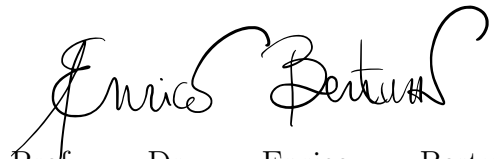


Universidade de São Paulo
Instituto de Física

Espalhamento Elástico Coerente de
Neutrino-Núcleo & Modelos Z' Leve: Formalismo,
Fenomenologia e Projeção de Sensitividade

Lucas Magno Dantas Ramos



Orientador: Prof. Dr. Enrico Bertuzzo

Tese de mestrado apresentada ao Instituto de Física da
Universidade de São Paulo, como requisito parcial para
a obtenção do título de Mestre em Ciências.

Banca Examinadora:

Prof. Dr. Enrico Bertuzzo - Orientador (Universidade de São Paulo)

Prof. Dr. Alex Gomes Dias (UFABC)

Prof. Dr. Ricardo Matheus D'Elia (IFT-UNESP)

São Paulo

2022

FICHA CATALOGRÁFICA
Preparada pelo Serviço de Biblioteca e Informação
do Instituto de Física da Universidade de São Paulo

Ramos, Lucas Magno Dantas

Espalhamento elástico coerente de neutrino-núcleo & modelos Z' leve: formalismo, fenomenologia e projeção de sensibilidade. / Coherent elastic neutrino-nucleus scattering & light Z' models: formalism, phenomenology and projected sensitivity. São Paulo, 2022.

Dissertação (Mestrado) – Universidade de São Paulo. Instituto de Física. Depto. de Física Matemática.

Orientador: Prof. Dr. Enrico Bertuzzo
Área de Concentração: Física

Unitermos: 1. Neutrinos; 2. Teoria Quântica de Campo; 3. Física de Partículas.

USP/IF/SBI-022/2022

University of São Paulo
Physics Institute

Coherent Elastic Neutrino-Nucleus Scattering & Light Z' Models: Formalism, Phenomenology and Projected Sensitivity

Lucas Magno Dantas Ramos

Supervisor: Prof. Dr. Enrico Bertuzzo

Thesis submitted to the Physics Institute of the University of São Paulo in partial fulfillment of the requirements for the degree of Master of Science.

Examining Committee:

Prof. Dr. Enrico Bertuzzo - Supervisor (Universidade de São Paulo)

Prof. Dr. Alex Gomes Dias (UFABC)

Prof. Dr. Ricardo Matheus D'Elia (IFT-UNESP)

São Paulo

2022

Acknowledgements

The content of this thesis is derived from research developed for the past two years under the supervision of Prof. Dr. Enrico Bertuzzo and collaboration with Dr. Giovanni Grilli di Cortona, which has already been published as an ArXiv Pre-print under the title “Probing light vector mediators with coherent scattering at future facilities” (code [arxiv:2112.04020](https://arxiv.org/abs/2112.04020)) and submitted to the Journal of High-Energy Physics, where it is currently under review. I’d also like to acknowledge the financial support from Comissão Nacional de Pesquisa (CNPq) throughout the entire duration of the work, under grant 131297/2020-1.

The world has been through a dramatic change during the period in which this thesis was developed, with countless lives lost and a radical breakdown of lifestyle brought about by a global pandemic. This work would have never been possible without the support of my family through these trying times, first and foremost. I’d also like to acknowledge and thank the support of all of my colleagues and professors at the Particle Physics group of the DFMA, as well as of the Neutrino Physics community as a whole, for their role in keeping our scientific activities afloat worldwide and rapidly adapting to the new reality of online environments. All the events we took part in were unique opportunities for learning and self-improvement, and I can confidently say I would not have been able to finish this work had I not gone through those.

I’d like to dedicate this thesis to all those who supported me, including aforementioned ones, my friends from T24 and from CCM and USP as a whole, all of my past teachers and professors who believed and pushed me in the way of Science, my long-time friends from Brasília who always offered their unyielding support, the MD and SL community, and countless more people who helped shape this journey to any degree.

I’d also like to dedicate this work to Favo 22, which will forever be my second home, and all the cherished memories that it housed, which probably shaped who I am today as a person and as a Physicist more so than any other.

Abstract

The confirmation of non-zero Neutrino masses defines Neutrino Physics phenomena as test cases of central importance to probe scenarios beyond the Standard Model. In particular, future experiments dedicated to the detection of Coherent Elastic Neutrino-Nucleus Scattering may be powerful tools for constraining light new physics. In this work, we introduce a complete approach in order to quantify predictions and derive sensitivity on light Z' mediators for two proposed experiments: a directional low-pressure Time Projection Chamber detector, ν BDX-DRIFT, that will utilize neutrinos produced at the Long Baseline Neutrino Facility; and several possible experiments to be installed at the European Spallation Source. We compare the results obtained with existing limits from fixed-target, accelerator, solar neutrino and reactor experiments. Furthermore, we show that these experiments have the potential to test unexplored regions that, in some case, could explain the anomalous magnetic moment of the muon or peculiar spectral features in the cosmic neutrino spectrum observed by IceCube.

