Title
Searches for Physics Beyond the Standard Model: From Colliders to Dark Matter Plasmas

ISBN
9798184183855

Author
Giffin, Pierce Walter

Publication Date
2026-05-15

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

SEARCHES FOR PHYSICS BEYOND THE STANDARD MODEL: FROM COLLIDERS TO DARK MATTER PLASMAS

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

PHYSICS

by

Pierce Giffin

June 2026

The Dissertation of Pierce Giffin
is approved:

____

Professor Stefania Gori, Chair

____

Professor Stefano Profumo

____

Professor Wolfgang Altmannshofer

____

Peter Biehl
Vice Provost and Dean of Graduate Studies

Contents

Figures and Tables **vi**

Abstract **viii**

Acknowledgments **ix**

1 Introduction **1**
1.1 Searches for Heavy Neutral Leptons 2
1.2 Motivation for Dark Matter Plasmas 4

2 Heavy Neutral Leptons at Beam Dump Experiments of Future Lepton Colliders **6**
2.1 Beam dump Experiments at Future e^+e^- Colliders 8
2.1.1 ILC beam dump Configuration 8
2.1.2 Beam dump Experiments at other e^+e^- Colliders: C^3, CLIC, FCC-ee, CEPC 10
2.1.3 Meson Production at high-energy electron beam dumps 12
2.2 Heavy Neutral Leptons 14
2.2.1 HNL Production at an Electron Beam Dump Experiment 15
2.2.2 HNL Decays 18
2.2.3 Detector Efficiency and Acceptance 20
2.3 Sensitivity for Heavy Neutral Leptons 23
2.3.1 Sensitivity of Lepton Collider Beam Dump Experiments 23
2.3.2 Modifying the Experimental Setup 27
2.3.3 Possible Backgrounds 29
2.4 Summary 31

3 Exploring Invisible New Physics with Exotic Pion Decays **32**
3.1 Exotic Pion Decays at PIENU and PIONEER 33
3.2 Exotic Two-Body Decays of Pions to Sterile Neutrinos 38
3.2.1 Sterile neutrinos 39
3.2.2 Experimental signatures at pion decay experiments 40
3.2.3 Reproducing the existing constraints on $|U_{eN}|$ and $|U_{\mu N}|$ 43

3.2.4 Expected sensitivities at PIONEER 45
3.3 Exotic Three-Body Decays of Pions 54
3.3.1 Simplified dark sector models and their pion decay spectra . . . 54
3.3.2 Experimental signatures at pion decay experiments 59
3.3.3 Recast of existing pion decay constraints 62
3.3.4 Expected sensitivities at PIONEER 64
3.3.5 Comparison to other dark sector searches 66
3.3.6 Sensitivity to dark sector model parameters 73
3.4 Conclusions 76

4 Dark Plasmas in the Nonlinear Regime: Constraints from Particle-in-Cell Simulations 79
4.1 Dark U(1) Plasma 79
4.2 The Bullet Cluster 83
4.3 Simulations 84
4.4 Results 86
4.5 Conclusions 90

5 Structure Formation with Dark Magnetohydrodynamics 92
5.1 Gravitational Collapse and classical Jeans length 93
5.2 Evolution of Dark Magnetized Plasma 94
5.2.1 Parallel propagation 97
5.2.2 Perpendicular propagation 99
5.3 Dark Magnetic Fields: Generation, Cosmological Constraints, and Caveats 100
5.3.1 Dark magnetic field evolution 100
5.3.2 Constraints on standard model primordial Magnetic Fields . . . 103
5.4 Effects on Matter Power Spectrum 106
5.5 Discussion and Conclusions 112

6 Conclusion 116

Appendix 119

A Simulation Details 119
A.1 Meson Productions 119
A.2 Three-Body Decays of HNLs 122

B Details on the PIENU Sterile Neutrino Search Recast 123

C Details on the PIENU Three-Body Pion Decay Recast 128

D Plasma review: Linear and non-linear regimes 133
D.1 Linear regime 134
D.1.1 Two-stream instability 135
D.1.2 Weibel instability 136
D.2 Nonlinear regime 137

E Simulations 138
E.1 Setup 138
E.2 Growth rates 139
E.3 Saturation energies 141
E.4 Additional diagnostic plots 142
E.5 Limitations of 2D3V simulations 144

F Derivation of Generalized Ohm's Law 148

G Derivation of Dispersion Relation 150

References 157

Figures and Tables

List of Figures

2.1 Layout of ILC beam dump experiment 9
2.2 Number of HNLs produced at beam dump experiments 17
2.3 Branching ratios of HNLs at beam dump experiments 19
2.4 Geometric acceptance and efficiency of beam dump experiments 21
2.5 Electron-mixed HNL reach curves for proposed beam dump experiments 24
2.6 Muon-mixed HNL reach curves for proposed beam dump experiments . 24
2.7 τ-mixed HNL reach curves for proposed beam dump experiments . . . 25
2.8 Electron-mixed HNL reach curve for various muon shield lengths . . . 28

3.1 Constraints on $|U_{eN}|^2$ for electron mixed HNLs 44
3.2 Constraints on $|U_{\mu N}|^2$ for muon mixed HNls 46
3.3 Positron enegy spectrum of PIONEER 48
3.4 Sensitivity curves for $|U_{eN}|^2$ at PIONEER 51
3.5 Electron energy spectrum of $\pi^+ \to e^+\nu_e X$ at PIONEER 60
3.6 Muon energy spectrum of $\pi^+ \to \mu^+\nu_\mu X$ at PIONEER 61
3.7 Branching ratio sensitivity to $\pi^+ \to e^+\nu_e X$ at PIENU and PIONEER . 63
3.8 Branching ratio sensitivity to $\pi^+ \to \mu^+\nu_\mu X$ at PIENU and PIONEER . 64
3.9 Constraints on dark sector particles for $\pi^+ \to e^+\nu_e X$ 74
3.10 Constraints on dark sector particles for $\pi^+ \to \mu^+\nu_\mu X$ 75

4.1 Constraints of dark charge-to-mass ratio of dark $U(1)$ electromagnetism 81
4.2 Fraction of simulated particles scattered during PIC simulation 87
4.3 Velocity distributions of Bullet Cluster subclusters in PIC simulation . . 91

5.1 Constraints on primordial dark magnetic fields 105
5.2 Effective phase speed of magnetized dark plasma as a function of dark magnetic field strength and dark charge-to-mass ratio 107
5.3 Linear matter power spectrum with dark plasma effects 111
5.4 Transfer function with dark plasma effects 111
5.5 Derived constraints on dark charge-to-mass ration and dark magnetic field strength parameter space 114

A.1 Energy and p_T of mesons produced at ILC beam dump experiment . . . 122

B.1 Expected number of detected $\pi^+ \to e^+ N$ decays at PIENU 124
B.2 Expected number of detected $\pi^+ \to \mu^+ N$ decays at PIENU 127

C.1 Expected number of detected $\pi^+ \to e^+ \nu_e X$ decays at PIENU 129
C.2 Upper limits on the $\pi^+ \to e^+ \nu_e X$ and $\pi^+ \to \mu^+ \nu_\mu X$ branching ratios at PIENU 130
C.3 Expected number of detected $\pi^+ \to \mu^+ \nu_\mu X$ decays at PIENU 132

E.1 Electric and magnetic field energy densities of PIC simulation of the Bullet Cluster 140
E.2 Energy lost to numerical heating during PIC simulation of the Bullet Cluster 142
E.3 Growth rate of longitudinal electrostatic instability in the Bullet Cluster 143
E.4 Growth rate of transverse electromagnetic instability in the Bullet Cluster 144
E.5 Longitudinal electric field from PIC simulations 145
E.6 Power spectrum of longitudinal electric field from PIC simulations . . . 146
E.7 Transverse magnetic field from PIC simulations 146
E.8 Power spectrum of transverse magnetic field from PIC simulations . . . 147

List of Tables

2.1 Beam parameters for e^+e^- colliders 10
2.2 Mesons and τ leptons produced at beam dump experiments 14

3.1 Summary of key experimental parameters relevant for exotic pion decay searches at PIENU and PIONEER 38

4.1 Initial conditions and resulting constraints on dark charge-to-mass ratio from different simulations of the Bullet Cluster 90

A.1 Number of mesons produced per generated Monte Carlo soft QCD event at different beam energies 120

Abstract

Searches for Physics Beyond the Standard Model: From Colliders to Dark Matter Plasmas

by

Pierce Giffin

In this dissertation, we provide a collection of works dedicated to various searches for physics beyond the Standard Model of particle physics. For Earth-based searches, we perform dedicated analyses of future beam dump and pion decay experiments with special emphasis on probing heavy neutral leptons. We find that these experiments show great promise in their discovery capabilities and provide exciting new opportunities to probe a variety of models. In the astrophysical context, we extend the typical single particle interaction paradigm of dark matter to incorporate collective modes generated from plasma instabilities. With the use of dedicated particle-in-cell simulations of the Bullet Cluster and semi-analytic models of the Jeans instability, we rule out large regions of parameter space for dark matter with a secluded $U(1)_D$ gauge group previously inaccessible.

Acknowledgements

First, I would like to thank the outstanding faculty on my defense committee, Stefania Gori, Stefano Profumo, and Wolfgang Altmannshofer. You have all been amazing advisors, teachers, and collaborators. The guidance you have given me as a graduate students at UCSC has truly been invaluable.

Thank you to all of the graduate students at UCSC who helped me grow and learn as a physicist. I am grateful to Sri Aditya Gadam, Pankaj Munbodh, and M. Grant Roberts for valuable discussion and collaborations. Additionally, I would like to thank all of the fellow banana slugs within the department who helped form a wonderful community around physics that I will be very sad to leave behind.

To my amazing parents, Paul and Maria Giffin, thank you for your support and encouragement throughout every stage of my academic career. Your guidance and passion for education inspired me to pursue my dreams from a young age and to be curious about the world around me.

Lastly, thank you to my wife, Hannah Giffin. Your unwavering support and compassion has meant the world to me during this journey. From Kansas to California, it's always been amazing having you at my side. Rock Chalk.

Chapter 1

Introduction

In 1903, Albert Michelson claimed in one of his lectures that "our future discoveries must be looked for in the sixth place of decimals" [1, pp. 23–24]. At the time, it was believed that all of physics was well understood. New developments in the field would arise only from the repetition of experiments at ever higher levels of precision. Two years later, Albert Einstein published his theory of special relativity in 1905 [2], followed by his theory of general relativity in 1915 [3]. Furthermore, the development of quantum mechanics by Werner Heisenberg [4] and Erwin Schrödinger [5] in the following decade revealed that many aspects of the universe remained poorly understood.

100 years later, the search for new physics continues. Today, many researchers look for signatures of new physics at the high-energy frontier of particle physics. Though the Standard Model (SM) of particle physics has served as an incredibly robust theory of nature, it still leaves many open questions unexplained. On Earth, many experiments aim to produce beyond SM (BSM) particles at high-energy particle colliders. By creating extremely large samples of particle collisions and decays, statistical methods can be used to search for increasingly rare BSM interactions.

In astrophysical settings, BSM particles may exhibit hints of their properties in extreme environments. Throughout the cosmos, many objects exist on larger physical scales and at much higher temperatures and densities than we can currently produce in laboratory settings. Often motivated by the search for particle dark matter (DM), observations of these environments allow us to probe for physics on scales otherwise inaccessible.

In Sec. 1.1, we discuss the BSM model of heavy neutral leptons (HNLs) in the context of beam dump and pion decay experiments. In the astrophysical context, many BSM searches aim to detect signatures of single-particle interactions of dark matter particles. In Sec. 1.2, we motivate the extension of these searches to include collective modes from long-range interactions, borrowing techniques from plasma physics to probe the dark sector.

§ 1.1 Searches for Heavy Neutral Leptons

HNLs are well-motivated extensions of the SM [6–12]. The observation of neutrino oscillations is evidence for non-zero neutrino masses. A straightforward way of generating them is to add to the SM right-handed gauge-singlet fermions, N_i, called HNLs. In addition, HNLs naturally appear in many other extensions of the SM, since they can be connected to the matter-antimatter asymmetry problem and to the nature of DM (see, e.g., [13–16].) A rich experimental program has been conducted in recent years to probe HNL parameter space both at high intensity and at high energy experiments. This includes neutrino experiments (e.g., CHARM [17], PS191 [18, 19], NOMAD [20], T2K [21, 22], MicroBooNE [23, 24], ArgoNeuT [25], the Fermilab SBN detectors [26], and the DUNE near detector complex [27–29]), proposed auxiliary LHC detectors that target long-lived particles (e.g. FASER [30], MATHUSLA [31–34], CODEX-b [35, 36]), and proton fixed-target beam dump experiments (e.g. DarkQuest [37], SHiP [38–42], and NA62 [43, 44]). In addition, there are searches at high energy proton-proton colliders such as the LHC, including its ATLAS [45] and LHCb detector [46], proposals for future searches at the High-Luminosity LHC (HL-LHC) [47], and at existing and future e^+e^- colliders, including Belle [48], Belle II [49], and the Future Circular Collider (FCC) [50]. HNLs can be produced in several different processes in these experiments. In particular, HNLs are produced from SM meson decays at neutrino and proton beam dump experiments, from Drell-Yan processes mediated by the weak interactions, and from Z or W decays at the LHC and e^+e^- colliders. In Chapter 2 we perform a dedicated study of HNL production at future beam dump experiments, including the ILC, C^3, and CLIC. We find that these experiments are especially capable of detecting HNLs for masses smaller than about 10 GeV, providing very complementary searches to proton beam dump experiments, neutrino experiments, and LHC auxiliary detectors.

In addition to beam dump experiments, precision measurements of rare meson decays offer a powerful avenue to test the SM and search for HNLs as well as other light BSM particles. Among these, charged pion decays play a special role [51]. They are theoretically exceptionally clean and experimentally accessible. Moreover, the pion's relatively small decay width makes them sensitive to a broad class of new light particles that may couple feebly to the Standard Model and that can be produced from pion decays. The proposed PIONEER experiment at PSI [52, 53] will collect an unprecedented sample of pions decaying at rest and presents a unique opportunity to test light new particles. In

Chapter 3, we show that in addition to providing an increased sensitivity to HNLs, the PIONEER experiment is capable of probing couplings roughly an order of magnitude smaller than preceding experiments for invisible scalars, axion-like particles (ALPs), and dark vector particles.

§ 1.2 Motivation for Dark Matter Plasmas

Vast theoretical and experimental effort has been devoted to discovering the nature of dark matter, often under the assumption that dark matter consists of a new fundamental particle with a low interaction cross-section with the visible sector. As a result, the phenomenology of dark matter has relied almost exclusively on single-particle interactions. This is in direct contrast to the visible sector of the Universe, in which the vast majority of visible matter exists in the form of plasma, the behavior of which is governed not by single-particle interactions but by collective effects mediated by long-range interactions. The most dramatic of these effects arise from *plasma instabilities*, in which small initial perturbations grow exponentially and produce nonlinear structures even on scales far smaller than the particle scattering mean free path. If, in analogy to the visible sector, self-interactions in the dark sector are dominated by collective effects, then the associated instabilities can lead to observable effects in astrophysical systems.

Despite this, little work has been done on characterizing dark sectors governed by collective interactions. Under the assumption that collective effects would make dark matter effectively collisional, Heinkinheimo et al. [54–56] performed hydrodynamic simulations of dissociative cluster mergers with an artificial viscosity that was tuned to best reproduce observed lensing maps. Growth rates for common instabilities have been computed in the linear regime for a variety of astrophysical settings [57, 58]. Simulations of interpenetrating electron-positron pair plasmas have also discussed possible implications for the dark sector [59], and analytic arguments for the existence of limits on particular models have also appeared in the literature [60]. In Chapter 4, we explore such instabilities in a dark sector with a secluded $U(1)_D$ gauge symmetry. By performing the world's first-ever dedicated particle-in-cell simulations of DM plasma instabilities, we show that these effects can be dramatic. In the context of the Bullet Cluster system, our results show that dark charge-to-mass ratios over 10 orders of magnitude smaller than previous constraints can be ruled out as unphysical.

However, one critical ingredient remains largely unexplored: the role of gravity in the formation and evolution of dark-sector plasma instabilities. Understanding its interplay with dark plasma dynamics is essential for determining the impact on both primordial density perturbations and late-time structure formation. In Chapter 5, we incorporate gravitational effects into the same secluded $U(1)_D$ DM model. We derive a system of dark magnetohydrodynamic equations that model the evolution of the dark plasma in the presence of a background magnetic field. If large-scale dark magnetic fields are present during early structure formation, we show that this can lead to an anisotropic alteration of the classic Jeans instability. We find that current observations of the linear matter power spectrum cannot yet constrain viable dark magnetic fields, as cosmic microwave background (CMB) tensor modes mostly provide more stringent constraints. Nevertheless, forthcoming high-resolution probes of the matter power spectrum (CMB-HD lensing [61], HERA [62], and EDGES [63]) will be able to test these predictions and are sensitive to dark charge-to-mass ratios in the range $10^{-20}\,\mathrm{GeV}^{-1} \lesssim q_\chi/m_\chi \lesssim 10^{-14}\,\mathrm{GeV}^{-1}$.

Chapter 2

Heavy Neutral Leptons at Beam Dump Experiments of Future Lepton Colliders

This chapter is based on [64]. In this work, we propose to use the beam dump of future electron-positron colliders, such as the International Linear Collider (ILC), the Cool Copper Collider (C^3) [65], and FCC-ee [50, 66], to search for HNLs. As we will show, adding instruments after the electron beam dump of these experiments could lead to the exploration of new HNL parameter space. While the main goal of these next-generation e^+e^- machines is to use the main collision to perform precision studies of SM particles (such as the Higgs boson and the top quark) and to produce light New Physics (NP) particles, they could also perform searches for the production of new dark sector (i.e. gauge singlet) states produced in the electron beam dump. More specifically, once the electrons and positrons pass the collision point, they will be discarded in the beam dump. Thanks to the large number of electrons on target, this has the potential to produce a large amount of light dark sector particles from the interactions of the electrons or positrons with the beam dump. These dark sector particles can subsequently decay into visible SM particles, that could be detected by a suitable apparatus installed behind the beam dump. It has been shown that an ILC beam dump experiment can probe currently unexplored parameter space of many well-motivated BSM theories (dark photons [67, 68], axion-like particles (ALPs) and light scalar particles [68, 69], and leptophilic Gauge Bosons [70]).

In this paper, we discuss the possibility of searching for HNLs at high-energy electron beam dump experiments where HNLs can be produced in charged-current electron-proton scattering and in the decays of mesons produced after the electrons collide with the beam dump. Once produced the HNLs can travel a macroscopic distance before decaying into SM particles. These particles can be detected if a suitable detector is installed behind the beam dump. Depending on the specific setup of the experiment, this type of signatures can have a very low background. Hence, these searches will be complementary to the HNL searches that can be performed using the main e^+e^- collider [50, 71–74]. Our study is also relevant for future electron-ion colliders [75] as they share similar production channels, especially the charged-current scattering for electron-mixed HNL discussed in our work.

Future electron collider beam dump experiments will provide exciting opportunities due to their unprecedented high energy and high- intensity beams. Compared to existing and near-future proton beam dump experiments, electron collider beam dumps can reach comparable or higher center-of-mass energies and can thus produce heavier dark sector particles. Far forward experiments at hadron colliders have higher energies [76, 77], but they have lower intensity compared to electron collider beam dumps due to their geometric acceptance and lower luminosities. The main beam dumps of hadron collider could also be interesting for probing dark sector particles, but they suffer from significant radiation that can easily overwhelm the detectors if the detectors are directly placed downstream of the dump [78].

This paper is organized as follows. In section 2.1, we give an overview of the beam dump experiments that could be performed at future e^+e^- colliders, including, ILC, C^3, and FCC-ee. In section 2.2, we discuss the phenomenology of HNLs. This is followed by the prospects for HNL searches at electron collider beam dump experiments in section 2.3. We conclude in section 2.4. In the Appendix, we report the details of our simulations for meson production at the beam dump. The meson production rate will be the main source of uncertainty in the calculation of the reach of the HNL parameter space.

§ 2.1 Beam dump Experiments at Future e^+e^- Colliders

We begin by describing the experimental setup under consideration, following the proposal for a beam dump experiment at the ILC in [67, 69]. We highlight that similar setups could be implemented at other future e^+e^- linear colliders like C^3 and CLIC, as well as at future e^+e^- circular colliders like FCC-ee and CEPC.

2.1.1 ILC beam dump Configuration

The International Linear Collider (ILC) experiment is a proposed future collider that uses high-energy collisions of electron (e^-) and positron (e^+) beams. The ILC is planned to initially run with a center-of-mass energy of 250 GeV, and will then be upgraded to run at 350 GeV, 500 GeV, and 1 TeV [79]. In addition to its central collider, the ILC can host a number of additional detectors, including detectors for beam dump experiments.

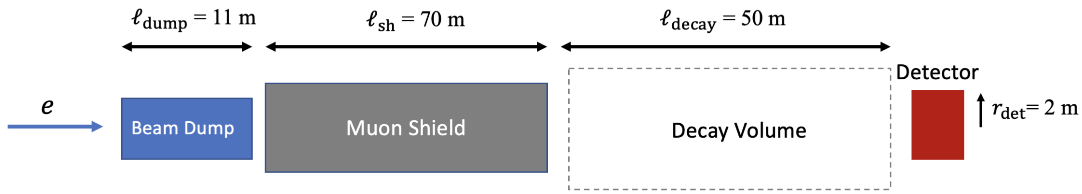

Figure 2.1: Schematic layout of the beam dump experiment for the ILC. A similar experiment can be designed for other e^+e^- colliders as C^3, CLIC, FCC-ee, and CEPC.

The main dumps accept the full power of the ILC beam. We consider the configuration described in [69] (see also section 11.2 of the Snowmass report [79]). The length of the main dump which uses water as the absorber is $\ell_{dump} = 11$ m. This is followed by a muon shield made of lead and concrete of length $\ell_{sh} = 70$ m along the beam direction, and by a decay volume of length $\ell_{decay} = 50$ m. The detector is assumed to be cylindrical with a radius $r_{det} = 2$ m and its axis aligned with the beam axis. We will make the simple assumption that the detector has a 100% detection efficiency, as long as SM particles are in geometric acceptance. The sensitivity is subject to $\mathcal{O}(1)$ adjustments depending on the exact determination of the detection efficiency as the beam-dump experiment is being constructed. A schematic layout of the experiment is shown in figure 2.1.

In the following sections, we will focus and compare the physics opportunities that arise using the 125 GeV ($\sqrt{s} = 250$ GeV) and 500 GeV ($\sqrt{s} = 1$ TeV) ILC electron beams, which we will denote by "ILC-250" and "ILC-1000", respectively. For both stages, we will consider the initial ILC-250 configuration which can deliver 4×10^{21} electrons-on-target (EOT) per year (see the ILC beam parameters discussed in table 4.1 of [79]). After the first ILC stage, the number of EOT/year can be increased up to 10^{22}

Collider-$\sqrt{s}$ [GeV]	ILC-250/1000	C^3-250	C^3-3000	CLIC-3000
Bunches/Train	1312	133	75	312
Train Rep. Rate [Hz]	5	120	120	50
Bunch Charge [nC]	3.2	1	1	0.6
Effective Luminosity [cm^{-2} s^{-1}]	1.6×10^{39}	1.2×10^{39}	6.9×10^{38}	6.9×10^{38}
EOT/Year	4.1×10^{21}	3.1×10^{21}	1.8×10^{21}	1.8×10^{21}

Table 2.1: Beam parameters and total number of electrons-on-target (EOT) per year for different e^+e^- colliders assuming one year of operation. The numbers for the beam parameters are taken from table 4.1 of [79] for ILC, table 3 of [65] for C^3, and table 4.1 of [80] for CLIC. The "effective luminosity", $\mathcal{L}_{\text{eff}}$, refers to the instantaneous luminosity of electron-nucleus collisions in the first radiation length of the dump (and not to the instantaneous luminosity of the corresponding primary e^+e^- collision).

EOT/year in the most optimistic scenario [79]. Thanks to these large luminosities, a sizable flux of mesons will be created. Most of these mesons are stopped, absorbed, or decay in the dump. Particles surviving the dump (e.g. muons) are shielded by the muon shield. New physics long-lived particles, like heavy neutral leptons, can be produced from meson decays in the dump, travel a macroscopic length, and decay in the decay volume. The visible decay products can be searched for by the detector placed at the end of the decay volume.

2.1.2 Beam dump Experiments at other e^+e^- Colliders: C^3, CLIC, FCC-ee, CEPC

In addition to the ILC, several next-generation e^+e^- colliders have been proposed for precision SM and BSM physics studies, such as the Cool Copper Collider (C^3), the Compact Linear Collider (CLIC), the Future Circular Collider (FCC), and the Circular

Electron Positron Collider (CEPC). In this section, we discuss how these experiments compare to the ILC.

C^3 is a recently proposed e^+e^- linear collider [65, 81, 82]. C^3 will initially run at a 250 GeV center-of-mass energy and has the possibility of an upgrade to 3 TeV. We denote these two runs by "C^3-250" and "C^3-3000", respectively. The beam configurations of these two stages differ only in the number of bunches-per-train. The numbers of EOT/year are 3.1×10^{21} and 1.8×10^{21}, at C^3-250 and C^3-3000, respectively. These numbers are similar to the number of EOT at an ILC beam dump experiment (see table 2.1).

CLIC is a proposed multi-TeV linear collider designed to operate in three stages with center-of-mass energies of 380 GeV, 1.5 TeV, and 3 TeV. Using the beam parameters reported in [80], we find a luminosity of 1.8×10^{21} EOT/year for the 3 TeV stage, denoted "CLIC-3000" in table 2.1. The luminosities for the other stages are similar.

FCC is a proposed next-generation circular collider, which includes a high luminosity e^+e^- storage ring collider (FCC-ee), with center-of-mass energies between 90 and 350 GeV [83]. The planned collider luminosity will be higher than at the ILC. However, given the nature of circular colliders, the beams are reused and dumped much less frequently, yielding a lower number of EOT per year. Approximately 10^{17} electrons will be dumped per year at FCC-ee [67]. Similarly, CEPC is an electron-positron circular collider that is planned to run with a center-of-mass energy of 240 GeV. Since CEPC is a circular collider, it will also lead to a much smaller number of EOT per year. For this reason, in this paper, we will focus on the reach of the HNL parameter space at linear e^+e^- colliders, focusing on the ILC-250, ILC-1000, and C^3/CLIC-3000 experiments.* We summarize the parameters of these experiments in table 2.1.

In the literature, there has not been any discussion of beam dump experiments at C^3 or CLIC. In the following sections, we will assume that the main beam dump of CLIC and C^3 is based on the ILC beam dump design [80, 84]. In section 2.3.2, we will discuss how different beam dump configurations will affect the reach of the HNL parameter space.

We hope our study will provide strong motivation to consider beam dump experiments at these future electron-positron colliders, and will inform the optimal beam dump configuration to maximize the reach on the HNL parameter space, and, more in general, on long-lived dark sector particles produced from meson decays.

2.1.3 Meson Production at high-energy electron beam dumps

Because of the high energy of the electron beam and the large number of electrons on target, mesons will be substantially produced in the beam dump. To determine the production rate of a given meson, we use `Pythia 8.3` [85–87] to simulate 125 GeV, 500 GeV, and 1.5 TeV electron beams striking a fixed target using the `SoftQCD:all` event class. In Appendix A, we collect the details of our simulations. We only consider the mesons produced in the first radiation length of the dump. This will lead to a conservative estimate of the number of mesons produced. Unfortunately, there is no past data on meson production in electron-nucleus collisions at these energies that can be used to normalize the rates. For this reason, the number of mesons produced will be the main source of uncertainty in the estimate of the reach on the HNL parameter space.

*As shown in table 2.1, from the point of view of the energy and the luminosity, CLIC running at $\sqrt{s} = 3$ TeV is identical to C^3-3000, and C^3-250 is identical to ILC-250.

Pions and kaons have large production rates in electron-nucleus collisions. Given their relatively long proper lifetime, they have a decay length that is much larger than the characteristic interaction length in the beam dump material. As a result, a sizable fraction of these mesons can be absorbed by the beam dump before they decay into HNLs. To take this into account, we estimate the number of mesons that decay before the first interaction length in liquid water as [37]

$$N_M \approx N_{\text{EOT}}\, n_M \Gamma_M \langle \gamma_M^{-1} \rangle \lambda_M \,, \tag{2.1}$$

where N_{EOT} is the number of electrons on target, n_M is the number of mesons produced in the first radiation length of the dump per electron on target (see eq. A.1 and table A.1), $\langle \gamma_M^{-1} \rangle$ is the mean inverse Lorentz boost, and λ_M is the meson interaction length in liquid water. We assume that the interaction length of pions and kaons in liquid water are similar and use $\lambda_M \approx 115$ cm [88].

D and B mesons are short-lived, and, therefore, the number of these mesons responsible for HNL production is simply given by $N_M = N_{\text{EOT}}\, n_M$. A large number of τ leptons are also produced, mainly from D_s meson decays, and the total number is given by $N_\tau = N_{\text{EOT}}\, n_{D_s} \times \text{BR}(D_s \to \tau \nu_\tau)$.

In table 2.2, we summarize the number of mesons and τ leptons produced at the several electron beam dump experiments per year of running. In the case of light mesons, the ILC-1000 numbers are slightly larger than the corresponding ones at C^3/CLIC-3000, despite its lower energy. This is a reflection of the lower number of EOT/year of C^3/CLIC-3000 compared to the ILC, which compensates for the larger production of mesons per EOT (see table A.1). The kinematical suppression of $B^\pm$ production at ILC-1000 is more relevant. In fact, C^3/CLIC-3000 will be able to produce a larger

	ILC-250	ILC-1000	C^3-250	C^3/CLIC-3000
$\pi^\pm$	1.5×10^{16}	6.0×10^{16}	1.2×10^{16}	4.8×10^{16}
$K^\pm$	3.9×10^{15}	1.5×10^{16}	3.0×10^{15}	1.2×10^{16}
K_L^0	3.8×10^{14}	1.5×10^{15}	2.9×10^{14}	1.2×10^{15}
$D^\pm$	1.4×10^{14}	7.8×10^{14}	1.0×10^{14}	8.6×10^{14}
$D_s^\pm$	3.8×10^{13}	2.5×10^{14}	2.9×10^{13}	2.6×10^{14}
$B^\pm$	1.0×10^{11}	2.2×10^{12}	7.5×10^{10}	3.4×10^{12}
$\tau^\pm$	2.2×10^{12}	1.4×10^{13}	1.5×10^{12}	1.5×10^{13}

Table 2.2: Total number of mesons and τ leptons produced per year at the several beam dump experiments, summing over both charged states. For pions and kaons, we show the number of mesons that decay before the first interaction length, not the total number of produced mesons.

sample of $B^\pm$ mesons, thanks to its higher center-of-mass energy ($\sqrt{s} \sim 55$ GeV vs. $\sqrt{s} \sim 30$ GeV of the ILC-1000).

§ 2.2 Heavy Neutral Leptons

HNLs can interact with the SM lepton doublet via the neutrino portal operator

$$-\mathcal{L}_{\text{HNL}} \supset Y_{ij} \hat{L}_i H N_j + \text{h.c.}\,, \tag{2.2}$$

where H is the SM Higgs doublet, $\hat{L}_i = (\nu_i, \ell_i)^T$ is the SM lepton doublet of flavor i, and Y_{ij} are the ij elements of the Yukawa matrix. After electroweak symmetry breaking, the HNLs will mix with the SM neutrinos with a 3×3 mixing matrix U_{ij}. These mixings induce the coupling of HNLs to electroweak gauge bosons $W^\pm, Z$ given by

$$\mathcal{L} \supset \frac{g_2}{\sqrt{2}} U_{ij} W_\mu^- \ell_i^\dagger \bar{\sigma}^\mu N_j + \frac{g_2}{2\cos\theta_W} U_{ij} Z_\mu \nu_i^\dagger \bar{\sigma}^\mu N_j + \text{h.c.}\,, \tag{2.3}$$

where g_2 is the $SU(2)_L$ gauge coupling constant and θ_W is the Weinberg angle.

In our analysis, we will take a phenomenological approach, often adopted in the literature, and assume the existence of one HNL that mixes with only one SM neutrino flavor. We will separately study the cases in which the HNL mixes with the SM electron neutrino with a mixing angle denoted by U_e, with the SM muon neutrino with U_μ, and with the SM τ neutrino with U_τ. The phenomenology is completely determined once the HNL mass, m_N, and mixing angle, U_α, are fixed. We assume that the HNL is a Dirac fermion for the remainder of the paper.

2.2.1 HNL Production at an Electron Beam Dump Experiment

The mixings generated by the Lagrangian in eq. (2.2) allow for HNLs to be produced in any process where a SM neutrino is produced. This allows for the production of HNLs at an electron beam dump experiment in: (1) charged-current electron-proton scattering; (2) two- or three-body (semi-) leptonic decays of mesons (e.g. $M^\pm \to \ell^\pm N$, $M^\pm \to M'^0 \ell^\pm N$ with M, M' being different mesons); (3) secondary production from the decays of $\tau^\pm$ produced from the decays of heavy mesons such as $D_s^\pm$.

The first process only happens to leading order in the case of electron-mixed HNLs. The corresponding production cross-section is given by (see also [89], where we neglect the contributions from heavy quark PDFs)

$$\sigma(ep \to Nn) \approx \frac{G_F^2 M_W^2}{2\pi} |V_{ud}|^2 |U_e|^2 \int_{m_N^2/s}^{1} dx \frac{(\hat{s} - m_N^2)^2}{\hat{s}(\hat{s} - m_N^2 + M_W^2)} f_u(x, q^2) \,, \tag{2.4}$$

where p and n denote a proton and a neutron, $\hat{s} = xs$ with $\sqrt{s} = \sqrt{2 m_p E_e}$ being the center-of-mass energy and x the fraction of the proton momentum carried by the initial state quark. The function $f_u(x, q^2)$ is the up-quark parton distribution function.

We have numerically checked that `MadGraph5_aMC@NLO` [90] gives approximately the same cross section.

To determine the number of electron-mixed HNLs produced directly through the charged-current exchange, we estimate the effective instantaneous luminosity $\mathcal{L}_{\text{eff}}$ of the electron-proton system in the first radiation length of the beam dump as

$$\mathcal{L}_{\text{eff}} \simeq \frac{Z(X_0/g)N_A}{A}\, EOT_s\,, \tag{2.5}$$

where $A(Z)$ is the atomic mass (number) of the target material, X_0 is the radiation length (36.08 g$/cm^2$ for liquid water), N_A is the Avogadro's number, and EOT$_s$ is the number of electrons-on-target per second. The values of $\mathcal{L}_{\text{eff}}$ for the different experiments are reported in table 2.1.

In the top left panel of figure 2.2 we show the number of HNLs produced in charge-current scattering for one year of operation of ILC-250 (solid curve), ILC-1000 (dashed curve), and C^3/CLIC-3000 (dotted curve) assuming a mixing angle of $|U_e|^2 = 10^{-6}$. As we will discuss, this production rate is generally much lower than the one obtained from pion, kaon, and D-meson decays. It is, instead, comparable to or larger than the production rate from B meson decays.

For all flavor alignments, HNLs can be copiously produced in two- or three-body decays of mesons, as well as secondary production from the decays of $\tau^\pm$ produced in the decay of $D_s^\pm$ meson if the HNL mixes exclusively with the τ neutrino. The number of HNLs produced is given by

$$N_N = N_M \text{BR}(M \to N + X)\,, \tag{2.6}$$

where N_M is the number of mesons or $\tau^\pm$ produced (see eq. (2.1) and the subsequent

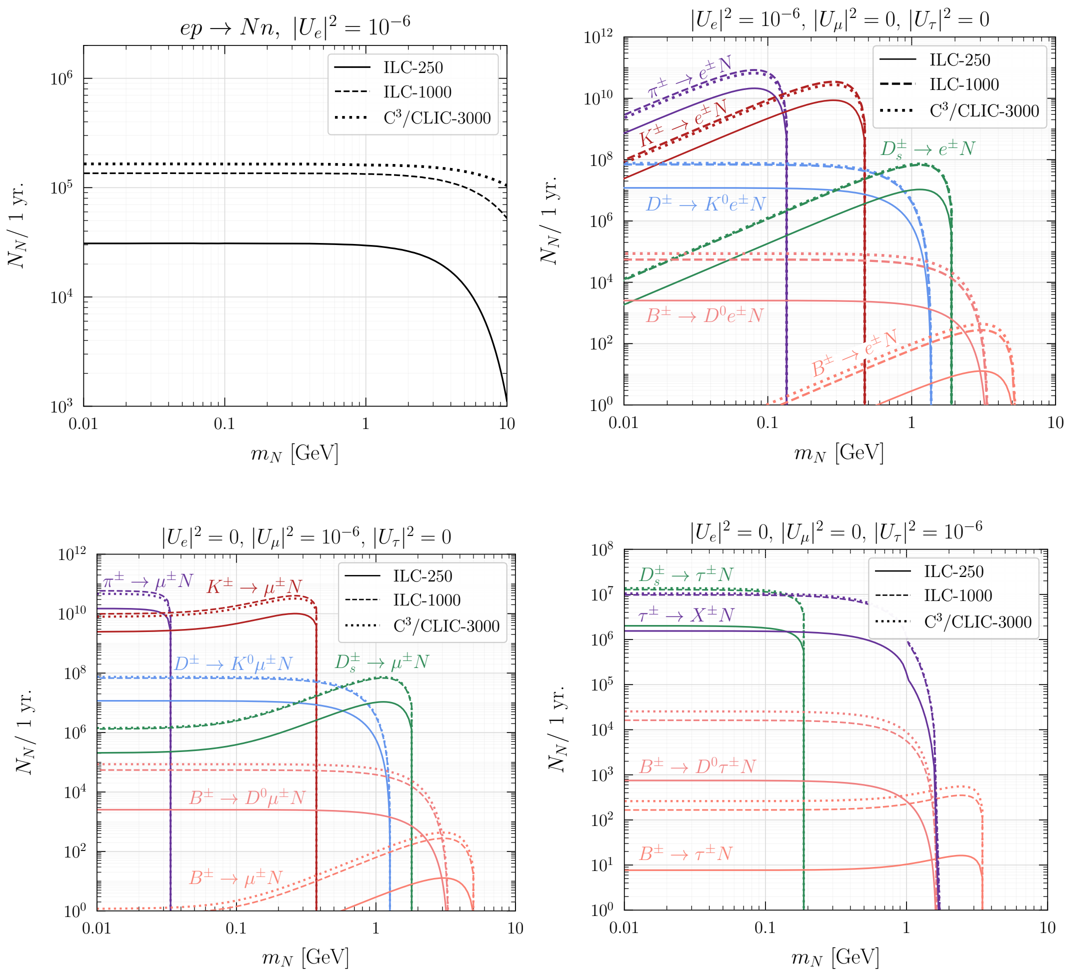

Figure 2.2: Number of HNLs produced per year at the several experiments, having fixed a mixing angle $|U_i|^2 = 10^{-6}$. *Top left:* Number of electron-mixed HNLs produced from charged-current electron-proton scattering. *Top right:* Number of electron-mixed HNLs produced from the decays of mesons. *Bottom*: Number of muon-mixed (left panel) and τ-mixed (right panel) HNLs produced from the decays of mesons and τ leptons.

discussion) and BR($M \to N + X$) is the branching ratio for the meson or $\tau^\pm$ decay to a HNL and final states X. The full expressions of all branching ratios can be found in [15, 91–93].

The total number of HNLs produced from meson and τ decays in an ILC and C^3/CLIC beam dump experiment is shown in the top right, bottom left, and bottom right panels of figure 2.2 for electron-mixed, muon-mixed, and τ-mixed HNLs, respectively, assuming one year of operation and a mixing angle of $|U_\alpha|^2 = 10^{-6}$. For all production channels, we sum over positively and negatively charged mesons using their respective production given in table 2.2.

The production of electron-mixed HNLs from meson decays are typically orders of magnitude larger than the direct production in charged-current scattering, and we expect the charged-current scattering to place weaker constraints for $m_N \lesssim m_B$. At higher masses, charged-current scattering becomes a relevant production mechanism. The production of HNLs at ILC-1000 and C^3/CLIC-3000 is generally larger than at ILC-250, especially if it arises from B meson decays, as well as charged current interactions.

2.2.2 HNL Decays

Once produced, HNLs will decay via the weak interactions in eq. (2.3) into two-body final states, $N \to \ell^\pm M$, $N \to \nu M$, where M is a meson, and three-body final states, $N \to f f' \nu_i$ for final state fermions f, f'. The partial widths for the several decays are given in [15, 94–96]. In figure 2.2.2 we show the branching ratios as a function of m_N for an electron-mixed (top-left), muon-mixed (top right), and τ-mixed (bottom) HNL. For HNL masses below 1 GeV, the total hadronic decay rate is given by the sum of the exclusive decays to mesons, while above 1 GeV we use the inclusive decay to quarks.

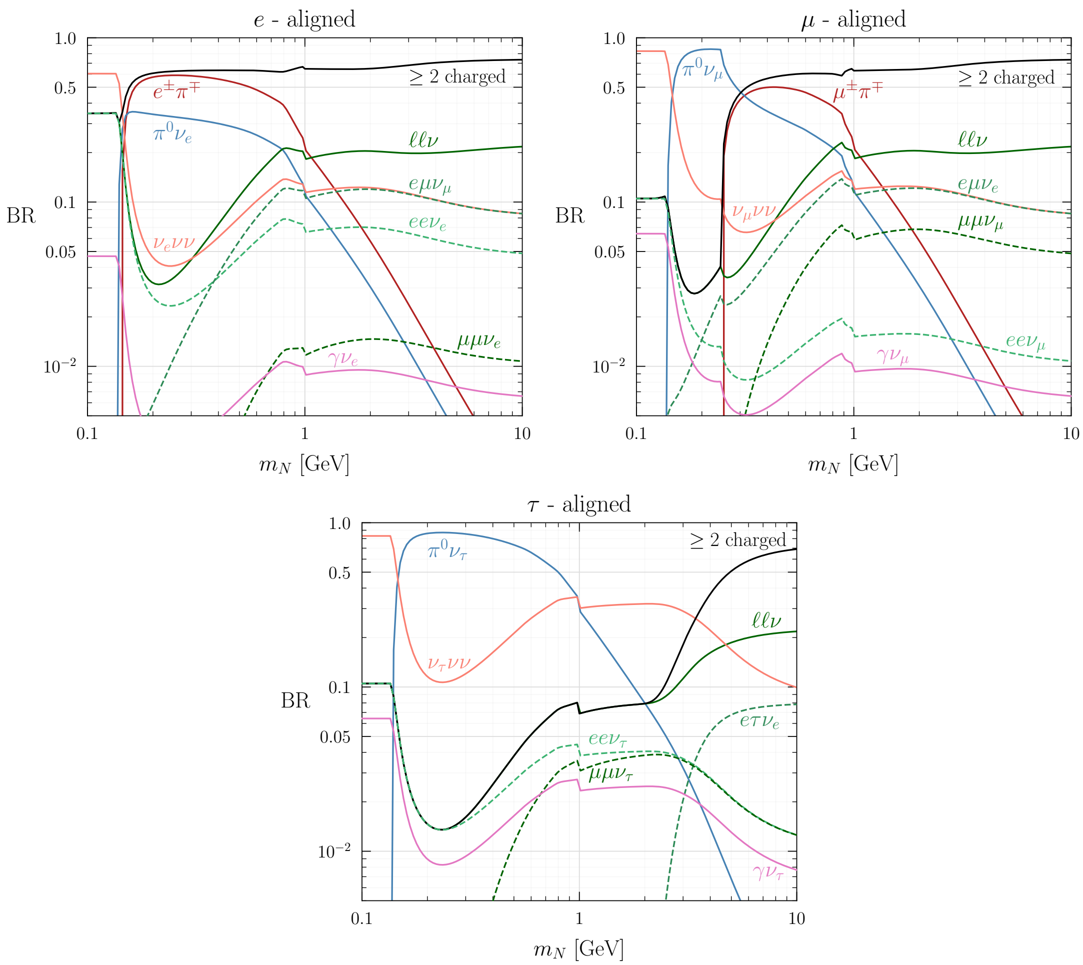

Figure 2.3: Branching ratios of an electron-mixed (left upper panel), muon-mixed (right upper panel), and τ-mixed (lower panel) HNL into different final states as a function of the HNL mass, m_N. The thick black curve represents the sum of the branching ratios into two or more charged particles in the final state, while the solid green curve labeled $\ell\ell\nu$ is the sum of the branching ratios into charged lepton final states.

The several branching ratios do not depend on the mixing angles $|U_i|^2$.

Among the possible decays, the most promising channels are those involving two (or more) charged particles in the final state, as those channels with only neutral particles (e.g., $\pi^0\nu$) will have larger SM backgrounds and are harder to be detected. To compute the reach on the HNL parameter space, we will take a conservative approach and only consider the HNL decay mode with at least two charged particles and with the largest branching ratio. For an electron- or muon-mixed HNL the most important decay for $m_N \lesssim m_\pi$ is $N \to ee\nu$ (dashed light-green curve), while for $m_\pi \lesssim m_N \lesssim 1$ GeV it is given by $N \to \ell^\pm\pi^\mp$ ($\ell = e, \mu$) (red curves). Above $m_N = 1$ GeV, the decay to two charged leptons will be the most relevant, with $N \to e\mu\nu$ being the dominant one. The most relevant decays of the τ-mixed HNL are to two charged leptons, as shown by the green curves in the bottom plot of figure 2.2.2.

2.2.3 Detector Efficiency and Acceptance

Given the experimental setup, we can determine the geometric acceptance and efficiency to detect the visibly decaying HNL. We require that the visible final states reach the detector after the HNL has decayed anywhere in the decay volume (see figure 2.1). For a given decay position, z, there will be a requirement on the direction of the visible final states with respect to the beam axis. We require that the transverse displacement of each of the two final states r_f is smaller than the radius of the detector, $r_{\text{det}} = 2$ m. This condition is given by

$$r_f = \sqrt{x_f^2 + y_f^2} < r_{\text{det}}, \quad x_f, y_f = \frac{p^H_{x,y}}{p^H_z} z + \frac{p^f_{x,y}}{p^f_z}(\ell_{\text{dump}} + \ell_{\text{sh}} + \ell_{\text{decay}} - z), \tag{2.7}$$

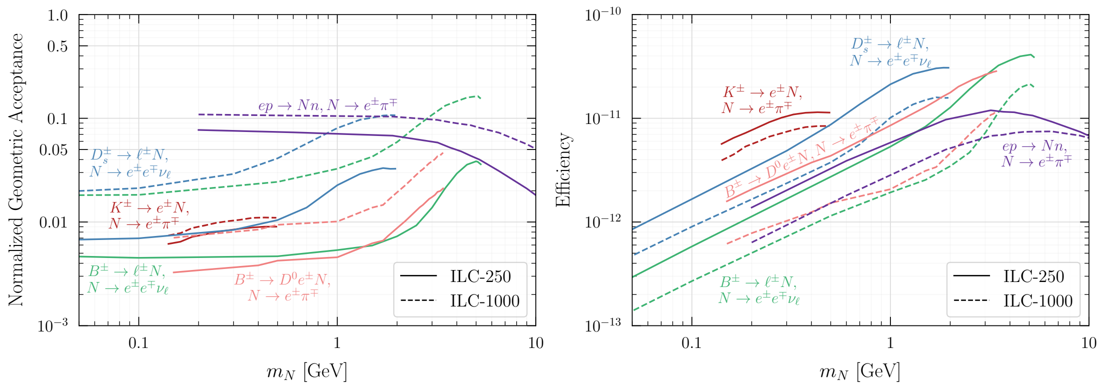

Figure 2.4: *Left:* Normalized geometric acceptance as a function of the HNL mass in the long lifetime regime (i.e., having fixed the proper lifetime to 30 seconds) for ILC-250 (solid curves) and ILC-1000 (dashed curves). *Right:* Efficiency as a function of HNL mass. See text for more details.

where $p^f_{x,y,z}$ are the fermion momentum components, and $p^H_{x,y,z}$ are the HNL momentum components. In fact, visible decay products with a transverse displacement $r_f > r_{\text{det}}$ will not reach the detector. Signal events are chosen to be those in which the HNL decays to two charged particles that are within the fiducial decay region enclosed by $z \in (z_{\text{min}}, z_{\text{max}})$, where $z_{\text{min}} = \ell_{\text{dump}} + \ell_{\text{sh}} = 81$ m, $z_{\text{max}} = \ell_{\text{dump}} + \ell_{\text{sh}} + \ell_{\text{decay}} = 131$ m, and satisfy eq. (2.7).

