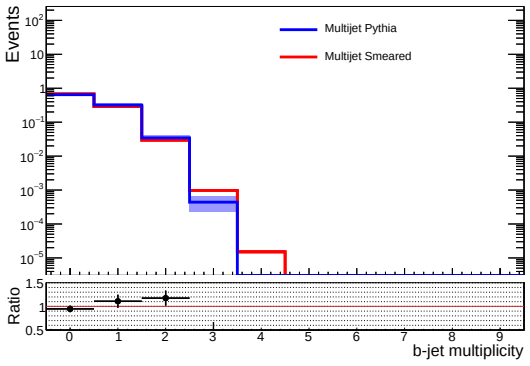

(c) n_b

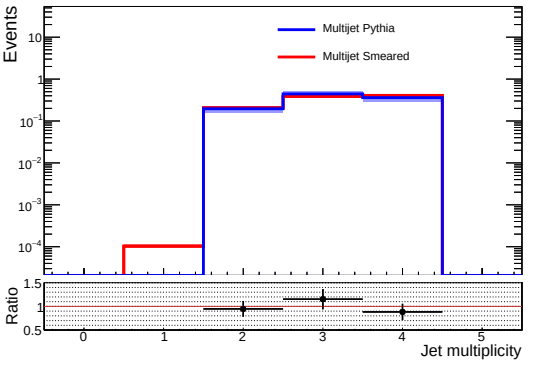

(d) n_j

Figure 5.39: Comparison of kinematic variables between different estimates of the multijet background after the b-tag requirement in the CR**multijet** region

5.6.5 $Z^0 \to \ell^{\pm}\ell^{\mp}$

The $Z \to \ell\ell$ background is very small in the SR, due to both the lower cross-section compared with $W \to \ell\nu$ and the lower efficiency of the lepton veto. After the b-jet p_T cut, it makes up less than 0.1% of the expected SM yield and even less after the $\not{E}_T$ cuts. However, one of the $Z \to \nu\nu$ estimates is produced with data corresponding to $Z \to \mu\mu$ events, so our understanding of this background should still be validated. Figs. 5.40-5.43 show relevant kinematic variables plotted at difference stages in the CR**zmm** selection, and Table 5.40 lists the total yields after each cut. Although the statistics are limited, the agreement between data and MC in this region is excellent.

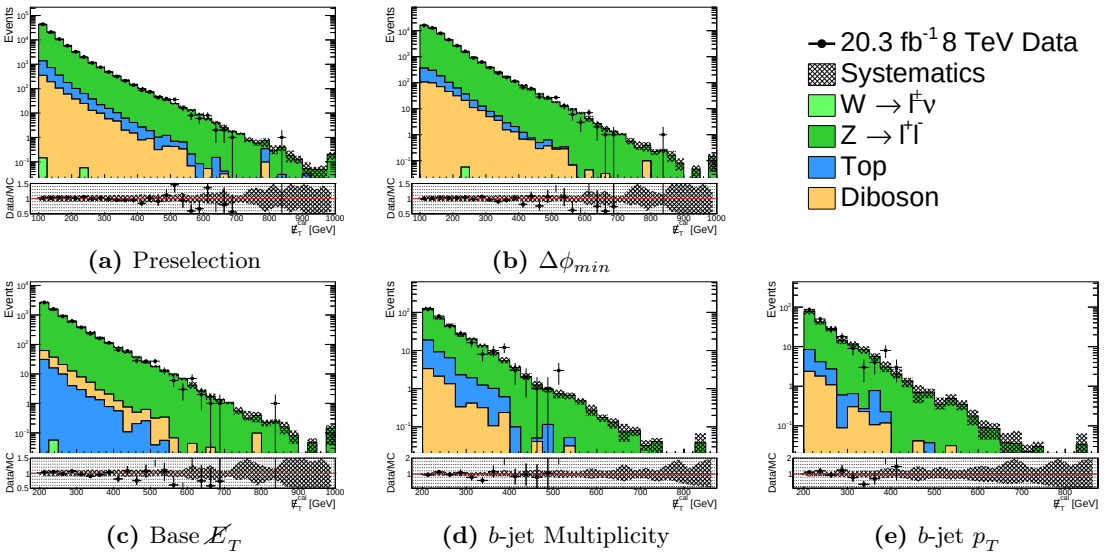

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.40: $\not{E}_T^{cal}$ distributions after each of the CR**zmm** cuts.

5.6.6 $\gamma + \mathbf{jets}$

As with the $Z \to \ell\ell$ background, the $\gamma +$ jets background is negligible in the SR. The nominal $Z \to \nu\nu$ estimate used in the analysis is derived from the $\gamma +$ jets CR though, so it's very important to validate this background. Figs. 5.44-5.47 show relevant kinematic variables plotted at difference stages in the CR**gamma** selection, and Table 5.41 lists the total yields after each cut. As mentioned above, there is known mismodeling of $p_T(\gamma_0)$ in

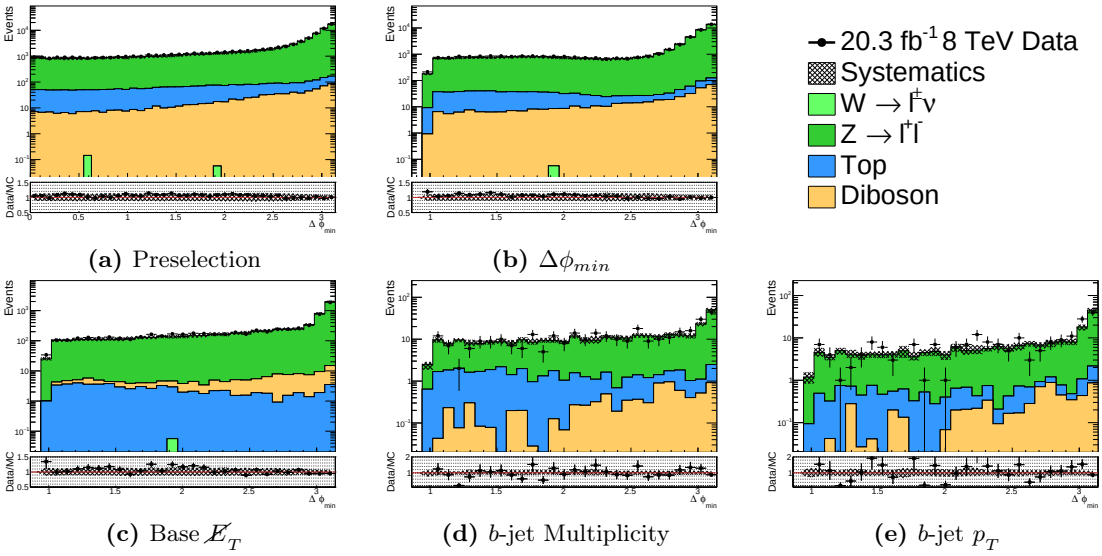

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.41: $\Delta\phi^{cal}_{min}$ distributions after each of the CR**zmm** cuts.

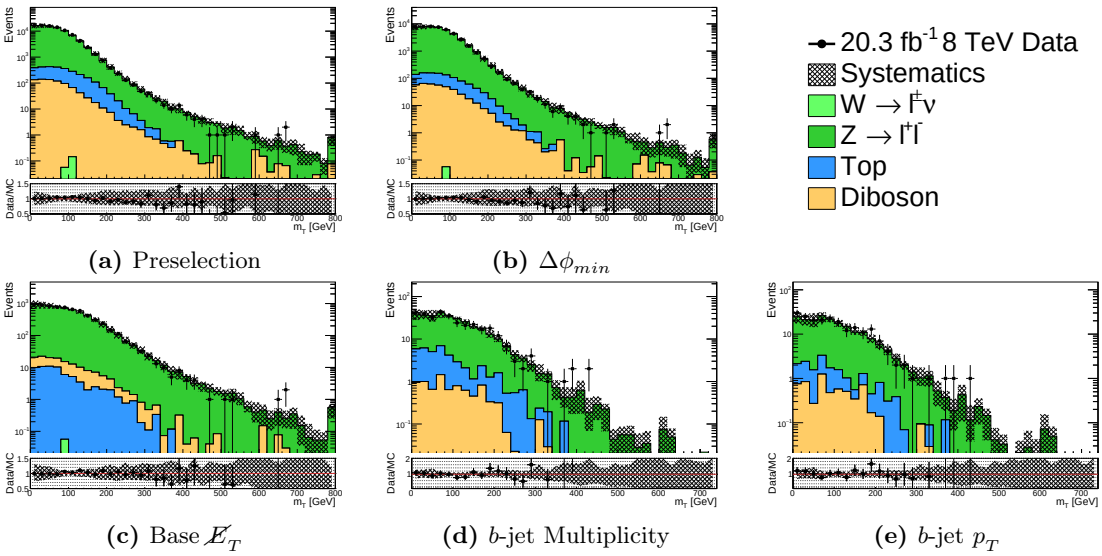

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.42: m^{μ}_T distributions after each of the CR**zmm** cuts.

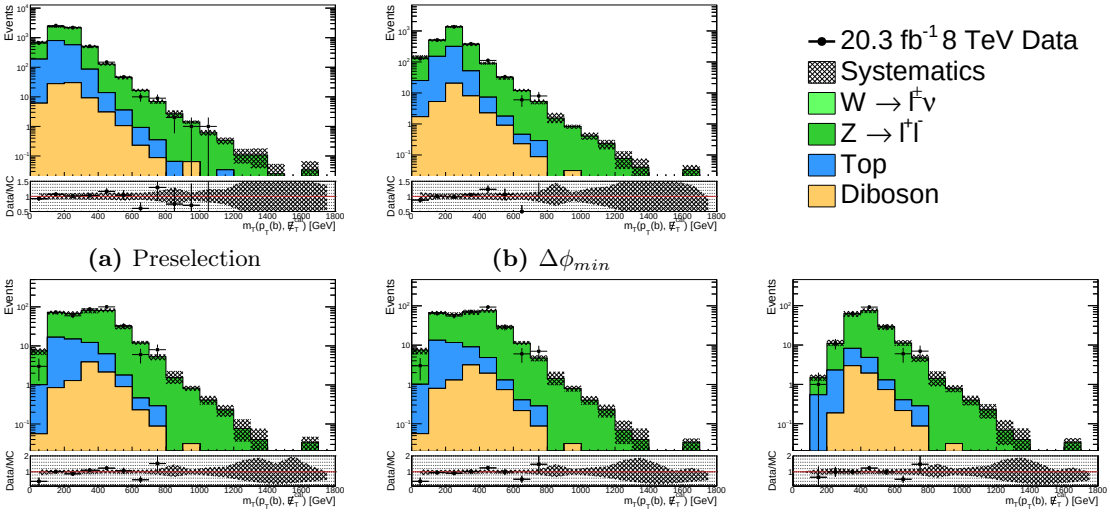

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.43: m_T^b distributions after each of the CR**zmm** cuts.

the Sherpa MC samples, but this effect is expected to cancel out in the TF, after dividing by the Sherpa $Z \to \nu\nu$ samples. Alpgen samples were also available, and comparisons between the two are plotted in Fig. 5.48.

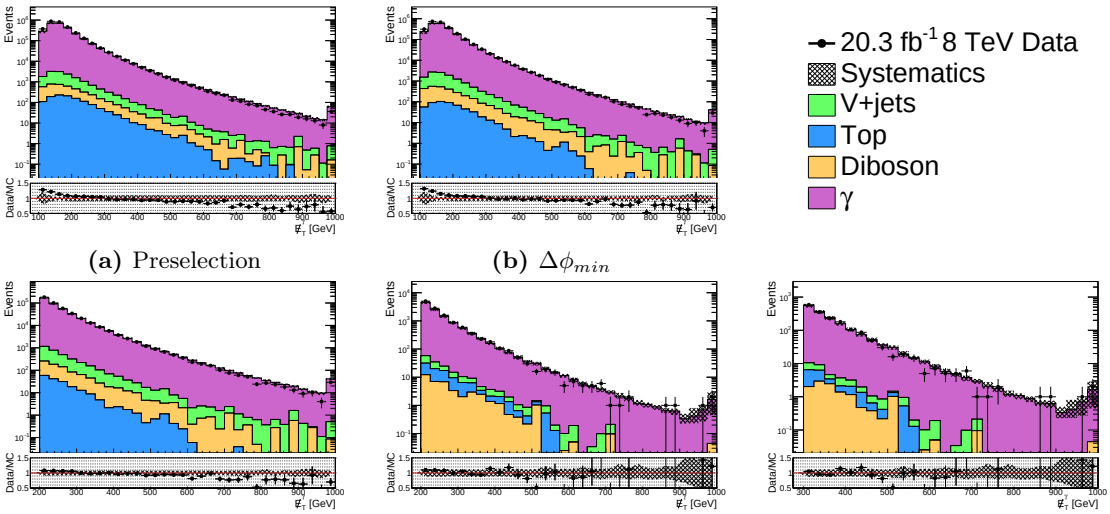

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet p_T **(e)** SR300

Figure 5.44: $\not{E}_T^\gamma$ distributions after each of the CR**gamma** cuts.

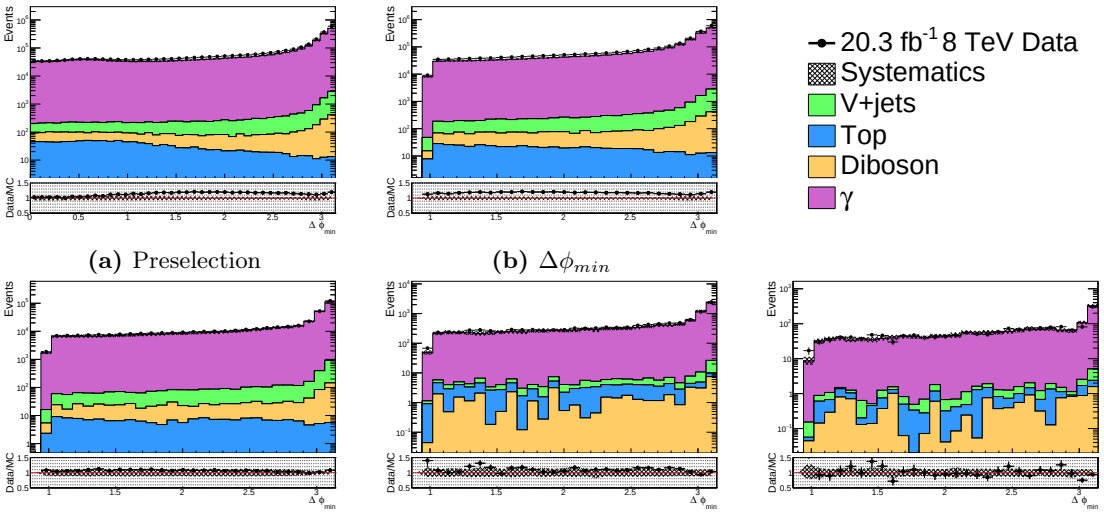

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet p_T **(e)** SR300

Figure 5.45: $\Delta\phi^{\gamma}_{min}$ distributions after each of the CR**gamma** cuts.

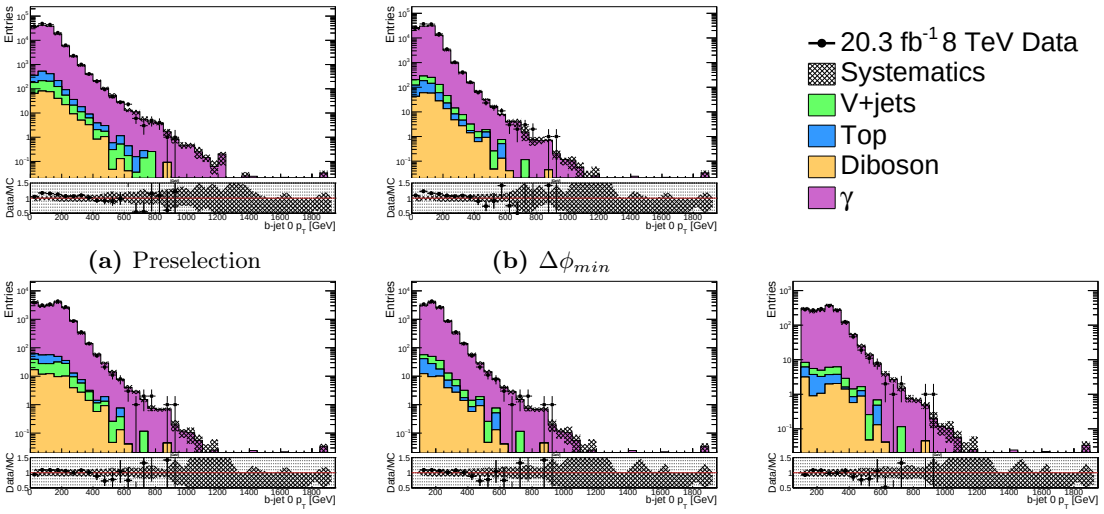

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet p_T **(e)** SR300

Figure 5.46: $p_T(b_0)$ distributions after each of the CR**gamma** cuts.

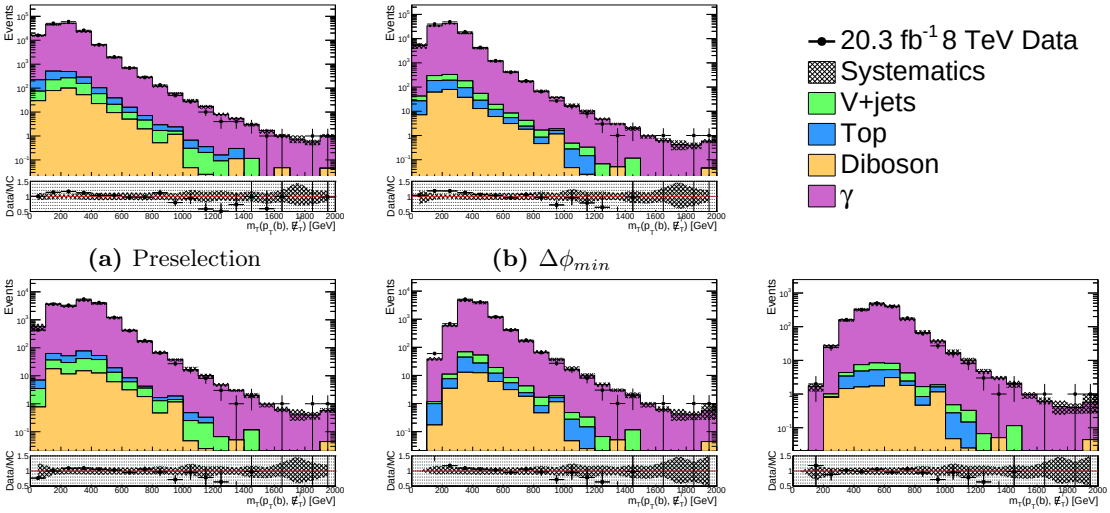

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Base $\not{E}_T$ **(d)** b-jet p_T **(e)** SR300

Figure 5.47: m_T^b distributions after each of the CR**gamma** cuts.

5.6.7 Diboson and Single top

The only SM processes without dedicated CRs are diboson and single top production, because of their very small cross-sections and similarities to more important backgrounds. Single top production contributes about 15% of the expected yield in CR**top**, and only 3% in the SR. The diboson processes have similar contributions of at most 3% in any region. Both of these backgrounds are normalized by their respective cross-sections, for which very conservative systematics are used. Additionally, any mismodeling in these backgrounds is likely to be, at least partially, compensated for during the normalization of the $t\bar{t}$, $W \to \ell\nu$, and $Z \to \nu\nu$ processes, due to their significant overlap.

5.6.8 Validation Regions

Figs. 5.49-5.54 show relevant kinematic variables plotted at different stages in the two VRs, and Tables 5.43 and 5.44 list the total yields after each cut. The agreement is very good for all of these, providing further evidence that the background estimates are accurate.

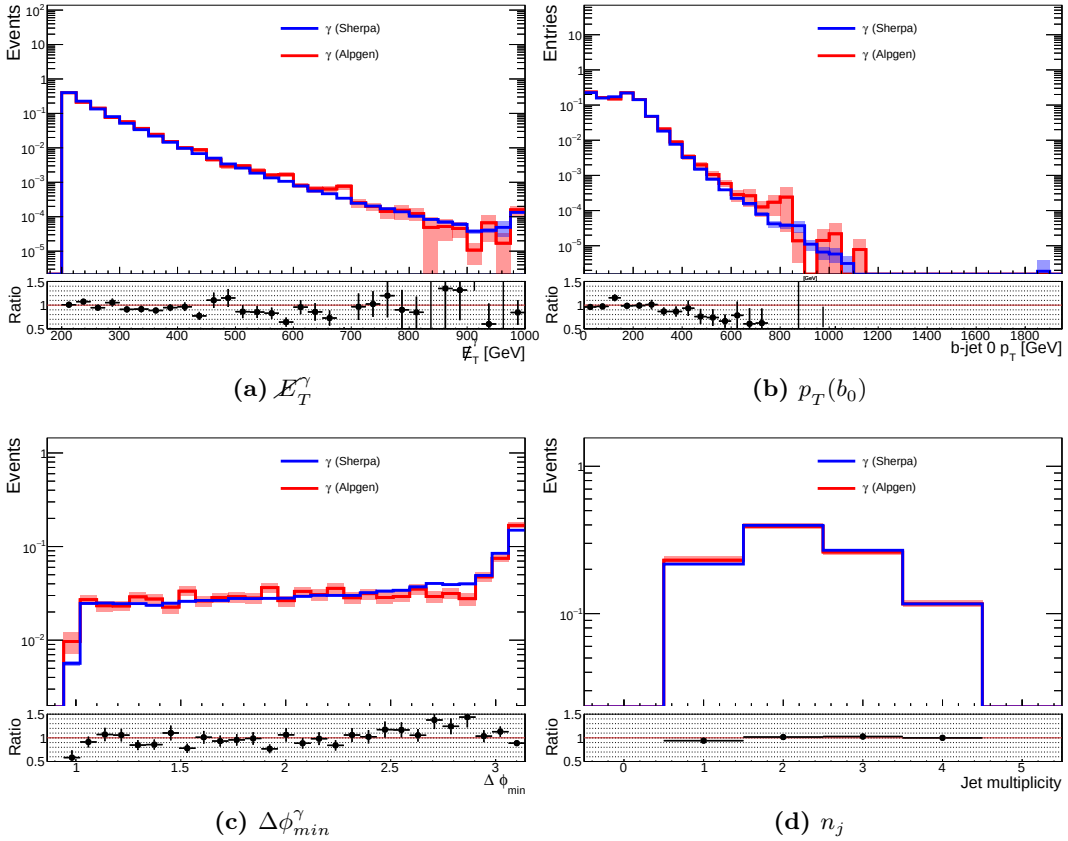

(a) $\not{E}_T^{\gamma}$ **(b)** $p_T(b_0)$

(c) $\Delta\phi_{min}^{\gamma}$ **(d)** n_j

Figure 5.48: Comparison of kinematic variables between different estimates of the γ + jets background after the b-tag requirement in the CR**gamma** region

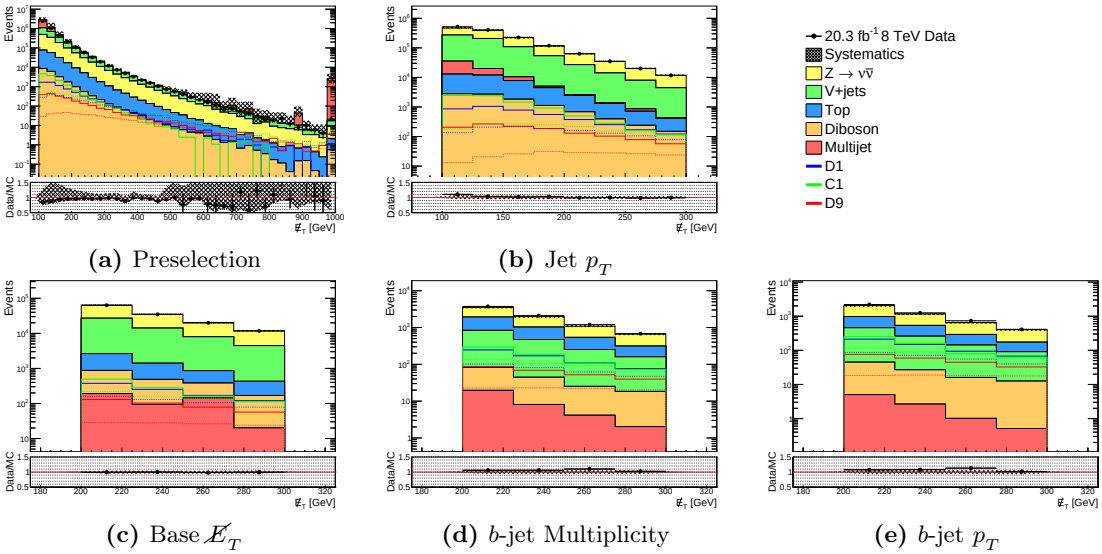

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.49: $\not{E}_T$ distributions after each of the VR**met** cuts.

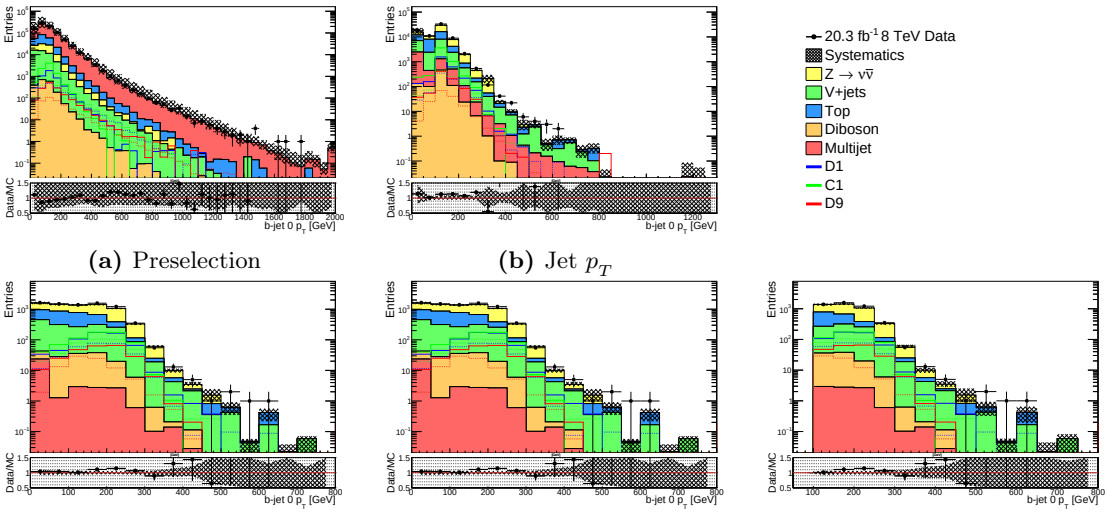

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.50: $p_T(b_0)$ distributions after each of the VR**met** cuts.

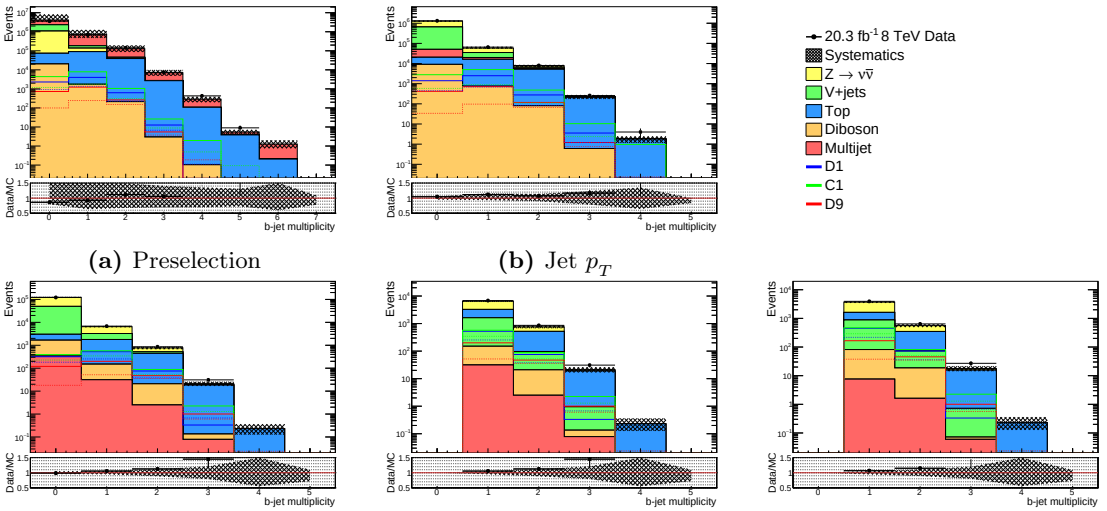

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure 5.51: n_b distributions after each of the VR**met** cuts.

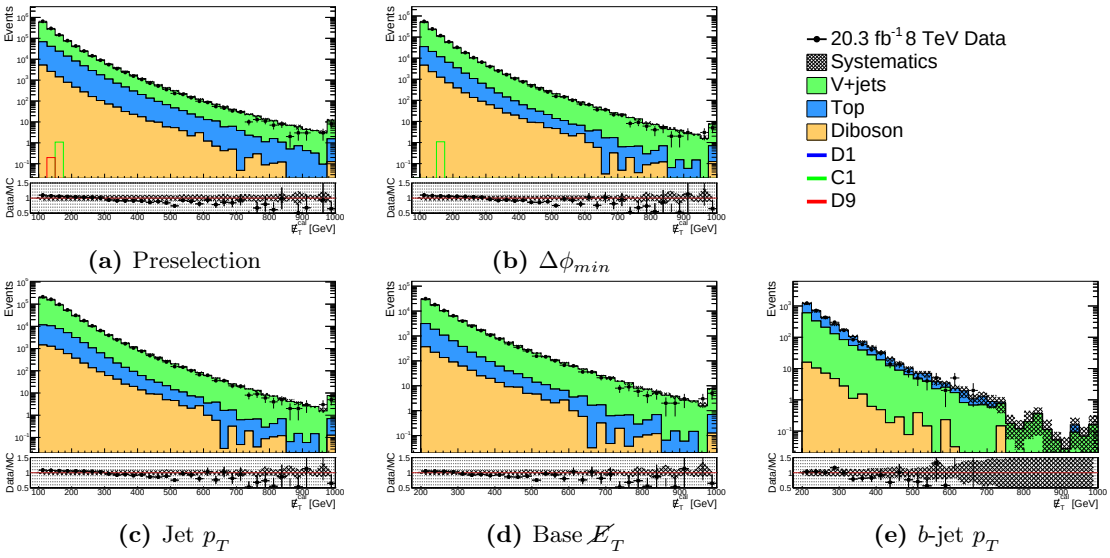

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.52: $\not{E}_T^{\mu}$ distributions after each of the VR**lep** cuts.

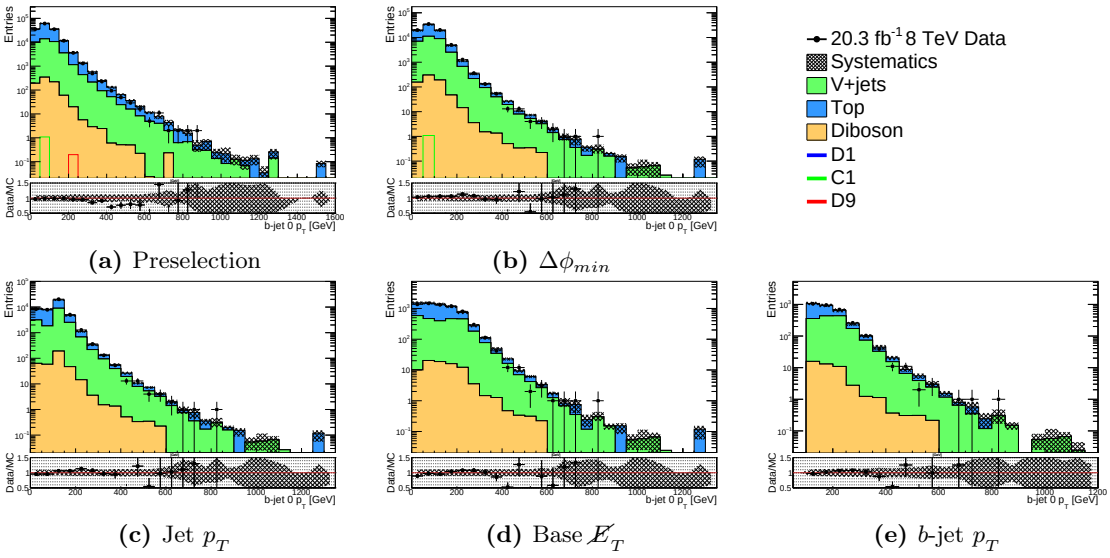

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.53: $p_T(b_0)$ distributions after each of the VR**lep** cuts.

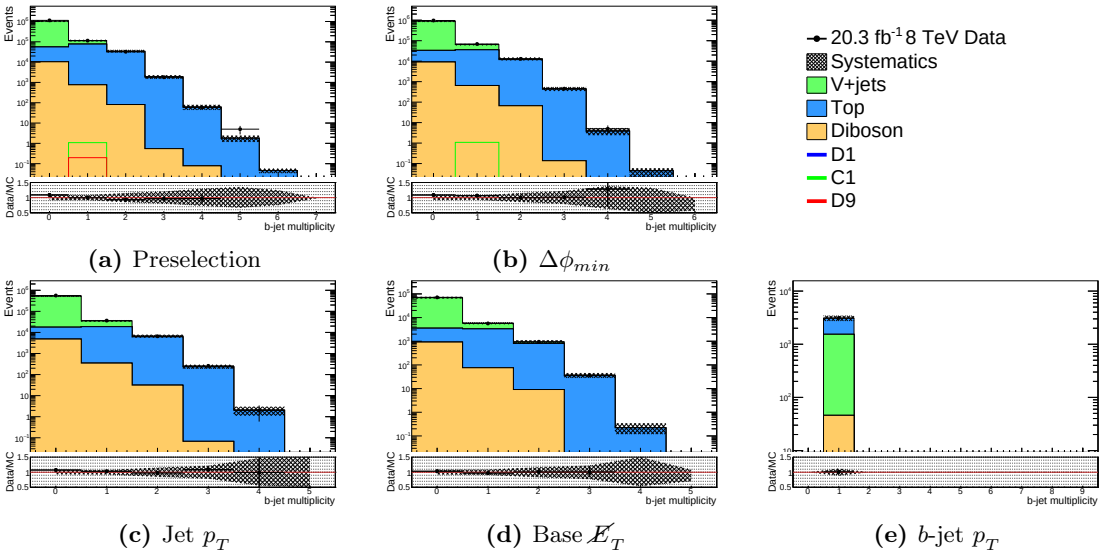

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** Base $\not{E}_T$ **(e)** b-jet p_T

Figure 5.54: n_b distributions after each of the VR**lep** cuts.

Cut	VV	Top	V + jets	Total SM	Data
Preselection	80,300 ± 7,200	507,000 ± 35,000	15,200,000 ± 620,000	15,800,000 ± 670,000	22,800,383 ± 4,775
$\not{E}_T > 100$ GeV	11,200 ± 1,300	157,000 ± 14,000	1,030,000 ± 96,000	1,190,000 ± 110,000	1,283,437 ± 1,133
Jet Multiplicity	11,000 ± 1,300	111,000 ± 8,000	1,000,000 ± 91,000	1,120,000 ± 100,000	1,223,808 ± 1,106
$\Delta\phi_{min}$	9,800 ± 1,100	72,100 ± 5,200	889,000 ± 79,000	971,000 ± 85,000	1,052,223 ± 1,026
Jet p_T	5,220 ± 590	37,700 ± 2,800	526,000 ± 40,000	569,000 ± 43,000	608,542 ± 780
Base $\not{E}_T$	1,010 ± 120	6,720 ± 630	68,300 ± 5,800	76,000 ± 6,600	78,139 ± 280
b-jet Multiplicity	75.8 ± 12	3,250 ± 310	2,480 ± 280	5,810 ± 600	5,651 ± 75
b-jet p_T	45.9 ± 8.3	1,550 ± 160	1,490 ± 200	3,080 ± 360	3,122 ± 56
SR300	7.95 ± 2.4	234 ± 25	262 ± 33	504 ± 60	442 ± 21

Table 5.43: Yields for VRlep region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature

Cut	Multijet	VV	Top	V + jets	$Z \to \nu\nu$	Total SM	Data
Preselection	8,490,000 ± 13,000,000	28,000 ± 2,400	240,000 ± 17,000	1,760,000 ± 30,000	1,340,000 ± 93,000	11,900,000 ± 13,000,000	13,920,151 ± 3,731
$\not{E}_T > 100$ GeV	2,470,000 ± 4,500,000	22,300 ± 2,300	184,000 ± 16,000	1,290,000 ± 81,000	1,060,000 ± 66,000	5,030,000 ± 4,600,000	4,409,360 ± 2,100
Jet Multiplicity	1,510,000 ± 2,700,000	21,100 ± 2,100	105,000 ± 7,700	1,230,000 ± 72,000	1,020,000 ± 64,000	3,890,000 ± 2,900,000	3,959,143 ± 1,990
$\Delta\phi_{min}$	38,000 ± 79,000	16,600 ± 1,600	54,600 ± 4,100	920,000 ± 50,000	889,000 ± 57,000	1,920,000 ± 180,000	1,991,835 ± 1,411
Jet p_T	34,200 ± 66,000	9,910 ± 1,100	32,200 ± 2,400	619,000 ± 44,000	636,000 ± 30,000	1,330,000 ± 140,000	1,397,988 ± 1,182
Base $\not{E}_T$	678 ± 800	1,460 ± 180	3,430 ± 340	48,300 ± 4,400	77,200 ± 3,600	131,000 ± 9,100	129,849 ± 360
b-jet Multiplicity	33.9 ± 220	137 ± 21	2,080 ± 230	1,550 ± 190	3,460 ± 190	7,260 ± 800	7,731 ± 88
b-jet p_T	9.25 ± 46	90.9 ± 15	1,030 ± 110	852 ± 110	2,280 ± 130	4,260 ± 400	4,601 ± 68

Table 5.44: Yields for VRmet region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

5.7 Systematic Uncertainties

In addition to the statistical uncertainties in any finite sized samples, there are also systematic effects intrinsic to the analysis that create further uncertainty in the results. Statistical uncertainties in this analysis are simply described by Poisson distributions, but for systematic uncertainties, it's much less straightforward to calculate both their values and correlations. The usual strategy for estimating systematics is to take very conservative measures, so that the final results will be weakened rather than strengthened due to any poorly understood effects. An analysis with overestimated systematics will still produce valid results, while an analysis with underestimated systematics will not. Since by their very nature systematics can't be directly measured, it can sometimes be very difficult to know if they're being underestimated or not. Therefore in this analysis, the lowest value for each systematic was chosen that can still arguably be an upper bound on the "true" value, which provides a reasonable balance between accuracy and precision.