Keywords: Neutrino Physics; CE ν NS; Simplified Models; Quantum Field Theory; Beyond Standard Model

Resumo

A confirmação de massas não-nulas para neutrinos define fenômenos relacionados à Física de Neutrinos como casos-teste de primordial importância para sondar cenários além do Modelo-Padrão. Em particular, experimentos futuros dedicados à detecção do Espalhamento Elástico Coerente de Neutrino-Núcleo podem ser ferramentas poderosas para impor limites em nova física leve. Neste trabalho, introduzimos uma abordagem completa para quantificar previsões e obter a sensibilidade de mediadores Z' leve para dois experimentos propostos: um detector direcional de Câmara de Projeção Temporal de baixa pressão, ν BDX-DRIFT, que utilizará neutrinos da Long Baseline Neutrino Facility; e vários possíveis experimentos a serem instalados na European Spallation Source. Comparamos os resultados obtidos com limites já existentes de alvo-fixo, aceleradores, neutrinos solares e experimentos de reator. Além disso, mostramos que esses experimentos têm o potencial de testar regiões inexploradas que, em certos casos, poderiam explicar o momento magnético anômalo do múon ou características peculiares do espectro cósmico de neutrinos observado pelo detector IceCube.

Palavras-chave: Física de Neutrinos; $CE\nu$ NS; Modelos Simplificados; Teoria Quântica de Campos; Além do Modelo-Padrão

List of Tables

Table 1	– Summary of properties for the detectors possibly employed at the ESS: target, mass, recoil energy threshold, width for the smearing and maximum recoil energy, steady-state background. The steady-state background includes the 4×10^{-2} reduction by the ESS duty factor.	83
Table 2	– Isotopes and respective nuclear radius values, taken from [45].	83
Table 3	– Standard Model prediction for the total number of events per year for all of the detector proposals.	83

List of Figures

Figure 1	– Particle content of the Standard Model (<i>Wikimedia Commons: MissMJ/CC-BY-SA-3.0</i>)	11
Figure 2	– Illustration of the neutrino oscillation phenomenon, for a 2-flavor case. θ parametrizes the misalignment between eigenstates, L is the baseline distance and $P_{e\ell}$ is the probability to observe a neutrino initially produced as ν_e as ν_ℓ . (Source: [17].)	13
Figure 3	– Effect of varying the baryonic density (c) and total matter density with fixed baryonic component (d) on the temperature power spectrum. The latter is equivalent to varying the dark matter density. Fiducial basis model presented in the original paper (Source: [26].)	15
Figure 4	– Planck 2018 Temperature Power Spectrum fit. (Source: [25])	15
Figure 5	– Some current limits on Spin-Independent Dark Matter Direct Detection Searches, scaled to a local Dark Matter density of $0.3 \text{ GeV}/c^3$. Plot built using the SuperCDMS Dark Matter Limit Plotter Tool. Data taken from [31–38].	17
Figure 6	– Schematic view of the $\text{CE}\nu\text{NS}$ process, an incident neutrino scatters off a nuclear target T , which recoils as if it was a single entity	26
Figure 7	– Feynman diagram for the $\text{CE}\nu\text{NS}$ process in the Fermi-like theory.	29
Figure 8	– Feynman diagram for the $\text{CE}\nu\text{NS}$ process in the Standard Model	33
Figure 9	– Schematic view of the collapse of the propagator for an Effective interaction in terms of the Feynman diagrams	39
Figure 10	– Diagram for the loop-induced mixing between Z' and photon, giving rise to the proton neutral current charge modification in the $Z'_{\mu-\tau}$ model.	47

Figure 11 – Schematic view of the CENNS-10 detector, including the fiducial LAr volume, overall detector vessel, vacuum chamber and shielding. Source: [48].	60
Figure 12 – Representation of the two scintillation pathways of Argon; direct dimer excitation (above) and ionized dimer formation and recombination (below). Source: [68].	61
Figure 13 – The COHERENT Collaboration experiments in the Neutrino Alley, at the SNS. Source: [47].	63
Figure 14 – Schematics of a pion decay at rest at the SNS. Source: [69].	63
Figure 15 – Components of the π DAR neutrino flux, given in arbitrary units. Source: [62].	65
Figure 16 – Expected number of events for the models of interest compared to the Standard Model. Only the first 6 bins out of 12 are shown, with the remaining having a total 0 of expected events.	67
Figure 17 – Sketch of the BDX-DRIFT-1m detector, where z denotes the beam axis; it was originally planned to detect Dark Matter (χ) produced via a beam dump experiment. Source: [73].	69
Figure 18 – Comparison between the Sulfur and the Carbon Form Factors, assuming the Helm Parametrization. Values for the nuclear mass and radii are taken from [45].	72
Figure 19 – Predicted Neutrino flux at the Near Detector Facility of the DUNE complex, for several positions with respect to the beam axis. Source: [72].	73
Figure 20 – Expected number of CE ν NS events in ν BDX-DRIFT as a function of the pressure configuration (in Torr) for a data-taking period of 7 years. The vertical grey line shows the pressure value where the Carbon contribution overtakes the Sulfur one. The number of events is computed in the SM and, for illustrative purposes, we do not impose any cut on the maximum recoil energy.	75
Figure 21 – Expected number of CE ν NS events in ν BDX-DRIFT as a function of the pressure configuration (in Torr) for a data-taking period of 7 years, for benchmark choices of $g_{Z'}$, $m_{Z'}$ in each of the models of interest.	76