Since the geometric acceptance depends on the particular location of the HNL decay, the total efficiency is given by

$$\text{eff} = m_N \Gamma \int_{z_{\text{min}}}^{z_{\text{max}}} dz \sum_{\text{events} \in \text{geom.}} \frac{e^{-z m_N \Gamma / p_z}}{N_{\text{MC}} p_z}\,, \tag{2.8}$$

where m_N, Γ, and p_z are the mass, decay width, and $z-$component of the momentum of the HNL, respectively. The sum is over the events that fall within the geometric acceptance of the detector, and N_{MC} is the total number of simulated events.

The efficiency as a function of the HNL mass is shown in the right panel of figure 2.4

for a few representative HNL production and decay channels at the ILC-250 and ILC-1000. All of these curves apply to an electron-mixed HNL while the blue and green curves apply also to a muon-mixed HNL. In computing these curves, we assume a large proper lifetime of the HNL, 30 seconds, and, therefore, a decay length much larger than the experimental setup. In this limit, we can approximate the efficiency as

$$\text{eff} \approx m_N \Gamma (z_{\text{max}} - z_{\text{min}}) \sum_{\text{events} \,\in\, \text{geom.}} \frac{1}{N_{\text{MC}} p_z} \,. \tag{2.9}$$

As expected, in this regime the efficiency depends linearly on the $|U_\alpha|^2$. In the right plot of figure 2.4 we can see that the efficiency is smaller for ILC-1000 (dashed curves) compared to ILC-250 (solid curves). This is because, at higher energies, the HNLs produced are more boosted in the forward direction, which gives a smaller probability that the HNL will decay within the decay volume of the experiment (in the long lifetime regime). For brevity, we omit the corresponding C^3/CLIC-3000 curves from figure 2.4. The same conclusions hold due to its higher energy: the efficiency is lower than ILC-1000.

The geometric acceptance shows the opposite behavior. In the left panel of figure 2.4, we show the acceptance of ILC-250 (solid lines) and ILC-1000 (dashed lines). We compute the acceptance by taking the ratio between the efficiency in eq. (2.8) with or without the requirement to be in geometric acceptance (eq. (2.7)). The geometric acceptance is larger at ILC-1000 compared to ILC-250 reflecting that the final states are more forward, and, therefore, the number of events that satisfy the geometric acceptance requirement is higher. Overall, the geometric acceptance is relatively large, thanks to the sizable boost of the HNLs produced in these experiments.

§ 2.3 Sensitivity for Heavy Neutral Leptons

2.3.1 Sensitivity of Lepton Collider Beam Dump Experiments

The expected reach of a high-energy electron beam dump experiment can be found using the production, decay, and experimental efficiency information discussed in the previous section. The total number of signal events is given by

$$N_{\text{sig}} = N_N \times \text{BR}(N \to i) \times \text{eff}_i\,, \tag{2.10}$$

where N_N is the number of HNLs produced in a particular production mode, BR$(N \to i)$ is the branching ratio for N decaying to a final state i, and eff$_i$ is the efficiency for detecting the final states. For each mass point, we consider only the HNL decay mode that leads to the best sensitivity, as discussed in section 2.2.2. To determine the expected reach, we require 10 signal events in the detector. The reach is, therefore, relatively conservative (other studies for the reach of the ILC beam dump on dark sector parameter space assume zero background and require 3 signal events ADD PAST PAPERS (ref.48-51 i think. check if they all use 3 events).)

In figures 2.5, 2.6, and 2.7, we show the reach of high-energy electron beam dump experiments for electron-mixed, muon-mixed, and τ-mixed HNLs, respectively, assuming one year of operation. In the left plot of each figure the solid-black, dashed-black, and dotted-black curves represent the reach of the ILC-250, ILC-1000, and C^3/CLIC-3000, respectively. In the right plot of figure 2.5 and 2.6, we decompose the bounds into contributions from the decay of kaons (red curves), D mesons (blue curves), B mesons (green curves), and charged-current scattering (orange curves, only for figure 2.5). For figure 2.7 we decompose the bounds into contributions from D mesons (blue curves),

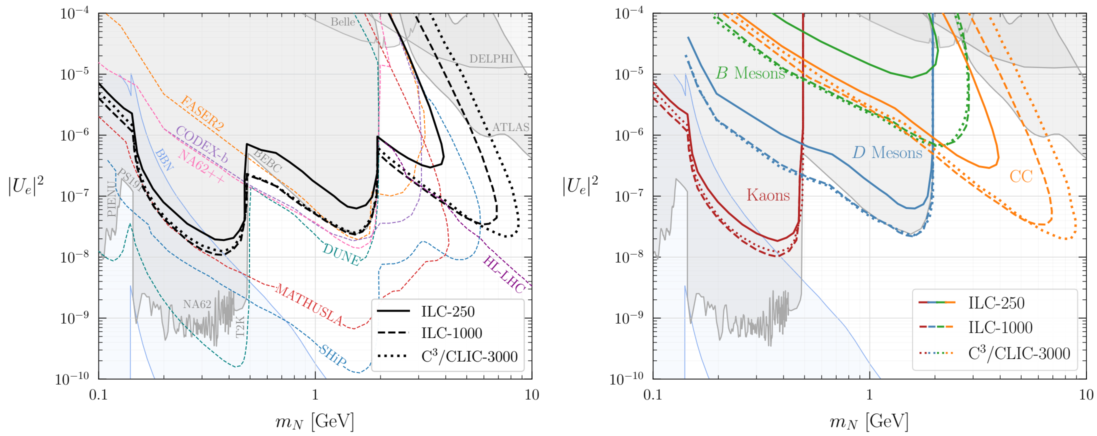

Figure 2.5: *Left:* The reach of beam dump experiments for electron-mixed HNLs in the $|U_e|^2$ vs m_N plane for ILC-250 (black, solid curve), ILC-1000 (black, dashed curve), and C^3/CLIC-3000 (black, dotted curve) assuming one year of operation. The limits are set by requiring 10 signal events in the detector. Existing experimental limits are depicted by the gray shaded region, while constraints from BBN are depicted by the blue shaded region. Limits from proposed experiments are depicted by the different dashed, colored curves. *Right:* Same as the left plot where we decompose the bounds into contributions from the decay of kaons (red curves), D mesons (blue curves), B mesons (green curves), and charged-current scattering (orange curves).

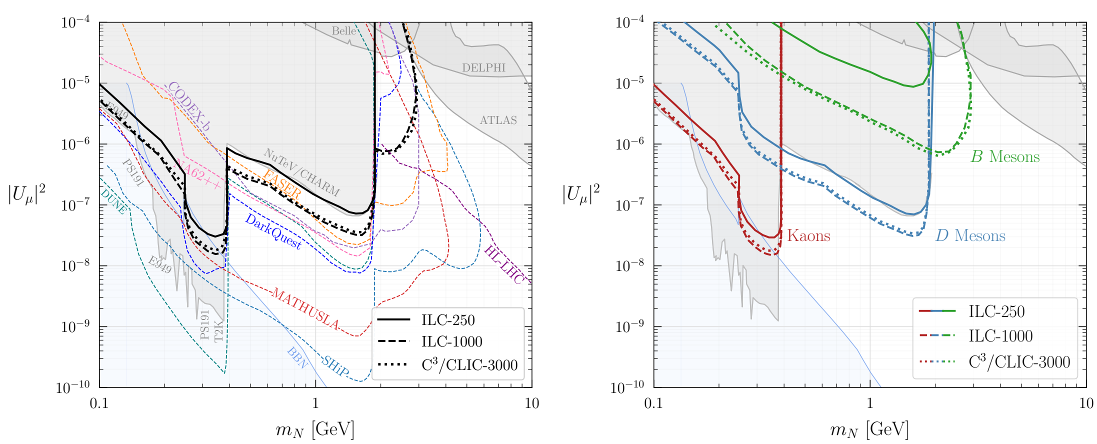

Figure 2.6: Same as in figure 2.5 but for muon-mixed HNLs.

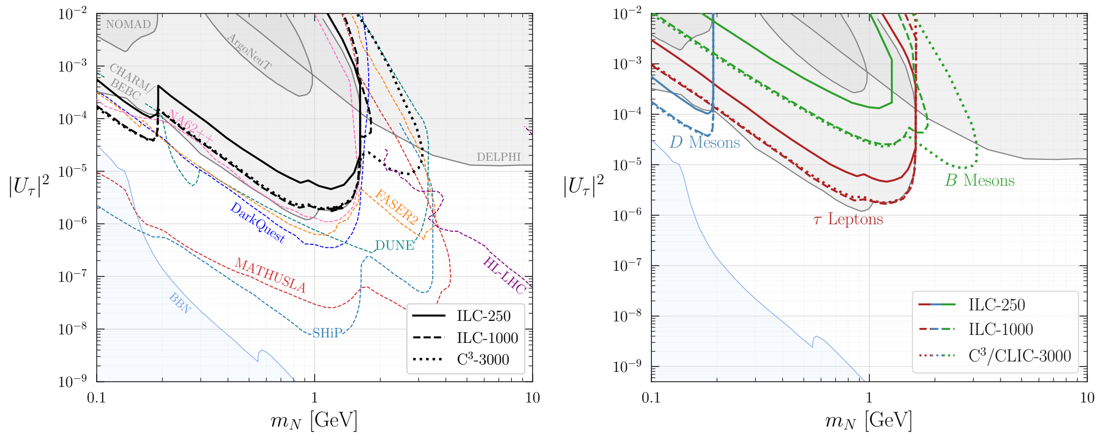

Figure 2.7: Same as in figure 2.5 but for τ-mixed HNLs.

$\tau^\pm$ (red curves), and B mesons (green curves).

The gray shaded regions are the leading existing constraints from PIENU [97, 98], PS191 [18, 19], NA62 [99], CHARM [100], BEBC [101, 102], Belle [48], DELPHI [103], and ATLAS [104] for electron-mixed HNLs; E949 [105], PS191 [18, 19], T2K [21, 22], NuTeV [106], CHARM [107], Belle [48], DELPHI [103], and ATLAS [104] for muon-mixed HNLS; and NOMAD [20], CHARM [17, 20, 108], BEBC [102], ArgoNeuT [25], and DELPHI [103] for τ-mixed HNLs. The shaded blue regions at small mixing angles are additional constraints from Big Bang nucleosynthesis (BBN) studies [109, 110]. HNLs can, in fact, be thermally produced in the early Universe. The decays of these HNLs in the plasma affect the abundance of light elements as well as the effective number of relativistic species, N_{eff}. Thus, the HNL parameter space can be constrained by BBN measurements. In all mixing scenarios, the ILC-1000, and C^3/CLIC-3000 will cover unexplored regions of parameter space. In the case of the electron-mixed scenario, also the ILC-250 will cover sizable new regions of parameter space.

We also compare the sensitivity of the ILC and C^3/CLIC beam dump experiments to the sensitivity of several proposed experiments such as CODEX-b [36], DarkQuest [37], DUNE [27–29], FASER [30], MATHUSLA [32, 34], NA62 [43], and SHiP [38–42] (see the colored dashed lines in the left plot of the figures).

For all HNL mixings, the sensitivity on $|U_\alpha|^2$ for $m_N \lesssim m_{D_s^\pm}, m_\tau$ improves only by a factor of a few when going from the ILC-250 to the ILC-1000 due the an increase in the number of mesons produced per electron-on-target. However, the ILC-1000 bound is not a simple rescaling with the number of mesons since the efficiency of ILC-1000 is lower compared to ILC-250, as shown in the right plot of figure 2.4. The HNLs are more boosted in the forward direction, resulting in a smaller probability that the HNL will decay in the decay volume. C^3/CLIC-3000 places similar limits as ILC-1000.

For $m_N > m_{D_s^\pm}$ using a higher energy electron beam leads to significant improvements in the bounds from HNL production in B decays. Approximately 10 times more B mesons are produced at ILC-1000 and C^3/CLIC-3000 compared to ILC-250. This helps increase the sensitivity for small mixing angles, $|U_{e,\mu}|^2 \lesssim 10^{-6}$ and $|U_\tau|^2 \lesssim 10^{-4}$. For larger mixing angles, $|U_{e,\mu}|^2 \gtrsim 10^{-6}$ and $|U_\tau|^2 \gtrsim 10^{-4}$, there are significant gains in sensitivity when going to higher energies. In this regime, the bound is determined by the location of the beginning of the decay volume: if the mixing angle is too large, the HNL will decay before it enters the decay volume. Then, for higher energy electron beams, the HNLs will have a larger boost factor and, therefore, a longer lifetime. Therefore, a larger mixing angle will be needed for the HNL to decay at the beginning of the decay volume.

Interestingly, the charged-current production of electron-mixed HNLs places the dominant sensitivity for $m_N > m_{D_s^\pm}$, and can significantly exceed the sensitivity from

B-meson decays due to the much lower B meson production rate (see the right plot of figure 2.5). Thanks to the charged-current production, HNLs with a mass up to $\sim 3, 7, 9$ GeV could be probed at ILC-250, ILC-1000, and C^3/CLIC-3000, respectively. This is a unique advantage of electron beam dump experiments over proton beam dump experiments where HNLs are mainly produced via meson decays and are limited by the heaviest mesons that can be produced at a given center-of-mass energy. Furthermore, contrary to the production from kaon and D meson decays, increasing the beam energy leads to a significant gain in sensitivity from charged-current production. This can be understood from the upper left plot of figure 2.2. ILC-1000 and C^3/CLIC-3000 can provide the leading sensitivity for electron-mixed HNLs in the 5 – 10 GeV mass range, providing a stronger reach than proton beam dump experiments.

In the left plots of Fig. 2.5-2.7 we also show the sensitivity of the HL-LHC [111], depicted by the dashed purple curves. For the electron-mixed case, we can see that the charged-current production of HNLs at ILC and C^3/CLIC beam dump experiments will be complementary to the HL-LHC. For the muon- and tau-mixed HNLs, the HL-LHC will have the strongest sensitivity beyond $m_N \simeq 2$ GeV.

2.3.2 Modifying the Experimental Setup

In the previous section, we analyzed the sensitivity of an ILC and C^3/CLIC beam dump experiment using the ILC beam configuration described in section 2.1.1. We have seen that increasing the electron beam energy leads to gains in the sensitivity for both small mixing angles (lower part of the curves) and for large mixing angles at high HNL masses (upper part of sensitivity curves). Here we show the impact of different parameters of the experimental setup in the reach of HNL parameter space.

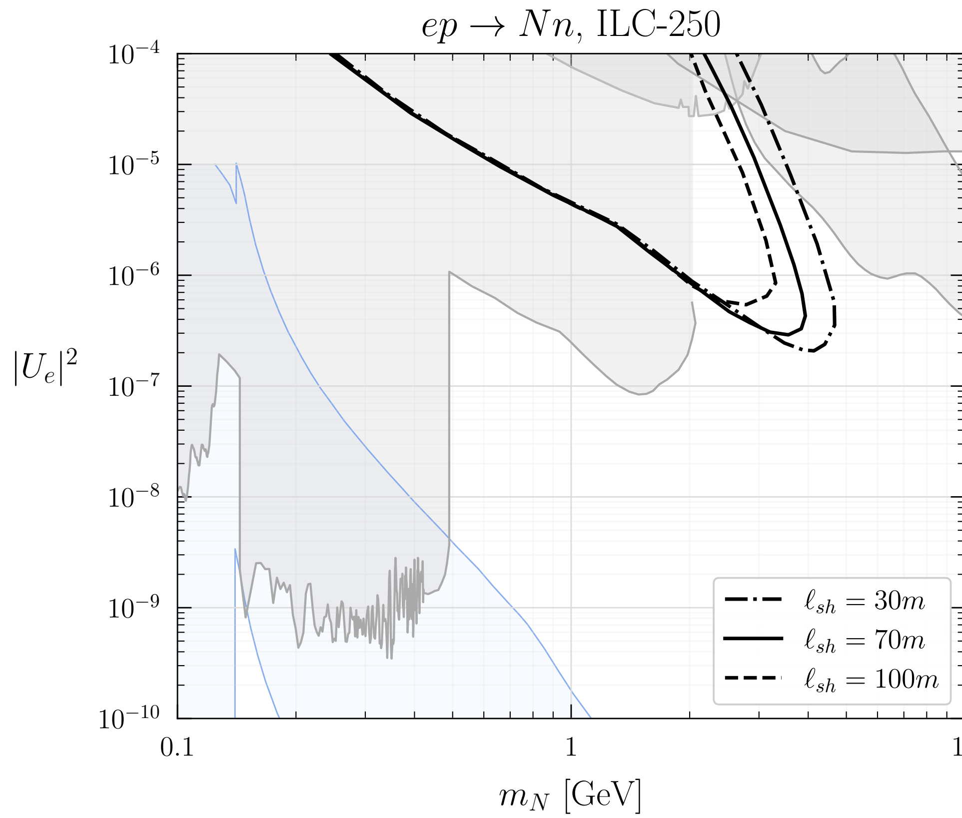

Figure 2.8: Variation of the bound on an electron-mixed HNL parameter space for different lengths of the muon shield ℓ_{sh} (fixing the distance to the dump ℓ_{dump}, the decay length ℓ_{decay}, and the energy of the beam, 125 GeV). For the purpose of the plot, we only consider the charged-current production of electron-mixed HNLs for ILC-250.

An adjustable parameter of the experimental setup is the length of the muon shield, ℓ_{sh}. Varying the length of the shield does not affect the bounds at small values of the mixing angle (the lower part of the bounds). In fact, in this regime, the efficiency varies linearly with the size of the decay length (see eq. (2.9)) which we keep fixed. On the other hand, varying the length of the shield will change the sensitivity for HNLs that are relatively short-lived (the upper part of the bounds). As a result, we only consider the direct production of electron-mixed HNLs from charged-current scattering, which has the best sensitivity in that region of parameter space.

We consider $\ell_{\text{sh}} = 30$ m, 70 m (i.e. the configuration described in section 2.1.1), and 300 m as benchmarks, while keeping all other parameters the same as the nominal beam dump configuration (see figure 2.1). The resulting bounds for ILC-250 are shown in figure 2.8 for the example of an electron-mixed HNL. The dot-dashed, solid, and dashed curves are for a shield length of 30 m, 70 m, and 300 m, respectively. As expected, changing the length of the shield only affects large values of the mixing angle, and we can see that a shorter shield length improves the sensitivity while a longer shield length decreases the sensitivity, compared to the nominal beam dump set up with $\ell_{\text{sh}} = 70$ m. This seems to hint that a small muon shield would be beneficial to obtain a better reach with a lower energy electron beam, potentially closing the gap in parameter space between 2–5 GeV. However, as we discuss in section 2.3.3, a muon shield that is too thin can lead to sizable SM backgrounds, depleting, therefore, the overall reach.

2.3.3 Possible Backgrounds

Long-lived mesons, especially K_L^0, could be background to the searches proposed in this paper. With the experimental setup adopted in our paper, $\mathcal{O}(10^{16} - 10^{17})$ K_L^0 will be produced in the first radiation length of the dump. The water dump will be able to absorb some of these kaons, reducing the number to $\mathcal{O}(10^{12} - 10^{13})$. As discussed in [69], it is assumed that the inner parts of the muon shield are made of lead. Considering the pion and kaon interaction length in lead (20 cm), we find that roughly 6 meters of muon shield will be able to reduce the kaon rates to negligible values. In all our estimates of the reach on HNL parameter space, we assume $l_{\text{sh}} \gtrsim 30$ m, and therefore the K_L^0 background should be negligible.

In addition, a large flux of muons will be produced. Approximately $\mathcal{O}(10^{18})$ muons will be produced in the first radiation length of the dump. The average energies of these muons are roughly 10 GeV, 20 GeV and 30 GeV at ILC-250, ILC-1000, and C^3/CLIC-3000, respectively. The muon energy distribution is broad, with sizable populated tails reaching 100 GeV, even at the ILC-250. The muons will lose only a few GeV energy in the water dump. In the lead muon shield, the muons will lose approximately 1 GeV/m. These considerations show that additional veto counters placed after the muon shield would be needed to reduce the muon background to a negligible rate for a muon shield of length $\ell_{\rm sh}$ = 30m or 70m. Here I cancel the discussion. not sure it is necessary. We could add: In addition, veto counters surrounding the detector will serve to reject cosmic rays. See [67] for more details.

The specific design of the detector will also affect the background consideration. As a phenomenological paper, we give a simple description and guidelines of the preferred detector, given examples of historically successful experiments. The auxiliary CHARM/NuCal decay detectors [100, 112, 113], are perfect examples for our proposal. To be more specific, a low density is preferred for the decay region, to reduce the SM backgrounds. A magnet to help separate the charged particle decay product and to bend background muons out is beneficial, but not necessary. Timing information is preferred to reduce potential background, so a fast (the ability to identify each beam bunch collision) detector is preferred. The technology is completely available today.

§ 2.4 Summary

We have investigated the sensitivity on the heavy neutral lepton parameter space of beam dump experiments that utilize the high energy electron beam of future e^+e^- colliders. We focused on the 250 GeV and 1000 GeV ILC stages, as well as on the 3 TeV stages of C^3 and CLIC experiments. In fact, linear colliders offer a better opportunity for high-energy electron beam experiments compared to circular colliders thanks to their much larger number of electrons dumped per second.

We pointed out that HNLs can be copiously produced in the beam dump of high energy, high intensity e^+e^- colliders either from meson decays, or from charged current electron-proton scattering. These production mechanisms are complementary to the production processes that occur in the main e^+e^- collider. Once produced the HNLs can travel a macroscopic distance before decaying into SM particles, that can be detected if a suitable detector is installed behind the beam dump. These signatures are again complementary to the signatures that can be looked for at the main detectors of the e^+e^- colliders, and, depending on the specific experimental apparatus, could have a very small or negligible SM background.

We showed that, thanks to their unprecedented energies and intensities, these electron beam dump experiments will probe a large range of HNL parameter space for HNL masses below ~ 10 GeV, not yet probed by past searches. These experiments will be complementary to other proposed experiments such as proton beam dump experiments, neutrino experiments, and LHC auxiliary detectors. Particularly, in the case of an electron-mixed HNL, they will be able to reach higher masses thanks to the charged current electron-proton scattering process.

Chapter 3

Exploring Invisible New Physics with Exotic Pion Decays

This chapter is based on [114]. In this work, we investigate the potential of the PIONEER experiment to discover or constrain exotic decays of charged pions to invisible particles. We consider both two-body decays such as $\pi^+ \to \ell^+ N$, where N is a sterile neutrino, and three-body decays involving other dark sector particles in the final state, $\pi^+ \to \ell^+ \nu_\ell X$. Such three-body decays are a general prediction of well-motivated extensions of the SM, including models with axion-like particles (ALPs), light scalars, and dark vector bosons. The considered decay topologies capture a wide class of new physics scenarios in which light, weakly coupled states are produced in pion decays and escape detection, see e.g. [115–133]. We compare the reach of the PIONEER experiment with that of previous pion experiments, including the PIENU experiment [97, 98, 134–137], as well as the one of kaon factories, such as NA62 [99, 138, 139]. We demonstrate that PIONEER will have access to significant new regions of parameter space of invisible dark particles.

The paper is structured as follows: in Sec. 3.1, we give an overview of the PIENU and

PIONEER experiments, their main physics goals, and the possible searches for new light particles. In Sec. 3.2, we discuss exotic two-body decays of charged pions into sterile neutrinos. In Sec. 3.3, we analyze exotic three-body decays of charged pions using a simplified model framework to describe decays into a scalar, an axion-like particle, or a dark vector. In all cases, we compare the projected PIONEER sensitivity with existing experimental constraints on these simplified models. Finally, Sec. 3.4 summarizes our conclusions. Details on the recast of PIENU searches for exotic two-body and three-body decays are given in Appendices B and C, respectively.

§ 3.1 Exotic Pion Decays at PIENU and PIONEER

The main goal of Phase I of the proposed PIONEER experiment is the precision measurement of the ratio of the charged pion branching ratios to electrons and muons

$$R_\pi \equiv \frac{\text{BR}(\pi^+ \to e^+\nu_e(\gamma))}{\text{BR}(\pi^+ \to \mu^+\nu_\mu(\gamma))}\ . \tag{3.1}$$

This observable is extremely sensitive to a wide variety of new physics effects and is one of the most accurate probes of lepton flavor universality. The ratio can be predicted with remarkable precision in the SM, as uncertainties from the CKM matrix element V_{ud} and the pion decay constant f_π cancel, and radiative corrections (indicated by Δ_{rad} in the equation below) are well understood. The main uncertainty in the SM prediction is due to structure dependent radiative corrections that are estimated at next-to-leading order in chiral perturbation theory. The SM prediction reads [52, 140–143]

$$R_\pi^{\text{SM}} = \frac{m_e^2}{m_\mu^2}\frac{(m_\pi^2 - m_e^2)^2}{(m_\pi^2 - m_\mu^2)^2}\big(1 + \Delta_{\text{rad}}\big) = 1.23524(15) \times 10^{-4}\ . \tag{3.2}$$

On the experimental side, the most precise measurement of R_π comes from the PIENU experiment [97]. Combining the PIENU result with previous measurements [144–146], the PDG gives the following world average [147]

$$R_\pi^{\text{exp}} = 1.2327(23) \times 10^{-4} \,. \tag{3.3}$$

Combining with the SM prediction from above, this gives the following 90% C.L. limits on new physics contributions to R_π

$$-5.1 \times 10^{-3} < R_\pi / R_\pi^{\text{SM}} - 1 < 1.0 \times 10^{-3} \,. \tag{3.4}$$

The goal of PIONEER Phase I is to measure R_π with a precision of 0.01%, a factor of ~ 20 more precisely than the PIENU measurement. Assuming the central value will coincide with the SM prediction, we find the following sensitivity to new physics contributions at 90% C.L.

$$-2.6 \times 10^{-4} < R_\pi / R_\pi^{\text{SM}} - 1 < 2.6 \times 10^{-4} \,. \tag{3.5}$$

The R_π measurement by itself can be used to constrain regions of parameter space of models that predict exotic decays of charged pions to light new physics particles. In many scenarios however, dedicated searches for the distinct signatures of exotic pion decays can lead to a even better sensitivity. Stopped pion experiments combine very large pion yields with a simple and known initial state (the stopped pion), giving access to potentially very small exotic branching ratios.

As emphasized for example in [118], searching for exotic decays of charged pions is very well motivated, considering the small charged pion decay width, $\Gamma_{\pi^+} \simeq 2.5 \times 10^{-17}$ GeV, and the fact that the prominent SM decay modes $\pi^+ \to e^+\nu_e$ [97] and

$\pi^+ \to \pi^0 e^+ \nu_e$ [148] are helicity or phase space suppressed and therefore have very small branching ratios. Even very feeble couplings of a new physics particle to the pion or leptons could lead to a detectable exotic branching ratio.

Exotic two-body decays into a sterile neutrino $\pi^+ \to \ell^+ N$ are particularly clean, as they yield a mono-energetic charged lepton whose energy shifts if N has a mass. This makes peak searches in the lepton energy spectrum extremely powerful. As we will discuss in section 3.2, the best limits on certain sterile neutrinos in the mass range below the pion mass come indeed from charged pion decay experiments. Also the production of new light particles in three-body decays, $\pi^+ \to \ell^+ \nu_e X$, modifies the charged lepton energy spectrum. However, in contrast to the SM decays $\pi^+ \to \ell^+ \nu_\ell$ and the decays into sterile neutrinos $\pi^+ \to \ell^+ N$, the three-body decays give a continuous spectrum of charged lepton energies. Precise measurements of the charged lepton energy spectra can set stringent constraints on the parameter space of new physics models.

In addition to exotic pion decays, pion factories can also look for exotic decays of secondary muons given the timing separation between the pion stop and the muon decay. Thanks to the very large statistics of muons, as well as to the small muon decay width, $\Gamma_\mu \sim 3 \times 10^{-19}$ GeV, pion decay experiments looking for $\mu^+ \to e^+ X$ can have sensitivity to tiny flavor changing couplings of the muon.

For all these reasons, the search for light new particles has been broadly pursued by the PIENU experiment, which conducted searches for sterile neutrinos $\pi^+ \to e^+ N$ [98] and $\pi^+ \to \mu^+ N$ [135] as well as for exotic three-body decay modes of the charged pion $\pi^+ \to \ell^+ \nu_\ell X$ [137], and the two-body decay of the muon $\mu^+ \to e^+ X$ [136], with a light new particle X that escapes detection. The PIENU search for $\mu^+ \to e^+ X$ achieved a sensitivity that rivals the ones of past dedicated experiments [149, 150] (but cannot compete with future experiments searching for $\mu \to e$ transitions [151–154]).

Many searches for exotic pion decays can be performed in the future at the PIONEER experiment. One can classify the new exotic decay modes according to their multiplicity. We consider final states that contain two particles, as in $\pi^+ \to \ell^+ N$, where the new state N has to be electrically neutral and a fermion (to be specific, we will consider an electrically neutral lepton); or three particles, as in $\pi^+ \to \ell^+ \nu_\ell X$, with SM leptons $\ell = e, \mu$ and where the new state X has to be an electrically neutral boson. In the interesting region of parameter space, the sterile neutrinos are either long-lived and decay far outside the detector, or they decay into other invisible states.* In either case, they do not give any visible signature in the detector. For the boson X there are several options:

(i) X does not give any visible signatures because it is either long lived on detector scales or it decays into invisible particles;

(ii) X decays into fully visible resonant final states, in particular the two possible two-body modes $X \to \gamma\gamma$ or $X \to e^+e^-$;

(iii) X decays into semi-visible final states, e.g. three-body $X \to \gamma\gamma Y$ or $X \to e^+e^- Y$ with another invisible state Y.

The visible and semi-visible decays can be either prompt or displaced. Displacements up to a few mm might be accessible to the PIONEER experiment whose active target is 6 mm thick [52].

Systematically studying PIONEER's sensitivity to the full set of signatures is an important task. In this article, we will focus on two-body and three-body pion decays

*As discussed in section 3.2.1, regions of parameter space where sterile neutrinos produced from pion decays decay promptly or slightly displaced to visible SM particles are already completely probed.

into invisible new physics particles. We leave the study of visible and semi-visible exotic decays for future work [155].

The main features of the PIENU and PIONEER experiments that are relevant for our study are reported in Table 3.1. The PIENU experiment relied on a target built from a 8 mm thick scintillator where the incoming pion beam was stopped. Muons from pion decays at rest have a kinetic energy of approximately 4 MeV, thus mostly remain inside the target. A cylindrical NaI crystal calorimeter with a length of 19 radiation lengths was used to measure the energy of the emitted positrons. CsI crystals surrounded the NaI crystal to better contain electromagnetic showers. The setup is described in more details in [97].

Based on the lessons learned from PIENU, the PIONEER experiment will use a highly segmented active target (ATAR) with 4800 LGAD channels read out instead of a single 8 mm scintillator. This enables tracking inside the stopping target and thus it will be possible to better match incoming pions to outgoing positrons which in turn allows for higher pion beam rates that result in larger statics. Additionally, it becomes possible to observe the detailed stopping behavior of pions and muons, thus providing additional handles to distinguish between $\pi^+ \to e^+\nu(\gamma)$ events and $\pi \to \mu \to e$ decay chains.

Additionally, the PIONEER experiment aims for a spherical state of the art calorimeter. As for PIENU, it was designed to be 19 radiation lengths deep. The spherical shape is a major improvement as it covers a larger solid angle and provides a more uniform detector response as a function of the polar angle. This reduces the variation in the $\pi^+ \to e^+\nu(\gamma)$ line shape. Moreover, the PIONEER calorimeter will feature a much faster decay time compared to the PIENU detector, which is crucial to support the higher anticipated beam rate. Combined with the improved ATAR capabilities, the PIONEER

	PIENU	PIONEER Phase I
$N_{\pi\to e}$ reconstructed	10^7	2×10^8
$N_{\pi\to e}$ selected	$\approx 1.4\times10^6$	0.5×10^8
Signal to Background below $56\,\mathrm{MeV}$	≈ 0.3	≈ 10
Resolution Function (e^+)	$\frac{\sigma_E}{E}\approx 3\%$ @ $40\,\mathrm{MeV}$	$\frac{\sigma_E}{E}\approx\left(0.43+\frac{4.95}{\sqrt{E}}+\frac{41.3}{E}\right)\%$
$N_{\pi\to\mu\to e}$ selected	9×10^6	$\mathcal{O}\left(10^{10}\right)$
Resolution (μ^+ at $4\,\mathrm{MeV}$)	$\approx 4\%$	5%

Table 3.1: Summary of key experimental parameters relevant for exotic pion decay searches at the PIENU experiment and the PIONEER Phase I setup, including event yields, background levels, and energy resolutions for positrons and muons. The PIENU parameters are estimated from [98, 135]. (In the case of the muon decays, they correspond to the analysis with muon energies above 1.2 MeV.)

experiment aims to achieve a $\pi^+\to e^+\nu(\gamma)$ event selection that provides a sample pure enough for a direct line shape measurement across the detector.

§ 3.2 Exotic Two-Body Decays of Pions to Sterile Neutrinos

In this section, we investigate the two-body decays $\pi^+\to\ell^+N$ as probes of light sterile neutrinos that mix with the active neutrino flavors. We begin by outlining the theoretical framework and the relevant parameters governing sterile neutrino production in pion decays. We then reproduce existing constraints from the PIENU experiment and derive projected sensitivities for the PIONEER experiment.

3.2.1 Sterile neutrinos

Sterile neutrinos are a well-motivated extension of the Standard Model, introduced in many frameworks to account, e.g., for neutrino masses, dark matter, or anomalies in neutrino oscillation experiments. Pion decay experiments can search for sterile neutrinos in the mass range below the charged pion mass [91, 92]. See e.g. [16, 156, 157] for reviews.

Sterile neutrinos can mix with the SM neutrinos through the renormalizable neutrino portal. This mixing determines the strength of the weak interactions of the sterile neutrinos. The interactions with the W boson can be written as

$$\mathcal{L} \supset -\frac{g}{\sqrt{2}} U_{jN} W_\mu^- \bar{\ell}_j \gamma^\mu N \ + \ \text{h.c.}\,, \tag{3.6}$$

where g is the $SU(2)_L$ gauge coupling, N is the sterile neutrino, and ℓ_j, with $j = 1, 2, 3$, are the three generations of left-handed charged leptons of the SM. In this simplified model, the sterile neutrino parameter space consists of the three mixing angles U_{jN} and the mass of the sterile neutrino m_N. The sterile neutrino can be a Dirac fermion or a Majorana fermion.

In this minimal setup, sterile neutrinos decay only via their mixing with active neutrinos, with a decay width scaling as $\Gamma_N \propto |U_{jN}|^2 m_N^5/v^4$, with v the vacuum expectation value of the SM Higgs. For sterile neutrino masses below the pion mass, this scaling implies that even sizable mixing angles completely probed by past experiments, $|U_{jN}| \sim 0.1$, correspond to macroscopic decay lengths of the order of $\gtrsim 10$ m. Sterile neutrinos in this regime therefore escape the detector and manifest as missing energy in pion decay experiments. As we will mention in more detail in section 3.2.4, Big Bang Nucleosynthesis (BBN) places stringent bounds on such long-lived states, essentially excluding the minimal scenario in this mass range for interesting mixing angles. A way to evade the BBN bound is to allow additional decay modes of the sterile neutrino into dark-sector states (see e.g., [158]). These channels can reduce the lifetime to the level required by BBN ($\tau_N \lesssim 0.1$ s), while still ensuring that sterile neutrinos appear as missing energy in laboratory experiments.

A sterile neutrino of Majorana nature could plausibly be involved in the mechanism that provides masses to the SM neutrinos. Parametrically, the see-saw mechanism predicts that the SM neutrinos acquire a mass of the order of $m_\nu \sim |U_{jN}|^2 m_N$. From neutrino oscillations we know that at least two neutrinos are massive with masses $m_\nu \gtrsim 0.008\,\text{eV}$ and $m_\nu \gtrsim 0.05\,\text{eV}$ [159]. The sum of neutrino masses is constrained by cosmology to be below $\sum m_\nu \lesssim \mathcal{O}(0.1\,\text{eV})$ [160–162]. For these reasons, as we will adopt in the next section, a good representative choice for the see-saw region is $0.01\,\text{e}V < |U_{jN}|^2 m_N < 0.1\,\text{e}V$. We emphasize, however, that these boundaries are not sharply defined and should be regarded as indicative rather than exact.

3.2.2 Experimental signatures at pion decay experiments

The mixing of the active neutrinos with the sterile neutrino has two effects. First, it reduces the standard leptonic decay rates of the pion into a positron and electron-neutrino, or muon and muon-neutrino

$$\frac{\Gamma(\pi^+ \to e^+\nu_e)}{\Gamma(\pi^+ \to e^+\nu_e)_{\text{SM}}} = 1 - |U_{eN}|^2 \, , \quad \frac{\Gamma(\pi^+ \to \mu^+\nu_\mu)}{\Gamma(\pi^+ \to \mu^+\nu_\mu)_{\text{SM}}} = 1 - |U_{\mu N}|^2 \, . \tag{3.7}$$

Second, the mixing introduces new two-body decay modes of the pion into the sterile neutrino, $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$. If those decays are kinematically allowed, i.e. if $m_N < m_\pi - m_e$ or $m_N < m_\pi - m_\mu$, the corresponding branching ratios are to a very good approximation

$$\frac{\text{BR}(\pi^+ \to e^+ N)}{\text{BR}(\pi^+ \to e^+ \nu_e)} = \frac{|U_{eN}|^2}{1-|U_{eN}|^2} \frac{\rho(\delta_e^\pi, \delta_N^\pi)}{\delta_e^\pi (1-\delta_e^\pi)^2} \tag{3.8}$$

$$= \frac{|U_{eN}|^2}{1-|U_{eN}|^2} \frac{m_N^2}{m_e^2} \left(1 - \frac{m_e^2}{m_\pi^2}\right)^{-2} \sqrt{\lambda\left(1, \frac{m_N^2}{m_\pi^2}, \frac{m_e^2}{m_\pi^2}\right)} \times \left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_e^2}{m_N^2} + \frac{2m_e^2}{m_\pi^2} - \frac{m_e^4}{m_N^2 m_\pi^2}\right) \tag{3.9}$$

$$\simeq \frac{|U_{eN}|^2}{1-|U_{eN}|^2} \frac{m_N^2}{m_e^2} \left(1 - \frac{m_N^2}{m_\pi^2}\right)^2 , \tag{3.10}$$

where $\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + ac + bc)$. The first line is a compact form introduced in [91, 92] that uses the dimensionless variables $\delta_e^\pi = m_e^2/m_\pi^2$ and $\delta_N^\pi = m_N^2/m_\pi^2$ and the function $\rho(x, y) = [x + y - (x - y)^2]\ [\lambda(1, x, y)]^{1/2}$. The last line holds under the assumption that the electron is much lighter than the sterile neutrino, $m_e \ll m_N$. For the muonic decay $\pi^+ \to \mu^+ N$, expressions analogous to equations (3.8) and (3.9) hold, but the analog to equation (3.10) is never a good approximation because of the large muon mass.

The experimental signature of the new decay modes is a peak in the positron or muon energy spectrum at

$$\bar{E}_e = \frac{m_\pi}{2}\left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_e^2}{m_\pi^2}\right) , \quad \bar{E}_\mu = \frac{m_\pi}{2}\left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_\mu^2}{m_\pi^2}\right) . \tag{3.11}$$

If the positron or muon energy spectra of the SM processes $\pi^+ \to e^+\nu_e$ and $\pi^+ \to \mu^+\nu_\mu$ (including bremsstrahlung and detector effects) and all relevant backgrounds can be modeled reliably, one can perform a bump hunt for the sterile neutrino signal.

Alternatively, one can simply consider the precisely measured ratio R_π and compare it to the SM prediction. One has to distinguish three mass regimes:

(1) $m_N < m_\pi - m_\mu$: in this regime both the $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$ decays are kinematically open, and one has

$$
\begin{aligned}
R_\pi = \frac{\text{BR}(\pi^+ \to e^+ + \text{invisible})}{\text{BR}(\pi^+ \to \mu^+ + \text{invisible})} &= \frac{\text{BR}(\pi^+ \to e^+\nu_e) + r_e\, \text{BR}(\pi^+ \to e^+ N)}{\text{BR}(\pi^+ \to \mu^+\nu_\mu) + r_\mu\, \text{BR}(\pi^+ \to \mu^+ N)} \\
= R_\pi^{\text{SM}} \times \Bigg(1 - |U_{eN}|^2 \Bigg[1 - \frac{m_N^2}{m_e^2}\left(1 - \frac{m_e^2}{m_\pi^2}\right)^{-2} &\sqrt{\lambda\left(1, \frac{m_N^2}{m_\pi^2}, \frac{m_e^2}{m_\pi^2}\right)} \\
\times \left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_e^2}{m_N^2} + \frac{2m_e^2}{m_\pi^2} - \frac{m_e^4}{m_N^2 m_\pi^2}\right)\Bigg]\Bigg) \\
\times \Bigg(1 - |U_{\mu N}|^2 \Bigg[1 - \frac{m_N^2}{m_\mu^2}\left(1 - \frac{m_\mu^2}{m_\pi^2}\right)^{-2} &\sqrt{\lambda\left(1, \frac{m_N^2}{m_\pi^2}, \frac{m_\mu^2}{m_\pi^2}\right)} \\
\times \left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_\mu^2}{m_N^2} + \frac{2m_\mu^2}{m_\pi^2} - \frac{m_\mu^4}{m_N^2 m_\pi^2}\right)\Bigg]\Bigg)^{-1} . \qquad (3.12)
\end{aligned}
$$

The SM prediction R_π^{SM} is given in equation (3.2). The coefficients r_e and r_μ are analysis acceptance ratios of the decays into sterile neutrinos versus the decays into SM neutrinos. We have set them to $r_e = r_\mu = 1$ in the second line. Deviations from 1 may arise if strong selection cuts are used in the analysis.

(2) $m_\pi - m_\mu < m_N < m_\pi - m_e$: for such sterile neutrino masses one has instead

$$
\begin{aligned}
R_\pi = R_\pi^{\text{SM}} \times \Bigg(1 - |U_{eN}|^2 \Bigg[1 - \frac{m_N^2}{m_e^2}\left(1 - \frac{m_e^2}{m_\pi^2}\right)^{-2} &\sqrt{\lambda\left(1, \frac{m_N^2}{m_\pi^2}, \frac{m_e^2}{m_\pi^2}\right)} \\
\times \left(1 - \frac{m_N^2}{m_\pi^2} + \frac{m_e^2}{m_N^2} + \frac{2m_e^2}{m_\pi^2} - \frac{m_e^4}{m_N^2 m_\pi^2}\right)\Bigg]\Bigg) \frac{1}{1 - |U_{\mu N}|^2} , \qquad (3.13)
\end{aligned}
$$

where we again set the acceptance ratio $r_e = 1$. The ratio r_μ is not relevant as the $\pi^+ \to \mu^+ N$ decay is not open.

(3) $m_\pi - m_e < m_N$: in this mass regime, the only observable effect is the modifica-

tion of the $\pi^+ \to e^+\nu_e$ and $\pi^+ \to \mu^+\nu_\mu$ rates due to the mixing, cf. equation (3.7),

$$R_\pi = R_\pi^{\rm SM} \times \frac{1-|U_{eN}|^2}{1-|U_{\mu N}|^2}\,. \tag{3.14}$$

Measurements of R_π can be used to constrain the sterile neutrino parameters. In this approach, one does not need to precisely know the positron and muon energy spectra of $\pi^+ \to e^+\nu_e$ and $\pi^+ \to \mu^+\nu_\mu$, and the derived results are arguably more robust. However, without spectral information, one does not disentangle the effects in the electron and muon channels and one cannot distinguish a sterile neutrino effect from the effect of other new physics that modifies R_π.

3.2.3 Reproducing the existing constraints on $|U_{eN}|$ and $|U_{\mu N}|$

Bump searches for $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$ have been performed at PIENU [98, 135] (See [163–173] for previous sterile neutrino searches in pion decays). Upper limits on the mixing parameter $|U_{eN}|^2$ were set in the ballpark of 10^{-7} to 10^{-8}, for sterile neutrino masses in the range of $62-134$ MeV, corresponding to a positron energy of $E_e \lesssim 56$ MeV. For smaller masses, the sterile neutrino peak moves to larger positron energies and it becomes increasingly challenging to precisely model the SM spectrum (see the discussion in section 3.2.4 below). Similar constraints on $|U_{eN}|^2$ were recently obtained by the NA62 experiment in the mass range $95-126$ MeV [139] (See [105, 174–176] for earlier sterile neutrino searches in kaon decays at higher masses.). The mixing parameter $|U_{\mu N}|^2$ is constrained by PIENU at the level of $\sim 10^{-5}$ for sterile neutrino masses between 16 MeV and 34 MeV.

Before producing sensitivity estimates for PIONEER, we first reproduced the PIENU results from the bump searches, both in $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$ to validate our approach. Details about the PIENU recast that closely mirrors our PIONEER sensitivity study are given in appendix B.