For systematics that only directly affect a single measurement used in the analysis and are fully correlated from event to event, the treatment is pretty straightforward. The measured value, X, is taken to be a random variable with mean μ_x at its nominal value, and an uncertainty σ_x, which represents the systematic in question. The final result of the analysis can be considered a function of X, $f(X)$, with a mean value μ_f and an error, due to X, of σ_f. A Taylor expansion of $f(X)$ around μ_x gives,

$$f(X) = f(\mu_x) + \left.\frac{\partial f}{\partial X}\right|_{\mu_x} (X - \mu_x) + \dots \tag{5.3}$$

which can be truncated after the second term for small variations of X around μ_x. The mean value of $f(X)$ is easily shown to be $\mu_f = f(\mu_x)$, and the error due to X is,

$$\sigma_f = \sqrt{\left\langle f^2 \right\rangle - \langle f \rangle^2} \approx \sigma_x \left.\frac{\partial f}{\partial X}\right|_{\mu_x} \tag{5.4}$$

This linear approximation is known as the "delta method" [124], and holds for sufficiently

small σ_x. Plugging $X = \mu_x \pm \sigma_x$ into the Taylor expansion gives,

$$f(\mu_x \pm \sigma_x) \approx f(\mu_x) \pm \sigma_x \left.\frac{\partial f}{\partial X}\right|_{\mu_x} \approx f(\mu_x) \pm \sigma_f, \tag{5.5}$$

from which we get the linear approximation, $\sigma_f \approx |f(\mu_x) - f(\mu_x \pm \sigma_x)|$, that can be calculated by simply repeating the analysis with $X \to X \pm \sigma_x$. A generalization to asymmetric errors can also be made where,

$$\sigma_f^{\pm} \approx \left| f(\mu_x) - f\left(\mu_x \pm \sigma_x^{\pm}\right) \right|. \tag{5.6}$$

Although σ_x has an obvious meaning for some systematics, for others it's more obscure. For example, a common source of systematic uncertainty comes from inaccuracies in the MC generators. *A priori* there is no reason to suspect one MC generator provides more accurate results than another, but they all give slightly different predictions. One way to estimate this uncertainty is to repeat the analysis with each generator and take the envelope of the results as the $f(X) \pm \sigma_f$ band. However, in this case X refers to the type of generator used, and therefore isn't a random variable and σ_x isn't even well defined. The rationale though, is that X can also represent the theoretical cross-section, which has some unknown systematic uncertainty, σ_x, due to inaccuracies intrinsic to the generator. Varying the generator used in the analysis, can give *conservative* bounds on $f(X \pm \sigma_x)$, despite the fact that σ_x is unknown. Other types of systematics can only be estimated as a single modification to the nominal analysis, such as resolution uncertainties, which are estimated by randomly smearing objects in every simulated event. These are often *symmetrized* by simply using the difference between the nominal and varied yields as both σ^+ and σ^-. This is just another conservative measure that overestimates the uncertainty in order to simplify the analysis.

Systematic uncertainties are only calculated for the background and signal estimates. For MC simulations, there are a number of known sources of systematic uncertainties, which

are estimated using the techniques mentioned above. There are also uncertainties associated with each of the data-driven background estimates that need to be considered. Although systematics are never placed directly on the data, the pseudo-data samples contain negative weight MC events and are also corrected by functions derived from MC. Therefore, the MC systematics must be propagated from both of these components to the final pseudo-data yields.

In general, systematic uncertainties can be correlated to each other in complex, hard to measure ways. However, some simplifying assumptions can be made. First, each individual systematic is treated as fully *correlated* between different events and samples, to reflect the fact that systematics describe correlated flaws in the experiment, rather than statistical effects. The second assumption is that the systematics are fully *uncorrelated* with *each other*. This tends to be a good approximation, mostly because of the way in which we choose to split up the systematics and samples. Samples and systematics are already organized into roughly uncorrelated groups, simply because that makes them easier for us to understand. With these two assumptions, the total uncertainty on any measurement can be approximated by adding all of the individual systematic and statistical uncertainties in quadrature.

5.7.1 Hadronic Uncertainties

Flavor Tagging One of the largest sources of uncertainty in the MC samples used by this analysis comes from the requirement of a b-tagged jet [111]. To compensate for deviations in b-tagging efficiency from data, every simulated jet is assigned a SF, which has an intrinsic uncertainty propagated from its estimation. This uncertainty is split into three systematics corresponding to bottom, charm, and light jets. As mentioned above, the splitting of this systematic is related to the uncorrelated nature between jets of different flavor. The fractional uncertainties after the b-jet p_T cut, which are around 2 % for the combine SM background, are plotted in Fig. 5.55. For the lower threshold SRs, the b flavor systematic dominates with a 3 %–5 % effect on the MC samples and a 1 % effect on the $Z \to \nu\nu$

estimate. The c and *light flavor* uncertainties have a 5 % effect on the $W \to \ell\nu$ background, but an opposite effect for $Z \to \nu\nu$, which tends to suppress them.

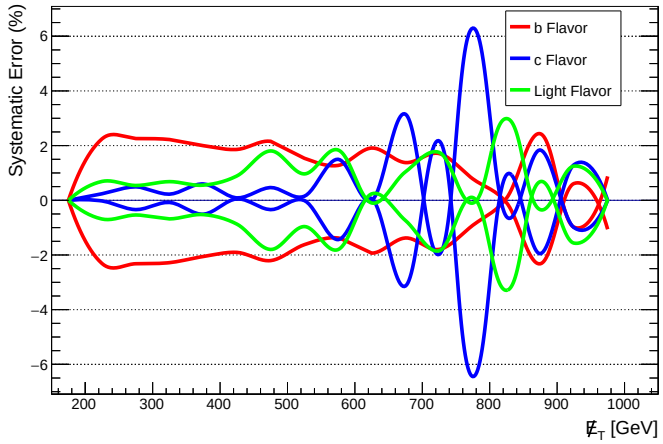

Figure 5.55: Fractional systematic uncertainties for the six flavor tagging uncertainties as functions of $\not{E}_T$. These are computed in the SR after the b-jet p_T cut.

Jet Energy Another major source of uncertainty in hadronic events comes from the jet measurements made by the calorimeters [125]. Both energy scale and resolution discrepancies between data and MC contribute to this, splitting it into two uncorrelated systematics. The fractional uncertainties after the b-jet p_T cut are around 3 % for the total SM expectation, and are plotted in Fig. 5.56. The *jet energy resolution* (JER) is well modeled by MC, so no smearing is applied nominally for simulated jets. To estimate the associated systematic though, they're randomly smeared in order to degrade the average energy resolution of the MC samples to one standard deviation lower than the nominal resolution. This effect, as mentioned before, is then symmetrized to get an upper bound. Below 500 GeV this uncertainty is a 1 % effect, so it doesn't have a huge impact on the analysis.

The JES uncertainty comes from the calibration procedure applied to *all* jets, which has uncertainties originating from the measurements that were used to define the calibrations. In principle there are over 60 nuisance parameters recommended by the JetEtMiss group, each of which should be treated as an individual systematic. However, a standard procedure used by many ATLAS analyses is to combine these into a single uncertainty by adding them

all in quadrature. Every simulated jet is then smeared by the combined effect of the JES systematics to get an estimate of the total uncertainty.

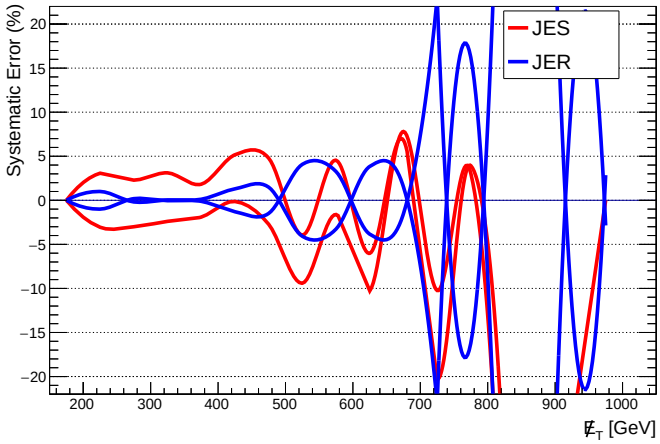

Figure 5.56: Fractional systematic uncertainties for the jet uncertainties as functions of $\not{E}_T$. These are computed in the SR after the b-jet p_T cut.

Soft Jets Because this analysis makes heavy use of the $\not{E}_T$ calculation, it's sensitive to the reconstruction of soft jets [120]. As with high energy jets, soft jets introduce energy scale and resolution systematics. Additionally though, there are two systematics originating from pileup corrections, that don't need to be applied to high p_T jets [126]. The first is introduced during the JVF cut applied in Section 5.3, and the second is from the pileup reweighting mentioned in Section 5.4. The fractional uncertainties of these four systematics are all plotted, after the b-jet p_T cut, in Fig. 5.57.

The soft term systematics are estimated in the same way as those for JES and JER, except that they act on objects with a much looser selection than jets. The uncertainty due to the pileup reweighting procedure is estimated by varying the pileup SF, applied to the measured μ in data, by its error ± 0.04, and the effect of the JVF cut is estimated by varying the cut value ± 0.03. Not surprisingly, the pileup systematics have a very small effect, of about 0.3 %, throughout all SRs The soft term systematics have a larger effect, since they do directly influence the $\not{E}_T$ of every event, but they're still much smaller than the jet uncertainties below 600 GeV.

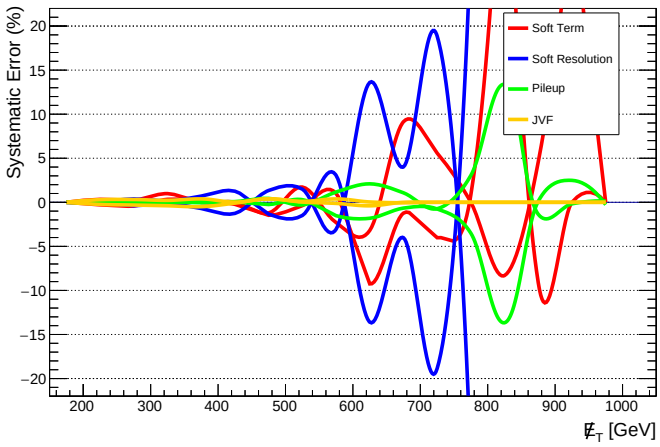

Figure 5.57: Fractional systematics for the soft uncertainties as functions of $\not{E}_T$. These are computed in the SR after the b-jet p_T cut.

5.7.2 Muon Uncertainties

Although there are no muons in the final SR, they're used in four of the CRs, and have five systematics associated with them [106]. As with jets, an energy scale uncertainty is introduced, along with two systematics coming from uncertainties in the p_T smearing applied to correct for mismodeling of the ID and MS resolutions. However, because muons are always removed from the $\not{E}_T$ calculation, these systematics are negligible everywhere in the analysis, as illustrated in Fig. 5.58. A systematic is also included to account for uncertainties in the trigger SFs applied to muon events. This and the fifth systematic, from the muon identification and reconstruction efficiency, contribute a $\sim 0.5\%$ uncertainty to events containing muons. All five of these effects are plotted in Fig. 5.59, using the $Z \to \nu\nu$ from $Z \to \mu\mu$ estimate after the nominal b-jet p_T cut.

5.7.3 Photon Uncertainties

Although photons are not considered in the nominal analysis, the $Z \to \nu\nu$ estimate makes use of high p_T prompt photon production, subtracted from the $\not{E}_T$ calculation. There are seven systematics from the photon in these events, which need to be considered [117]. Five of these are due to uncertainties in the calorimeter energy measurement, which become

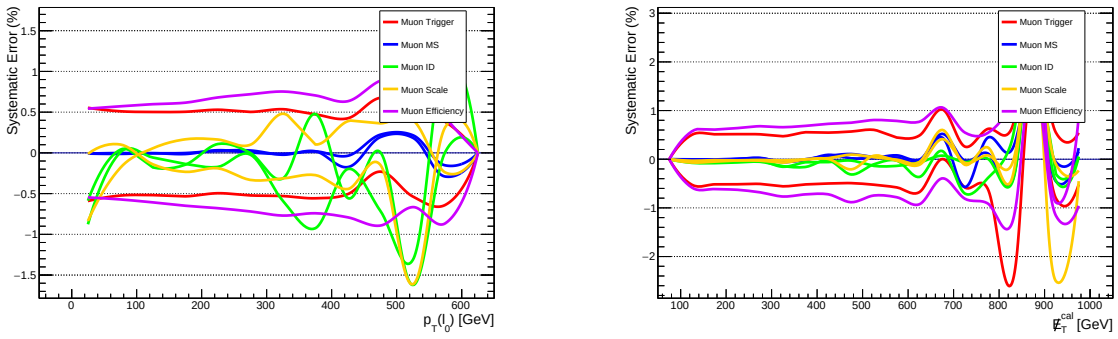

Figure 5.58: Cancelation of many of the muon systematics after subtraction from $\not{E}_T$. On the left is the effect of muon uncertainties on the leading muon's p_T, and on the right is the $\not{E}_T$ calculated without the muon's contributions. Since the only quantity relevant to this analysis is $\not{E}_T^{\mu}$, uncertainties on the measured muon momentum become irrelevant.

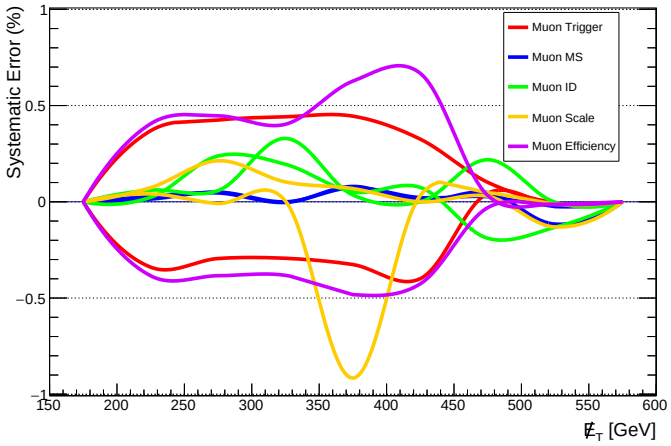

Figure 5.59: Fractional systematics for the muon uncertainties as functions of $\not{E}_T$, using the di-muon $Z \to \nu\nu$ estimate. All systematics are computed after the b-jet p_T cut.

negligible after subtraction from the $\not{E}_T$, as illustrated in Fig. 5.60. The SFs applied to each photon for the identification and reconstruction efficiency produce very small, but non-negligible, uncertainties on the $Z \to \nu\nu$ estimate. The largest uncertainty though, comes from the photon isolation requirement. Following the mono-photon analysis [48], a 4% systematic is applied to all events in which a tight photon is selected, in order to cover the uncertainty from this cut. All seven of these systematics are plotted in Fig. 5.61.

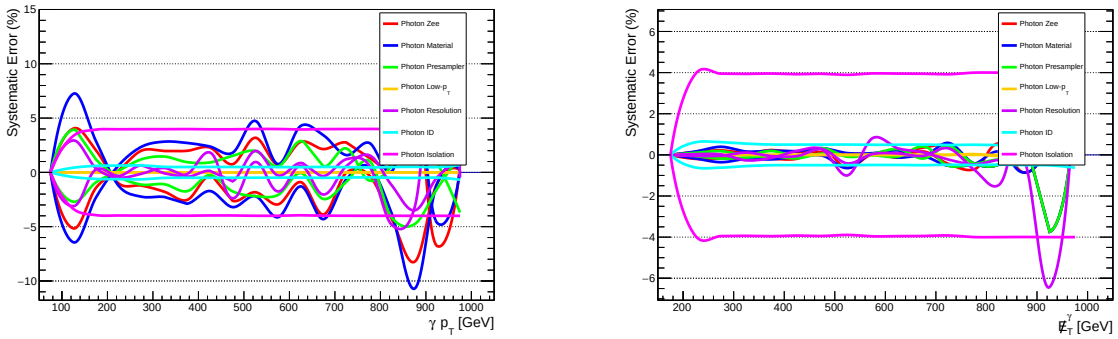

Figure 5.60: Cancelation of many of the photon systematics after subtracting. On the left is the effect of photon uncertainties on the photon's p_T, and on the right is the $\not{E}_T$ calculated without the photon contributions. Since the only quantity relevant to this analysis is $\not{E}_T^{\gamma}$, energy scale and resolution uncertainties on the photon become irrelevant.

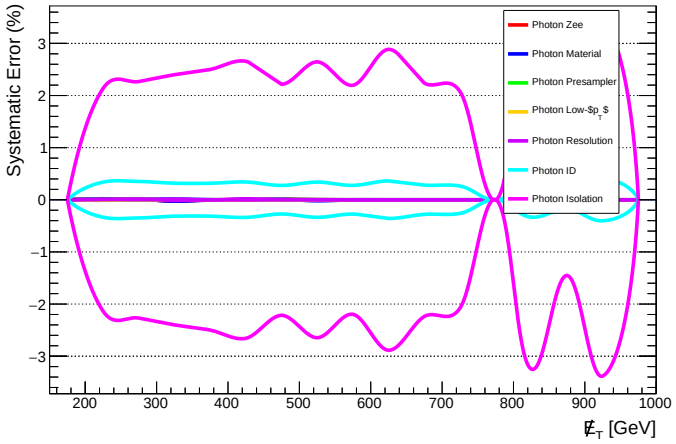

Figure 5.61: Fractional systematics for the photon uncertainties as functions of $\not{E}_T$. All systematics are computed after the b-jet p_T cut.

5.7.4 Theory Uncertainties

In addition to uncertainties arising from the ATLAS detector, there are also theoretical uncertainties from the MC generation process. A large number of approximations are made throughout the generation of MC samples, as was discussed in Section 5.1, and these must be considered to get a reliable result. Fig. 5.62 shows the fractional effect on the total SM estimate in the SR for the generator uncertainties discussed below.

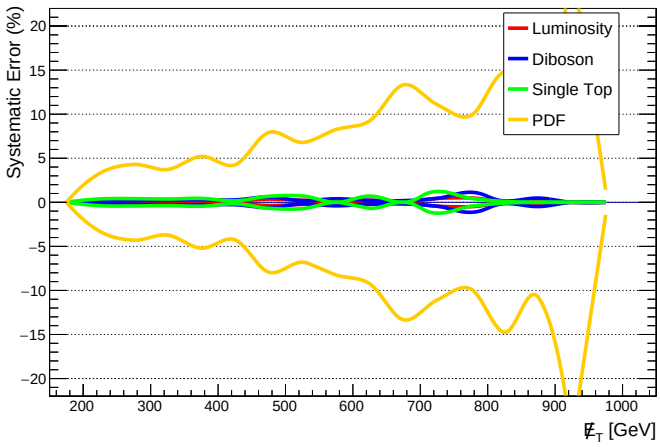

Figure 5.62: Fractional systematics for the generator uncertainties as functions of $\not{E}_T$. All systematics are computed after the b-jet p_T cut.

Luminosity MC samples are typically normalized to the integrated luminosity corresponding to the data under consideration. The amount of analysis-quality data measured by ATLAS in the 8 TeV run, was estimated to be $20.3\,\text{fb} \pm 2.8\%$ [100]. Since the exact luminosity is unknown, any uncertainty in the nominal value will propagate through the analysis. This is treated as a 2.8% uncertainty on all samples normalized by the luminosity. This systematic doesn't affect samples fit to data, since the fitting process compensates for it. In this analysis, the only processes without a dedicated CR are the diboson, single top, and EFT samples, all of which gain a 2.8% systematic from this uncertainty.

Cross-section The cross-sections of the simulated SM processes have various uncertainties associated with their calculation. Most of the samples in this analysis were simulated at LO, and given k-factors corresponding to NLO corrections to their total cross-sections. Typically, one would normalize each sample with the LO cross-section and the NLO k-factor, multiplied by the integrated luminosity. In this analysis though, all of the major backgrounds are normalized to data in dedicated CRs. This generates statistical uncertainties from the data and systematics from processes that contaminate the fit, so for these backgrounds, no additional systematic is needed on the cross-section.

While the major backgrounds are normalized to data, the diboson, single top, and

signal samples are not, so the uncertainty on these cross-sections have to be considered. For diboson and single top production, the uncertainties on the *measured* cross-section at ATLAS are used as conservative upper bounds on the theoretical uncertainties, which are typically much lower. These measurements are summarized in Table 5.45, where systematic and statistical uncertainties are combined in quadrature, and luminosity uncertainties have been excluded to avoid double counting. An additional conservative measure is taken for the diboson processes, by using the maximum of these three uncertainties, 8.2%, as the systematic uncertainty on the entire background.

Process	Cross-section	Systematic
W^+W^-	$71.4^{+1.2}_{-1.2}(\text{stat})^{+5.0}_{-4.4}(\text{syst})$ pb	7.2%
$W^\pm Z^0$	$20.3^{+0.8}_{-0.7}(\text{stat})^{+1.2}_{-1.1}(\text{syst})$ pb	7.1%
Z^0Z^0	$7.1 \pm 0.5(\text{stat}) \pm 0.3(\text{syst})$ pb	8.2%
single t	$68 \pm 2(\text{stat}) \pm 8(\text{syst})$ pb	12%

Table 5.45: The diboson and single top production cross-sections measured at ATLAS [127–130]. The maximal fractional 1σ uncertainty for each process is given in the last column. These are used as conservative upper bounds on the cross-section uncertainty.

Normalization As mentioned earlier, the backgrounds fit in dedicated CRs don't have luminosity and cross-section uncertainties, but there are still errors associated with the fitting process. These are mainly statistical in nature, although the systematics on the *other* processes in the region do contribute. These were listed back in Table 5.36, and their effects on the predicted SM yield in the SR are plotted in Fig. 5.63.

Factorization and Renormalization Scales In calculating the amplitudes of the hard interactions, divergences need to be accounted for. UV divergences refer to loop diagrams that involve integration over arbitrarily high momenta and must be renormalized, as discussed in Chapter 2. Renormalization results in the running of the *strong coupling*, α_s, as a function of the *renormalization scale*, μ_R. Another problem arises in QFT, where any number of additional gluons are emitted in an interaction that are either below the energy

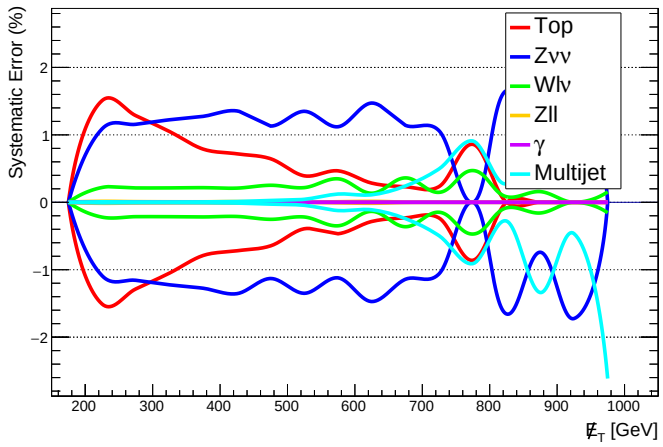

Figure 5.63: Fractional systematics for the normalization uncertainties as functions of $\not\!\!E_T$. All systematics are computed after the b-jet p_T cut.

threshold of the detector or collinear with one of the primary particles. These should all contribute to the amplitude, since they're indistinguishable, resulting in divergent terms with arbitrarily high numbers of gluons. These *infrared* (IR) divergences are dealt with through *factorization*, which connects non-perturbative and perturbative QCD at some *factorization scale*, μ_F. In principle, no physical measurement should depend on either μ_R or μ_F. However, truncating the perturbative expansion of an amplitude adds a dependence on both, so that in LO samples, the choice of μ_R and μ_F can have a significant effect.

The EFT signal simulations can potentially have very large factorization and renormalization scale uncertainties, which need to be estimated. The generator used to produce these samples was Madgraph [88], which dynamically chooses appropriate μ_F and μ_R for each interaction. To estimate the uncertainties, each of the EFT signals was simulated with the default factorization and renormalization scales chosen by Madgraph scaled by 0.5 and 2.0, as in [47, 48]. The resulting systematic variations are listed in Table 5.46.

Beam Energy The average energy of each proton beam during the 8 TeV run was $(3{,}988 \pm 26)$ GeV [131]. Comparing this with the 4 TeV used by Madgraph for the production of the EFT samples, the uncertainty on the beam energy should be propagated through the analysis as another systematic. To estimate the effect, simulations were rerun

Mass (GeV)	C1	D1	D9
10	$^{+3}_{-3}$	$^{+4}_{-4}$	$^{+3}_{-3}$
50	$^{+3}_{-3}$	$^{+4}_{-4}$	$^{+3}_{-3}$
100	$^{+3}_{-3}$	$^{+4}_{-4}$	$^{+3}_{-4}$
200	$^{+4}_{-4}$	$^{+5}_{-4}$	$^{+4}_{-4}$
400	$^{+4}_{-4}$	$^{+5}_{-5}$	$^{+4}_{-4}$
700	$^{+6}_{-6}$	$^{+7}_{-6}$	$^{+5}_{-5}$
1000	$^{+8}_{-7}$	$^{+8}_{-8}$	$^{+6}_{-6}$
1300	$^{+9}_{-8}$	$^{+9}_{-9}$	$^{+8}_{-8}$

Table 5.46: The factorization scale uncertainties for the signal samples.

at 3,962 GeV, 3,988 GeV, and 4,014 GeV for each signal sample. The fractional differences in the cross-sections, listed in Table 5.47, were then treated as additional systematics on the nominal samples used in the analysis.

Mass (GeV)	C1	D1	D9
10	$^{+3}_{-0}$	$^{+2}_{-0}$	$^{+6}_{-1}$
50	$^{+3}_{-0}$	$^{+1}_{-0}$	$^{+6}_{-0}$
100	$^{+2}_{-0}$	$^{+1}_{-0}$	$^{+6}_{-1}$
200	$^{+2}_{-1}$	± 0	$^{+6}_{-1}$
400	± 0	$^{+2}_{-1}$	$^{+7}_{-2}$
700	$^{+3}_{-2}$	$^{+4}_{-2}$	$^{+11}_{-4}$
1000	$^{+5}_{-3}$	$^{+6}_{-3}$	$^{+14}_{-5}$
1300	$^{+6}_{-4}$	$^{+6}_{-4}$	$^{+16}_{-6}$

Table 5.47: The beam energy uncertainties for the signal samples.

PDF+α_s The choice of PDF and strong coupling, α_s, in the MC generation can have a significant impact on an analysis, introducing further systematic uncertainties. This is especially true for our signal samples, produced with the CTEQ6L1 PDF set, which is known to perform poorly in events with large momentum transfers [47]. The official

recommendation for ATLAS analyses comes from the PDF4LHC group [132], which gives a recipe for estimating the uncertainty on NLO PDFs. CTEQ6L1 is a LO PDF, so for the EFT samples some modifications had to be made to get reasonable results.

The PDF4LHC prescription is to use a combination of three NLO PDFs, NNPDF2.3 [133], CT10 [134], and MSTW2008 [134], all of which are fit to a large collection of data, from experiments including HERA [135] and the Tevatron [75]. The groups responsible for these have provided sets of varied PDFs for estimating each of the individual errors associated with the PDF fit and the α_s value used. These are combined using each groups recommendations, and then the envelope of these three errors is used as the total PDF+α_s uncertainty. The combined error can be defined by the envelope,

$$E^+ = \max_s\{\mathcal{O}_s^0 + \sigma_s^+\}, \tag{5.7}$$

$$E^- = \min_s\{\mathcal{O}_s^0 - \sigma_s^-\}, \tag{5.8}$$

where s is an index running over the three PDF sets, $\mathcal{O}_s^0$ is the central value for the observable of interest using PDF set s, and σ_s is its uncertainty. The error band is then scaled to the nominal value used by the analysis, $\mathcal{O}_0$, to give a final systematic error of,

$$\sigma_{\text{PDF}+\alpha_s} \equiv 2\mathcal{O}_0\frac{E^+ - E^-}{E^+ + E^-}. \tag{5.9}$$

The errors are all defined by 68% confidence intervals, and the uncertainty of α_s is fixed at 0.0012.

To calculate the CT10 uncertainty, 47 PDFs are used. 44 of these correspond to up/down variations of PDF parameters at a 90% confidence level, 2 to a variation of $\alpha_s \pm 0.0012$, and one used to calculate $\mathcal{O}^0_{\text{CT10}}$. The recommended PDF and α_s uncertainty are then given by,

$$\sigma^{\pm}_{\text{PDF}} = 0.608\sqrt{\sum_i^{22}\max\{\pm\mathcal{O}_{2i-1} \mp \mathcal{O}_0, \pm\mathcal{O}_{2i} \mp \mathcal{O}_0\}} \tag{5.10}$$

$$\sigma^{\pm}_{\alpha_s} = \frac{6}{5}\left(\mathcal{O}^{\pm}_{\alpha_s} - \mathcal{O}_0\right) \tag{5.11}$$

where $\mathcal{O}_i$ is the value of $\mathcal{O}$ using PDF i, $\mathcal{O}^{\pm}_{\alpha_s}$ is the value using $\alpha_s = 0.118 \pm 0.0012$, and the two correction factors are to bring the uncertainties to 68% confidence. The two upper errors and two lower errors can be added in quadrature to obtain the overall asymmetric errors $\sigma^{\pm}_{\mathrm{CT10}}$. The MSTW error band is calculated in an almost identical way, except that there are only 40 PDF variations and they're given at 68% confidence. The downward variation on α_s, σ^{-}_{α}, is given at 79% confidence, so it's given a factor of 0.8, while σ^{+}_{α} receives no correction.

The NNPDF uncertainty is calculated very differently. *Replicas* of the PDF fit are produced at different α_s scales, spaced apart by 0.001, and this is treated as a statistical ensemble. To get the 68% confidence interval, replicas need to be selected in a Gaussian manner around the central α_s value of 0.119. Here we use a total of 100 replicas, with (2, 8, 24, 32, 24, 8, 2) replicas chosen from the α_s bins (0.116, 0.117, 0.118, 0.119, 0.120, 0.121, 0.122) respectively. The total σ_{NNPDF} uncertainty is then simply given by the standard deviation of $\mathcal{O}$ within the chosen population of replicas.

In principle, the entire analysis, starting with the MC simulation, should be repeated for each PDF variation. Considering that there are 190 variations used here, this is an intractable task. Therefore, an approximation known as *PDF reweighting* is used to estimate the uncertainty. A tool called LHAPDF [136] has been developed to calculate a weight for each initial state parton as a function of its momentum fraction x, the energy transfer Q, and particle type. The product of the weights for both partons in an event, w_{PDF}, are representative of the probability of that event occurring with a given PDF. Each event is then given a SF of w_{PDF}/w_0, where w_0 is the weight corresponding to the nominal PDF used to generate the sample. In most cases this provides a reasonable alternative to regenerating the sample with a different PDF, and is much less computationally intensive.

One technicality that has not been mentioned is the change in cross-section of a process under a new PDF. This mixes uncertainty in the theoretical cross-section into the PDF

uncertainty and can very significantly overestimate them. The weighting procedure does not preserve the cross-section of the process, so each sample is normalized to match the nominal sample *before any selection cuts*. This effect is even more pronounced when we're using NLO PDFs to estimate the uncertainty of LO ones, since any k-factors associated with the NLO cross-section would propagate to the reweighted samples. To avoid this effect and estimate *only* the uncertainty due to the choice of PDF and α_s, the *width* of the envelope is used and the central value is discarded. The PDF uncertainty calculations can be visualized, as in in Fig. 5.64, and the breakdown of the uncertainties for the signal samples in the SR are listed in Tables 5.48, 5.49, and 5.50.

Mass (GeV)	CT10	MSTW	NNPDF	PDF4LHC
10	$^{+10\%}_{-8\%}$	$\pm 5\%$	$\pm 6\%$	$\pm 16\%$
50	$^{+15\%}_{-8\%}$	$\pm 5\%$	$\pm 6\%$	$\pm 18\%$
100	$^{+10\%}_{-6\%}$	$\pm 4\%$	$\pm 5\%$	$\pm 13\%$
200	$^{+8\%}_{-6\%}$	$\pm 3\%$	$\pm 5\%$	$\pm 13\%$
400	$^{+12\%}_{-6\%}$	$\pm 3\%$	$\pm 5\%$	$\pm 16\%$
700	$^{+10\%}_{-6\%}$	$\pm 2\%$	$\pm 6\%$	$\pm 15\%$
1000	$^{+14\%}_{-6\%}$	$^{+2\%}_{-1\%}$	$\pm 8\%$	$\pm 18\%$
1300	$^{+12\%}_{-5\%}$	$\pm 1\%$	$\pm 9\%$	$\pm 17\%$

Table 5.48: The PDF+α_s uncertainties for the D1 samples, for each PDF error set individually and the combination using the PDF4LHC prescription.

In addition to the PDF+α_s uncertainties on the signal samples, we can repeat this procedure on each of the background samples to get an estimate of their generator-level systematics. Table 5.51 shows the calculated uncertainties on the *simulated* samples in the seven selection regions. For the data-driven $Z \to \nu\nu$ estimate, detailed in Appendix B, some care has to be taken. The production of Z^0 bosons and γ-rays is very similar, but their ratio is highly flavor-dependent. For couplings to up-type quarks, the Z/γ ratio is naively about 0.93, while for down-type quarks it's 4.80. Since the PDF determines the flavor structure of the initial state protons, any mismodeling could have a huge impact on the final result.

Mass (GeV)	CT10	MSTW	NNPDF	PDF4LHC
10	$^{+10\%}_{-8\%}$	$\pm5\%$	$\pm5\%$	$\pm12\%$
50	$^{+9\%}_{-7\%}$	$\pm5\%$	$\pm5\%$	$\pm12\%$
100	$^{+9\%}_{-6\%}$	$\pm4\%$	$\pm4\%$	$\pm11\%$
200	$^{+9\%}_{-6\%}$	$\pm3\%$	$\pm4\%$	$\pm12\%$
400	$^{+10\%}_{-6\%}$	$\pm3\%$	$\pm5\%$	$\pm14\%$
700	$^{+7\%}_{-4\%}$	$\pm2\%$	$\pm5\%$	$\pm11\%$
1000	$^{+8\%}_{-5\%}$	$\pm1\%$	$\pm8\%$	$\pm15\%$
1300	$^{+11\%}_{-6\%}$	$\pm1\%$	$\pm11\%$	$\pm18\%$

Table 5.49: The PDF+α_s uncertainties for the C1 samples, for each PDF error set individually and the combination using the PDF4LHC prescription.

To estimate this uncertainty we follow the same LHAPDF procedure, except that the final observable is the predicted yield *after* the TF has been applied. This means that for each systematic variation, we need to recalculate the TF that is applied to the $\gamma+$jets data. The TFs are simply the ratio of the simulated $Z \to \nu\nu$ to $\gamma +$ jets yields in any region, varied together by the systematic. In Fig. 5.65 the inputs to the transfer function, as well as the resulting ratio are plotted. As it turns out, the systematics are heavily correlated between the two samples and the final uncertainty is significantly reduced.

Jet Smearing There are three systematics related to the jet smearing method used for the multijet estimate described in Appendix C. The largest of these is the uncertainty on the non-Gaussian tails of the jet response function, which is an $\mathcal{O}(100\%)$ effect. This systematic can be estimated by repeating the analysis with different jet response functions, corresponding to up and down variations of the fit. There is also an uncertainty on the Gaussian part of the response function, but this is much smaller because the high $\not\!\!E_T$ region is dominated by the rare tail events. This systematic is estimated by performing a secondary smearing of each jet, similar to other resolution uncertainties. The smallest systematic is to account for bias introduced by the seed selection, which shows up as mismodeling of the leading jet p_T in the smeared pseudo-data, and is corrected for by introducing a $\pm5\%$ shift

Mass (GeV)	CT10	MSTW	NNPDF	PDF4LHC
10	$^{+5\%}_{-4\%}$	$^{+3\%}_{-2\%}$	$\pm 3\%$	$\pm 8\%$
50	$^{+4\%}_{-3\%}$	$\pm 2\%$	$\pm 2\%$	$\pm 6\%$
100	$^{+4\%}_{-3\%}$	$\pm 2\%$	$\pm 2\%$	$\pm 6\%$
200	$^{+3\%}_{-3\%}$	$^{+3\%}_{-2\%}$	$\pm 2\%$	$\pm 4\%$
400	$^{+5\%}_{-5\%}$	$^{+4\%}_{-3\%}$	$\pm 3\%$	$\pm 6\%$
700	$^{+8\%}_{-7\%}$	$^{+4\%}_{-3\%}$	$\pm 5\%$	$\pm 10\%$
1000	$^{+9\%}_{-7\%}$	$\pm 4\%$	$\pm 9\%$	$\pm 13\%$
1300	$^{+8\%}_{-10\%}$	$\pm 3\%$	$\pm 12\%$	$\pm 15\%$

Table 5.50: The PDF+α_s uncertainties for the D9 samples, for each PDF error set individually and the combination using the PDF4LHC prescription.

Sample	SR	CR**znn**	CR**wjets**	CR**top**	CR**zmm**	CR**gamma**	CR**multijet**
$Z^0 \to \nu\bar{\nu}$	$\pm 5\%$	$\pm 2\%$					$\pm 5\%$
$Z^0 \to \mu^{\pm}\mu^{\mp}$	$\pm 9\%$	$\pm 7\%$	$\pm 8\%$	$\pm 7\%$	$\pm 8\%$	$\pm 7\%$	$\pm 10\%$
$W^{\pm} \to \ell^{\pm}\nu$	$\pm 4\%$	$\pm 3\%$			$\pm 0\%$	$\pm 7\%$	$\pm 5\%$
$\gamma +$ jets						$\pm 3\%$	$\pm 11\%$
$t\bar{t}$	$\pm 5\%$	$\pm 9\%$	$\pm 6\%$	$\pm 1\%$	$\pm 1\%$	$\pm 6\%$	$\pm 5\%$
single t	$\pm 37\%$	$\pm 37\%$	$\pm 38\%$	$\pm 36\%$	$\pm 37\%$	$\pm 44\%$	$\pm 36\%$
VV	$\pm 3\%$	$\pm 3\%$	$\pm 3\%$	$\pm 3\%$	$\pm 4\%$	$\pm 5\%$	$\pm 3\%$
Multijet	$\pm 6\%$	$\pm 6\%$					$\pm 7\%$

Table 5.51: The PDF+α_s uncertainties for the simulated backgrounds in the SR and CRs. Note that the $Z \to \nu\nu$ and multijet systematics here are for the *simulated* samples. The data-driven estimates have much lower uncertainties.

to the mean p_T of every smeared jet [137]. The fractional uncertainties in CR**multijet** and the SR are plotted in Figs. 5.66 and 5.67, for both the total background estimate and the multijet estimate alone.

5.7.5 Combined Uncertainties

To estimate the combined systematic effect, every source of systematic uncertainty is treated as fully correlated across each sample, and fully uncorrelated with the other systematics. This translates to summing the background estimates under each systematic

variation, and adding the resulting uncertainties in quadrature. Both the individual and combined asymmetric uncertainties on the SM backgrounds are listed in Table 5.52. The same is done for all of the signal samples, in Tables 5.53, 5.54 and 5.55.

Systematic	Multijet	VV	Single t	$t\bar{t}$	$Z \to \ell\ell$	$W \to \ell\nu$	$Z \to \nu\nu$	Total
b Flavor		$^{+3.6}_{-3.6}$	$^{+6.0}_{-6.1}$	$^{+5.6}_{-5.7}$	$^{+3.9}_{-4.0}$	$^{+2.9}_{-2.9}$	$^{+0.8}_{-0.9}$	$^{+2.1}_{-2.1}$
c Flavor		$^{+7.0}_{-7.0}$	$^{+0.8}_{-0.8}$	$^{+0.7}_{-0.7}$	$^{+6.4}_{-6.4}$	$^{+5.4}_{-5.5}$	$^{-1.7}_{+1.9}$	$^{+0.3}_{-0.2}$
Light Flavor		$^{+1.8}_{-1.8}$	$^{+0.2}_{-0.2}$	$^{+0.0}_{-0.0}$	$^{+0.1}_{-0.1}$	$^{+5.0}_{-5.0}$	$^{-0.4}_{+0.4}$	$^{+0.7}_{-0.7}$
JES		$^{+11.2}_{-5.3}$	$^{+3.9}_{-4.5}$	$^{+5.7}_{-8.2}$	$^{+6.7}_{-11.3}$	$^{+9.7}_{-6.7}$	$^{-0.4}_{+0.4}$	$^{+2.8}_{-2.4}$
JER		± 0.0	± 3.0	± 1.3	± 29.2	± 0.6	± 0.2	± 0.0
Soft Term		$^{-0.0}_{-0.7}$	$^{+0.7}_{+0.8}$	$^{+0.1}_{-0.1}$	$^{+0.0}_{+0.0}$	$^{+0.2}_{-0.4}$	$^{+0.8}_{+0.5}$	$^{+0.6}_{+0.2}$
Soft Resolution		± 0.8	± 1.0	± 0.3	± 0.0	± 0.1	± 0.0	± 0.0
Pileup		$^{+0.7}_{-1.0}$	$^{-1.1}_{+1.5}$	$^{-0.2}_{+0.1}$	$^{+0.7}_{+0.2}$	$^{+0.5}_{-0.5}$	$^{+0.1}_{+0.0}$	$^{+0.1}_{-0.1}$
JVF		$^{+0.4}_{+0.0}$	$^{+1.2}_{+0.0}$	$^{+0.8}_{-0.5}$	$^{+4.4}_{-0.0}$	$^{+0.2}_{-0.2}$	$^{+0.1}_{-0.4}$	$^{+0.2}_{-0.3}$
PDF		± 8.9	± 39.2	± 5.3	± 13.8	± 5.1	± 2.3	± 4.6
Normalization	± 4.4	± 8.7	± 12.3	± 6.7	± 6.6	± 1.2	± 2.0	± 1.6
Data-Driven	$^{+75.7}_{-9.5}$						$^{+4.0}_{-4.0}$	$^{+2.5}_{-2.5}$
$\sum \sigma^+$	$+75.8\%$	$+18.6\%$	$+41.9\%$	$+11.8\%$	$+34.8\%$	$+13.6\%$	$+5.6\%$	$+6.6\%$
$\sum \sigma^-$	-10.5%	-15.8%	-41.9%	-13.2%	-35.7%	-11.7%	-5.4%	-6.4%

Table 5.52: Breakdown of systematic uncertainties on the background estimates in SR**300**.