- Figure 22 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the universal Z' model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green regions are excluded at 90% C.L. by CONNIE [79] and COHERENT [56] assuming dominant Z' decays to SM states. The dashed black curve shows the sensitivity of ν IOLETA [86]. The gray shaded areas are excluded at 90% C.L. by BaBar [87] and NA64 [88, 89] assuming that the Z' decays dominantly in invisible dark sector states. 86
- Figure 23 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $B-L$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green regions are excluded at 90% C.L. by NA64 [95], E141 [96], Orsay [97], U70 [98], COHERENT [56], NA48/2 [99], KLOE [100, 101], BaBar [102] and LHCb [103] assuming dominant Z' decays to SM states. 89
- Figure 24 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $B-L$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The gray shaded areas are excluded at 90% C.L. by BaBar [87] and NA64 [88, 89] assuming that the Z' decays dominantly in invisible dark sector states. 91
- Figure 25 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $L_\mu - L_\tau$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green shaded areas are excluded at 95% C.L. by BaBar [107], ATLAS [108, 109], CMS [110], CCFR [111, 112] and at 90% C.L. by Borexino [113, 114] assuming that the Z' decays to muons. In the red region the model explains at 2σ the anomalous magnetic moment of the muon [7, 115, 116]. The yellow band shows the region that can explain the Xenon1T excess [117] in some specifically extended $L_\mu - L_\tau$ model [118], while the black diamonds refer to a model [119] that explain the cosmic neutrino spectrum features observed by IceCube [120]. 92

Contents

	Acknowledgements	1
	Abstract	2
	Resumo	3
	List of Tables	4
	List of Figures	4
	Contents	7
I	CEνNS & LIGHT Z' MODELS - INTRODUCTION, FORMALISM AND OBJECTIVES	9
1	INTRODUCTION	10
1.1	The Standard Model and the Future	10
1.2	Coherent Scattering	17
1.3	CEνNS	23
2	CEνNS FORMALISM	26
2.1	Process Kinematics	26
2.2	Standard Model Cross Section for CEνNS	28
2.3	A Brief Interlude on Effective Field Theory	38
2.4	Measurable Quantity - Event Rate	40
3	LIGHT Z' MODELS	41
3.1	Definition and Assumptions	41
3.2	Neutrino Non-Standard Interactions	42
3.3	Simplified Model Approach	44
3.4	Models of Interest and Revised Cross Section	45
4	HYPOTHESIS TESTING AND WILKS' THEOREM	49
4.1	Definitions and Formalism - Brief overview	49
4.2	Testing Algorithm	56

II	PHENOMENOLOGICAL APPLICATIONS - CEνNS EXPERIMENTS	58
5	BENCHMARK CASE - CENNS-10 DETECTOR	59
5.1	Detector properties	59
5.2	Neutrino Flux	62
5.3	Event Rate computation	65
6	EXPERIMENTAL PROPOSAL - νBDX-DRIFT	68
6.1	Detector Properties	68
6.2	Neutrino Flux and Region of Interest	72
6.3	Event Rate	73
7	EXPERIMENTAL PROPOSAL - ESS-CEνNS DETECTORS	77
7.1	ESS vs SNS	77
7.2	Proposed Detector Setups	79
7.3	Neutrino Flux and Region of Interest	81
7.4	Event Rate	82
8	SENSITIVITY ANALYSIS	84
8.1	Sensitivity for the Universal Z' Model	85
8.2	Sensitivity for the $B - L$ model	90
8.3	Sensitivity for the $L_\mu - L_\tau$ model	91
9	CONCLUSION	93
	Bibliography	95

Part I

CE ν NS & Light Z' Models - Introduction, Formalism and Objectives

1 Introduction

1.1 The Standard Model and the Future

The past decade has been special for the Particle Physics community, with the culmination of a near century-long program to formalize fundamental interactions of nature under a single framework of Quantum Field Theory. While it is difficult to define a single starting point, one could trace a crucial benchmark to Paul Dirac’s work in the late 1920’s on a consistent quantum theory of emission, absorption and radiation [1]. In this seminal paper, Dirac formalized a consistent description of interactions between charges and the electromagnetic field in first order of perturbation theory, while describing both as quantum objects, explaining the phenomenon of spontaneous emission. Furthermore, it also is the first use of the term “Quantum Electrodynamics”, coining the name for the theory later described by Richard Feynman as “the jewel of Physics” [2] for the accuracy of its predictions, and effectively inaugurating the field which would be the basis for all subsequent development of the currently known Standard Model of Particle Physics. Several decades later, the discovery of the Higgs Boson by the Large Hadron Collider collaborations in 2012 [3, 4] marked the final piece in this gargantuan effort.

The Standard Model is a gauge theory based on a $SU(3)_C \times SU(2)_L \times U(1)_Y$ symmetry group, bringing into a unified framework both strong and electroweak interactions, describing dynamics under these forces for all known elementary particles. Below the electroweak scale ($\Lambda_{EW} \approx O(100 GeV)$), the theory undergoes Spontaneous Symmetry Breaking as the Higgs field acquires a vacuum expectation value, such that the ground state of the theory is now symmetric under $SU(3)_C \times U(1)_{em}$, reproducing the familiar interactions of electromagnetism via the unbroken $U(1)$ group. The mechanism also generates mass terms for all massive fermions and for the weak bosons (W, Z), which were forbidden in the unbroken phase by $SU(2)_L$ gauge invariance.