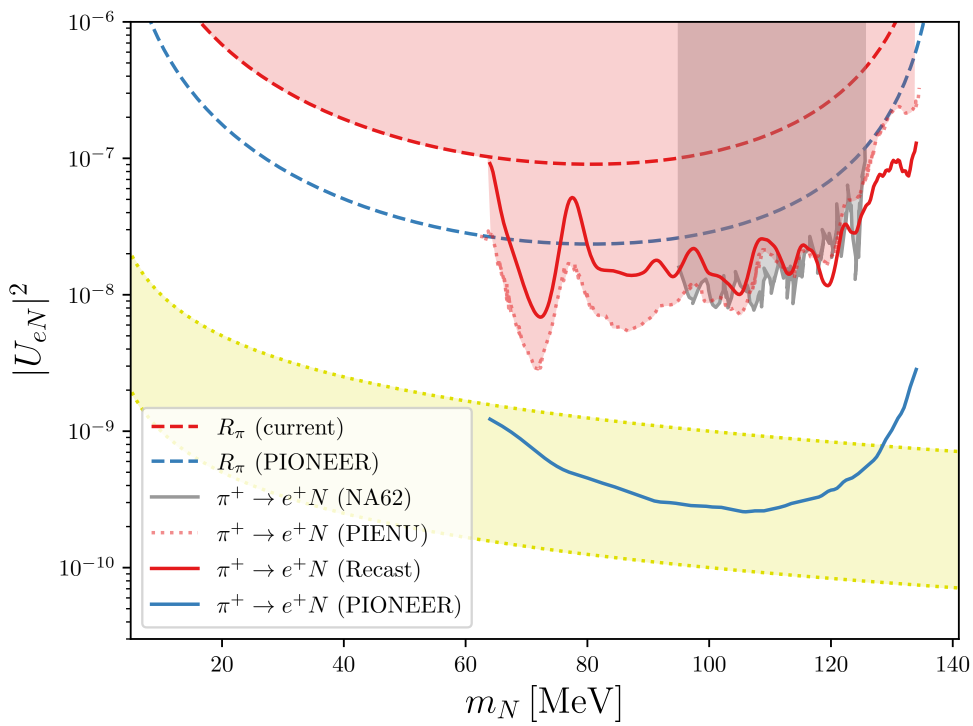

Figure 3.1: Constraints from pion decays $\pi^+ \to e^+N$ on the sterile neutrino mixing angle $|U_{eN}|^2$ at 90% C.L. as a function of the sterile neutrino mass m_N. The PIENU result from [98] is shown by the dotted red line. Our reproduction of the PIENU bound is shown in solid red. The corresponding limit from NA62 [139] is shown in gray. The bound from the R_π measurement is shown in dashed red. The red and gray shaded regions are excluded. The solid blue and the dashed blue curves correspond to our sensitivity estimates of a PIONEER search for sterile neutrinos and the PIONEER R_π measurement, respectively. The yellow shaded region indicates a see-saw target (see text for more details). This region can be reached by PIONEER for 64 MeV $\lesssim m_N \lesssim$ 130 MeV.

From the $\pi^+ \to e^+N$ analysis, we obtain the 90% C.L. bound on the sterile neutrino mixing angle with electron neutrinos, $|U_{eN}|^2$, shown in figure 3.1. Our reproduction of the bound is given by the solid red curve as a function of the sterile neutrino mass in the range between 62 MeV and 134 MeV. Our reproduction agrees reasonably well with the bound obtained by the PIENU analysis [98] shown in dotted red. The bound we find is slightly weaker for low masses and slightly stronger for high masses. For comparison, the plot also shows the limit from NA62 [139] in gray. We also show, as a dashed red curve, the bound we derived from existing measurements of R_π [177] which is in good agreement with the one obtained in [157]. To find this bound, we assume that the sterile neutrino mixes only with the electron neutrino, and that no additional new physics affects R_π. For sterile neutrino masses between 64 MeV and 135 MeV, the bump hunt provides constraints that are more stringent by approximately one order of magnitude.

Similarly, the limits on $|U_{\mu N}|^2$ that we reproduce from PIENU's $\pi^+ \to \mu^+ N$ analysis are shown in figure 3.2. The solid red curves correspond to our reproduction of the 90% C.L. bounds from PIENU shown in dotted red [135]. Overall, we find remarkable good agreement. The constraint from the R_π measurement (assuming no additional new physics than a sterile neutrino mixing with a muon neutrino) is shown by the red dashed curve.

3.2.4 Expected sensitivities at PIONEER

After successfully reproducing the limits on the sterile neutrino mixing angles from PIENU, we next estimate the sensitivity that can be expected at the PIONEER experiment.

To assess the sensitivity of PIONEER to sterile neutrinos, we follow the strategy of bump hunts in the charged lepton energy spectrum, similar to the searches performed

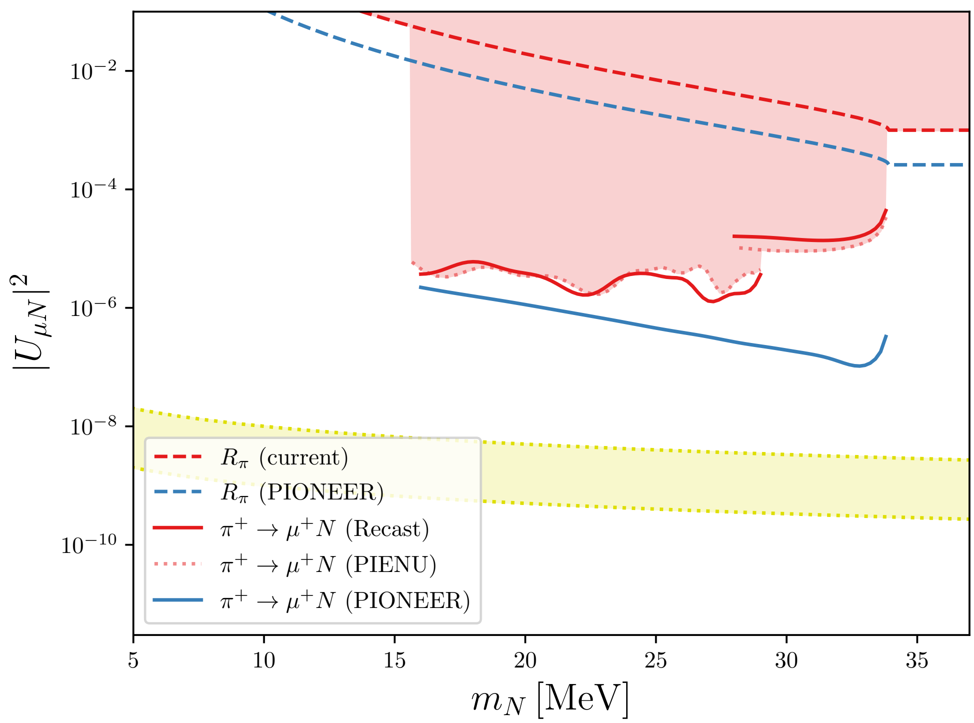

Figure 3.2: Constraints on the sterile neutrino mixing angle $|U_{\mu N}|^2$ at 90% C.L. as a function of the sterile neutrino mass m_N. The PIENU result from [135] is shown by the dotted red lines. Our reproduction of the PIENU bound is shown in solid red. The bound from the R_π measurement is shown in dashed red. The shaded regions are excluded. The solid blue and the dashed blue curves correspond to our sensitivity estimates of a PIONEER search for sterile neutrinos and the PIONEER R_π measurement, respectively. The see-saw target (indicated with the yellow band) is out of the PIONEER reach.

at PIENU. In particular, we simulate bump searches in the channels $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$.

The $\pi^+ \to e^+ N$ search: The search for sterile neutrinos in the decay $\pi^+ \to e^+ N$ at PIONEER profits from the increased statistics, the much better suppression of the muon decay backgrounds at small positron energy, and the better energy resolution (see Table 3.1).

We used simulated data including both inclusive $\pi^+ \to e^+ \nu_e(\gamma)$ events as well as muon decay in flight events ($\pi^+ \to \mu^+ \nu_\mu$ where the $\mu^+ \to e^+ \nu_e \bar{\nu}_\mu$ follows before the muon comes to rest, within less than ≈ 13 ps). Preliminary studies suggest that other backgrounds can be suppressed well enough using time differences and the very characteristic 4 MeV energy deposited by a stopping muon along with other selection criteria involving the segmentation of the active target. Therefore, those backgrounds are considered negligible for this study. This assumption will need to be tested by the PIONEER experiment.

While the PIONEER experiment requires at least 2×10^8 reconstructed $\pi^+ \to e^+ \nu_e(\gamma)$ events in an unbiased fashion to measure the branching ratio $R_{e/\mu}$ to the desired precision, a set of background rejections will be used to measure the $\pi \to e\nu\,(\gamma)$ line shape. Until more detailed exotic search strategies are developed by the collaboration, it is reasonable to assume that those searches will follow a similar set of background rejections to find features without a corresponding SM explanation in this line shape. Early reconstruction and background suppression studies suggest that a sample with only 8.8% muon decay in flight contamination in the search region can be achieved while preserving 0.5×10^8 of all $\pi^+ \to e^+ \nu(\gamma)$ events [178, 179]. The corresponding positron energy spectra are shown in figure 3.3. The shapes of both the $\pi^+ \to e^+\nu_e$ decay and the muon decay in flight background are taken directly from Monte Carlo simulations [178, 179].

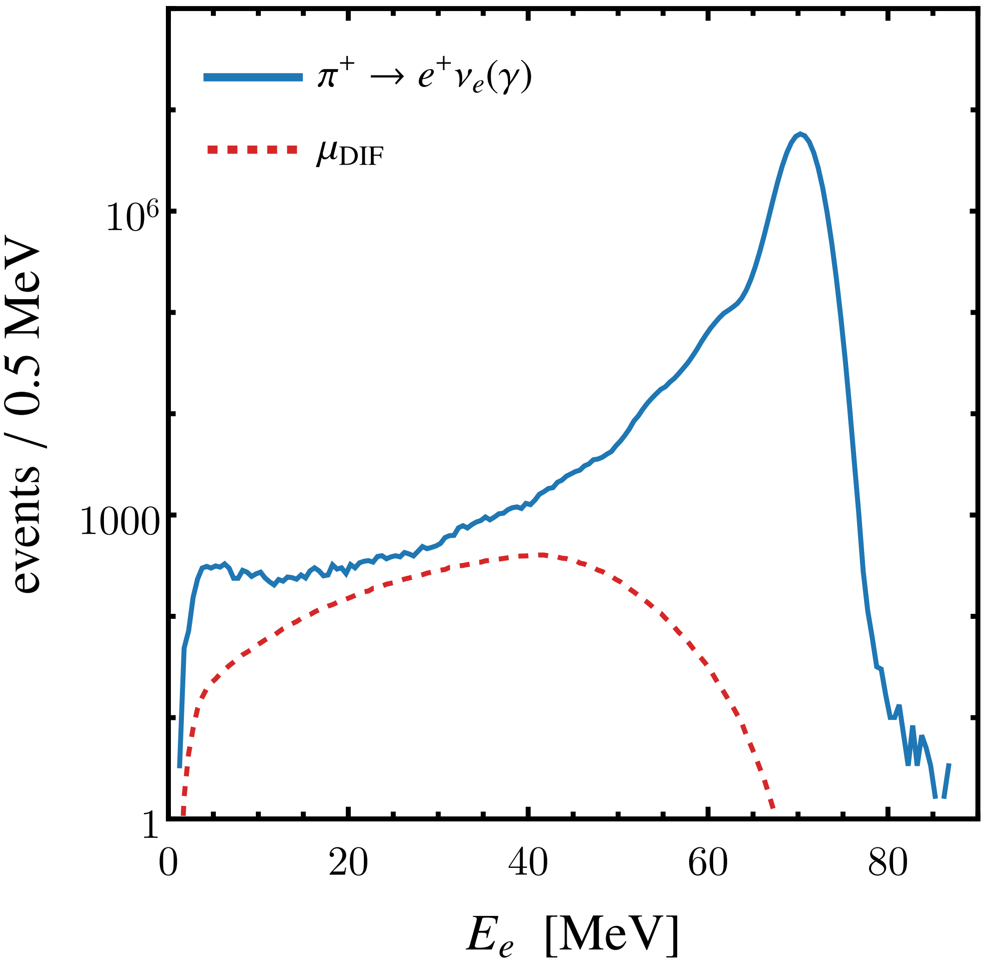

Figure 3.3: The expected positron energy spectrum at phase I of PIONEER after background suppression cuts. Shown in blue are 0.5×10^8 $\pi^+ \to e^+\nu_e(\gamma)$ events and in red dashed an 8.8% contamination from muon decays in flight, $\mu_{\rm DIF}$.

For the hypothetical sterile neutrino signal, we model the bump by a Gaussian peak whose width is determined by the expected energy resolution of the calorimeter. The energy dependence of the resolution is implemented using a parameterized curve and varies from $\sim 1.5\%$ for a positron energy of $E_e = 70\,\text{MeV}$ to $\sim 6\%$ for a positron energy of $E_e = 10\,\text{MeV}$ (see Table 3.1). The parameters are obtained using the full simulation chain which convolutes experimental calorimeter crystal responses taken from beam measurements [180], proper dead material models and early calibration strategies.*

The sensitivity of the sterile neutrino search is evaluated through a binned χ^2 fit of the simulated positron energy spectra. We use $0.5\,\text{MeV}$ wide bins and include all bins in the range $(4-56)\,\text{MeV}$, following the approach of the PIENU analysis. The χ^2 function accounts for statistical uncertainties only, based on an assumed total of 0.5×10^8 reconstructed $\pi^+ \to e^+\nu_e$ events. In the fit, the overall normalizations of the hypothetical sterile neutrino signal, the $\pi^+ \to e^+\nu_e$ component, and the muon decay in flight background are allowed to float while their shapes are kept fixed.

The $\pi^+ \to \mu^+ N$ search: We follow a similar approach for the sterile neutrino search in the decay $\pi^+ \to \mu^+ N$ at PIONEER. Our analysis is based on a simulated data set consisting of $\pi^+ \to \mu^+\nu_\mu$ events that produce muons with a kinetic energy of approximately $4\,\text{MeV}$ with a tail to lower energies. Additional background components are expected to be negligible as such low-energy muons are highly ionizing and stop within a few strips of the active target. This results in a characteristic Bragg peak in the energy loss rate per unit distance that can be used as selection criteria. The $4\,\text{MeV}$ peak has a width of about $0.2\,\text{MeV}$, and we model the sterile neutrino signal as a Gaussian peak of the same fixed width, independent of its position. The χ^2 fit is performed using $0.06\,\text{MeV}$ wide bins over the energy range $0\,\text{MeV}$ - $3.2\,\text{MeV}$, similar to

*As a robustness check, we have tested an alternative modeling of the signal shape by using the shape of the SM $\pi^+ \to e^+\nu_e$ peak and rescaling its width according to the relative energy resolution. This approach yields results consistent with the Gaussian modeling, confirming that our sensitivity projections are not very sensitive to the precise choice of signal parameterization.

the PIENU analysis. Only statistical uncertainties on the event counts in each bin are included, assuming a total of 10^{10} reconstructed $\pi^+ \to \mu^+\nu_\mu$ events. This is an early order-of-magnitude estimate based on few 10^{12} pion decays required to obtain 2×10^8 $\pi \to e\nu(\gamma)$ events, combined with a prescale factor of 50 and the fact that about 20% of pion decays will satisfy the $\pi \to \mu \to e$ trigger conditions on their own. A so far unknown amount of $\pi^+ \to \mu^+\nu_\mu$ events will have their trigger completed by accidental background, adding to the number. In the fit, the normalizations of the $\pi^+ \to \mu^+ N$ signal and the $\pi^+ \to \mu^+\nu_\mu$ component are allowed to float while their shapes are kept fixed.

The sensitivity projections that we obtain from the described bump searches for $\pi^+ \to e^+ N$ and $\pi^+ \to \mu^+ N$ are shown by the solid blue curves in figures 3.1 and 3.2. Under the assumption of no signal, we show the 90% C.L. upper limits on the mixing angle of the sterile neutrino with the electron neutrino $|U_{eN}|^2$ and the muon neutrino $|U_{\mu N}|^2$. The sensitivities of the expected PIONEER R_π measurement, see equation (3.5), are shown in the same plots by the dashed blue curves. We observe that the bump hunt at PIONEER will be able to improve the sensitivity to the squared mixing angles by more than an order of magnitude compared to PIENU. The improvement when using the expected R_π measurement is approximately a factor of four, see the upper limits quoted in equations (3.4) and (3.5). The see-saw target regions are also shown in the plots of figures 3.1 and 3.2. Interestingly, the PIONEER Phase I sensitivity of $\pi^+ \to e^+ N$ reaches into the shown see-saw target region.

Assuming systematic uncertainties can be kept under control, the sensitivities of the bump-hunt analyses scale, to good approximation, with the square root of the total number of events. As mentioned above, we consider 0.5×10^8 $\pi^+ \to e^+\nu_e$ events, corresponding approximately to the expected sample size during PIONEER Phase I. Improved sensitivity to $|U_{eN}|^2$ could be achieved in later phases of PIONEER with larger data sets. For the $\pi^+ \to \mu^+ N$ search, we have assumed 10^{10} $\pi^+ \to \mu^+\nu_\mu$ events. However, already during Phase I, up to 10^{12} $\pi^+ \to \mu^+\nu_\mu$ decays are expected. If a larger fraction of these events can be recorded and reconstructed, the sensitivity to $|U_{\mu N}|^2$ could be significantly improved already at Phase I.

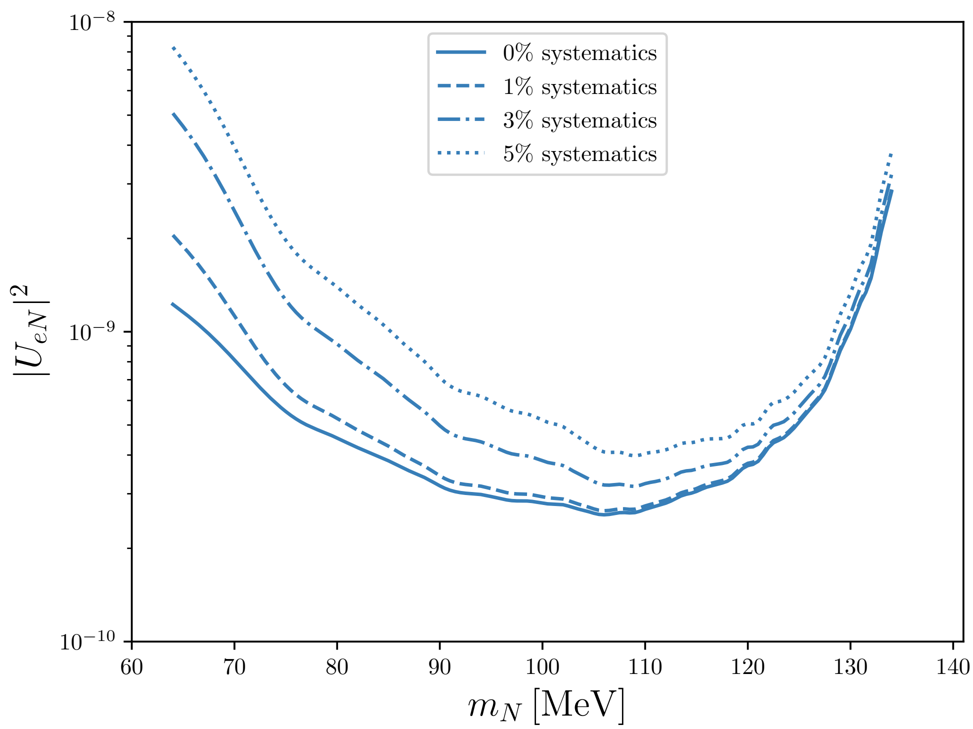

Figure 3.4: Sensitivity to the sterile neutrino mixing angle $|U_{eN}|^2$ at 90% C.L. as a function of the sterile neutrino mass m_N from a bump search at PIONEER. The solid line takes into account statistical uncertainties on the event numbers and allows for a floating normalization of background keeping the background shape fixed. The dashed, dot-dashed, and dotted lines include in addition uncorrelated systematic uncertainties of 1%/3%/5% in each positron energy bin.

As already observed by PIENU, and expected to be even more relevant for PIONEER, photonuclear effects in the calorimeter leave a characteristic imprint on the tail of the positron energy spectrum in the $\pi^+ \to e^+\nu_e$ decay, in particular a prominent feature at $E_e \simeq 60\,\text{MeV}$, see figure 3.3. Accurately modeling these effects is challenging. In our default bump-hunt analysis described above we therefore only include data below $E_e \leq 56\,\text{MeV}$. The background shape is taken from simulation and kept fixed in the fit, with only its overall normalization allowed to float. To assess the potential impact of imperfect background modeling at lower energies, we introduced additional uncorrelated systematic uncertainties of 1%, 3%, and 5% on the background event counts in each energy bin. This treatment is conservative, as it allows the background yield to fluctuate independently from bin to bin. The corresponding results are shown in figure 3.4, which compares the sensitivity to $|U_{eN}|^2$ from our default analysis with that obtained when including these additional uncertainties. We find that the impact of background shape uncertainties is most significant at lower sterile neutrino masses, corresponding to higher positron energies. In this case, the bound of the analysis that includes a 5% systematics can differ by almost an order of magnitude if compared to the bound with no systematics. This highlights the importance of a precise understanding of the line shape and, consequently, of a careful calibration of the PIONEER calorimeter response, which requires dedicated positron-beam measurements.

Finally, we mention that several other complementary probes can be sensitive to similar regions in sterile neutrino parameter space:

- Searches for kaon decays $K \to eN$ and $K \to \mu N$ in [99, 138] by the NA62 collaboration constrain mixing angles at the level of $|U_{eN}|^2 \lesssim 10^{-9}$ and $|U_{\mu N}|^2 \lesssim 10^{-8}$. However, limits are only reported for sterile neutrino masses above the pion mass, outside the region shown in figures 3.1 and 3.2.

- Big bang nucleosynthesis (BBN) gives relevant constraints both on electron and muon mixed sterile neutrinos and can exclude large regions of the parameter space shown in figures 3.1 and 3.2 [109, 181]. The BBN constraints can be avoided if the sterile neutrinos have a sufficiently short lifetime $\tau_N \lesssim 0.02\,\mathrm{s} - 0.1\,\mathrm{s}$. Additional invisible decay modes of the sterile neutrino are required to achieve such a sterile neutrino lifetime in the regions of parameter space that can be covered by PIONEER. Note that not all scenarios with additional invisible decay modes are cosmologically viable, as discussed in [158, 182].

- A single sterile Majorana neutrino that mixes with the electron neutrino is strongly constrained by neutrinoless double beta decay [183–186]. For masses around 100 MeV, the constraints are around $|U_{eN}|^2 \lesssim 10^{-9}$. The constraints can be much weaker if one has a pair of Majorana neutrinos with a small mass splitting that form a quasi-Dirac state. If the sterile neutrino is Dirac, there is no constraint from neutrinoless double beta decay.

- Interesting sensitivity to sterile neutrino parameter space can also be achieved at the DUNE near detector, that can look for the visible decays of sterile neutrinos that are produced from meson decays at the DUNE target station [27–29, 187]. In the relevant mass range of minimal models, DUNE can reach a sensitivity that is comparable to that of the PIENU search for $\pi^+ \to e^+ N$, and might be competitive with PIONEER's $\pi^+ \to \mu^+ N$ search. However, in contrast to PIENU and PIONEER, DUNE is not sensitive to short lived sterile neutrinos that avoid the BBN constraints via decays into invisible final states.

§ 3.3 Exotic Three-Body Decays of Pions

In this section we study the three-body pion decays $\pi^+ \to \ell^+\nu_\ell X$ as probes of light dark sector particles. We first introduce the benchmark models under consideration and describe their expected experimental signatures. We then present a recast of existing PIENU searches and a projection of the sensitivity achievable at PIONEER. Finally, we compare the resulting sensitivities with complementary probes, including anomalous magnetic moments of the electron and muon, monophoton searches at B factories, and beam dump experiments, and discuss the combined implications for the various model scenarios. In this section, we focus on bounds derived from laboratory-based experiments. We note that observations of SN1987A might also provide complementary constraints on the parameter space, generally at significantly smaller values of the coupling between the new dark particles and leptons (see e.g. [188, 189]).

3.3.1 Simplified dark sector models and their pion decay spectra

One can consider a variety of dark sector states X that can be produced in pion decays. In the following, we will focus on a class of models in which dark spin 0 or spin 1 states couple to the charged leptons. More specifically, we consider scalars, axion-like particles, as well as vectors with a variety of couplings. In contrast to the masses of axion-like particles and vectors, masses of scalars are susceptible to large radiative corrections. Light scalars require either a protection mechanism or are fine-tuned.

Scalars: The first example we consider is a light scalar s interacting with electrons or muons with the following Lagrangian

$$\mathcal{L}_{\text{scalar}} = g_s s(\bar{\ell}\ell) \ . \tag{3.15}$$

Invariance under the SM gauge group implies that this interaction arises from a dimension 5 operator and is therefore suppressed. A typical expectation would be $g_s \sim m_\ell/\Lambda$ with some high new physics scale Λ. The coupling could be as large as $g_s \sim v/\Lambda$, with the Higgs vacuum expectation value $v \simeq 246\,\text{GeV}$.

The corresponding double differential decay rate of $\pi^+ \to \ell^+\nu_\ell\, s$ as a function of the lepton energy E_ℓ and the scalar energy E_s is given by

$$\begin{aligned}\frac{d\text{BR}(\pi^+ \to \ell^+\nu_\ell\, s)}{\text{BR}(\pi^+ \to \ell^+\nu_\ell)} &= \frac{g_s^2}{4\pi^2}\frac{dE_\ell dE_s}{m_\pi^2} \\ &\times \frac{1}{(1-x_\ell)^2}\left[\frac{x_{\ell\nu}x_{\ell s}}{x_\ell(x_{\ell s}-x_\ell)} + \frac{x_{\ell s}(3+x_s-4x_{\ell s})-x_s+x_\ell}{(x_{\ell s}-x_\ell)^2}\right] \ ,\end{aligned} \tag{3.16}$$

where we used a convenient normalization to the branching ratio of the corresponding two-body decay $\pi^+ \to \ell^+\nu_\ell$. In the above expression we introduced the following shorthand notation for squared mass ratios

$$x_\ell = \frac{m_\ell^2}{m_\pi^2} \ , \quad x_s = \frac{m_s^2}{m_\pi^2} \ , \tag{3.17}$$

and ratios of squared invariant masses to the squared pion mass

$$x_{\ell\nu} = \frac{1}{m_\pi^2}(p_\ell+p_\nu)^2 = 1+\frac{m_s^2}{m_\pi^2}-\frac{2E_s}{m_\pi} \ , \quad x_{\ell s} = \frac{1}{m_\pi^2}(p_\ell+p_s)^2 = \frac{2E_s}{m_\pi}+\frac{2E_\ell}{m_\pi}-1 \ . \tag{3.18}$$

We note that the branching ratio in equation (3.16) contains a term proportional to $1/x_\ell = m_\pi^2/m_\ell^2$, signaling that the three-body decay into a scalar can overcome the helicity suppression of the two-body decay, $\pi^+ \to \ell^+ \nu_\ell$.

Axion-like particles: Next, we consider an axion-like particle (ALP) a with the following derivative interactions

$$\mathcal{L}_{\text{ALP}} = \frac{\partial_\alpha a}{2m_\ell}\Big(\bar{g}_a(\bar{\ell}\gamma^\alpha \ell) + g_a(\bar{\ell}\gamma^\alpha\gamma_5 \ell) + g_{a,\nu}(\bar{\nu}\gamma^\alpha P_L \nu)\Big)\,. \tag{3.19}$$

The generic expectation for the ALP couplings is $g_a, \bar{g}_a, g_{a,\nu} \sim m_\ell/f_a$, with the axion decay constant f_a. One also expects a relation due to $SU(2)_L$ invariance, $\bar{g}_a - g_a - g_{a,\nu} \sim v^2/\Lambda^2 \ll 1$, with a new physics scale Λ not necessarily related to f_a. We refer to an ALP with $\bar{g}_a - g_a - g_{a,\nu} \neq 0$ as a "weak violating ALP", i.e., an ALP with different couplings to SM neutrinos and left-handed charged leptons (see [127] for more details and possible UV completions of such an ALP).

The corresponding normalized differential branching ratio for the three-body decay $\pi^+ \to \ell^+ \nu_\ell\, a$ is

$$\begin{aligned}\frac{d\text{BR}(\pi^+ \to \ell^+\nu_\ell\, a)}{\text{BR}(\pi^+ \to \ell^+\nu_\ell)} &= \frac{1}{4\pi^2}\frac{dE_\ell dE_a}{m_\pi^2} \\ &\times \frac{1}{x_\ell(1-x_\ell)^2}\Bigg[g_a^2\,\frac{x_{\ell a}x_a(x_{\ell a}-1)}{(x_{\ell a}-x_\ell)^2} + g_a(\bar{g}_a - g_{a,\nu})\frac{x_{\ell a}(x_{\ell\nu}-x_a)+x_a-x_\ell}{x_{\ell a}-x_\ell} \\ &\qquad + (g_a - \bar{g}_a + g_{a,\nu})^2\frac{1}{4x_\ell}\big(x_{\ell a}(x_{\nu a}-x_\ell) - x_a + x_\ell\big)\Bigg]\,,\end{aligned} \tag{3.20}$$

where x_a, $x_{\ell\nu}$, and $x_{\ell a}$ are defined analogously to equations (3.17) and (3.18), and we have in addition

$$x_{\nu a} = \frac{1}{m_\pi^2}(p_\nu + p_a)^2 = 1 + \frac{m_\ell^2}{m_\pi^2} - \frac{2E_\ell}{m_\pi}\,. \tag{3.21}$$

The $\pi^+ \to \ell^+ \nu_\ell\, a$ branching ratio overcomes the helicity suppression if the ALP is weak violating (see the last term in equation (3.20)).

Vectors: We also consider a light vector V with the following set of couplings to the SM leptons

$$\mathcal{L}_{\text{vector}} = V_\alpha \Big(g_V (\bar{\ell}\gamma^\alpha \ell) + g_A (\bar{\ell}\gamma^\alpha \gamma_5 \ell) + (g_V - g_A)(\bar{\nu}\gamma^\alpha P_L \nu) \Big) + \frac{V_{\alpha\beta}}{\Lambda} \Big(g_T (\bar{\ell}\sigma^{\alpha\beta}\ell) + g_{T5} (\bar{\ell}\sigma^{\alpha\beta} i\gamma_5 \ell) \Big) \,. \quad (3.22)$$

The vector coupling g_V could for example arise if the vector V is the gauge boson of an abelian gauge symmetry that includes electron or muon number. The axial vector coupling g_A (or combinations of g_A and g_V) can for example be engineered in the effective Z' construction from [190]. Note that we have fixed the coupling of neutrinos using $SU(2)_L$ gauge invariance. Exploring "weak violating vectors" is beyond the scope of this work.

The second line of the above Langrangian shows dipole interactions with the field strength tensor $V_{\alpha\beta} = \partial_\alpha V_\beta - \partial_\beta V_\alpha$. It is plausible that such dipole interactions are loop suppressed. Due to $SU(2)_L$ invariance, we expect $g_T, g_{T5} \lesssim v/(16\pi^2\Lambda)$, with a typical expectation of $g_T, g_{T5} \sim m_\ell/(16\pi^2\Lambda)$. The simultaneous presence of both g_T and g_{T5} leads to 1-loop contributions of electric dipole moments, which are very strongly constrained. In our numerical analysis, we will therefore consider only one of them at a time.

Taking into account only the vector and axial vector couplings, we find the following expression for the $\pi^+ \to \ell^+ \nu_\ell V$ differential branching ratio

$$\frac{d\mathrm{BR}(\pi^+ \to \ell^+ \nu_\ell V)}{\mathrm{BR}(\pi^+ \to \ell^+ \nu_\ell)} = \frac{1}{2\pi^2} \frac{dE_\ell dE_V}{m_\pi^2} \frac{1}{(1-x_\ell)^2} \Bigg[g_V^2 \bigg(2 + \frac{x_{\ell V} - 2 + x_\ell}{x_{\nu V}} - \frac{x_V(1-x_\ell)}{x_{\nu V}^2} + \frac{2(1-x_\ell)^2 + 2x_\ell x_V}{(x_{\ell V} - x_\ell)x_{\nu V}} - \frac{2(1-x_\ell) - x_{\nu V}}{x_{\ell V} - x_\ell} - \frac{(2x_\ell + x_V)(1-x_\ell)}{(x_{\ell V} - x_\ell)^2} \bigg) + g_A^2 \bigg(\frac{x_{\ell V} + 2 - 3x_\ell}{x_{\nu V}} - \frac{2(x_{\ell V} - 1 + x_{\nu V})}{x_V} - \frac{(1-x_\ell)x_V}{x_{\nu V}^2} + \frac{2x_\ell}{x_{\ell V} - x_\ell} \Big(\frac{x_V}{x_{\nu V}} - \frac{x_{\nu V}}{x_V} \Big) - \frac{2(1-x_\ell)^2}{(x_{\ell V} - x_\ell)x_{\nu V}} + \frac{2 + x_{\nu V} - 6x_\ell}{x_{\ell V} - x_\ell} - \frac{(1-x_\ell)(x_V - 4x_\ell)}{(x_{\ell V} - x_\ell)^2} \bigg) - 2g_V g_A \bigg(\frac{x_{\ell V} - x_\ell}{x_{\nu V}} - \frac{x_V(1-x_\ell)}{x_{\nu V}^2} - \frac{x_{\nu V} + 2x_\ell}{x_{\ell V} - x_\ell} + \frac{2x_\ell x_V}{(x_{\ell V} - x_\ell)x_{\nu V}} + \frac{x_V(1+x_\ell) - 2x_\ell x_{\nu V}}{(x_{\ell V} - x_\ell)^2} \bigg) \Bigg] . \quad (3.23)$$

Interestingly, this three-body decay features the same helicity suppression as the two-body decay $\pi^+ \to \ell^+ \nu_\ell$.

If instead we only consider the dipole interactions, we find

$$\frac{d\mathrm{BR}(\pi^+ \to \ell^+ \nu_\ell V)}{\mathrm{BR}(\pi^+ \to \ell^+ \nu_\ell)} = \frac{1}{\pi^2} \frac{dE_\ell dE_V}{m_\pi^2} \frac{m_\pi^2}{\Lambda^2} \frac{1}{(1-x_\ell)^2} \times \Bigg[g_T^2 \bigg(2 + 8x_V - 2x_{\ell V} - \frac{x_V(8 - 16x_\ell - 2x_V + x_{\nu V})}{x_{\ell V} - x_\ell} - \frac{(1-x_\ell)x_V(8x_\ell + x_V)}{(x_{\ell V} - x_\ell)^2} \bigg) + g_{T5}^2 \bigg(2 - 4x_V - 2x_{\ell V} + \frac{x_V(4 - 8x_\ell + 2x_V - x_{\nu V})}{x_{\ell V} - x_\ell} + \frac{(1-x_\ell)x_V(4x_\ell - x_V)}{(x_{\ell V} - x_\ell)^2} \bigg) + (g_T^2 + g_{T5}^2) \frac{1}{x_\ell} \Big(x_{\ell V}(2x_{\nu V} - x_V) - x_V(1 + x_{\nu V} - x_V) \Big) \Bigg] . \quad (3.24)$$

Both the g_T and the g_{T5} couplings provide a term enhanced by $\sim m_\pi^2/m_\ell^2$.

In all the cases discussed in this section, the phase space boundaries of the lepton energy E_ℓ and the energy of the dark sector particle E_X with $X = s, a, V$, are given by the expressions

$$E_- < E_X < E_+ \,, \quad m_\ell < E_\ell < \frac{m_\pi}{2}\left(1 - x_X + x_\ell\right) \,, \tag{3.25}$$

$$E_\pm = \frac{m_\pi}{2}\left[1 - \frac{E_\ell}{m_\pi} \pm \sqrt{\frac{E_\ell^2}{m_\pi^2} - x_\ell + \frac{x_X\left(1 - E_\ell/m_\pi \mp \sqrt{E_\ell^2/m_\pi^2 - x_\ell}\right)}{1 + x_\ell - 2E_\ell/m_\pi}}\right] \,. \tag{3.26}$$

In figures 3.5 and 3.6 we show example spectra for the decays $\pi^+ \to e^+\nu_e X$ and $\pi^+ \to \mu^+\nu_\mu X$, for scalars, ALPs, and vectors with representative choices of couplings and various choices of X masses. Each scenario gives its own characteristic signal shape.

3.3.2 Experimental signatures at pion decay experiments

There are several types of experimental signatures of the $\pi^+ \to \ell^+\nu_\ell X$ decays, depending on the lifetime and the decay modes of X. The most relevant are

$$\pi^+ \to \ell^+\nu_\ell + \text{invisible} \,, \quad \pi^+ \to \ell^+\nu_\ell + \gamma\gamma \,, \quad \pi^+ \to \ell^+\nu_\ell + e^+e^- \,. \tag{3.27}$$

The photons and the electron-positron pair could be prompt or displaced. As already mentioned in section 3.1, in this work we focus on the signature $\pi^+ \to \ell^+\nu_\ell + \text{invisible}$ which resembles most closely the SM decay $\pi^+ \to \ell^+\nu_\ell$ but features a non-standard continuous lepton energy spectrum. The shape of the lepton energy spectrum can be predicted in a given model and it depends on the mass of X, its spin, and the type of coupling to the SM particles. Searches for a non-standard component in the lepton energy spectrum from $\pi^+ \to \ell^+\nu_\ell X$ decays are therefore necessarily model-dependent.

Without directly using any kinematic information, we can still bound the branching ratios into light exotic states X, by making use of the R_π measurement. Assuming that

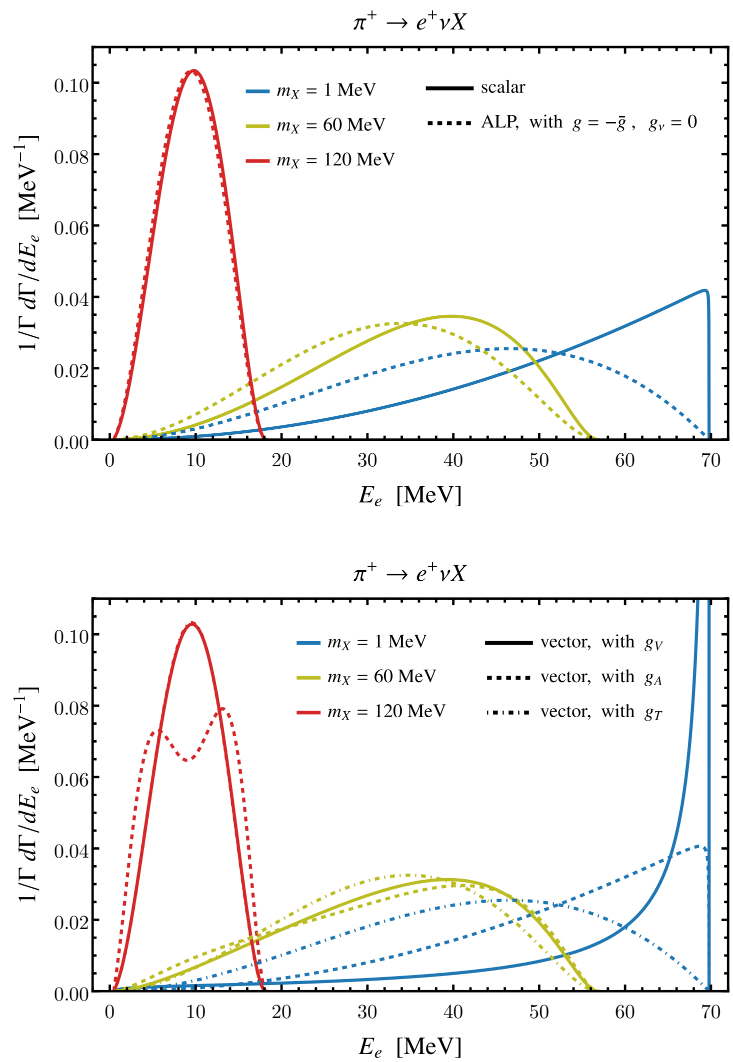

Figure 3.5: Representative examples of normalized signal shapes of the $\pi^+ \to e^+\nu_e X$ decay as a function of the positron energy. Top: scalar and ALP models; bottom: various vector models. For a mass of $m_X = 120\,\mathrm{MeV}$, the spectra of the vector models with g_V coupling and g_T coupling are almost indistinguishable. The spectra of the vector model with g_{T5} coupling look virtually identical to the one with g_T coupling and are therefore not shown.

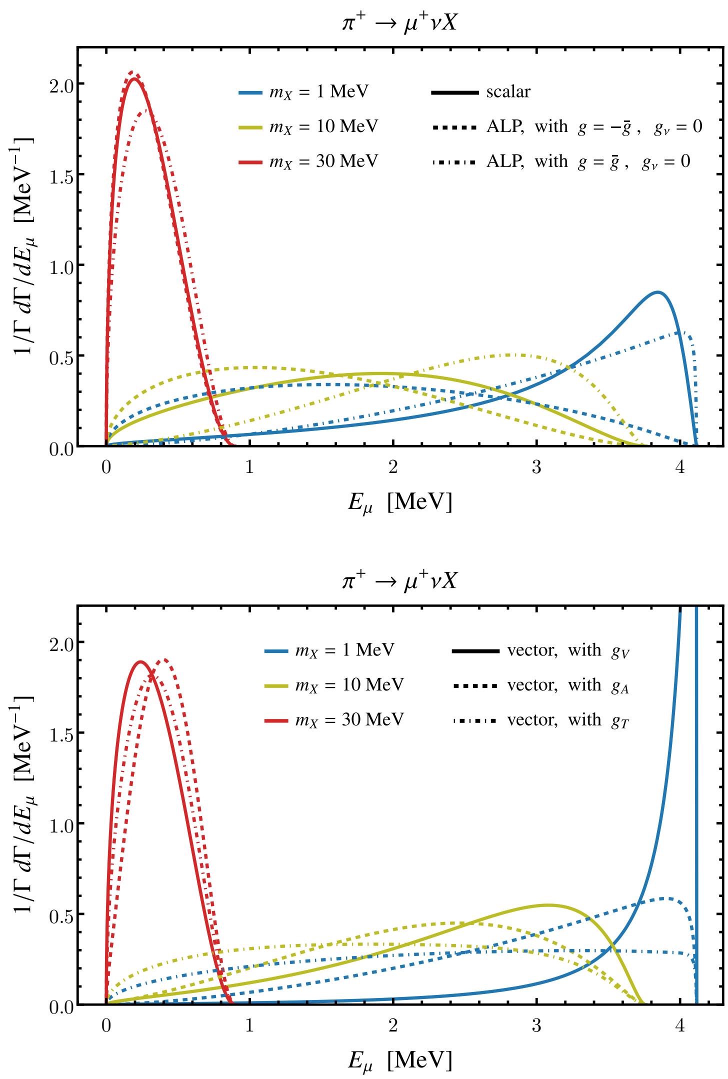

Figure 3.6: Representative examples of normalized signal shapes of the $\pi^+ \to \mu^+\nu_\mu X$ decay as a function of the muon kinetic energy. Top: scalar and ALP models; bottom: various vector models.

the only new physics effects are the $\pi^+ \to e^+\nu_e\,X$ or $\pi^+ \to \mu^+\nu_\mu\,X$ decays, the leading new physics correction to R_π is given by

$$\frac{R_\pi}{R_\pi^{\mathrm{SM}}} \simeq 1 + r_e \times \frac{\mathrm{BR}(\pi^+ \to e^+\nu_e\,X)}{\mathrm{BR}(\pi^+ \to e^+\nu_e)_{\mathrm{SM}}} - r_\mu \times \frac{\mathrm{BR}(\pi^+ \to \mu^+\nu_\mu\,X)}{\mathrm{BR}(\pi^+ \to \mu^+\nu_\mu)_{\mathrm{SM}}}\,, \tag{3.28}$$

where r_e and r_μ are efficiency ratios due to the different kinematics of the three-body decays into the dark sector state as compared to the two-body SM decays. We will assume $r_e = r_\mu = 1$.

3.3.3 Recast of existing pion decay constraints

The PIENU experiment has performed a search for the three-body decay $\pi^+ \to \ell^+\nu_\ell X$ with an invisible dark sector particle X, as reported in [137]. The signal shape used in their analysis is based on the model from [116], and closely resembles the contribution from a scalar or a weak-preserving axion-like particle, as shown in equations (3.16) and (3.20). PIENU set upper limits on the branching ratio $\mathrm{BR}(\pi^+ \to e^+\nu_e X)$ at the level of 10^{-7} to 10^{-6} for dark sector masses $m_X \lesssim 120\,\mathrm{MeV}$. For decays involving muons, $\pi^+ \to \mu^+\nu_\mu X$, the corresponding branching ratio limits reach approximately 10^{-5} for masses $m_X \lesssim 35\,\mathrm{MeV}$.

As described in detail in appendix C we were able to successfully reproduce the PIENU limits on the $\mathrm{BR}(\pi^+ \to e^+\nu_e X)$ and $\mathrm{BR}(\pi^+ \to \mu^+\nu_\mu X)$ branching ratios using the signal model from [116]. As a next step, we recast the searches in the context of all the alternative models considered in our study. We substitute the signal shape from [116] with the predicted spectra of other models from equations (3.16), (3.20), (3.23), and (3.24), and apply the same analysis pipeline. We obtain the branching ratio limits as a function of the X mass as shown in figures 3.7 and 3.8.

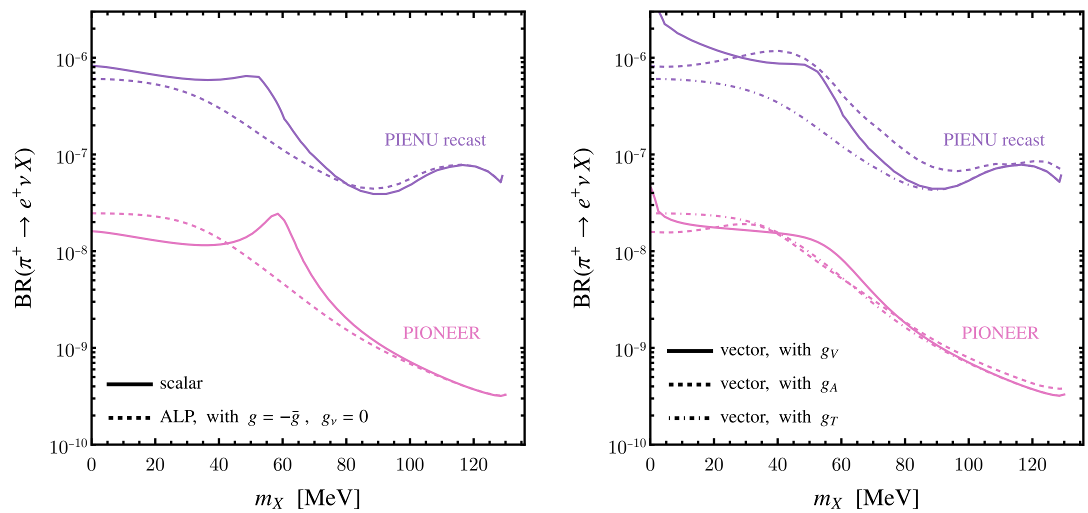

Figure 3.7: Recasted PIENU limits and expected PIONEER sensitivities to the $\pi^+ \to e^+\nu_e X$ branching ratio for various signal models as described in section 3.3.1. The PIONEER sensitivity assumes 0.5×10^8 recorded $\pi^+ \to e^+\nu_e$ events.

The limits on the $\pi^+ \to e^+\nu_e X$ branching ratio range from approximately 10^{-6} at low masses $m_X \ll m_\pi$ to approximately 10^{-7} for the largest masses $m_X \simeq m_\pi$. Similarly, we find limits on the $\pi^+ \to \mu^+\nu_\mu X$ branching ratio that vary from $\sim 10^{-4}$ to $\sim 10^{-5}$ depending on the mass of X. The precise shape of the exclusion curves are model dependent. In the case of the decay $\pi^+ \to e^+\nu_e X$ we do not show the weak preserving axion-like particle and the vector with g_{T5} coupling. The corresponding curves are almost identical to the scalar and vector with g_T coupling, respectively. The relative differences in the respective differential decay rates are indeed proportional to m_e^2/m_π^2 and therefore tiny. In the case of the $\pi^+ \to \mu^+\nu_\mu X$ decay, the curves of all models differ visibly.