	D1 Mass (GeV)							
Systematic	10	50	100	200	400	700	1000	1300
b Flavor	$^{+7.4}_{-7.4}$	$^{+6.8}_{-6.9}$	$^{+7.1}_{-7.2}$	$^{+7.4}_{-7.4}$	$^{+7.7}_{-7.7}$	$^{+7.2}_{-7.3}$	$^{+7.5}_{-7.6}$	$^{+7.4}_{-7.5}$
c Flavor	$^{-0.0}_{+0.0}$	$^{+0.1}_{-0.1}$	$^{+0.1}_{-0.1}$	$^{+0.0}_{-0.0}$	$^{-0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{-0.0}_{+0.0}$	$^{-0.0}_{+0.0}$
Light Flavor	$^{-0.0}_{+0.0}$	$^{+0.3}_{-0.3}$	$^{-0.1}_{+0.1}$	$^{-0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{-0.1}_{+0.1}$	$^{-0.0}_{+0.0}$	$^{+0.0}_{-0.0}$
JES	$^{+6.2}_{-7.8}$	$^{+7.9}_{-4.5}$	$^{+4.4}_{-5.8}$	$^{+5.3}_{-4.4}$	$^{+4.1}_{-4.8}$	$^{+2.8}_{-3.3}$	$^{+2.3}_{-2.8}$	$^{+1.5}_{-3.7}$
JER	± 2.6	± 2.1	± 1.2	± 0.8	± 0.3	± 1.4	± 1.0	± 1.1
Soft Term	$^{-0.6}_{+0.0}$	$^{+0.2}_{-0.0}$	$^{+0.3}_{+0.2}$	$^{-0.0}_{-0.5}$	$^{+0.2}_{-0.1}$	$^{+0.1}_{+0.1}$	$^{-0.1}_{+0.1}$	$^{-0.3}_{-0.1}$
Soft Resolution	± 0.4	± 0.0	± 0.0	± 0.7	± 0.0	± 0.0	± 0.3	± 0.0
Pileup	$^{-0.4}_{+0.3}$	$^{+0.1}_{+0.1}$	$^{-0.3}_{+0.4}$	$^{-0.3}_{+0.5}$	$^{+0.6}_{-0.6}$	$^{+0.6}_{-0.4}$	$^{+0.1}_{-0.1}$	$^{+0.8}_{-0.8}$
JVF	$^{+0.3}_{-0.3}$	$^{+0.0}_{-0.0}$	$^{-0.0}_{-0.0}$	$^{+0.4}_{-0.0}$	$^{+0.1}_{-0.3}$	$^{+0.4}_{-0.2}$	$^{+0.2}_{-0.2}$	$^{+0.2}_{-0.1}$
PDF	± 14.7	± 17.4	± 13.1	± 13.5	± 14.9	± 16.2	± 18.9	± 25.8
Beam Energy	$^{+3.8}_{-3.8}$	$^{+3.6}_{-3.6}$	$^{+4.0}_{-3.7}$	$^{+4.6}_{-4.0}$	$^{+5.3}_{-4.9}$	$^{+6.7}_{-6.3}$	$^{+8.0}_{-7.7}$	$^{+9.5}_{-9.1}$
Factorization	$^{+1.6}_{-0.3}$	$^{+1.5}_{-0.4}$	$^{+1.3}_{+0.0}$	$^{+0.2}_{-0.2}$	$^{+2.0}_{-0.8}$	$^{+4.0}_{-2.1}$	$^{+5.7}_{-3.1}$	$^{+6.5}_{-3.7}$
Normalization	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8
$\sum \sigma^+$	+18.5%	+21.0%	+16.4%	+17.2%	+18.4%	+19.8%	+22.9%	+29.4%
$\sum \sigma^-$	−19.1%	−19.9%	−16.8%	−16.8%	−18.4%	−19.5%	−22.3%	−29.0%

Table 5.53: Summary of systematic uncertainties on the D1 samples in SR**300**.

	D9 Mass (GeV)							
Systematic	10	50	100	200	400	700	1000	1300
b Flavor	$^{+6.6}_{-6.7}$	$^{+6.1}_{-6.3}$	$^{+6.4}_{-6.5}$	$^{+6.0}_{-6.2}$	$^{+5.9}_{-6.1}$	$^{+4.8}_{-5.2}$	$^{+4.7}_{-5.1}$	$^{+4.9}_{-5.2}$
c Flavor	$^{+0.1}_{-0.1}$	$^{+0.0}_{+0.0}$	$^{+0.2}_{-0.2}$	$^{+0.1}_{-0.1}$	$^{+0.0}_{-0.0}$	$^{+0.0}_{-0.0}$	$^{+0.0}_{-0.0}$	$^{-0.0}_{+0.0}$
Light Flavor	$^{+0.0}_{-0.0}$	$^{-0.0}_{+0.0}$	$^{+0.1}_{-0.1}$	$^{-0.1}_{+0.1}$	$^{+0.0}_{-0.0}$	$^{-0.1}_{+0.1}$	$^{-0.1}_{+0.1}$	$^{-0.1}_{+0.1}$
JES	$^{+3.1}_{-4.3}$	$^{+2.1}_{-5.2}$	$^{+2.7}_{-4.4}$	$^{+2.6}_{-3.1}$	$^{+1.5}_{-2.2}$	$^{+1.3}_{-1.8}$	$^{+2.1}_{-2.0}$	$^{+2.1}_{-1.7}$
JER	± 1.1	± 0.3	± 1.4	± 0.6	± 0.1	± 0.8	± 0.5	± 0.2
Soft Term	$^{+0.3}_{-0.5}$	$^{-0.0}_{-0.3}$	$^{+0.2}_{-0.4}$	$^{-0.3}_{-0.2}$	$^{+0.0}_{-0.0}$	$^{-0.2}_{+0.1}$	$^{+0.1}_{+0.1}$	$^{+0.1}_{+0.0}$
Soft Resolution	± 0.3	± 0.0	± 0.2	± 0.3	± 0.1	± 0.3	± 0.1	± 0.1
Pileup	$^{+0.9}_{-1.5}$	$^{+0.8}_{-0.8}$	$^{+7.0}_{-7.5}$	$^{+0.3}_{-0.1}$	$^{+0.4}_{-0.6}$	$^{+0.0}_{+0.1}$	$^{+0.1}_{-0.3}$	$^{-0.1}_{+0.0}$
JVF	$^{+0.1}_{-0.2}$	$^{+0.4}_{-0.1}$	$^{-0.0}_{-0.2}$	$^{+0.2}_{-0.1}$	$^{+0.2}_{-0.1}$	$^{+0.1}_{-0.2}$	$^{+0.2}_{-0.2}$	$^{+0.1}_{-0.1}$
PDF	± 8.5	± 8.8	± 10.0	± 6.9	± 10.7	± 15.6	± 23.2	± 26.5
Beam Energy	$^{+3.1}_{-3.1}$	$^{+3.3}_{-3.5}$	$^{+2.9}_{-3.5}$	$^{+3.8}_{-3.7}$	$^{+4.3}_{-4.2}$	$^{+5.4}_{-6.3}$	$^{+6.4}_{-6.3}$	$^{+8.1}_{-7.8}$
Factorization	$^{+6.3}_{-0.6}$	$^{+6.0}_{-0.5}$	$^{+5.5}_{-0.8}$	$^{+5.8}_{-1.3}$	$^{+6.9}_{-2.4}$	$^{+10.6}_{-3.6}$	$^{+13.7}_{-4.7}$	$^{+16.0}_{-5.7}$
Normalization	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8
$\sum \sigma^+$	+13.6%	+13.2%	+15.7%	+12.1%	+15.0%	+20.5%	+28.4%	+32.5%
$\sum \sigma^-$	−12.6%	−12.9%	−15.5%	−10.9%	−13.7%	−18.3%	−25.3%	−28.9%

Table 5.54: Summary of systematic uncertainties on the D9 samples in SR**300**.

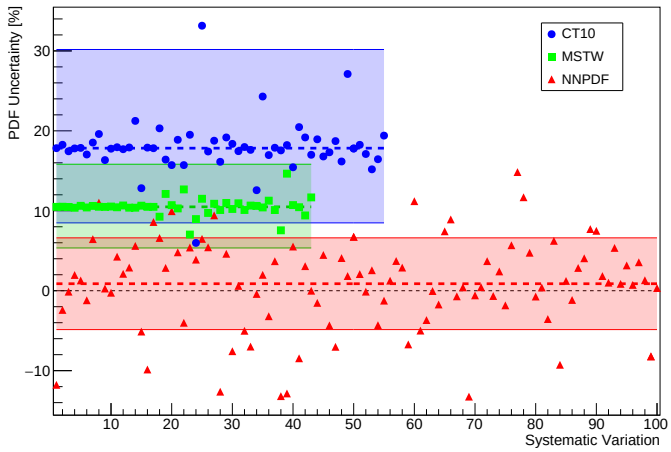

(a) D1, 10 GeV

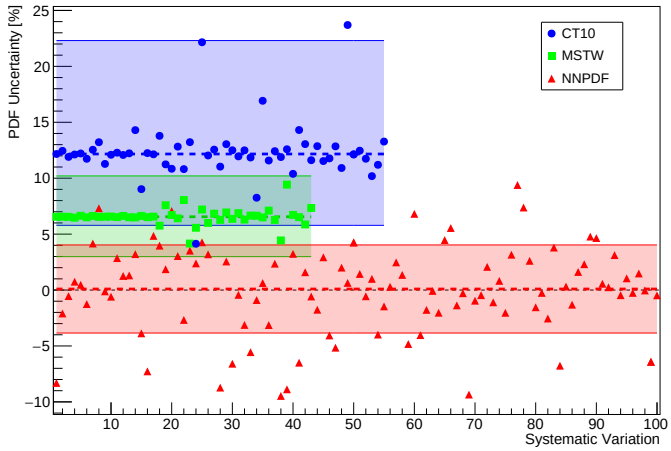

(b) C1, 200 GeV

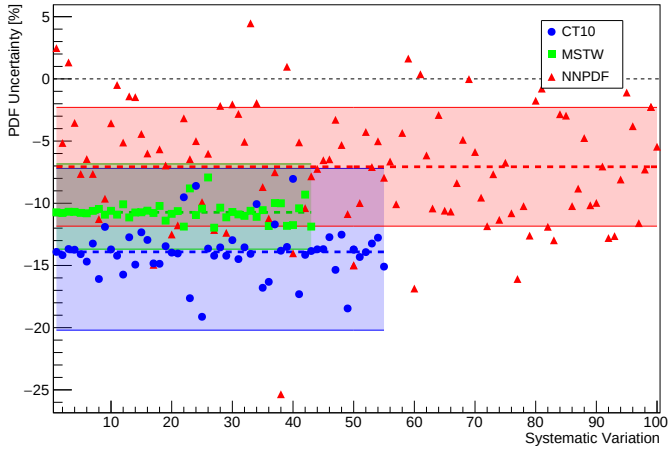

(c) D9, 700 GeV

Figure 5.64: PDF uncertainties in SR**300** for the three signal operators at various m_χ. Each point corresponds to the predicted yield using a different PDF, and the three bands correspond to the individual PDF set errors. The total PDF error is estimated as the envelope around all three bands.

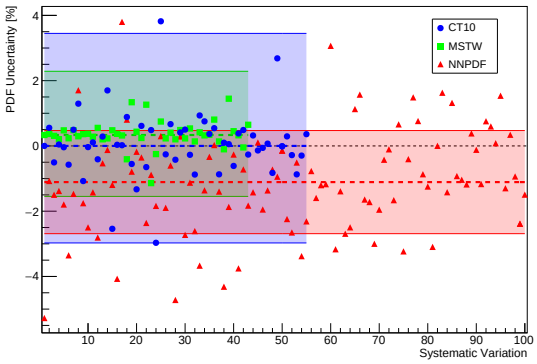

(a) $\gamma + \text{jets}$

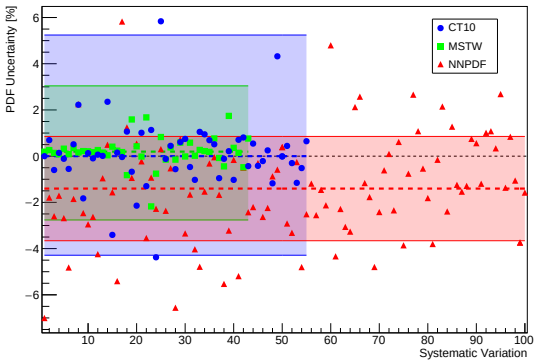

(b) $Z \to \nu\nu$ Sherpa

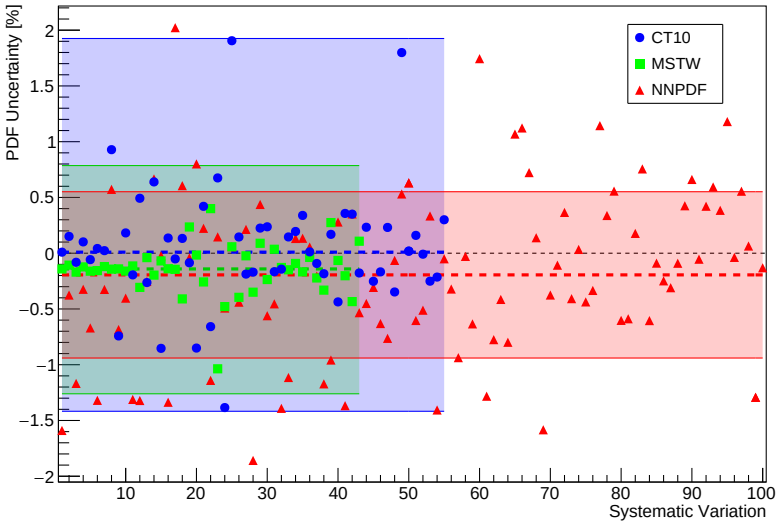

(c) $Z \to \nu\nu$ from γ

Figure 5.65: PDF uncertainties in SR**300** for the final data-driven $Z \to \nu\nu$ estimate and the two samples used to derive the transfer function. Each point corresponds to the predicted yield using a different PDF, and the three bands correspond to the individual PDF set errors. The total PDF error is estimated as the envelope around all three bands.

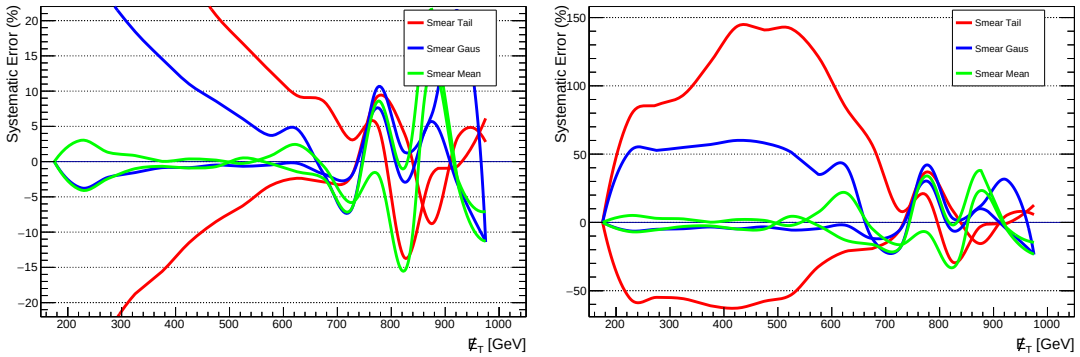

(a) Total Background **(b)** Multijet

Figure 5.66: Fractional systematics for the jet smearing uncertainties after the b-jet p_T cut of CR**multijet**, as functions of $\not{E}_T$.

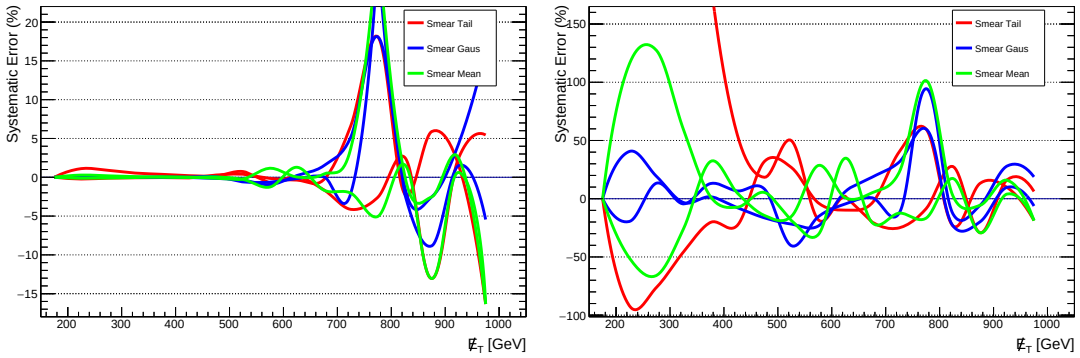

(a) Total Background **(b)** Multijet

Figure 5.67: Fractional systematics for the jet smearing uncertainties after the b-jet p_T cut, as functions of $\not{E}_T$.

	C1 Mass (GeV)							
Systematic	10	50	100	200	400	700	1000	1300
b Flavor	$^{+6.8}_{-7.0}$	$^{+7.5}_{-7.5}$	$^{+7.6}_{-7.6}$	$^{+7.5}_{-7.5}$	$^{+7.4}_{-7.5}$	$^{+7.1}_{-7.2}$	$^{+7.2}_{-7.3}$	$^{+7.6}_{-7.6}$
c Flavor	$^{-0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{+0.1}_{-0.1}$	$^{-0.0}_{+0.0}$	$^{-0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{+0.0}_{-0.0}$	$^{+0.1}_{-0.1}$
Light Flavor	$^{-0.1}_{+0.1}$	$^{-0.1}_{+0.1}$	$^{-0.1}_{+0.1}$	$^{-0.0}_{+0.0}$	$^{-0.1}_{+0.1}$	$^{+0.0}_{-0.0}$	$^{-0.0}_{+0.0}$	$^{-0.1}_{+0.1}$
JES	$^{+6.2}_{-7.7}$	$^{+6.8}_{-3.3}$	$^{+5.1}_{-4.6}$	$^{+5.8}_{-3.7}$	$^{+5.8}_{-3.6}$	$^{+3.4}_{-2.5}$	$^{+2.1}_{-4.3}$	$^{+4.1}_{-2.3}$
JER	± 1.8	± 1.6	± 1.1	± 0.3	± 0.3	± 0.1	± 1.5	± 0.5
Soft Term	$^{+1.0}_{-0.1}$	$^{-0.2}_{+1.0}$	$^{+0.3}_{-0.3}$	$^{+0.4}_{+0.4}$	$^{+0.6}_{-0.2}$	$^{+0.1}_{-0.0}$	$^{+0.2}_{-0.4}$	$^{+0.3}_{+0.3}$
Soft Resolution	± 0.0	± 0.9	± 1.6	± 0.0	± 0.6	± 0.4	± 0.4	± 0.2
Pileup	$^{+0.6}_{-1.0}$	$^{+0.3}_{-0.1}$	$^{+0.2}_{-0.4}$	$^{+0.3}_{-0.1}$	$^{-0.1}_{-0.1}$	$^{+0.3}_{-0.5}$	$^{-0.3}_{+0.1}$	$^{+0.3}_{-0.5}$
JVF	$^{+0.0}_{+0.0}$	$^{-0.0}_{+0.0}$	$^{+0.2}_{-0.3}$	$^{+0.0}_{-0.2}$	$^{+0.1}_{+0.0}$	$^{+0.1}_{-0.1}$	$^{+0.1}_{-0.2}$	$^{+0.1}_{-0.1}$
PDF	± 12.3	± 12.6	± 10.8	± 12.3	± 14.5	± 12.9	± 16.8	± 23.1
Beam Energy	$^{+3.3}_{-2.9}$	$^{+3.2}_{-3.0}$	$^{+3.4}_{-3.4}$	$^{+3.5}_{-3.6}$	$^{+4.2}_{-4.2}$	$^{+5.9}_{-5.8}$	$^{+7.6}_{-6.9}$	$^{+8.7}_{-8.5}$
Factorization	$^{+3.3}_{+0.0}$	$^{+3.3}_{+0.0}$	$^{+2.4}_{-0.2}$	$^{+1.6}_{-1.1}$	$^{+0.4}_{+0.0}$	$^{+2.9}_{-1.7}$	$^{+4.9}_{-2.9}$	$^{+6.1}_{-3.6}$
Normalization	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8	± 2.8
$\sum \sigma^+$	+16.4%	+17.1%	+15.1%	+16.3%	+18.0%	+16.8%	+20.7%	+27.0%
$\sum \sigma^-$	−16.7%	−15.6%	−14.7%	−15.6%	−17.5%	−16.4%	−20.5%	−26.3%

Table 5.55: Summary of systematic uncertainties on the C1 samples in SR**300**.

5.8 Results

In the absence of a discovery, limits can still be set on M_* that will constrain any theory that meets the EFT and MFV assumptions. The standard procedure in ATLAS analyses is to calculate an upper bound on the signal strength parameter μ, using the CLs method [138], excluding regions of the signal parameter space at 90% or 95% confidence. The μ parameter is simply a normalization factor on the signal cross-section, meaning that for samples generated at $M_* = 1\,\text{TeV}$ it's related to the observed M_* by,

$$\mu = \left(\frac{1\,\text{TeV}}{M_*}\right)^y, \tag{5.12}$$

where y is 4 or 6 for scalar or fermionic WIMPs respectively. Upper limits on μ can then be easily converted to lower limits on M_*, the only parameter for our signals.

The standard tool used throughout ATLAS is known as HistFitter [122], which implements the CLs method as a flexible tool. HistFitter first constructs the likelihood function,

$$L\left(\boldsymbol{n}, \boldsymbol{\theta^0}|\mu, \boldsymbol{b}, \boldsymbol{\theta}\right) = P\left(n_{\text{SR}}|\mu, \boldsymbol{b}, \theta\right) \times \prod_{i\in\text{CR}} \times P\left(n_i|\mu, \boldsymbol{b}, \boldsymbol{\theta}\right) \prod_{j\in\text{sys}} P\left(\theta_j|\theta_j^0\right), \tag{5.13}$$

which represents the likelihood of the measured yields, $\boldsymbol{n}$, for a set of background predictions, $\boldsymbol{b}$, signal strength μ, and auxiliary measurements, $\boldsymbol{\theta}$. The auxiliary measurements, or nuisance parameters, are used to model the systematic uncertainties, with central values $\boldsymbol{\theta^0}$ and observed value $\boldsymbol{\theta}$. The probability distributions $P(\theta_j|\theta_j^0)$ are usually assumed to be Gaussian, with central values $\theta_j^0 = 0$ and widths corresponding to the corresponding systematics. Once this function is constructed, the profile *log likelihood ratio* (LLR) can be defined as,

$$q(\mu) = -2\text{log}\left(\frac{\text{max}_{\hat{\boldsymbol{\theta}}'}\, L\left(\boldsymbol{n}, \boldsymbol{\theta^0}|\mu, \boldsymbol{b}, \hat{\boldsymbol{\theta}}'\right)}{\text{max}_{\hat{\boldsymbol{\mu}}\geq\boldsymbol{0},\hat{\boldsymbol{\theta}}}\, L\left(\boldsymbol{n}, \boldsymbol{\theta^0}|\hat{\mu}, \boldsymbol{b}, \hat{\boldsymbol{\theta}}\right)}\right), \tag{5.14}$$

where $\hat{\boldsymbol{\theta}}'$ maximizes the likelihood for the given μ, and $\hat{\boldsymbol{\theta}}$ maximizes it for *any* $\hat{\mu} \geq 0$. This function $q(\mu)$ is used as a test statistic which, using Wilks' theorem, follows a χ^2

distribution with 1 degree of freedom for sufficiently large sample sizes. This can then be used to calculate a p-value for any given hypothesis. For discovery searches, the test statistic $q(0)$ is used, so that the p-value constrains the background model. For exclusion tests, $q(\mu)$ is calculated as a function of μ and the corresponding p-values are used to exclude possible values of μ. The standard at ATLAS is to use 95% confidence levels, meaning that any μ is excluded if $P(q(\mu)) < 0.05$. Dark matter searches typically use 90% confidence levels though, so both are calculated and listed in Tables 5.56 and 5.57. The SR chosen for each sample corresponds to whichever has the maximum *expected* sensitivity, since the analysis was blinded when SRs were being assigned.

Figs. 5.68-5.70 plot the 95% CLs limits on M_* for each of the six representative operators, D1, D1u, D9, D9u, C1, and C1u. The values of M_* and m_χ that reproduce the observed relic abundance of DM are calculated using the results of Appendix E.1, and validity constraints are plotted using the techniques discussed in Section 2.3.2. Shaded contours show the regions in which 99%, 95%, 68%, 38%, and 0% of the signal events meet the validity requirements, and the darkest shaded region corresponds to the minimal $M_\phi > 2m_\chi$ validity condition.

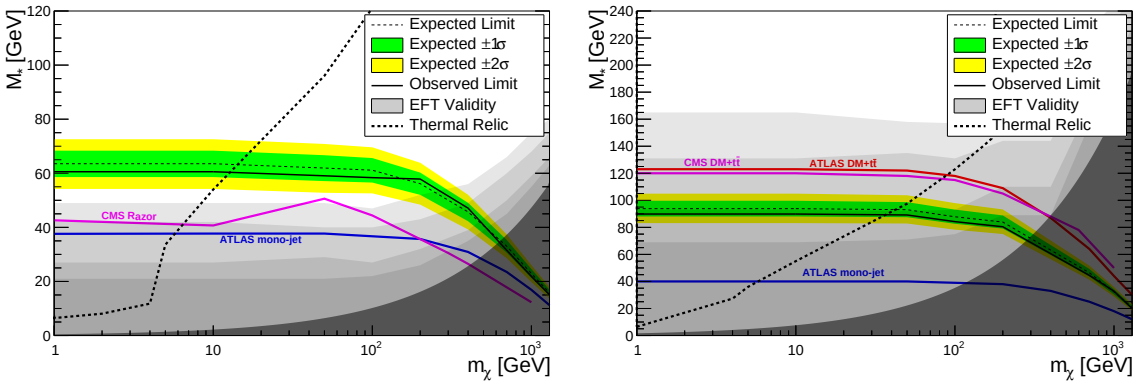

(a) D1 **(b)** D1u

Figure 5.68: 95% confidence limits on M_* for the D1 operators. Validity contours mark regions where 99%, 95%, 68%, 38%, and 0% of the signal events meet the EFT validity criteria. Results from other collider experiments are plotted for reference, using the techniques in Appendix E [22, 47, 139, 140].

Figs. 5.71-5.78 plot the 90% CLs limits on M_* for each of the signal operators, compared

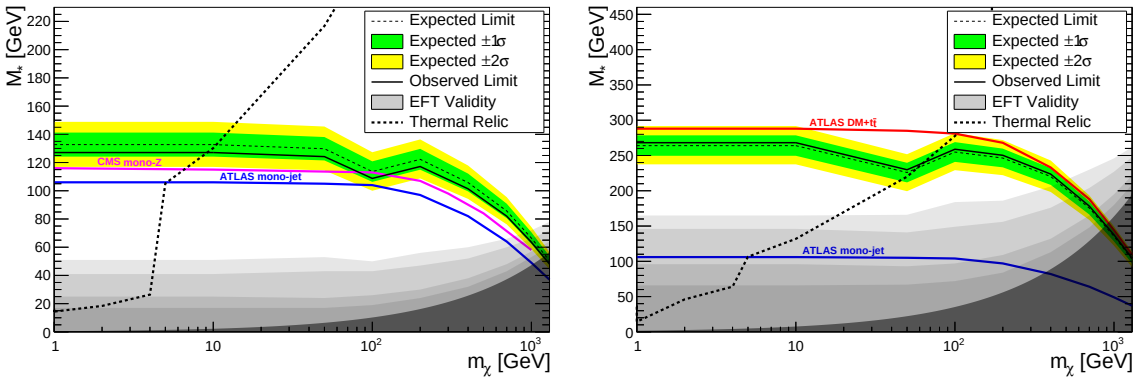

(a) D9 **(b)** D9u

Figure 5.69: 95% confidence limits on M_* for the D9 operators. Validity contours mark regions where 99%, 95%, 68%, 38%, and 0% of the signal events meet the EFT validity criteria. Results from other collider experiments are plotted for reference, using the techniques in Appendix E [22, 47, 141].

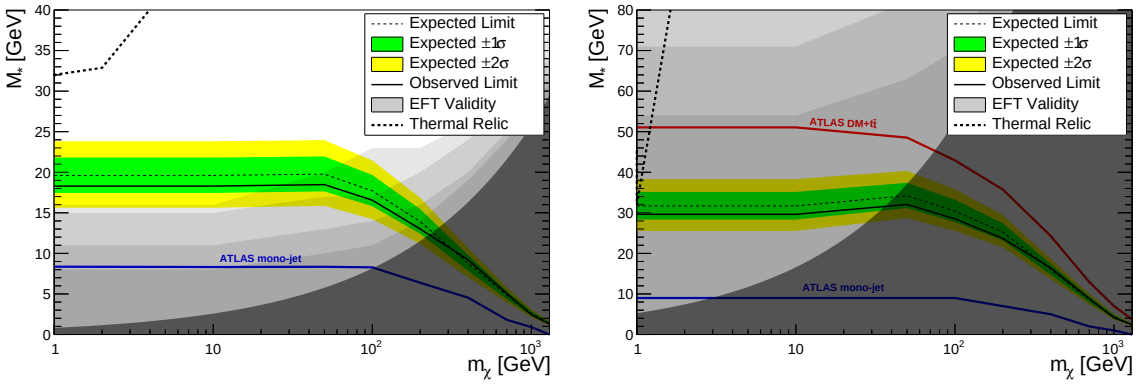

(a) C1 **(b)** C1u

Figure 5.70: 95% confidence limits on M_* for the C1 operators. Validity contours mark regions where 99%, 95%, 68%, 38%, and 0% of the signal events meet the EFT validity criteria. Results from other collider experiments are plotted for reference, using the techniques in Appendix E [22, 47].

with the results of a representative set of other dark matter searches. The M_* values of each operator that would reproduce the observed relic abundance is plotted again for each, and serves as an upper bound on M_*. Every other curve represents a lower bound on M_*, and correspond to various limits set by other experiments. The direct detection results for LUX [16], XENON100 [15], and CDMS [14] results are plotted for the spin-independent interactions, and COUPP [19] and PICASSO [21] for spin-dependent. The latest IceCube [13] and Fermi-LAT [12] results are also compared, along with the regions

that would reproduce a potential DM signal observed in the galactic center. The calculations used to transform the results of these different experiments to our parameter space of M_* and m_χ are detailed in Appendix E. This analysis sets the strongest limits at masses below about 5 GeV for all of the operators, where it's well within the validity region for down-type quark couplings. For universal couplings, the validity constraints are much tighter, where only D9u and D10u meet them at any of the mass points.

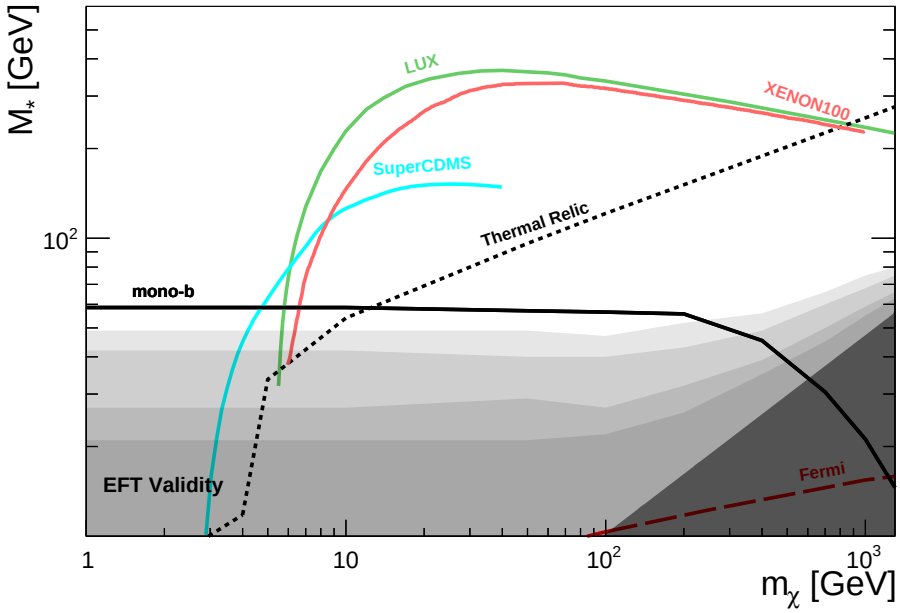

Figure 5.71: 90% confidence limits set on M_* for the D1 operator, compared to other experiments using the techniques in Appendix E [12, 14–16].

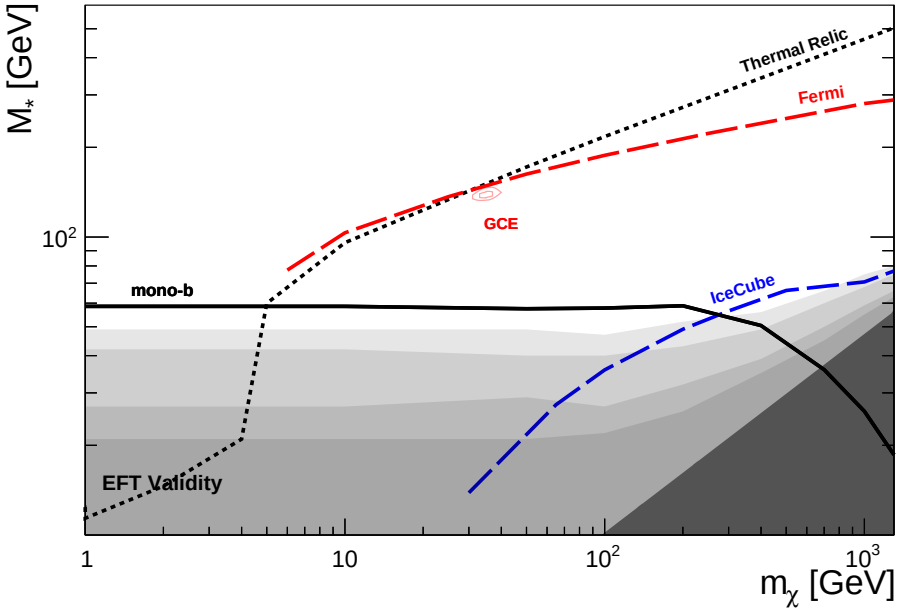

Figure 5.72: 90% confidence limits set on M_* for the D2 operator, compared to other experiments using the techniques in Appendix E [12, 13, 142].

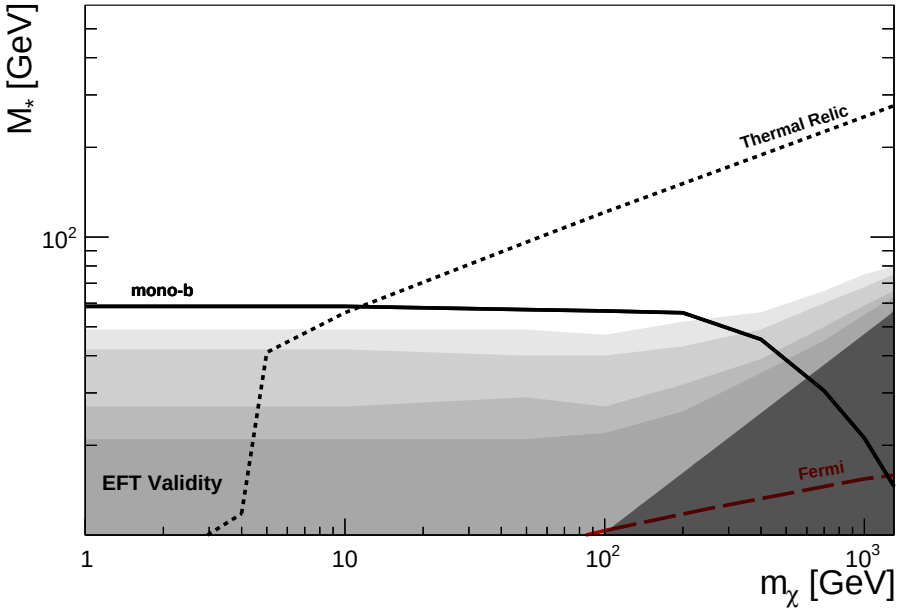

Figure 5.73: 90% confidence limits set on M_* for the D3 operator, compared to other experiments using the techniques in Appendix E [12].

Operator	m_χ	SR	M_*^{90}	M_*^{95}
D1	10	300	61 (64)	59 (61)
D1	50	300	59 (62)	57 (60)
D1	100	300	58 (61)	57 (59)
D1	200	400	58 (56)	56 (54)
D1	400	400	47 (46)	45 (44)
D1	700	450	32 (33)	30 (32)
D1	1000	450	22 (23)	21 (22)
D1	1300	450	15 (16)	15 (15)
D9	10	450	127 (133)	123 (128)
D9	50	450	124 (130)	120 (125)
D9	100	450	109 (114)	105 (110)
D9	200	450	117 (122)	114 (118)
D9	400	450	101 (106)	98 (102)
D9	700	450	82 (86)	79 (83)
D9	1000	450	64 (66)	62 (64)
D9	1300	450	48 (50)	47 (49)
C1	10	300	18 (20)	17 (19)
C1	50	300	18 (20)	18 (19)
C1	100	300	17 (18)	16 (17)
C1	200	300	13 (14)	12 (13)
C1	400	400	9 (9)	9 (8)
C1	700	400	5 (5)	5 (5)
C1	1000	450	2 (3)	2 (2)
C1	1300	450	1 (1)	1 (1)

Table 5.56: Observed (expected) M_* limits placed on the mono-b operators for each signal sample at 90% and 95% confidence.

Operator	m_χ	SR	M_*^{90}	M_*^{95}
D1u	10	450	90 (94)	87 (91)
D1u	50	450	89 (93)	87 (90)
D1u	100	450	84 (88)	82 (85)
D1u	200	450	80 (84)	78 (81)
D1u	700	450	44 (46)	43 (45)
D1u	1000	500	33 (32)	32 (31)
D1u	1300	450	20 (21)	19 (20)
D9u	10	500	268 (264)	259 (255)
D9u	50	500	230 (226)	221 (217)
D9u	100	500	259 (255)	250 (246)
D9u	200	500	250 (246)	242 (238)
D9u	400	500	224 (220)	216 (213)
D9u	700	500	178 (175)	172 (170)
D9u	1000	500	138 (135)	133 (131)
D9u	1300	500	104 (103)	101 (99)
C1u	10	300	30 (32)	28 (30)
C1u	50	450	32 (34)	31 (33)
C1u	100	450	29 (30)	30 (29)
C1u	200	450	24 (25)	23 (24)
C1u	400	500	16 (16)	16 (15)
C1u	700	500	9 (9)	8 (8)
C1u	1000	450	4 (4)	4 (4)
C1u	1300	500	2 (2)	2 (2)

Table 5.57: Observed (expected) M_* limits placed on the operators with universal couplings, for each signal sample at 90% and 95% confidence.