This is all summarized (in a somewhat compact notation) in the Standard Model Lagrangian:

$$\begin{aligned}
 \mathcal{L} = & \sum_{\psi} \bar{\psi}_i i \gamma^{\mu} D_{\mu} \psi_i - \sum_F \frac{1}{4} F_{\mu\nu}^a F^{a,\mu\nu} \text{ (Gauge sector)} \\
 & + |D_{\mu}(H)|^2 - \mu^2 |H^{\dagger} H| - \lambda^2 |H^{\dagger} H|^2 \text{ (Higgs sector)} \\
 & - Y_{i,j}^u \bar{Q}^i \tilde{H} u_R^j - Y_{i,j}^d \bar{Q}^i H d_R^j - Y_{i,j}^l \bar{L}^i H e_R^j + h.c. \text{ (Yukawa sector)}
 \end{aligned} \tag{1.1}$$

where ψ runs over all chiral fermion components, F runs through the three types of gauge vector fields (G, W, B) and i, j are flavor indices summed over the three generations. The covariant derivatives D_{μ} contain appropriate terms for each fermion and the Higgs




Figure 1 – Particle content of the Standard Model (*Wikimedia Commons: MissMJ/CC-BY-SA-3.0*)

according to their respective gauge charges. The entire particle content and its respective electric charges and interactions of the fermions with the gauge fields are summarized in Figure 1.

Throughout its history, the Standard Model proved to be an almost unparalleled success both in the accuracy of its predictions and its own robustness. However, this did not mark the end of the journey. By design, the Standard Model was never meant to be a final “theory of everything”, since it does not include gravity in its description; it was always meant to have a range of validity restricted at least by the Planck scale, where gravity interactions become significant and one cannot neglect their effects. Even setting aside gravity, there are still many open problems in the field (such as Baryon Asymmetry [5], Strong CP problem [6], Anomalous magnetic dipoles [7]), emphasizing that there is still more to explore. And yet, for the most part, so far the community has consistently failed to gauge exactly what lies beyond it.

Precision tests and intricacies of Flavor Physics point to some promising roads in which the Standard Model may eventually fail [8, 9], but the evidence so far is not only scarce, but also usually dependent on assumptions or details that may still not tell the full story. The recently observed 4.2σ deviation of the predicted value of anomalous magnetic moment of the muon [7] is such an example. One of the strongest deviations and

most promising results in the last 10 years, it is still put into question due to different results for some of the crucial hadronic contributions for the Standard Model prediction put forward by lattice groups [10], which could significantly reduce its significance.

Many extensions for the Standard Model have been proposed in order to solve some of its most pressing concerns, like the Hierarchy Problem [11], besides other aforementioned unsolved puzzles. However, experimental results in general fail to favor any of them consistently over the Standard Model predictions. Nevertheless, so far we've been at best scratching the surface of all the possible scenarios, and a growing number of options and parameter choices are discarded on a daily basis. A prominent example would be Supersymmetry-based extensions [11]. One of the most popular solutions to the Hierarchy problem due to providing a natural way to keep radiative corrections to the Higgs mass under control, new "superpartner" particles predicted by them were expected to show up in the TeV scale, probed in the Large Hadron Collider runs. However, no definitive evidence of their discovery has been observed so far [12, 13], with further searches gradually excluding more of the diminishing available phase space.

This is all to say that it is crucial to look at concrete evidence in which the Standard Model is not adequate, in order to probe which is the way ahead to tread. While there are many micro-anomalies within 3σ significance range [8], whose relative importance varies considerably among different research groups, at present there are only two prevalent topics which are consensus for an absolute majority of the community: neutrino masses and the existence of particle-based Dark Matter. Both of these are not included in the known Standard Model framework, with the former already confirmed by experiments and the latter strongly implied.

Neutrino Masses

In the Standard Model, neglecting the contribution of quark condensation phenomena, all fermion masses originate from Yukawa terms once the Higgs hits its vacuum expectation value:

$$\mathcal{L}_{\text{Yukawa}} \supset -Y_{ij}^f \bar{f}_L^i H f_R^j + h.c. \xrightarrow{H \rightarrow \langle H \rangle} -\frac{v}{\sqrt{2}} Y_{ij}^f \bar{f}_L^i f_R^j + h.c., \quad (1.2)$$

where Y are the Yukawa couplings, f is a particular fermion, i and j are flavor indices and L,R denote the chirality. This originates Dirac mass terms for all fermions in the Standard Model with one exception: the neutrinos, which do not have a right-handed component. Thus, within the framework, they should be massless. That, however, is not the case in nature, as shown by the 2015 Nobel Prize Winner works on neutrino oscillations [14, 15].




Figure 2 – Illustration of the neutrino oscillation phenomenon, for a 2-flavor case. θ parametrizes the misalignment between eigenstates, L is the baseline distance and $P_{e\ell}$ is the probability to observe a neutrino initially produced as ν_e as ν_ℓ . (Source: [17].)

The phenomenon is characterized by a non-conservation of lepton number through the propagation of a neutrino; in simple terms, this means that a neutrino, produced initially in a particular leptonic flavor state (electron, muon or tau), is observed in a different one after it has propagated. At its core, this means that there is a misalignment between the interaction eigenstates (which are the flavor ones) and the propagation eigenstates (which are the mass eigenstates), which can only exist if at least two out of the three neutrino species are massive.