As we discuss in section 3.3.6 below, these branching ratio limits can be translated into constraints on the underlying coupling constants of the respective models, providing

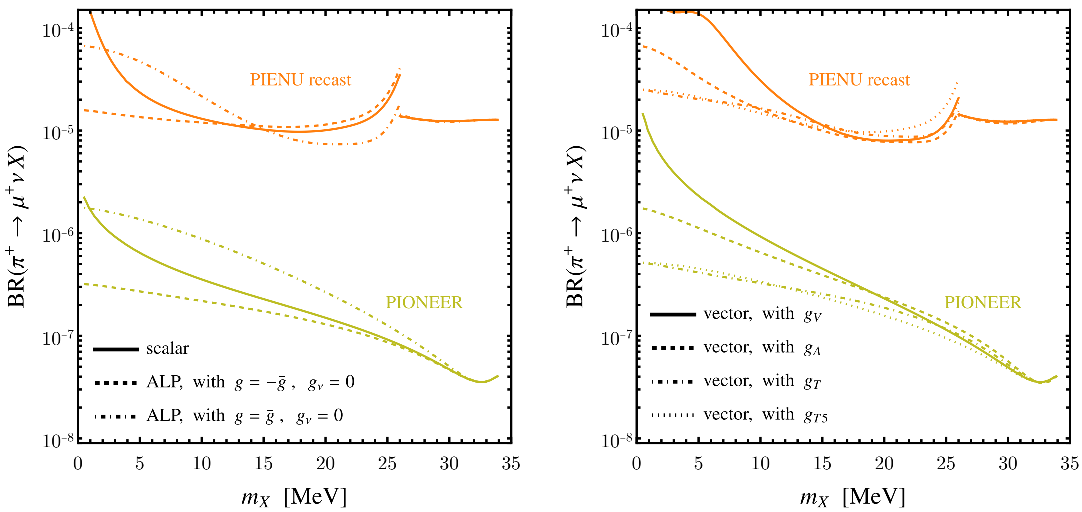

Figure 3.8: Recasted PIENU limits and expected PIONEER sensitivities to the $\pi^+ \to \mu^+\nu_\mu X$ branching ratio for various signal models as described in section 3.3.1. The PIONEER sensitivity assumes 10^{10} recorded $\pi^+ \to \mu^+\nu_\mu$ events.

a model-specific interpretation of the experimental sensitivity.

3.3.4 Expected sensitivities at PIONEER

Having validated our analysis framework through a successful reproduction of the PIENU search results for $\pi^+ \to \ell^+\nu_\ell X$ decays, we now turn to estimating the future sensitivity of PIONEER. To assess PIONEER's discovery potential, we analyze the same set of benchmark models and use simulated PIONEER data and an analysis pipeline closely mirroring that employed for the PIENU recast. In the following, we discuss the projected sensitivities for the $\pi^+ \to e^+\nu_e X$ and $\pi^+ \to \mu^+\nu_\mu X$ channels.

The $\pi^+ \to e^+\nu_e X$ search: Our sensitivity study of $\pi^+ \to e^+\nu_e X$ is based on the same simulated dataset and analysis framework as the sterile neutrino search described in section 3.2.4. The simulation includes 0.5×10^8 reconstructed $\pi^+ \to e^+\nu_e$ events and a subdominant component of approximately 8.8% from muon decay in flight. Other backgrounds are negligible.

The main difference from the sterile neutrino analysis lies in the modeling of the signal. Instead of a narrow peak in the positron energy spectrum, the decay $\pi^+ \to e^+\nu_e X$ leads to a continuous positron energy distribution that depends on the mass, spin, and couplings of the dark particle X. For each model under consideration, the theoretical positron spectrum is convolved with a Gaussian resolution function, using the same energy-dependent width as in the sterile neutrino study.

Sensitivities are obtained from a binned χ^2 fit of the simulated positron energy spectrum using 0.5 MeV bins in the range $(4-56)$ MeV. The normalizations of the $\pi^+ \to e^+\nu_e X$ signal, as well as the $\pi^+ \to e^+\nu_e$ and muon decay in flight components are allowed to float while their shapes are kept fixed. Assuming no signal, we derive projected 90% C.L. upper limits on BR$(\pi^+ \to e^+\nu_e X)$ for all the benchmark models discussed above.

The $\boldsymbol{\pi^+ \to \mu^+\nu_\mu X}$ search: The sensitivity to the decay $\pi^+ \to \mu^+\nu_\mu X$ is evaluated using simulated $\pi^+ \to \mu^+\nu_\mu$ events corresponding to 10^{10} reconstructed decays, with negligible background contributions. The dark sector signal is a three-body spectrum convolved with a Gaussian resolution function, analogous to the treatment of $\pi^+ \to e^+\nu_e X$. The width of the Gaussian is fixed to 0.2 MeV. A binned χ^2 fit is performed over the muon kinetic energy range 0–3.2 MeV with 0.06 MeV bins. The normalizations of the signal and the standard $\pi^+ \to \mu^+\nu_\mu$ component are allowed to float, while their shapes are fixed. Under the assumption of no signal, we derive projected 90% C.L.

upper limits on BR$(\pi^+ \to \mu^+\nu_\mu X)$ for the benchmark models introduced above.

The PIONEER sensitivities to the $\pi^+ \to e^+\nu_e X$ and $\pi^+ \to \mu^+\nu_\mu X$ branching ratios are shown in figures 3.7 and 3.8 and compared to our recast of the existing PIENU limits. We find that PIONEER will be able to improve over PIENU by one to two orders of magnitude. For X masses around $100\,$MeV, $\pi^+ \to e^+\nu_e X$ branching ratios below 10^{-9} are in reach. In the case of $\pi^+ \to \mu^+\nu_\mu X$, PIONEER will be able to probe branching ratios as low as few $\times\ 10^{-8}$.

To a good approximation, the sensitivity to the $\pi^+ \to \ell^+\nu_\ell X$ branching ratios scales with the square root of events. Therefore, further improved sensitivity to the branching ratios can be expected in later phases of PIONEER with larger data sets. Already during Phase I, a better sensitivity to $\pi^+ \to \mu^+\nu_\mu X$ could be achieved if more than 10^{10} $\pi^+ \to \mu^+\nu$ decays can be recorded and reconstructed.

3.3.5 Comparison to other dark sector searches

The dark sector scenarios introduced in section 3.3.1 are subject to various additional constraints beyond those that can be found from pion decays. The most relevant constraints in the mass range that we are considering are the anomalous magnetic moments of the electron and the muon, mono-photon searches at the B factories, and searches at electron beam dumps. Also the rare kaon decays $K \to \pi +$ invisible give constraints, which are, however, to some extent model dependent.

Anomalous magnetic moments

On the experimental side, the anomalous magnetic moments of the electron and the muon are measured with extremely high precision. In the case of the electron the most precise determination is given in [191]. For the muon we use the world average provided by the Fermilab muon g-2 collaboration [192]. Comparing with SM predictions [193], one finds

$$\Delta a_e = a_e^{\rm exp} - a_e^{\rm SM} = \begin{cases} (-10.0 \pm 2.6) \times 10^{-13}\,, & \text{for } \alpha_{\rm em} \text{ from Cs [194]}\,, \\ (+3.5 \pm 1.6) \times 10^{-13}\,, & \text{for } \alpha_{\rm em} \text{ from Rb [195]}\,, \end{cases} \tag{3.29}$$

$$\Delta a_\mu = a_\mu^{\rm exp} - a_\mu^{\rm SM} = (38.5 \pm 63.7) \times 10^{-11} \text{ [192, 193]}\,. \tag{3.30}$$

For the Standard Model prediction of the electron $g-2$, the latest updates on the 5-loop QED contribution are included [196–198] and we quote two values corresponding to the two most precise determination of the fine structure constant from Rubidium atom recoil measurements [195] and Cesium atom recoil measurements [194], respectively. For the muon, we take the latest SM prediction by the Muon g-2 Theory Initiative [193] which is based on lattice determinations of the hadronic vacuum polarization contribution and which updates the result in their first white paper [199].

In the case of the electron, in view of the large discrepancy of the fine structure constant determinations using Rubidium and Cesium, we impose the 2σ upper bound from Rb and the 2σ lower bound from Cs. For the muon we use the 2σ bound. In our numerical analysis, possible new physics contributions are thus bounded by

$$-15.2\times 10^{-13} < \Delta a_e < 6.7\times 10^{-13}\,, \quad -0.89\times 10^{-9} < \Delta a_\mu < 1.66\times 10^{-9}\,. \tag{3.31}$$

Next, we discuss the new physics contributions to the anomalous magnetic moments in the scenarios that we consider. We calculated the contributions of the scalar, the ALP, and the vector and find the following expressions

$$\Delta a_\ell^s = \frac{g_s^2}{16\pi^2}\Big(3 - 2z_s + z_s(z_s-3)\log z_s - 2(z_s-1)(z_s-4)\Phi(z_s)\Big)\,, \tag{3.32}$$

$$\Delta a_\ell^a = -\frac{g_a^2}{16\pi^2}\Big(1 + 2z_a - z_a(z_a-1)\log z_a + 2z_a(z_a-3)\Phi(z_a)\Big)\,, \tag{3.33}$$

$$\begin{aligned}\Delta a_\ell^V = {}& \frac{g_V^2}{8\pi^2}\Big(1 - 2z_V + z_V(z_V-2)\log z_V - 2(2+(z_V-4)z_V)\Phi(z_V)\Big) \\ & -\frac{g_A^2}{8\pi^2}\left(\frac{2}{z_V} - 5 + 2z_V - (2+z_V(z_V-4))\log z_V \right. \\ & \left. +2(z_V-2)(z_V-4)\Phi(z_V)\right),\end{aligned} \tag{3.34}$$

with the function

$$\Phi(z) = \begin{cases} \frac{1}{\sqrt{1-4/z}}\ \log\left(\frac{\sqrt{z}+\sqrt{z-4}}{2}\right)\,, & \text{for } z > 4\,, \\ \frac{1}{\sqrt{4/z-1}}\ \arccos\left(\frac{\sqrt{z}}{2}\right)\,, & \text{for } z < 4\,, \end{cases} \tag{3.35}$$

and the mass ratios $z_X = m_X^2/m_\ell^2$ with $X = s, a, V$. The term $\propto 1/z_V$ is unique to the axial-vector scenario.

In the above expression for the vector contributions we have only considered vector and axial vector couplings g_V, g_A. If the vector boson instead couples to leptons through the dipole interactions g_T, g_{T5} we find

$$\begin{aligned}\Delta a_\ell^V &= \frac{g_T^2}{4\pi^2}\frac{m_\ell^2}{\Lambda^2}\Bigg(6\log\left(\frac{\Lambda^2}{m_\ell^2}\right) + z_V(9+2z_V) + z_V(8-z_V(z_V+3))\log z_V \\ &\qquad -2z_V(16-(z_V+1)z_V)\Phi(z_V)\Bigg) \\ &\quad -\frac{g_{T5}^2}{4\pi^2}\frac{m_\ell^2}{\Lambda^2}\Bigg(2\log\left(\frac{\Lambda^2}{m_\ell^2}\right) - z_V(3-2z_V) + z_V^2(3-z_V)\log z_V \\ &\qquad +2z_V(z_V-1)(z_V-4)\Phi(z_V)\Bigg) \quad (3.36)\end{aligned}$$

To the best of our knowledge, this is a new result. The contributions are logarithmically divergent and we have set the renormalization scale equal to the new physics scale Λ that parameterizes the size of the dipole coupling. The large logarithms could be re-summed using rernormalization group equations, but this is beyond the scope of our work. We note that in the presence of both g_T and g_{T5} the vector contributes to electric dipole moments of leptons which are strongly constrained. We do not consider such a scenario in the following.

For electrons there is a large region of parameter space for which it is a good approximation to expand in $m_e^2 \ll m_X^2$. Doing so, one finds

$$\Delta a_e^s \simeq \frac{1}{8\pi^2}\frac{m_e^2}{m_s^2}\, g_s^2\left[\log\left(\frac{m_s^2}{m_e^2}\right) - \frac{7}{6}\right], \quad (3.37)$$

$$\Delta a_e^a \simeq -\frac{1}{8\pi^2}\frac{m_e^2}{m_a^2}\, g_a^2\left[\log\left(\frac{m_a^2}{m_e^2}\right) - \frac{11}{6}\right], \quad (3.38)$$

$$\Delta a_e^V \simeq \frac{1}{12\pi^2}\frac{m_e^2}{m_V^2}\left(g_V^2 - 5g_A^2\right), \quad (3.39)$$

$$\Delta a_e^V \simeq \frac{1}{2\pi^2}\frac{m_e^2}{\Lambda^2}\left[\left(3g_T^2 - g_{T5}^2\right)\log\left(\frac{\Lambda^2}{m_V^2}\right) + \frac{1}{6}\left(g_T^2 - 7g_{T5}^2\right)\right]. \quad (3.40)$$

Mono-photon searches

Mono-photon searches at the BaBar experiment [200], $e^+e^- \to \gamma\, X$, place relevant constraints on each of the considered models that feature couplings to electrons. Analogous searches at the Belle II experiment are expected to further improve the reach [201].

The BaBar analysis [200] targets specifically invisibly decaying dark photons and sets limits on the kinetic mixing parameter ϵ as a function of the dark photon mass. In the following, we describe how we recast the results for our scenarios. We first determine the differential cross sections of $e^+e^- \to \gamma X$, with $X = s, a, V$ in the models we consider. We find

$$\frac{d\sigma(e^+e^- \to \gamma s)}{d\cos\theta} = \frac{e^2 g_s^2}{8\pi s}\frac{1}{\sin^2\theta}\left(1+\frac{m_s^4}{s^2}\right)\left(1-\frac{m_s^2}{s}\right)^{-2}, \tag{3.41}$$

$$\frac{d\sigma(e^+e^- \to \gamma a)}{d\cos\theta} = \frac{e^2 g_a^2}{8\pi s}\frac{1}{\sin^2\theta}\left(1+\frac{m_a^4}{s^2}\right)\left(1-\frac{m_a^2}{s}\right)^{-2}, \tag{3.42}$$

$$\frac{d\sigma(e^+e^- \to \gamma V)}{d\cos\theta} = \frac{e^2\,(g_V^2+g_A^2)}{8\pi s}\left[\frac{2}{\sin^2\theta}\left(1+\frac{m_V^4}{s^2}\right)\left(1-\frac{m_V^2}{s}\right)^{-2} - 1\right] \tag{3.43}$$

$$\begin{aligned}\frac{d\sigma(e^+e^- \to \gamma V)}{d\cos\theta} &= \frac{e^2\,(g_T^2+g_{T5}^2)}{2\pi\Lambda^2}\\ &\quad\times\left[\frac{1}{\sin^2\theta}\frac{m_V^2}{s}\left(1+\frac{m_V^4}{s^2}\right)\left(1-\frac{m_V^2}{s}\right)^{-2} + 1 - \frac{m_V^2}{s}\right]\end{aligned} \tag{3.44}$$

In the above expressions, $\sqrt{s} \simeq 10.58\,\text{GeV}$ is the center of mass energy at BaBar and θ is the angle between the emitted photon and the beamline in the center of mass frame. We have taken the limit $m_e \to 0$ which is an excellent approximation.

As the BaBar analysis imposes the cut $-0.4 < \cos\theta < 0.6$, we require that the cross sections of our models in the same region of $\cos\theta$ does not exceed the corresponding cross section of a dark photon

$$\int_{-0.4}^{0.6} d\cos\theta \, \frac{d\sigma(e^+e^- \to \gamma X)}{d\cos\theta} < \int_{-0.4}^{0.6} d\cos\theta \, \frac{d\sigma(e^+e^- \to \gamma A')}{d\cos\theta} \,, \tag{3.45}$$

with the differential dark photon cross section given by

$$\frac{d\sigma(e^+e^- \to \gamma A')}{d\cos\theta} = \frac{\epsilon^2 e^4}{8\pi s}\left[\frac{2}{\sin^2\theta}\left(1+\frac{m_{A'}^4}{s^2}\right)\left(1-\frac{m_{A'}^2}{s}\right)^{-2} - 1\right] . \tag{3.46}$$

Beam dump constraints

The NA64 experiment at CERN places constraints on the production of invisible particles coupled to electrons [202]. More specifically, NA64 is an electron/positron beam dump experiment utilizing 100 GeV beams scattering with an active target and looking for missing energy. In their latest analysis based on 9.37×10^{11} electrons on target [203], NA64 set the most stringent constraint on the parameter space of an invisible dark photon in the $(10^{-3} - 0.35)$ GeV mass range*. Analogous searches at the future LDMX missing-momentum experiment are expected to further improve these bounds [205].

Signal efficiencies and yields are expected to change by an $\mathcal{O}(1)$ amount for other light dark sector particles. Instead of recasting the NA64 dark photon search, we use the NA64 results from [206] which provide directly limits on the couplings of scalars, pseudo-scalars, vectors, and axial-vectors in the mass range from 1 MeV to 1 GeV.

Note that the results in [206] are based on approximately one third of the statistics used in [203]. A recast of the full current data set might improve the limits on the

*The NA64 collaboration also sets interesting bounds on the dark photon parameter space using a positron beam run [204], which exploits dark photon production via resonant annihilation of positrons with atomic electrons in the target nuclei. However, these bounds are not relevant to our analysis, as they are only competitive at higher dark particle masses, $m \gtrsim 0.1$ GeV.

couplings by a factor of approximately $\sqrt{3} \sim 1.7$.

Kaon decays

Similar to the $\pi^+ \to \ell^+ \nu_\ell X$ decays, also the analogous kaon decays $K^+ \to \ell^+ \nu_\ell X$ are sensitive probes of light dark sectors. The NA62 experiment has searched for the decay $K^+ \to \mu^+ \nu_\mu X$, with an invisible X [138]. In the mass range that overlaps with the pion decays, $m_X \lesssim m_\pi - m_\mu$, NA62 has set limits on the $K^+ \to \mu^+ \nu_\mu X$ branching ratio in the range between 2×10^{-6} and 2×10^{-5} for a scalar and a vector model.

Also flavor-changing neutral current decays of kaons, the $K \to \pi X$ decay in particular, potentially give relevant constraints in the models we are considering. On the experimental side, the strongest constraints on the $K^+ \to \pi^+ X$ rate have been obtained at BNL-E949 and NA62 [207–210] (see also [211] for a recast of the NA62 measurement of the $K^+ \to \pi^+ \nu \bar{\nu}$ branching ratio as a constraint on ALPs). Even if our dark sector particles couple at tree level only to charged leptons, couplings to other SM particles are generated at the loop level.

This has been extensively explored in the context of axion-like particles, see e.g. [121, 123]. In particular, ALP couplings to left-handed (right-handed) charged leptons induce flavor changing couplings to strange and down quarks at the 2-loop (3-loop) level, which are strongly constrained by rare kaon decays. However, some of these loop effects depend logarithmically on an unknown UV cut-off scale. Moreover, additional contributions to flavor changing couplings can arise from loops involving couplings to quarks or gauge bosons. Models might even contain tree-level flavor changing couplings. Without making specific assumptions about many couplings, it is thus not possible to model independently assess the constraints from $K \to \pi X$.

As an example, we note that in the leptophilic ALP models considered in [127], the searches for $K^+ \to \pi^+ X$ constrain the couplings of the ALP to electrons at the level of $g_a \sim 10^{-4}$ to 10^{-6}.

3.3.6 Sensitivity to dark sector model parameters

In figures 3.9 and 3.10 we compare the pion decay sensitivity to the sensitivities from the other dark sector searches, the anomalous magnetic moments, mono-photon searches, and beam dump searches in particular.

We show the results for a representative selection of models. In both figures, the top left panels correspond to a scalar with coupling g_s as defined in equation (3.15). The top right panels show axion-like particles with weak violating couplings. In the notation of equation (3.19) this corresponds to $g_a = -\bar{g}_a = g_{ee}$ or $g_{\mu\mu},\ g_{a,\nu} = 0$. The bottom left and bottom right panels correspond to spin-1 bosons with vector couplings g_V or axial-vector couplings g_A as defined in equation (3.22). We do not show a weak preserving ALP, as the sensitivities from pion decays turn out to be very similar to the scalar case. We also do not show spin-1 bosons with dipole interactions. Setting the scale Λ that enters the couplings in (3.22) to 1 TeV, we find constraints from pion decays are in the highly non-perturbative regime $g_T, g_{T5} \gg 4\pi$.

Across all the models we consider, the anomalous magnetic moment of the electron and the BaBar mono-photon searches put limits on the coupling of the dark sector particles to electrons at the level of approximately 10^{-3} to 10^{-4}. The beam dump constraints from NA64 tend to be consistently stronger and reach down to $\sim 10^{-5}$ for low masses $m_X \ll m_\pi$. The anomalous magnetic moment of the muon constrains the scalar, axion-like particle, and vector couplings to muons at around 10^{-3}. In the case of

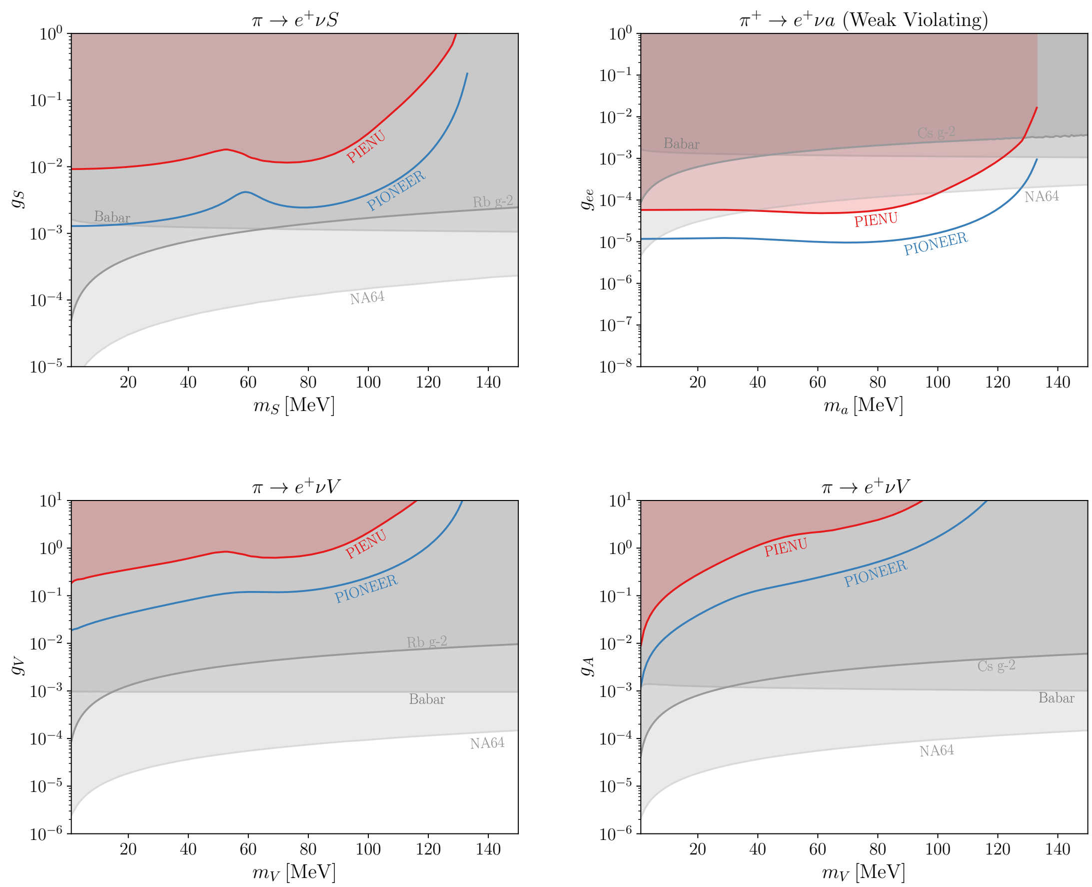

Figure 3.9: Limits on the couplings of dark sector particles to electrons as a function of their mass. **Top left**: scalar; **top right**: axion-like particle with weak violating couplings; **bottom left**: spin-1 with vector coupling; **bottom right**: spin-1 with axial-vector coupling. Existing constraints from the anomalous magnetic moment of the electron, from mono-photon searches at BaBar and from beam dump searches at NA64 are shown in gray. The pion decay constraints from PIENU are shown in red and the PIONEER sensitivity (assuming 0.5×10^8 recorded $\pi^+ \to e^+\nu_e$ events) in blue.

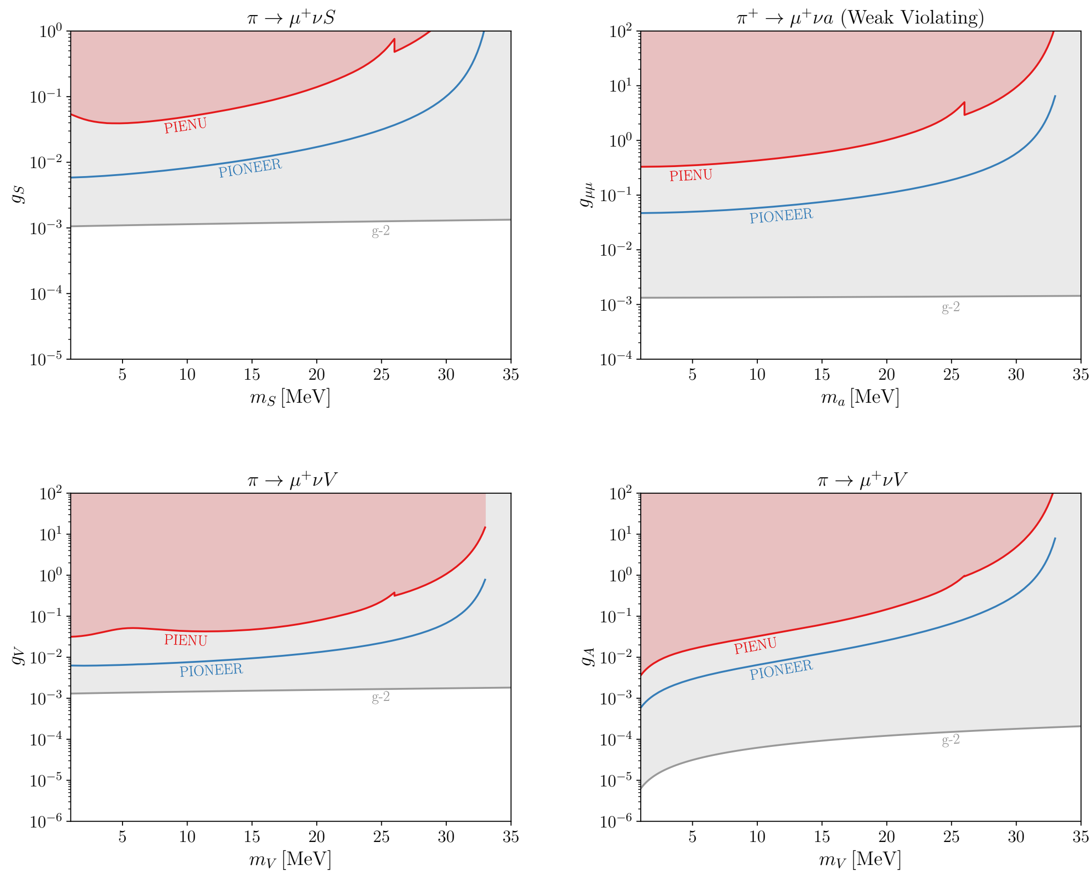

Figure 3.10:]
Limits on the couplings of dark sector particles to muons as a function of their mass. **Top left**: scalar; **top right**: axion-like particle with weak violating couplings; **bottom left**: spin-1 with vector coupling; **bottom right**: spin-1 with axial-vector coupling. Existing constraints from the anomalous magnetic moment of the muon are shown in gray. The pion decay constraints from PIENU are shown in red and the PIONEER sensitivity (assuming 10^{10} recorded $\pi^+ \to \mu^+ \nu_\mu$ events) in blue.

the axialvector, this constraint is stronger by at least an order of magnitude. We do not show constraint from the kaon decays $K \to \pi X$ due to their strong model dependence, see the discussion in section 3.3.5.

The existing PIENU searches for $\pi^+ \to e^+ \nu_e X$ cover otherwise unexplored parameter space of the weak violating axion-like particle model. We find that PIONEER can improve over PIENU by approximately one order of magnitude in coupling sensitivity. However, in all but the weak violating axion-like particle model, the other dark sector probes turn out to be more constraining. In the case of dark sector particles coupled to muons, we find that the anomalous magnetic moment of the muon gives stronger constraints for all models we consider. In principle, the constraint from $(g-2)_\mu$ can be avoided by invoking a tuned cancellation with other unrelated loop contributions. We do not consider such a scenario.

The above discussion highlights the weak violating axion-like particle coupled to electrons as a useful benchmark model for dark sector searches at PIONEER.

Another class of models that is very weakly constrained by other probes are light dark sector particles that couple exclusively to neutrinos. So-called Majorons are a prominent example of such particles [212–215]. A study of the PIONEER sensitivity to light neutrinophilic scalars can be found in [216].

§ 3.4 Conclusions

In this work we have presented a systematic study of light dark sector searches in exotic pion decays at the PIENU experiment and the forthcoming PIONEER experiment. Exploiting the strong helicity or phase-space suppression of Standard Model pion decays, we showed that even very weakly coupled new particles can lead to observable signatures in both two-body and three-body decay channels. We considered a number of dark sector scenarios, including sterile neutrinos produced in two-body decays $\pi^+ \to \ell^+ N$, as well as invisible scalars, axion-like particles, or dark vectors contributing to the three-body decays $\pi^+ \to \ell^+ \nu_\ell X$.

The sterile neutrinos yield monochromatic positrons from $\pi^+ \to e^+ N$ or muons from $\pi^+ \to \mu^+ N$ and can be searched for using bump hunts. We validated our sterile neutrino analysis framework by successfully reproducing existing PIENU limits. Building on this, we have estimated the sensitivity of PIONEER using detailed simulations based on its current detector design. The main sterile neutrino results are summarized in figures 3.1 and 3.2. Already in phase 1, PIONEER will be able improve over PIENU by at least one order of magnitude and will start to reach into see-saw motivated parameter space.

Light, invisible dark sector particles can be produced in $\pi^+ \to e^+ \nu_e X$ and $\pi^+ \to \mu^+ \nu_\mu X$ decays and result in continuous distortions in the low-energy tail of the positron and muon energy spectra. We have successfully reproduced the existing PIENU limits on the $\pi^+ \to \ell^+ \nu_\ell X$ branching ratios in the benchmark scenario used by the experiment. We have then recast the PIENU limits to dark sector scenarios with invisible scalars, axion-like particles, and dark vectors. Due to the different shape of the charged lepton energy spectrum in each model, we find that the pion decay constraints vary by $\mathcal{O}(1)$ between the different scenarios. Using simulated data we found that PIONEER will extend sensitivity to significantly smaller branching ratios, improving the PIENU bounds by one to two orders of magnitude. These results are summarized in figures 3.7 and 3.8.

By comparing with complementary probes, including lepton anomalous magnetic moments, BaBar mono-photon searches, and NA64 beam-dump searches, we determined the role of pion decays in testing regions of model parameter space. We have found that pion decays into weak-violating axion-like particles emerge as a well-motivated benchmark, with PIENU already probing novel parameter space and PIONEER expected to improve sensitivity by roughly an additional order of magnitude.

Our analysis has focused on fully invisible dark sector particles, where missing-energy signatures dominate. An important extension of this program will be the exploration of visible decay signatures from short-lived dark sector states, which could further enhance the discovery potential of PIONEER. We will study these visible signatures in a follow up paper. Overall, our results stress that exotic pion decays form an important component of PIONEER's program and provide a relevant avenue for probing light dark sectors.

Chapter 4

Dark Plasmas in the Nonlinear Regime: Constraints from Particle-in-Cell Simulations

This chapter is based on [217]. In this work, we perform the fist ever dedication particle-in-cell simulations of dark matter plasma instabilities in order to constrain dark charge-to-mass ratios. In Sec. 4.1, we define the secluded dark $U(1)_D$ model and the beam plasma scenario. In Sec. 4.2, we outline the relevant parameters of the Bullet Cluster system and highlight key features used in our simulation. In Sec. 4.3, we describe the simulation methods used. In Sec. 4.4, we show the results of our simulations and derive a conservative constraint on the dark charge-to-mass ratio. We conclude and summarize our results in Sec. 4.5.

§ 4.1 Dark U(1) Plasma

The SM electromagnetic sector exhibits a rich phenomenology of plasma behaviors; a dark sector with a massless $U(1)$ mediator is therefore a natural place to look for similar effects. As such, we will take as our fiducial model that of "dark electromagnetism," namely a sector composed of "dark" electrons and positrons interacting via a massless $U(1)$ gauge boson. This model has been explored previously in the literature [218–220] and has been found to be able to produce the correct relic abundance of dark matter [220]. The Lagrangian is given by

$$\mathcal{L} = \mathcal{L}_{\text{SM}} - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \bar{\chi}(\gamma^\mu(i\partial_\mu - q_\chi A'_\mu) - m_\chi)\chi \tag{4.1}$$

where χ, $\bar{\chi}$ are dark fermion fields with mass m_χ and charge q_χ under a new $U(1)$ gauge symmetry mediated by a massless "dark photon" with the four-potential A'_μ. Here, $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the dark electromagnetic field strength tensor, $\mathcal{L}_{\text{SM}}$ is the standard model Lagrangian, and γ^μ are the Dirac matrices. While the particle content is similar to the SM electromagnetic sector, the charge q_χ and mass m_χ of the fermions are taken to be free parameters. We restrict our attention to the case of a massless mediator and assume throughout this paper that the dark sector is completely decoupled from the Standard Model, i.e., there is no kinetic mixing term $\epsilon F_{\mu\nu}F'^{\mu\nu}$. As discussed in Ref. [218], this is a self-consistent choice, as if ϵ is set to zero at some high scale, $\epsilon = 0$ is preserved under renormalization group evolution.

Neglecting plasma effects, the strongest existing constraints on this model have been found to arise from the survival of elliptical galaxies on cosmological timescales [218, 220]. Dissociative cluster mergers also bound the self-interactions of this model, however, when purely considering two-to-two scattering, the associated bounds are several orders of magnitude weaker than the ellipticity bound (see Fig. 4.1) [218–220]. Despite this, as we show in Sec. 4.4, taking into account *collective* effects allows such mergers to constrain couplings as low as ten orders of magnitude below the ellipticity bound (Fig. 4.1).

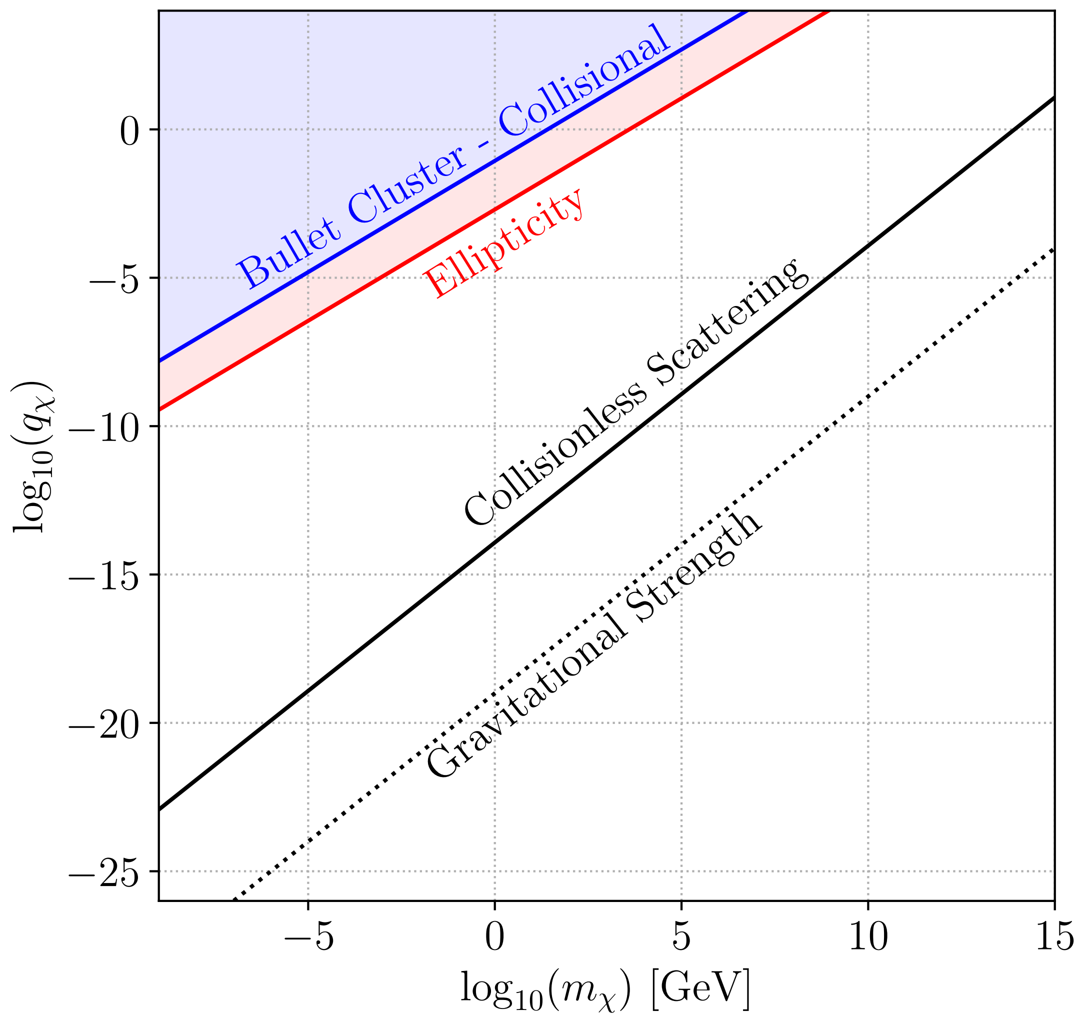

Figure 4.1: Constraints on dark $U(1)$ electromagnetism in the $q_\chi - m_\chi$ plane. Bullet Cluster constraints from two-to-two scattering are shown in blue. Constraints from the survival of elliptical galaxies are shown in red [218–220]. The results presented in this work place constraints above the solid black line.

This is because at the microphysical level, the interactions of a dark $U(1)$ sector in the Bullet Cluster can be well-modeled by two counter-streaming beams of net-neutral pair plasma moving at a relative velocity v_0. We can confirm that such a plasma approximation is valid by checking that the "plasma parameter," which is the number of particles within a sphere of radius λ_D (the effective screening length in the plasma), is much greater than 1. This screening length is set by the ratio of the particle velocity dispersion σ to the plasma frequency $\omega_{\rm pl} = \sqrt{\frac{q_\chi^2 n_\chi}{m_\chi}}$ where n_χ is the number density of the dark matter particles and we adopt units such that $\hbar = c = \varepsilon_0 = \mu_0 = 1$. In the region of parameter space between existing ellipticity bounds and our constraint, λ_D varies between $\sim 10^{-14} - 10^{-3}$ kpc for a 1 GeV particle, corresponding to a plasma parameter of $\Lambda = 4\pi n_\chi \lambda_D^3 \sim 10^{23} - 10^{56} \gg 1$, which indicates that the plasma approximation is valid and the system is dominated by collective effects.

As described in detail in the Appendix (see also [221–225, 225–235]), this system is subject to two primary instabilities: the electrostatic *longitudinal instability* [236] and the electromagnetic *Weibel instability* [237]*. These instabilities act to slow the relative motion of the two beams while heating each individually, even in a fully collisionless regime [59]. The longitudinal instability grows at a rate proportional to the plasma frequency which we find to be

$$\omega_{\rm pl} = 0.5\,{\rm kyr}^{-1} \left(\frac{q_\chi/m_\chi}{2\times 10^{-14}\,{\rm GeV}^{-1}}\right)\left(\frac{\rho_{\rm DM}}{0.035\,{\rm GeV/cm^3}}\right)^{1/2} \tag{4.2}$$

where $\rho_{\rm DM}$ is the mass density of the dark matter particles. The Weibel instability grows more slowly, at a rate proportional to

$$\omega_{\rm pl} = 5\,{\rm Myr}^{-1} \left(\frac{q_\chi/m_\chi}{2\times 10^{-14}\,{\rm GeV}^{-1}}\right)\times\left(\frac{\rho_{\rm DM}}{0.035\,{\rm GeV/cm^3}}\right)^{1/2}\left(\frac{v_0}{3000\,{\rm km/s}}\right), \tag{4.3}$$

however it saturates later and at larger electromagnetic field values, hence can play a

*Additional electromagnetic oblique modes can occur in asymmetric beam plasmas. For beams that are approximately symmetric, as in this work, their growth rates and saturation mechanisms only vary slightly and can be determined numerically. See [223] for further discussion.

significant role in the dynamics [58]. Prior to saturation, the system can be solved analytically in certain limits, however, once the system saturates, it becomes nonlinear and requires dedicated plasma simulations to model. Note that the Vlasov-Maxwell system that governs the plasma dynamics (see Appendix E) can be recast in a form in which the only unitful quantity is the plasma frequency, hence results for particular values of (q_χ, m_χ) can be generalized to other regions in parameter space at which the plasma takes on the same $\omega_{\rm pl}$.

§ 4.2 The Bullet Cluster

The Bullet Cluster (1E 0657-56) [238–241] is a well-known dissociative cluster merger that has been used extensively to constrain dark matter self-interactions [242–246]. The system consists of a large main galaxy cluster of mass $\approx 10^{15}\ M_\odot$ and a smaller subcluster (the "Bullet") of mass $\approx 10^{14}\ M_\odot$ that merged at a high relative velocity. The primary observable provided by weak lensing maps of the system is on the observed offset between the mass in dark matter enclosed in the central 150 kpc of the Bullet and the shocked Standard Model gas, which trails behind, which has been used to constrain the collisional cross-section of self-interacting dark matter with mass m to be $\sigma/m \lesssim 1 - 4\,{\rm cm}^2/{\rm g}$ [244–247]. Though there is tension on the precise value of this constraint, in this paper, we adopt $4\,{\rm cm}^2/{\rm g}$ as a conservative limit.

Existing limits on self-interactions are derived under the assumption of individual particle scattering. However, the Bullet Cluster is also ideal for studying the effects of plasma instabilities on dark sectors charged under a long-range interaction. At small scales on which the dark matter density of each cluster varies negligibly, the merger is well-modeled by the two-stream system described in Section 4.1. Two-stream and Weibel instabilities may arise, leading to the formation of small-scale electromagnetic fields that efficiently scatter DM particles through a wide variety of diffusion mechanisms [235, 248–256] inducing an effective collisional cross-section despite the plasma being largely collisionless. We can compute this effective collisional cross-section via dedicated simulations of the formation and saturation of plasma microinstabilities, as discussed in the following Section. Having extracted these collision rates from our simulations, we can rescale the existing constraints on two-to-two scattering to place constraints on collective interactions as well.*

§ 4.3 Simulations

We simulate two counter-streaming beams of net neutral plasma using `Smilei`, an open-source particle-in-cell simulation framework [257]. To connect our results to the Bullet Cluster, we adopt values taken to be consistent with the analytic Hernquist profiles of the merging clusters provided in [245], taking $r = 150\,\mathrm{kpc}$ for both clusters, corresponding to the maximal radius at which dark matter self-interactions affect observations by weak lensing (see Sec. 4.2) [244]. This choice yields velocity dispersions of $\sigma_{\mathrm{Main}} = 1080$ km/s for the Main cluster and $\sigma_{\mathrm{Bullet}} = 630\,\mathrm{km/s}$ in the Bullet sub-cluster and a relative number density of $n_{\mathrm{Main}}/n_{\mathrm{Bullet}} = 1.91$. Since both higher temperatures and larger density ratios tend to inhibit instability growth, the values we have selected are conservative. To test this hypothesis, we performed two additional simulations,

*Note that in this work, we do not explore the formation or stability of dark matter halos, instead assuming that for very weak couplings, the standard picture of structure formation is not appreciably altered. If, however, this is not the case, then stronger constraints may arise from the existence of dark matter halos themselves. We leave a thorough exploration of this to future work.

adopting the parameters that the authors of [245] used to explore the sensitivity of their simulations to initial conditions. Specifically, we increased the mass density of the main cluster to correspond to a halo with concentration parameter $c = 1.94$, rather than $c = 3$, yielding a density ratio of 1.19 (Table 4.1, Simulation R3). We also varied the mass of the Bullet sub-cluster, increasing it by a factor of two in keeping with the mass suggested by strong lensing observations [245]. This yielded a density ratio of 0.96 (Simulation $R4$ in Table 4.1.)

The beams are initialized with a relative bulk velocity of $v_0 = 3000\,\mathrm{km/s}$, which is consistent with the findings of [258] that the DM relative velocity may be much less than the observed gas shock velocity (≈ 4700 km/s). Furthermore, we note that as larger relative velocities only serve to enhance the microinstabilities, this choice of a low v_0 is also conservative. We demonstrate this quantitatively by performing an additional simulation in which v_0 is increased to 4000 km/s, which we find significantly decreases the time to saturation. (See Table 4.1, R2.) All other parameters can be absorbed into the plasma frequency as discussed in Section 4.1. Our choice of parameters places the system in the "warm" regime where the thermal velocity of each stream is comparable to the streams' relative velocity. Warm, non-relativistic pair plasmas are rare in the Standard Model, hence their evolution is less well-studied than the cold and hot regimes, making dedicated simulations even more necessary.

Our simulation consists of a 2D3V* Cartesian geometry with periodic boundary conditions. Each simulation was run for $10^3\,\omega_{pl,B}^{-1}$, where $\omega_{pl,B}$ is the plasma frequency corresponding to the Bullet sub-cluster. Two species were initialized, χ and $\bar{\chi}$, with equal and opposite charge-to-mass ratios and relative bulk velocity such that the simulation is

*For a discussion of how this choice affects our results, see Appendix E.5.

in the center of momentum frame. See Appendix E for further discussion of simulation setup.

§ 4.4 Results

Throughout the simulation, we record the trajectories of 10^4 test particles. At regular intervals, we calculate for each particle

$$\Delta v_t = \frac{\vec{v_0} \cdot \vec{v_t}}{|\vec{v_0}|^2} \tag{4.4}$$

where $\vec{v_0}$ is the particle's velocity vector at the start of the simulation and $\vec{v_t}$ is the particle's velocity vector at time t. If Δv_t changes by a factor of e, we then consider the particle to have undergone a hard scatter, similar to the studies performed in [259–261]. We can then determine the *two-to-two* scattering cross-section that this fraction of scattered particles would correspond to through [262]

$$\frac{\sigma}{m} = -\Sigma^{-1} \log(1 - p), \tag{4.5}$$

where p is the fraction of particles originating from the Bullet sub-cluster that have scattered and $\Sigma \sim 0.33\,\mathrm{g/cm^2}$ is the projected mass surface density of the main cluster at a radius of 150 kpc [245]. This relation allows us to connect our results to existing limits on DM self-interactions from large-scale simulations of the Bullet Cluster.