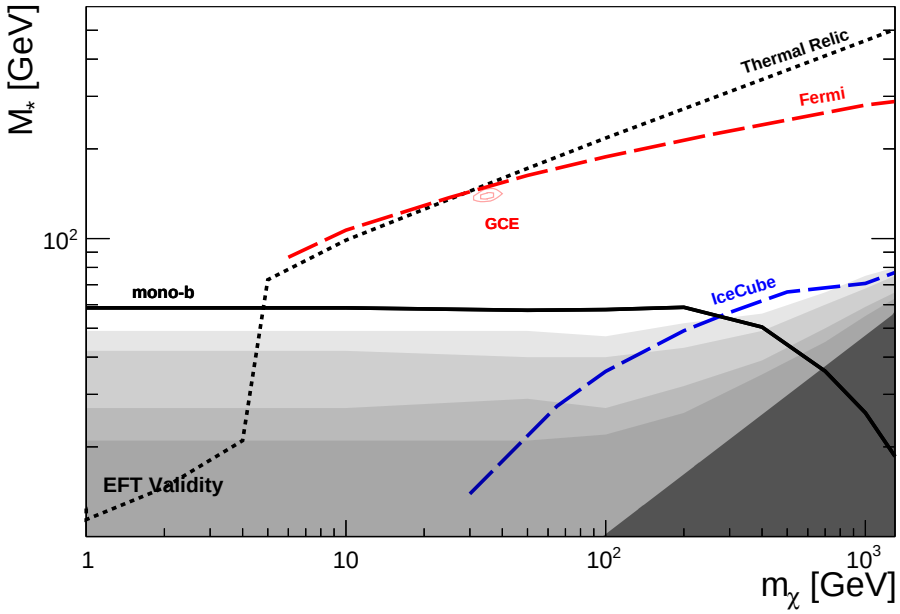

Figure 5.74: 90% confidence limits set on M_* for the D4 operator, compared to other experiments using the techniques in Appendix E [12, 13, 142].

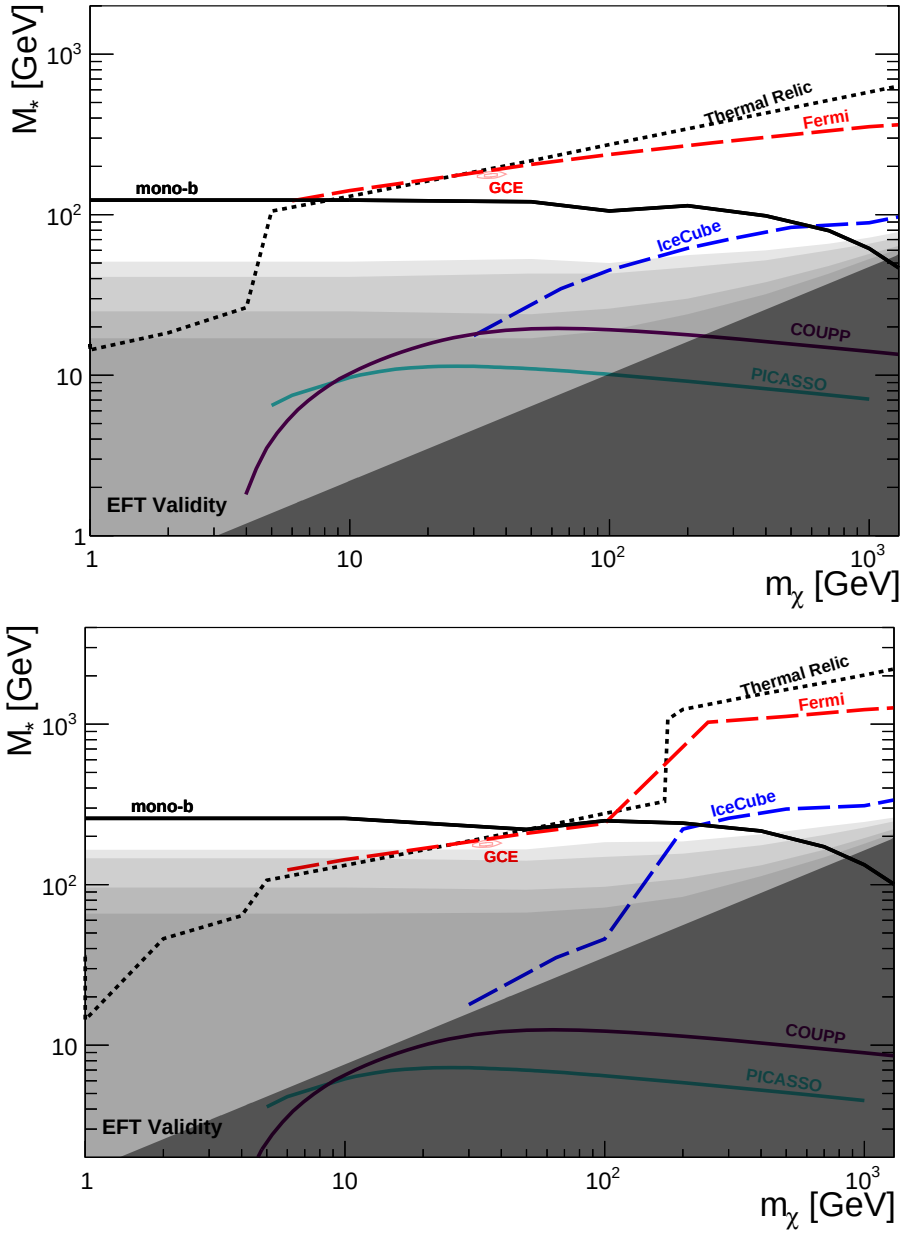

Figure 5.75: 90% confidence limits set on M_* for the D9 (top) and D9u (bottom) operators, compared to other experiments using the techniques in Appendix E [12, 13, 19, 21, 142].

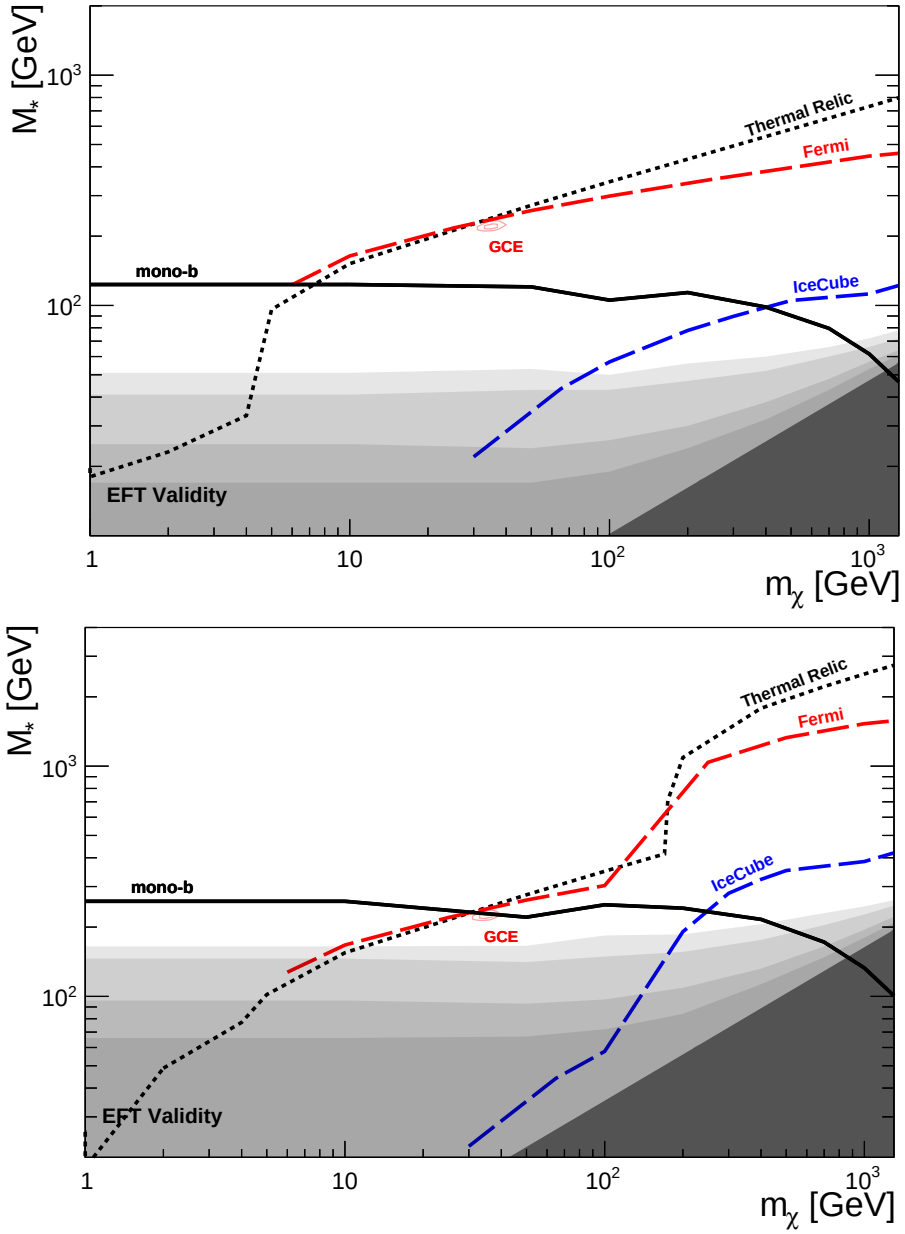

Figure 5.76: 90% confidence limits set on M_* for the D10 (top) and D10u (bottom) operators, compared to other experiments using the techniques in Appendix E [12, 13, 142].

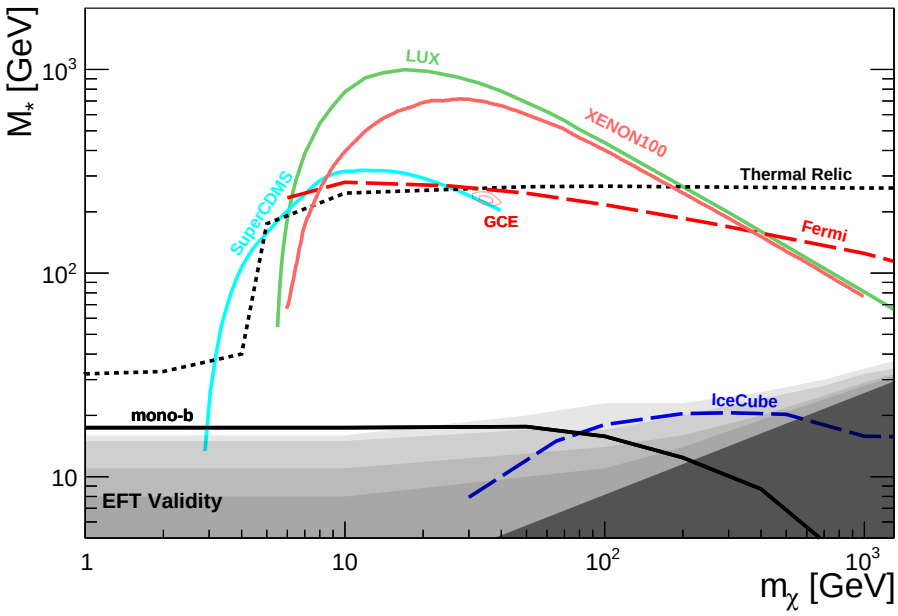

Figure 5.77: 90% confidence limits set on M_* for the C1 operator, compared to other experiments using the techniques in Appendix E [12–16, 142].

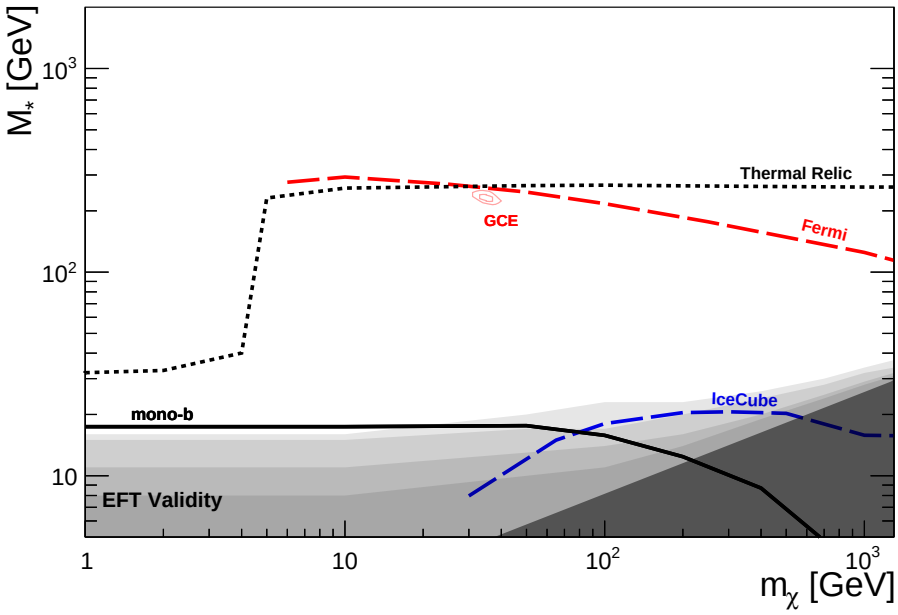

Figure 5.78: 90% confidence limits set on M_* for the C2 operator, compared to other experiments using the techniques in Appendix E [12, 13, 142].

Chapter 6

Conclusion

In this thesis I have laid out a search for a specific class of particles that are capable of explaining the cosmological phenomena attributed to dark matter. Using data collected by ATLAS, limits were set on generic WIMP-quark interactions that have mass dependent couplings. Within the EFT validity region, where the details of the heavy mediator become irrelevant, lower bounds were placed on the energy scale, M_*, for 16 different interactions. These limits are currently the strongest available below about 5 GeV.

Of the 16 operators considered, direct detection experiments only have enhanced sensitivity to the four spin-independent D1 and C1 interactions, and to a lesser degree, the two spin-dependent D9 interactions. On the other hand, indirect detection experiments provide strong limits on *all* operators down to about 5 GeV, with the exception of D1 and D3. The D3 interaction is unique, in that only collider searches are sensitive to it, and its mass dependence under MFV gives this analysis significantly better sensitivity than any search in the past. For the other interactions, this analysis provides complementary limits in the low mass region, where other search techniques lose sensitivity. Additionally, by taking cosmological measurements into account, thermally produced fermionic WIMPs can be completely ruled out as the sole constituent of dark matter, up to around 5 GeV–10 GeV, depending on the operator.

Although the processes searched for here are assumed to have mass dependent couplings, the large mass difference between quark flavors makes this well approximated by interactions with *only* bottom or top quarks. The signal samples were generated to reflect this, requiring the production of at least one bottom or two top quarks. Therefore, the results of this

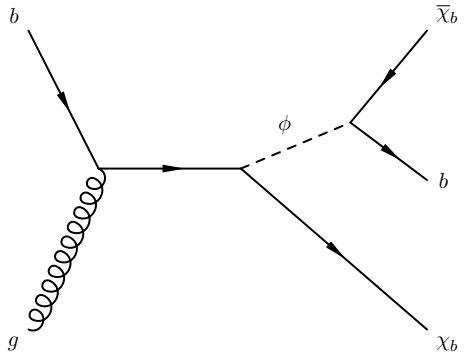

Figure 6.1: Example WIMP production in $b-$FDM [143].

analysis can be easily extended to generic flavor-dependent couplings between heavy quarks and WIMPs. A concrete example of this is the *bottom flavored dark matter* ($b-$FDM) theory, which adds a new fermionic WIMP for each quark flavor, and a scalar mediator with color charge [143]. The bottom flavored WIMP is chosen to be the lightest of these, with the dominant production mechanism shown in Fig. 6.1. The results of this analysis can then be applied to $b-$FDM to set limits on the masses of the two particles, m_ϕ and m_χ. An example of this is plotted in Fig. 6.2, where a region of the phase space has been ruled out at 95% confidence.

One recent hint at flavor-dependent couplings comes from the Fermi-LAT experiment. An excess of γ-rays was observed in the galactic center that shows up as the extended region in Fig. 6.3 [144]. This signal is consistent with WIMP annihilations into bottom pairs, suggesting that dark matter might have preferential couplings to bottom quarks. Although the recent Fermi-LAT survey of dwarf galaxies appears to rule this scenario out at 90% confidence [145], theories are still being proposed that can accommodate both results [146]. While ATLAS isn't sensitive to this excess yet, future analyses such as this one will be able to provide new insights. Run II has already begun at the LHC, and over the next 3 years is scheduled to produce over $100\,\text{fb}^{-1}$ of 13 TeV–14 TeV proton-proton collisions. This 75 % increase in energy, and 400 % increase in integrated luminosity will greatly extend the reach of searches for physics beyond the Standard Model.

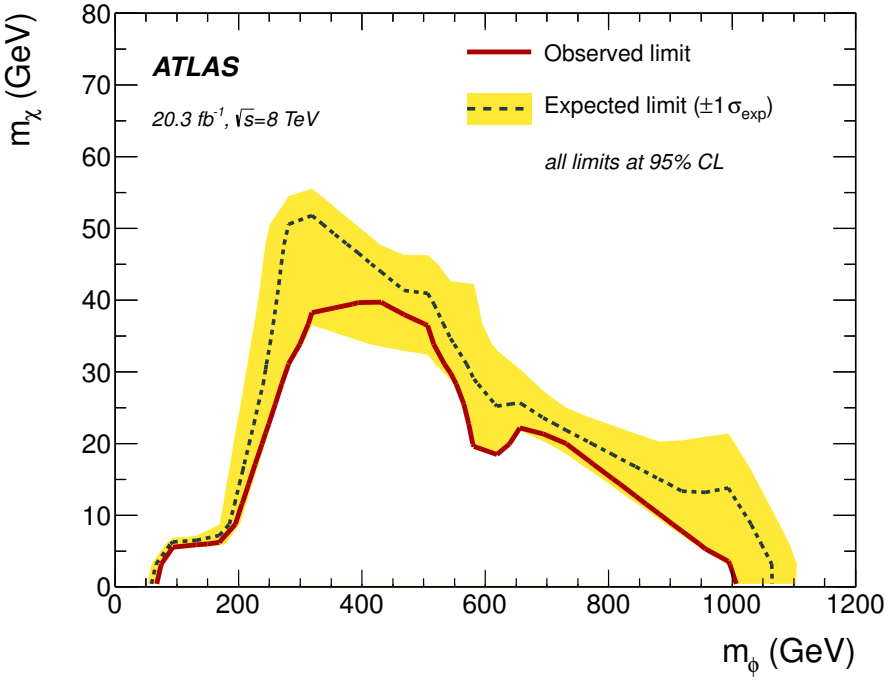

Figure 6.2: The limits set on the $b-$FDM model using SR1 and SR2 of [22].

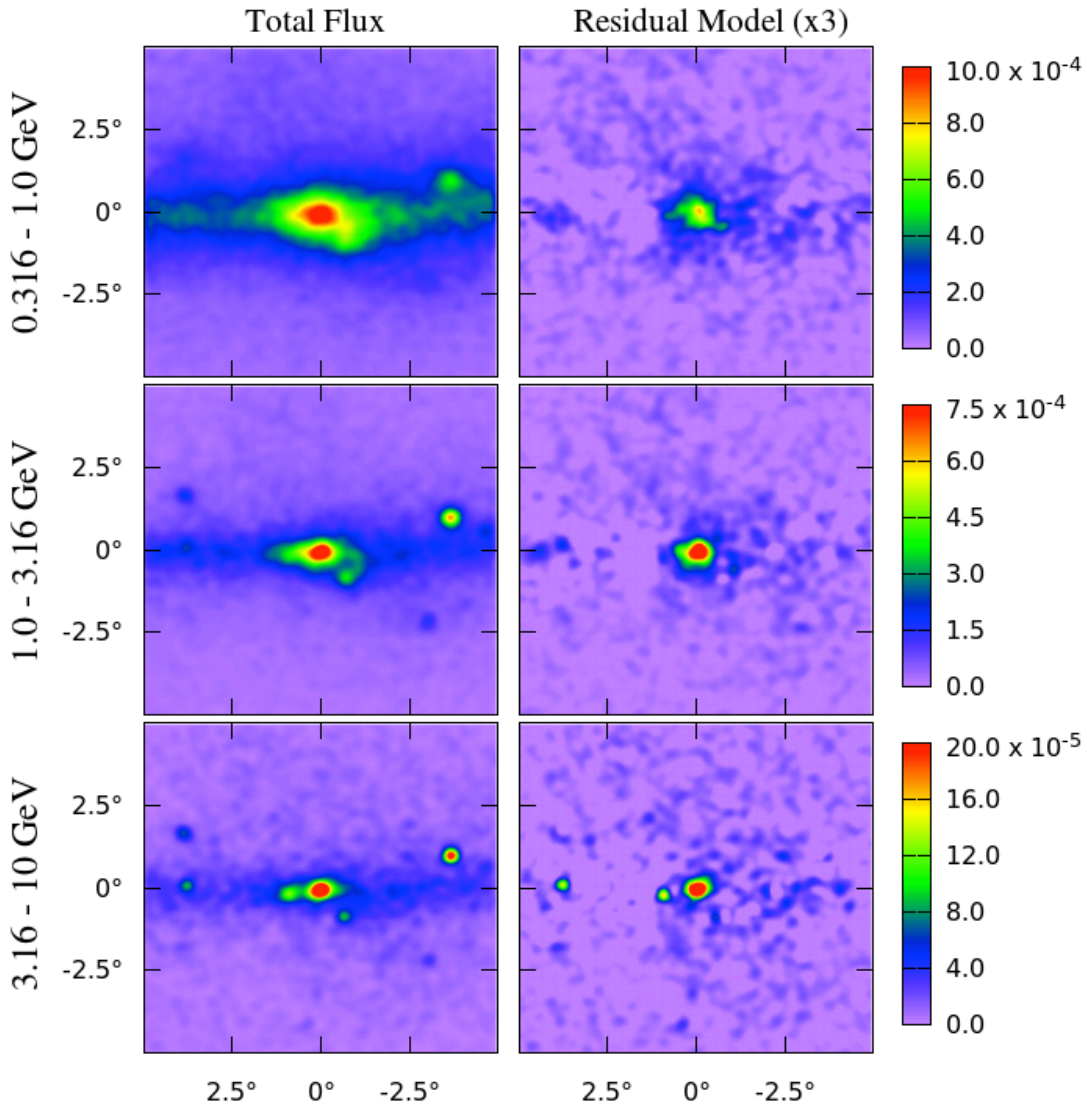

Figure 6.3: Fermi-LAT γ-ray images of the galactic center excess. On the left are the total γ-ray flux measured for three energy ranges, and on the right are the residuals after the background from known sources is subtracted. [142]

Appendices

Appendix A

Selection Optimization

A.1 Method

In order to perform *any* experiment, we need to fix a *null hypothesis* (H_0), and pick an *alternative hypothesis* (H_a) to test. Generally, a search for NP can be performed using either exclusion or discovery approaches, which differ entirely in the definition of H_0 and H_a. In a discovery search, H_a can either be a specific signal, or any generic excess over the expected background, which is used as H_0. The results of the experiment will be interpreted as the probability of fluctuations in H_0 mimicking an observation of H_a. If H_a is correct, optimizing with this approach provides the best results, and as such it's used in analyses where H_a is expected to be true.

In an exclusion test, where you don't expect to observe anything new, the goal is to rule out regions of some parameter space defining your signal. Of course, the discovery approach will still work in these cases, but it doesn't produce optimal results. Instead, we switch the definitions of H_0 and H_a, so that the results are interpreted as the probability of observing no NP signal *if* NP exists. This is the approach taken in the typical WIMP search, since we don't *expect* to find anything in such a specific region of the parameter space.

Let us define N_a and N_0 as the total number of events predicted for H_a and H_0 respectively, estimated using any method. We can define the statistical significance, $S^{\pm}$, of observing H_a as,

$$S^{\pm} \equiv \pm \frac{N_a - N_0}{\sqrt{(\sigma^{\pm}_{\text{data}})^2 + (\sigma^{\pm}_0)^2}} \tag{A.1}$$

where $\sigma^{\pm}_{\text{data}}$ is the *expected* uncertainty on an observation of H_0, and $\sigma^{\pm}_0$ is the uncertainty on the N_0 estimate. Because the ultimate measurement in this analysis is event yields, N_0 will always be described by Poisson statistics, meaning that $\sigma^{\pm}_{\text{data}} = \sqrt{N_0}$. The significance is always positive, by definition, so that S^+ corresponds to the case $N_a > N_0$ and S^- to $N_0 > N_a$. The typical discovery search looks for an *excess* over the expected background, and would therefore use S^+, while S^- is used for exclusion tests such as this analysis.

Now let us define R to be some region in the parameter space used by the analysis. N_s, N_b, $\sigma^{\pm}_s$, $\sigma^{\pm}_b$ are all defined as functions of R, that give the predicted number of signal and background events, and the overall asymmetric uncertainties on each. Next, we define a signal strength parameter, μ, which is a freely floating parameter that determines the normalization of the signal relative to the predicted N_s. The number of expected signal events for any value of μ then becomes $\mu N_s \pm \mu\sigma_s$. To optimize an exclusion test, we can explicitly define the earlier quantities as $N_a \equiv N_b(R)$ and $N_0 \equiv N_b(R) + \mu N_s(R)$. The test statistic then, as a function of R and μ, becomes,

$$S^-(\mu, R) = \frac{\mu N_s(R)}{\sqrt{N_b(R) + \mu N_s(R) + \left(\sigma^-_b(R)\right)^2 + \left(\mu\sigma^-_s(R)\right)^2}}. \tag{A.2}$$

For the EFT signals used in this analysis, the signal normalization μ is directly related to the effective coupling parameter M_* by,

$$\mu = \left(\frac{M^0_*}{M_*}\right)^y, \tag{A.3}$$

where $M^0_* = 1\,\text{TeV}$ is just the coupling used in the production of the simulated samples and y is 4 and 6 for scalar and fermionic WIMPs respectively. To optimize a signal region then, we need to fix the significance S^-, and adjust R to maximize M_*, arriving at M^{lim}_*. M^{lim}_* then corresponds to an exclusion of $M_* > M^{\text{lim}}_*$, with a significance determined by

the fixed value of S^-. We can use the above formulae to solve for $M_*^{\lim}$, arriving at,

$$M_*^{\lim}(S,R) = M_*^0 \left[\frac{2\left(\frac{N_s^2}{S^2} - (\sigma_s^-)^2\right)}{N_s + \sqrt{N_s^2 + 4\left(\frac{N_s^2}{S^2} - (\sigma_s^-)^2\right)(N_b + (\sigma_b^-)^2)}} \right]^{\frac{1}{y}}. \tag{A.4}$$

The standard used by ATLAS analyses is 95% confidence, which roughly sets $S^- \equiv 2$. For completeness, this can be compared to the much simpler form obtained using the discovery approach,

$$M_*^{\lim}(S,R) = M_*^0 \left(\frac{N_s}{S\sqrt{N_b + (\sigma_b^+)^2}} \right)^{\frac{1}{y}}. \tag{A.5}$$

Note that for $\sigma_s = 0$, $\sigma_b^+ = \sigma_b^-$, if $N_s \ll N_b$, both of these methods become equivalent. In an exclusion test, constraints are set on downward fluctuations and in general, the uncertainty on the signal model is crucial. On the other hand, discovery searches constrain upward fluctuations, and the uncertainty on the signal model is irrelevant. In any region, $M_*^{\lim}$ will always exist using the discovery method, while for exclusion tests $M_*^{\lim}$ becomes undefined when $S\sigma_s^- > N_s$. This is expected behavior, since if $S\sigma_s^- > N_s$, even a signal yield of $N_s = 0$ can't be ruled out with significance S.

A.2 Application

Now that we can find a lower bound on M_* in any region, R, we can optimize the analysis cuts that define R, in order to maximize that bound. Any R can be represented by a volume in some N-dimensional vector space, $\mathcal{V}$, over the analysis variables. Let us define M functions, $C_i(R)$, that transform a region R into some new region in $\mathcal{V}$. To represent simple cuts, such as the ones used here, $C_i(R) \to C_i^{\pm}(R,x)$, parameterized by a single number, x, corresponding to either a lower bound or upper bound applied to R along some

axis, v_i. An analysis will then be defined as a collection of these bounds, $\boldsymbol{x}$, which specify a SR $R(\boldsymbol{x}) \equiv C^{\pm}_{i_1}(C^{\pm}_{i_2}(\ldots C^{\pm}_{i_M}(\mathcal{V}, x_M)\ldots, x_2), x_1)$.

This approach to the analysis allows us to define a new function,

$$I^{\pm}_i(\boldsymbol{x_0}, x, \pm) = M^{\lim}_* \left(S, C^{\pm}_i\left(R(\boldsymbol{x_0}), x\right)\right), \tag{A.6}$$

which gives the $M^{\lim}_*$ value corresponding to the limit set by applying cut i, with a bound of x, to the region $R(\boldsymbol{x_0})$. For any region $R(\boldsymbol{x_0})$, we now have one-dimensional functions for $M^{\lim}_*$, over any of the analysis variables. These can be easily visualized, as in Fig. A.1, which shows examples of the optimization plots we will be using. This method gives a very good estimate of the $M^{\lim}_*$ expected from any SR, but its *absolute* value is irrelevant to optimization. In order to compare multiple signals with varying cross-sections, each $I^{\pm}_i(\boldsymbol{x_0}, x)$ is normalized to its value in the absence of cut i, $M^{\lim}_*(S, R(\boldsymbol{x_0}))$, giving the fractional effect of any cut on $M^{\lim}_*$. Error bands on each curve are estimated by propagating the statistical errors on the estimates N_s and N_b, and can point out when the MC statistics are too limited for a cut to be reliable.

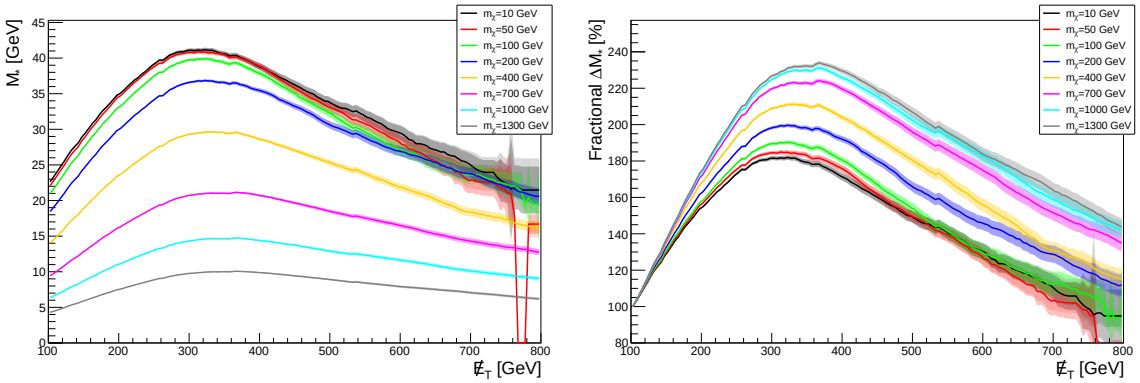

(a) $M^{\lim}_*$ **(b)** Normalized

Figure A.1: Example limit curves used for optimization.

The optimal value for a cut on some variable, v_i, applied to a region $R(\boldsymbol{x_0})$, is then

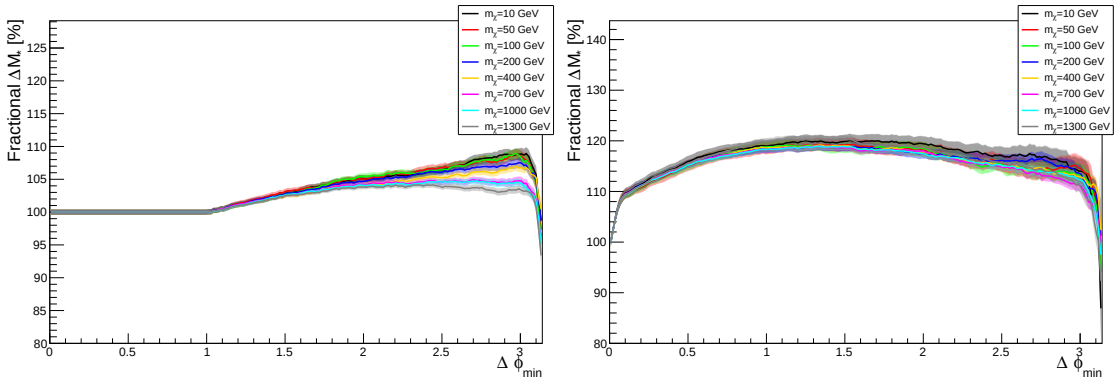

(a) b-jet Multiplicity **(b)** SR300

Figure A.2: An example of the effect of correlated variables on optimization.

given by

$$x_i^{\max} = \arg\max_x I_i^{\pm}(\boldsymbol{x_0}, x). \tag{A.7}$$

If all of the cut variables were completely independent, we could optimize the entire analysis by simply calculating $x_i^{\max}$ for each event variable, in *any* region. However, in reality, the variables used in physics analyses tend to be heavily correlated, leading to a region-dependent $x_i^{\max}$ that will change after every cut. This is illustrated in Fig. A.2, where the variable $\Delta\phi_{min}$ is optimized both before and after the $\not{E}_T > 300\,\text{GeV}$ cut. Going by the earlier stage plot, we would cut on $\Delta\phi_{min} > 2.0$ or even $\Delta\phi_{min} > 3.0$. However, it's clear that after applying the $\not{E}_T$ cut this could actually reduce sensitivity. Due to correlations such as these, simply fitting each cut in an uncorrelated way is likely to produce sub-optimal results. A more reliable method would be to do an N-dimensional search over *all* the analysis variables simultaneously, to find a *global* maximum for $M_*^{\lim}$. Ideally, this would be done with multivariate fitting algorithms, but trial and error is an effective solution too. Once we are satisfied that the SR is, at least nearly, at the global maximum, we can go back to looking at the individual cuts to test its stability. The result will then be a stable, local maximum of $M_*^{\lim}$, which we believe is nearly as large as the global maximum.

Seven variables were chosen that discriminate well between the signal samples and one or more of the SM backgrounds. The variables n_b and $\not{E}_T$ clearly follow from the description

of the signal, and a cut on $n_j < 5$ was necessary to avoid overlap with the DM $+ t\bar{t}$ channel. The MV1_{max} of an event corresponds to the maximum MV1 score of any reconstructed jet, and allows us to optimize the b-tag working point using the above methods. To separate the multijet background, a number of angular variables were considered, but a single lower bound on $\Delta\phi_{min}$ was found to be sufficient. To cut out the multijet background further, a cut on $p_T(j_0)$ was originally included, but, because of correlations with the $\not{E}_T$ cut it was later found to be redundant. It was left as a preselection cut though, because it does remove poorly modeled regions in earlier cut stages, and the analysis had already been unblinded. A cut on $p_T(b_0)$ was also found to be very effective in discriminating against most of the SM backgrounds, making the $p_T(j_0)$ cut completely ineffective. The final variable considered was m_T^b, the transverse mass of the leading b-jet and the $\vec{\not{p}}_T$. This is likely to be the most effective single discriminant for the signal, even though it wasn't used in the final analysis. Fig. A.3 shows histograms of these seven variables in the optimization regions where all of the final selection cuts have been applied, except for the one plotted.

The nominal selection described in Chapter 5 was decided on using an iterative procedure of trial and error. Cuts were made by examining histograms that compare signal to background, and also by using the single variable optimization curves. Once a final selection was chosen, the methods above were used to validate that it was a stable maximum. The resulting curves, in their respective optimization regions, are plotted in Figs. A.4-A.10, showing that the selection chosen roughly maximizes sensitivity to within a few percent for all signal samples. Because of the very large discrepancy in optimal $\not{E}_T$ cut between each signal, the original $\not{E}_T$ cut of 300 GeV was split into five cuts on $\not{E}_T$, ranging from 300 GeV–500 GeV in 50 GeV increments. From these results it's evident that for lower mass WIMPs, the n_b cut is the strongest discriminant considered, followed by the $\not{E}_T$ and $\Delta\phi_{min}$ cuts. For the higher mass WIMPs, the $\not{E}_T$ cut has a much larger impact, although the analysis isn't particularly sensitive to these signals to begin with. The remaining cuts are productive, but only improve the final sensitivity by up to a few percent.

The m_T^b variable combines three of the major variables in the analysis ($\not{E}_T$, $p_T(b_0)$, and

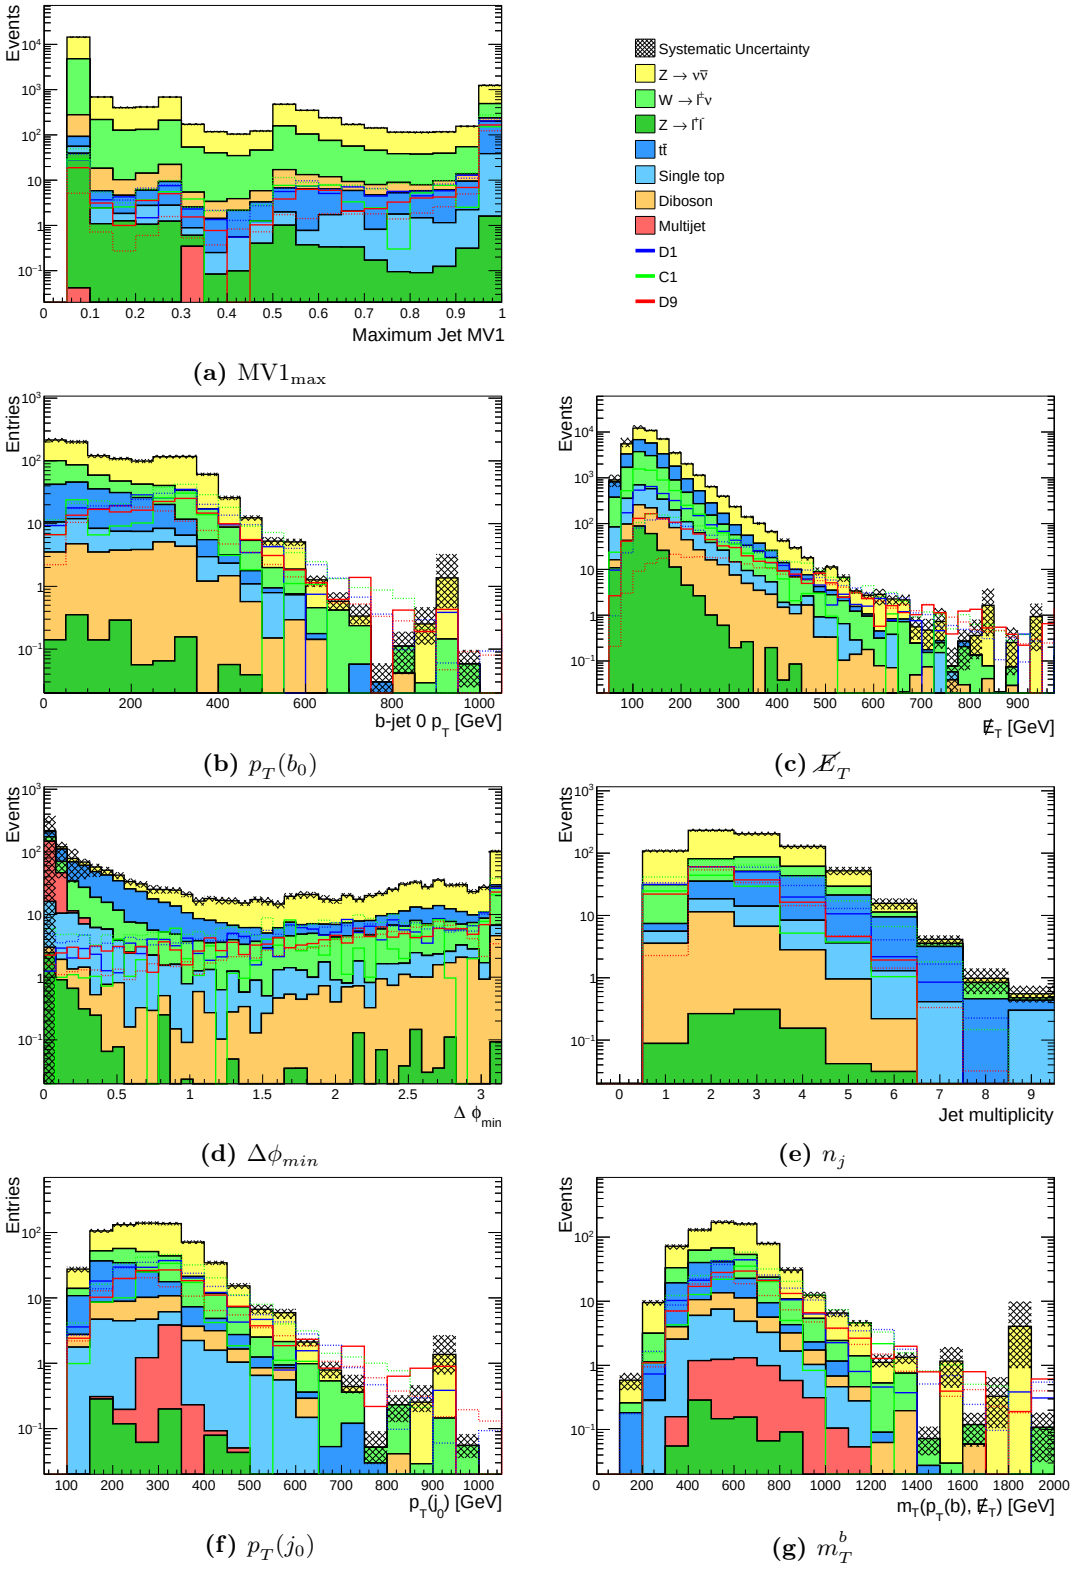

(a) $\mathrm{MV1_{max}}$ **(b)** $p_T(b_0)$ **(c)** $\not{E}_T$

(d) $\Delta\phi_{min}$ **(e)** n_j

(f) $p_T(j_0)$ **(g)** m_T^b

Figure A.3: Predicted distributions for kinematic variables used for optimization.