In this case, each flavor is produced as a coherent mixture of different mass eigenstates, with each component propagating with a distinct evolution. As a result, the initial mixture changes, affecting the probability of detecting the neutrino in a given (flavor) state. This is illustrated for a simplified scenario with only two active neutrino flavors propagating through vacuum in Figure 2. A formal characterization of the theory and the many inherent subtleties in its treatment is beyond the scope of this work, but an excellent review can be found in [16].

Having been confirmed repeatedly in a multitude of scenarios ranging from solar observations to controlled baseline experiments, it is no exaggeration to say that the existence of neutrino masses is currently the most important topic in Particle Physics with respect to what may lie beyond the current Standard Model. While many mechanisms have been proposed to explain the origin of neutrino masses (Seesaw Mechanisms being the most popular approach [17]), no experiment so far has been able to confirm a particular model. As such, any new information on deviations from the Standard Model concerning the neutrino sector are incredibly valuable to guide the way for future developments in the field.

Dark Matter

In the same vein, another crucial component of nature which has not been contemplated within the Standard Model is Dark Matter. Initially proposed by the work of Fritz Zwicky in 1937 [18] concerning the Virial relation in galaxy clusters, and later expanded by Vera Rubin's and collaborators' work on galaxy rotation curves [19] in 1980, it was characterized as a massive component of matter which does not emit in the electromagnetic spectrum. Further studies with cluster collision events such as the Bullet Cluster [20], observing the mismatch between the visible matter distribution and the gravitational lensing one, showed it to also be nearly collisionless (both with itself and with ordinary/"barionic" matter).

While many proposals to explain the phenomenon came up in the 80's and later to explain the findings - such as primordial black holes (PBHs)[21], massive compact dense objects (MACHOs)[22], modifications to newtonian mechanics (MOND)[23] and modified gravity theories (Scalar-Tensor, f(R) and TeVeS, to name a few) [24] - detailed observations of the power spectrum of the CMB revealed themselves incredibly consistent with the existence of a particle-like, massive and collisionless component in the early universe [25]. Furthermore, it also made possible to estimate the contribution of such a component to the overall matter-energy content of the universe, latest observations yielding surprising 27% of Dark Matter against just 5% of barionic matter [25]. Figure 3, taken from Plate 4 of [26], illustrates the influence of varying the matter content on the CMB temperature Power Spectrum, which leads to the fit with observed Planck data in Figure 4.

The remainder is dubbed "Dark Energy", and while also outside of the Standard Model framework, we shall not discuss it here due to the lack of a clear direction of the community in how to treat it.

Currently, besides the aforementioned characteristics, implications for the formation of halos and galaxies in the Standard Cosmological Model (Lambda-CDM) also require that dark matter is cold and stable (a brief overview and alternative scenarios can be found in [27]). Should its production be due to a thermal mechanism, the former means it should be non-relativistic at the time it decouples from the primordial bath, and the later that it does not decay at timescales around the current age of the universe. The first condition in particular discards Standard Model neutrinos, the only neutral, massive and weakly-interacting fermion, as viable candidates.

Those conditions lead to an interesting realization, if we keep the assumption of Dark Matter production in the early universe following a thermal scenario. At first, it would be in equilibrium with the Standard Model particles in the primordial bath at temperatures above its mass scale, with its number density fluctuating due to a balance




Figure 3 – Effect of varying the baryonic density (c) and total matter density with fixed baryonic component (d) on the temperature power spectrum. The latter is equivalent to varying the dark matter density. Fiducial basis model presented in the original paper (Source: [26].)

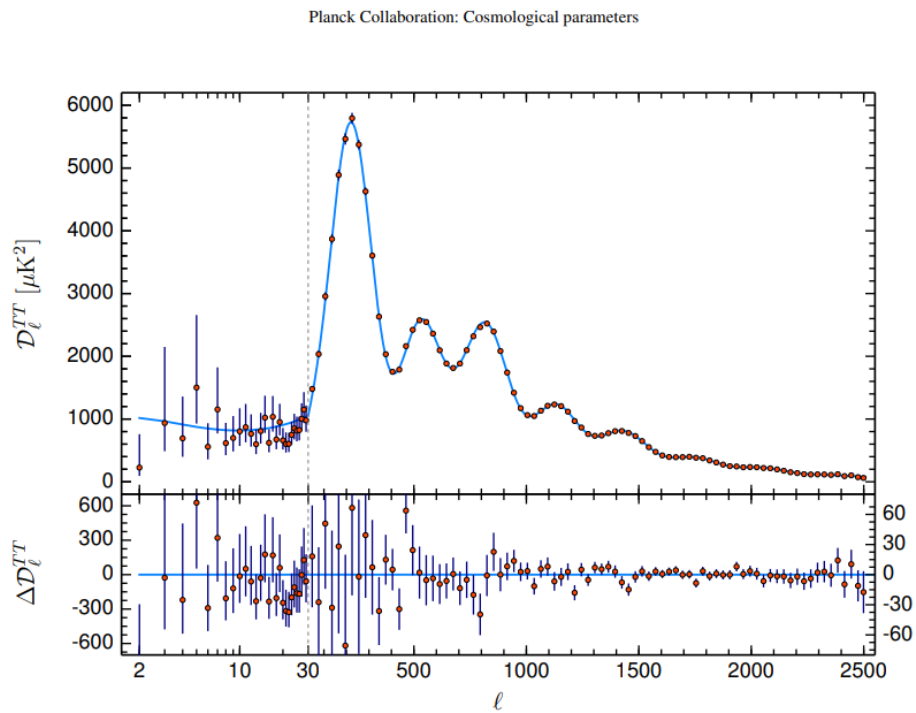


Figure 4 – Planck 2018 Temperature Power Spectrum fit. (Source: [25])

between annihilation to Standard Model states and production in the reverse process. However, as the universe expands and cools down, there would be a point where the expansion rate dominates the rate of interactions, making it so that the collision rate goes to zero and the Dark Matter candidate decouples from the bath. Given its stability and collisionless nature, the number abundance of a single dark matter species would be “frozen-out” as it decouples, and remain the same henceforth. Since it is possible to track this quantity using the CMB power spectrum, one can define a “Relic Abundance” that a theory must correctly predict, if it admits a thermal production mechanism. Following this “Freeze-out” scenario, the abundance is set by the thermally-averaged annihilation cross section.