Fig. 4.2 shows a cumulative count of the percentage of test particles that have undergone a momentum change equivalent to a hard scatter as a function of time. Even though the plasma is collisionless, we find that particles have experienced a significant change in momentum mimicking that of a hard collisional scattering process. Previous

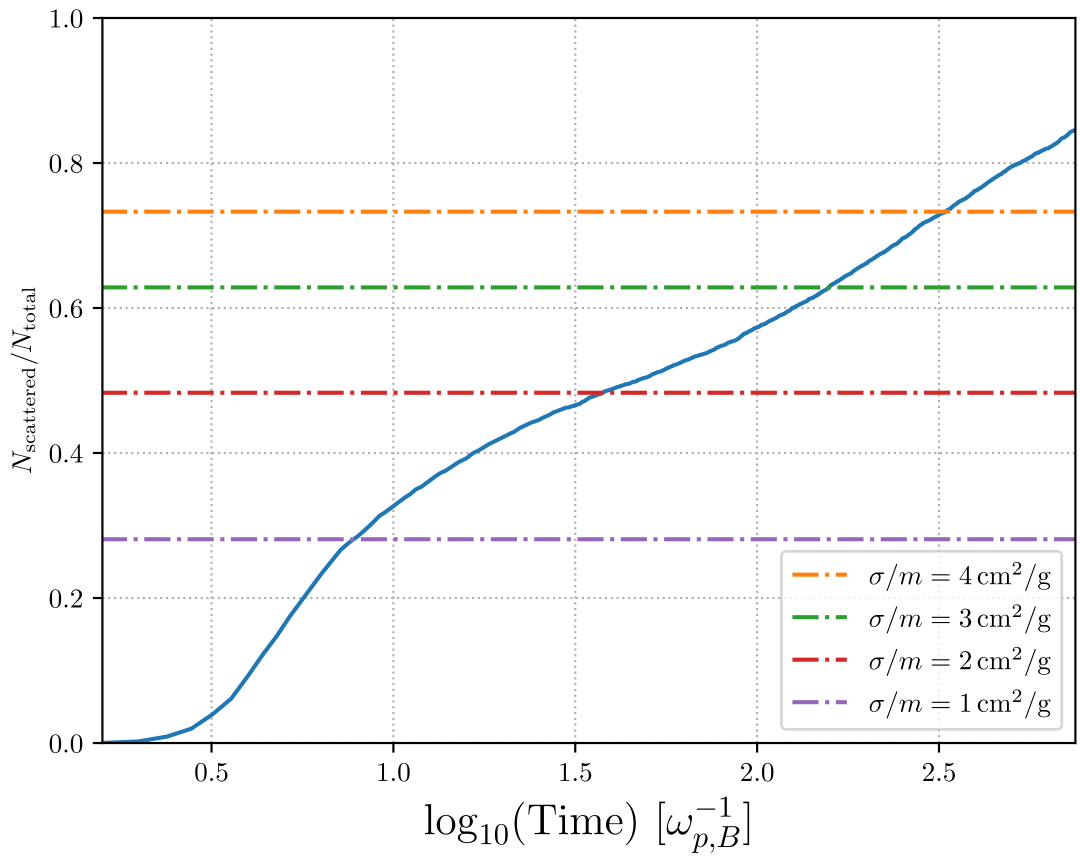

Figure 4.2: Fraction of tracked macroparticles that have undergone an $\mathcal{O}(1)$ momentum change as a function of time (cf. Eq. 4.4). The dash-dotted lines show the fraction of particles that existing simulations show undergo *two-to-two* scattering events during the merger for varying cross-sections [262]. At 330 $\omega_{pl,B}^{-1}$, 73% of particles have been scattered, corresponding to an effective cross-section of $4\,\mathrm{cm}^2/\mathrm{g}$.

studies agree that a collisional cross-section of $\sigma/m \sim 4\,\mathrm{cm}^2/\mathrm{g}$ is ruled out by the centroid offsets of the Bullet Cluster's dark matter halo and galaxies [244–247]. By Eq. 4.5, this corresponds to 73.3% of particles experiencing a significant change in momentum. We therefore place a constraint on all charge-to-mass ratios that effectively scatter 73.3% of particles in less than 1% of the Bullet Cluster crossing time, where we have chosen 1% of crossing as a conservative scale over which the halo density varies negligibly ($\Delta\rho/\rho < 5\%$) hence our simulations remain an appropriate treatment of the system. We estimate the crossing time t_{cross} as twice the time it takes for the Hernquist mass distributions described in Sec. 4.2 to move from a centroid separation of 300 kpc to 0 kpc assuming the halos were infalling from rest at infinite separation and that the infall does not significantly disrupt the halo shape. This calculation yields $t_{\text{cross}} = 1.07 \times 10^8$ yr. This leads to a constraint on the dark charge-to-mass ratio of

$$q_\chi < 1.2 \times 10^{-14} \left(\frac{m_\chi}{\text{GeV}}\right). \tag{4.6}$$

This corresponds to a length scale of the fastest growing modes of $\lambda_{\text{TS}} = 0.05$ kpc for the two-stream instability and $\lambda_{\text{W}} = 6.5$ kpc for the Weibel instability. We further verify our results by explicitly computing the slowdown of the bulk velocity of the counter-streaming beams as a function of time. We show the evolution of the velocity distributions in Fig. 4.3, where the left-hand (center) panel shows the velocity distribution in the longitudinal (transverse) direction at three different times and the right-hand panel shows the bulk velocities of the two beams as a function of time. We find that the bulk velocity decreases by roughly an order of magnitude by the time of Weibel saturation, with the dominant slowdown being due to diffusion of particles by late-time Weibel modes. By the end of our simulation, the two beams have effectively isotropized in the center-of-momentum frame, as can be seen in the dotted distribution corresponding to $t = 750\,\omega_{pl,B}^{-1}$ in the left and center panels. We find these results to be consistent with the observed slowdown found in other similar simulations of nonrelativistic pair plasmas with $v_0 \sim 0.01c$ and length scales of $\sim 10c/\omega_{pl}$. The simulation parameters of R_3 in [59] match closest to those of our fiducial simulation. After 250 ω_{pl}^{-1}, the authors found that the relative bulk velocity had reduced to $0.442\,v_0$. From our simulation, we find that after the same amount of time, the relative bulk velocity had decreased to $0.444\,v_0$, which we find to be in good agreement.

As discussed in Sec. 4.3, we performed three additional simulations with varying initial conditions to test the robustness of our results. We varied the relative velocity of the halos, as well as their density ratios, in keeping with the analysis in [245]. The parameter choices and resulting time of scattering 73.3% of particles, $t_{73.3}$, and charge-to-mass ratio constraint for each simulation are shown in Table 4.1. We find that simulation $R4$ produces the most conservative constraint on the dark charge-to-mass ratio. Though the initial conditions chosen for $R4$ are less observationally motivated than those that produced our fiducial limit (Eq. 4.6), in the interest of being conservative, we adopt

$$q_\chi < 2.0 \times 10^{-14} \left(\frac{m_\chi}{\text{GeV}}\right). \tag{4.7}$$

as our final constraint, corresponding to the bound set by simulation R4.

This constraint is shown in Fig. 4.1 as a solid black line. We additionally show with the effective charge-to-mass ratio of the gravitational force

$$q_G = \sqrt{4\pi G} \left(\frac{m_\chi}{\text{GeV}}\right) \tag{4.8}$$

with a dotted black line for comparison. The proximity of this curve demonstrates that

Run	v_0 [km/s]	$n_{\rm Main}/n_{\rm Bullet}$	$\sigma_{\rm Main}$ [km/s]	$\sigma_{\rm Bullet}$ [km/s]	$t_{73.3}$ $\left[\omega_{pl,B_F}^{-1}\right]$	q_χ/m_χ $\left[{\rm GeV}^{-1}\right]$
Fiducial	3000	1.91	1080	630	330	1.2×10^{-14}
R2	4000	1.91	1080	630	20	1.0×10^{-15}
R3	3000	1.19	900	630	440	1.6×10^{-14}
R4	3000	0.96	1080	880	540	2.0×10^{-14}

Table 4.1: Initial conditions (left four columns) and resulting constraints (right two columns) from different simulations. R2 utilizes a less conservative estimate of the relative bulk velocity v_0 between the two sub-clusters. R3 (R4) utilizes an alternative mass density of the Main (Bullet) sub-cluster. Here, ω_{pl,B_F} is the plasma frequency of the Bullet sub-cluster from the fiducial simulation. See Sec. 4.3 and 4.4 for definitions of each parameter.

we are probing exceptionally weak couplings; however, as our constraint is still several orders of magnitude above this line, gravitational forces should have negligible effects on the evolution of plasma instabilities in the region of parameter space we constrain. Existing constraints from the ellipticity of dwarf galaxies [218, 220] (red) and two-to-two scattering limits from the Bullet Cluster (blue) are shown as well. Even for our conservative choice of parameters, our results extend existing limits by over ten orders of magnitude.

§ 4.5 Conclusions

In this work, we have shown that long-range collective effects can have a dramatic impact on the large-scale behavior of dark matter in dissociative cluster mergers. We have performed dedicated particle-in-cell simulations of counter-propagating beams of dark plasma to characterize the growth rate, saturation, and late-time behavior of the beams. We find that instabilities produce electromagnetic inhomogeneities that act as scattering sites for dark matter. This process leads to an effective collisionality in the

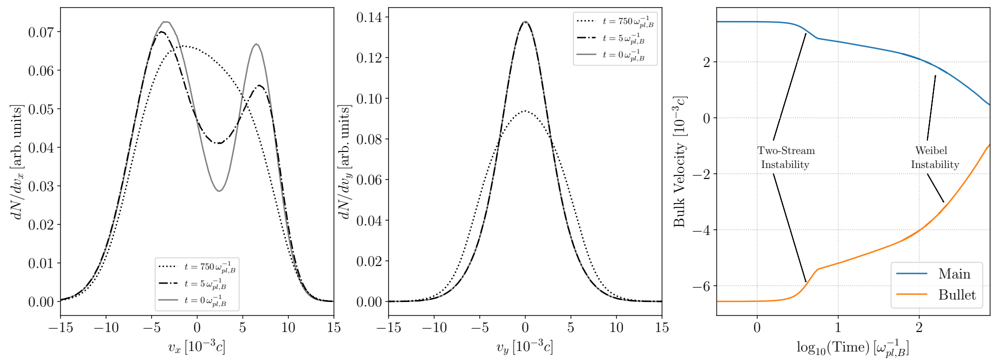

Figure 4.3: Left: Center of momentum velocity distribution of particles in the longitudinal direction at $t = 0\,\omega_{pl,B}^{-1}$ (solid line), $t = 5\,\omega_{pl,B}^{-1}$ (dot-dashed line), and $t = 750\,\omega_{pl,B}^{-1}$ (dotted line). Middle: Center of momentum velocity distribution of particles in the transverse direction at $t = 0\,\omega_{pl,B}^{-1}$ (solid line), $t = 5\,\omega_{pl,B}^{-1}$ (dot-dashed line), and $t = 750\,\omega_{pl,B}^{-1}$ (dotted line). Right: Bulk velocity of particles in the longitudinal direction for particles originating in the Main sub-cluster beam (blue) and Bullet sub-cluster beam (orange).

dark sector, which we then connect to existing limits on DM self-interactions derived from simulations of the Bullet Cluster. The associated results allow us to place new constraints at couplings orders of magnitude below existing limits.

Chapter 5

Structure Formation with Dark Magnetohydrodynamics

This chapter is based on [263]. In this work, we incorporate gravitational effects into the analysis of dark-sector plasma instabilities for the first time. We characterize the alterations to primordial structure formation resulting from the formation of dark plasma instabilities. The chapter will be structured as follows:

1. We begin in Sec. 5.1 with a brief review of the classical gravitational instability without any alterations from dark sector plasma effects. This is a primer for the work done in Sec. 5.2.

2. In Sec. 5.2 we introduce the dark $U(1)_D$ model that we consider throughout this work, and couple dark sector plasma effects to the classical Jeans instability, demonstrating a directionally dependent alteration to both the Jeans length and group velocity in the presence of an external dark magnetic field.

3. Section 5.3 explores the possible generation mechanisms of background dark sector magnetic fields. We introduce the primordial dark magnetic field power

spectra and provide conservative constraints on their present-day amplitudes.

4. We then present the alterations to the linear matter power spectrum under this model in Sec. 5.4. Scanning over parameter space and performing a χ^2 analysis against current and future measurements of the primordial matter power spectrum, we also show our constraints on the background dark magnetic field, as well as the charge-to-mass ratio for this model for various different spectral indices for the primordial magnetic field spectrum.

5. Finally, with Sec. 5.5 we conclude the paper with a discussion of not only the constraints we derived, but also possible future extensions of our work, such as the consideration of kinetic mixing with the Standard Model photon and the possibility of halo triaxiality as a result of anisotropic collapse due to the directionally dependent alteration to the Jeans length.

§ 5.1 Gravitational Collapse and classical Jeans length

The classical Jeans length is given by $\lambda_J = \frac{2\pi}{k_J}$, where k_J is given by

$$k_J = \left(\frac{4\pi G\rho_0}{c_s^2}\right)^{1/2} \tag{5.1}$$

Here, c_s denotes the speed of sound in the medium, and ρ_0 is the unperturbed mass density of the ensemble. This defines the critical length at which the gas becomes stable against gravitational collapse. Specifically, only perturbations with wavelengths that exceed the Jeans length $\lambda > \lambda_J$ will trigger instabilities [264].

While the traditional Jeans length formula $\lambda_J = \sqrt{\frac{\pi c_s^2}{G\rho}}$ breaks down for CDM due to its zero sound speed [265], an effective Jeans length can be defined that varies with scale. At small scales, this effective Jeans length is negligibly small, allowing CDM to cluster on all scales. On larger scales, the expansion of the Universe acts as an effective pressure, leading to a comoving Jeans length approximated by $\lambda_J \approx \frac{2\pi c_s}{Ha\sqrt{\Omega_m}}$, where H is the Hubble parameter, a is the scale factor, and Ω_m is the matter density parameter [266]. The effective sound speed c_s is related to the velocity dispersion $\sigma(R)$ at a given scale R by $c_s^2 \approx \frac{\sigma^2(R)}{3}$. The velocity dispersion itself scales with the enclosed mass as $\sigma_{200} = A[h(z)\frac{M_{200}}{10^{15}M_{\odot}}]^{\alpha}$, where $h(z)$ is the dimensionless Hubble parameter, $A \approx 1000$ km/s, $\alpha \approx 0.35 - 0.36$ [267, 268]. This leads to a scale-dependent Jeans length that scales approximately linearly with R: $\lambda_J(R) \propto R$. This scale-dependent Jeans length is crucial for understanding structure formation in CDM models, supporting the bottom-up scenario of hierarchical clustering. Unlike warm or hot DM, CDM lacks a definitive cutoff scale for structure formation, which contributes to the small-scale challenges in ΛCDM cosmology [269].

§ 5.2 Evolution of Dark Magnetized Plasma

To allow for collective plasma effects, we consider a dark sector consisting of dark positrons and electrons of equal abundance that interact through a massless $U(1)_D$ gauge boson, which we will refer to as a “dark photon." The dark sector Lagrangian for this model takes the form

$$\mathcal{L}_{\text{Dark}} = -\frac{1}{4}F'^{\mu\nu}F'_{\mu\nu} + \bar{\chi}(\gamma^{\mu}(i\partial_{\mu} - q_{\chi}A'_{\mu}) - m_{\chi})\chi \tag{5.2}$$

where χ, $\bar{\chi}$ are dark electrons and positrons, respectively, with charge q_χ and mass m_χ, and A' is the dark photon. An additional term of the form $\frac{\epsilon}{2}F'^{\mu\nu}F_{\mu\nu}$ would allow for the possibility of kinetic mixing between the dark photon and the Standard Model photon. In this work, we will only consider the case where kinetic mixing does not occur, giving rise to a secluded dark sector. Note that if ϵ is set to zero at some high energy scale, it will remain zero under renormalization group evolution unless there are heavy states charged under both $U(1)_D$ and $U(1)_{\mathrm{E}M}$ [218]. We leave the study of $\epsilon \neq 0$ for future work and restrict the model to the scenario of a massless dark photon.

Previous works have utilized the evolution of beam instabilities in plasmas to derive constraints on such a model [58, 217, 270]. For the purpose of characterizing the effects of plasma instabilities and how they alter structure formation in the Universe, beam-type instabilities do not provide the correct picture for this scenario. Instead, we consider a homogeneous and isotropic, net-neutral plasma of particles to characterize the criteria for gravitational collapse in the presence of collective plasma effects.

The evolution of a plasma can be well-approximated by a fluid model. Though this assumption is typically made in highly collisional systems, it can also provide a framework to study the evolution of collisionless systems whose velocity distributions remain approximately Maxwellian. In a collisional plasma, particle collisions drive the distribution function toward a Maxwellian, as explained by Boltzmann's H-theorem. Even in collisionless plasmas, collective kinetic processes such as phase mixing, Landau damping, wave-particle scattering, and microinstabilities can produce distributions that are close to Maxwellian over macroscopic scales, despite the absence of true collisions. These systems may, however, evolve away from Maxwellians (e.g., toward power-laws) on longer timescales [271]. We leave a more detailed treatment of these effects to future work. Treating the plasma as a fluid yields the system of partial differential equations often referred to as the magnetohydrodynamic (MHD) equations. Adopting a system of units in which $\hbar = c = \epsilon_0 = \mu_0 = 1$, the system of equations is

$$\nabla \cdot \vec{E} = \rho_c \qquad \nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{5.3}$$

$$\nabla \cdot \vec{B} = 0 \qquad \nabla \times \vec{B} = \vec{J} + \frac{\partial \vec{E}}{\partial t} \tag{5.4}$$

$$\frac{\partial \rho_m}{\partial t} + \nabla \cdot (\rho_m \vec{U}) = 0 \qquad \rho_m \frac{d\vec{U}}{dt} = \vec{J} \times \vec{B} - \nabla \cdot \overset{\leftrightarrow}{P} \tag{5.5}$$

$$\frac{d}{dt}\left(\frac{p_\perp B^2}{\rho_m^3}\right) = 0 \qquad \frac{d}{dt}\left(\frac{p_\parallel}{\rho_m B}\right) = 0 \tag{5.6}$$

$$\vec{E} + \vec{U} \times \vec{B} = \frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho_m}\frac{\partial \vec{J}}{\partial t} \tag{5.7}$$

Here $\vec{U}$ is the bulk velocity of the DM fluid, ρ_c and ρ_m are the charge and mass density of the fluid, respectively. $\overset{\leftrightarrow}{P}$ is the pressure tensor containing only diagonal elements of $p_\perp$ in the directions perpendicular to the magnetic field and $p_\parallel$ in the direction parallel to the magnetic field. Additionally, we have assumed the Chew-Goldberger-Low (CGL) equation of state for cylindrically symmetric magnetized plasmas [272]. The last equation is a specific scenario of the generalized Ohm's law discussed in Appendix F.

To incorporate gravitational effects into our system of equations, we only need to include Poisson's equation,

$$\nabla^2 V = 4\pi G \rho_m, \tag{5.8}$$

where V is the gravitational potential, along with a correction to the conservation of momentum equation

$$\rho_m \frac{d\vec{U}}{dt} = \vec{J} \times \vec{B} - \nabla \cdot \overset{\leftrightarrow}{P} - \rho_m \nabla V. \tag{5.9}$$

We consider the case of an infinitely large homogeneous self-gravitating plasma in the presence of a uniform background magnetic field. After perturbing this initial state, small perturbations become unstable and trigger the Jeans instability. Considering electromagnetic effects, we find that such perturbations obey the following dispersion relation in the linear regime

$$\left[\rho\omega^2 - (1+\beta)k_\perp^2 p - \beta B^2 k^2 + 4\pi G\rho^2 \frac{k_\perp^2}{k^2}\right]\left(\rho\omega^2 - 3k_\parallel^2 p + 4\pi G\rho^2 \frac{k_\parallel^2}{k^2}\right)$$
$$= k_\perp^2 k_\parallel^2 \left[\left(p - \frac{4\pi G\rho^2}{k^2}\right)^2 + 2\left(p - \frac{4\pi G\rho^2}{k^2}\right)(1-\beta)\,p\right] \tag{5.10}$$

with

$$\beta = \frac{1}{1 + \frac{k^2}{2\rho}\frac{m_X^2}{q_X^2}} \tag{5.11}$$

Here p is the unperturbed initial isotropic pressure. A detailed derivation is given in Appendix G. It is assumed that all fields are proportional to $\exp i(\vec{k}\cdot\vec{r} - \omega t)$ and the perpendicular and parallel directions are measured with respect to the background magnetic field, B. Note that we assume ρ_c is negligibly small due to the large time and length scales of the system of interest. Hence, we drop the m subscript from ρ_m for ease of notation.

5.2.1 Parallel propagation

We first consider waves oriented parallel to the magnetic field ($k_\perp = 0$); Eq.(5.10) then reduces to

$$\left[\rho\omega^2 - k_\parallel^2 \beta B^2\right] \times \left(\rho\omega^2 - 3k_\parallel^2 p + 4\pi G\rho^2 \frac{k_\parallel^2}{k^2}\right) = 0\,. \tag{5.12}$$

This leads to two potential solutions. The first

$$\frac{\omega^2}{k_\parallel^2} = \frac{\beta B^2}{\rho}, \tag{5.13}$$

only contains solutions with real-valued ω. Hence, no unstable modes are present. The second solution

$$\frac{\omega^2}{k_\parallel^2} = \frac{3p_\parallel}{\rho} - \frac{4\pi G\rho}{k_\parallel^2}, \tag{5.14}$$

shows that the Jean's instability can trigger if

$$k_\parallel < \left(\frac{4\pi G\rho^2}{3p}\right)^{1/2} \tag{5.15}$$

which is independent of β. This implies that the speed of sound in the parallel direction is unaltered by the background magnetic field.

$$c_{s,\parallel}^2 \equiv \frac{3\,p}{\rho}. \tag{5.16}$$

Using the CGL equation of state, we can rewrite the sound speed in terms of the dispersive velocity of the DM particles,

$$p_\parallel = nk_BT = \rho v_{\mathrm{th}}^2, \tag{5.17}$$

allowing us to simplify the effective parallel group velocity to the following form

$$c_{s,\parallel}^2 = 3v_{th}^2\,. \tag{5.18}$$

5.2.2 Perpendicular propagation

For modes perpendicular to the background magnetic field ($k_\parallel = 0$), the dispersion relation contains both perpendicular pressure and electromagnetic support,

$$\left[\rho\omega^2 - k_\perp^2\left(p(1+\beta) + \beta B^2\right) + 4\pi G\rho^2\right] = 0. \tag{5.19}$$

These modes only contain a single solution for ω

$$\frac{\omega^2}{k_\perp^2} = \frac{p(1+\beta)+\beta B^2}{\rho} - \frac{4\pi G\rho}{k_\perp^2}\,. \tag{5.20}$$

Here, we only have a Jean's-type instability that triggers when

$$k_\perp < \left(\frac{4\pi G\rho^2}{p(1+\beta)+\beta B^2}\right)^{1/2}\,. \tag{5.21}$$

This leads to an effective perpendicular group velocity that receives contributions from the Alfvén group velocity, $v_A = B/\sqrt{\rho}$

$$c_{s,\perp}^2 \;\equiv\; \frac{(1+\beta)\,p}{\rho} + \beta\frac{B^2}{\rho}. \tag{5.22}$$

We can simplify the perpendicular group velocity using the CGL equation of state for the perpendicular component,

$$p_\perp = \frac{\rho v_{\text{th}}^2}{2}, \tag{5.23}$$

giving us

$$c_{s,\perp}^2 = \frac{v_{\text{th}}^2}{2} + \beta\left(\frac{v_{\text{th}}^2}{2} + \frac{B^2}{\rho}\right)\,. \tag{5.24}$$

Note that in the limit of $q_\chi/m_\chi \to 0$, we have $\beta \to 0$. This shows that the plasma decouples from the background magnetic field in the small charge limit, and we retain the classic Jeans length for a cylindrically symmetric system. In the following sections, we investigate the modification of the sound speed in motivated magnetic field models.

§ 5.3 Dark Magnetic Fields: Generation, Cosmological Constraints, and Caveats

If DM particles carry charge under this hidden gauge group, the dark plasma can support dark electric and magnetic fields, collective excitations, and MHD phenomena analogous to those in the visible sector. A central open question is whether dark magnetic fields can be dynamically significant during cosmic structure formation. In such scenarios, gravitational collapse competes with magnetic pressure and tension, modifying the DM power spectrum and the internal structure of halos. In the following subsections, we review the evolution of SM cosmological magnetic fields as well as specific constraints that should additionally provide constraints on dark magnetic fields.

5.3.1 Dark magnetic field evolution

Translating constraints at a given cosmological epoch (e.g., big bang nucleosynthesis (BBN) or recombination) to one at structure formation requires understanding how the magnetic field evolves under cosmological evolution. The evolution differs considerably depending on whether the magnetic field is turbulent or ordered. Heuristically, the transition from order to turbulence depends on the magnetic Reynolds number, $\mathrm{Re}_M \equiv v_\ell \ell/\eta$, where v_ℓ is the velocity of the (dark) plasma on scale ℓ, and η is the (dark) magnetic diffusivity, defined as $\eta \equiv 1/\sigma$, where σ is the (dark) conductivity. We assume the dark sector is fully ionized at all relevant epochs, which is expected since the dark photon temperature is much larger than the binding energy of dark positronium in the relevant parameter space. The dark conductivity is then

$$\sigma_\chi = \frac{q_\chi^2 n_\chi}{m_\chi \nu_{\chi\chi}} = \frac{12\pi^{3/2} T_\chi^{3/2}}{q_\chi^2 m_\chi^{1/2} \ln \Lambda}, \tag{5.25}$$

where $\nu_{\chi\chi} = q_\chi^4 n_\chi \ln \Lambda / 12\pi^{3/2} m_\chi^{1/2} T_\chi^{3/2}$ is the Coulomb scattering frequency of χ particles*, $\ln \Lambda$ is the Coulomb logarithm factor, and T_χ is the DM temperature. As we will be considering very small dark charge, we will typically have $\mathrm{Re}_M \gg 1$, which means the evolution of magnetic fields is well-described by ideal MHD:

$$\frac{\partial \vec{B}_{\text{com}}}{\partial t} = \frac{1}{a} \vec{\nabla} \times (\vec{v}_b \times \vec{B}_{\text{com}}) \tag{5.26}$$

where $\vec{B}_{\text{com}} = a^2 \vec{B}_{\text{phys}}$ is the comoving magnetic field, $\vec{B}_{\text{p}hys}$ is the physical field, and $\vec{v}_b$ is the bulk velocity of the plasma. On large comoving scales, $\ell \gg v_b/aH$, the comoving magnetic field remains constant, giving rise to the standard evolution of the physical field, $\vec{B}_{\text{p}hys} \propto a^{-2}$. Below this scale, MHD turbulence develops, suppressing both density perturbations and magnetic fields [273–277]. The transition scale is

$$\ell \sim \frac{\left\langle B_{\text{p}hys}^2 \right\rangle^{1/2}}{aH\sqrt{\rho_m}} \approx 0.1\,\text{M}pc \left(\frac{B}{1\,\text{n}G} \right), \tag{5.27}$$

where $\langle ... \rangle$ represents an ensemble average. Assuming no helical fields, the PMF power spectrum is given by

*Exclusion of $\chi, \bar{\chi}$ collisions does not alter our conclusions.

$$\left\langle B_i(\vec{k})B_j^*(\vec{k}')\right\rangle = (2\pi)^3\delta^3(\vec{k}-\vec{k}')\left(\delta_{ij}-\frac{k_ik_j}{k^2}\right)\frac{P_B(k)}{2}, \tag{5.28}$$

where $B_i(\vec{k}) \equiv \int d^3x B(\vec{x})\exp(i\vec{k}\cdot\vec{x})$ is the Fourier transform of the magnetic field. The power spectrum is typically parametrized as follows:

$$P_B(k) = Ak^n e^{-k^2\ell^2}, \tag{5.29}$$

where n is the model-dependent spectral tilt, and ℓ is the transition scale in (5.27). It is assumed that MHD dissiption is the only cutoff in the PMF power spectrum. PMFs generated during inflation result in nearly scale-invariant spectra with $n \approx -3$. PMFs generated after inflation must be causally connected which constraints $n = 2$ [278]. Constraints on the magnetic field are typically quoted in terms of the comoving field strength, B_{com}, at some reference scale λ_B, typically 1 Mpc

$$\begin{aligned} B_{com}(\lambda) \equiv \left\langle B^2(\lambda)\right\rangle^{1/2} &\propto \left[\int_0^\infty k^2 P_B(k)dk\right]^{1/2} \\ &= B(\lambda_B)\left(\frac{\lambda}{\lambda_B}\right)^{-\frac{n+3}{2}}\Gamma^{1/2}\left[(n+3)/2\right], \end{aligned} \tag{5.30}$$

where $\Gamma(x)$ is the gamma function. The physical RMS magnetic field at comoving scale λ_{com} and redshift z can be represented in terms of the comoving magnetic field at reference scale λ_B as

$$B_{\text{phys}}(\lambda_{\text{com}}, z) = B_{\text{com}}(\lambda_B) \left(\frac{\lambda_{com}}{\lambda_B}\right)^{-\frac{n+3}{2}} (1+z)^2. \quad (5.31)$$

Throughout the paper, we do not assume any particular mechanism to generate PMFs. However, in order to avoid the development of turbulence on scales of interest, we assume that the energy injection scale is much smaller than the transition scale in Eq. (5.27). In the following subsection, we summarize current constraints on SM PMFs and show how these bounds can be mapped onto constraints on dark PMFs.

5.3.2 Constraints on standard model primordial Magnetic Fields

Like SM magnetic fields, dark magnetic fields may also exist on large scales. For these magnetic fields to have appreciable effects on dark sector dynamics, they must be seeded by some mechanism in the early Universe and may be enhanced through dynamo effects as the Universe expands. Such mechanisms fall into two main categories: formation of primordial magnetic fields in the early Universe, and late-Universe dynamo amplification of weak seed fields generated by battery mechanisms.

A common mechanism for primordial seed magnetic field generation is the so-called Biermann battery, in which the plasma pressure gradient and temperature gradient are not aligned [279, 280]. In SM plasmas, the misalignment arises from electron-ion drift effects originating from the mass hierarchy between the two charged species. The Biermann battery is inoperative for charged species with equal masses, such as in the dark $U(1)_D$ model under consideration. In this instance, the contributions to the pressure gradient from the positive and negative species cancel. For more details see Appendix F. Kinetic plasma instabilities, such as the Weibel instability [237] can also generate and amplify magnetic fields in plasmas with anisotropic velocity distributions [281, 282]. In the SM sector, this is typically studied in the context of merging and collapsing galaxies. These mechanisms could be extended to the dark sector as well if significant velocity anisotropies can be created.

Alternatively, more exotic mechanisms have been considered to generate seed primordial magnetic fields. During the inflation, primordial magnetic fields can be generated by explicitly breaking the conformal invariance of the electromagnetic action. References [283–293] consider direct couplings of the electromagnetic fields to other fields that explicitly break the conformal symmetry during inflation. Additionally, models have been suggested that add an $RA_\mu A^\mu$ term to the action but are strongly disfavored since they contain ghosts [294].

Lastly, phase transitions in the early Universe could lead to the generation of magnetic fields. In the SM sector, electroweak phase transition and the QCD transition are often considered [273, 295]. The dark $U(1)_D$ could be the result of some other phase transition in the dark sector or from couplings to SM. First-order phase transitions provide ideal generation mechanisms as colliding bubble walls can seed turbulent fields. Reference [296] also claims that the gradient of the Higgs vacuum expectation value can induce electromagnetic fields at an efficient rate during electroweak phase transition.

After seed field generation, the primordial magnetic fields may be enhanced exponentially by dynamo processes. For turbulent fields, the fluctuation dynamo [297] provides the most promising growth mechanism. Given we are in a regime of very high magnetic Reynolds number, the seed field can grow rapidly [298] but may not grow optimally due to the DM's compressible nature [299].

Any magnetic field (SM or dark) will contribute to the Hubble parameter governing the expansion rate of the Universe. If magnetic fields are too large, then the correct abundance of ^{4}He will not be produced. Ref. [300] finds that an upper limit of $B_{\text{phys}} < 2 \times 10^{11}$ Gauss can be placed at $T = 10^9$ K for magnetic fields that are uniform over scales larger than the Hubble radius. This corresponds to $B_{\text{phys}} < 0.7\,\mu$G today.

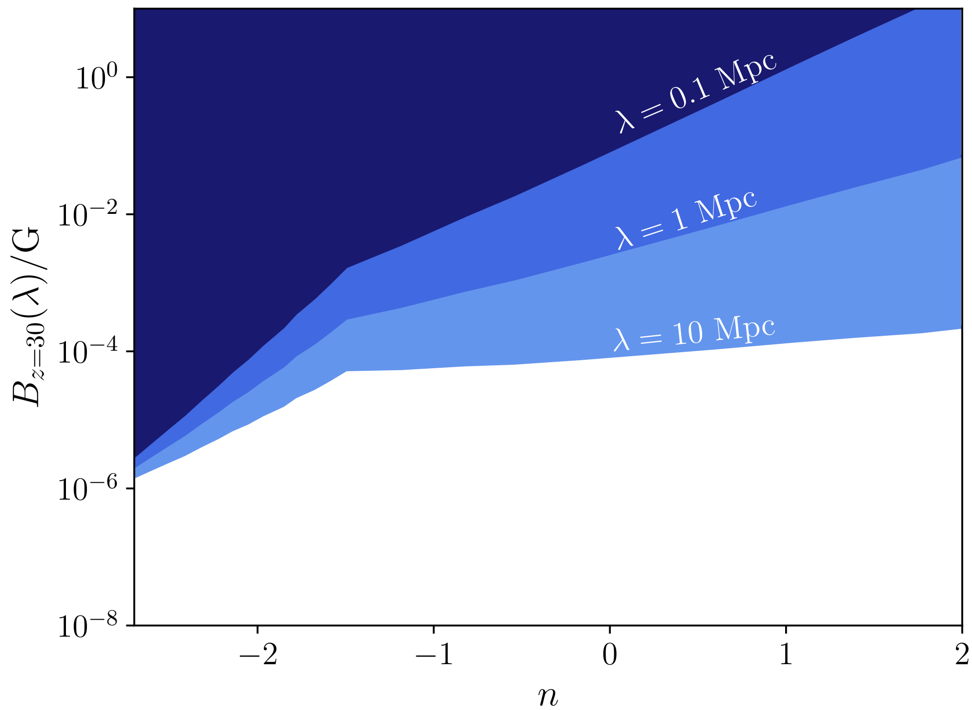

Figure 5.1: Constraints on the magnetic field strength at different comoving scales: $\lambda = 0.1$ Mpc (dark blue), 1 Mpc (medium blue), and 10 Mpc (light blue) at redshift $z = 30$ as a function of spectral tilt of the magnetic power spectrum, n, reproduced from Ref. [301].

Random magnetic fields may also generate tensor perturbations leading to anisotropies in the CMB through their coupling to gravitational waves. Ref. [301] find that for a scale invariant spectrum on galactic scales, $B_{\text{phys}} < 1$ nG. For a spectral index $n < -3/2$, $B_\lambda < 7.9e^{3n}\,\mu$G and for a spectral index of $n > -3/2$, $B_\lambda < 95e^{0.37n}$ nG where B_λ is the amplitude at $\lambda = 0.1\,h^{-1}$ Mpc. These constraints, translated to constraints at $z = 30$

using Eq. (5.31), are shown in Fig. 5.1.

§ 5.4 Effects on Matter Power Spectrum

The primordial matter spectrum $P(k)$ is a power spectrum that is used to quantify the fluctuations in the DM density in the Universe following cosmological inflation with respect to the inverse length scale k, with units $h/\mathrm{M}pc$ in a comoving reference frame. Alternative DM models beyond the ΛCDM paradigm can potentially alter the matter power spectrum; thus, in this section, we present the alteration of the matter power spectrum due to the presence of a magnetized dark plasma.

We have demonstrated that through the alteration of the Jeans length in the directions perpendicular to a uniform background dark magnetic field, collective effects in the dark sector as a result of long-range interaction produce an anisotropic pressure support controlled by the dark magnetic field oriented in the direction $\hat{\boldsymbol{b}}$ with coherence length L_B. Linear density modes then obey

$$\ddot{\delta}_{\boldsymbol{k}} + 2H\dot{\delta}_{\boldsymbol{k}} + \left[\frac{c_s^2(\mu)k^2}{a^2} - 4\pi G\bar{\rho}_m\right]\delta_{\boldsymbol{k}} = 0, \tag{5.32}$$

with

$$c_s^2(\mu) = c_{s,\parallel}^2\,\mu^2 + c_{s,\perp}^2(1-\mu^2), \qquad \mu \equiv \hat{\boldsymbol{k}}\cdot\hat{\boldsymbol{b}}, \tag{5.33}$$

where $\bar{\rho}_m$ is the average matter density, it is defined $\bar{\rho}_m = \Omega_m\rho_{\text{crit}}$ where $\Omega_m \sim 0.3$ today, which implies a directional Jeans scale

$$k_J^2(\mu, a) = \frac{4\pi G\bar{\rho}_m(a)\,a^2}{c_s^2(\mu)}. \tag{5.34}$$

For this section, we set the matter density, ρ, equal to $\bar{\rho}_m$. The derivation in Sec. 5.2

shows that the restoring forces differ for perturbations parallel and perpendicular to the background field direction $\hat{\boldsymbol{b}}$.

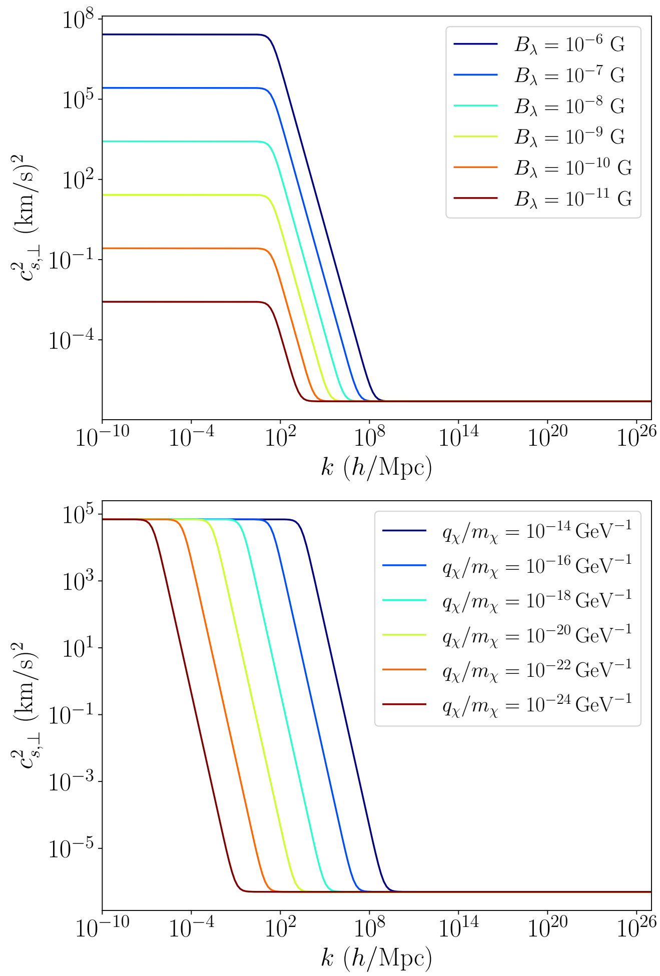

Figure 5.2: In the plots above, we show the perpendicular group velocity as a function of wave number. We set our reference scale for the magnetic field spectrum to be $\lambda = 0.1\,h^{-1}\,\mathrm{M}pc$. Left: the charge-to-mass ratio has been fixed to $10^{-16}\,\mathrm{G}eV^{-1}$ and we vary over different strengths of the dark magnetic field. Right: here, the dark magnetic field has been fixed to $B_\lambda = 10^{-6}$ G and we vary over different charge-to-mass ratios.

While assuming that the coherence length of the dark magnetic field is larger than the MHD turbulence scale defined in Eq. (5.27) ($L_B \gg \ell$), the effective group velocity,

c_s^2 from Eq. (5.33), acquires a k- and angle-dependent pressure term, via Eq. (5.18) and Eq. (5.24):

$$c_s^2(\mu, k, z) = 3v_{\text{th}}^2\,\mu^2 + \left[\frac{1+\beta}{2}v_{\text{th}}^2 + \beta\frac{B_{\text{com}}^2}{\rho(z)}\left(\frac{k}{k_B}\right)^{n+3}(1+z)^4\right](1-\mu^2); \quad (5.35)$$

We choose fiducial values of v_{th} and ρ to be ~ 18 cm s^{-1} and $\sim 40\ M_\odot$ kpc^{-3} at $z = 0$. In addition, we set the reference scale $\lambda = 0.1\,h^{-1}$ Mpc which translates to $k_B = \frac{2\pi}{\lambda} \sim 63\,h/$Mpc for the rest of the paper. In Fig. 5.2 (left) we plot the perpendicular sound speed for various values of the magnetic field today, B_λ, with a fixed value of q_χ/m_χ and in Fig. 5.2 (right) we plot the perpendicular sound speed for various values of q_χ/m_χ with a fixed value of B_λ. In both cases, the magnetic field couples to the dark plasma for sufficiently large scales and completely decouples for sufficiently small scales. This can also be seen in the definition of β in Eq. (5.11), where $0 < \beta < 1$ for all parameters. We can see that changing the magnetic field alters the sound speed's amplitude, whereas the charge-to-mass ratio determines the transition scale at which the dark plasma couples to the magnetic field and behaves as an ideal MHD fluid.

For the purpose of comparing altered matter power spectra from alternative DM models to that of CDM, it is convenient to define the transfer function, given by

$$T(k)^2 = \frac{P(k)}{P_{\Lambda\text{CDM}}(k)}, \quad (5.36)$$

where $P(k)$ is the power spectrum under an alternative DM model, and $P_{\Lambda CDM}(k)$ is the isotropic CDM matter power spectrum. Here, we only study the effects of the linear evolution of density perturbations in Eq. (5.32). Though we expect our model to additionally impact the nonlinear matter power spectrum, such effects would require dedicated simulation which we consider to be beyond the scope of this paper. In our model, to compare with the isotropic matter power spectrum, we need to integrate over the directionality parameter μ. To do so, we expand the power spectrum in terms of Legendre polynomials, $\mathcal{L}_\ell(\mu)$,

$$P(k,\mu) = \sum_{\ell=0} P_\ell(k)\mathcal{L}_\ell(\mu), \tag{5.37}$$

which admits Legendre multipoles (for $\ell = 0, 2, \ldots$),

$$P_\ell(k,z) = \frac{2\ell+1}{2}\int_{-1}^{1} d\mu\; P(k,\mu,z)\;\mathcal{L}_\ell(\mu). \tag{5.38}$$

Thus, the $\ell = 0$ multipole will be the isotropic matter power spectrum from our model. In Fig. 5.3, we plot several examples of $P_0(k,0)$ for various values of the spectral index n. We show the current data from Planck 2018 [302, 303], DES Y1 [304], SDSS DR7 [305], eBOSS DR14 [306], MW Satellites [307], as well as projected future measurements from CMB-HD lensing [61], and the HERA and EDGES 21-cm projections [62, 63, 308]. In Fig. 5.4 we plot the corresponding transfer function of each of the curves from Fig. 5.3. We note that, unlike warm DM or SIDM models, we do not see pure exponential suppression or dark acoustic oscillations (DAOs), but rather a smooth reduction/alteration in power, followed by a continued power-law fall-off, with DAOs occurring at much larger wave numbers. In principle, there is a scale such that all power is suppressed, though we do not probe these scales. Additionally, it has recently been shown that the solutions to Eq. (5.33) closely match those from state-of-the-art numerical MHD simulations for scales larger than the MHD turbulent scale, $\lambda \gg \ell$ [309].

To place constraints on magnetic fields in the secluded $U(1)_D$ gauge sector, we scan over a 2D parameter space of the magnetic field today at reference scale $\lambda = 0.1\,h^{-1}$ Mpc, and the DM charge-to-mass ratio. At each point in the scan, we numerically solve Eq. (5.33) to generate $P_0(k, 0)$ and calculate the χ^2 statistic against the current and future power spectrum data given in Fig. 5.3. The resulting parameters that can be excluded at the 68%, 95%, and 99.73% confidence levels are shown in the solid pink, purple, and red colored regions, respectively. Along with these contours, we also show projected constraints from future observations given by the dashed lines—the color scheme has the same meaning.

In each plot, we also fix the spectral index, n, to several representative values. For each index, the region is constrained from above by limits from CMB anisotropies [301], noting that the strongest constraints are for the scale invariant power spectrum, $n = -3$. Simulations of dark plasma instabilities in the Bullet Cluster additionally constrain the DM charge-to-mass ratio $q_\chi/m_\chi < 2 \times 10^{-14}\,\mathrm{GeV}^{-1}$ [217]. Lastly, the weak gravity conjecture [310] mandates that theories of quantum gravity with a $U(1)_D$ gauge field must satisfy $q_\chi/m_\chi > 1/M_P$, where M_P is the reduced Planck mass. This additionally constrains the dark charge-to-mass ratio from below. We find that with current data, the dark magnetic fields that have not already been excluded by measurements of the CMB anisotropies (with the small exception of the $n = -2$ case, see Fig. 5.5) do not significantly alter the linear matter power spectrum, and are therefore consistent with current cosmology. With future measurements at higher wave numbers, we expect to be able to probe deviations from the CDM power spectrum.

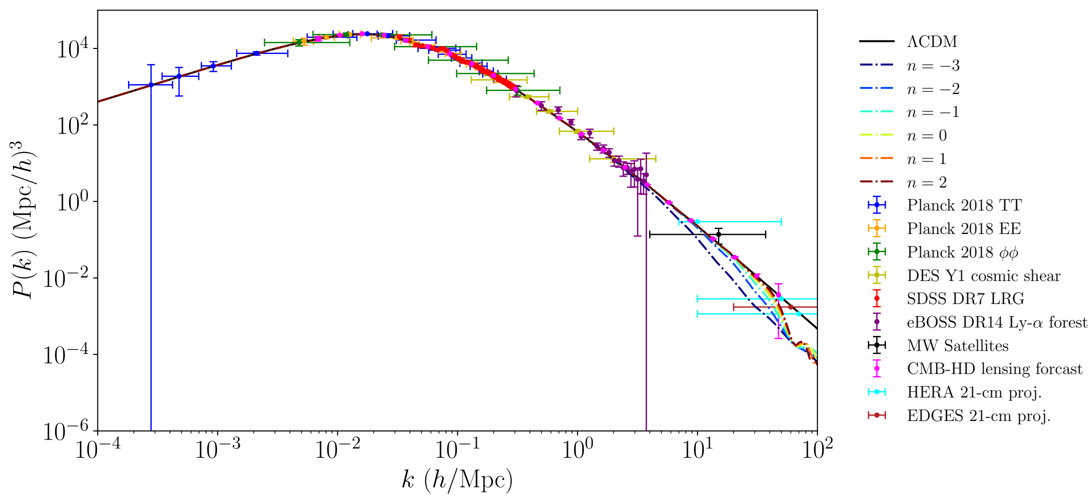

Figure 5.3: Here we show the isotropic power spectrum for a fixed charge-to-mass ratio, $q_\chi/m_\chi = 10^{-18}$ GeV^{-1}, and a fixed background dark magnetic field value of $B_\lambda = 5 \times 10^{-9}$ Gauss for different spectral indices against the current measurements as well as future/planned surveys (see text for details). We plot using different colored dotted lines to denote various scenarios of a background dark magnetic field, based on constraints derived in the previous section.