$\Delta\phi_b$) in a non-trivial way, so it has the potential to improve sensitivity. Fig. A.11 shows the optimization curves for m_T^b after the nominal $\Delta\phi_{min}$, n_j, and $\not{E}_T$ cuts. By comparing the absolute M_* sensitivities, it turns out that while the $\Delta\phi_{min}$ variable still improves sensitivity over a cut on m_T^b, both the $\not{E}_T$ and $p_T(b_0)$ cuts can be replaced by a single cut on m_T^b, increasing sensitivity to certain signals by up to 5%. Unfortunately, this was discovered too late, and this variable could not be included *this* analysis.

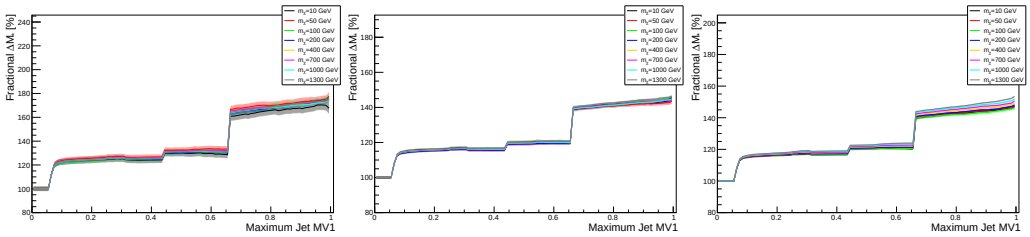

(a) C1 **(b)** D1 **(c)** D9

Figure A.4: The optimization curves for MV1_{max}, used to determine the best b-tag working point. These are generated after all of the selection cuts *except* for the cut on n_b.

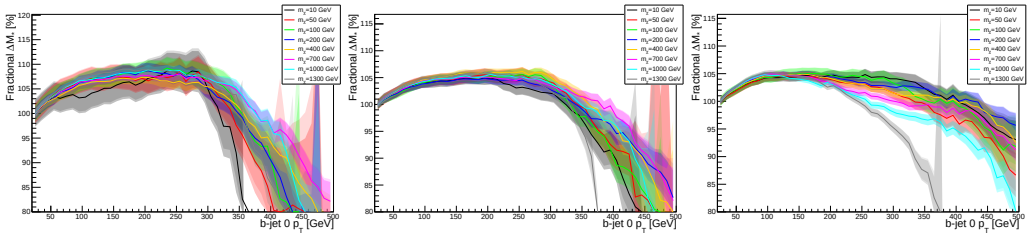

(a) C1 **(b)** D1 **(c)** D9

Figure A.5: The optimization curves for $p_T(b_0)$, generated after all of the selection cuts *except* for the cut on $p_T(b_0)$.

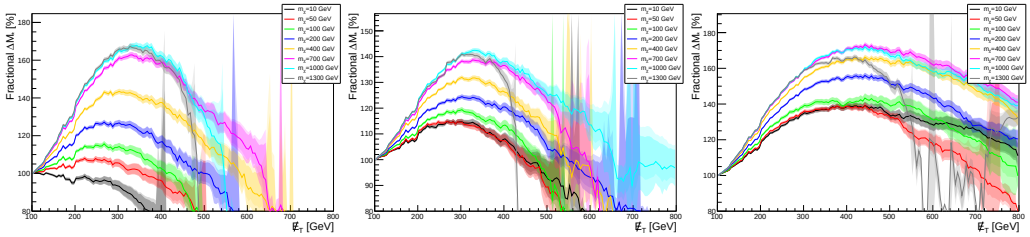

(a) C1 **(b)** D1 **(c)** D9

Figure A.6: The optimization curves for $\not{E}_T$, generated after all of the selection cuts *met* for the $\not{E}_T$ cuts.

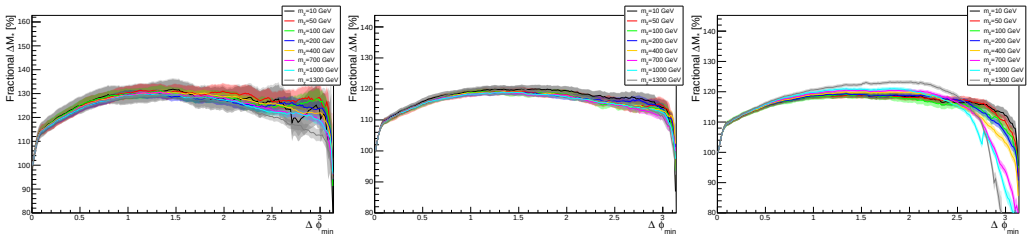

(a) C1 **(b)** D1 **(c)** D9

Figure A.7: The optimization curves for $\Delta\phi_{min}$, generated after all of the selection cuts *except* for the $\Delta\phi_{min}$ cut.

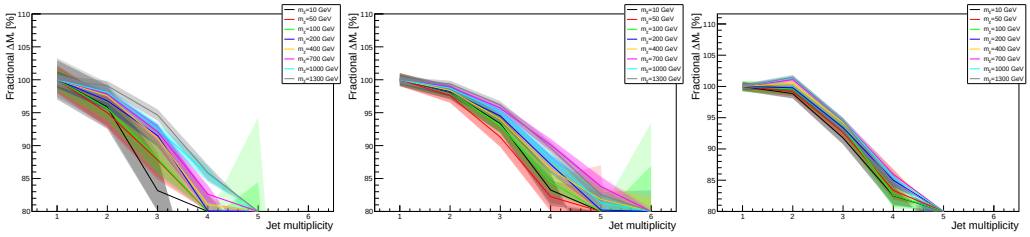

(a) C1 **(b)** D1 **(c)** D9

Figure A.8: The optimization curves for a lower bound on n_j, generated after all of the selection cuts *except* for the n_j cut.

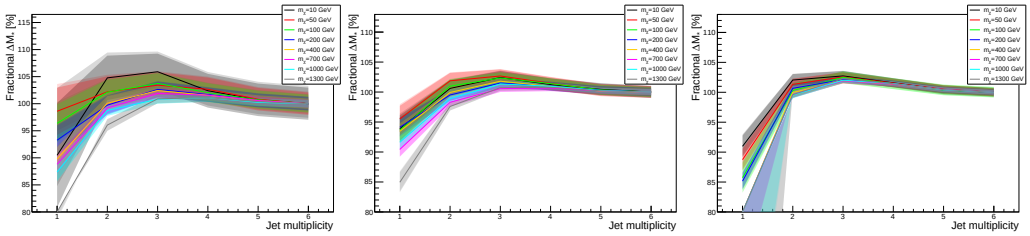

(a) C1 **(b)** D1 **(c)** D9

Figure A.9: The optimization curves for an upper bound on n_j, generated after all of the selection cuts *except* for the n_j cut.

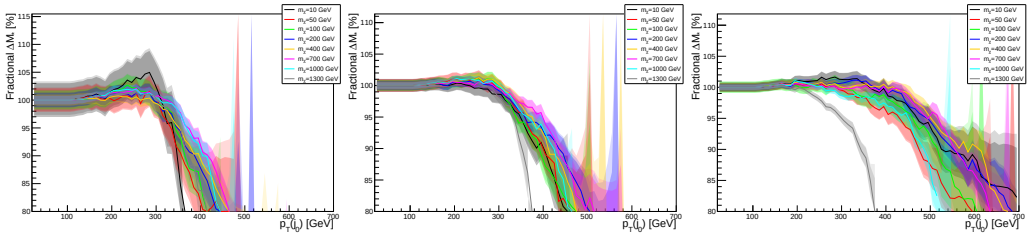

(a) C1 **(b)** D1 **(c)** D9

Figure A.10: The optimization curves for $p_T(j_0)$, generated after all of the selection cuts *except* for the $p_T(j_0)$ cut.

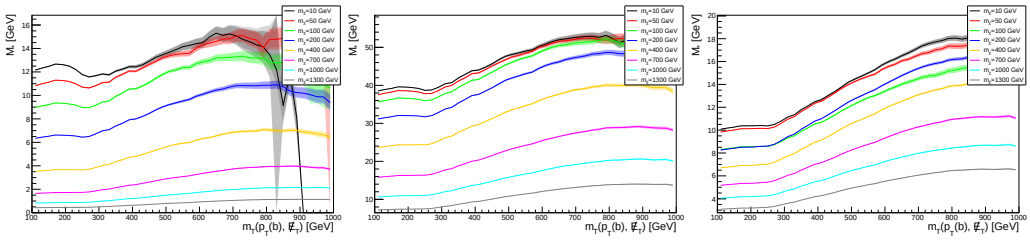

(a) C1, Jet Multiplicity **(b)** D1, Jet Multiplicity **(c)** D9, Jet Multiplicity

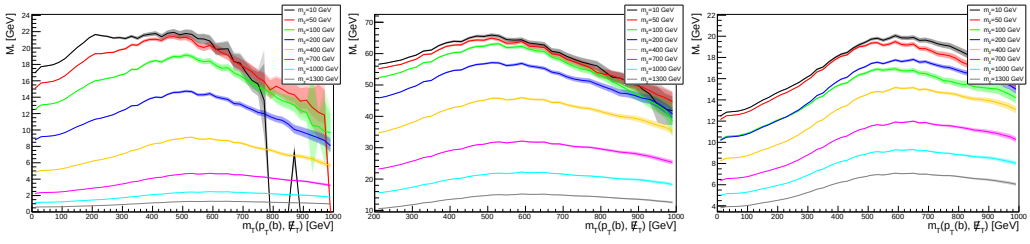

(d) C1, $\Delta\phi_{min}$ **(e)** D1, $\Delta\phi_{min}$ **(f)** D9, $\Delta\phi_{min}$

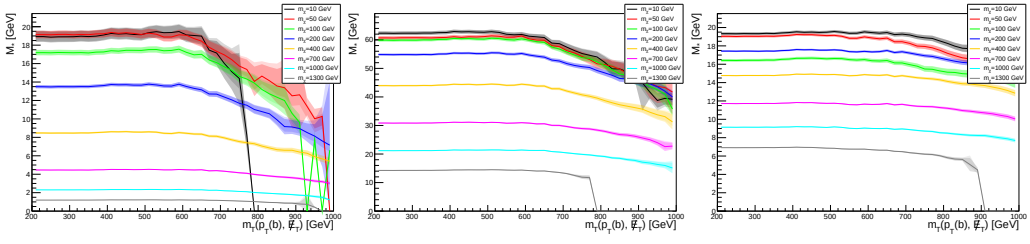

(g) C1, SR300 **(h)** D1, SR300 **(i)** D9, SR300

Figure A.11: Evolution of the optimization curves for a cut on m_T^b

Appendix B

$Z^0 \to \nu\bar{\nu}$ Estimate

Being the dominant background process to the signal, it's very important to have a precise method of modeling $Z \to \nu\nu$. While MC simulations provide an adequate approximation in many cases, they typically have large systematic uncertainties that reduce the sensitivity of an analysis. For the largest backgrounds, it's usually best to develop data-driven methods for estimating their contribution to the SR, in order to reduce these effects. In this analysis, two different data-driven methods are used to generate a very precise $Z \to \nu\nu$ estimate across the full range of $\not{E}_T$.

B.1 $Z^0 \to \nu\bar{\nu}$ from $Z^0 \to \mu^{\pm}\mu^{\mp}$

The simplest method of estimating the $Z \to \nu\nu$ process is by making use of simpler Z^0 decays, such as $Z \to \mu\mu$. Not only do muons have a higher reconstruction efficiency and fewer backgrounds than any other particle, but they are also minimally ionizing particles which behave similarly to neutrinos throughout much of ATLAS. Excluding the MS and ID muon tracks, which are easily isolated, these two processes are nearly identical. The enhancement of this process near the Z^0 boson mass resonance allows a very pure sample of real $Z \to \mu\mu$ events to be collected and then transformed into $Z \to \nu\nu$ pseudo-data.

In order to generate this estimate, the di-muon preselection from Section 5.4 is used to isolate a high purity set of $Z \to \mu\mu$ events. Once two muons have been selected, the $\not{E}_T$ is recalculated without their contributions, effectively removing them from the entire event. These pseudo-data events are then put through the nominal cuts to provide a first

pass at the $Z \to \nu\nu$ estimate. Although this CR is nearly 100% pure at preselection, after b-tagging the top backgrounds become non-negligible, and must be subtracted. This is done by repeating the selection and modification of di-muon events on each of the MC samples, excluding $Z \to \mu\mu$, and assigning them a negative event weight. By including these simulated events with the pseudo-data, the contributions from all other SM processes will cancel out. The pseudo-data events, before subtraction, are plotted in Figs. B.1 and B.2, along with the MC samples put through the same process. The excellent agreement between data and MC in this region is what allows us to subtract the backgrounds and treat the result as a pure data-driven estimate of the $Z \to \mu\mu$ process.

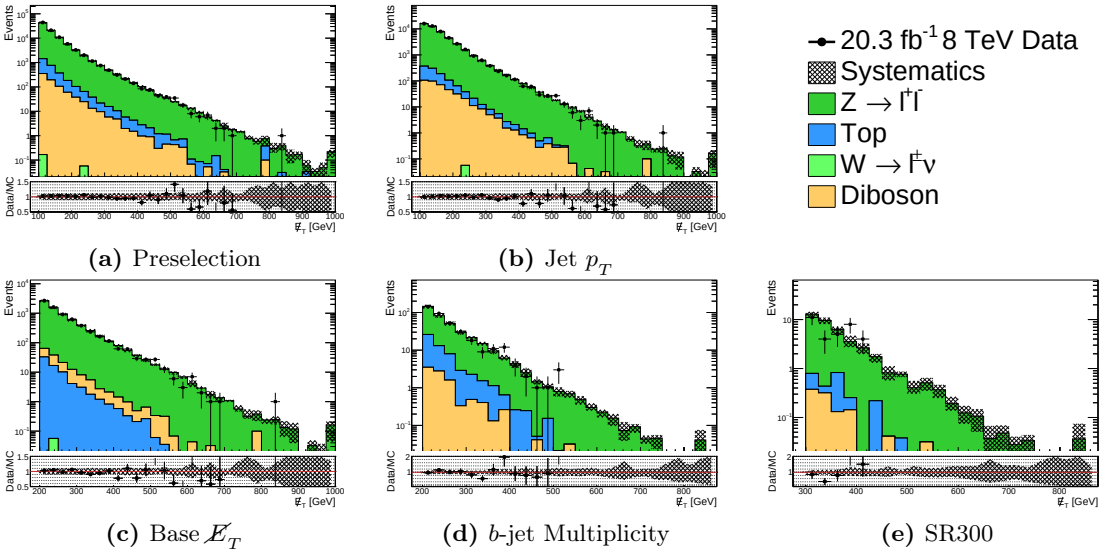

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** SR300

Figure B.1: $\not{E}_T$ distributions of the $Z \to \nu\nu$ pseudo-data after each of the nominal SR cuts.

Despite the similarities between $Z \to \nu\nu$ and $Z \to \mu\mu$, there are still efficiency, acceptance, and branching ratio differences between the two processes that need to be taken into account. For example, the branching ratios of the two decays add a weight of about 6 to each pseudo-data event. Additionally, the $\not{E}_T$ calculation in a typical $Z \to \nu\nu$ decay is measured using calorimeter data, while muons are identified and reconstructed using the MS. These different subsystems have different η coverage and identification efficiencies, leading to a nontrivial relation between the $Z \to \mu\mu$ sample and the $Z \to \nu\nu$ we want to simulate.

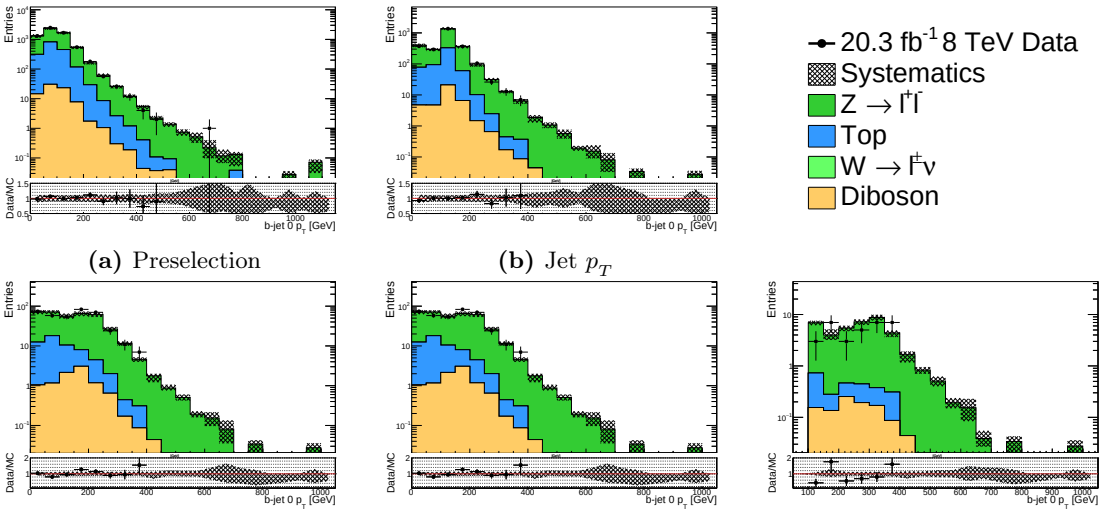

(a) Preselection **(b)** Jet p_T

(c) Base $\not{E}_T$ **(d)** b-jet Multiplicity **(e)** SR300

Figure B.2: $p_T(b_0)$ distributions of the $Z \to \nu\nu$ pseudo-data after each of the nominal SR cuts.

This relation is known as a *transfer factor* (TF), and in general can vary as a function of many event variables.

For any region being considered, the TF is calculated by taking the ratio $R = \sigma(Z \to \nu\nu)/\sigma(Z \to \mu\mu)$, derived from MC in the same region. Examples of the TF in different regions are plotted in Fig. B.3. The total yield of the $Z \to \mu\mu$ sample in this region is then weighted by the TF, giving the final data-driven estimate of the $Z \to \nu\nu$ yield. In practice, we typically look at histograms that are split up into bins. Every bin can be treated as a sub-region of whichever region the histogram is being plotted in, and a separate TF can be calculated. To accomplish this, for every data-driven $Z \to \nu\nu$ histogram we want to produce, four histograms are necessary. Calculating the TF requires separate histograms for the simulated $Z \to \mu\mu$ and $Z \to \nu\nu$ processes, while the data-driven $Z \to \mu\mu$ estimate requires histograms for the $Z \to \mu\mu$ pseudo-data and modified MC backgrounds. The TF can then be applied to the data-driven $Z \to \mu\mu$ sample on a bin-by-bin basis, to get a histogram for the $Z \to \nu\nu$ estimate. All of the histograms in Chapter 5 use this method for plotting the data-driven $Z \to \nu\nu$ estimate.

The advantage of the TF method is that it's the ratio of two heavily correlated MC

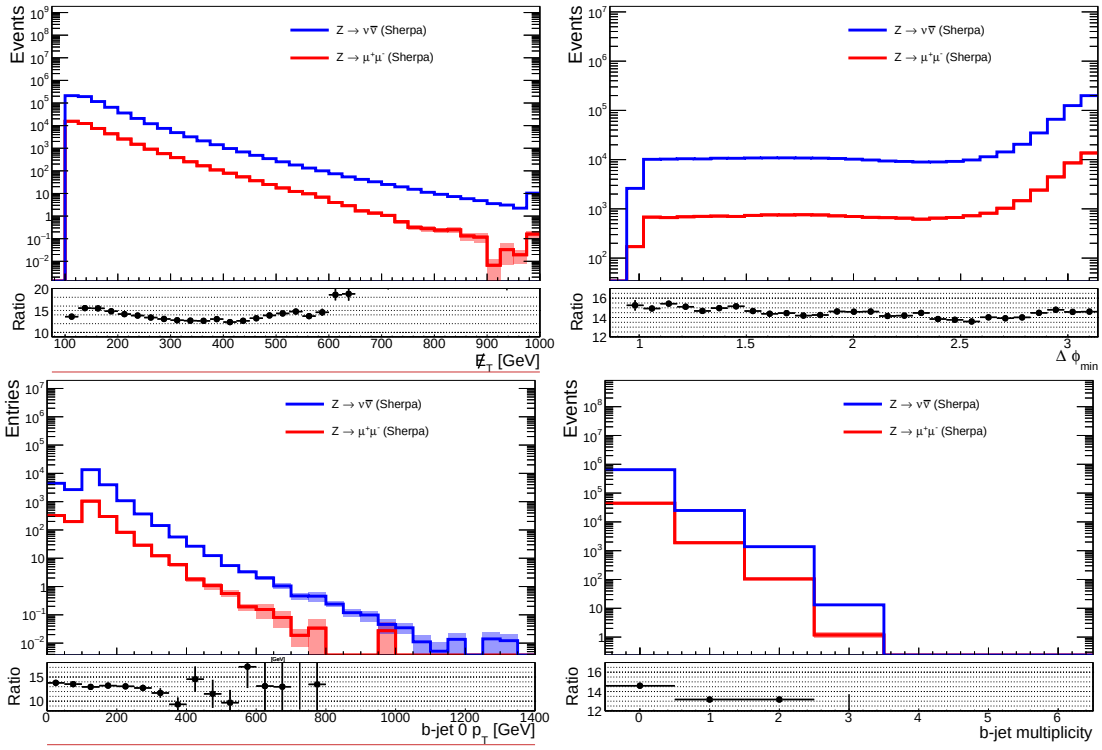

Figure B.3: Transfer factors of the $Z \to \nu\nu$ from $Z \to \mu\mu$ estimate for different kinematic variables after the Jet p_T cut

samples. Many of the systematic uncertainties in these simulations affect the two samples in the same way, partially canceling out any effect on their ratio. A comparison of the systematic uncertainties on the two Sherpa samples and the two pseudo-data samples is given in Table B.1. Although the major uncertainties from the two Sherpa samples almost completely cancel each other out, the large contributions of the $t\bar{t}$ process in this region produce additional uncertainties on the $Z \to \mu\mu$ pseudo-data, which *aren't* canceled out by the ratio. This results in about a 9% systematic uncertainty, in addition to a 20% statistical uncertainty in SR300. While this method does manage to reduce systematics slightly, the total uncertainty ends up being much larger in the high $\not{E}_T$ regions.

	Monte Carlo		Pseudo-Data	
Systematic	$Z \to \nu\nu$	$Z \to \mu\mu$	$Z \to \mu\mu$	$Z \to \nu\nu$
b Flavor	$^{+4.2}_{-4.3}$	$^{+4.1}_{-4.2}$	$^{-0.2}_{+0.2}$	$^{-0.0}_{+0.2}$
c Flavor	$^{+4.0}_{-4.0}$	$^{+4.3}_{-4.4}$	$^{-0.3}_{+0.3}$	$^{-0.5}_{+0.5}$
Light Flavor	$^{+3.2}_{-3.2}$	$^{+3.2}_{-3.2}$	$^{-0.0}_{+0.1}$	$^{+0.4}_{-0.2}$
JES	$^{+7.6}_{-7.5}$	$^{+6.7}_{-7.6}$	$^{-0.5}_{+0.8}$	$^{-4.5}_{-3.5}$
JER	± 0.1	± 1.1	± 0.1	± 4.0
Soft Term	$^{+0.0}_{-0.4}$	$^{-0.0}_{+0.2}$	$^{+0.0}_{+3.4}$	$^{+1.3}_{+3.2}$
Soft Resolution	± 0.1	± 0.1	± 6.8	± 0.8
Pileup	$^{+0.5}_{-0.4}$	$^{+0.4}_{-0.4}$	$^{+0.1}_{-0.1}$	$^{+1.0}_{-1.0}$
JVF	$^{+0.2}_{-0.1}$	$^{-0.0}_{-0.1}$	$^{-0.0}_{+0.0}$	$^{+0.2}_{-0.0}$
PDF	± 4.9	± 7.3	± 3.4	± 6.8
Muon Trigger	$^{+0.0}_{+0.0}$	$^{+0.5}_{-0.6}$	$^{-0.1}_{+0.1}$	$^{-0.5}_{+0.7}$
Muon MS	$^{+0.0}_{+0.0}$	$^{-0.0}_{-0.0}$	$^{-0.0}_{-0.0}$	$^{+0.0}_{+0.0}$
Muon ID	$^{+0.0}_{+0.0}$	$^{-0.2}_{-0.2}$	$^{-0.0}_{+0.0}$	$^{+0.2}_{+0.3}$
Muon Scale	$^{+0.0}_{+0.0}$	$^{-0.0}_{-0.1}$	$^{+0.0}_{+0.1}$	$^{-0.5}_{+0.1}$
Muon Efficiency	$^{+0.0}_{+0.0}$	$^{+0.7}_{-0.8}$	$^{-0.1}_{+0.1}$	$^{-0.7}_{+0.9}$
Normalization	± 0.0	± 0.0	± 0.0	± 0.0
$\sum \sigma^+$	+11.3%	+12.1%	+8.4%	+8.8%
$\sum \sigma^-$	−11.2%	−12.7%	−7.6%	−9.9%

Table B.1: Summary of systematic uncertainties of the $Z \to \nu\nu$ from $Z \to \mu\mu$ estimate in SR**300**. The two MC samples are used to calculate the TF, which is applied to the $Z \to \mu\mu$ pseudo-data to arrive at the final data-driven $Z \to \nu\nu$ estimate.

B.2 $Z^0 \to \nu\bar{\nu}$ from γ

Although the method described in the previous section very accurately models the $Z \to \nu\nu$ process, the branching ratio for the di-muon decay is about 6 times lower than that for the neutrino decay. Taking into account the efficiency differences turns this into a factor of about 14. In regions where the number of expected background events is very low, the statistical errors on this estimate will become very significant and even exceed the systematic errors of MC simulations. This was demonstrated above, where even in the loosest SR, the statistical uncertainties were over twice as large as the systematics.

Therefore, in the SRs used here, that estimate isn't an adequate replacement for MC.

Using methods such as those described in [147], the $Z \to \nu\nu$ process can also be estimated by taking advantage of the similarities between Z^0 boson and γ-ray production. At high p_T, these two processes become increasingly similar, with the main difference coming from flavor couplings. It was found that for $p_T > 200\,\text{GeV}$, γ-ray production becomes a suitable estimate of Z^0 production. Because of the high p_T requirement, this can't *completely* replace the above estimate and instead, the two are stitched together around the $\not{E}_T = 200\,\text{GeV}$ point. For events below this threshold, $Z \to \mu\mu$ is used and for higher $\not{E}_T$ events γ-rays are used. This allows the low statistic area of the $Z \to \mu\mu$ sample to be replaced by the higher statistic γ-ray sample, which has fewer backgrounds and no branching ratio suppression.

The selection of prompt photons makes use of CR**gamma**, which singles out events with one high p_T photon. Its contributions are removed from the $\not{E}_T$ calculation, and the other SM processes are subtracted using MC, exactly as done earlier. The $Z \to \nu\nu$ Sherpa samples are then split into a high $\not{E}_T$ and low $\not{E}_T$ sample, around the $\not{E}_T = 200\,\text{GeV}$ point. The TF for the $Z \to \mu\mu$ method is recalculated using the low $\not{E}_T$ sample, and the TF for the γ method is calculated using the high $\not{E}_T$ sample. Examples of the γ to $Z \to \nu\nu$ TF are plotted in Fig B.4.

Unlike the $Z \to \mu\mu$ pseudo-data, the $\gamma + \text{jets}$ pseudo-data is nearly 100% pure, as was shown in Section 5.6. This means that the MC subtraction step has a negligible effect, and does not introduce any significant systematics. Table B.2 shows a comparison of the systematic uncertainties for the two MC samples used to calculate the TF, along with the pseudo-data sample before and after reweighting in SR300. The $\gamma + \text{jets}$ pseudo-data is almost entirely real data, meaning that the only significant systematics come from the TF, which are very suppressed, and the photon isolation requirement. Additionally, because photons are stable, there is no branching ratio suppression, and the statistical uncertainties using this method are only 3.2% in SR300, and 13% in SR500.

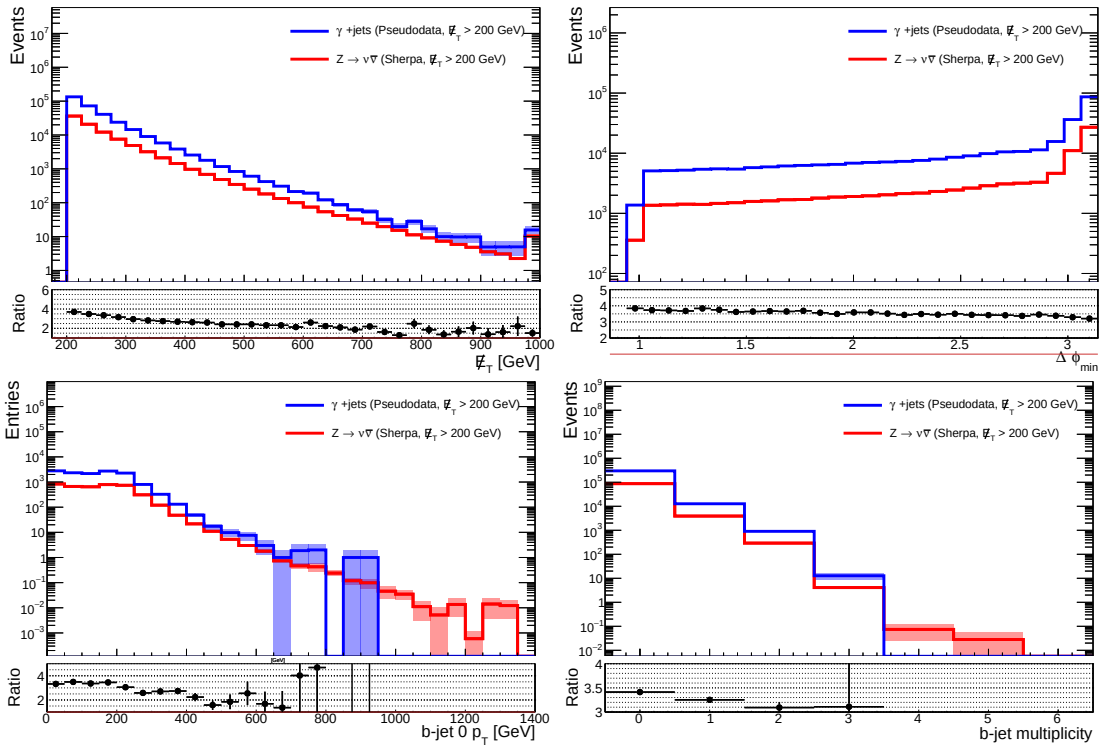

Figure B.4: Transfer factors of the $Z \to \nu\nu$ from γ estimate for different kinematic variables after the Base $\not\!\!E_T$ cut.

	Monte Carlo		Pseudo-Data	
Systematic	$Z \to \nu\nu$	γ + jets	γ + jets	$Z \to \nu\nu$
b Flavor	$^{+4.2}_{-4.3}$	$^{+3.4}_{-3.4}$	$^{-0.0}_{+0.0}$	$^{+0.8}_{-0.9}$
c Flavor	$^{+4.0}_{-4.0}$	$^{+5.7}_{-5.8}$	$^{-0.1}_{+0.1}$	$^{-1.7}_{+1.9}$
Light Flavor	$^{+3.2}_{-3.2}$	$^{+3.6}_{-3.6}$	$^{-0.0}_{+0.0}$	$^{-0.4}_{+0.4}$
JES	$^{+7.6}_{-7.5}$	$^{+8.2}_{-7.7}$	$^{-0.1}_{+0.0}$	$^{-0.4}_{+0.4}$
JER	± 0.1	± 0.0	± 0.0	± 0.2
Soft Term	$^{+0.0}_{-0.4}$	$^{+0.3}_{-0.3}$	$^{+0.9}_{+0.4}$	$^{+0.8}_{+0.5}$
Soft Resolution	± 0.1	± 0.1	± 0.4	± 0.0
Pileup	$^{+0.5}_{-0.4}$	$^{+0.5}_{-0.5}$	$^{+0.0}_{-0.0}$	$^{+0.1}_{+0.0}$
JVF	$^{+0.2}_{-0.1}$	$^{+0.2}_{-0.2}$	$^{+0.2}_{-0.4}$	$^{+0.1}_{-0.4}$
PDF	± 4.9	± 3.4	± 0.2	± 2.3
Photon Zee	$^{+0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{+0.0}_{+0.0}$	$^{-0.0}_{+0.0}$
Photon Material	$^{-0.0}_{+0.0}$	$^{+0.0}_{+0.0}$	$^{+0.0}_{+0.0}$	$^{-0.0}_{+0.0}$
Photon Presampler	$^{+0.0}_{+0.0}$	$^{-0.0}_{+0.0}$	$^{+0.0}_{+0.0}$	$^{+0.0}_{-0.0}$
Photon Low-p_T	$^{+0.0}_{+0.0}$	$^{+0.0}_{-0.0}$	$^{+0.0}_{+0.0}$	$^{-0.0}_{-0.0}$
Photon Resolution	$^{+0.0}_{+0.0}$	$^{+0.0}_{+0.0}$	$^{-0.0}_{+0.0}$	$^{-0.0}_{+0.0}$
Photon ID	$^{+0.0}_{+0.0}$	$^{+0.5}_{-0.5}$	$^{-0.0}_{+0.0}$	$^{-0.5}_{+0.5}$
Photon Isolation	$^{+0.0}_{+0.0}$	$^{+4.0}_{-4.0}$	$^{+4.0}_{-4.0}$	$^{+4.0}_{-4.0}$
Normalization	± 0.0	± 0.0	± 0.0	± 0.0
$\sum \sigma^+$	+11.3%	+12.3%	+4.2%	+5.2%
$\sum \sigma^-$	−11.2%	−12.1%	−4.1%	−5.0%

Table B.2: Summary of systematic uncertainties of the $Z \to \nu\nu$ from γ estimate in SR**300**. The two MC samples are used to calculate the TF, which is applied to the γ + jets pseudo-data to arrive at the final data-driven $Z \to \nu\nu$ estimate.

Appendix C

Multijet Estimate

Multijet production is one of the most challenging SM backgrounds to model in hadron collisions. Because billions of protons are being smashed into each other every 50 ns, low energy QCD processes end up being recorded in nearly every event. Even if the pileup effects can be ignored, the cross-section for the hard production of jets in a proton-proton collision is many orders of magnitude above even the least interesting electroweak processes. If these jets were perfectly measured, they could easily be filtered out at the trigger level. The problem though, is that no detector is perfect, and rare multijet events can occasionally mimic interesting processes. Since the cross-section is so many orders of magnitude higher, the result is the product of a very large and very small number, making the contributions from this background difficult to predict. MC simulations usually do a poor job of simulating these effects and therefore data-driven methods are almost always required to get a reliable estimate. Many methods have been developed at ATLAS to handle this, but the one that was used for this analysis is known as jet smearing [137].

C.1 Single Jet Triggers

The first step in the jet smearing method is to obtain a high purity sample that are representative of the typical multijet event, which will be modified to simulate atypical events. This is referred to as the *seed sample*, and its selection requires the use of low threshold single jet triggers. Because of the gigantic cross-section for the production of low p_T jets, these triggers are heavily prescaled, and their effective luminosity is significantly

lower than the $20.3\,\text{fb}^{-1}$ collected by ATLAS. Getting a large seed sample requires some careful stitching together of different triggers, to maximize the effective luminosity, thereby minimizing statistical uncertainties.

The trigger selection for the seed events involves choosing a trigger based on the offline leading jet p_T. For any event, the highest threshold trigger is used that is predicted to be 99% efficient in triggering on the leading jet. Implementing this, of course, requires a measurement of the trigger efficiencies as a function of the leading jet's p_T and η. Ideally, trigger efficiencies are calculated by using a second unbiased trigger, as in Appendix D. However, for the wide range of trigger thresholds used here, it's difficult to find an appropriate one. Therefore, a method known as *bootstrapping* is used instead, to extract an unbiased efficiency from a set of biased triggers. The only requirement is that the *lowest* threshold trigger must have enough overlap with an unbiased trigger to extract its efficiency. This can be a problem, since these are typically heavily prescaled triggers that will have almost no overlap. Rather than checking whether or not the target trigger *actually* fires then, we check if it *should* have fired on the data that was accessible to it at the time. For the single jet triggers this is pretty straightforward, since the triggers only cut on the η and E_T of the jet object available to it, which is stored for offline trigger studies such as this. The thresholds for each of the single jet trigger chains used in bootstrapping are listed in Table C.1, along with their reference trigger chain and total prescale value. The reference trigger for the three lowest threshold triggers is a zero bias trigger, which fires randomly. Because of the lower efficiency and acceptance of the forward jet triggers, only central jet triggers were used, with a cut of $|\eta| < 3.2$. The 45 GeV and 55 GeV triggers were excluded as well because the 35 GeV trigger had approximately the same prescale value, and was more reliable throughout the 8 TeV run.

Bootstrapping is an application of Bayes' theorem, $P(R|T)P(T) = P(T|R)P(R)$, on a pair of triggers to calculate the efficiency of a target T, $\epsilon_T(p_T)$, using the efficiency of a reference R, $\epsilon_R(p_T)$. Since, in this case, R is just a lower threshold version of T, $P(R|T) = 1$, and therefore $\epsilon_T(p_T) = C(p_T)\epsilon_R(p_T)$, where $C(p_T) = P(T|R)$ is the efficiency

of T on events which pass R. This quantity, $C(p_T)$, is the biased efficiency we actually measure from the data. The bootstrapping process begins at the zero bias trigger, which is completely independent of the collision and therefore has a constant 100% efficiency with respect to the leading jet p_T. For a fixed leading jet p_T then, the efficiency of the 15 GeV trigger is simply the fraction of events that have a trigger jet object with $p_T > 15$ GeV and $|\eta| < 3.2$. The 80 GeV EF trigger can then be calculated since ϵ_R, which is the efficiency of the 15 GeV trigger, is now known.