It turns out that one can get the correct relic abundance, within few orders of magnitude, with a candidate with mass and coupling of the order of the electroweak scale (~ 100 GeV). These proposed candidates were named “WIMPs” (Weakly Massive Interacting Particles), and popularly this coincidence was dubbed the “WIMP Miracle”. Due to the inherent familiarity of working with the electroweak sector in the Standard Model, much of the research in the area realized in the last decade has focused on WIMPs, initially trying to consistently introduce a Electroweak-like extra sector with small interactions with its Standard Model counterpart.

A more in-depth look is provided in [28], and while there are both alternative thermal (eg. SIMPs [29]) and non-thermal (Axion, ALPs, FIMPs, etc [30]) scenarios, WIMP + Freeze-Out has been the most well-described and explored approach in the literature.

Many detection techniques were introduced to try and detect such candidates, which are roughly grouped within direct detection (detection of a Dark Matter particle), indirect detection (observation of Standard Model astroparticle signals from a Dark Matter process) and collider searches (production of Dark Matter). While developments have been made in all three groups, direct detection has been a particularly active field both theoretically and experimentally, and efforts have already discarded a large part of the available phase space for WIMP scenarios, as illustrated by Figure 5. While sensitivity thresholds keep getting pushed, the yellow region in the plot defines a region where experiments could be able to detect solar neutrinos, configuring an irreducible background for most dark matter searches.

Our Goals

In summary, we are currently at a crossroads in the field of Particle Physics, celebrating a resounding success from continued efforts formalizing and verifying predictions but with no clear picture of what goes beyond. In this situation, neutrino physics and dark matter become invaluable guiding beacons in order to light what may lie ahead. Both of




Figure 5 – Some current limits on Spin-Independent Dark Matter Direct Detection Searches, scaled to a local Dark Matter density of $0.3 \text{ GeV}/c^3$. Plot built using the [SuperCDMS Dark Matter Limit Plotter Tool](#). Data taken from [31–38].

these, however, are impaired by the massive challenge presented in designing experiments able to detect signals despite their reduced coupling to Standard Model states. In order to probe this type of “Beyond the Standard Model” Physics, we shall introduce in the next section the notion of Coherent Scattering, justifying how it might prove itself useful for such a goal.

1.2 Coherent Scattering

In Physics, the notion of Coherence is rooted in the capacity of two waves to interfere, due to a constant phase difference between them. Fully coherent waves can produce a standing wave state. While primarily applied to the field of Optics and as a general wave dynamics concept, it is also deeply significant for Quantum Mechanics and beyond, due to the wave-like nature of quantum states.

A pedagogical example

Likewise, one can interpret the notion of scattering within this context in terms of coherence - whether there is an interference among the partial amplitudes of a given process. This is particularly significant when describing a scenario of a elementary probe colliding into a composite target. A simplified view, described by Freedman et al. in [39], is very useful to present some intuition on the process:

Suppose a target made up of A components/constituents, respectively at positions

$\{\vec{x}_j\}, j = 1, \dots, A$, and an incident particle with momentum \vec{k} , which will scatter with outgoing momentum \vec{k}' . From the superposition principle, the amplitude of the process $M(\vec{k}, \vec{k}')$ can be written as a sum over the contributions due to each component:

$$M(\vec{k}, \vec{k}') = \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_j} \quad (1.3)$$

where each partial amplitude is weighted by a factor which takes into account the relative phase difference of the wave scattering at each \vec{x}_j . We are assuming there are no double-scatters in this scenario, which will be a reasonable assumption in the cases we'll focus on later.

The differential cross section of the process is given by this scattering amplitude, as

$$\frac{d\sigma}{d\Omega} = |M(\vec{k}, \vec{k}')|^2. \quad (1.4)$$

A key realization is that if the relative phase difference is either zero or very small, or constant and periodic for all of the partial amplitudes, they can be added coherently (in this case interfering constructively), such that, defining $\bar{M} = \frac{1}{A} \sum_{i=1}^A M_j$:

$$\frac{d\sigma}{d\Omega} = A^2 |\bar{M}(\vec{k}, \vec{k}')|^2. \quad (1.5)$$

In the case where the target is made of only one type of constituent, this further reduces to:

$$\frac{d\sigma}{d\Omega} = A^2 |M_j(\vec{k}, \vec{k}')|^2 = A^2 \left(\frac{d\sigma}{d\Omega} \right)_{\text{Point-like}} \quad (1.6)$$

which effectively means we are boosting the cross section with respect to the one for a single component by the number of components squared. If there are multiple types of constituents, this scaling isn't exact, but barring cases with a very fine cancellation between the amplitudes M_j of each constituent, one would still expect a significant enhancement.