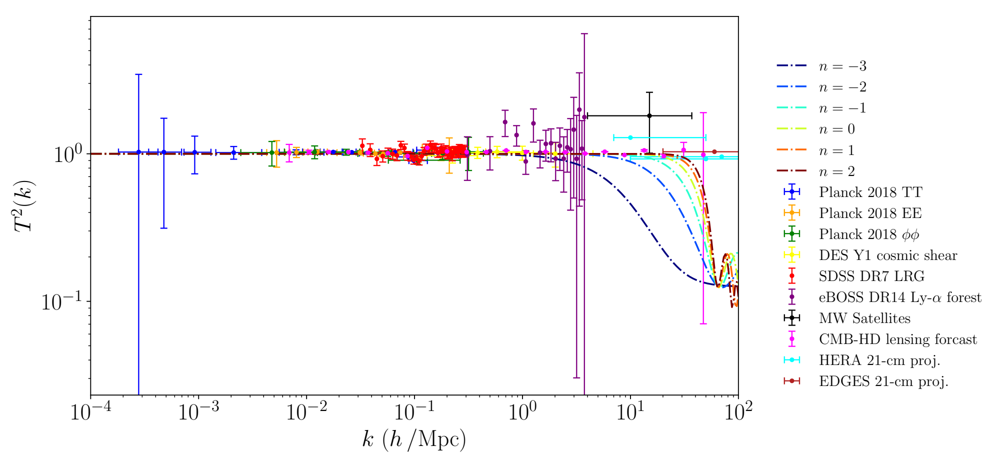

Figure 5.4: Here we show the transfer function at a fixed charge-to-mass ratio, $q_\chi/m_\chi = 10^{-18}$ GeV^{-1}, and a fixed background dark magnetic field value of $B_\lambda = 5 \times 10^{-9}$ Gauss for different spectral indices. We have overlaid current measurements and future/planned surveys (see text for details).

§ 5.5 Discussion and Conclusions

In this work, we show that in the presence of large-scale dark magnetic fields, DM halos charged under a secluded $U(1)_D$ gauge field experience magnetic pressure that delays gravitational collapse via the Jeans instability. For correlation lengths on scales much larger than the turbulent scale, ℓ, we solve Eq. (5.33) to map the effects onto the linear matter power spectrum. Figure 5.5 shows that current measurements of the linear matter power spectrum provide constraints weaker than constraints set by tensor modes of the CMB. However, we expect future observations to probe into unconstrained regions of parameter space, further closing the gap between the constraint set by Ref. [217] and the weak gravity conjecture.

This work can be extended in several new directions to incorporate well-motivated models of DM. Kinetic mixing between the SM photon and the dark photon can induce a small effective electric charge for particles in the dark sector. As a result, such particles, commonly referred to as millicharged DM, acquire suppressed couplings to SM electromagnetic fields in addition to their interactions with dark sector forces. This scenario has garnered significant interest due to the potential for millicharged DM to interact with both visible and dark sector electromagnetic fields.

Previous literature has examined the interactions of millicharged DM with SM particles and electromagnetic fields. In particular, the evolution of millicharged DM in the presence of plasma instabilities has been used as a powerful tool to constrain its properties, as demonstrated in studies of supernova remnants [270] and the Bullet Cluster [58]. Considering a similar framework in which millicharged DM couples solely to SM electromagnetic fields, the model used in this work can be extended to incorporate baryonic effects by utilizing a two-fluid MHD model that includes gravity. Though SM primordial fields are generally better understood in the literature, the two-fluid model is sufficiently complex that the system likely needs to be solved numerically. We intend to consider this scenario in future work.

Another future direction is the consideration of the evolution of anisotropic collapse and halo triaxiality due to dark plasma effects. As a result of the directionally dependent alteration to the Jeans length, halo triaxiality and its radial evolution can serve as a direct macroscopic tracer of microscopic dark sector plasma effects. The anisotropic screening mechanism developed here predicts coherent, orientation-dependent modifications to collapse along and across the local dark-field direction $\hat{\boldsymbol{b}}$, leading to systematic biases in intrinsic shape distributions. Relative to the CDM baseline, dark-plasma screening enhances the contrast between inner and outer halo shapes and induces characteristic axis twists that preserve memory of the underlying field geometry. This "radial memory" manifests as steeper inner-outer ellipticity gradients and alignment signatures that cannot be replicated by baryonic condensation or isotropic self-interaction models. Triaxiality therefore emerges as a sensitive diagnostic of anisotropic support in the dark sector, complementing traditional small-scale power-spectrum constraints.

Several independent probes can access these shape signatures across mass and redshift. Weak-lensing quadrupoles and projected ellipticity profiles offer stackable, statistical tests of enhanced inner–outer contrast; strong-lensing reconstructions at $r \sim 0.05$–$0.2\,R_{\rm vir}$ provide direct measures of central isopotential flattening and axis twists; and x-ray or SZ isophote ellipticities trace the same anisotropy in the intracluster gas. Correlations between halo orientations and large-scale filaments, as well as anisotropies in satellite and splashback distributions, further constrain the coherence and preferred direction of $\hat{\boldsymbol{b}}$.

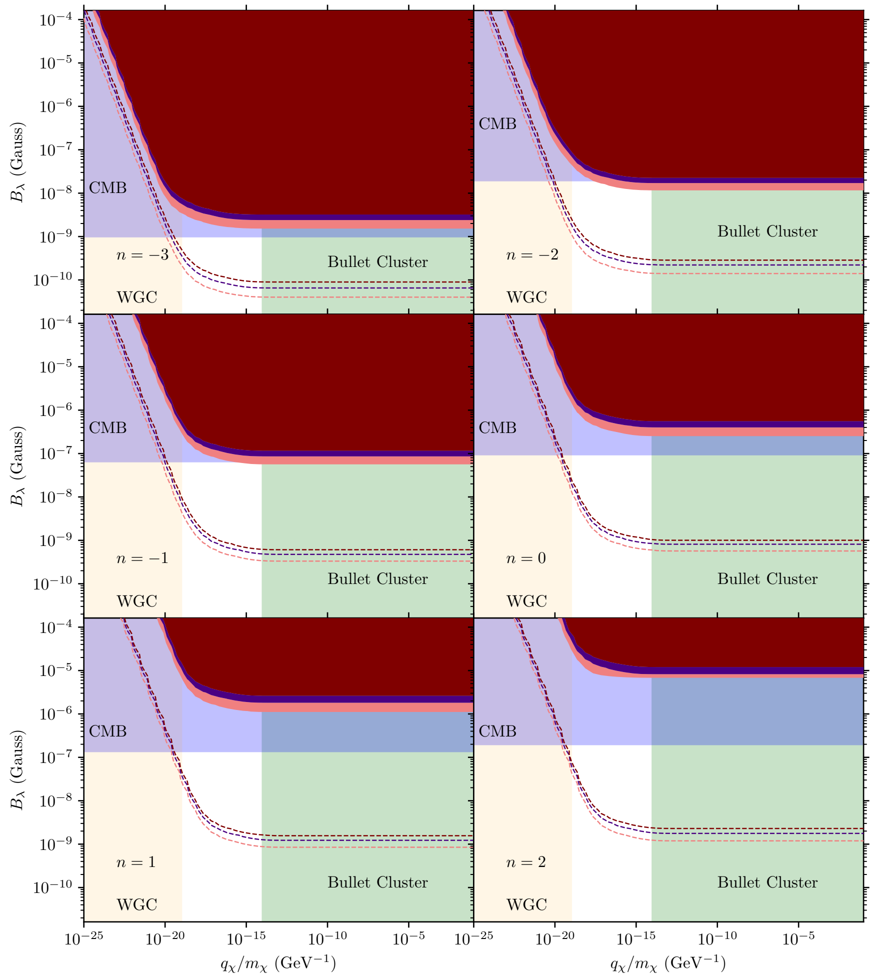

Figure 5.5: Here we show the resultant parameter space for scans using our effective group velocity Eq. (5.35) to grow the linear density modes via Eq. (5.33) for current (shaded regions) and future (dashed lines) measurements of the linear matter power spectrum shown in Fig. 5.3 (see text for the data details). The region shown in pink is excluded at the 1σ level, the purple is excluded at the 2σ level, and the red region is excluded by a minimum of 3σ. We also show constraint regions from the CMB [301] (blue), Bullet Cluster [217] (green), and the weak gravity conjecture (yellow) [310].

In Milky Way–like and more massive disk galaxies, however, a strong stellar bar can imprint nonaxisymmetric signatures on both the luminous and dark components, including boxy/peanut bulges, isophotal twists, and apparent inner–outer misalignments [311–315]. These bar-driven distortions are largely confined to radii $\lesssim (1\text{–}2)\, R_{\rm d}$ and are tied to the presence, pattern speed, and orientation of the stellar disk, in contrast with the large-scale, approximately radius-independent alignment predicted for the anisotropic-screening axis $\hat{\boldsymbol{b}}$. At dwarf and low-mass-galaxy scales, hydrodynamic simulations and H I kinematic studies show that baryons can produce both oblate and mildly prolate dark-matter haloes, with the halo flattening and triaxiality correlating with gas fraction and the development of a stellar disk [316–318]. Gas-rich dwarfs with $M_{\rm gas}/M_{\rm baryon} \gtrsim 0.5$ tend to host oblate haloes, whereas more stellar-dominated systems span a wide range of shapes from significantly flattened to nearly spherical or slightly prolate [316]. In ΛCDM, baryonic condensation in more massive dwarfs and spirals also tends to round the central halo relative to dark-matter–only expectations [319], while ultrafaint, gas-poor dwarfs retain the strongly prolate shapes of collisionless haloes [317]. Our plasma-mediated scenario instead predicts a coherent preferred axis for the inner halo that can persist even in gas-poor systems, so a joint detection of (i) an unusually steep ellipticity gradient or twist relative to ΛCDM expectations and (ii) a consistent alignment axis across dwarfs, massive disks (with and without bars), and clusters would be difficult to reproduce with purely baryonic or elastic-SIDM rounding alone [320]. Such cross-comparisons would establish halo triaxiality as a powerful, multiobservable avenue for testing plasma-mediated structure formation in the dark sector.

Chapter 6

Conclusion

Throughout these chapters, we show that there exist many exciting opportunities in the near future to search for BSM physics. In the collider setting, there are new prospects of probing unexplored parameter space. Future beam-dump and pion-decay experiments provide unique opportunities to discover HNLs with smaller mixing angles than previously achievable. In the astrophysical setting, dark matter plasma instabilities and collective effects open a new window for studying DM models. Compared with previous studies, DM models with long-range interactions can be ruled out for couplings approaching the limit set by the weak gravity conjecture.

In Chapter 2, we perform a dedicated study of the production of HNLs in beam dump experiments at future e^+e^- colliders. We specifically study the sensitivity of the ILC, C^3, and CLIC experiments and find that the results are highly complementary to the reaches of collider and proton beam dump experiments, being able to probe HNLs with masses less than $\sim 10\,$GeV. Due to the high number of electrons on target, the linear colliders provide an even greater opportunity when compared to their circular collider counterparts.

In Chapter 3, we perform a dedicated study on the new physics capabilities of the PIONEER experiment. We analyze the scenario of two-body decays of $\pi^+ \to \ell^+ N$ as well as three-body decays $\pi^+ - > \ell^+ \nu_\ell X$ for invisible scalars, ALPs, and dark vector particles. To validate our analysis, we reproduce the constraints from PIENU. In all models considered, we find that the PIONEER experiment can probe regions of parameter space with couplings at least an order of magnitude smaller than those constrained by existing limits.

In Chapter 4, we perform the first ever dedicated particle-in-cell simulations modeling plasma instabilities in a secluded dark sector with a dark $U(1)_D$ gauge symmetry. Modeling the dissociative cluster merger of the Bullet Cluster as a two-beam system, we show that both longitudinal electrostatic and transverse electromagnetic instabilities are required to scatter enough particles to be inconsistent with current observations. This leads to a constraint on the dark charge-to-mass ratio of $q_\chi / m_\chi < 2.0 \times 10^{-14}\,\text{GeV}^{-1}$, over ten orders of magnitude below previous studies.

In Chapter 5, we continue studying the same secluded dark $U(1)_D$ model in the context of magnetized plasmas during structure formation. We derive a set of dark MHD equations to show that the classical Jeans' instability is significantly altered in the presence of large-scale dark magnetic fields. These effects will alter the linear matter power spectrum by increasing the effective Jeans length from the dark magnetic pressure. Although most unconstrained parameters are consistent with current observations of the linear matter power spectrum, significant portions of the parameter space can produce a unique signature in future experiments.

Appendix

A Simulation Details

§ A.1 Meson Productions

In electron-proton collisions, mesons are mainly produced when a photon is emitted from the incoming electron and interacts with partons in the proton. The photon can interact directly with quarks inside the proton or it can fluctuate into a hadronic state and interact with quarks and gluons in the proton. The total cross-section for photon-hadron collisions is dominated by soft QCD processes. To simulate the production of mesons, we use `Pythia 8.3` [85, 86] using the `SoftQCD` event class which generates all processes having significant contributions to the total hadronic cross-section. We point out that we are pushing Pythia beyond its nominal operation, so there will be some significant uncertainty associated with our calculations.

To determine the total production rate of mesons we run `Pythia 8.3` using the flags `SoftQCD:all = on`, `PDF:beamA2gamma = on`, and `Photon:ProcessType = 0`. The result of this simulation is given in table A.1 where we show the number of mesons produced per generated soft QCD Monte Carlo event for an incoming electron beam with an energy of 125 GeV, 500 GeV, and 1.5 TeV. The total number of mesons

Meson Type	125 GeV	500 GeV	1.5 TeV
π^+	2.95	3.43	3.87
π^-	2.65	3.13	3.58
K^+	0.26	0.31	0.37
K^-	0.21	0.27	0.32
K_L^0	0.22	0.27	0.32
D^+	3.81×10^{-4}	5.76×10^{-4}	8.19×10^{-4}
D^-	4.11×10^{-4}	6.16×10^{-4}	8.48×10^{-4}
D_s^+	1.08×10^{-4}	1.79×10^{-4}	2.52×10^{-4}
D_s^-	1.18×10^{-4}	1.93×10^{-4}	2.58×10^{-4}
B^+	4.60×10^{-7}	1.80×10^{-6}	3.30×10^{-6}
B^-	1.30×10^{-7}	1.50×10^{-6}	3.30×10^{-6}

Table A.1: The number of mesons produced per generated Monte Carlo soft QCD event ($N_{M,\text{MC}}/N_{\text{MC}}^{\text{tot}}$) at the three different beam energies. To obtain the number of mesons per EOT, we have to multiply these numbers by $\sigma_{\text{SoftQCD}}/\sigma_{eN}$, as described in the text. Note that here we are quoting the electron beam energy instead of the center-of-mass energy of the corresponding collider.

produced per EOT is determined by

$$n_M \equiv \frac{N_M}{\text{EOT}} = \frac{\sigma_{\text{SoftQCD}}}{\sigma_{eN}} \frac{N_{M,\text{MC}}}{N_{\text{MC}}^{\text{tot}}} , \tag{A.1}$$

where σ_{eN} is the total electron-nucleus cross section (~ 46 mb) and σ_{SoftQCD} is the total hadronic cross section estimated by Pythia ($\sim 2 \times 10^{-3}$ mb, $\sim 7 \times 10^{-3}$ mb, $\sim 10^{-2}$ mb for the 125 GeV, 500 GeV, and 1500 GeV beam energy, respectively). $N_{M,\text{MC}}$ is the total number of mesons produced by Pythia and $N_{\text{MC}}^{\text{tot}}$ is the total number of Monte Carlo events that were simulated. n_M is the input for eq. (2.6) when we calculate the number of HNLs produced from meson decays.

To determine the geometric acceptance and efficiency of the electron beam dump experiment, we also need to know the four-momenta of the mesons produced in the electron collision with the beam dump. The charged meson momentum distribution of Pythia was compared to HERA results in [321], which aimed to tune the photoproduction framework of Pythia to match HERA data. We used this tune to compute the meson distributions. To generate a large sample of mesons, we use different settings for the QCD production as follows:

- For $\pi^\pm$, $K^\pm$: "`SoftQCD:all = on`" is used.
- For $D^\pm, D_s^\pm$: "`HardQCD:all = on`" and "`PhotonParton:all=on`" are used.
- For $B^\pm$: "`HardQCD:qqbar2bbbar = on`" and "`PhotonParton:ggm2bbar = on`" are used.

The `HardQCD` and `PhotonParton` settings were used for $D^\pm_{(s)}$ and $B^\pm$ mesons because with the `SoftQCD` event class the number of these mesons produced per electron on target is very small (cf. table A.1). While `SoftQCD` will give a more realistic meson spectrum, generating a reasonable number of events to analyze will be computationally very time-consuming, e.g., $\mathcal{O}(10^{11})$ events will need to be generated to produce $\sim$10k B mesons. The number of mesons produced per electron on target using the `HardQCD` event class is significantly larger. Therefore, we use it to obtain a large sample D and B mesons. The `HardQCD` event class generates mesons with slightly higher p_T and larger energy, as shown in figure A.1. This leads to a slightly lower geometric acceptance. However, we have checked that the limits on the HNL parameter space do not vary appreciably if one uses the spectrum of mesons produced with `SoftQCD` or `HardQCD` event class.

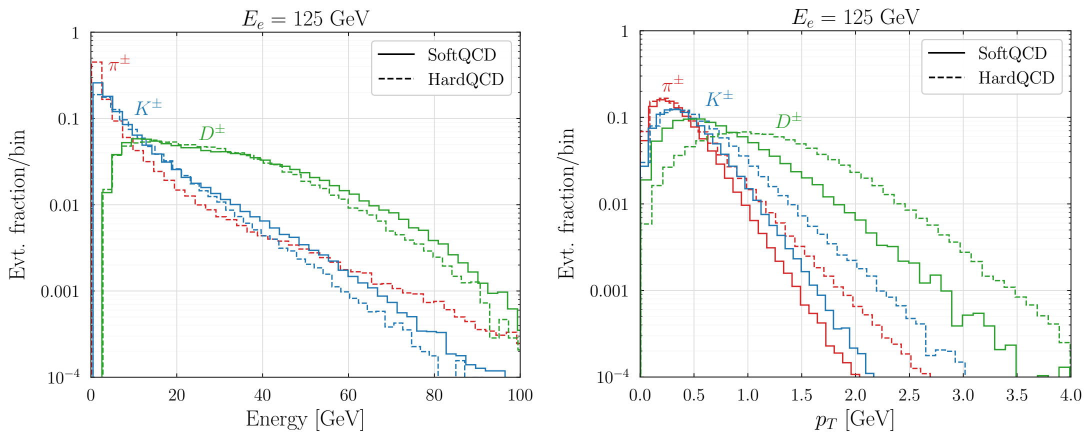

Figure A.1: Energy (left) and p_T (right) distributions of mesons produced at an ILC beam dump experiment with an electron beam with energy $E_e = 125$ GeV hitting a proton at rest. Events are generated with `Pythia` using the `SoftQCD` event class (solid curves) and the `HardQCD` event class (dashed curves). The distributions are normalized such that the total bin count sums to 1.

§ A.2 Three-Body Decays of HNLs

To determine the sensitivity to the HNL parameter space requires simulating the decays of HNLs to two- and three-body final states and using the resulting four-momenta to determine the geometric acceptance of the detector. For two-body HNL decays (e.g., $N \to \ell^\pm \pi^\mp$), this can be done analytically, starting from the rest frame of the HNL and boosting into the lab frame. For the three-body HNL decays to charged leptons, we use the publicly available code `muBHNL` [24, 322, 323], which uses the differential decay distributions of the HNL to generate a weighted sample of final state leptons. Events are generated in the rest frame of the HNL. They are then boosted to the lab frame using the 4-momenta of mesons generated by Pythia. We then use eq. (2.8) to compute the efficiency and eq. (2.10) to calculate the reach on the HNL parameter space.

B Details on the PIENU Sterile Neutrino Search Recast

The PIENU experiment performed a bump search for sterile neutrinos in $\pi^+ \to e^+ N$ decay for sterile neutrino masses in the range of $(62 - 134)$ MeV [98]. Here we provide a detailed description of our reproduction of the search.

In [98], the PIENU collaboration shows the observed positron energy spectrum, the fitted model of the spectrum without a sterile neutrino, and the residuals of the fit for positron energies between 4 MeV and 56 MeV. This range of positron energies corresponds to sterile neutrino masses between approximately 62 MeV and 134 MeV (cf. equation (3.11)). The residuals shown by PIENU are rebinned and need to be interpreted with care. We re-extracted the residuals ourselves from the shown data and the fitted model that consists of the $\pi^+ \to e^+ \nu_e$ decay, as well as backgrounds from muon decay at rest and in flight. Our residuals, including statistical uncertainties only, are shown in figure B.1.

The PIENU paper [98] also shows the positron energy spectrum of the $\pi^+ \to e^+ N$ signal with a sterile neutrino mass $m_N \simeq 90\,$MeV. For this mass, the spectrum has a peak at $\bar{E}_e \simeq 40\,$MeV and is smeared out, for example, due to the calorimeter response. To obtain the signal shape for different sterile neutrino masses, we assume that the width

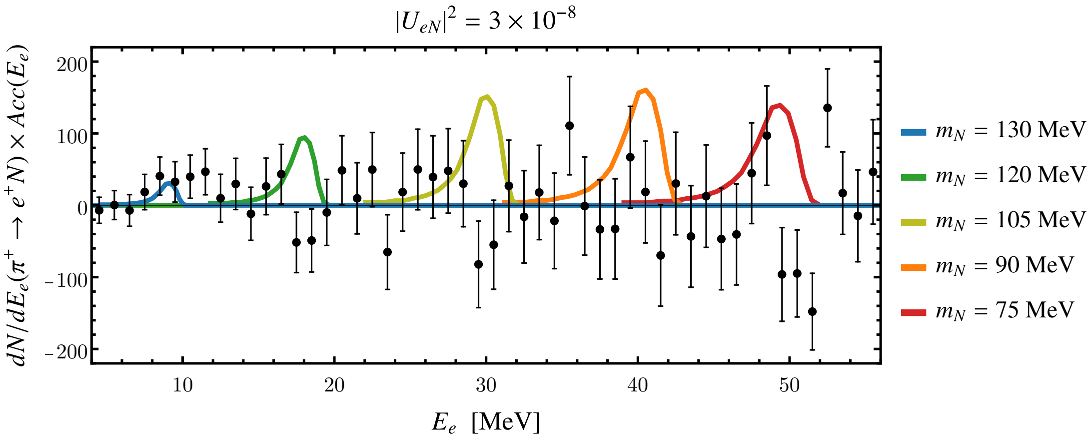

Figure B.1: Expected number of detected $\pi^+ \to e^+N$ decays at the PIENU experiment as a function of the positron energy for several values of the sterile neutrino mass. The sterile neutrino mixing angle is set to $|U_{eN}|^2 = 3 \times 10^{-8}$. For comparison, the residuals extracted from [98] are shown in black.

of the shown spectrum scales with the square root of the positron energy. Such a scaling of the energy resolution with the square root of the energy was used in parts of the PIENU $\pi^+ \to \mu^+N$ analysis [135].*

To obtain the number of $\pi^+ \to e^+N$ events as a function of the sterile neutrino model parameters, we use

$$\frac{N(\pi^+ \to e^+N)}{N(\pi^+ \to e^+\nu_e)} = \frac{\mathrm{BR}(\pi^+ \to e^+N)}{\mathrm{BR}(\pi^+ \to e^+\nu_e)} \,, \tag{B.1}$$

with the ratio of branching ratios given in equations (3.8)-(3.9). From the fitted $\pi^+ \to e^+\nu_e$ component of the positron energy spectrum shown in [98], we extract $N(\pi^+ \to e^+\nu_e) \simeq 1.4 \times 10^6$.

*Since this scaling assumption is approximate, we performed several cross-checks. In correspondence with the PIENU collaboration [324], we obtained an approximate expression for the PIENU energy resolution as a function of the positron energy. As a first check, we modeled the signal as a Gaussian with a width corresponding to this provided resolution. As a second check, we retained the shape of the sterile neutrino signal shown by PIENU, but rescaled its width using the relative energy resolution from the same expression instead of the $\sqrt{E_e}$ scaling. All methods yielded very similar results, confirming the robustness of our approach.

In figure B.1, we show the expected spectra of a sterile neutrino signal for several choices of the sterile neutrino mass. In the plot, we include the detector acceptance correction $Acc(E_e)$ given in [98] as a function of the positron energy. For illustration, we set the mixing angle to $|U_{eN}|^2 = 3 \times 10^{-8}$ and compare the spectra to our extracted residuals.

To obtain bounds on the mixing angle $|U_{eN}|^2$, we fit the energy spectrum of the sterile neutrino signal to our extracted residuals. Emulating the approach of the PIENU analysis, we include in the fit also components from the SM decay $\pi^+ \to e^+\nu_e$ and backgrounds from muon decay at rest and muon decay in flight. We keep the shape of these components fixed to the ones shown in [98], but let their normalizations float. The uncertainties of the event numbers across the energy bins are taken to be uncorrelated.

Imposing the physical prior BR$(\pi^+ \to e^+N) \geq 0$, we obtain the 90% C.L. bound on the sterile neutrino mixing angle $|U_{eN}|^2$ shown in figure 3.1. We also attempted an alternative approach to constrain $|U_{eN}|^2$ without making use of any background shape information. Demanding that the predicted number of sterile neutrino events in a given energy bin does not exceed the observed number of events gives, in principle, a very robust bound. However, we find that it is approximately 1 to 2 orders of magnitude weaker than the constraint from the bump hunt.

We follow a similar procedure to reproduce the PIENU searches for sterile neutrinos in $\pi^+ \to \mu^+N$ decays. The PIENU analysis [135] considered separate signal selections with low muon energy and high muon energy, corresponding to sterile neutrino masses of 28 MeV to 34 MeV and 16 MeV to 29 MeV, respectively. We extract residuals between the provided event data and the fitted model that consists of the $\pi^+ \to \mu^+\nu_\mu$ decay, as well as backgrounds from radiative muon decay (high energy analysis) and pion decay in flight (low energy analysis). Our residuals are shown in figure B.2.

The number of expected sterile neutrino signal events is determined by

$$\frac{N(\pi^+ \to \mu^+ N)}{N(\pi^+ \to \mu^+ \nu_\mu)} = \frac{\mathrm{BR}(\pi^+ \to \mu^+ N)}{\mathrm{BR}(\pi^+ \to \mu^+ \nu_\mu)} , \tag{B.2}$$

with $N(\pi^+ \to \mu^+ \nu_\mu) = 9.1 \times 10^6$ in the high energy analysis and $N(\pi^+ \to \mu^+ \nu_\mu) = 1.3 \times 10^8$ in the low energy analysis [137]. The signal shape was chosen to be a Gaussian. In the low energy analysis, we fix the width of the Gaussian to 0.6 MeV, which corresponds to the width of the main $\pi^+ \to \mu^+ \nu_\mu$ peak. In the high energy analysis, we use an energy dependent width as detailed in [135]: at the main $\pi^+ \to \mu^+ \nu_\mu$ peak ($E_\mu \simeq 4.1$ MeV) the width corresponds to the width of the peak which is approximately 0.16 MeV. At lower energies, the width scales with the square root of the energy. The resulting signal shapes are shown in the plots of figure B.2 for a few example choices of the sterile neutrino mass for the high energy (top) and low energy (bottom) analysis.

To obtain bounds on the sterile neutrino mixing angle with muon neutrinos, $|U_{\mu N}|^2$, we perform a very similar analysis as for the electron events discussed above, including the additional components from $\pi^+ \to \mu^+ \nu_\mu$ and radiative decay backgrounds. The pion decay in flight background in the low energy analysis is parameterized by a second order polynomial with three free parameters that are allowed to vary in the fit. The 90% C.L. bound on the mixing angle $|U_{\mu N}|^2$ is shown in figure 3.2, after imposing the physical prior $\mathrm{BR}(\pi^+ \to \mu^+ N) \geq 0$.

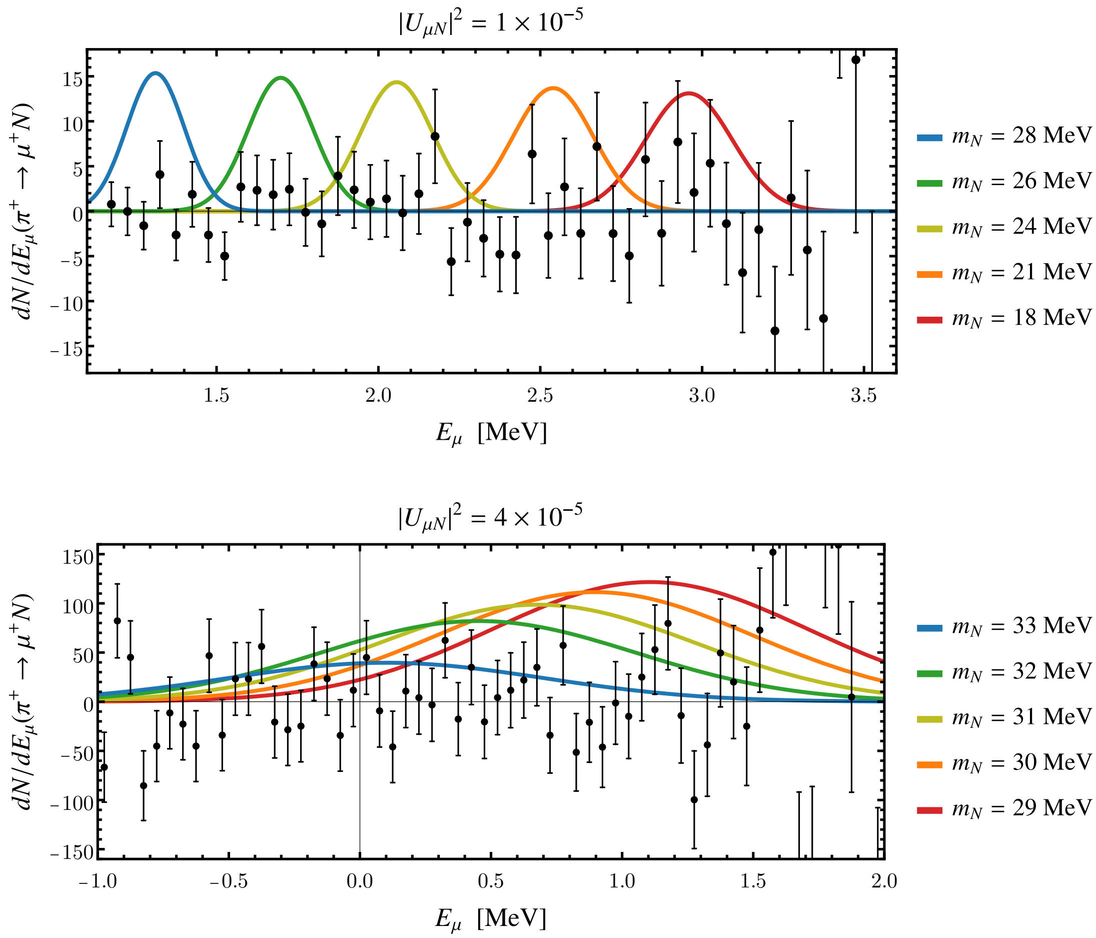

Figure B.2: Expected number of detected $\pi^+ \to \mu^+ N$ decays at the PIENU experiment as a function of the muon kinetic energy for several values of the sterile neutrino mass. The top shows the high muon energy analysis, the bottom the low muon energy analysis. The sterile neutrino mixing angles are set to $|U_{\mu N}|^2 = 1 \times 10^{-5}$ and $|U_{\mu N}|^2 = 4 \times 10^{-5}$, respectively. For comparison, the residuals extracted from [135] are shown in black.

C Details on the PIENU Three-Body Pion Decay Recast

In this appendix, we describe our reproduction of the PIENU search [137] for a dark sector state X in three-body pion decays, focusing first on the $\pi^+ \to e^+\nu_e X$ channel, and subsequently addressing the muon mode $\pi^+ \to \mu^+\nu_\mu X$.

We extracted the experimental data from the two data sets used for the $\pi^+ \to e^+\nu_e X$ search and presented in the plots of [137], comprising approximately 5×10^5 and 8×10^5 $\pi^+ \to e^+\nu_e$ events, respectively. To model the data, PIENU included background components from the low-energy tail of the Standard Model $\pi^+ \to e^+\nu_e$ decay, as well as from muon decays at rest and muon decays in flight.

We convoluted the signal spectrum from [116] with a signal spread function to account for detector resolution effects. Specifically, we assumed that the signal spread at the positron energy of $E_e = 69.8\,\mathrm{MeV}$ (corresponding to the peak of the SM $\pi^+ \to e^+\nu_e$ decay) is described by the observed $\pi^+ \to e^+\nu_e$ line shape. For other positron energies, the spread was scaled proportionally to the square root of the positron energy, similar to our approach in reproducing the PIENU sterile neutrino search described in appendix B. We find that this convolution procedure slightly broadens the signal spectrum and induces a modest downward shift of the peak position by a few MeV. As noted in [137],

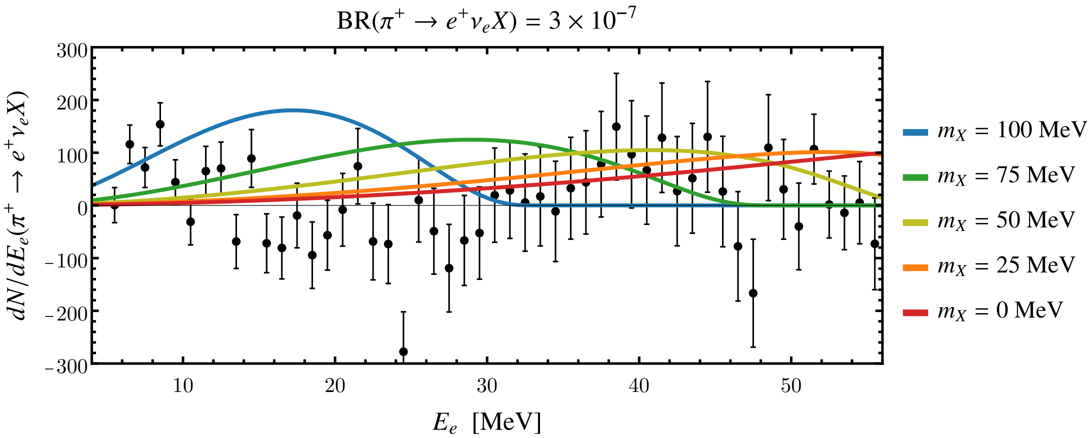

Figure C.1: Expected number of detected $\pi^+ \to e^+\nu_e X$ decays at the PIENU experiment as a function of the positron energy for several values of the X mass. The signal model is taken from [116]. The $\pi^+ \to e^+\nu_e X$ branching ratio is set to 3×10^{-7}. For comparison, the residuals extracted from [137] are shown in black.

no energy-dependent acceptance corrections were applied in the analysis, and we follow the same approach in our reproduction.

We find excellent agreement between our reconstructed signal spectra and those presented in [137] for a dark sector mass of $m_X = 80\,$MeV. As an illustration, figure C.1 shows our signal spectra for various X masses, assuming a fixed branching ratio BR$(\pi^+ \to e^+\nu_e X) = 3 \times 10^{-7}$. The event numbers and shown residuals correspond to the sum of the two data sets from [137].

To extract the branching ratio limits, we performed fits of the signal spectrum to the residuals in the positron energy range $4.5\,$MeV $< E_e < 56\,$MeV, including the three background components with floating normalizations. The 90% C.L. limits on the branching ratio for the $\pi^+ \to e^+\nu_e X$ decay are shown in figure C.2 by the solid purple curve. They are in good agreement with those reported by the PIENU collaboration shown by the dashed curve, validating our reproduction of the analysis.

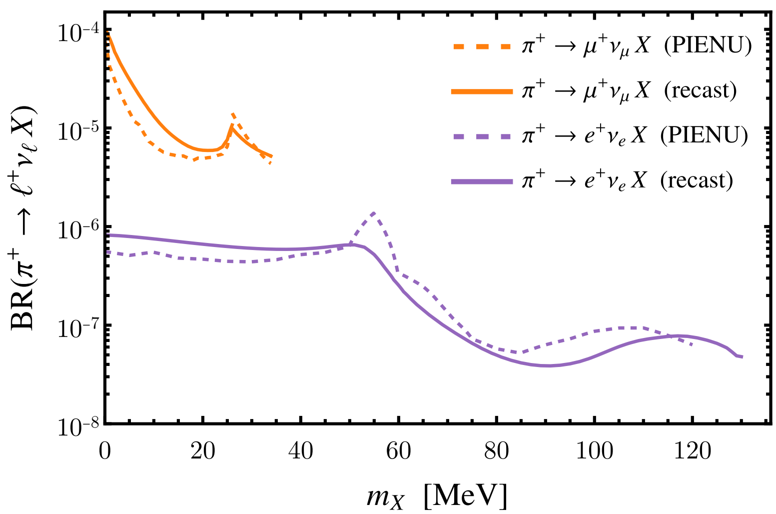

Figure C.2: Upper limits on the $\pi^+ \to e^+\nu_e X$ (purple) and $\pi^+ \to \mu^+\nu_\mu X$ (orange) branching ratios as a function of the X mass for the signal model from [116]. The PIENU result from [137] is shown in dashed. Our recast is shown by the solid lines.

Our reproduction of the $\pi^+ \to \mu^+\nu_\mu X$ search by PIENU [137] follows a similar procedure. Analogous to the search for sterile neutrinos in $\pi^+ \to \mu^+ N$ decays (described in appendix B), the $\pi^+ \to \mu^+\nu_\mu X$ search is split into two muon energy regions. At low energies, the data contain approximately 1.3×10^8 muon events and it is well described by the SM $\pi^+ \to \mu^+\nu_\mu$ peak and a second order polynomial in E_μ that models contributions from pion decay in flight. The data in the high energy region is comprised of approximately 9.1×10^6 muon events, and it is well described by the SM peak and a tail from radiative decays $\pi^+ \to \mu^+\nu_\mu\gamma$ with a shape obtained from Monte Carlo simulations.

To model the signal, we convolute the $\pi^+ \to \mu^+\nu_\mu X$ spectrum with a Gaussian spread function. Analogous to the sterile neutrino search, in the low energy analysis, we fix the width of the Gaussian to $0.6\,\mathrm{MeV}$, while in the high energy analysis, we use a width that scales with the square root of the muon kinetic energy and is $0.16\,\mathrm{MeV}$ for $E_\mu = 4.1\,\mathrm{MeV}$. We find good agreement between our final signal shapes and the shapes shown in [137] for X masses of $15\,\mathrm{MeV}$ and $33.9\,\mathrm{MeV}$. For illustration we show in figure C.3 our signal spectra for several choices of the X mass and compare them to the residuals of the PIENU low energy and high energy analysis [137].

We then fit the signal spectra to the residuals. In the low energy analysis, we also vary all three parameters of the polynomial that models the pion decay in flight, while for the high energy analysis we vary the normalization of the radiative decay background, keeping its shape fixed. We show our 90% C.L. limits on the $\pi^+ \to \mu^+ \nu_\mu X$ branching ratio in figure C.2 by the orange lines. Our recast (solid) agrees very well with the official limit reported by PIENU (dashed).

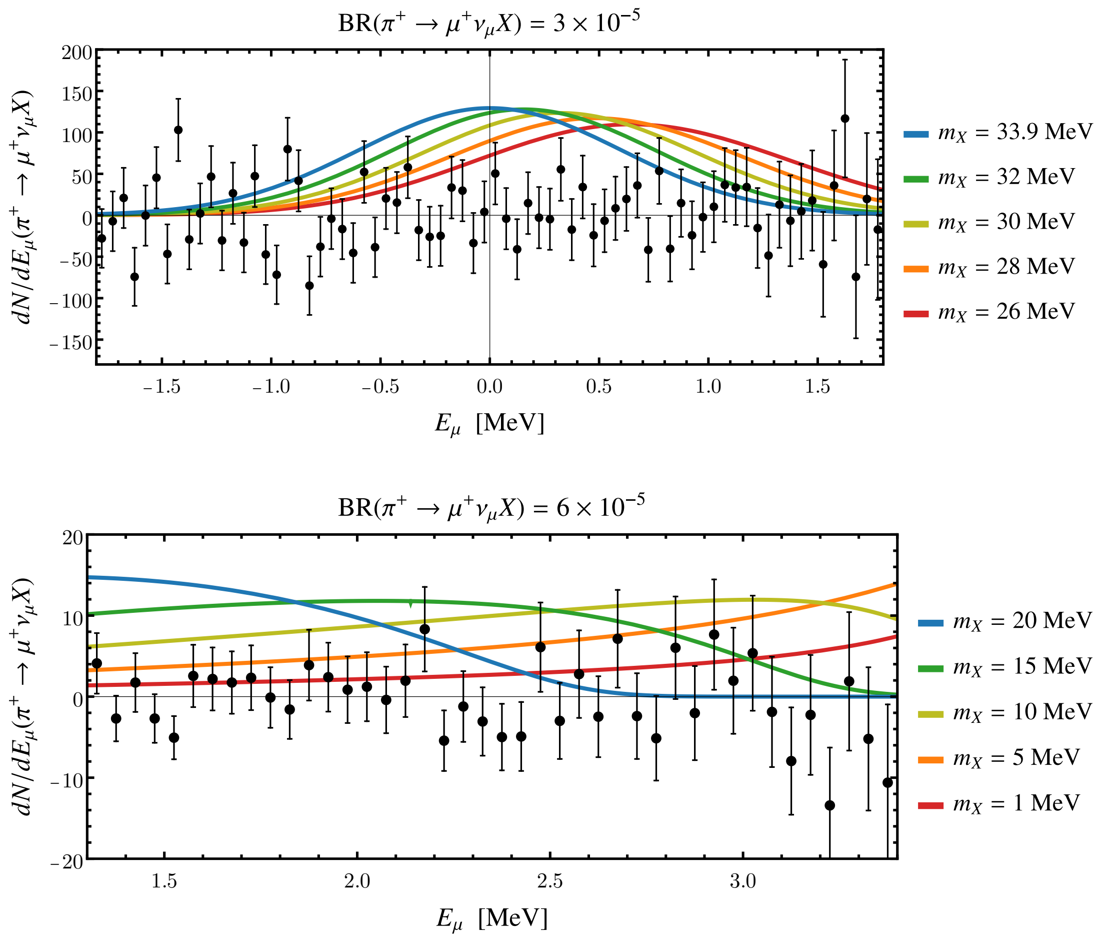

Figure C.3: Expected number of detected $\pi^+ \to \mu^+\nu_\mu X$ decays at the PIENU experiment as a function of the muon kinetic energy for several values of the X mass. The signal model is taken from [116]. The top and bottom plots show the low and high energy analysis with the $\pi^+ \to \mu^+\nu_\nu X$ branching ratio set to 3×10^{-5} and 6×10^{-5}, respectively. For comparison, the residuals extracted from [137] are shown in black.

D Plasma review: Linear and non-linear regimes

Plasmas are quasineutral mixtures of charged particles in which oppositely charged particles are not bound to each other. When particle collisions can be neglected, the evolution of the plasma is governed by the Vlasov-Maxwell equations [228],

$$\frac{\partial f_s}{\partial t} + \vec{v} \cdot \nabla f_s + \frac{q}{m}\left[\vec{E} + \vec{v} \times \vec{B}\right] \cdot \nabla_v f_s = 0 \tag{D.1}$$

$$\nabla \cdot \vec{E} = \rho; \qquad \nabla \cdot \vec{B} = 0;$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}; \qquad \nabla \times \vec{B} = \vec{J} + \frac{\partial \vec{B}}{\partial t}$$

where f_i represents the number density of each species of particle in phase space,

$$\rho = \sum_s \int q_s f_s(\vec{x}, \vec{v}, t) d^3 v, \tag{D.2}$$

and

$$\vec{J} = \sum_s \int q_s \vec{v} f_s(\vec{x}, \vec{v}, t) d^3 v. \tag{D.3}$$

This system can be recast in a form in which the only unitful quantity is the *plasma* *frequency*, defined as

$$\omega_{\mathrm{pl},s} = \sqrt{\frac{q_s^2 n_s}{m_s}} \tag{D.4}$$

where n is the number density and s indexes species. Here, we use units in which $\hbar = c = \epsilon_0 = \mu_0 = 1$. This allows the behavior of plasma at one set of (q_s, m_s) parameters to be generalized to other regions in parameter space at which the plasma takes on the same $\omega_{\mathrm{pl},s}$.

The Vlasov-Maxwell system is inherently nonlinear; however, when perturbations are small, the equation can be linearized and a dispersion equation computed. When a solution to the dispersion relation contains a positive imaginary part, small perturbations grow exponentially, giving rise to an instability. The growth rate of instabilities in the linear regime can be computed analytically in certain simplified settings; however, once the perturbation has grown sufficiently large that the linear approximation breaks down, the dynamics enter the *nonlinear regime*, and numerical simulations are needed to characterize the subsequent evolution. In this regime, the exponential growth back-reacts on the background plasma, leading to a saturation of the instability and, often, the formation of non-linear structures such as collisionless shocks [221–226].

In the following subsections, we will focus on the linear and nonlinear behavior of one astrophysically-relevant system, namely counter-streaming non-magnetized pair plasmas.

§ D.1 Linear regime

A simple system in which instabilities arise consists of two beams of plasma with comparable densities streaming through one another. Though simple, this system captures the essential features of the astrophysical settings we discuss in Section III of the main text. Two primary instabilities dominate this system. The first is the *two-stream instability* [236], in which longitudinal modes grow exponentially as a result of electrostatic interactions. The second is the *Weibel instability* [237], in which modes transverse to the beam's direction of motion are amplified by electromagnetic interactions with the other beam.

D.1.1 Two-stream instability

The electrostatic instability arises due to the presence of an unstable longitudinal mode in the counter-streaming plasmas. The growth rate of perturbations can be calculated in the linear regime via the plasma dispersion relation [236]

$$\Gamma_{\rm TS} \approx \frac{1}{2\sqrt{2}}\omega_{\rm pl} \tag{D.5}$$

with the fastest growing mode having $k_{x,\rm TS}^{\rm crit} \approx \frac{\sqrt{3/2}}{v_0}\omega_{\rm pl}$, where the subscript TS denotes "two-stream" and v_0 is the initial relative bulk velocity of the plasma beams. This growth rate has been computed under the assumption that the relative velocity of the beams (v_0) is significantly larger than their velocity dispersions (the "cold limit"). The inclusion of a non-zero temperature only serves to weaken the instability and suppress the growth rate.

The instability saturates when the potential becomes large enough that approaching particles become trapped in the troughs between regions of high potential. This increases the charge density in the trough, raising its potential and reducing the amplitude of the mode.

D.1.2 Weibel instability

While the two-stream instability is driven entirely by interactions with the electric field, hence is "electrostatic," there is an additional electromagnetic instability in the system known as the Weibel instability.* The growth rate of perturbations can be calculated in the linear regime and is given by [237]

$$\Gamma_{\rm W} \approx v_0 \omega_{\rm pl} \tag{D.6}$$

with the instability approaching its maximum for modes with $k \gtrsim k_{y,\rm W}^{\rm crit} \approx \omega_{\rm pl}$ [224] in the zero-temperature limit. Note that the Weibel instability grows at a rate suppressed by v_0 with respect to the two-stream instability, as is expected given that the two-stream instability arises due to the $q\mathbf{E}$ term in the Vlasov equation while the Weibel instability arises due to the $q\mathbf{v} \times \mathbf{B}$ term.