Generally, for each event where a reference trigger with known efficiency, $\epsilon_R(p_T)$ has fired, a histogram of the leading jet's reconstructed p_T is filled, $N_R(p_T)$. If the EF, L2, *and* L1 levels of the target chain contain jet objects with $\eta < 3.2$ and that meet the E_T cuts listed in Table C.1, a second histogram is filled that represents $N_T(p_T) = N_R(p_T)C(p_T)$. Dividing this by the first histogram gives $C(p_T)$, which can be multiplied by $\epsilon_R(p_T)$ to apply Bayes' theorem and arrive at a histogram of the unbiased target efficiency, $\epsilon_T(p_T)$. The efficiency curves for all of the relevant triggers are plotted together in Fig. C.1. For simplicity, the reference trigger for each target was chosen so that the plateau region of the target had $\sim 100\%$ efficiency, so that $\epsilon_R \sim 1$ in the region of interest. The point at which each of the efficiency curves reach 99% are displayed in the legend, and these are used to define the stitching point for each triggers.

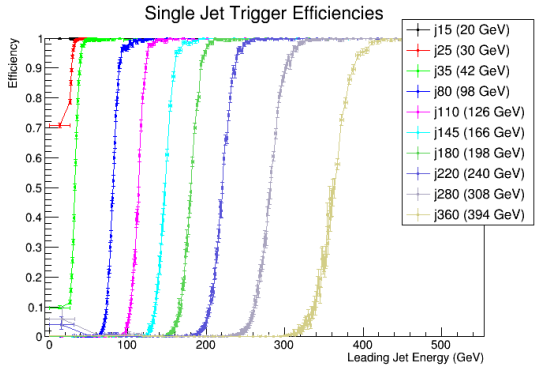

Figure C.1: Efficiencies of the single jet triggers used to select the seed sample for the multijet estimate. 99% efficiency point for each trigger is given in the legend. Note that the 15 GeV trigger plateau's below the 20 GeV p_T requirement in the jet selection.

EF	L2	L1	Reference	Avg. Prescale
15 GeV	Random	Random	Random	1,784,732
25 GeV	Random	Random	Random	1,784,732
35 GeV	Random	Random	Random	1,784,732
80 GeV	75 GeV	30 GeV	15 GeV	11,009,691
110 GeV	105 GeV	50 GeV	25 GeV	2,068,008
145 GeV	140 GeV	75 GeV	35 GeV	358,303
180 GeV	165 GeV	75 GeV	80 GeV	8,754
220 GeV	165 GeV	75 GeV	110 GeV	2,067
280 GeV	165 GeV	75 GeV	145 GeV	559
360 GeV	165 GeV	75 GeV	180 GeV	257

Table C.1: Central jet trigger chains, along with the reference trigger used for efficiency calculations. The prescale of the reference trigger is shown on the right. The reference triggers are chosen so that they're nearly 100% efficient in the target triggers plateau region.

C.2 Seed Selection

Table C.2 shows single jet trigger requirements used in the multijet seed selection, based on the results of the previous section. The p_T bound of each trigger represents the lower bound on the leading jet p_T, below which the trigger isn't used. Each event is assigned the highest threshold trigger consistent with the reconstructed p_T of its leading jet. The efficiencies of each trigger are over 99% in this selection, so trigger SFs can be safely ignored. However, the prescales still need to be accounted for, so to properly stitch these bins together, each event is weighted by the average prescale of its corresponding trigger.

In addition to the single jet trigger, cuts are placed on $\not{E}_T^{sig}$ and $p_T(j_0)$. Cutting on $\not{E}_T^{sig} < 0.7$ ensures that the $\not{E}_T$ of the event is low, without biasing the selection against heavy flavor jets that frequently produce real $\not{E}_T$, and are crucial to this analysis. Additionally, the leading jet's p_T is required to be over 130 GeV, to avoid the very high prescale weights that are placed on the lower threshold triggers. This cut was carefully selected by validating the shape of the multijet estimate against data in CR**multijet**.

p_T bound	Trigger	$\int Ldt$	Avg. Prescale
20 GeV	EF_j15_a4tchad	$1.8\,\mathrm{nb}^{-1}$	11,009,691
30 GeV	EF_j25_a4tchad	$9.8\,\mathrm{nb}^{-1}$	2,068,008
45 GeV	EF_j35_a4tchad	$57\,\mathrm{nb}^{-1}$	358,303
100 GeV	EF_j80_a4tchad	$2.3\,\mathrm{pb}^{-1}$	8,754
130 GeV	EF_j110_a4tchad	$9.8\,\mathrm{pb}^{-1}$	2,067
170 GeV	EF_j145_a4tchad	$36\,\mathrm{pb}^{-1}$	559
200 GeV	EF_j180_a4tchad	$79\,\mathrm{pb}^{-1}$	257
250 GeV	EF_j220_a4tchad	$261\,\mathrm{pb}^{-1}$	77.6
310 GeV	EF_j280_a4tchad	$1.2\,\mathrm{fb}^{-1}$	17.4
400 GeV	EF_j360_a4tchad	$20\,\mathrm{fb}^{-1}$	1.00

Table C.2: Central jet trigger selection if leading jet has $|\eta| < 3.2$.

When events with leading jets under 130 GeV were included, significant mismodeling and large statistical uncertainties were found after jet smearing. These problems associated with soft jets mainly affected the $\not\!\!E_T$ and $p_T(j_0)$ distributions, but removing them resulted in the excellent agreement that was shown in Chapter 5.

C.3 Jet Smearing

The jet smearing method follows the techniques described in [137] very closely, to estimate the multijet background of high $\not\!\!E_T$ signals. Multijet events can produce $\not\!\!E_T$ through two mechanisms. Fake $\not\!\!E_T$ is a phenomenon related to the imperfect resolution of the calorimeters, which can result in the mismeasurement of a jet. Since $\not\!\!E_T$ is calculated by taking the vector sum of all the objects in an event, this will produce fake $\not\!\!E_T$ either parallel or antiparallel to the jet, depending on whether the energy is underestimated or overestimated. The second source of $\not\!\!E_T$ is real $\not\!\!E_T$ from neutrinos produced in heavy flavor jets. This $\not\!\!E_T$ will typically be parallel to the jet, since the jet energy is underestimated.

The jet response function of the calorimeters can be estimated from data, producing a

Sample	Base Factor	Flavor Factor	Nominal Factor
Multijet Data	10	500	2
Multijet MC	1	100	1
Top MC	50	1	15
Other MC	4	100	1

Table C.3: The number of times seed samples were smeared.

probability density function parameterizing random fluctuations in the measured jet energy. The jets of well measured multijet events, from the seed sample, can be randomly smeared by this function to simulate both the real and fake sources of $\not\!\!E_T$. Each seed event is smeared multiple times, producing arbitrarily large samples. Eventually, of course, double counting will occur and the statistical uncertainties on the estimate will become underestimated. Because of the computational requirements of this process, the number of iterations was reduced in less important regions. For each sample, there is a base factor applied to all seed events, a flavor factor applied to seed events with $n_b \geq 1$ and $p_T(j_0) > 100\,\text{GeV}$, and a nominal factor applied to nominal samples, as opposed to samples produced under systematic variations. Each event's weight is divided by the appropriate factors, and the number of times it's independently smeared is equal to their product. The factors used in the analysis are listed in Table C.3, and were chosen to maximize statistics in important regions without wasting time smearing others.

The seed selection, followed by smearing, was applied to data and simulations for each SM process, in addition to the three systematics discussed in Section 5.7, each with up/down variations. The validation of the final estimate was already done in Chapter 5, but it's also important to verify that we know which processes contribute to the seed sample, and that it's well modeled. Fig. C.2 shows 14 kinematic variables in the seed and smeared selections, comparing data to MC. From these plots it's clear that, as expected, the seed selection is well over 99% pure in multijet events, even in the extreme regions. The statistical uncertainties become very high for the Pythia sample after applying the nominal $\Delta\phi_{min}$ cut, making

it difficult to compare shapes, so the $\not\!\!E_T$ distributions in the CR**multijet** region are also plotted in Fig. C.3. Additionally, Tables C.4 and C.5 list the overall yields in SR300 and CR**multijet** respectively.

As is evident from the plots, the agreement appears to be significantly better than some of the large uncertainties would suggest. This is mostly from statistical fluctuations of the Pythia sample while varying the systematics, where there are so few events in each bin, that each systematic appears to have a much larger effect than it actual would. However, as long as the uncertainties aren't *underestimated*, the $\mathcal{O}(100\%)$ uncertainty has very little effect on the final results. The expected yield from multijets in the *loosest* SR is less than 1 %, with an uncertainty of only 50%.

Cut	Electroweak	Top	Multijet	Total SM	Data
Preselection	$1{,}070 \pm 350$	696 ± 170	$8{,}220{,}000 \pm 2{,}400{,}000$	$8{,}220{,}000 \pm 2{,}400{,}000$	$7{,}086{,}484 \pm 15{,}568$
$\not\!\!E_T > 100\,\text{GeV}$	348 ± 120	281 ± 66	$2{,}170{,}000 \pm 600{,}000$	$2{,}170{,}000 \pm 600{,}000$	$2{,}063{,}082 \pm 10{,}442$
Jet Multiplicity	143 ± 53	40.9 ± 9.8	$1{,}000{,}000 \pm 270{,}000$	$1{,}000{,}000 \pm 270{,}000$	$1{,}261{,}167 \pm 7{,}948$
$\Delta\phi_{min}$	3.13 ± 6.3	1.11 ± 0.56	$23{,}100 \pm 27{,}000$	$23{,}100 \pm 27{,}000$	$32{,}180 \pm 1{,}898$
Jet p_T	2.90 ± 5.6	0.967 ± 0.53	$18{,}200 \pm 29{,}000$	$18{,}200 \pm 29{,}000$	$29{,}008 \pm 1{,}798$
Base $\not\!\!E_T$	0.216 ± 0.73	0.220 ± 0.22	$1{,}150 \pm 5{,}600$	$1{,}150 \pm 5{,}600$	$1{,}015 \pm 324$
b-jet Multiplicity	0.0246 ± 0.13	0.0535 ± 0.058	20.3 ± 54	20.4 ± 54	34 ± 2
b-jet p_T	0.0131 ± 0.053	0.0228 ± 0.024	6.18 ± 22	6.22 ± 22	12 ± 1
SR300	0.000480 ± 0.0021	0.0171 ± 0.026	0.611 ± 2.2	0.628 ± 2.2	4 ± 0

Table C.4: Yields for smeared samples in SR region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

Cut	Electroweak	Top	Multijet	Total SM	Data
Preselection	$1{,}070 \pm 350$	696 ± 170	$8{,}220{,}000 \pm 2{,}400{,}000$	$8{,}220{,}000 \pm 2{,}400{,}000$	$7{,}086{,}484 \pm 15{,}568$
$\not\!\!E_T > 100\,\text{GeV}$	348 ± 120	281 ± 66	$2{,}170{,}000 \pm 600{,}000$	$2{,}170{,}000 \pm 600{,}000$	$2{,}063{,}082 \pm 10{,}442$
Jet Multiplicity	143 ± 53	40.9 ± 9.8	$1{,}000{,}000 \pm 270{,}000$	$1{,}000{,}000 \pm 270{,}000$	$1{,}261{,}167 \pm 7{,}948$
$\Delta\phi_{min}$	136 ± 56	38.1 ± 9.1	$936{,}000 \pm 270{,}000$	$936{,}000 \pm 270{,}000$	$1{,}184{,}730 \pm 7{,}419$
Jet p_T	139 ± 55	39.1 ± 9.3	$954{,}000 \pm 270{,}000$	$954{,}000 \pm 270{,}000$	$1{,}213{,}737 \pm 7{,}634$
Base $\not\!\!E_T$	3.41 ± 8.3	1.24 ± 0.51	$12{,}600 \pm 7{,}200$	$12{,}600 \pm 7{,}200$	$15{,}102 \pm 483$
b-jet Multiplicity	0.909 ± 1.2	0.909 ± 0.41	$3{,}010 \pm 1{,}100$	$3{,}010 \pm 1{,}100$	$4{,}790 \pm 14$
b-jet p_T	0.737 ± 0.31	0.723 ± 0.40	$2{,}220 \pm 740$	$2{,}220 \pm 740$	$3{,}540 \pm 6$

Table C.5: Yields for smeared samples in CRmultijet region. Quoted errors are the combination of statistical and systematic uncertainties added in quadrature.

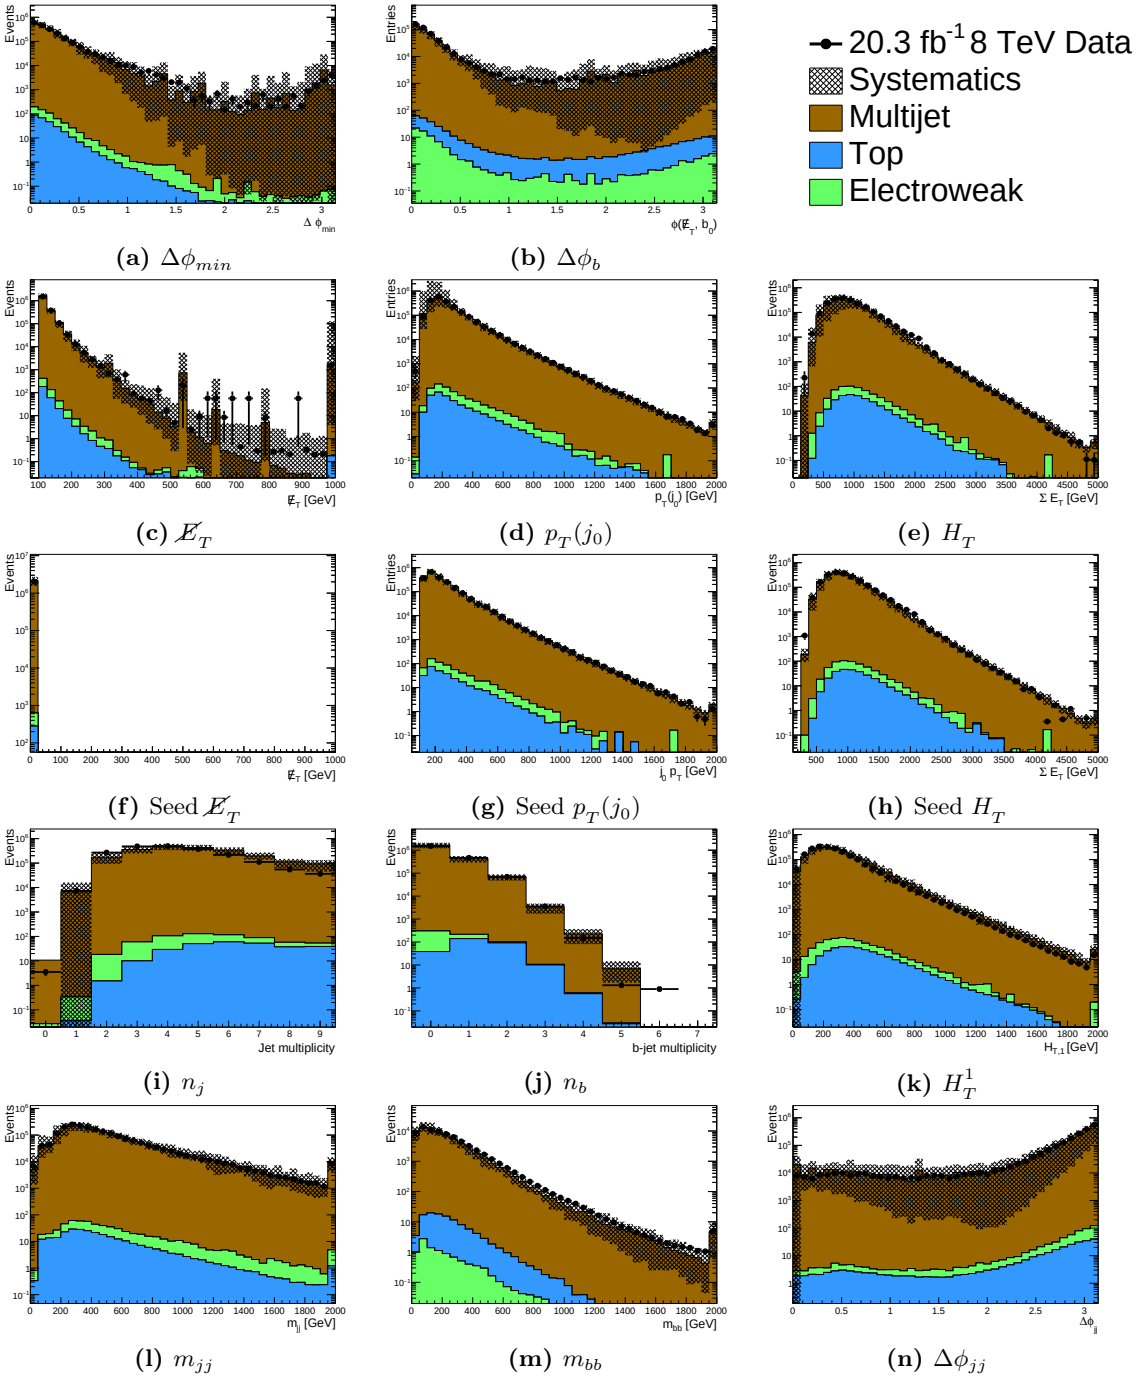

(a) $\Delta\phi_{min}$ **(b)** $\Delta\phi_b$

(c) $\not{E}_T$ **(d)** $p_T(j_0)$ **(e)** H_T

(f) Seed $\not{E}_T$ **(g)** Seed $p_T(j_0)$ **(h)** Seed H_T

(i) n_j **(j)** n_b **(k)** H_T^1

(l) m_{jj} **(m)** m_{bb} **(n)** $\Delta\phi_{jj}$

Figure C.2: Kinematic distributions after the 0 lepton preselection cuts for the smeared samples.

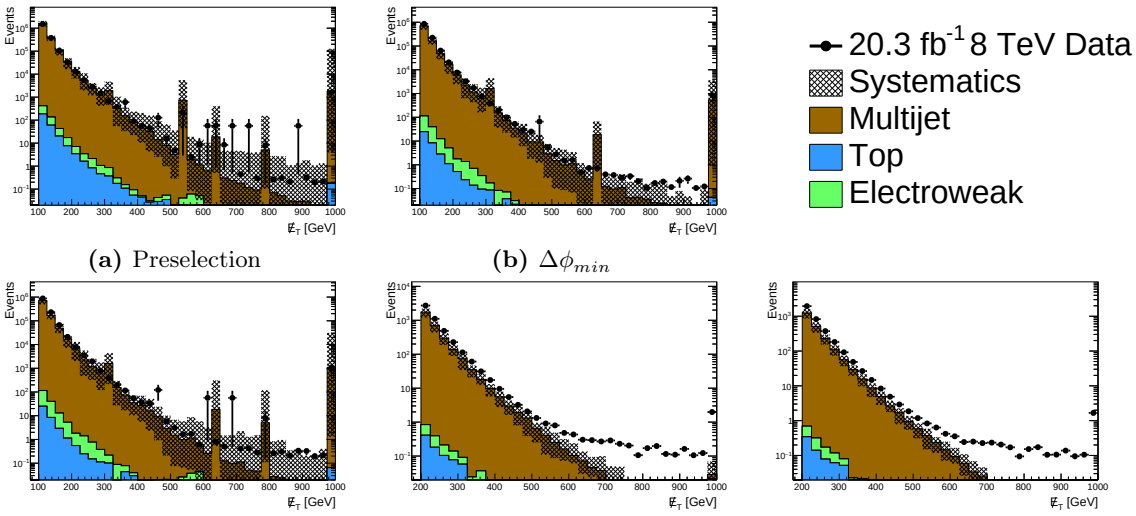

(a) Preselection **(b)** $\Delta\phi_{min}$

(c) Jet p_T **(d)** b-jet Multiplicity **(e)** b-jet p_T

Figure C.3: $\not\!\!E_T$ distributions after each of the CRmultijet cuts for the smeared samples. The is some clear mismodeling in the tail region between the pseudo-data and the smeared MC samples. The exact cause of this was not discovered, but it's very well covered by the smearing systematics, not displayed here, which are $\mathcal{O}(100\%)$.

Appendix D

Missing E_T Triggers

Because this analysis makes use of $\not{E}_T$ triggers, their efficiency with respect to $\not{E}_T$ needs to be well understood, and the low $\not{E}_T$ region that depends on the trigger configuration needs to be defined. Once the trigger efficiencies are known, it extends the $\not{E}_T$ range of the analysis, making it easier to validate background estimates. In this analysis three $\not{E}_T$ triggers were considered: EF_xe80_tclcw, EF_xe80_tclcw_loose, and EF_xe80T_tclcw_loose, which were the lowest unprescaled $\not{E}_T$ triggers in 2012. Although EF_xe80_tclcw_loose and EF_xe80T_tclcw_loose have higher acceptances, they were not enabled for the entire 2012 data taking period, whereas the tighter EF_xe80_tclcw recorded the full $20.3\,\text{fb}^{-1}$. It was found that all three of these triggers plateau at around $\not{E}_T \sim 160\,\text{GeV}$ with $> 99\%$ efficiency, so the looser criteria simply complicates the analysis without improving sensitivity. Therefore, only EF_xe80_tclcw was used for regions in which a $\not{E}_T$ trigger is required.

One fortunate aspect of the $\not{E}_T$ triggers is that their online calculation of $\not{E}_T$ is done without any muon contributions. This allows a tag-and-probe method to calculate their efficiency, similar to what was done by [148]. Using the single muon preselections defined in Section 5.4, we can isolate very pure collections of muon events and calculate $\not{E}_T^{\mu}$, as an estimate of the online trigger $\not{E}_T$ that the 80 GeV threshold is applied to. The efficiency of the trigger can then be calculated, as a function of $\not{E}_T^{\mu}$, simply by looking at whether or not a given $\not{E}_T$ trigger fired. By comparing the efficiencies of the MC samples with data, SFs can be obtained to correct the simulated trigger.

This calculation is done in three orthogonal regions, in order to both validate the method

and estimate systematics. Events with one muon and no b-jets isolate the $W \to \ell\nu$ process, events with one muon and two or more b-jets isolate top quark production, and the dimuon preselection isolates the $Z \to \mu\mu$ process, giving three very pure regions. To calculate the efficiency in data, the muon stream is used, while for MC, the Sherpa $W \to \ell\nu$, $Z \to \ell\ell$ samples and the Powheg $t\bar{t}$ samples are used for their respective selections. Fig. D.1 shows the six efficiency curves obtained and the SFs derived from their ratio. The plateau points at which the trigger becomes 99% efficient are calculated for all of these, and listed in Table D.1. All three of these points are much lower than the 300 GeV $\not{E}_T$ cut, as expected, and should not affect the results of this analysis.

	99% Efficiency (GeV)	
Region	Data	MC
$W^{\pm} \to \ell^{\pm}\nu$	150	144
$Z^0 \to \mu^{\pm}\mu^{\mp}$	151	142
$t\bar{t}$	151	145

Table D.1: The 99% efficiency points for the $\not{E}_T$ triggers, in data and MC for each of the three processes.

The nominal efficiencies and SFs use the $W \to \ell\nu$ data, because it has the lowest statistical uncertainties. Any MC event with $\not{E}_T < 160$ GeV is assigned a SF to correct for mismodeling in the $\not{E}_T$ trigger, and data-driven estimates that simulate the $\not{E}_T$ trigger, such as the multijet and $Z \to \nu\nu$ samples, are weighted by the efficiency measurements. Although it has a negligible effect on the SR, this will introduce an addition uncertainty in the lower $\not{E}_T$ regions. The $t\bar{t}$ and $Z \to \mu\mu$ efficiency measurements are used as systematic variations to estimate this uncertainty, reported in Section 5.7, although even in the low $\not{E}_T$ regions it's a very small effect.

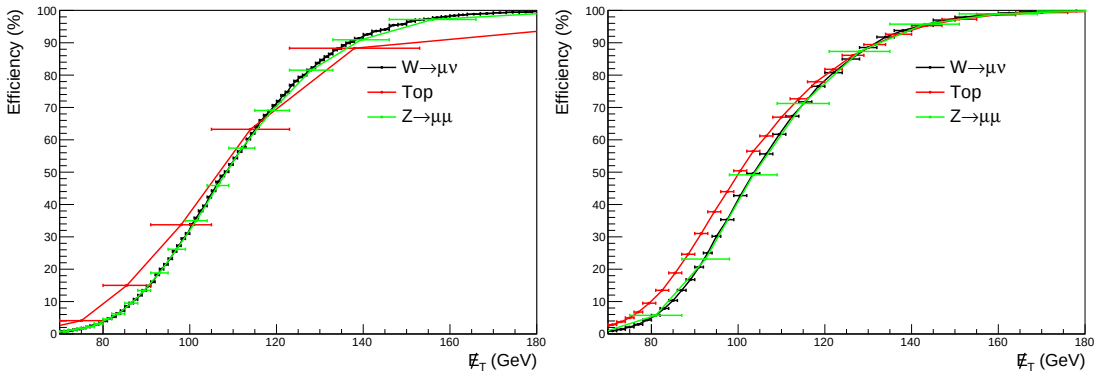

(a) Data

(b) Monte Carlo

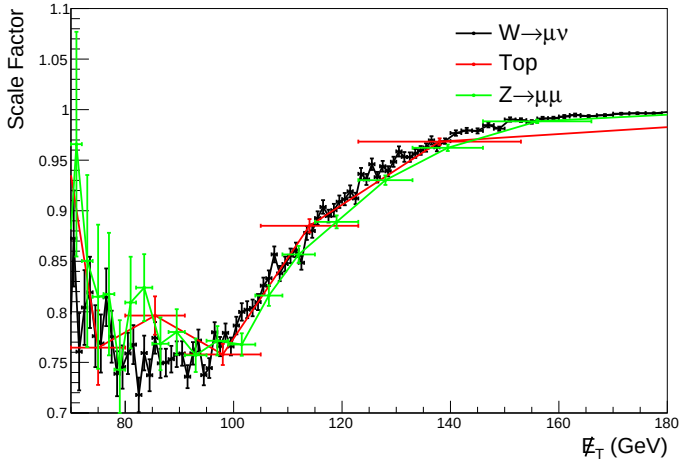

(c) Scale Factor

Figure D.1: Efficiency of the $\not{E}_T$ trigger in early selection stages three different CRs. The data efficiencies were derived from the muon stream, and the MC efficiencies from the nominal samples corresponding to each CR (i.e. Powheg $t\bar{t}$ samples for CR**top**, and Sherpa $W \to \ell\nu$ and $Z \to \ell\ell$ samples for CR**wjets** and CR**zmm** respectively). The SF curves are just the ratio between the efficiency in data and MC, rebinned for clarity.

Appendix E

Experimental Constraints on Dark Matter

There are four different methods that have been developed to study DM. Cosmological observation is the only one of these that has actually observed it so far, and has given us a lot of information about DM on the cosmological scale. By making various assumptions, these observations can be used to place constraints on WIMP models, but in general, details about the particle nature of DM is impossible to probe at these large scales, which is where the other three methods come in. As mentioned in Chapter 1, these correspond to rotations in space-time of arbitrary WIMP-SM interactions. Indirect detection experiments search for the SM byproducts of WIMP annihilations in distant regions dense in DM, direct detection looks for local WIMP collisions with a target nucleus and finally, collider searches attempt to observe the creation of WIMP pairs in collisions between SM particles. All of these methods have a variety of strengths and weaknesses, and complement each other very well. Usually, at least one of these methods will be sensitive to any specific WIMP model across a large region of the parameter space. However, comparing their results can be difficult, since model-dependent choices will always have to be made. In this section, we will use the EFT/MFV assumptions to derive methods that transform the results of different searches, in order to make comparisons.

E.1 Cosmological Constraints

E.1.1 Relic Density

The strongest cosmological constraints come from the observed relic abundance of dark matter, $\Omega h^2 = 0.1188 \pm 0.0010$. If we assume that dark matter is composed entirely of a single WIMP, whose relic abundance is a result of thermal freeze-out, we can use this measurement to place limits on its self-annihilation cross-section. Going back and slightly rearranging the Boltzmann equation from Chapter 3, we start with,

$$\frac{dy}{dx} = -\frac{\langle \sigma v \rangle}{x^2}\left(y^2 - y_{\rm eq}^2\right)\left(\frac{sx}{H}\right). \tag{E.1}$$

The typical thermal WIMP, pretty much by definition, will freeze-out during the radiation dominated period and will be non-relativistic at the time. Since $s \propto x^{-3}$ and $H \propto x^{-2}$ at this time, the term in parentheses on the right, labeled as λ, is constant with respect to x. The non-relativistic nature of the WIMPs allows us to expand the thermally averaged annihilation cross-section as $\langle \sigma v \rangle \approx a + bv^2 + \mathcal{O}(v^4)$, where we will only consider the leading (s-wave) and sub-leading (p-wave) terms. The Boltzmann equation is notoriously difficult to solve, even numerically. However, we are only interested in estimating the relic abundance, which is defined as the density as $x \to \infty$, and various approximations can be made [58]. Well after freeze-out, $x \gg \langle \sigma v \rangle$, y asymptotically approaches a constant value, and $y_{\rm eq}$ continues to drop exponentially. Therefore, we can assume $y \gg y_{\rm eq}$, and arrive at the equation,

$$\frac{dy}{dx} \approx -\lambda \frac{y^2}{x^2}\left(a + 6\frac{b}{x}\right), \tag{E.2}$$

where v^2 has been replaced by its thermal average $6x^{-1}$. This new equation can be solved exactly to get,

$$y = -\frac{1}{3\lambda\dfrac{b}{x^2} + \lambda\dfrac{a}{x} - \dfrac{1}{y_\infty}}, \tag{E.3}$$

where y_∞ is the abundance as $x \to \infty$, and is roughly the relic abundance we observe today. We can fix some arbitrary freeze-out time, x_f, at which point the abundance is defined as $y_f \equiv y(x_f)$, and solving for y_∞ we find,

$$y_\infty = \frac{1}{\dfrac{1}{y_f} + 3\lambda\dfrac{b}{x_f^2} + \lambda\dfrac{a}{x_f}}. \tag{E.4}$$

The y_f^{-1} term can typically be neglected, since $y_\infty \ll y_f$, leaving the relic abundance as a function of a, b, and x_f, which still need to be determined. To extract the physical density Ωh^2 from y, we need an estimate of the radiation entropy today, which is almost entirely contained in the CMB at a known temperature of $(2.718 \pm 0.021)\,\mathrm{K}$ [53]. Carrying out the calculation,

$$\Omega h^2 = 0.1188 \pm 0.0010 = \left(4.23 \times 10^{-11}\,\mathrm{GeV}^{-2}\right) \frac{x_f}{\sqrt{g_*}} \left(a + 3\frac{b}{x_f}\right)^{-1}, \tag{E.5}$$

which constrains the three parameters a, b, and x_f.

In the non-relativistic limit, the equilibrium density obeys,

$$n_{\mathrm{eq}} \approx g \left(\frac{T^2 x}{2\pi}\right)^{\frac{3}{2}} e^{-x}, \tag{E.6}$$

where g is the number of internal degrees of freedom for the WIMP. We can then define x_f as the point at which $y(x_f) = (c+1)y_{\mathrm{eq}}(x_f)$, for some arbitrary c. Well before this freeze-out point, $|y_{\mathrm{eq}} - y| \ll 1$, reducing the Boltzmann equation to,

$$\frac{dy_{\mathrm{eq}}}{dx} \approx -\lambda \frac{\langle \sigma v \rangle}{x^2} c(2+c) y_{\mathrm{eq}}^2. \tag{E.7}$$

Solving this for x and extrapolating to x_f, we find that,

$$e^{x_f} \approx m_\chi c(c+2) \frac{g}{\sqrt{g_*}} \sqrt{\frac{45}{4\pi^5 x_f}} \left(a + 6\frac{b}{x_f}\right), \tag{E.8}$$

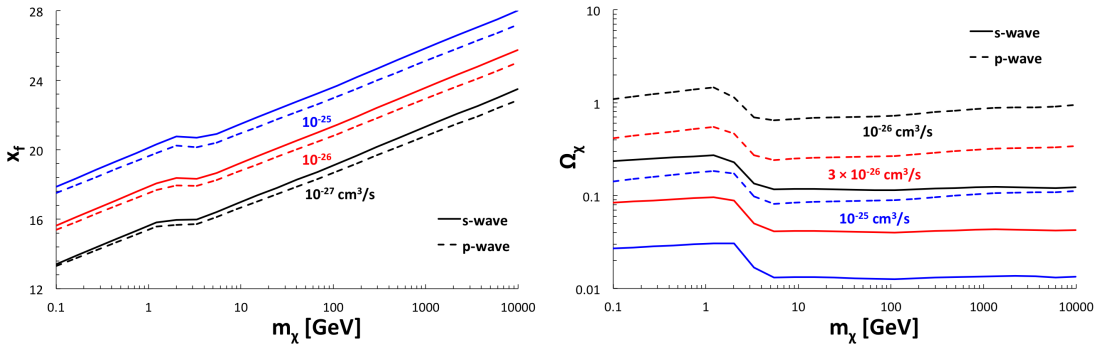

Figure E.1: The freeze-out temperature and relic abundance over a wide range of m_χ and $\langle\sigma v\rangle$ values consistent with the WIMP hypothesis. The model-dependence of freeze-out is very small, and is mostly determined by $\langle\sigma v\rangle$.

which can then be solved numerically for any given a, b, and c. The value of c shown to give the best results ranges from 1, if $a \gg b$, to 2, if $b \gg a$ [58], and the freeze-out point of typical WIMPs is usually in the range $20 < x_f < 30$. Fig. E.1 plots the values of x_f and Ω_x for a variety of different interaction strengths and WIMP masses.

E.1.2 Annihilation Cross-sections

Now that we can constrain the annihilation cross-sections in any WIMP model, we need to calculate them using model parameters. For each of the EFT operators of interest, the thermally averaged annihilation cross-section can be easily estimated from the tree-level Feynman diagrams. Although these diagrams have been calculated in other papers [44, 149, 150], there is no consensus on the results. Therefore, they will be derived again here, using the Feynman rules of the EFT model.

Each fermion line in the Feynman diagram corresponds to a plane wave solution to the Dirac equation [29]. The Feynman rules for external fermions introduce 4-component spinors, u_i, for each fermion i which satisfy a form of the Dirac equation,

$$(\not{p}_i \pm m_i)\, u_i(p_i, s_i) = 0. \tag{E.9}$$

Here m_i is the mass of particle i, s_i is its spin, and $\not{p}_i$ is the Feynman slash notation for

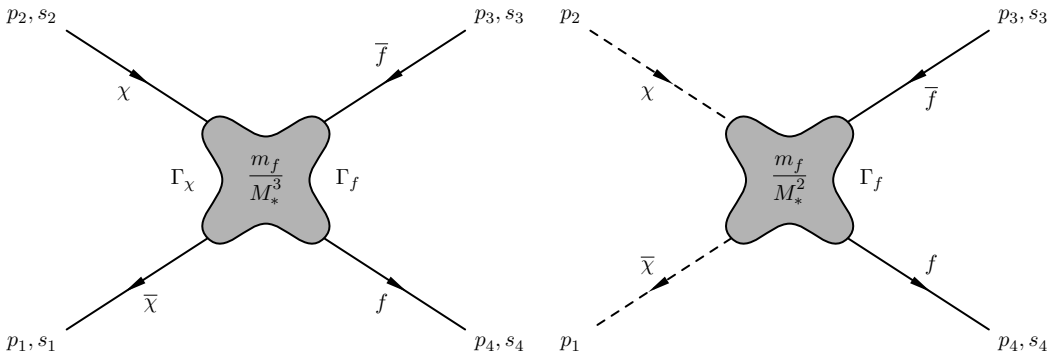

Figure E.2: The tree-level Feynman diagrams for WIMP annihilation to fermions for Dirac (left) and scalar (right) WIMPs.

its 4-momentum contracted with the gamma matrices, $\gamma_\mu p_i^\mu$. Fermions have two possible spins $\pm 1/2$, and we define the usual anti-particle spinors $v_i(p_i, s_i) \equiv u_i(-p_i, -s_i)$ to use for the negative energy solutions. These spinors satisfy the completeness relations,

$$\sum_s u(p,s)\overline{u}(p,s) = \not{p} + m, \tag{E.10}$$

$$\sum_s v(p,s)\overline{v}(p,s) = \not{p} - m, \tag{E.11}$$

where $\overline{u}_i \equiv \gamma^0 u_i^\dagger$. Incoming fermion (anti-fermion) lines and outgoing anti-fermion (fermion) lines are represented by u_i and v_i ($\overline{v}_i$ and $\overline{u}_i$) respectively. Generally, unless dealing with polarized beams, observations are averaged over the initial spin states and summed over the possible final spin states. These sums, combined with the completeness relations, can be rewritten in terms of traces over combinations of gamma matrices, which can then be evaluated and rewritten as scalar expressions. Incoming and outgoing scalar particles have none of these complexities, and are represented by simple factors of 1 in Feynman diagrams.