There are two ways to satisfy the conditions for this type of effect, which we'll name as a "coherence enhancement":

- either the target has some organized periodic substructure imposed on the position vectors, such that exist integers $N(j, l)$ that satisfy

$$i(\vec{k}' - \vec{k}) \cdot \vec{x}_j = i [(\vec{k}' - \vec{k}) \cdot \vec{x}_l + 2\pi N(j, l)], \forall (j, l) = \{1, \dots, A\}, \quad (1.7)$$

since this implies

$$\begin{aligned}
M(\vec{k}, \vec{k}') &= \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_j} \\
&= \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i[(\vec{k}' - \vec{k}) \cdot \vec{x}_i + 2\pi N(j,l)]} \\
&= \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} e^{i2\pi N(j,l)} \\
&= e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i2\pi N(j,l)} \xrightarrow{1} \\
\implies \frac{d\sigma}{d\Omega} &= A^2 \left| e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \bar{M} \right|^2 = A^2 |\bar{M}|^2 \left| e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \right|^2 \xrightarrow{1}
\end{aligned} \tag{1.8}$$

which could happen, for example, in the case of scattering on a crystal lattice using a source with momentum fine-tuned to reproduce equation (1.7);

- or, more generally, if the relative phase is small enough to be effectively zero for all terms, as it implies

$$\begin{aligned}
M(\vec{k}, \vec{k}') &= \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_j} \\
&= e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \left(M_i(\vec{k}, \vec{k}') + \sum_{j=1, j \neq i}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot (\vec{x}_j - \vec{x}_i)} \right) \\
\stackrel{\text{if } |(\vec{k}' - \vec{k}) \cdot (\vec{x}_j - \vec{x}_i)| \approx 0}{\longrightarrow} &\approx e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \left(M_1(\vec{k}, \vec{k}') + \sum_{j=1, j \neq i}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot (\vec{x}_j - \vec{x}_i)} \right) \xrightarrow{1} \\
&= e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \left(\sum_{j=1}^A M_j(\vec{k}, \vec{k}') \right) \\
\implies \frac{d\sigma}{d\Omega} &= A^2 \left| e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \bar{M} \right|^2 = A^2 |\bar{M}|^2 \left| e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_i} \right|^2 \xrightarrow{1}
\end{aligned} \tag{1.9}$$

like in the previous case.

Let us define $\vec{q} = \vec{k}' - \vec{k}$, the momentum exchanged in the process, and $R = \max_{j,l} |\vec{x}_j - \vec{x}_l|$, the size of the target. Clearly,

$$|\vec{q} \cdot (\vec{x}_j - \vec{x}_l)| \leq |\vec{q}| |\vec{x}_j - \vec{x}_l| \leq |\vec{q}| R \equiv qR \tag{1.10}$$

and thus, the previous condition is equivalent to

$$qR \ll 1 \tag{1.11}$$

in appropriate units. We shall define this expression as the **Coherency Condition** for scattering. Since the value of R is fixed for a given target, in practice this means that

we can expect for coherent scattering to take place for low momentum regimes, with the scale defined by the inverse of the target size, $q \lesssim \frac{1}{R}$.

This simple case can be pushed a bit further still, with a small generalization; for simplicity, let's consider a target made of only one type of component, so $M_j \equiv M_0$ for all j . We can rewrite (1.3) for a single constituent type in a more general way introducing the "component density" $\rho(\vec{x}) = \sum_{j=1}^A \delta(\vec{x} - \vec{x}_j)$, so it becomes

$$\begin{aligned} M(\vec{k}, \vec{k}') &= \sum_{j=1}^A M_j(\vec{k}, \vec{k}') e^{i(\vec{k}' - \vec{k}) \cdot \vec{x}_j} \\ &= M_0(\vec{k}, \vec{k}') \sum_{j=1}^A e^{i\vec{q} \cdot \vec{x}_j} \\ &= M_0(\vec{k}, \vec{k}') \int d^3x e^{i\vec{q} \cdot \vec{x}} \rho(\vec{x}) \equiv M_0(\vec{k}, \vec{k}') F(\vec{q}) \end{aligned} \quad (1.12)$$

which tells that we're dealing with the Fourier transform of some component density in the target, which is essentially the concept of a Form Factor, which we denote as a function of the exchanged momentum. Using this, we can write

$$\frac{d\sigma}{d\Omega} = \frac{d\sigma}{d\Omega} \Big|_{\text{Point-like}} |F(q)|^2 \quad (1.13)$$

The equation hints about some intuition on the nature of the process. When the exchanged momentum tends to zero, entering in the fully coherent regime, the form factor just results in its normalization, which in our toy case it's simply the number of components - that is to say, the experiment cannot say anything about the substructure of the target, just its size. In other words, the target behaves like a single object, with all pieces recoiling together.

This is entirely consistent with what would be expected when framing the question of coherence as the probing of inner structure; from Optics, we know that the resolution of an experiment is limited by the diffraction limit, so one cannot probe distances much smaller than the wavelength of the incident light. For a matter probe, the relevant wavelength is the de Broglie wavelength $\lambda = \frac{1}{|\vec{q}|}$, in natural units, so if we impose the coherency condition, we obtain

$$\lambda = \frac{1}{q} \xrightarrow{q=\lambda^{-1}} qR \ll 1 \implies \frac{R}{\lambda} \ll 1 \quad (1.14)$$

which is equivalent to requiring that the probe has wavelength larger than the size of the target. Thus, in this simplified example, Eq. (1.13) tells us that, up to a normalization which can be extracted from F , we are parametrizing the process as a point-like interaction modulated by some correction encoded in the Form Factor which addresses the substructure of the target.