While the Weibel instability grows more slowly than the two-stream instability at non-relativistic velocities, it still has a significant effect on late-time dynamics. The instability saturates when the magnetic field becomes sufficiently strong such that the gyroradius of particles becomes comparable to the spatial scale of the perturbation. At this stage, deflected particles can no longer penetrate through the magnetic filaments formed during instability growth, and the instability mechanism is halted. Saturation of the Weibel instability generally occurs at a higher electromagnetic field density than the two-stream instability hence, at late times, it is the nonlinear behavior of this instability that dominantly drives the evolution of the plasma.

*Additional electromagnetic oblique modes can occur in asymmetric beam plasmas. For beams that are approximately symmetric, as in this work, their growth rates and saturation mechanisms only vary slightly and can be determined numerically. See [223] for further discussion.

§ D.2 Nonlinear regime

The growth of the two-stream instability will cause the formation of "collisionless shocks" [221, 223, 224, 227–229, 325, 326], shock fronts between the beams supported by collective electromagnetic interactions despite the mean-free-path of an individual particle being much longer than the scale of the shock. The formation of this structure halts when the exponential growth rate becomes comparable to the bounce frequency of the streaming particles [223]

$$\Gamma_{TS}^2 \approx \omega_{b,E}^2 = \frac{q_\chi E k_x}{m_\chi}. \tag{D.7}$$

This process interrupts the bulk motion of the counter-steaming plasma, slowing down the relative bulk velocity while heating up the plasma [227].

While electrostatic modes initially generate collisionless shocks before dissipating quickly, Weibel filaments will continue to grow and support these shocks on much longer timescales, further contributing to the slowdown of the beams. The Weibel instability saturates when the exponential growth rate becomes comparable to the magnetic trapping frequency [223]

$$\Gamma_W^2 \approx \omega_{b,B}^2 = \frac{q_\chi B v_0 k_y}{m_\chi} \tag{D.8}$$

which in general corresponds to a larger electromagnetic field strength than at two-stream saturation.

After saturation, the Weibel filaments will undergo macroscopic instabilities which lead to further turbulence and slowdown. See Appendix (see also references [223, 230–235, 327] therein) for further discussion.

E Simulations

§ E.1 Setup

Our simulations consist of a 2D3V Cartesian geometry of dimensions $L_x = 7\, c/\omega_{pl,B}$ in the longitudinal direction and $L_y = 175\, c/\omega_{pl,B}$ in the transverse direction with periodic boundary conditions across all boundaries. Here, $\omega_{pl,B}$ is the plasma frequency corresponding to the Bullet sub-cluster. We choose a grid spacing of $\Delta x = 0.0035\, c/\omega_{pl,B}$ in the longitudinal direction and $\Delta y = 0.18\, c/\omega_{pl,B}$ in the transverse direction to ensure proper resolutions of both two-stream and Weibel instabilities. Each simulation was run for $10^3\, \omega_{pl,B}^{-1}$ with a temporal resolution of $\Delta t = 0.98 \Delta x$. Two species were initialized, χ and $\bar{\chi}$, with equal and opposite charge-to-mass ratios, 64 total macroparticles per cell were created with second order interpolation schemes for the particle shape function, and bulk velocities $v_{\text{Bullet}} = -n_{\text{Main}} v_0/(n_{\text{Bullet}} + n_{\text{Main}})$ and $v_{Main} = n_{\text{Bullet}} v_0/(n_{\text{Bullet}} + n_{\text{Main}})$ such that the simulation is in the center of momentum frame. Additionally, five multi-pass binomial current filters [328] were applied in each dimension after each time step for increased numerical stability. We found this setup to properly resolve all of the relevant length and time scales and produced results with negligible numerical heating.

We undertook several studies to ensure the numerical stability of our simulation. In

Fig. E.2, we show the amount of energy increase from numerical heating, a known numerical effect arising from truncation errors and finite-sized time and space domains [230–232]. Throughout our simulation, this remained below 0.1% of the total initial kinetic and electromagnetic energy, indicating stability.

Additionally, we compared the length scales of the two-stream and Weibel instabilities in the asymmetric warm case to ensure that our choice of cell size accurately resolved the relevant instability in each direction. Fig. E.3 displays the exponential growth rate for the two-stream instability for the Bullet Cluster system. We find that the maximal growth rate is $\Gamma_{TS} = 0.397\,\omega_{pl,B}$ at $k_x = 152\,\omega_{pl,B}/c$. This corresponds to a wavelength of $\lambda_{TS} = 0.04\,c/\omega_{pl,B}$, which is well above our resolution in the x direction of $\Delta x = 0.0035\,c/\omega_{pl,B}$.

Similarly, in the transverse direction, Fig. E.4 shows the exponential growth rate of the Weibel instability. For the Bullet Cluster system, the instability modes reach a maximum of $0.00995\,\omega_{pl,B}$ and display a rapid drop-off for $k < k_{y,\mathrm{W}}^{\mathrm{crit}} \sim 1\,\omega_{pl,B}/c$. This corresponds to a maximal wavelength of $\lambda_W = 2\pi\,c/\omega_{pl,B}$, which is also well above our resolution in the y direction of $\Delta y = 0.18\,c/\omega_{pl,B}$.

§ E.2 Growth rates

Fig. E.1 shows the energy densities of the longitudinal electric field (red), the transverse magnetic field (blue), and total electromagnetic energy density (purple) from our simulation normalized to the initial kinetic energy of the two-beam system as a function of time. As discussed in Section D.1.1, at early times, we expect the two-stream instability to dominate; in Fig. E.1, we see that the energy density of the longitudinal electric fields

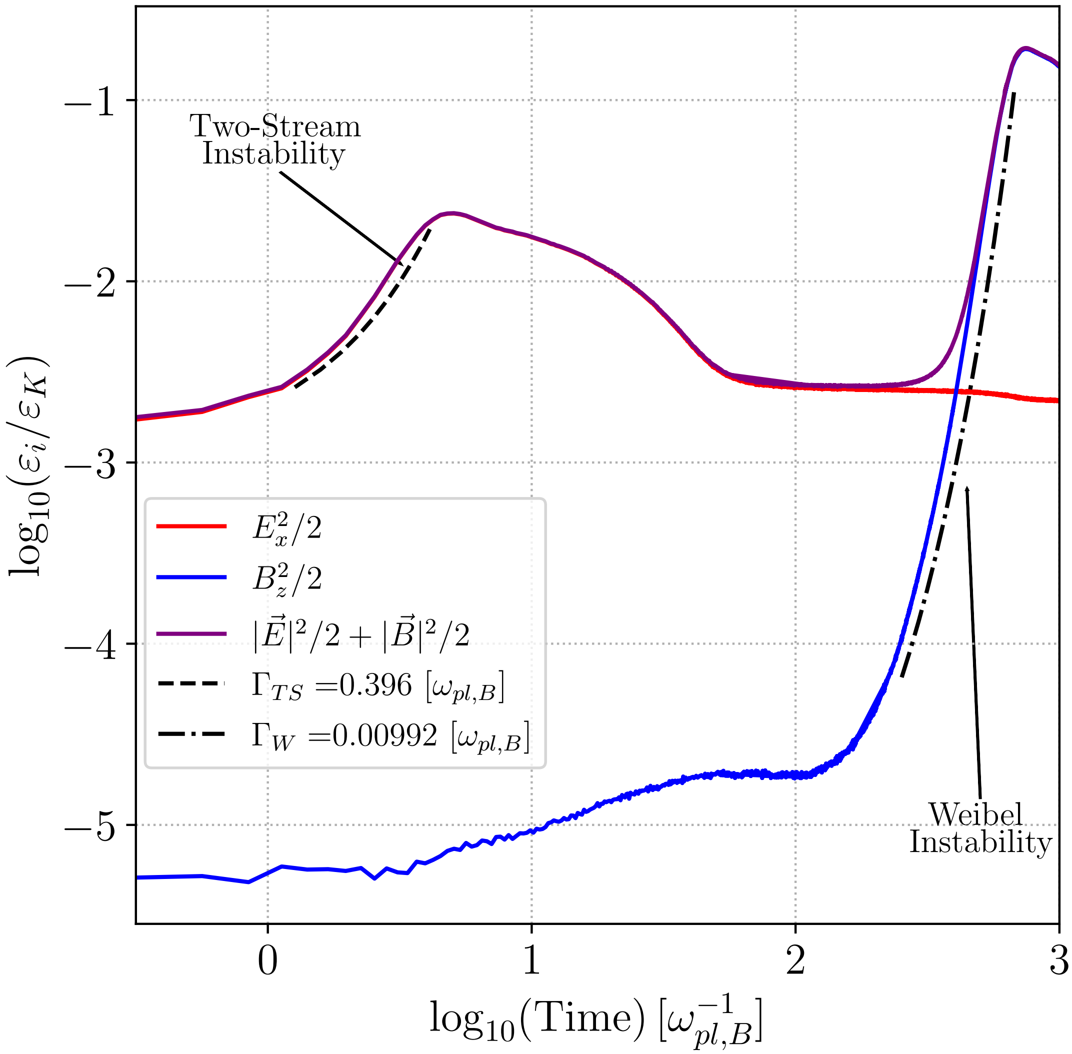

Figure E.1: Energy densities from simulation normalized to the initial kinetic energy of the two-beam system. The energy stored in the longitudinal electric fields (red), the transverse magnetic field (blue), and total electromagnetic energy (purple) are shown. Fits to the exponential growth associated with the two-stream and Weibel instabilities are shown as dashed and dashed-dotted lines respectively.

grows exponentially with a growth rate of $\Gamma_{TS} = 0.396\,\omega_{pl,B}$, but saturates rapidly. This growth rate agrees well with numerical estimates. After the electrostatic instability saturates, the Weibel instability continues to grow at $0.00992\,\omega_{pl,B}$, then saturates at $\sim 750\,\omega_{pl,B}^{-1}$ (Fig. E.1).

§ E.3 Saturation energies

From the growth rate in the linear regime, we calculate the approximate saturation energy of the two-stream instability, which yields

$$\frac{\varepsilon_E}{\varepsilon_K} = \frac{E_{\text{sat}}^2/2}{\frac{1}{2}(4n_0)m_\chi(v_0/2)^2} \approx \frac{1}{96}, \tag{E.1}$$

where ε_K is the initial kinetic energy density of the plasma. Similarly, we determine the saturation energy for the Weibel instability, which yields

$$\frac{\varepsilon_B}{\varepsilon_K} = \frac{B_{\text{sat}}^2/2}{\frac{1}{2}(4n_0)m_\chi(v_0/2)^2} \approx \left(\frac{v_W}{v_0}\right)^2, \tag{E.2}$$

where v_W is the relative bulk velocity of the plasma during saturation. This value is often less than the initial relative bulk velocity, v_0, due to the decrease in bulk kinetic energy during the formation of the electrostatic shock by the two-stream instability. At regular intervals we sample the velocity distribution of all particles and compute the mean velocity. We find $v_W = 0.0014\,c$. We explicitly measure the saturation levels of the longitudinal electric field and the transverse magnetic field in our simulation in order to compare to these analytic predictions, finding $\varepsilon_E/\varepsilon_K = 0.024$ and $\varepsilon_B/\varepsilon_K = 0.19$, respectively. We find these values to be within reasonable agreement with Eqs. E.1 and Eq. E.2.

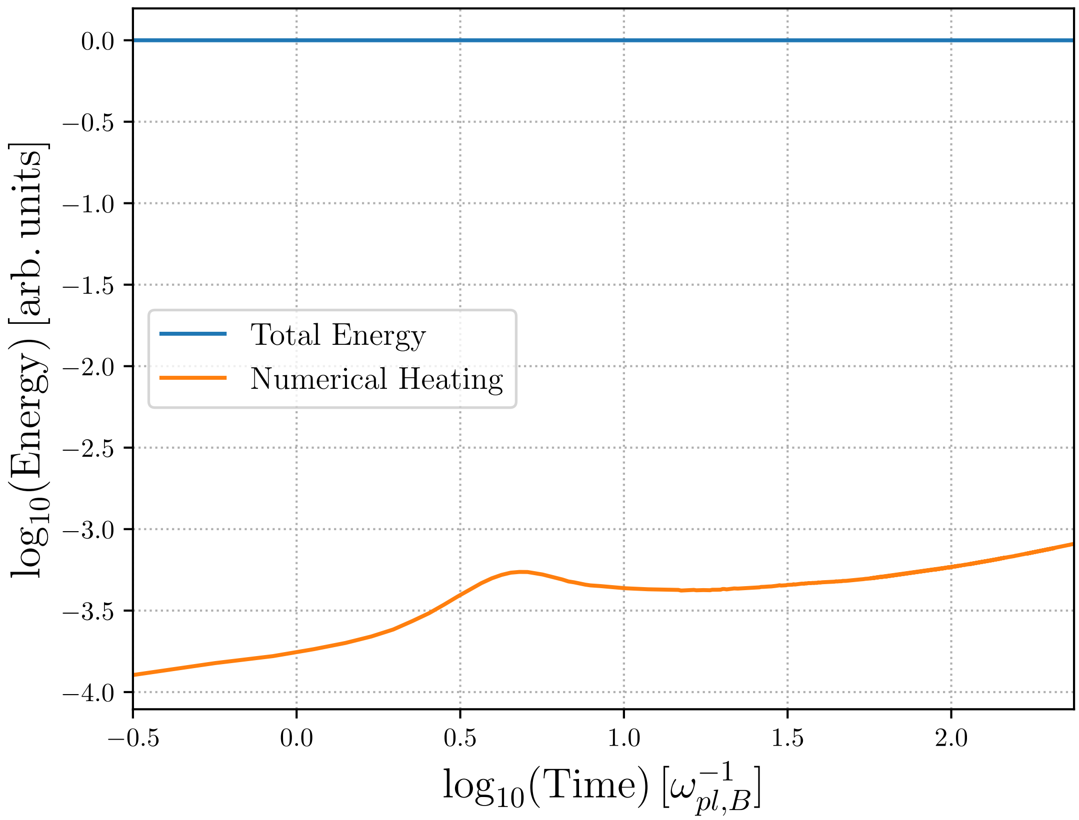

Figure E.2: Combined particle and field energy from simulation (blue) and energy lost due to numerical heating (orange).

§ E.4 Additional diagnostic plots

Here, we show some additional diagnostic plots from our simulation to highlight some of the key features of the evolution of the two-stream and Weibel instabilities.

In Fig. E.5 (E.6), we show a plot of the (Fourier transform of the) electric field along the direction of beam propagation as a function of position for three different times. In the left-most panels, corresponding to $t = 0.3\,\omega_{pl,B}^{-1}$, the simulation has properly initialized, and all electromagnetic fields are relatively small and are generated purely by thermal noise effects. At $t = 5\,\omega_{pl,B}^{-1}$ (center panels), the two-stream instability has

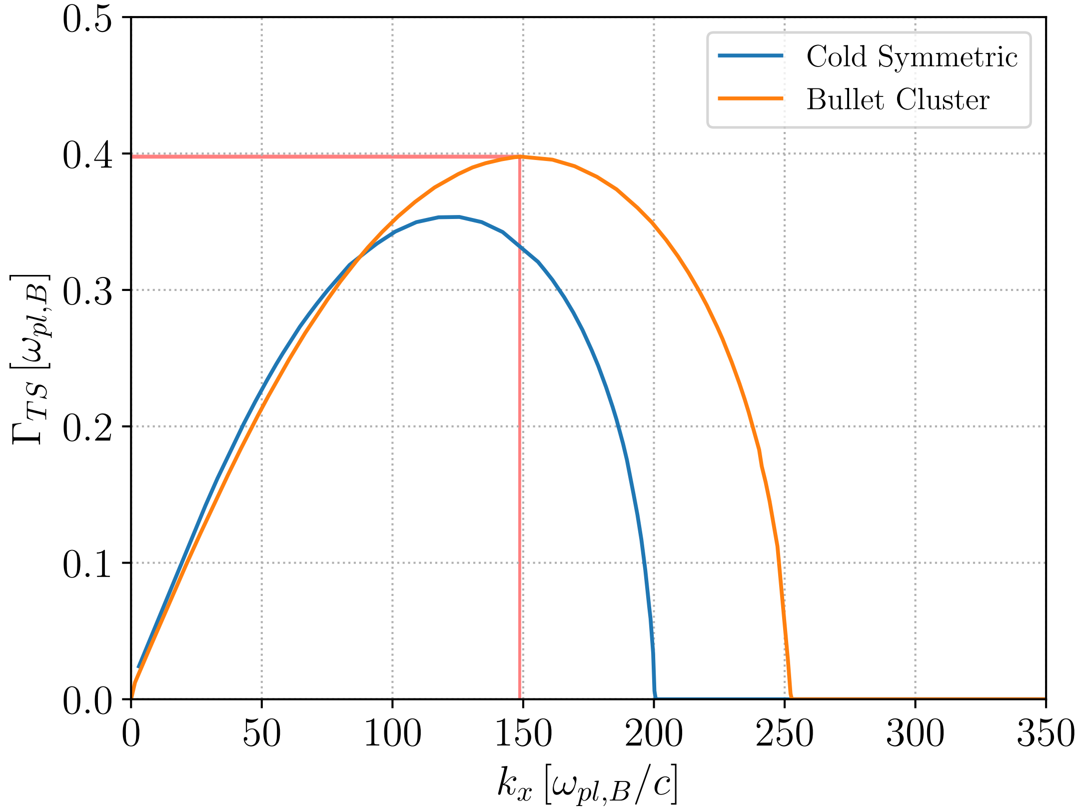

Figure E.3: Exponential growth rate of longitudinal two-stream instability in the case of cold symmetric beams (blue) and for the parameters of the Bullet Cluster system (orange). The red lines show the fastest growing mode and corresponding instability rate.

reached saturation, as is indicated by the presence of small vertical strips alternating polarity corresponding to the fastest growing longitudinal mode in Fig. E.5 and the sharp peak around the critical k_x mode in Fig. E.6. At $t = 200\,\omega_{pl,B}^{-1}$ (right-most panels), the two-stream instability has dissipated.

In Fig. E.7 (E.8), we show a plot of the (Fourier transform of the) magnetic field in the transverse direction as a function of position at the same three times as above. However, while the two-stream instability has saturated by $t = 750\,\omega_{pl,B}^{-1}$, Fig. E.7 shows that the Weibel filaments are beginning to form and are nearing saturation at this time. Fig. E.8 shows that the transverse magnetic filaments peak around the critical k_y mode with increasing intensity as the fields reach saturation.

§ E.5 Limitations of 2D3V simulations

Though our simulation is 2-dimensional, it is still able to conservatively estimate the macrophysical properties of the 3-dimensional Bullet Cluster system. First, we note

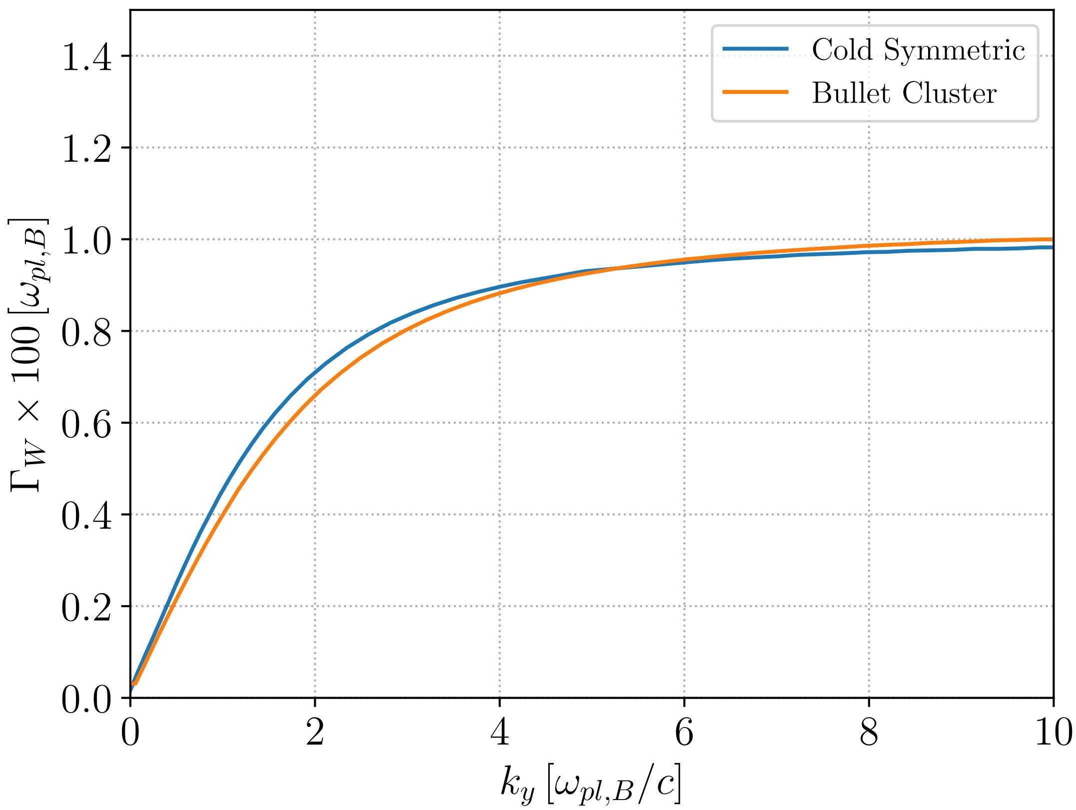

Figure E.4: Exponential growth rate of transverse Weibel instability in the case of cold symmetric beams (blue) and for the parameters of the Bullet Cluster system (orange).

that the Weibel instability has previously been studied in 1 and 2-dimensional systems [233, 327], and it has been shown that the growth rates and magnetic field saturation levels agree with analytical models. After Weibel saturation, we expect our simulation to no longer capture the relevant dynamics, as at this time, the Weibel filaments should enter a late-stage merging process that requires the current filaments to move around in 3-dimensional space. Though we cannot capture this effect, it has been previously studied in the literature [223, 234], and the magnetic field energy is expected to grow approximately linearly with time. These filaments will eventually become unstable to the z-pinch instability, which generally works to slow down the flow of particles and isotropize the magnetic field [235]. Though these late-time phenomena likely only further slow the beams, we chose to set our limit prior to Weibel saturation, long before these merging events are expected to begin, in an effort to be as conservative as possible. Furthermore, the correspondence of our results with existing literature that simulates both 2D3V and 3D3V systems (see Section IV of main text) provides strong evidence that our results are trustworthy up through Weibel saturation.

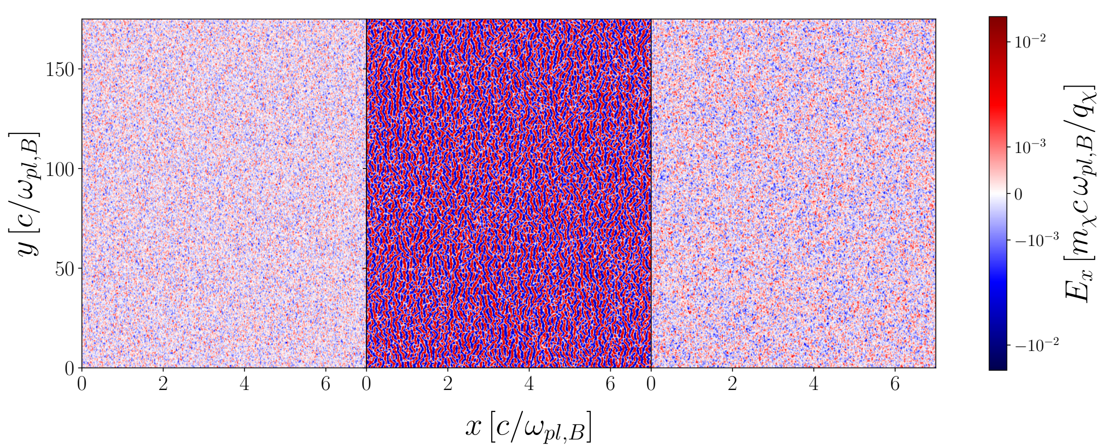

Figure E.5: Longitudinal component of the electric field at $t = 0.3\,\omega_{pl,B}^{-1}$ (left), $t = 5\,\omega_{pl,B}^{-1}$ (middle), and $t = 750\,\omega_{pl,B}^{-1}$ (right).

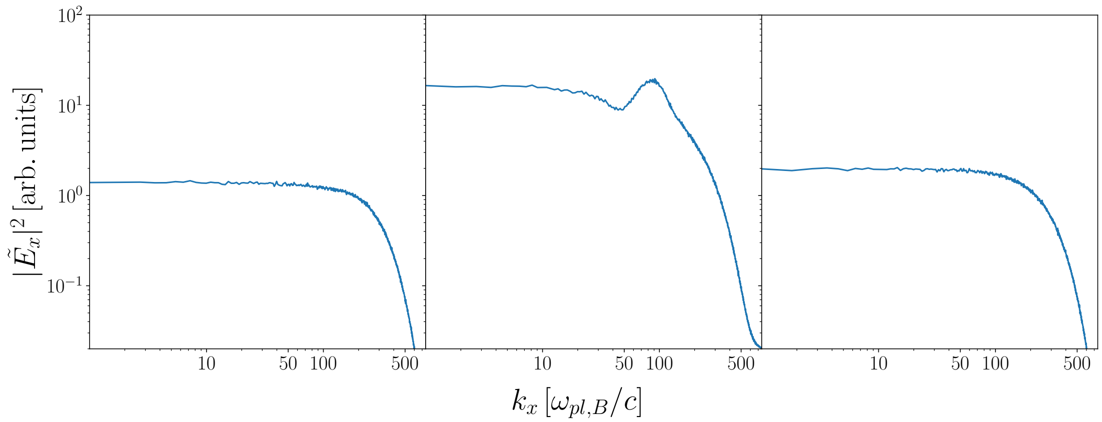

Figure E.6: Power spectrum of the longitudinal electric field at $t = 0.3\,\omega_{pl,B}^{-1}$ (left), $t = 5\,\omega_{pl,B}^{-1}$ (middle), and $t = 750\,\omega_{pl,B}^{-1}$ (right).

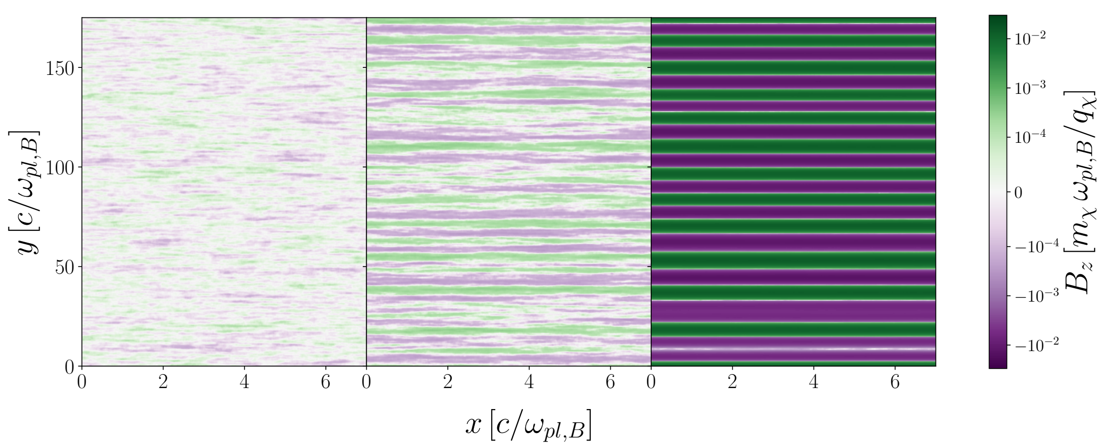

Figure E.7: Transverse component of the magnetic field at $t = 0.3\,\omega_{pl,B}^{-1}$ (left), $t = 200\,\omega_{pl,B}^{-1}$ (middle), and $t = 750\,\omega_{pl,B}^{-1}$ (right).

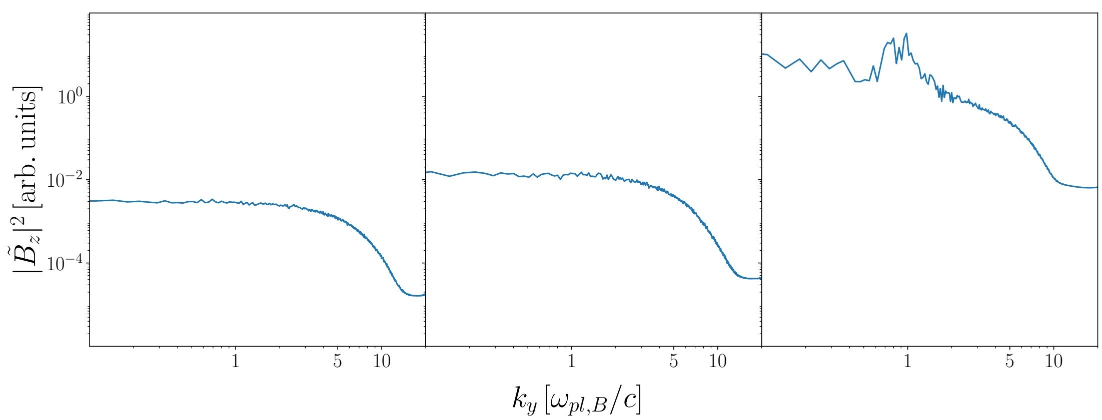

Figure E.8: Power spectrum of the transverse magnetic field at $t = 0.3\,\omega_{pl,B}^{-1}$ (left), $t = 200\,\omega_{pl,B}^{-1}$ (middle), and $t = 750\,\omega_{pl,B}^{-1}$ (right).

F Derivation of Generalized Ohm's Law

We now aim to justify our choice of closure of the typical MHD system of equations through the generalized Ohm's law. The equations of motion of the positive and negative species of particles can be described by the following differential equations

$$\partial_t(\rho_+\vec{U}_+) + \nabla\cdot(\rho_+\vec{U}_+\vec{U}_+) = \rho_+\frac{q_\chi}{m_\chi}\left(\vec{E}+\vec{U}_+\times\vec{B}\right) \\ - \nabla\cdot \overset{\leftrightarrow}{P}_+ - \rho_+\nu_\pm(\vec{U}_+ - \vec{U}_-), \tag{F.1}$$

$$\partial_t(\rho_-\vec{U}_-) + \nabla\cdot(\rho_-\vec{U}_-\vec{U}_-) = -\rho_-\frac{q_\chi}{m_\chi}\left(\vec{E}+\vec{U}_-\times\vec{B}\right) \\ - \nabla\cdot \overset{\leftrightarrow}{P}_- - \rho_-\nu_\pm(\vec{U}_- - \vec{U}_+), \tag{F.2}$$

where $\rho_{+(-)}$ is the mass density of the positive (negative) species, $\vec{U}_{+(-)}$ is the average velocity of the positive (negative) species, and $\nu_\pm$ is the frequency of Coulomb scatterings between the positive and negative species. Taking the difference between these two equations we arrive at the generalized Ohm's law for pair plasmas.

$$\vec{E}+\vec{U}\times\vec{B} = \eta\vec{J} + \frac{1}{2}\frac{m_\chi^2}{q_\chi^2\rho}\left[\partial_t\vec{J} + \nabla\cdot(\vec{U}\vec{J}+\vec{J}\vec{U})\right] + \frac{m_\chi}{q_\chi\rho}\nabla\cdot\left(\overset{\leftrightarrow}{P}_+ - \overset{\leftrightarrow}{P}_-\right) \tag{F.3}$$

where we have now defined the fluid mass density $\rho = \rho_+ + \rho_-$, the fluid bulk velocity

$\vec{U} = (\vec{U}_+ + \vec{U}_-)/2$, and the current $\vec{J} = q_\chi \rho(\vec{U}_+ - \vec{U}_-)/m_\chi$ as well as the resistivity of the plasma

$$\eta = \frac{m_\chi^2 \nu_\pm}{q_\chi^2 \rho}. \tag{F.4}$$

We now consider the scenario under study. Given that we expect the positively changed species to be in thermal equilibrium with the negatively charged species, $\overleftrightarrow{P}_- = \overleftrightarrow{P}_+$ and the last term in Eq. (F.3) vanishes. Next, we assume that the unperturbed plasma initially has no net current or bulk velocity. Hence, after linearizing Eq. (F.3), $\nabla \cdot (\vec{U}\vec{J} + \vec{J}\vec{U})$ will only contain second order terms which we neglect in this work. Lastly, due to the strong constraints already placed on the dark charge-to-mass ratio by [**?**], we expect $\nu_\pm \ll T^{-1}$, where T is the timescale of gravitational collapse of the plasma. Hence, the resistive term of Eq. (F.3) is negligible compared to the $\partial_t \vec{J}$ term. Thus, we chose to implement the following Ohm's law for our study

$$\vec{E} + \vec{U} \times \vec{B} = \frac{1}{2} \frac{m_\chi^2}{q_\chi^2 \rho} \partial_t \vec{J} \tag{F.5}$$

Note that in the limit of $q_\chi/m_\chi \to \infty$ we obtain the Ohms law for an ideal MHD plasma

$$\vec{E} + \vec{U} \times \vec{B} = 0. \tag{F.6}$$

G Derivation of Dispersion Relation

In this Appendix, we explicitly show the full derivation of the dispersion relation in Eq. (5.10). These calculations follow very similarly to those performed in Ref. [329]. However, we consider here effects from nonideal MHD conditions. In the presence of an external magnetic field, a self-gravitating plasma obeys the following system equations

$$\nabla \cdot \vec{E} = \rho_c \qquad \nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{G.1}$$

$$\nabla \cdot \vec{B} = 0 \qquad \nabla \times \vec{B} = \vec{J} + \frac{\partial \vec{E}}{\partial t} \tag{G.2}$$

$$\frac{\partial \rho_m}{\partial t} + \nabla \cdot (\rho_m \vec{U}) = 0 \qquad \rho_m \frac{d\vec{U}}{dt} = \vec{J} \times \vec{B} - \nabla \overset{\leftrightarrow}{P} - \rho_m \nabla V \tag{G.3}$$

$$\frac{d}{dt}\left(\frac{p_\perp B^2}{\rho_m^3}\right) = 0 \qquad \frac{d}{dt}\left(\frac{p_\parallel}{\rho_m B}\right) = 0 \tag{G.4}$$

$$\vec{E} + \vec{U} \times \vec{B} = \frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho_m}\frac{\partial \vec{J}}{\partial t} \qquad \nabla^2 V = 4\pi G \rho_m \tag{G.5}$$

Here we have utilized our choice of Ohm's Law discussed in Appendix F as well as the Chew-Goldberger-Low equation of state for cylindrically symmetric magnetized plasmas [272]. As in typical MHD studies, we assume that the plasma remains net neutral during its evolution and neglect effects due to separation of charges by setting $\rho_c = 0$. Additionally, we assume that the system is varying slowly enough in time such that displacement current $\partial \vec{E}/\partial t$, can be neglected when compared to $\vec{J}$. This allows $\vec{J}$

to be eliminated from the system of equations

$$\nabla \cdot \vec{E} = 0, \qquad \nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}, \tag{G.6}$$

$$\nabla \cdot \vec{B} = 0, \qquad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{U}) = 0, \tag{G.7}$$

$$\rho \frac{d\vec{U}}{dt} = (\nabla \times \vec{B}) \times \vec{B} - \nabla \overleftrightarrow{P} - \rho \nabla V, \tag{G.8}$$

$$\frac{d}{dt}\left(\frac{p_\perp B^2}{\rho^3}\right) = 0, \qquad \frac{d}{dt}\left(\frac{p_\parallel}{\rho B}\right) = 0, \tag{G.9}$$

$$\vec{E} + \vec{U} \times \vec{B} = \frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho}\frac{\partial}{\partial t}\nabla \times \vec{B}, \qquad \nabla^2 V = 4\pi G \rho, \tag{G.10}$$

Here we have defined $\rho = \rho_m$ for ease of notation. It is important to note that $\nabla \cdot \vec{E} = 0$ and $\nabla \cdot \vec{B} = 0$ only need to be satisfied in the initial conditions. Ensuring the remaining Maxwell's equations are satisfied during the time evolution of the plasma ensures that the initial conditions remain satisfied. As a homogeneous plasma in a homogeneous magnetic field satisfies these conditions, we no longer require them. Further eliminating the electric field from the system yields

$$\nabla \times \left(\vec{U} \times \vec{B} - \frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho}\frac{\partial}{\partial t}\nabla \times \vec{B}\right) = \frac{\partial \vec{B}}{\partial t} \tag{G.11}$$

$$\nabla^2 V = 4\pi G \rho, \qquad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{U}) = 0, \tag{G.12}$$

$$\rho \frac{d\vec{U}}{dt} = (\nabla \times \vec{B}) \times \vec{B} - \nabla \overleftrightarrow{P} - \rho \nabla V, \tag{G.13}$$

$$\frac{d}{dt}\left(\frac{p_\perp B^2}{\rho^3}\right) = 0, \qquad \frac{d}{dt}\left(\frac{p_\parallel}{\rho B}\right) = 0. \tag{G.14}$$

Note that in the $q_\chi / m_\chi \to \infty$ limit, we obtain the same system of equations studied under ideal MHD conditions in Ref. [329].

We now define the system which we wish to study. Here, we consider an infinitely large homogeneous plasma with a uniform magnetic field oriented in the z direction.

We define

$$\vec{B} = (0, 0, B_0) \tag{G.15}$$

and a pressure tensor

$$\overleftrightarrow{P} = \begin{bmatrix} p_\perp & 0 & 0 \\ 0 & p_\perp & 0 \\ 0 & 0 & p_\parallel \end{bmatrix} = p_\perp I + (p_\parallel - p_\perp)\hat{z}\hat{z}. \tag{G.16}$$

Initially, the plasma is at rest with $\vec{U} = 0$ with some finite mass density $\rho = \rho_0$. We consider the plasma to initially have isotropic pressure $p_\perp = p_0$ and $p_\parallel = p_0$. However, the plasma pressure may evolve anisotropically in directions parallel or perpendicular to the background magnetic field.

We now proceed to linearize the system, considering small perturbations we replace

$$\rho \to \rho_0 + \delta\rho, \quad \vec{U} \to \delta\vec{U}, \quad \overleftrightarrow{P} \to \overleftrightarrow{P}_0 + \delta\overleftrightarrow{P}, \tag{G.17}$$

$$\vec{B} \to \vec{B}_0 + \delta\vec{B}, \quad V \to \delta V_0. \tag{G.18}$$

Keeping only first order terms, we obtain the new system of equations

$$\frac{\partial}{\partial t}\delta\rho + \rho_0 \nabla \cdot \delta\vec{U} = 0 \tag{G.19}$$

$$\rho_0 \frac{\partial}{\partial t}\delta\vec{U} = -\nabla\cdot \delta\overleftrightarrow{P} + (\nabla \times \delta\vec{B}) \times \vec{B}_0 - \rho_0 \nabla\delta V \tag{G.20}$$

$$\frac{\delta p_\parallel}{p_0} + \frac{2\delta B}{B_0} = \frac{3\delta\rho}{\rho_0} \tag{G.21}$$

$$\frac{\delta p_\perp}{p_0} = \frac{\delta\rho}{\rho_0} + \frac{\delta B}{B_0} \tag{G.22}$$

$$\frac{\partial}{\partial t}\delta\vec{B} = -\frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho}\frac{\partial}{\partial t}\nabla \times (\nabla \times \delta\vec{B}) + \nabla \times (\delta\vec{U} \times \vec{B}_0) \tag{G.23}$$

$$\nabla^2 \delta V = 4\pi G \delta\rho \tag{G.24}$$

where we have defined the scalar quantity $\delta B = \vec{B}_0 \cdot \delta\vec{B}/|\vec{B}_0|$. We now assume all perturbed quantities are proportional to $\exp i(\vec{k}\cdot\vec{r} - \omega t)$ and replace derivative operators. Additionally, we substitute δU for the time derivative of a Lagrangian displacement of the fluid, ξ

$$\delta\vec{U} = \frac{\partial\vec{\xi}}{\partial t} = -i\omega\vec{\xi} \tag{G.25}$$

Beginning with Eq. (G.23)

$$-i\omega\delta\vec{B} = i\vec{k} \times \left(-i\vec{k} \times (-i\omega)\frac{1}{2}\frac{m_\chi^2}{q_\chi^2 \rho}\delta\vec{B} - i\omega\vec{\xi} \times B_0\hat{z} \right) \tag{G.26}$$

yields the solution

$$\delta\vec{B} = iB_0\beta(k_z\xi_x\hat{x} + k_z\xi_y\hat{y} - (k_x\xi_x + k_y\xi_y)\hat{z}). \tag{G.27}$$

where

$$\beta = \frac{1}{1 + \frac{k^2}{2\rho}\frac{m_\chi^2}{q_\chi^2}} \tag{G.28}$$

which implies the following form for the scalar δB,

$$\delta B = \frac{\vec{B}_0 \cdot \delta\vec{B}}{B_0} = -iB_0\beta(k_x\xi_x + k_y\xi_y) \tag{G.29}$$

Next, from Eq. (G.19)

$$\frac{\delta\rho}{\rho_0} = -i\vec{k}\cdot\vec{\xi} \tag{G.30}$$

Can be utilized along with Eqs. (G.21) and (G.22) to obtain

$$\frac{\delta p_\parallel}{p_0} = -i(3 - 2\beta)(k_x\xi_x + k_y\xi_y) - 3i(k_z\xi_z) \tag{G.31}$$

and

$$\frac{\delta p_\perp}{p_0} = -i(1+\beta)(k_x\xi_x + k_y\xi_y) - i(k_z\xi_z) \tag{G.32}$$

Next, we first start by determining $\overset{\leftrightarrow}{\delta P}$

$$\overset{\leftrightarrow}{\delta P} = \delta p_\perp I + (\delta p_\parallel - \delta p_\perp)\hat{z}\hat{z} \tag{G.33}$$

We then find the x, y, and z components of $\nabla\cdot \overset{\leftrightarrow}{\delta P}$ to be

$$(i\vec{k}\cdot \overset{\leftrightarrow}{\delta P})_x = \left[(1+\beta)(k_x\xi_x + k_y\xi_y) + k_z\xi_z\right] k_x p_0 \tag{G.34}$$

$$(i\vec{k}\cdot \overset{\leftrightarrow}{\delta P})_y = \left[(1+\beta)(k_x\xi_x + k_y\xi_y) + k_z\xi_z\right] k_y p_0 \tag{G.35}$$

$$(i\vec{k}\cdot \overset{\leftrightarrow}{\delta P})_z = [(3-2\beta)(k_x\xi_x + k_y\xi_y) + 3(k_z\xi_z)]k_z p_0 \tag{G.36}$$

For the next term in Eq. (G.20) we have

$$\begin{aligned}(\nabla \times \delta\vec{B}) \times \vec{B}_0 = &- B_0^2\beta[(k_x(k_x\xi_x + k_y\xi_y) + k_z^2\xi_x)\hat{x} \\ &+ (k_y(k_x\xi_x + k_y\xi_y) + k_z^2\xi_y)\hat{y}]\end{aligned} \tag{G.37}$$

From Eq. (G.24) we find

$$\delta V = i\frac{4\pi G\rho_0}{k^2}(\vec{k}\cdot\vec{\xi}) \tag{G.38}$$

Together, this yields the following equations for each component of Eq. (G.20)

$$\begin{aligned}-\rho_0\omega^2\xi_x = &- \left[(1+\beta)(k_x\xi_x + k_y\xi_y) + k_z\xi_z\right] k_x p_0 \\ &- B_0^2\beta(k_x(k_x\xi_x + k_y\xi_y) + k_z^2\xi_x) \\ &+ 4\pi G\rho_0^2\frac{k_x}{k^2}(k_x\xi_x + k_y\xi_y + k_z\xi_z)\end{aligned} \tag{G.39}$$

$$
\begin{aligned}
-\rho_0\omega^2\xi_y = & - \left[(1+\beta)(k_x\xi_x + k_y\xi_y) + k_z\xi_z\right] k_y p_0 \\
& - B_0^2\beta(k_y(k_x\xi_x + k_y\xi_y) + k_z^2\xi_y) \\
& + 4\pi G\rho_0^2\frac{k_y}{k^2}(k_x\xi_x + k_y\xi_y + k_z\xi_z)
\end{aligned} \tag{G.40}
$$

$$
\begin{aligned}
-\rho_0\omega^2\xi_z = & - \left[(3-2\beta)(k_x\xi_x + k_y\xi_y) + 3(k_z\xi_z)\right] k_z p_0 \\
& + 4\pi G\rho_0^2\frac{k_z}{k^2}(k_x\xi_x + k_y\xi_y + k_z\xi_z)
\end{aligned} \tag{G.41}
$$

Having rid of all of the perturbed quantities, we now drop the 0 subscript for ease of notation. We also define

$$\xi_\parallel = \xi_z, \quad \xi_\perp = (\xi_x, \xi_y), \quad k_\parallel = k_z, \tag{G.42}$$

$$k_\perp = (k_x, k_y), \quad k_\perp^2 = k_x^2 + k_y^2. \tag{G.43}$$

We now multiply Eq. (G.39) by k_x and Eq. (G.40) by k_y and add the resulting equations together to yield

$$
\begin{aligned}
& \left[\rho\omega^2 - (1+\beta)k_\perp^2 p - \beta B^2 k^2 + 4\pi G\rho^2\frac{k_\perp^2}{k^2}\right] k_\perp \cdot \xi_\perp \\
& \qquad = \left(k_\perp^2 p - 4\pi G\rho^2\frac{k_\perp^2}{k^2}\right) k_\parallel\xi_\parallel
\end{aligned} \tag{G.44}
$$

Equation (G.41) yields

$$\left(\rho\omega^2 - 3k_\parallel^2 p + 4\pi G\rho^2\frac{k_\parallel^2}{k^2}\right) k_\parallel\xi_\parallel \tag{G.45}$$

$$= \left((3-2\chi)k_\parallel^2 p - 4\pi G\rho^2\frac{k_\parallel^2}{k^2}\right) k_\perp \cdot \xi_\perp \tag{G.46}$$

Combining the two above results gives the dispersion relation

$$\left[\rho\omega^2 - (1+\beta)k_\perp^2 p - \beta B^2 k^2 + 4\pi G\rho^2 \frac{k_\perp^2}{k^2}\right]\left(\rho\omega^2 - 3k_\parallel^2 p + 4\pi G\rho^2 \frac{k_\parallel^2}{k^2}\right)$$
$$= k_\perp^2 k_\parallel^2 \left[\left(p - \frac{4\pi G\rho^2}{k^2}\right)^2 + 2\left(p - \frac{4\pi G\rho^2}{k^2}\right)(1-\beta)\,p\right] \quad \text{(G.47)}$$

References

[1] A. A. Michelson, *Light Waves and Their Uses* (The University of Chicago Press, Chicago, 1903).

[2] A. Einstein, Annalen der Physik **322**, 891 (1905), https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19053221004 .

[3] A. Einstein, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften , 844 (1915).

[4] W. Heisenberg, Zeitschrift für Physik **33**, 879 (1925).

[5] E. Schrödinger, Annalen der Physik **384**, 361 (1926), https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19263840404 .

[6] P. Minkowski, Phys. Lett. **67B**, 421 (1977).

[7] M. Gell-Mann, P. Ramond, and R. Slansky, *Supergravity Workshop Stony Brook, New York, September 27-28, 1979*, Conf. Proc. **C790927**, 315 (1979), arXiv:1306.4669 [hep-th] .

[8] T. Yanagida, *Proceedings: Workshop on the Unified Theories and the Baryon Number in the Universe: Tsukuba, Japan, February 13-14, 1979*, Conf. Proc. **C7902131**, 95 (1979).

[9] S. L. Glashow, NATO Sci. Ser. B **61**, 687 (1980).