For the EFT operators used in this analysis, Fig. E.2 shows the general Feynman diagrams for fermion or scalar WIMP annihilations, where the Γ factors correspond to the same gamma matrices that appear in the operators from Table 2.1. For each diagram this

leads to the matrix elements,

$$\mathcal{M} = \frac{m_q^2}{M_*^6}\overline{u}_4\Gamma_f v_3\overline{v}_1\Gamma_\chi u_2, \quad \text{(fermion)} \tag{E.12}$$

$$\mathcal{M} = \frac{m_q^2}{M_*^4}\overline{u}_4\Gamma_f v_3. \quad \text{(scalar)} \tag{E.13}$$

The differential cross-section of an annihilation is given by,

$$\frac{d\sigma}{d\Omega} = \frac{1}{256\pi^2 s}\frac{|\vec{p}_3|}{|\vec{p}_1|}\sum_{s_1,s_2,s_3,s_4}|\mathcal{M}|^2, \tag{E.14}$$

which will be integrated to get the total cross-section σ, once the angular dependence of the matrix element is established. Two useful definitions,

$$\Sigma_f \equiv \mathrm{Tr}\left\{(\not{p}_4 + m_4)\,\Gamma_f\,(\not{p}_3 - m_3)\,\gamma^0\Gamma_f^\dagger\gamma^0\right\}, \tag{E.15}$$

$$\Sigma_\chi \equiv \mathrm{Tr}\left\{(\not{p}_1 - m_1)\,\Gamma_\chi\,(\not{p}_2 + m_2)\,\gamma^0\Gamma_\chi^\dagger\gamma^0\right\}, \tag{E.16}$$

can be extracted by evaluating the sum over spin states and rearranging the terms within the trace. Re-expressing the cross-section in the center of mass frame then, we find that,

$$\sigma = \frac{1}{256\pi^2 s}\sqrt{\frac{s-4m_f^2}{s-4m_\chi^2}}\frac{m_f^2}{M_*^6}\int\left|\Sigma_f\Sigma_\chi\right|d\Omega, \quad \text{(fermion)} \tag{E.17}$$

$$\sigma = \frac{1}{256\pi^2 s}\sqrt{\frac{s-4m_f^2}{s-4m_\chi^2}}\frac{m_f^2}{M_*^4}\int\left|\Sigma_f\right|d\Omega. \quad \text{(scalar)} \tag{E.18}$$

The Σ_i terms are straightforward to evaluate for each operator, using the trace properties of the gamma matrices [29]. The values for scalar and pseudoscalar couplings can be calculated very easily as,

$$\Sigma_i^S = 2s - 8m_i^2, \quad (\text{scalar}, \Gamma_i \equiv 1) \tag{E.19}$$

$$\Sigma_i^P = -2s, \quad (\text{pseudoscalar}, \Gamma_i \equiv \gamma^5) \tag{E.20}$$

while the tensor and pseudotensor terms are much more complicated. Because they share indices that are summed over, they aren't really separable. However, they can all be represented by a rank-4 tensor,

$$\mathcal{M}_{\mu\nu\lambda\rho} \equiv \mathrm{Tr}\left\{\left(\not{p}_4 + m_4\right)\sigma_{\mu\nu}\left(\not{p}_3 - m_3\right)\sigma_{\lambda\rho}\right\}. \tag{E.21}$$

The products of Σ_f and Σ_χ then, are simply the only unique contractions of $\mathcal{M}_{\mu\nu\lambda\rho}$ with itself that result in a nontrivial Lorentz scalar,

$$\Sigma_f^T \Sigma_\chi^T \equiv \mathcal{M}_{\mu\nu\lambda\rho}\mathcal{M}^{\mu\nu\lambda\rho}, \quad \text{(tensor)} \tag{E.22}$$

$$\Sigma_f^{PT} \Sigma_\chi^{PT} \equiv \epsilon^{\mu\nu\sigma\delta}\epsilon^{\lambda\rho\gamma\epsilon}\mathcal{M}_{\mu\nu\lambda\rho}\mathcal{M}_{\sigma\delta\gamma\epsilon}. \quad \text{(pseudotensor)} \tag{E.23}$$

After the tedious exercise of evaluating these sums, which the reader will not be subjected to here, these can be re-expressed as,

$$\Sigma_f^T \Sigma_\chi^T = 512 m_f^2 m_\chi^2 + 16 s^2 + 16\left(s - 4m_f^2\right)\left(s - 4m_\chi^2\right)\cos 2\theta, \quad \text{(tensor)} \tag{E.24}$$

$$\Sigma_f^{PT} \Sigma_\chi^{PT} = 64\left[s^2 + s\left(s - 4m_\chi^2 - 4m_f^2\right)\cos 2\theta - 32 m_f^2 m_\chi^2 \sin^2\theta\right]. \quad \text{(pseudotensor)} \tag{E.25}$$

The final step in calculating the annihilation cross-sections is to take the non-relativistic limit where $v \ll 1$, and expand $\langle\sigma v\rangle$ in powers of v^2. Plugging Eqns. E.19, E.20, E.24, and E.25 in to Eqns. E.17 and E.18, and keeping only the leading order terms, we arrive at,

$$\left\langle \sigma^{D1} v \right\rangle = \frac{3 m_\chi^2 v^2}{8\pi M_*^6} \sum_f m_f^2 \left(1 - \frac{m_f^2}{m_\chi^2}\right)^{\frac{3}{2}}, \tag{E.26}$$

$$\left\langle \sigma^{D2} v \right\rangle = \frac{3 m_\chi^2}{2\pi M_*^6} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}}\left[\left(1 - \frac{m_f^2}{m_\chi^2}\right) + \frac{1}{8}\left(1 + 2\frac{m_f^2}{m_\chi^2}\right) v^2\right], \tag{E.27}$$

$$\left\langle \sigma^{D3} v \right\rangle = \frac{3 m_\chi^2 v^2}{8\pi M_*^6} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}}, \tag{E.28}$$

$$\left\langle \sigma^{D4} v \right\rangle = \frac{3m_\chi^2}{2\pi M_*^6} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}} \left[1 + \frac{v^2}{8} \frac{2m_\chi^2 - m_f^2}{m_\chi^2 - m_f^2} \right], \tag{E.29}$$

$$\left\langle \sigma^{D9} v \right\rangle = \frac{6m_\chi^2}{\pi M_*^6} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}} \left[2 + \frac{m_f^2}{m_\chi^2} + \frac{v^2}{24} \frac{4m_\chi^4 - 11m_f^2 m_\chi^2 + 16m_f^4}{m_\chi^2 - m_f^2} \right], \tag{E.30}$$

$$\left\langle \sigma^{D10} v \right\rangle = \frac{24m_\chi^2}{\pi M_*^6} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}} \left[1 - \frac{m_f^2}{m_\chi^2} + \frac{1}{24} \left(11m_f^2 + 4m_\chi^2 \right) v^2 \right], \tag{E.31}$$

$$\left\langle \sigma^{C1} v \right\rangle = \frac{3}{4\pi M_*^4} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}} \left[1 - \frac{m_f^2}{m_\chi^2} + \frac{3m_f^2}{8m_\chi^2} v^2 \right], \tag{E.32}$$

$$\left\langle \sigma^{C2} v \right\rangle = \frac{3}{4\pi M_*^4} \sum_f m_f^2 \sqrt{1 - \frac{m_f^2}{m_\chi^2}} \left[1 + \frac{v^2}{8} \sqrt{\frac{m_f^2}{m_\chi^2 - m_f^2}} \right]. \tag{E.33}$$

The sums are only over fermions contributing to the interaction, and a color factor of 3 has been included for quarks. Note that for each of these it's assumed that there is only a single WIMP with a single set of fermion interactions. Combining Eq. E.8 and E.5 with these cross-sections, the M_* necessary to reproduce the observed relic abundance for any m_χ can be calculated numerically. This serves as a good lower bound on the strength of each operator (i.e. an upper bound on M_*), since any given WIMP does not need to account for 100% of dark matter, but it surely can't have a *larger* density.

E.2 Indirect Detection Constraints

The goal behind indirect detection experiments is to observe the annihilations of WIMPs in distant regions of space, with a high concentration of dark matter. Since we have assumed that WIMPs were produced thermally in the Big Bang, many of the annihilation products should be SM particles, which produce visible signatures regardless of any model specifics. Most indirect searches specialize in γ-ray, anti-matter, or neutrino detection, and the leading detectors for each type of particle are Fermi-LAT [12] for γ-rays, AMS [10] and PAMELA [11] for anti-matter, and IceCube [13] for neutrinos. Typical targets for these experiments are the Milky Way galactic center and its satellite dwarf galaxies, which are known to

contain large amounts of dark matter.

The thermally averaged cross-section of WIMP annihilation through the EFT interactions are precisely those used for the relic abundance calculations, Eqns. E.26-E.33. The primary difference is that the temperature of the dark matter is colder and much less uniform today then it was at freeze-out, suppressing p-wave annihilation. A conservative estimate of $100\,\mathrm{km\,s^{-1}}$ is used here for the current dark matter velocity, which is roughly the maximum near the galactic center [151]. This will tend to overestimate the sensitivity of the indirect detection experiments we will be comparing this analysis to, since they're all focused on the galactic center and satellite dwarf galaxies. However, the precise velocity used has very little impact on the results, since it's so much smaller than the speed of light and heavily suppresses any velocity-dependent terms in the annihilation cross-section. This severely reduces the sensitivity of these experiments to the D1 and D3 operators, but for the others they can set very strong limits.

E.2.1 Cosmic Rays

There are significant populations of highly energetic particles propagating throughout our galaxy. Their primary origin is suspected to be from supernovae, but this hasn't been conclusively determined. While the overwhelming majority of these are protons, plenty of electrons, helium nuclei, and heavier atomic nuclei have been observed, with the abundance falling off rapidly with mass [152]. In addition to ordinary matter, there are also a significant number positrons, and anti-matter nuclei. Since WIMPs are not expected to be charged under any SM force, an equal amount of matter and anti-matter should be produced from any of their annihilations. Any ordinary matter produced would be completely obscured by the enormous background of non-exotic sources. However, these sources don't produce any primary anti-particles, and the anti-matter cosmic rays we observe are believed to be secondary products from collisions of primary cosmic rays with the *interstellar medium* (ISM). With their much lower background, they're excellent targets for WIMP annihilation searches.

One of the major goals of the AMS and PAMELA experiments is to observe an excess of positrons or antiprotons over the background expected from the current cosmological models. There are two key complications to these searches though. The first is intrinsic to the method, in that charged cosmic rays follow very chaotic paths while traveling through the galactic magnetic fields, in addition to various radiative effects they experience. There is absolutely no way to resolve *where* the particles came from, making it impossible to know if any observed excess came from dark matter-rich regions or from some other exotic baryonic source. The second problem, which has been recently illuminated by [153–155], is that the GALPROP model appears to have some fundamental issues with respect to charged cosmic rays. There are a huge number of free parameters, compared with the number of observable quantities, so that a very large portion of the model is fit to data. In the end, only a handful of quantities can be used as true predictions to verify the model, and of these, many show significant discrepancies between predictions and data. Although many people [156–158] have considered the possibility that these discrepancies are the result of dark matter annihilations, results from all the other dark matter searches make this increasingly unlikely. The simplest conclusion that can be drawn is that we don't yet have a sufficient understanding of the production and propagation of charged cosmic rays, and therefore these types of searches will not be considered here.

E.2.2 γ-Rays

A simpler alternative to charged cosmic rays is high energy γ-rays. These are not bent by magnetic fields, don't radiate, and for the most part, travel in straight lines from their source. Cosmic γ-rays can be produced in a number of ways, both primary and secondary. The main source of secondary γ-rays is cosmic ray electrons emitting γ-rays during interactions with the ISM (Bremmsstrahlung), the CMB (inverse-Compton scattering), or the galactic magnetic field (synchrotron radiation). This creates a large, diffuse background of gamma rays throughout the entire galaxy, on top of which primary γ-ray sources can be observed, such as supernova remnants and pulsars. An additional excess of γ-rays, originat-

ing from a region of space with a large amount of concentrated dark matter, would provide a clear signal of possible WIMP annihilations. Although WIMPs, which don't have electric charge, shouldn't annihilate directly to photons, most SM products would result in γ-rays being produced either through radiation or decay.

Fermi-LAT is a general purpose γ-ray telescope, which has recently had a number of interesting dark matter results. In 2009, using the public Fermi-LAT data, a significant gamma ray excess originating from the galactic center was found [144]. It shows up as an extended region around the center of the galaxy that emits photons in the energy range 0.2 GeV–10 GeV. This excess has been thoroughly studied, and the one thing that has become clear is that *something* exotic is producing it. Fig. E.3 shows a residual plot of the excess, along with some proposed explanations. It seems to be described very well by both WIMP annihilation [159–163] and an undiscovered population of millisecond pulsars [164–166], resulting in a large number of papers being published on both sides. The interesting thing about the dark matter interpretation though, is that it's best fit by 35 GeV WIMPs annihilating into $b\bar{b}$, as shown in Fig. E.4. Since all of the EFT operators considered here would predict $b\bar{b}$ production from WIMP annihilations, this observation becomes very relevant to this analysis. Even in the case of universal couplings, for WIMPs with mass less than the top mass of 173 GeV, the dominant annihilation channel would still be to $b\bar{b}$.

Our poor understanding of the galactic center prevents us from drawing any strong conclusions from an excess, so a much more reliable source comes from the dwarf galaxies that orbit the Milky Way. These satellite galaxies are predicted to orbit most galaxies, although $\Lambda - \mathrm{CDM}$ actually predicts many more than we see. This is an open problem, but given the discovery of ultra-faint dwarf galaxies it's possible that many of them have simply not collected enough baryonic matter to become visible [167]. Although the exact number of dwarf galaxies is uncertain, 25 have been definitively identified by Fermi-LAT [168]. Dwarf galaxies are unique in their extraordinarily high mass-to-light ratios, which suggests a larger dark matter population than the typical galaxy, in comparison to the visible baryonic matter. They are much less massive than the Milky Way though, meaning that the dark

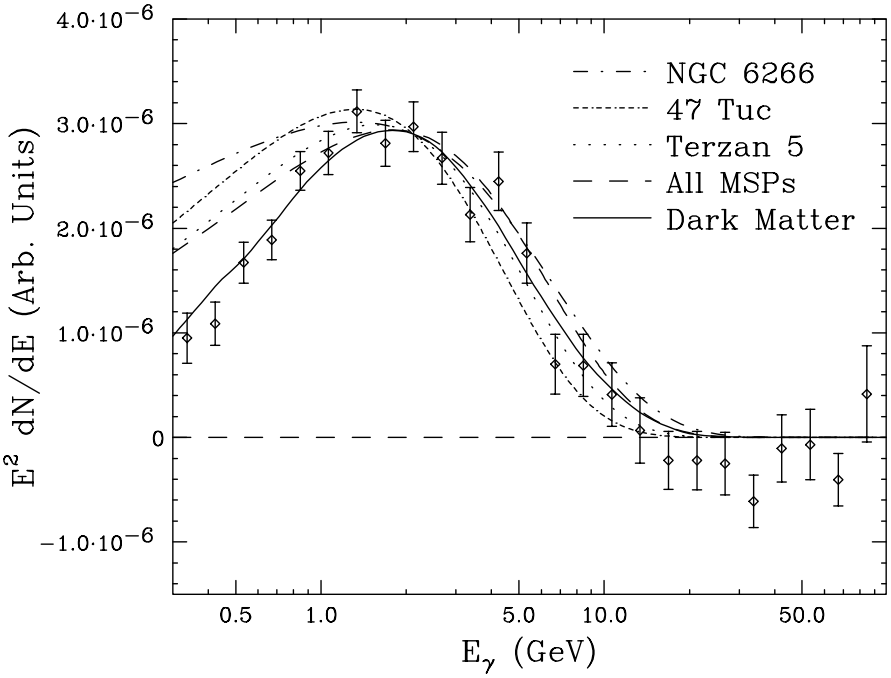

Figure E.3: Galactic center excess over expected background. [142]

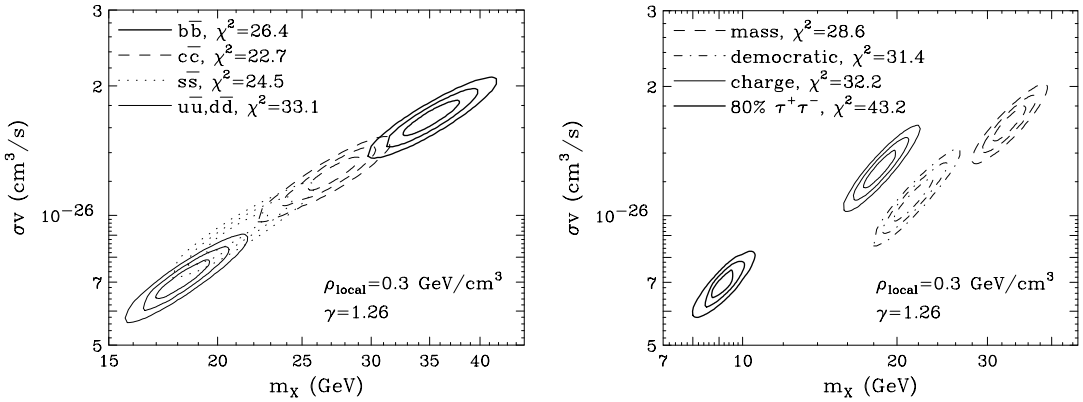

Figure E.4: WIMP annihilation modes fit to the galactic center excess. [142]

matter is less densely distributed and moves at slower velocities, which tends to make the observation of WIMP annihilations more difficult. The background γ-ray sources though, are also significantly suppressed, making them easier to model. The reduced uncertainty from our understanding of the baryonic matter sources, compensates well for the reduced signal, leading to very high sensitivity.

A survey of γ-rays coming from these regions was performed by the Fermi-LAT collaboration after 6 years of data collecting [145]. Although the 4 year data showed an excess similar to the one observed in the galactic center [168], it has dissapeared and the results set very strong limits on the annihilation cross-section of WIMPs. The limits set on the $b\bar{b}$ channel (and the harder $\tau^+\tau^-$ channel) are shown in Fig. E.5. Although the best-fit WIMP signal reported by [142] is ruled out a 90% confidence, there are plenty of similar models which are not.

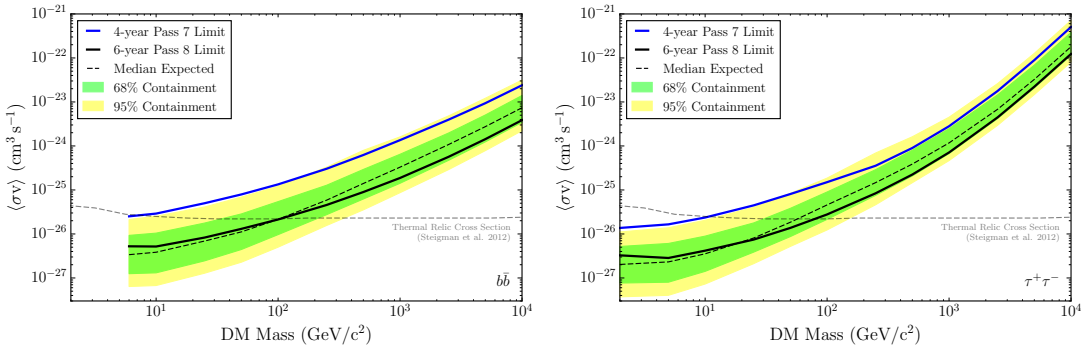

Figure E.5: WIMP annihilation limits set by Fermi LAT, showing both the 4-year results and the 6-year results. These were obtained from the combined results of 15 selected satellite dwarf galaxies. [145]

E.2.3 Neutrinos

A third option for indirect detection comes from neutrino telescopes. Neutrinos only interact via the weak force, so they're notoriously difficult to detect. However, with a large enough detector, and enough neutrinos, rare conversions into charged leptons can be observed and tracked. Muons are the most common product focused on, since unlike electrons or tau leptons, they can travel large distances with minimal energy loss. Although there is a large background source of muons produced in atmospheric cosmic ray collisions, which are far more common than neutrino-induced events, it's easily removed. A common technique used by these telescopes is to "direct" them downward into the earth, meaning that only upward-going tracks are considered. Since neutrinos are the only known particles that can travel through the entire planet with negligible energy loss, this provides an excellent

filter for charged lepton backgrounds, which typically come from atmospheric interactions above the telescope. IceCube is the general purpose neutrino telescope that currently has the strongest limits on the production of neutrinos through WIMP annihilation in the galactic center [13].

E.3 Direct Detection Constraints

Direct detection experiments search for atomic nuclei recoiling off rare interactions with WIMPs passing through the detector. Most of the allowed EFT interactions are suppressed by the WIMP velocity or momentum transfer, both of which are very low, and direct detection search typically have little to no sensitivity to these. Despite this shortcoming, these experiments have unmatched sensitivity to the interactions they *can* search for, when $m_\chi \gtrsim 10\,\mathrm{GeV}$. Given the simplicity of interpreting an observed signal, these experiments are typically treated as a kind of "gold standard" to which all other searches are compared.

The results from direct detection experiments are frequently presented as model independent limits on the cross-section of WIMP-nucleon interactions, examples of which are shown in Figs. E.6 and E.7. There are two main classes of such interactions: *spin-independent* (SI) interactions, which couple to the total number of nucleons, and *spin-dependent* (SD), which couple to the total spin of the nucleus. The sensitivity of SI interactions can be significantly increased by using heavy nuclei, while SD interactions have no such enhancement.

Following the methods in [150], the SI cross-section for WIMP interaction with a nucleus, of charge Z and mass number A, can be written out as,

$$\sigma_{\mathrm{SI}} = \frac{4\mu^2}{g^2\pi}\left(Zf_p + (A-Z)f_n\right)^2 \tag{E.34}$$

where μ is the reduced mass of the WIMP-nucleon system, g is the number of internal degrees of freedom for the WIMP, and $f_{n,p}$ are the effective nucleon couplings for neutrons

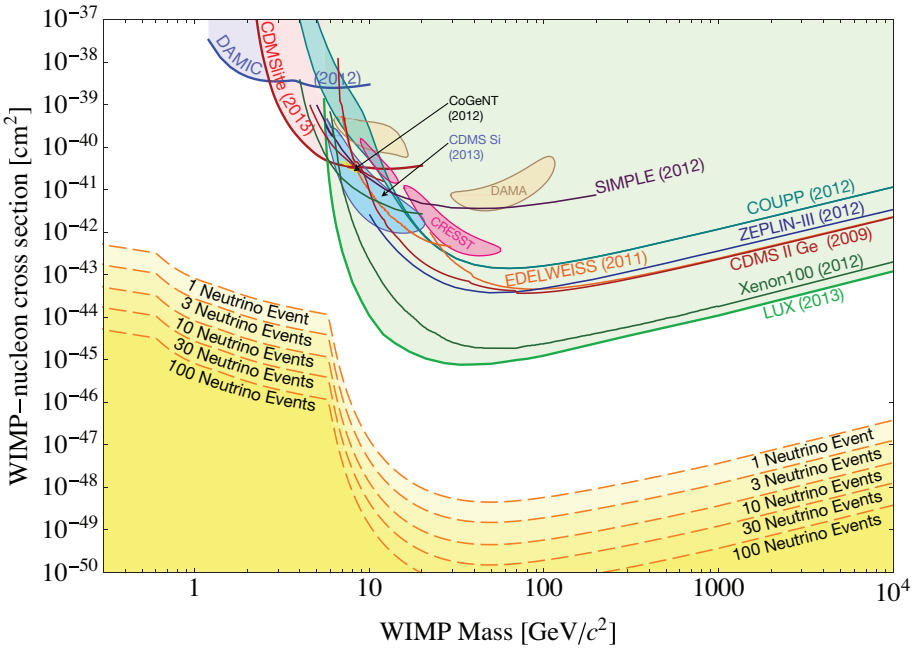

Figure E.6: Combined direct detection results for spin-independent WIMP-nucleon interactions. Exclusion curves on the cross-section, at 90% confidence, are displayed as lines for each experiment, with the total exclusion area filled. Previous hints at a WIMP signal from DAMA, CRESST, CoGeNT, and CDMS are also included as shaded regions, although they have been completely ruled out at 90% confidence. The neutrino floor is shown at the bottom, which represents the point at which direct detection experiments will begin face significant backgrounds. [169]

and protons given by,

$$f_{p,n} = \sum_{q=u,d,s} g_q f_q^{p,n} \frac{m_{p,n}}{m_q} + \frac{2}{27} f_{TG}^{p,n} \sum_{q=c,b,t} g_q \frac{m_{p,n}}{m_q}. \tag{E.35}$$

The parameters g_q are the WIMP-quark couplings from the EFT Lagrangian, while $f_q^{p,n}$ and $f_{TG}^{p,n}$ represent the contribution from each quark flavor within the nucleon. These values are obtained from lattice QCD calculations [171] and are given in Table E.1. Generally, f_p and f_n are nearly equal, so only f_p is reported. This makes σ_{SI} proportional to A^2, which can simply be factored out to obtain a WIMP-nucleon cross-section that is independent of the nuclei used by the experiment.

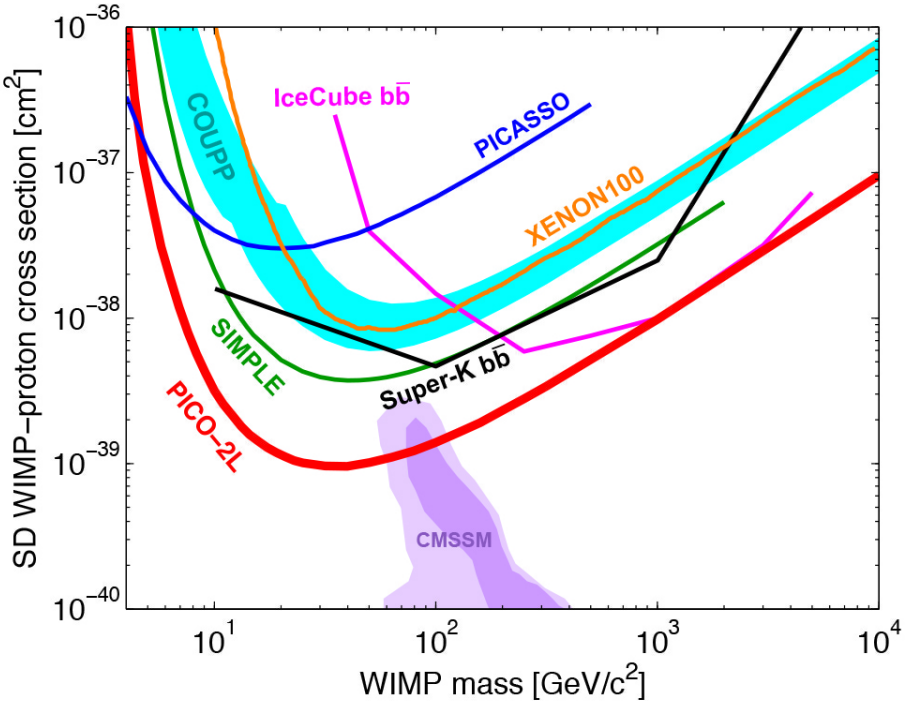

Figure E.7: Combined direct detection results for spin-independent WIMP-nucleon interaction. Exclusion curves on the cross-section, which are at 90% confidence, are displayed as lines for each experiment. For reference, the 90% exclusion curves of the indirect detection experiments Super-K and IceCube are plotted for the case of WIMP annihilation into bottom quarks. The shaded region at the bottom represents the theoretically preferred region for a supersymmetric model that is beginning to be ruled out. [170]

Form Factor	Value
f_u^p	0.028
f_d^p	0.028
f_s^p	0.0689
f_{TG}^p	0.8751

Table E.1: The contributions to the WIMP-proton interaction from each quark flavor. [171]

The only operators with unsuppressed SI cross-sections are D1 and C1, giving,

$$\sigma_{\rm SI}^{\rm D1} \approx \frac{\mu^2}{M_*^6}\left(2.86\times10^{-30}\,{\rm GeV}^4{\rm cm}^2\right), \tag{E.36}$$

$$\sigma_{\rm SI}^{\rm D1u} \approx \frac{\mu^2}{M_*^6}\left(1.12\times10^{-29}\,{\rm GeV}^4{\rm cm}^2\right), \tag{E.37}$$

$$\sigma_{\mathrm{SI}}^{\mathrm{C1}} \approx \frac{\mu^2}{M_*^4 m_\chi^2} \left(7.16 \times 10^{-31}\,\mathrm{GeV}^4\mathrm{cm}^2\right), \tag{E.38}$$

$$\sigma_{\mathrm{SI}}^{\mathrm{C1u}} \approx \frac{\mu^2}{M_*^4 m_\chi^2} \left(2.79 \times 10^{-30}\,\mathrm{GeV}^4\mathrm{cm}^2\right). \tag{E.39}$$

The tensor operator D9 does not have any velocity suppression either, but it's a SD interaction and requires different treatment. The general cross-section can be written as,

$$\sigma_{\mathrm{SD}} = 4\frac{\mu^2}{\pi} \left(a_p \left\langle S_p \right\rangle + a_n \left\langle S_n \right\rangle\right)^2 \frac{J+1}{J}, \tag{E.40}$$

where $\left\langle S_p \right\rangle$ and $\left\langle S_n \right\rangle$ are the total expected spin for the protons and neutrons in the nucleus, J is the total angular momentum of the nucleus, and $a_{n,p}$ are the nucleon couplings given by [172],

$$a_{p,n} = \sum_q g_q \Delta_q^{p,n}. \tag{E.41}$$

The parameter $\Delta_q^{p,n}$ is the spin contribution of quark type q to the nucleon, given in Table E.2, and g_q is the Lagrangian coupling. Unlike in the SI case, a_n and a_p can vary widely from each other. For most nuclei though, only one of a_p or a_n contribute significantly to the cross-section, depending on whether they contain more unpaired protons or neutrons. Each type of nucleus will usually only probe one of two possible interactions, and a separate limit must be calculated for protons and neutrons. The predicted D9 WIMP-nucleon cross-sections are then,

$$\sigma_p^{\mathrm{D9}} \approx \frac{x\mu^2}{M_*^6} \left(4.19 \times 10^{-32}\,\mathrm{GeV}^4\mathrm{cm}^2\right), \tag{E.42}$$

$$\sigma_p^{\mathrm{D9u}} \approx \frac{x\mu^2}{M_*^6} \left(2.80 \times 10^{-32}\,\mathrm{GeV}^4\mathrm{cm}^2\right), \tag{E.43}$$

$$\sigma_n^{\mathrm{D9}} \approx \frac{x\mu^2}{M_*^6} \left(7.59 \times 10^{-33}\,\mathrm{GeV}^4\mathrm{cm}^2\right), \tag{E.44}$$

$$\sigma_n^{\mathrm{D9u}} \approx \frac{x\mu^2}{M_*^6} \left(1.13 \times 10^{-32}\,\mathrm{GeV}^4\mathrm{cm}^2\right). \tag{E.45}$$

Proton	Neutron	Spin
Δ_u^p	Δ_d^n	0.84
Δ_d^p	Δ_u^n	-0.43
Δ_s^p	Δ_s^n	-0.09

Table E.2: The contribution to the overall proton spin from the three light quarks. [173]

There have been a huge number of direct detection searches performed recently, many of which have claimed to see hints of a WIMP signal [14, 17, 18, 174]. However, in 2013 the LUX collaboration released what are currently the leading direct detection limits, ruling out almost all these [16]. LUX is a time-projection chamber that uses 368 kg of liquid xenon as a target for WIMP collisions. Liquid xenon is scintillating, very dense, and has no naturally occurring radioactive isotopes. These properties make it the ideal material for a direct detection search, which is typically plagued by background radiation from both outside sources and the detector itself. The detector is capable of reconstructing the three-dimensional location of any event using arrays of photomultiplier tubes, providing further background rejection. LUX is scheduled to release new results of 300 live days, on top of the 86 live day run in 2013, later this year. The next generation of detectors will be sensitive to cross-sections about 100 times lower than LUX, and will begin to hit the *neutrino floor* that poses a number of technical problems for future searches.

Although direct detection experiments have incredible sensitivity to SI interactions, they're not nearly as strong for SD ones, and usually require dedicated detectors to place competitive limits. Before 2015, COUPP, SIMPLE, PICASSO, and XENON100 were the leading SD detectors, each with leading limits in different regions of WIMP masses. In 2013 though, COUPP and PICASSO combined their efforts to create the PICO collaboration, which in 2015 released the currently leading SD limits from the PICO-2L detector [170]. Similar to PICASSO and COUPP, PICO-2L is a 2 liter bubble chamber that uses super-heated liquid to detect any WIMP interactions within. Nuclear recoils cause the production

of bubbles in the liquid, which are then photographed and analyzed to find events caused by WIMP collisions. These experiments typically use lighter nuclei, so the SI limits set are much weaker.

E.4 Collider Constraints

Since the EFT framework is standard at ATLAS and CMS, comparing to other collider searches would usually be very simple. However, there are some complications involving the modified operators considered here. Most analyses assume universal couplings between up-type and down-type quarks, which doubles the number of terms for each interaction, and the mass dependence of the operators becomes inconsistent across analyses. On top of that, the D9 operator is almost *always* used with mass-independent couplings, making it a significantly different operator than the one considered here, even with universal couplings. The scalar interactions are often considered with mass-independent couplings too, but most analyses include the MFV variants. A limit set on an M_* value here does not always translate well to other analyses, similarly to the way M_* values between different operators can't be compared.

To demonstrate this, take a limit $M_* > M$, placed on an operator, $\mathcal{O}$, by some mono-X analysis. We wish to translate this to a new limit $M_* > M'$ placed on another operator $\mathcal{O}'$, by the same analysis. The limit placed on the cross-section of $\mathcal{O}$, in terms of M, is,

$$\epsilon_{\mathcal{O}}\sigma_{\mathcal{O}}\left(\frac{1\,\text{TeV}}{M}\right)^{y} < \frac{N}{L}, \tag{E.46}$$

where N is the maximum number of signal events consistent with the observation, L is the total integrated luminosity, $\epsilon_{\mathcal{O}}$ is the efficiency of the signal selection, $\sigma_{\mathcal{O}}$ is the cross-section of the operator at $M_* = 1\,\text{TeV}$, and y is either 4 or 6 depending on the operator. Since the limit on N is model-independent, we can use this to also set a limit on the cross-section of

$\mathcal{O}'$. Comparing these, we find that,

$$M' = \left(\frac{\epsilon_{\mathcal{O}'}\sigma_{\mathcal{O}'}}{\epsilon_{\mathcal{O}}\sigma_{\mathcal{O}}}\right)^{\frac{1}{y'}} \left(\frac{M}{1\,\mathrm{TeV}}\right)^{\frac{y}{y'}} \mathrm{TeV}. \tag{E.47}$$

While the ratio between the two cross-sections are easily estimated using a Monte Carlo generator such as Madgraph and $\epsilon_{\mathcal{O}}$ is usually reported by the analysis, $\epsilon_{\mathcal{O}'}$ is much more difficult to obtain. Short of rerunning complete simulations of $\mathcal{O}'$ with the same filters used on the $\mathcal{O}$ samples, and then repeating the entire analysis, there is no good method for getting a reliable estimate, in general. However, assumptions can be made that do allow some comparisons to be made of the results from different analyses.

As an example, take the D9 operator used by the mono-jet group, where it has universal couplings and no mass dependence. The dominant *visible* interaction for this operator contains a gluon radiating off a quark pair that annihilate to two WIMPs. In contrast, the dominant interaction for the D9 operator considered here is a $b\bar{b}$ pair produced together with a WIMP pair. The acceptance and efficiency for selecting a single light jet compared to two heavy jets is highly dependent on the analysis details. Simply knowing the fraction of *their* signal that makes it through the selection tells us nothing about the efficiency for *our* signal. However, the higher we set $\epsilon_{\mathcal{O}'}$ to be, the more optimistic we are about the sensitivity of their search towards our signal. A conservative approach for comparing this analysis with others then, is simply to maximize $\epsilon_{\mathcal{O}'}$ within some reasonable range. Since the other analyses are optimized to a similar but distinct process from our signal, it's generally safe to assume that $\epsilon_{\mathcal{O}'} < \epsilon_{\mathcal{O}}$, and therefore the dependence of M' on the two efficiencies can be eliminated by using the optimistic $\epsilon_{\mathcal{O}'} \equiv \epsilon_{\mathcal{O}}$. All of the comparisons to other LHC analyses in Section 5.8 use these methods, along with the results of Appendix F.

Appendix F

Signal Simulation

Only the samples listed in Section 5.1 were put through the official simulations of the detector and the underlying event, which aren't useful in examining variations to the model parameters. Various truth-level studies were performed on the signal samples, using Madgraph [88] and Pythia8 [175], in order to look more closely at the modifications to the interactions discussed here, and the systematic uncertainties arising from the generator. A general assumption, tested using Madgraph's simplified detector simulation, was made that the models used by the official samples were similar enough to the modified processes that detector-level effects wouldn't vary much. When this holds, we can use truth-level differences to estimate potential analysis-level variations.

Fig. F.1 plots the normalized $\not{E}_T$ distributions of all 8 EFT operators at four representative mass points. From this, it's inferred that in high energy collisions, these operators are degenerate and can be described using a reduced set of three. Although most of these operators are identical in total cross-section as well, the D2 and D4 operators require a different normalization than the D1 and D3 operators. The differences are summarized in Table F.1, which lists the total cross-section predicted by Madgraph and Pythia at each of the relevant mass points. To estimate the D2 or D4 signals then, the D1 samples simply need to be scaled by the ratio of their cross-sections. The remaining operators are summarized in Tables F.2 and F.3, with calculations for the cross-sections of the reduced set. Three different filters were used to quantify the flavor dependence of the operators, corresponding to light jet, b-jet, and $t\bar{t}$ production.

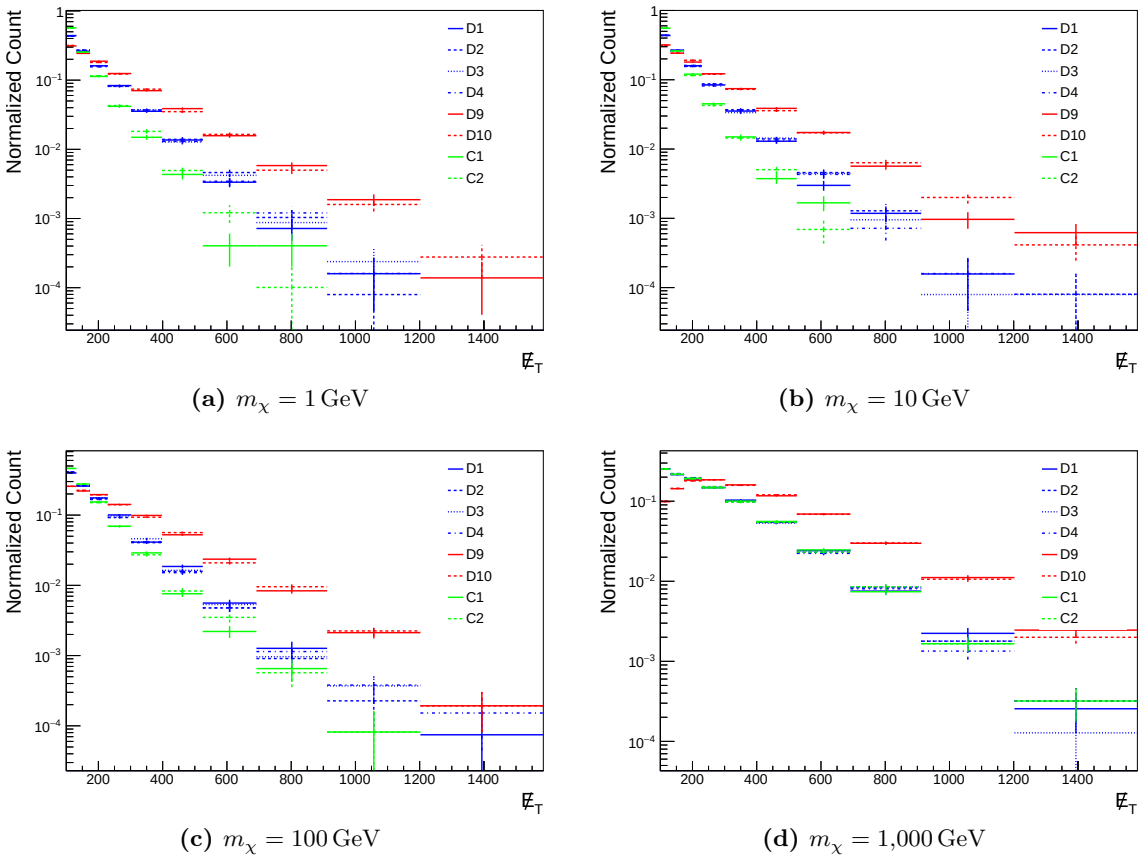

(a) $m_\chi = 1\,\text{GeV}$ **(b)** $m_\chi = 10\,\text{GeV}$

(c) $m_\chi = 100\,\text{GeV}$ **(d)** $m_\chi = 1{,}000\,\text{GeV}$

Figure F.1: Truth-level $\not\!\!E_T$ distributions for each operator. These plots show that all 8 signal operators have similar kinematics as either D1, D9, or C1.