[10] R. N. Mohapatra and G. Senjanovic, Phys. Rev. Lett. **44**, 912 (1980), [,231(1979)].

[11] J. Schechter and J. W. F. Valle, Phys. Rev. D **22**, 2227 (1980).

[12] R. Foot, H. Lew, X. G. He, and G. C. Joshi, Z. Phys. C **44**, 441 (1989).

[13] T. Asaka *et al.*, Phys. Lett. **B631**, 151 (2005), arXiv:hep-ph/0503065 [hep-ph] .

[14] T. Asaka and M. Shaposhnikov, Phys. Lett. **B620**, 17 (2005), arXiv:hep-ph/0505013 [hep-ph] .

[15] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, JHEP **11**, 032, arXiv:1805.08567 [hep-ph] .

[16] A. M. Abdullahi *et al.*, in *2022 Snowmass Summer Study* (2022) arXiv:2203.08039 [hep-ph] .

[17] J. Orloff, A. N. Rozanov, and C. Santoni, Phys. Lett. B **550**, 8 (2002), arXiv:hep-ph/0208075 .

[18] G. Bernardi *et al.*, Phys. Lett. **166B**, 479 (1986).

[19] G. Bernardi *et al.*, Phys. Lett. **B203**, 332 (1988).

[20] P. Astier *et al.* (NOMAD), Phys. Lett. B **506**, 27 (2001), arXiv:hep-ex/0101041 .

[21] K. Abe *et al.* (T2K), Phys. Rev. D **100**, 052006 (2019), arXiv:1902.07598 [hep-ex] .

[22] C. A. Argüelles, N. Foppiani, and M. Hostert, Physical Review D **105**, 10.1103/physrevd.105.095006 (2022).

[23] P. Abratenko *et al.* (MicroBooNE), Phys. Rev. D **101**, 052001 (2020), arXiv:1911.10545 [hep-ex] .

[24] K. J. Kelly and P. A. N. Machado, Phys. Rev. D **104**, 055015 (2021), arXiv:2106.06548 [hep-ph] .

[25] R. Acciarri *et al.* (ArgoNeuT), Phys. Rev. Lett. **127**, 121801 (2021), arXiv:2106.13684 [hep-ex] .

[26] P. Ballett *et al.*, JHEP **04**, 102, arXiv:1610.08512 [hep-ph] .

[27] P. Ballett, T. Boschi, and S. Pascoli, JHEP **03**, 111, arXiv:1905.00284 [hep-ph] .

[28] J. M. Berryman, A. de Gouvea, P. J. Fox, B. J. Kayser, K. J. Kelly, and J. L. Raaf, JHEP **02**, 174, arXiv:1912.07622 [hep-ph] .

[29] P. Coloma, E. Fernández-Martínez, M. González-López, J. Hernández-García, and Z. Pavlovic, Eur. Phys. J. C **81**, 78 (2021), arXiv:2007.03701 [hep-ph] .

[30] F. Kling and S. Trojanowski, Physical Review D **97**, 10.1103/physrevd.97.095016 (2018).

[31] C. Alpigiani *et al.* (MATHUSLA), A Letter of Intent for MATHUSLA: A Dedicated Displaced Vertex Detector above ATLAS or CMS. (2018), arXiv:1811.00927 [physics.ins-det] .

[32] D. Curtin *et al.*, Rept. Prog. Phys. **82**, 116201 (2019), arXiv:1806.07396 [hep-ph] .

[33] H. Lubatti *et al.* (MATHUSLA), JINST **15** (06), C06026, arXiv:1901.04040 [hep-ex] .

[34] C. Alpigiani *et al.* (MATHUSLA), An Update to the Letter of Intent for MATHUSLA: Search for Long-Lived Particles at the HL-LHC (2020), arXiv:2009.01693 [physics.ins-det] .

[35] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, Phys. Rev. D **97**, 015023 (2018), arXiv:1708.09395 [hep-ph] .

[36] G. Aielli *et al.*, Eur. Phys. J. C **80**, 1177 (2020), arXiv:1911.00481 [hep-ex] .

[37] B. Batell, J. A. Evans, S. Gori, and M. Rai, JHEP **05**, 049, arXiv:2008.08108 [hep-ph] .

[38] S. N. Gninenko, D. S. Gorbunov, and M. E. Shaposhnikov, Adv. High Energy Phys. **2012**, 718259 (2012), arXiv:1301.5516 [hep-ph] .

[39] W. Bonivento *et al.*, (2013), arXiv:1310.1762 [hep-ex] .

[40] S. Alekhin *et al.*, Rept. Prog. Phys. **79**, 124201 (2016), arXiv:1504.04855 [hep-ph] .

[41] C. Ahdida *et al.* (SHiP), JHEP **04**, 077, arXiv:1811.00930 [hep-ph] .

[42] D. Gorbunov, I. Krasnov, Y. Kudenko, and S. Suvorov, Phys. Lett. B **810**, 135817 (2020), arXiv:2004.07974 [hep-ph] .

[43] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, JHEP **07**, 105, arXiv:1801.04207 [hep-ph] .

[44] E. J. Chun, A. Das, S. Mandal, M. Mitra, and N. Sinha, Phys. Rev. D **100**, 095022 (2019), arXiv:1908.09562 [hep-ph] .

[45] G. Aad *et al.* (ATLAS), JHEP **10**, 265, arXiv:1905.09787 [hep-ex] .

[46] R. Aaij *et al.* (LHCb), Eur. Phys. J. C **81**, 248 (2021), arXiv:2011.05263 [hep-ex] .

[47] K. Cheung, Y.-L. Chung, H. Ishida, and C.-T. Lu, Phys. Rev. D **102**, 075038 (2020), arXiv:2004.11537 [hep-ph] .

[48] D. Liventsev *et al.* (Belle), Phys. Rev. D **87**, 071102 (2013), [Erratum: Phys.Rev.D 95, 099903 (2017)], arXiv:1301.1105 [hep-ex] .

[49] S. Dreyer *et al.*, (2021), arXiv:2105.12962 [hep-ph] .

[50] A. Blondel, E. Graverini, N. Serra, and M. Shaposhnikov (FCC-ee study Team), Nucl. Part. Phys. Proc. **273-275**, 1883 (2016), arXiv:1411.5230 [hep-ex] .

[51] D. Bryman and R. Shrock, (2025), arXiv:2502.18384 [hep-ph] .

[52] W. Altmannshofer *et al.* (PIONEER), (2022), arXiv:2203.01981 [hep-ex] .

[53] W. Altmannshofer *et al.* (PIONEER), in *Snowmass 2021* (2022) arXiv:2203.05505 [hep-ex] .

[54] M. Heikinheimo, M. Raidal, C. Spethmann, and H. Veermäe, Phys. Lett. B **749**, 236 (2015), arXiv:1504.04371 [hep-ph] .

[55] M. Heikinheimo, M. Raidal, C. Spethmann, and H. Veermae, Plasma Phys. Control. Fusion **60**, 014011 (2017), arXiv:1707.03662 [astro-ph.CO] .

[56] C. Spethmann, H. Veermäe, T. Sepp, M. Heikinheimo, B. Deshev, A. Hektor, and M. Raidal, Astron. Astrophys. **608**, A125 (2017), arXiv:1603.07324 [astro-ph.CO] .

[57] R. Lasenby, JCAP **11**, 034, arXiv:2007.00667 [hep-ph] .

[58] A. Cruz and M. McQuinn, Journal of Cosmology and Astroparticle Physics **2023** (04), 028.

[59] N. Shukla, K. Schoeffler, J. Vieira, R. Fonseca, E. Boella, and L. O. Silva, Physical Review E **105**, 10.1103/physreve.105.035204 (2022).

[60] M. V. Medvedev and A. Loeb, Plasma constraints on the millicharged dark matter (2024), arXiv:2406.15750 [astro-ph.CO] .

[61] A. MacInnis and N. Sehgal, Journal of Cosmology and Astroparticle Physics **2025** (02), 048.

[62] D. R. DeBoer *et al.*, Publications of the Astronomical Society of the Pacific **129**, 045001 (2017).

[63] M. Leo, T. Theuns, C. M. Baugh, B. Li, and S. Pascoli, Journal of Cosmology and Astroparticle Physics **2020** (04), 004.

[64] P. Giffin, S. Gori, Y.-D. Tsai, and D. Tuckler, Journal of High Energy Physics **2023**, 10.1007/jhep04(2023)046 (2023).

[65] M. Bai *et al.*, C^3: A "cool" route to the higgs boson and beyond (2021), arXiv:2110.15800 [hep-ex] .

[66] M. Chrzaszcz, M. Drewes, and J. Hajer, Eur. Phys. J. C **81**, 546 (2021), arXiv:2011.01005 [hep-ph] .

[67] S. Kanemura, T. Moroi, and T. Tanabe, Phys. Lett. B **751**, 25 (2015), arXiv:1507.02809 [hep-ph] .

[68] K. Asai, S. Iwamoto, Y. Sakaki, and D. Ueda, Journal of High Energy Physics **2021**, 10.1007/jhep09(2021)183 (2021).

[69] Y. Sakaki and D. Ueda, Phys. Rev. D **103**, 035024 (2021), arXiv:2009.13790 [hep-ph] .

[70] K. Asai, T. Moroi, and A. Niki, Phys. Lett. B **818**, 136374 (2021), arXiv:2104.00888 [hep-ph] .

[71] P. Achard *et al.* (L3), Phys. Lett. B **517**, 75 (2001), arXiv:hep-ex/0107015 .

[72] S. Antusch and O. Fischer, JHEP **05**, 053, arXiv:1502.05915 [hep-ph] .

[73] S. Antusch, E. Cazzato, and O. Fischer, Int. J. Mod. Phys. A **32**, 1750078 (2017), arXiv:1612.02728 [hep-ph] .

[74] C. B. Verhaaren, J. Alimena, M. Bauer, P. Azzi, R. Ruiz, M. Neubert, O. Mikulenko, M. Ovchynnikov, M. Drewes, J. Klaric, A. Blondel, C. Rizzi, A. Sfyrla, T. Sharma, S. Kulkarni, A. Thamm, A. Blondel, R. G. Suarez, and L. Rygaard, Frontiers in Physics **10**, 10.3389/fphy.2022.967881 (2022).

[75] B. Batell, T. Ghosh, and K. Xie (nu-Test), in *2022 Snowmass Summer Study* (2022) arXiv:2203.06705 [hep-ph] .

[76] L. A. Anchordoqui *et al.*, Phys. Rept. **968**, 1 (2022), arXiv:2109.10905 [hep-ph] .

[77] J. L. Feng *et al.*, Journal of Physics G: Nuclear and Particle Physics **50**, 030501 (2023).

[78] K. J. Kelly, P. A. N. Machado, A. Marchionni, and Y. F. Perez-Gonzalez, JHEP **08**, 087, arXiv:2103.00009 [hep-ph] .

[79] A. Aryshev *et al.* (ILC International Development Team), (2022), arXiv:2203.07622 [physics.acc-ph] .

[80] CERN, Cern yellow reports: Monographs, vol 4 (2018): The compact linear collider (clic) – project implementation plan (2019).

[81] S. Dasu *et al.*, in *2022 Snowmass Summer Study* (2022) arXiv:2203.07646 [hep-ex] .

[82] E. A. Nanni *et al.*, in *2022 Snowmass Summer Study* (2022) arXiv:2203.09076 [physics.acc-ph] .

[83] A. Abada *et al.* (FCC), Eur. Phys. J. ST **228**, 261 (2019).

[84] M. Aicheler *et al.*, A Multi-TeV Linear Collider Based on CLIC Technology: CLIC Conceptual Design Report (2012).

[85] T. Sjostrand, S. Mrenna, and P. Z. Skands, JHEP **05**, 026, arXiv:hep-ph/0603175 [hep-ph] .

[86] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, Comput. Phys. Commun. **191**, 159 (2015), arXiv:1410.3012 [hep-ph] .

[87] C. Bierlich *et al.*, (2022), arXiv:2203.11601 [hep-ph] .

[88] Particle data group atomic and nuclear properties, `https://pdg.lbl.gov/2020/AtomicNuclearProperties/`.

[89] W. Buchmuller and C. Greub, Nucl. Phys. B **363**, 345 (1991).

[90] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079, arXiv:1405.0301 [hep-ph] .

[91] R. E. Shrock, Phys. Lett. B **96**, 159 (1980).

[92] R. E. Shrock, Phys. Rev. D **24**, 1232 (1981).

[93] R. E. Shrock, Phys. Rev. D **24**, 1275 (1981).

[94] D. Gorbunov and M. Shaposhnikov, JHEP **10**, 015, [Erratum: JHEP 11, 101 (2013)], arXiv:0705.1729 [hep-ph] .

[95] A. Atre, T. Han, S. Pascoli, and B. Zhang, JHEP **05**, 030, arXiv:0901.3589 [hep-ph] .

[96] J. M. Berryman, A. de Gouvêa, K. J. Kelly, and Y. Zhang, Phys. Rev. D **96**, 075010 (2017), arXiv:1706.02722 [hep-ph] .

[97] A. Aguilar-Arevalo *et al.* (PiENu), Phys. Rev. Lett. **115**, 071801 (2015), arXiv:1506.05845 [hep-ex] .

[98] A. Aguilar-Arevalo *et al.* (PIENU), Phys. Rev. D **97**, 072012 (2018), arXiv:1712.03275 [hep-ex] .

[99] E. Cortina Gil *et al.* (NA62), Phys. Lett. B **807**, 135599 (2020), arXiv:2005.09575 [hep-ex] .

[100] F. Bergsma *et al.* (CHARM), Phys. Lett. B **166**, 473 (1986).

[101] A. M. Cooper-Sarkar *et al.* (WA66), Phys. Lett. B **160**, 207 (1985).

[102] R. Barouki, G. Marocco, and S. Sarkar, SciPost Phys. **13**, 118 (2022), arXiv:2208.00416 [hep-ph] .

[103] P. Abreu *et al.* (DELPHI), Z. Phys. **C74**, 57 (1997), [Erratum: Z. Phys.C75,580(1997)].

[104] G. Aad *et al.*, Physical Review Letters **131**, 10.1103/physrevlett.131.061803 (2023).

[105] A. V. Artamonov *et al.* (E949), Phys. Rev. **D91**, 052001 (2015), [Erratum: Phys. Rev.D91,no.5,059903(2015)], arXiv:1411.3963 [hep-ex] .

[106] A. Vaitaitis *et al.* (NuTeV, E815), Phys. Rev. Lett. **83**, 4943 (1999), arXiv:hep-ex/9908011 [hep-ex] .

[107] P. Vilain *et al.* (CHARM II), Phys. Lett. **B343**, 453 (1995), [Phys. Lett.B351,387(1995)].

[108] I. Boiarska, A. Boyarsky, O. Mikulenko, and M. Ovchynnikov, Phys. Rev. D **104**, 095019 (2021), arXiv:2107.14685 [hep-ph] .

[109] A. Boyarsky, M. Ovchynnikov, O. Ruchayskiy, and V. Syvolap, Phys. Rev. D **104**, 023517 (2021), arXiv:2008.00749 [hep-ph] .

[110] N. Sabti, A. Magalich, and A. Filimonova, JCAP **11**, 056, arXiv:2006.07387 [hep-ph] .

[111] M. Drewes and J. Hajer, JHEP **02**, 070, arXiv:1903.06100 [hep-ph] .

[112] J. Blumlein *et al.*, Z. Phys. C **51**, 341 (1991).

[113] Y.-D. Tsai, P. deNiverville, and M. X. Liu, Phys. Rev. Lett. **126**, 181801 (2021), arXiv:1908.07525 [hep-ph] .

[114] W. Altmannshofer, J. A. Dror, P. Giffin, S. Gori, O. Jackson, K. Luong, P. Schwendimann, and S. R. Seo, Exploring invisible new physics with exotic pion decays (2026), arXiv:2601.06254 [hep-ph] .

[115] Y. G. Aditya, K. J. Healey, and A. A. Petrov, Phys. Lett. B **710**, 118 (2012), arXiv:1201.1007 [hep-ph] .

[116] B. Batell, T. Han, D. McKeen, and B. Shams Es Haghi, Phys. Rev. D **97**, 075016 (2018), arXiv:1709.07001 [hep-ph] .

[117] D. S. M. Alves and N. Weiner, JHEP **07**, 092, arXiv:1710.03764 [hep-ph] .

[118] W. Altmannshofer, S. Gori, and D. J. Robinson, Phys. Rev. D **101**, 075002 (2020), arXiv:1909.00005 [hep-ph] .

[119] J. A. Dror, Phys. Rev. D **101**, 095013 (2020), arXiv:2004.04750 [hep-ph] .

[120] M. Hostert and M. Pospelov, Phys. Rev. D **105**, 015017 (2022), arXiv:2012.02142 [hep-ph] .

[121] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, JHEP **04**, 063, arXiv:2012.12272 [hep-ph] .

[122] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, Phys. Rev. Lett. **127**, 081803 (2021), arXiv:2102.13112 [hep-ph] .

[123] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, JHEP **09**, 056, arXiv:2110.10698 [hep-ph] .

[124] J. A. Gallo, A. W. M. Guerrera, S. Peñaranda, and S. Rigolin, Nucl. Phys. B **979**, 115791 (2022), arXiv:2111.02536 [hep-ph] .

[125] B. Dutta, D. Kim, A. Thompson, R. T. Thornton, and R. G. Van de Water, Phys. Rev. Lett. **129**, 111803 (2022), arXiv:2110.11944 [hep-ph] .

[126] G. Bickendorf and M. Drees, Eur. Phys. J. C **82**, 1163 (2022), arXiv:2206.05038 [hep-ph] .

[127] W. Altmannshofer, J. A. Dror, and S. Gori, Phys. Rev. Lett. **130**, 241801 (2023), arXiv:2209.00665 [hep-ph] .

[128] A. W. M. Guerrera and S. Rigolin, Fortsch. Phys. **71**, 2200192 (2023), arXiv:2211.08343 [hep-ph] .

[129] L. Di Luzio, A. W. M. Guerrera, X. P. Díaz, and S. Rigolin, JHEP **06**, 046, arXiv:2304.04643 [hep-ph] .

[130] M. Hostert and M. Pospelov, Phys. Rev. D **108**, 055011 (2023), arXiv:2306.15077 [hep-ph] .

[131] A. Eberhart, M. Fedele, F. Kahlhoefer, E. Ravensburg, and R. Ziegler, (2025), arXiv:2504.05873 [hep-ph] .

[132] L. Di Luzio, P. Paradisi, and N. Selimovic, Nucl. Phys. B **1021**, 117177 (2025), arXiv:2504.14014 [hep-ph] .

[133] S. Okawa and Y. Omura, (2025), arXiv:2511.08102 [hep-ph] .

[134] M. Aoki *et al.* (PIENU), Phys. Rev. D **84**, 052002 (2011), arXiv:1106.4055 [hep-ex] .

[135] A. Aguilar-Arevalo *et al.* (PIENU), Phys. Lett. B **798**, 134980 (2019), arXiv:1904.03269 [hep-ex] .

[136] A. Aguilar-Arevalo *et al.* (PIENU), Phys. Rev. D **101**, 052014 (2020), arXiv:2002.09170 [hep-ex] .

[137] A. Aguilar-Arevalo *et al.* (PIENU), Phys. Rev. D **103**, 052006 (2021), arXiv:2101.07381 [hep-ex] .

[138] E. Cortina Gil *et al.* (NA62), Phys. Lett. B **816**, 136259 (2021), arXiv:2101.12304 [hep-ex] .

[139] B. Bloch-Devaux *et al.* (NA62), Physics Letters B **872**, 140119 (2026), arXiv:2507.07345 [hep-ex] .

[140] W. J. Marciano and A. Sirlin, Phys. Rev. Lett. **71**, 3629 (1993).

[141] M. Knecht, H. Neufeld, H. Rupertsberger, and P. Talavera, Eur. Phys. J. C **12**, 469 (2000), arXiv:hep-ph/9909284 .

[142] V. Cirigliano and I. Rosell, Phys. Rev. Lett. **99**, 231801 (2007), arXiv:0707.3439 [hep-ph] .

[143] V. Cirigliano and I. Rosell, JHEP **10**, 005, arXiv:0707.4464 [hep-ph] .

[144] D. A. Bryman, M. S. Dixit, R. Dubois, J. A. Macdonald, T. Numao, B. Olaniyi, A. Olin, and J. M. Poutissou, Phys. Rev. D **33**, 1211 (1986).

[145] D. I. Britton *et al.*, Phys. Rev. Lett. **68**, 3000 (1992).

[146] G. Czapek *et al.*, Phys. Rev. Lett. **70**, 17 (1993).

[147] S. Navas *et al.* (Particle Data Group), Phys. Rev. D **110**, 030001 (2024).

[148] D. Pocanic *et al.*, Phys. Rev. Lett. **93**, 181803 (2004), arXiv:hep-ex/0312030 .

[149] A. Jodidio *et al.*, Phys. Rev. D **34**, 1967 (1986), [Erratum: Phys.Rev.D 37, 237 (1988)].

[150] R. Bayes *et al.* (TWIST), Phys. Rev. D **91**, 052020 (2015), arXiv:1409.0638 [hep-ex] .

[151] L. Calibbi, D. Redigolo, R. Ziegler, and J. Zupan, JHEP **09**, 173, arXiv:2006.04795 [hep-ph] .

[152] Y. Jho, S. Knapen, and D. Redigolo, JHEP **10**, 029, arXiv:2203.11222 [hep-ph] .

[153] R. J. Hill, R. Plestid, and J. Zupan, Phys. Rev. D **109**, 035025 (2024), arXiv:2310.00043 [hep-ph] .

[154] S. Knapen, K. Langhoff, T. Opferkuch, and D. Redigolo, JHEP **07**, 243, arXiv:2311.17915 [hep-ph] .

[155] W. Altmannshofer *et al.*, in preparation.

[156] A. de Gouvêa and A. Kobach, Phys. Rev. D **93**, 033005 (2016), arXiv:1511.00683 [hep-ph] .

[157] D. A. Bryman and R. Shrock, Phys. Rev. D **100**, 073011 (2019), arXiv:1909.11198 [hep-ph] .

[158] F. F. Deppisch, T. E. Gonzalo, C. Majumdar, and Z. Zhang, JCAP **02**, 054, arXiv:2410.06970 [hep-ph] .

[159] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, JHEP **12**, 216, arXiv:2410.05380 [hep-ph] .

[160] N. Aghanim *et al.* (Planck), Astron. Astrophys. **641**, A5 (2020), arXiv:1907.12875 [astro-ph.CO] .

[161] M. S. Madhavacheril *et al.* (ACT), Astrophys. J. **962**, 113 (2024), arXiv:2304.05203 [astro-ph.CO] .

[162] A. G. Adame *et al.* (DESI), JCAP **02**, 021, arXiv:2404.03002 [astro-ph.CO] .

[163] D. A. Bryman, R. Dubois, T. Numao, B. Olaniyi, A. Olin, M. S. Dixit, J. M. Poutissou, and J. A. Macdonald, Phys. Rev. Lett. **50**, 1546 (1983).

[164] G. Azuelos *et al.*, Phys. Rev. Lett. **56**, 2241 (1986).

[165] N. De Leener-Rosier *et al.*, Phys. Rev. D **43**, 3611 (1991).

[166] D. I. Britton *et al.*, Phys. Rev. D **46**, R885 (1992).

[167] R. Abela, M. Daum, G. H. Eaton, R. Frosch, B. Jost, P. R. Kettle, and E. Steiner, Phys. Lett. B **105**, 263 (1981), [Erratum: Phys.Lett.B 106, 513 (1981)].

[168] R. C. Minehart, K. O. H. Ziock, R. Marshall, W. A. Stephens, M. Daum, B. Jost, and P. R. Kettle, Phys. Rev. Lett. **52**, 804 (1984).

[169] M. Daum, B. Jost, R. M. Marshall, R. C. Minehart, W. A. Stephens, and K. O. H. Ziock, Phys. Rev. D **36**, 2624 (1987).

[170] M. Daum, R. Frosch, W. Hajdas, M. Janousch, P. R. Kettle, S. Ritt, and Z. G. Zhao, Phys. Lett. B **361**, 179 (1995).

[171] D. A. Bryman and T. Numao, Phys. Rev. D **53**, 558 (1996).

[172] K. Assamagan, C. Bronnimann, M. Daum, R. Frosch, P. R. Kettle, and C. Wigger, Phys. Lett. B **434**, 158 (1998).

[173] M. Daum, M. Janousch, P. R. Kettle, J. Koglin, D. Pocanic, J. Schottmuller, C. Wigger, and Z. G. Zhao, Phys. Rev. Lett. **85**, 1815 (2000), arXiv:hep-ex/0008014 .

[174] Y. Asano *et al.*, Phys. Lett. B **104**, 84 (1981).

[175] R. S. Hayano *et al.*, Phys. Rev. Lett. **49**, 1305 (1982).

[176] E. Cortina Gil *et al.* (NA62), Phys. Lett. B **778**, 137 (2018), arXiv:1712.00297 [hep-ex] .

[177] R. L. Workman *et al.* (Particle Data Group), PTEP **2022**, 083C01 (2022).

[178] P. Schwendimann, Pioneer: A next generation rare pion decay experiment relying on 5d tracking (Workshop on ACTS Tracking for Nuclear Physics 2025, LBNL May 12-15, 2025).

[179] P. Schwendimann, Overview and current status of the simulation, reco and analysis (PIONEER collaboration meeting, TRIUMF Jan 7-10, 2025).

[180] O. Beesley *et al.*, Nucl. Instrum. Meth. A **1075**, 170320 (2025), arXiv:2409.14691 [physics.ins-det] .

[181] K. Bondarenko, A. Boyarsky, J. Klaric, O. Mikulenko, O. Ruchayskiy, V. Syvolap, and I. Timiryasov, JHEP **07**, 193, arXiv:2101.09255 [hep-ph] .

[182] P. B. Dev, Q.-f. Wu, and X.-J. Xu, Physical Review D **113**, 10.1103/r7bw-jzjl (2026).

[183] P. D. Bolton, F. F. Deppisch, and P. S. Bhupal Dev, JHEP **03**, 170, arXiv:1912.03058 [hep-ph] .

[184] W. Dekens, J. de Vries, K. Fuyuto, E. Mereghetti, and G. Zhou, JHEP **06**, 097, arXiv:2002.07182 [hep-ph] .

[185] F. F. Deppisch, L. Graf, F. Iachello, and J. Kotila, Phys. Rev. D **102**, 095016 (2020), arXiv:2009.10119 [hep-ph] .

[186] J. de Vries, H. K. Dreiner, J. Groot, J. Y. Günther, and Z. S. Wang, JHEP **04**, 007, arXiv:2406.15091 [hep-ph] .

[187] M. Breitbach, L. Buonocore, C. Frugiuele, J. Kopp, and L. Mittnacht, JHEP **01**, 048, arXiv:2102.03383 [hep-ph] .

[188] P. Carenza and G. Lucente, Phys. Rev. D **104**, 103007 (2021), [Erratum: Phys.Rev.D 110, 049901 (2024)], arXiv:2107.12393 [hep-ph] .

[189] D. F. G. Fiorillo, T. Pitik, and E. Vitagliano, Phys. Rev. D **112**, 083008 (2025), arXiv:2503.15630 [hep-ph] .

[190] P. J. Fox, J. Liu, D. Tucker-Smith, and N. Weiner, Phys. Rev. D **84**, 115006 (2011), arXiv:1104.4127 [hep-ph] .

[191] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse, Phys. Rev. Lett. **130**, 071801 (2023), arXiv:2209.13084 [physics.atom-ph] .

[192] D. P. Aguillard *et al.* (Muon g-2), Phys. Rev. Lett. **135**, 101802 (2025), arXiv:2506.03069 [hep-ex] .

[193] R. Aliberti *et al.*, Phys. Rept. **1143**, 1 (2025), arXiv:2505.21476 [hep-ph] .

[194] R. H. Parker, C. Yu, W. Zhong, B. Estey, and H. Müller, Science **360**, 191 (2018), arXiv:1812.04130 [physics.atom-ph] .

[195] L. Morel, Z. Yao, P. Cladé, and S. Guellati-Khélifa, Nature **588**, 61 (2020).

[196] T. Aoyama, T. Kinoshita, and M. Nio, Atoms **7**, 28 (2019).

[197] S. Volkov, Phys. Rev. D **110**, 036001 (2024), arXiv:2404.00649 [hep-ph] .

[198] T. Aoyama, M. Hayakawa, A. Hirayama, and M. Nio, Phys. Rev. D **111**, L031902 (2025), arXiv:2412.06473 [hep-ph] .

[199] T. Aoyama *et al.*, Phys. Rept. **887**, 1 (2020), arXiv:2006.04822 [hep-ph] .

[200] J. P. Lees *et al.* (BaBar), Phys. Rev. Lett. **119**, 131804 (2017), arXiv:1702.03327 [hep-ex] .

[201] L. Aggarwal *et al.* (Belle-II), (2022), arXiv:2207.06307 [hep-ex] .

[202] S. N. Gninenko, N. V. Krasnikov, M. M. Kirsanov, and D. V. Kirpichnikov, Phys. Rev. D **94**, 095025 (2016), arXiv:1604.08432 [hep-ph] .

[203] Y. M. Andreev *et al.* (NA64), Phys. Rev. Lett. **131**, 161801 (2023), arXiv:2307.02404 [hep-ex] .

[204] Y. M. Andreev *et al.* (NA64), Phys. Rev. D **109**, L031103 (2024), arXiv:2308.15612 [hep-ex] .

[205] T. Akesson *et al.* (LDMX), (2025), arXiv:2508.11833 [hep-ex] .

[206] Y. M. Andreev *et al.* (NA64), Phys. Rev. Lett. **126**, 211802 (2021), arXiv:2102.01885 [hep-ex] .

[207] A. V. Artamonov *et al.* (BNL-E949), Phys. Rev. D **79**, 092004 (2009), arXiv:0903.0030 [hep-ex] .

[208] E. Cortina Gil *et al.* (NA62), JHEP **02**, 201, arXiv:2010.07644 [hep-ex] .

[209] E. Cortina Gil *et al.* (NA62), JHEP **03**, 058, arXiv:2011.11329 [hep-ex] .

[210] E. Cortina Gil *et al.* (NA62), JHEP **11**, 143, arXiv:2507.17286 [hep-ex] .

[211] D. Guadagnoli, A. Iohner, C. Lazzeroni, D. Martinez Santos, J. C. Swallow, and C. Toni, (2025), arXiv:2503.05865 [hep-ph] .

[212] Y. Chikashige, R. N. Mohapatra, and R. D. Peccei, Phys. Lett. B **98**, 265 (1981).

[213] G. B. Gelmini and M. Roncadelli, Phys. Lett. B **99**, 411 (1981).

[214] V. D. Barger, W.-Y. Keung, and S. Pakvasa, Phys. Rev. D **25**, 907 (1982).

[215] J. Schechter and J. W. F. Valle, Phys. Rev. D **25**, 774 (1982).

[216] C. H. de Lima, D. McKeen, J. Ng, and D. Tuckler, (2026), arXiv:2601.06248 [hep-ph] .

[217] W. DeRocco and P. Giffin, Physical Review D **111**, 10.1103/physrevd.111.095031 (2025).

[218] L. Ackerman, M. R. Buckley, S. M. Carroll, and M. Kamionkowski, Physical Review D **79**, 10.1103/physrevd.79.023519 (2009).

[219] J. L. Feng, M. Kaplinghat, H. Tu, and H.-B. Yu, Journal of Cosmology and Astroparticle Physics **2009** (07), 004.

[220] P. Agrawal, F.-Y. Cyr-Racine, L. Randall, and J. Scholtz, JCAP **05**, 022, arXiv:1610.04611 [hep-ph] .

[221] A. Bret, The Astrophysical Journal **900**, 111 (2020).

[222] A. Bret and C. Deutsch, Physics of Plasmas **12**, 082704 (2005).

[223] A. Bret, L. Gremillet, and M. E. Dieckmann, Physics of Plasmas **17**, 120501 (2010).

[224] M. E. Dieckmann and A. Bret, Journal of Plasma Physics **83**, 905830104 (2017).

[225] A. Bret, C. C. Haggerty, and R. Narayan, Journal of Plasma Physics **90**, 905900213 (2024).

[226] A. Bret, A. Stockem Novo, R. Narayan, C. Ruyer, M. E. Dieckmann, and L. O. Silva, Laser and Particle Beams **34**, 362–367 (2016).

[227] G. Livadiotis, The Astrophysical Journal **886**, 3 (2019).

[228] D. A. Gurnett and A. Bhattacharjee, *Introduction to Plasma Physics: With Space, Laboratory and Astrophysical Applications*, 2nd ed. (Cambridge University Press, Cambridge, 2017).

[229] A. Bret, A. Stockem, F. Fiúza, E. P. Álvaro, C. Ruyer, R. Narayan, and L. O. Silva, Laser and Particle Beams **31**, 487–491 (2013).

[230] S. Jubin, A. T. Powis, W. Villafana, D. Sydorenko, S. Rauf, A. V. Khrabrov, S. Sarwar, and I. D. Kaganovich, Physics of Plasmas **31**, 023902 (2024).

[231] P. Rambo, Journal of Computational Physics **118**, 152 (1995).

[232] M. D. Acciarri, C. Moore, and S. D. Baalrud, Physics of Plasmas **31**, 093903 (2024).

[233] A. Grassi, M. Grech, F. Amiranoff, F. Pegoraro, A. Macchi, and C. Riconda, Phys. Rev. E **95**, 023203 (2017).

[234] M. V. Medvedev, M. Fiore, R. A. Fonseca, L. O. Silva, and W. B. Mori, The Astrophysical Journal **618**, L75 (2004).

[235] C. Ruyer and F. Fiuza, Physical Review Letters **120**, 245002 (2018).

[236] D. Bohm and E. P. Gross, Phys. Rev. **75**, 1864 (1949).

[237] B. D. Fried, The Physics of Fluids **2**, 337 (1959).

[238] W. Tucker, P. Blanco, S. Rappoport, L. David, D. Fabricant, E. E. Falco, W. Forman, A. Dressler, and M. Ramella, The Astrophysical Journal **496**, L5–L8 (1998), aDS Bibcode: 1998ApJ...496L...5T.

[239] M. Markevitch, A. H. Gonzalez, L. David, A. Vikhlinin, S. Murray, W. Forman, C. Jones, and W. Tucker, The Astrophysical Journal **567**, L27–L31 (2002), aDS Bibcode: 2002ApJ...567L..27M.

[240] R. Barrena, A. Biviano, M. Ramella, E. E. Falco, and S. Seitz, Astronomy & Astrophysics **386**, 816–828 (2002).

[241] D. Paraficz, J.-P. Kneib, J. Richard, A. Morandi, M. Limousin, E. Jullo, and J. Martinez, Astronomy & Astrophysics **594**, A121 (2016).

[242] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, The Astrophysical Journal **648**, L109 (2006).

[243] D. Clowe, A. Gonzalez, and M. Markevitch, The Astrophysical Journal **604**, 596 (2004).

[244] S. W. Randall, M. Markevitch, D. Clowe, A. H. Gonzalez, and M. Bradač, The Astrophysical Journal **679**, 1173 (2008).

[245] A. Robertson, R. Massey, and V. Eke, Mon. Not. R. Astron. Soc. **465**, 569 (2017), arXiv:1605.04307 [astro-ph.CO] .

[246] F. Kahlhoefer, K. Schmidt-Hoberg, M. T. Frandsen, and S. Sarkar, Monthly Notices of the Royal Astronomical Society **437**, 2865–2881 (2014).

[247] S. Tulin and H.-B. Yu, Physics Reports **730**, 1 (2018), dark matter self-interactions and small scale structure.

[248] C. Ruyer, L. Gremillet, A. Debayle, and G. Bonnaud, Physics of Plasmas **22**, 032102 (2015).

[249] C. H. Jaroschek, H. Lesch, and R. A. Treumann, The Astrophysical Journal **616**, 1065 (2004).

[250] H. R. Kaufman, Journal of Vacuum Science & Technology B: Microelectronics Processing and Phenomena **8**, 107–108 (1990).

[251] H. Takabe, Physics of Plasmas **30**, 030901 (2023).

[252] A. Vanthieghem, M. Lemoine, and L. Gremillet, Physics of Plasmas **25**, 072115 (2018).

[253] P. Chang, A. Spitkovsky, and J. Arons, The Astrophysical Journal **674**, 378 (2008).

[254] M. D. McFarland and A. Y. Wong, Physics of Plasmas **8**, 110–121 (2001).

[255] M. D. McFarland and A. Y. Wong, Physics of Plasmas **8**, 122–131 (2001).

[256] P. H. Yoon, Physics of Plasmas **25**, 011603 (2017).

[257] J. Derouillat, A. Beck, F. Pérez, T. Vinci, M. Chiaramello, A. Grassi, M. Fle, G. Bouchard, I. Plotnikov, N. Aunai, J. Dargent, C. Riconda, and M. Grech, Computer Physics Communications **222**, 351–373 (2018).

[258] V. Springel and G. R. Farrar, Monthly Notices of the Royal Astronomical Society **380**, 911–925 (2007).

[259] M. Zhou, V. Zhdankin, M. W. Kunz, N. F. Loureiro, and D. A. Uzdensky, The Astrophysical Journal **960**, 12 (2023).

[260] **824**, 123 (2016).

[261] A. F. A. Bott, M. W. Kunz, E. Quataert, J. Squire, and L. Arzamasskiy (2025).

[262] S. Adhikari *et al.*, Astrophysical tests of dark matter self-interactions (2022), arXiv:2207.10638 [astro-ph.CO] .

[263] P. Giffin, A. Liu, J. Boucsein, A. Cruz, A. Prabhu, S. Profumo, and M. G. Roberts, Physical Review D **113**, 10.1103/4fkn-vx39 (2026).

[264] J. Binney and S. Tremaine, *Galactic dynamics*, 2nd ed., Princeton series in astrophysics (Princeton University Press, 2008).

[265] D. Matravers, General Relativity and Gravitation **41**, 1455 (2009).

[266] R. Barkana and A. Loeb, Physics Reports **349**, 125–238 (2001).

[267] A. E. Evrard *et al.*, Astrophys. J. **672**, 122 (2008), arXiv:astro-ph/0702241 .

[268] E. Munari, A. Biviano, S. Borgani, G. Murante, and D. Fabjan, Monthly Notices of the Royal Astronomical Society **430**, 2638–2649 (2013).

[269] J. S. Bullock and M. Boylan-Kolchin, Ann. Rev. Astron. Astrophys. **55**, 343 (2017), arXiv:1707.04256 [astro-ph.CO] .

[270] J.-T. Li and T. Lin, Physical Review D **101**, 10.1103/physrevd.101.103034 (2020).

[271] R. J. Ewart, M. L. Nastac, P. J. Bilbao, T. Silva, L. O. Silva, and A. A. Schekochihin, Proc. Nat. Acad. Sci. **122**, e2417813122 (2025), arXiv:2409.01742 [physics.plasm-ph] .

[272] G. F. Chew, M. L. Goldberger, and F. E. Low, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences **236**, 112 (1956).

[273] R. Banerjee and K. Jedamzik, Phys. Rev. D **70**, 123003 (2004).

[274] P. Trivedi, J. Reppin, J. Chluba, and R. Banerjee, Mon. Not. Roy. Astron. Soc. **481**, 3401 (2018), arXiv:1805.05315 [astro-ph.CO] .

[275] K. Jedamzik and A. Saveliev, Phys. Rev. Lett. **123**, 021301 (2019), arXiv:1804.06115 [astro-ph.CO] .

[276] E.-J. Kim, A. V. Olinto, and R. Rosner, The Astrophysical Journal **468**, 28 (1996).

[277] K. Subramanian and J. D. Barrow, Physical Review D **58**, 10.1103/physrevd.58.083502 (1998).

[278] R. Durrer and C. Caprini, Journal of Cosmology and Astroparticle Physics **2003** (11), 010–010.

[279] L. Mestel and I. W. Roxburgh, Astrophys. J. **136**, 615 (1962).

[280] L. Biermann, Zeitschrift Naturforschung Teil A **5**, 65 (1950).

[281] M. V. Medvedev, L. O. Silva, and M. Kamionkowski, The Astrophysical Journal **642**, L1 (2006).

[282] M. Lazar, R. Schlickeiser, R. Wielebinski, and S. Poedts, Astrophys. J. **693**, 1133 (2009).

[283] M. S. Turner and L. M. Widrow, Phys. Rev. D **37**, 2743 (1988).

[284] B. Ratra, Astrophys. J. Lett. **391**, L1 (1992).

[285] A. D. Dolgov, Phys. Rev. D **48**, 2499 (1993).

[286] M. Gasperini, M. Giovannini, and G. Veneziano, Phys. Rev. Lett. **75**, 3796 (1995).

[287] M. Giovannini, Phys. Rev. D **62**, 067301 (2000).

[288] K. Atmjeet, I. Pahwa, T. R. Seshadri, and K. Subramanian, Phys. Rev. D **89**, 063002 (2014).

[289] T. Fujita, R. Namba, Y. Tada, N. Takeda, and H. Tashiro, Journal of Cosmology and Astroparticle Physics **2015** (05), 054.

[290] A. Kandus, K. E. Kunze, and C. G. Tsagas, Physics Reports **505**, 1 (2011).

[291] K. Subramanian, Astronomische Nachrichten **331**, 110 (2010), https://onlinelibrary.wiley.com/doi/pdf/10.1002/asna.200911312 .

[292] K. Subramanian, Reports on Progress in Physics **79**, 076901 (2016).

[293] J. Martin and J. Yokoyama, Journal of Cosmology and Astroparticle Physics **2008** (01), 025.

[294] B. Himmetoglu, C. R. Contaldi, and M. Peloso, Physical Review D **80**, 10.1103/physrevd.80.123530 (2009).

[295] A. Brandenburg, K. Enqvist, and P. Olesen, Physical Review D **54**, 1291–1300 (1996).

[296] T. Vachaspati, Physics Letters B **265**, 258 (1991).

[297] A. P. Kazantsev, Soviet Journal of Experimental and Theoretical Physics **26**, 1031 (1968).

[298] N. I. Kleeorin, A. A. Ruzmaikin, and D. D. Sokoloff, in *Plasma Astrophysics*, ESA Special Publication, Vol. 251, edited by T. D. Guyenne and L. M. Zeleny (1986) pp. 557–561.

[299] A. P. Kazantsev, A. A. Ruzmaikin, and D. D. Sokolov, Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki **88**, 487 (1985).

[300] D. Grasso and H. R. Rubinstein, Physics Letters B **379**, 73–79 (1996).

[301] R. Durrer, P. G. Ferreira, and T. Kahniashvili, Phys. Rev. D **61**, 043001 (2000).

[302] N. Aghanim *et al.* (Planck), Astron. Astrophys. **641**, A1 (2020), arXiv:1807.06205 [astro-ph.CO] .

[303] S. Chabanier, M. Millea, and N. Palanque-Delabrouille, Mon. Not. Roy. Astron. Soc. **489**, 2247 (2019), arXiv:1905.08103 [astro-ph.CO] .

[304] Dark Energy Survey Collaboration, Phys. Rev. D **98**, 043528 (2018).

[305] B. A. Reid *et al.*, Monthly Notices of the Royal Astronomical Society **404**, 60 (2010), https://academic.oup.com/mnras/article-pdf/404/1/60/11177257/mnras0404-0060.pdf .

[306] B. Abolfathi *et al.*, The Astrophysical Journal Supplement Series **235**, 42 (2018).

[307] I. Esteban, A. H. G. Peter, and S. Y. Kim, Phys. Rev. D **110**, 123013 (2024).

[308] J. B. Muñoz, C. Dvorkin, and F.-Y. Cyr-Racine, Phys. Rev. D **101**, 063526 (2020), arXiv:1911.11144 [astro-ph.CO] .

[309] P. Ralegankar, E. Garaldi, and M. Viel, Journal of Cosmology and Astroparticle Physics **2025** (08), 011.

[310] N. Arkani-Hamed, L. Motl, A. Nicolis, and C. Vafa, Journal of High Energy Physics **2007**, 060–060 (2007).

[311] M. K. Seidel, J. Falcón-Barroso, I. Martínez-Valpuesta, S. Díaz-García, E. Laurikainen, H. Salo, and J. H. Knapen, Mon. Not. Roy. Astron. Soc. **451**, 936 (2015), arXiv:1504.08001 [astro-ph.GA] .

[312] S. Díaz-García, H. Salo, E. Laurikainen, and M. Herrera-Endoqui, Astron. Astrophys. **587**, A160 (2016), arXiv:1509.06743 [astro-ph.GA] .

[313] T. Kim, D. A. Gadotti, E. Athanassoula, A. Bosma, K. Sheth, and M. G. Lee, Mon. Not. Roy. Astron. Soc. **462**, 3430 (2016), arXiv:1607.08245 [astro-ph.GA] .

[314] A. Collier and A.-M. Madigan, Astrophys. J. **915**, 23 (2021), arXiv:2105.04698 [astro-ph.GA] .

[315] D. A. Marostica, R. E. G. Machado, E. Athanassoula, and T. Manos, Galaxies **12**, 27 (2024), arXiv:2405.17128 [astro-ph.GA] .

[316] M. Das, R. Ianjamasimanana, S. McGaugh, J. Schombert, and K. S. Dwarakanath, Proceedings of the International Astronomical Union **19**, 353–359 (2023).

[317] M. D. A. Orkney, E. Taylor, J. I. Read, M. P. Rey, A. Pontzen, O. Agertz, S. Y. Kim, and M. Delorme, Mon. Not. Roy. Astron. Soc. **525**, 3516 (2023), arXiv:2302.12818 [astro-ph.GA] .

[318] B. Keith, F. Munshi, A. M. Brooks, J. Van Nest, A. Engelhardt, A. Cruz, B. Keller, T. R. Quinn, and J. Wadsley, Astrophys. J. **986**, 138 (2025), arXiv:2501.16317 [astro-ph.GA] .

[319] J. Zavala and C. S. Frenk, Galaxies **7**, 81 (2019), arXiv:1907.11775 [astro-ph.GA] .

[320] A. H. G. Peter, M. Rocha, J. S. Bullock, and M. Kaplinghat, Mon. Not. Roy. Astron. Soc. **430**, 105 (2013), arXiv:1208.3026 [astro-ph.CO] .

[321] I. Helenius, PoS **DIS2018**, 113 (2018), arXiv:1806.07326 [hep-ph] .

[322] A. de Gouvêa, P. J. Fox, B. J. Kayser, and K. J. Kelly, Phys. Rev. D **104**, 015038 (2021), arXiv:2104.05719 [hep-ph] .

[323] K. J. Kelly and P. A. N. Machado, muBHNL, `https://github.com/kjkellyphys/muBHNL`.

[324] S. Ito, private communication.

[325] V. Skoutnev, A. Hakim, J. Juno, and J. M. TenBarge, The Astrophysical Journal Letters **872**, L28 (2019).

[326] S. P. Gary, The plasma dispersion function, in *Theory of Space Plasma Microinstabilities*, Cambridge Atmospheric and Space Science Series (Cambridge University Press, 1993) p. 170–171.

[327] A. Bret, A. Stockem, F. Fiuza, C. Ruyer, L. Gremillet, R. Narayan, and L. O. Silva, Physics of Plasmas **20**, 042102 (2013).

[328] J.-L. Vay, C. Geddes, E. Cormier-Michel, and D. Grote, Journal of Computational Physics **230**, 5908 (2011).

[329] J. E. C. Gliddon, Astrophys. J. **145**, 583 (1966).