Mass [GeV]	σ_{D1} [pb]	σ_{D2} [pb]	σ_{D3} [pb]	σ_{D4} [pb]
1	$(1.26 \pm 0.01) \times 10^{-9}$	$(1.27 \pm 0.01) \times 10^{-9}$	$(1.28 \pm 0.01) \times 10^{-9}$	$(1.27 \pm 0.01) \times 10^{-9}$
10	$(1.26 \pm 0.01) \times 10^{-9}$	$(1.25 \pm 0.01) \times 10^{-9}$	$(1.26 \pm 0.01) \times 10^{-9}$	$(1.27 \pm 0.01) \times 10^{-9}$
50	$(1.17 \pm 0.01) \times 10^{-9}$	$(1.22 \pm 0.01) \times 10^{-9}$	$(1.17 \pm 0.01) \times 10^{-9}$	$(1.24 \pm 0.01) \times 10^{-9}$
100	$(9.89 \pm 0.08) \times 10^{-10}$	$(1.13 \pm 0.01) \times 10^{-9}$	$(9.85 \pm 0.08) \times 10^{-10}$	$(1.13 \pm 0.01) \times 10^{-9}$
200	$(5.95 \pm 0.05) \times 10^{-10}$	$(8.24 \pm 0.07) \times 10^{-10}$	$(5.93 \pm 0.05) \times 10^{-10}$	$(8.24 \pm 0.07) \times 10^{-10}$
400	$(1.6 \pm 0.1) \times 10^{-10}$	$(3.11 \pm 0.03) \times 10^{-10}$	$(1.66 \pm 0.01) \times 10^{-10}$	$(3.11 \pm 0.03) \times 10^{-10}$
700	$(1.97 \pm 0.02) \times 10^{-11}$	$(5.3 \pm 0.5) \times 10^{-11}$	$(1.97 \pm 0.02) \times 10^{-11}$	$(5.35 \pm 0.05) \times 10^{-11}$
1000	$(2.15 \pm 0.02) \times 10^{-12}$	$(7.54 \pm 0.07) \times 10^{-12}$	$(2.16 \pm 0.02) \times 10^{-12}$	$(7.56 \pm 0.07) \times 10^{-12}$
1300	$(2.27 \pm 0.02) \times 10^{-13}$	$(9.56 \pm 0.09) \times 10^{-13}$	$(2.24 \pm 0.02) \times 10^{-13}$	$(9.7 \pm 0.9) \times 10^{-13}$

Table F.1: Cross-sections calculated by Madgraph and Pythia of the D1-like processes

$\mathcal{O}$	m_χ (GeV)	σ_{light} (pb)	σ_b (pb)	$\sigma_{t\bar{t}}$ (pb)
C1	1	$(6.40 \pm 0.04) \times 10^{-10}$	$(3.23 \pm 0.02) \times 10^{-9}$	$(5.65 \pm 0.03) \times 10^{-13}$
C1	10	$(6.33 \pm 0.04) \times 10^{-10}$	$(3.08 \pm 0.02) \times 10^{-9}$	$(5.56 \pm 0.03) \times 10^{-13}$
C1	50	$(5.53 \pm 0.04) \times 10^{-10}$	$(2.59 \pm 0.02) \times 10^{-9}$	$(4.73 \pm 0.03) \times 10^{-13}$
C1	100	$(4.07 \pm 0.03) \times 10^{-10}$	$(1.76 \pm 0.01) \times 10^{-9}$	$(3.49 \pm 0.02) \times 10^{-13}$
C1	200	$(1.87 \pm 0.01) \times 10^{-10}$	$(7.09 \pm 0.06) \times 10^{-10}$	$(1.62 \pm 0.01) \times 10^{-13}$
C1	400	$(3.30 \pm 0.02) \times 10^{-11}$	$(1.15 \pm 0.01) \times 10^{-10}$	$(3.02 \pm 0.02) \times 10^{-14}$
C1	700	$(2.46 \pm 0.02) \times 10^{-12}$	$(8.49 \pm 0.07) \times 10^{-12}$	$(2.40 \pm 0.02) \times 10^{-15}$
C1	1000	$(2.00 \pm 0.02) \times 10^{-13}$	$(6.66 \pm 0.06) \times 10^{-13}$	$(2.00 \pm 0.01) \times 10^{-16}$
C1	1300	$(1.72 \pm 0.01) \times 10^{-14}$	$(5.51 \pm 0.05) \times 10^{-14}$	$(1.77 \pm 0.01) \times 10^{-17}$
D1	1	$(3.20 \pm 0.02) \times 10^{-10}$	$(1.26 \pm 0.01) \times 10^{-9}$	$(2.85 \pm 0.02) \times 10^{-13}$
D1	10	$(3.21 \pm 0.02) \times 10^{-10}$	$(1.26 \pm 0.01) \times 10^{-9}$	$(2.86 \pm 0.02) \times 10^{-13}$
D1	50	$(3.05 \pm 0.02) \times 10^{-10}$	$(1.17 \pm 0.01) \times 10^{-9}$	$(2.70 \pm 0.02) \times 10^{-13}$
D1	100	$(2.59 \pm 0.02) \times 10^{-10}$	$(9.89 \pm 0.08) \times 10^{-10}$	$(2.35 \pm 0.01) \times 10^{-13}$
D1	200	$(1.64 \pm 0.01) \times 10^{-10}$	$(5.95 \pm 0.05) \times 10^{-10}$	$(1.49 \pm 0.01) \times 10^{-13}$
D1	400	$(4.63 \pm 0.04) \times 10^{-11}$	$(1.60 \pm 0.01) \times 10^{-10}$	$(4.44 \pm 0.03) \times 10^{-14}$
D1	700	$(5.74 \pm 0.04) \times 10^{-12}$	$(1.97 \pm 0.02) \times 10^{-11}$	$(5.79 \pm 0.04) \times 10^{-15}$
D1	1000	$(6.60 \pm 0.06) \times 10^{-13}$	$(2.15 \pm 0.02) \times 10^{-12}$	$(6.71 \pm 0.05) \times 10^{-16}$
D1	1300	$(7.42 \pm 0.06) \times 10^{-14}$	$(2.27 \pm 0.02) \times 10^{-13}$	$(7.54 \pm 0.06) \times 10^{-17}$
D9	1	$(2.45 \pm 0.02) \times 10^{-9}$	$(3.98 \pm 0.03) \times 10^{-8}$	$(4.61 \pm 0.03) \times 10^{-12}$
D9	10	$(2.44 \pm 0.02) \times 10^{-9}$	$(3.94 \pm 0.03) \times 10^{-8}$	$(4.59 \pm 0.02) \times 10^{-12}$
D9	50	$(2.28 \pm 0.02) \times 10^{-9}$	$(3.70 \pm 0.03) \times 10^{-8}$	$(4.28 \pm 0.02) \times 10^{-12}$
D9	100	$(1.91 \pm 0.01) \times 10^{-9}$	$(3.04 \pm 0.02) \times 10^{-8}$	$(3.62 \pm 0.02) \times 10^{-12}$
D9	200	$(1.17 \pm 0.01) \times 10^{-9}$	$(1.75 \pm 0.01) \times 10^{-8}$	$(2.17 \pm 0.01) \times 10^{-12}$
D9	400	$(3.71 \pm 0.03) \times 10^{-10}$	$(4.41 \pm 0.04) \times 10^{-9}$	$(6.17 \pm 0.04) \times 10^{-13}$
D9	700	$(5.93 \pm 0.05) \times 10^{-11}$	$(5.24 \pm 0.05) \times 10^{-10}$	$(8.53 \pm 0.05) \times 10^{-14}$
D9	1000	$(8.46 \pm 0.07) \times 10^{-12}$	$(5.96 \pm 0.06) \times 10^{-11}$	$(1.11 \pm 0.01) \times 10^{-14}$
D9	1300	$(1.16 \pm 0.01) \times 10^{-12}$	$(6.59 \pm 0.06) \times 10^{-12}$	$(1.42 \pm 0.01) \times 10^{-15}$

Table F.2: Cross-sections calculated by Madgraph and Pythia for each signal process, with filters on light jets, b-jets, and top pairs.

$\mathcal{O}$	m_χ (GeV)	σ_{light} (pb)	σ_b (pb)	$\sigma_{t\bar{t}}$ (pb)
C1u	1	$(1.15 \pm 0.01) \times 10^{-9}$	$(3.19 \pm 0.02) \times 10^{-9}$	$(1.42 \pm 0.01) \times 10^{-7}$
C1u	10	$(1.12 \pm 0.01) \times 10^{-9}$	$(3.09 \pm 0.02) \times 10^{-9}$	$(1.42 \pm 0.01) \times 10^{-7}$
C1u	50	$(9.54 \pm 0.07) \times 10^{-10}$	$(2.55 \pm 0.02) \times 10^{-9}$	$(1.10 \pm 0.01) \times 10^{-7}$
C1u	100	$(6.92 \pm 0.05) \times 10^{-10}$	$(1.76 \pm 0.01) \times 10^{-9}$	$(7.11 \pm 0.05) \times 10^{-8}$
C1u	200	$(3.09 \pm 0.02) \times 10^{-10}$	$(7.21 \pm 0.06) \times 10^{-10}$	$(2.79 \pm 0.02) \times 10^{-8}$
C1u	400	$(5.26 \pm 0.04) \times 10^{-11}$	$(1.16 \pm 0.01) \times 10^{-10}$	$(4.88 \pm 0.04) \times 10^{-9}$
C1u	700	$(3.97 \pm 0.03) \times 10^{-12}$	$(8.43 \pm 0.07) \times 10^{-12}$	$(3.90 \pm 0.04) \times 10^{-10}$
C1u	1000	$(3.11 \pm 0.03) \times 10^{-13}$	$(6.72 \pm 0.06) \times 10^{-13}$	$(3.14 \pm 0.03) \times 10^{-11}$
C1u	1300	$(2.64 \pm 0.02) \times 10^{-14}$	$(5.56 \pm 0.05) \times 10^{-14}$	$(2.33 \pm 0.02) \times 10^{-12}$
D1u	1	$(5.34 \pm 0.04) \times 10^{-10}$	$(1.28 \pm 0.01) \times 10^{-9}$	$(5.06 \pm 0.04) \times 10^{-8}$
D1u	10	$(5.29 \pm 0.04) \times 10^{-10}$	$(1.27 \pm 0.01) \times 10^{-9}$	$(5.07 \pm 0.04) \times 10^{-8}$
D1u	50	$(4.96 \pm 0.04) \times 10^{-10}$	$(1.18 \pm 0.01) \times 10^{-9}$	$(4.74 \pm 0.04) \times 10^{-8}$
D1u	100	$(4.30 \pm 0.03) \times 10^{-10}$	$(9.87 \pm 0.08) \times 10^{-10}$	$(3.89 \pm 0.03) \times 10^{-8}$
D1u	200	$(2.67 \pm 0.02) \times 10^{-10}$	$(5.93 \pm 0.05) \times 10^{-10}$	$(2.36 \pm 0.02) \times 10^{-8}$
D1u	400	$(7.48 \pm 0.06) \times 10^{-11}$	$(1.64 \pm 0.01) \times 10^{-10}$	$(7.17 \pm 0.06) \times 10^{-9}$
D1u	700	$(9.13 \pm 0.07) \times 10^{-12}$	$(1.96 \pm 0.02) \times 10^{-11}$	$(9.12 \pm 0.09) \times 10^{-10}$
D1u	1000	$(1.02 \pm 0.01) \times 10^{-12}$	$(2.18 \pm 0.02) \times 10^{-12}$	$(9.75 \pm 0.09) \times 10^{-11}$
D1u	1300	$(1.12 \pm 0.01) \times 10^{-13}$	$(2.26 \pm 0.02) \times 10^{-13}$	$(8.81 \pm 0.09) \times 10^{-12}$
D9u	1	$(7.39 \pm 0.05) \times 10^{-9}$	$(3.97 \pm 0.03) \times 10^{-8}$	$(2.40 \pm 0.02) \times 10^{-6}$
D9u	10	$(7.49 \pm 0.05) \times 10^{-9}$	$(3.98 \pm 0.03) \times 10^{-8}$	$(2.45 \pm 0.02) \times 10^{-6}$
D9u	50	$(6.93 \pm 0.05) \times 10^{-9}$	$(3.67 \pm 0.03) \times 10^{-8}$	$(2.34 \pm 0.02) \times 10^{-6}$
D9u	100	$(5.81 \pm 0.04) \times 10^{-9}$	$(3.06 \pm 0.02) \times 10^{-8}$	$(1.89 \pm 0.02) \times 10^{-6}$
D9u	200	$(3.40 \pm 0.02) \times 10^{-9}$	$(1.74 \pm 0.01) \times 10^{-8}$	$(1.09 \pm 0.01) \times 10^{-6}$
D9u	400	$(9.59 \pm 0.08) \times 10^{-10}$	$(4.40 \pm 0.04) \times 10^{-9}$	$(3.02 \pm 0.03) \times 10^{-7}$
D9u	700	$(1.28 \pm 0.01) \times 10^{-10}$	$(5.17 \pm 0.05) \times 10^{-10}$	$(3.53 \pm 0.03) \times 10^{-8}$
D9u	1000	$(1.67 \pm 0.01) \times 10^{-11}$	$(6.00 \pm 0.06) \times 10^{-11}$	$(3.86 \pm 0.04) \times 10^{-9}$
D9u	1300	$(2.10 \pm 0.02) \times 10^{-12}$	$(6.52 \pm 0.06) \times 10^{-12}$	$(3.77 \pm 0.04) \times 10^{-10}$

Table F.3: Cross-sections calculated by Madgraph and Pythia for each signal process with universal couplings, with filters on light jets, b-jets, and top pairs.

Bibliography

[1] F. Zwicky. The redshift of extragalactic nebulae. *Helvetica Physica Acta*, 6:110–127, 1933.

[2] Vera C. Rubin and Jr. Ford, W. Kent. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *The Astrophysical Journal*, 159:379–403, 1970.

[3] Douglas Clowe, Marusa Bradac, Anthony H. Gonzalez, et al. A direct empirical proof of the existence of dark matter. *The Astrophysical Journal*, 648:L109–L113, 2006.

[4] George R. Blumenthal, S.M. Faber, Joel R. Primack, et al. Formation of Galaxies and Large Scale Structure with Cold Dark Matter. *Nature*, 311:517–525, 1984.

[5] N. Jarosik, C. L. Bennett, J. Dunkley, et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. 192:14, February 2011.

[6] Supernova Cosmology Project Collaboration. New constraints on Omega(M), Omega(lambda), and w from an independent set of eleven high-redshift supernovae observed with HST. *The Astrophysical Journal*, 598:102, 2003.

[7] EROS-2 Collaboration. Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *Astronomy and Astrophysics*, 469:387–404, 2007.

[8] Mordehai Milgrom and Jacob Bekenstein. The modified newtonian dynamics as an alternative to hidden matter. In *Dark matter in the universe*, volume 117, pages 319–330, 1987.

[9] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Physical Review Letters*, 38:1440–1443, Jun 1977.

[10] AMS Collaboration. High statistics measurement of the positron fraction in primary cosmic rays of 0.5–500 gev with the alpha magnetic spectrometer on the international space station. *Physical Review Letters*, 113:121101, Sep 2014.

[11] O. Adriani, G.C. Barbarino, G.A. Bazilevskaya, et al. A new measurement of the antiproton-to-proton flux ratio up to 100 GeV in the cosmic radiation. *Physical Review Letters*, 102:051101, 2009.

[12] W. B. Atwood, A. A. Abdo, M. Ackermann, et al. The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission. *The Astrophysical Journal*, 697:1071–1102, June 2009.

[13] IceCube Collaboration. IceCube Search for Dark Matter Annihilation in nearby Galaxies and Galaxy Clusters. *Physical Review D*, 88(12):122001, 2013.

[14] CDMS Collaboration. Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *Physical Review Letters*, 111:251301, 2013.

[15] XENON100 Collaboration. Dark Matter Results from 225 Live Days of XENON100 Data. *Physical Review Letters*, 109:181301, 2012.

[16] LUX Collaboration. First results from the LUX dark matter experiment at the Sanford Underground Research Facility. *Physical Review Letters*, 112:091303, 2014.

[17] CRESST-II Collaboration. Results on low mass WIMPs using an upgraded CRESST-II detector. 2014.

[18] C. E. Aalseth, P. S. Barbeau, J. Colaresi, et al. CoGeNT: A search for low-mass dark matter using p-type point contact germanium detectors. *Physical Review D*, 88(1):012002, July 2013.

[19] COUPP Collaboration. First Dark Matter Search Results from a 4-kg CF_3I Bubble Chamber Operated in a Deep Underground Site. *Physical Review D*, 86:052001, 2012.

[20] M. Felizardo, T. A. Girard, T. Morlat, et al. The simple phase ii dark matter search. *Physical Review D*, 89:072013, Apr 2014.

[21] PICASSO Collaboration. Constraints on Low-Mass WIMP Interactions on ^{19}F from PICASSO. *Physics Letters B*, 711:153–161, 2012.

[22] ATLAS Collaboration. Search for dark matter in events with heavy quarks and missing transverse momentum in *pp* collisions with the ATLAS detector. *The European Physical Journal*, C75(2):92, 2015.

[23] Albert Einstein et al. On the electrodynamics of moving bodies. *Annalen der Physik*, 17(891):50, 1905.

[24] Jerry B Marion and Stephen T Thornton. Classical mechanics of particles and systems. *Saunders Colledge Pub., quarta ediçao*, 1995.

[25] David Jeffrey Griffiths and Reed College. *Introduction to electrodynamics*, volume 3. prentice Hall Upper Saddle River, NJ, 1999.

[26] J Preskill. *Field Theory Lecture Notes*. 1986.

[27] Julian Seymour Schwinger, WY Tsai, Lester L De Raad, et al. *Classical electrodynamics*. Perseus, 1998.

[28] M.P. Hobson, G. Efstathiou, A.N. Lasenby. *General Relativity: An Introduction for Physicists*. Cambridge University Press, 2006.

[29] M. E. Peskin, D. V. Shroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.

[30] E. Noether. Invariante variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1918:235–257, 1918.

[31] R. P. Feynman. The Theory of Positrons. *Physical Review*, 76:749–759, September 1949.

[32] E. Eichten, Kenneth D. Lane, and Michael E. Peskin. New Tests for Quark and Lepton Substructure. *Physical Review Letters*, 50:811–814, 1983.

[33] TOPAZ Collaboration. Measurement of the electromagnetic coupling at large momentum transfer. *Physical Review Letters*, 78:424–427, Jan 1997.

[34] Kenzo Nakamura, Particle Data Group, et al. Review of particle physics. *Journal of Physics G*, 37(7A):075021, 2010.

[35] Andrew Purcell. Go on a particle quest at the first CERN webfest. Le premier webfest du CERN se lance à la conquête des particules. (BUL-NA-2012-269. 35/2012):10, Aug 2012.

[36] Peter Ware Higgs. Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13:508–509, 1964.

[37] J. Goldstone, A. Salam, and S. Weinberg. Broken Symmetries. *Physical Review*, 127:965–970, August 1962.

[38] A. D. Sakharov. SPECIAL ISSUE: Violation of CP in variance, C asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi*, 34:392–393, May 1991.

[39] J. H. Christenson, J. W. Cronin, V. L. Fitch, et al. Evidence for the 2π Decay of the K20 Meson. *Physical Review Letters*, 13:138–140, July 1964.

[40] A. Alavi-Harati, I. F. Albuquerque, T. Alexopoulos, et al. Observation of Direct CP Violation in $K_{S,L} \to \pi\pi$ Decays. *Physical Review Letters*, 83:22–27, July 1999.

[41] V. Fanti and S. Palestini. A new measurement of direct CP violation in two pion decays of the neutral kaon. *Physics Letters B*, 465:335–348, November 1999.

[42] Ken Croswell. The alchemy of the heavens. 1, 1996.

[43] Jessica Goodman, Masahiro Ibe, Arvind Rajaraman, et al. Constraints on Dark Matter from Colliders. *Physical Review D*, 82:116010, 2010.

[44] John March-Russell, James Unwin, and Stephen M. West. Closing in on Asymmetric Dark Matter I: Model independent limits for interactions with quarks. *The Journal of High Energy Physics*, 08:029, 2012.

[45] Giorgio Busoni, Andrea De Simone, Enrico Morgante, et al. On the Validity of the Effective Field Theory for Dark Matter Searches at the LHC. *Physics Letters B*, 728:412–421, 2014.

[46] G. D'Ambrosio, G.F. Giudice, G. Isidori, et al. Minimal flavor violation: An Effective field theory approach. *Nuclear Physics*, B645:155–187, 2002.

[47] J Abdallah, K A Assamagan, D Berge, et al. Search for new phenomena with mono-jet plus missing transverse energy signature in proton-proton collisions at $\sqrt{s} = 8\,\text{TeV}$ with the ATLAS detector. Technical Report ATL-COM-PHYS-2013-1578, CERN, Geneva, Nov 2013.

[48] L Carminati, D Cavalli, M-H Genest, et al. Search for new phenomena with the ATLAS detector in monophoton events from proton-proton collisions at sqrt(s)=8TeV. Technical Report ATL-PHYS-INT-2014-036, CERN, Geneva, Nov 2014.

[49] ATLAS Collaboration. Search for new particles in events with one lepton and missing transverse momentum in *pp* collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *The Journal of High Energy Physics*, 1409:037, 2014.

[50] ATLAS Collaboration. Search for dark matter in events with a hadronically decaying W or Z boson and missing transverse momentum in *pp* collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Physical Review Letters*, 112(4):041802, 2014.

[51] ATLAS Collaboration. Search for dark matter in events with a Z boson and missing transverse momentum in pp collisions at $\sqrt{s}$=8 TeV with the ATLAS detector. *Physical Review D*, 90:012004, 2014.

[52] Jonathan Allday. *Quarks, leptons and the Big Bang*. CRC Press, 2001.

[53] Planck Collaboration. Planck 2015 results. XIII. Cosmological parameters. 2015.

[54] H. S. Leff. Teaching the photon gas in introductory physics. *The American Journal of Physics*, 70:792–797, August 2002.

[55] Gary Steigman, Basudeb Dasgupta, and John F. Beacom. Precise Relic WIMP Abundance and its Impact on Searches for Dark Matter Annihilation. *Physical Review D*, 86:023506, 2012.

[56] Barbara Ryden. *Introduction to cosmology*, volume 1. 2003.

[57] Alan H Guth and Paul J Steinhardt. The inflationary universe. *Scientific American*, 250(5):116–128, 1984.

[58] Robert J Scherrer and Michael S Turner. On the relic, cosmic abundance of stable, weakly interacting massive particles. *Physical Review D*, 33(6):1585, 1986.

[59] Kris Heyde. *Basic ideas and concepts in nuclear physics: an introductory approach*. CRC Press, 2004.

[60] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr. Extended rotation curves of high-luminosity spiral galaxies. IV - Systematic dynamical properties, SA through SC. *The Astrophysical Journal, Letters*, 225:L107–L111, November 1978.

[61] M. S. Roberts and R. N. Whitehurst. The rotation curve and geometry of M31 at large galactocentric distances. *The Astrophysical Journal*, 201:327–346, October 1975.

[62] MACHO Collaboration. The MACHO project: Microlensing results from 5.7 years of LMC observations. *The Astrophysical Journal*, 542:281–307, 2000.

[63] David S. Graff and Katherine Freese. Analysis of a hubble space telescope search for red dwarfs: limits on baryonic matter in the galactic halo. *The Astrophysical Journal*, 456:L49, 1996.

[64] VIRGO Collaboration. Halo and Galaxy Formation Histories from the Millennium Simulation: Public release of a VO-oriented and SQL-queryable database for studying the evolution of galaxies in the Lambda-CDM cosmogony. 2006.

[65] M Colless, GB Dalton, J Bland-Hawthorn, et al. The 2dfgrs galaxy redshift survey. *Monthly Notices of the Royal Astronomical Society*, 328(astro-ph/0106498):1039, 2001.

[66] SW Allen, RW Schmidt, and AC Fabian. Cosmological constraints from the x-ray gas mass fraction in relaxed lensing clusters observed with chandra. *Monthly Notices of the Royal Astronomical Society*, 334(2):L11–L15, 2002.

[67] D. J. Fixsen. The Temperature of the Cosmic Microwave Background. *The Astrophysical Journal*, 707:916–920, December 2009.

[68] D. J. Eisenstein, W. Hu, and M. Tegmark. Cosmic Complementarity: H_0 and Ω_m from Combining Cosmic Microwave Background Experiments and Redshift Surveys. *The Astrophysical Journal, Letters*, 504:L57–L60, September 1998.

[69] Bruno Leibundgut. Type Ia supernovae. *The Astronomy and Astrophysics Review*, 10:179, 2000.

[70] S. Chandrasekhar. The Maximum Mass of Ideal White Dwarfs. *The Astrophysical Journal*, 74:81, July 1931.

[71] S Perlmutter, G Aldering, G Goldhaber, et al. Supernova cosmology project. *The Astrophysical Journal*, 517:565, 1999.

[72] Lyndon Evans and Philip Bryant. LHC Machine. *The Journal of Instrumentation*, 3:S08001, 2008.

[73] John D Cockcroft and Ernest Thomas Sinton Walton. Disintegration of lithium by swift protons. *Nature*, 129(3261):649, 1932.

[74] S. Myers and E. Picasso. The Design, construction and commissioning of the CERN Large Electron Positron collider. *Contemporary Physics*, 31:387–403, 1990.

[75] Robert Rathbun Wilson. The Tevatron. *Physics Today*, 30N10:23–30, 1977.

[76] CMS Collaboration. The CMS experiment at the CERN LHC. *The Journal of Instrumentation*, 3:S08004, 2008.

[77] ALICE Collaboration. The ALICE experiment at the CERN LHC. *The Journal of Instrumentation*, 3:S08002, 2008.

[78] LHCb Collaboration. The LHCb Detector at the LHC. *The Journal of Instrumentation*, 3:S08005, 2008.

[79] Sue A. Cano. Proton substructure.

[80] A.D. Martin, W.J. Stirling, R.S. Thorne, et al. Parton distributions for the LHC. *The European Physical Journal*, C63:189–285, 2009.

[81] Stefan Höche. Introduction to parton-shower event generators. In *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders (TASI 2014) Boulder, Colorado, June 2-27, 2014*, 2014.

[82] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *The Journal of Instrumentation*, 3:S08003, 2008.

[83] ATLAS Collaboration. Performance of the ATLAS Trigger System in 2010. *The European Physical Journal*, C72:1849, 2012.

[84] J. Pumplin et al. New generation of parton distributions with uncertainties from global QCD analysis. *The Journal of High Energy Physics*, 07:012, 2002.

[85] Michelangelo L. Mangano, Mauro Moretti, Fulvio Piccinini, et al. Matching matrix elements and shower evolution for top-quark production in hadronic collisions. *The Journal of High Energy Physics*, 01:013, 2007.

[86] Stephen Mrenna and Peter Richardson. Matching matrix elements and parton showers with HERWIG and PYTHIA. *The Journal of High Energy Physics*, 05:040, 2004.

[87] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 Physics and Manual. *The Journal of High Energy Physics*, 0605:026, 2006.

[88] Johan Alwall et al. MadGraph/MadEvent v4: The New Web Generation. *The Journal of High Energy Physics*, 09:028, 2007.

[89] Borut Paul Kersevan and Elzbieta Richter-Was. The Monte Carlo event generator AcerMC versions 2.0 to 3.8 with interfaces to PYTHIA 6.4, HERWIG 6.5 and ARIADNE 4.1. *Computer Physics Communications*, 184:919–985, 2013.

[90] Michelangelo L. Mangano et al. ALPGEN, a generator for hard multiparton processes in hadronic collisions. *The Journal of High Energy Physics*, 0307:001, 2003.

[91] Paolo Nason. A New method for combining NLO QCD with shower Monte Carlo algorithms. *The Journal of High Energy Physics*, 0411:040, 2004.

[92] Stefano Frixione, Fabian Stoeckli, Paolo Torrielli, et al. The MCaNLO 4.0 Event Generator. 2010.

[93] T. Gleisberg, Stefan. Hoche, F. Krauss, et al. Event generation with SHERPA 1.1. *The Journal of High Energy Physics*, 0902:007, 2009.

[94] Stefan Hoche, Frank Krauss, Steffen Schumann, et al. QCD matrix elements and truncated showers. *The Journal of High Energy Physics*, 0905:053, 2009.

[95] Tanju Gleisberg and Stefan Hoche. Comix, a new matrix element generator. *The Journal of High Energy Physics*, 0812:039, 2008.

[96] Steffen Schumann and Frank Krauss. A Parton shower algorithm based on Catani-Seymour dipole factorisation. *The Journal of High Energy Physics*, 0803:038, 2008.

[97] G. Corcella, I.G. Knowles, G. Marchesini, et al. HERWIG 6: An Event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *The Journal of High Energy Physics*, 0101:010, 2001.

[98] GEANT4 Collaboration. GEANT4: A Simulation toolkit. *Nuclear Instruments and Methods in Physics Research*, A506:250–303, 2003.

[99] E. Barberio, J. Boudreau, B. Butler, et al. Fast shower simulation in the ATLAS calorimeter. *Journal of Physics: Conference Series*, 119(3):032008, July 2008.

[100] ATLAS Collaboration. Luminosity Determination in pp Collisions at sqrt(s)=7 TeV Using the ATLAS Detector at the LHC. *The European Physical Journal*, C71:1630, 2011.

[101] ATLAS Collaboration. Luminosity determination in pp collisions at sqrt(s) = 7 tev using the atlas detector in 2011.

[102] ATLAS Collaboration. Search for direct third-generation squark pair production in final states with missing transverse momentum and two b-jets in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. *The Journal of High Energy Physics*, 1310:189, 2013.

[103] D Damazio, T Kono, F Monticelli, et al. Performance of the ATLAS Electron and Photon Triggers in p-p Collisions at $\sqrt{s} = 8$ TeV in 2012. Technical Report ATL-COM-DAQ-2013-121, CERN, Geneva, Oct 2013.

[104] M Stoebe, S Chekanov, J Cantero, et al. Measurement of the inclusive isolated prompt photon cross section in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector using 20.2 fb^{-1}. Technical Report ATL-COM-PHYS-2014-744, CERN, Geneva, Jun 2014. Supporting note for the 2012 inclusive cross section measurement.

[105] J Almond, M Aoki, M Biglietti, et al. Performance of the ATLAS muon trigger in pp collisions at sqrt(s)=8 TeV. Technical Report ATL-COM-DAQ-2013-147, CERN, Geneva, Nov 2013.

[106] ATLAS Collaboration. Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data. 2014.

[107] ATLAS Collaboration. Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV. *The European Physical Journal*, C73(3):2304, 2013.

[108] Guennadi Pospelov and the Atlas Hadronic Calibration Group. The overview of the atlas local hadronic calibration. *Journal of Physics: Conference Series*, 160(1):012079, 2009.

[109] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti-k_t jet clustering algorithm. *The Journal of High Energy Physics*, 04:063, 2008.

[110] ATLAS collaboration. Tagging and suppression of pileup jets. 2014.

[111] ATLAS Collaboration. Commissioning of the ATLAS high-performance b-tagging algorithms in the 7 TeV collision data. Technical Report ATLAS-CONF-2011-102, CERN, Geneva, Jul 2011.

[112] ATLAS Collaboration. Calibration of the performance of b-tagging for c and light-flavour jets in the 2012 ATLAS data. Technical Report ATLAS-CONF-2014-046, CERN, Geneva, Jul 2014.

[113] ATLAS Collaboration. Calibration of b-tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment. 2014.

[114] ATLAS Collaboration. Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data. *The European Physical Journal*, C74(7):2941, 2014.

[115] ATLAS Collaboration. Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data. 2014.

[116] ATLAS Collaboration. Electron efficiency measurements with the ATLAS detector using the 2012 LHC proton-proton collision data. 2014.

[117] ATLAS Collaboration. Expected photon performance in the atlas experiment.

[118] ATLAS Collaboration. Measurement of the inclusive isolated prompt photon cross section in pp collisions at sqrt(s) = 7 TeV with the ATLAS detector. *Physical Review D*, 83:052005, 2011.

[119] ATLAS Collaboration. Measurements of the photon identification efficiency with the ATLAS detector using 4.9 fb^{-1} of pp collision data collected in 2011. 2012.

[120] ATLAS collaboration. Performance of Missing Transverse Momentum Reconstruction in ATLAS studied in Proton-Proton Collisions recorded in 2012 at 8 TeV. 2013.

[121] William Buttinger. Using Event Weights to account for differences in Instantaneous Luminosity and Trigger Prescale in Monte Carlo and Data. Technical Report ATL-COM-SOFT-2015-119, CERN, Geneva, May 2015.

[122] M. Baak, G. J. Besjes, D. Côte, et al. HistFitter software framework for statistical data analysis. *The European Physical Journal*, C75(4):153, 2015.

[123] ATLAS Collaboration. Search for Scalar Top Quarks in the final states with two charm jets and missing transverse momentum at sqrt(s)=8TeV. *Proceedings, 2013 European Physical Society Conference on High Energy Physics*, EPS-HEP2013:111, 2013.

[124] Gary W. Oehlert. A note on the delta method. *The American Statistician*, 46(1):pp. 27–29, 1992.

[125] ATLAS Collaboration. Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. 2014.

[126] Pile-up subtraction and suppression for jets in ATLAS. Technical Report ATLAS-CONF-2013-083, CERN, Geneva, Aug 2013.

[127] ATLAS Collaboration. Measurement of the W^+W^- production cross section in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. 2014.

[128] ATLAS Collaboration. A Measurement of WZ Production in Proton-Proton Collisions at sqrt(s)=8TeV with the ATLAS Detector. 2013.

[129] ATLAS Collaboration. Measurement of the total ZZ production cross section in proton-proton collisions at sqrt(s) = 8 TeV in 20 fb-1 with the ATLAS detector. 2013.

[130] ATLAS Collaboration. Comprehensive measurements of t-channel single top-quark production cross sections at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Physical Review D*, 90(11):112006, 2014.

[131] J Wenninger. Energy Calibration of the LHC Beams at 4 TeV. Technical Report CERN-ATS-2013-040, CERN, Geneva, May 2013.

[132] Michiel Botje et al. The PDF4LHC working group interim recommendations, 2011.

[133] Richard D. Ball, Valerio Bertone, Stefano Carrazza, et al. Parton distributions with LHC data. *Nuclear Physics*, B867:244–289, 2013.

[134] Hung-Liang Lai, Marco Guzzi, Joey Huston, et al. New parton distributions for collider physics. *Physical Review D*, 82:074024, 2010.

[135] H1 Collaboration. Observation of events at very high Q^2 in *ep* collisions at HERA. *Zeitschrift für Physik*, C74:191–206, 1997.

[136] M.R. Whalley, D. Bourilkov, and R.C. Group. The Les Houches accord PDFs (LHAPDF) and LHAGLUE. 2005.

[137] Simon Owen. *Multijet Background Estimation for Supersymmetry Searches Using the ATLAS Detector at the Large Hadron Collider*. PhD thesis, University of Sheffield, 2012.

[138] A L Read. Presentation of search results: the cl s technique. *Journal of Physics G*, 28(10):2693, 2002.

[139] CMS Collaboration. Search for dark matter direct production using razor variables in events with two or more jets in pp collisions at 8 TeV. 2015.

[140] CMS Collaboration. Search for the production of dark matter in association with top-quark pairs in the single-lepton final state in proton-proton collisions at sqrt(s) = 8 TeV. *JHEP*, 06:121, 2015.

[141] CMS Collaboration. Search for dark matter and unparticles produced in association with a Z boson in pp collisions at sqrt(s) = 8 TeV. Technical Report CMS-PAS-EXO-12-054, CERN, Geneva, 2015.

[142] Tansu Daylan, Douglas P. Finkbeiner, Dan Hooper, et al. The Characterization of the Gamma-Ray Signal from the Central Milky Way: A Compelling Case for Annihilating Dark Matter. 2014.

[143] Prateek Agrawal, Brian Batell, Dan Hooper, et al. Flavored Dark Matter and the Galactic Center Gamma-Ray Excess. *Physical Review D*, 90(6):063512, 2014.

[144] Lisa Goodenough and Dan Hooper. Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope. 2009.

[145] Fermi-LAT Collaboration. Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi-LAT Data. 2015.

[146] Dan Hooper and Tim Linden. On The Gamma-Ray Emission From Reticulum II and Other Dwarf Galaxies. 2015.

[147] Tanya Sandoval, M A Parker, and S Ask. Estimation of the $Z \to \nu\nu$ background to New Physics searches in ATLAS. May 2012. Presented 22 Jun 2012.

[148] ATLAS Collaboration. Search for the Standard Model Higgs boson in produced in association with a vector boson and decaying to bottom quarks with the ATLAS detector. (ATLAS-CONF-2012-161), Nov 2012.

[149] Maria Beltran, Dan Hooper, Edward W. Kolb, et al. Deducing the nature of dark matter from direct and indirect detection experiments in the absence of collider signatures of new physics. *Physical Review D*, 80:043509, 2009.

[150] Matthew R. Buckley. Using effective operators to understand CoGeNT and CDMS-Si signals. *Physical Review D*, 88(5):055028, 2013.

[151] P. Bhattacharjee, S. Chaudhury, S. Kundu, et al. Deriving the velocity distribution of Galactic dark matter particles from the rotation curve data. *Physical Review D*, 87(8):083525, April 2013.

[152] E.C. Stone, C.M.S. Cohen, W.R. Cook, et al. The cosmic-ray isotope spectrometer for the advanced composition explorer. pages 285–356, 1998.

[153] M Werner, R Kissmann, AW Strong, et al. Spiral arms as cosmic ray source distributions. 64:18–33, 2015.

[154] Ralf Kissmann, Michael Werner, Olaf Reimer, et al. Propagation in 3d spiral-arm cosmic-ray source distribution models and secondary particle production using picard. 70:39–53, 2015.

[155] Ralf Kissmann, Olaf Reimer, and Andrew W. Strong. Galactic cosmic ray propagation models using Picard. 2015.

[156] Yu-Heng Chen, Kingman Cheung, and Po-Yan Tseng. Dark Matter with Multi-Annihilation Channels and AMS-02 Positron Excess and Antiproton. 2015.

[157] Andi Hektor, Martti Raidal, Alessandro Strumia, et al. The cosmic-ray positron excess from a local dark matter over-density. *Physics Letters B*, 728:58 – 62, 2014.

[158] Keith R. Dienes, Jason Kumar, and Brooks Thomas. Dynamical dark matter and the positron excess in light of ams results. *Physical Review D*, 88:103509, Nov 2013.

[159] Dan Hooper and Lisa Goodenough. Dark matter annihilation in the galactic center as seen by the fermi gamma ray space telescope. *Physics Letters B*, 697(5):412 – 428, 2011.

[160] Seyda Ipek, David McKeen, and Ann E. Nelson. Renormalizable model for the galactic center gamma-ray excess from dark matter annihilation. *Physical Review D*, 90:055021, Sep 2014.

[161] P. Ko, Wan-Il Park, and Yong Tang. Higgs portal vector dark matter for gev scale γ-ray excess from galactic center. *The Journal of Cosmology and Astroparticle Physics*, 2014(09):013, 2014.

[162] Céline Bœhm, Matthew J. Dolan, and Christopher McCabe. Interpretation of the galactic center excess of gamma rays with heavier dark matter particles. *Physical Review D*, 90:023531, Jul 2014.

[163] Eder Izaguirre, Gordan Krnjaic, and Brian Shuve. The Galactic Center Excess from the Bottom Up. *Physical Review D*, 90(5):055002, 2014.

[164] Kevork N. Abazajian. The consistency of fermi-lat observations of the galactic center with a millisecond pulsar population in the central stellar cluster. *The Journal of Cosmology and Astroparticle Physics*, 2011(03):010, 2011.

[165] Qiang Yuan and Bing Zhang. Millisecond pulsar interpretation of the galactic center gamma-ray excess. *The Journal of High Energy Astrophysics*, 3–4:1 – 8, 2014.

[166] Richard Bartels, Suraj Krishnamurthy, and Christoph Weniger. Strong support for the millisecond pulsar origin of the Galactic center GeV excess. 2015.

[167] Joshua D. Simon and Marla Geha. The Kinematics of the Ultra-Faint Milky Way Satellites: Solving the Missing Satellite Problem. *The Astrophysical Journal*, 670:313–331, 2007.

[168] Fermi-LAT Collaboration. Dark matter constraints from observations of 25 Milky Way satellite galaxies with the Fermi Large Area Telescope. *Physical Review D*, 89:042001, 2014.

[169] J. Billard, L. Strigari, and E. Figueroa-Feliciano. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Physical Review D*, 89(2):023524, 2014.

[170] PICO Collaboration. Dark Matter Search Results from the PICO-2L C_3F_8 Bubble Chamber. *Physical Review Letters*, 114(23):231302, 2015.

[171] Joel Giedt, Anthony W. Thomas, and Ross D. Young. Dark matter, the CMSSM and lattice QCD. *Physical Review Letters*, 103:201802, 2009.

[172] Gerard Jungman, Marc Kamionkowski, and Kim Griest. Supersymmetric dark matter. 267:195–373, 1996.

[173] John R. Ellis, Keith A. Olive, and Christopher Savage. Hadronic Uncertainties in the Elastic Scattering of Supersymmetric Dark Matter. *Physical Review D*, 77:065026, 2008.

[174] DAMA/LIBRA Collaboration. New results from DAMA/LIBRA. *The European Physical Journal*, C67:39–49, 2010.

[175] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. A Brief Introduction to PYTHIA 8.1. *Computer Physics Communications*, 178:852–867, 2008.

Curriculum Vitae