Multipole Expansion and Form Factors

While quite simple and elegant, the previous pedagogical example helped to highlight some important features. If the coherency condition from (1.11) is satisfied, we can expect an enhancement for the cross section with respect to the one corresponding to a single component. Though, the exact expression will depend on the form factor, which modulates this enhancement while tracking to some degree the substructure of the target accessible by the probe, as a function of the exchanged momentum.

Things, however, are not quite so simple. In general, the target can have some complex substructure tying its components together, and the resulting amplitude will not be computable in a simple, closed form. The proper procedure to be followed is quite long and will not be approached in full here; excellent modern and complete descriptions of the formalism can be found in [40, 41], with applications to Dark Matter scattering, and a recent practical application to neutrino scattering is also shown in [42]. We present a brief overview of the procedure involved:

- For most practical applications, we're interested in the case where the exchanged momentum is small to take advantage of the coherency enhancement, and thus we can assume that $q \ll m_T$, with the latter being the mass of the target nucleus. We should then take the non-relativistic limit of the amplitude during the procedure, or alternatively work with an effective non-relativistic description of the operators;
- Decompose the relevant operators for the process of interest, in a basis of products of operators acting only on the internal structure of the target, times others acting only on the remaining degrees of freedom;
- Perform a multipole expansion expanding the resulting matrix elements over a basis of spherical tensors, assuming the target's internal state has well-defined angular momentum;
- Apply the Wigner-Eckart Theorem to reduce the matrix elements, and proceed to average over the spins;
- Write the operators in terms of Response functions, applying selection rules to cancel the irrelevant contributions and compute the rest of Clebsh-Gordan coefficients;
- Compute the reduced matrix elements, which will often require to assume a descriptive model for the target, in order to construct the total angular momentum eigenstates and get any other relevant quantum numbers.

While the procedure is quite complex, in general it is still possible to sum the contributions of the response functions to define a form factor (usually truncating the

expansion to the leading multipoles). This way, eq. (1.13) stills holds, albeit with modifications depending on the target composition, and with the caveat that $F(q)$ will also encode effects due to the operators related to the probe, and not only to the target. This means that, for example, one could not factor the cross section into a term with charges from the underlying interaction and another with the compositeness of the target encoded in the form factor in the most general scenario [42].

This formalism also reveals some non-trivial behavior concerning the loss of coherence encoded in the form factor. Among the response functions $R_{\mathcal{T}}(q)$ (where the subscript denotes the corresponding generating spherical tensor \mathcal{T}) in a general case, not all (once normalized) satisfy $R_{\mathcal{T}}(0) = 1$ with a decaying behavior as the exchanged momentum grows. In fact, some of the responses are fully non-coherent, in the sense that $\lim_{q \rightarrow 0^+} R_{\mathcal{T}}(q) = 0$. Still, they might interfere with fully coherent responses to produce a semi-coherent enhancement effect, as described in [42], and thus become relevant at higher momentum exchanges close to the coherency limit given by eq. (1.11).

Nuclear Targets and Helm Form Factor

Given all of those considerations about the procedure in the previous section, computing the necessary responses for a nuclear target can be incredibly difficult; intricacies of the formalism aside, at some point it will be necessary to assume a nuclear model and compute the density matrix (which is very non-trivial for most elements) besides endless scores of Clebsh-Gordan coefficients. Still, it might be worth the work involved, given that we expect a very large enhancement of the order of A^2 for isospin-invariant interactions, or at worst some function $f(A, Z)^2 > 1$, except in the case where there is some very fine cancellation. Furthermore, most of the relevant coherent scattering experiments will rely on detectors filled with nuclear targets. Thus, it would be convenient to obtain at least a practical approximation of the form factor in order to perform predictions.

In the literature, the one adopted most often is the **Helm Form Factor**[43], characterized by a radius parameter $c(A)$ (corresponding to the radius to half-maximum of the charge distribution) and a surface thickness s (corresponding to the thickness separating boundaries encompassing 10% to 90% of the maximum of the charge distribution). Those parameters were initially fitted using elastic scattering of electrons off even-even nuclei, following a Born approximation for the cross section, but have since been refined with other alternative parametrizations. Instead of $c(A)$, current literature usually works with the diffraction Radius R_0 , expressed in terms of the measured nuclear root-mean-squared (charge) radius R_N , as will be shown next.

The Helm form factor is particularly convenient due to its closed analytic form

$$F_{\text{Helm}}^N(q) = 3 \frac{j_1(qR_0^N)}{qR_0^N} e^{-q^2 s^2/2}, \quad (1.15)$$

with [44]

$$R_0^N = \sqrt{\frac{5}{3}R^N - 5s^2}, \quad (1.16)$$

$$s \simeq 0.9\text{fm} .$$

and j_1 being the spherical Bessel function of the first kind, $N = \{p, n\}$ defines a particular nucleon and R_N is the nuclear rms radius for a given nucleon distribution. Whenever necessary, we shall refer to [45] for the reference values.

Given that only the coherent responses, by definition, are maximal at zero-momentum exchange, it's expected that uncertainties due to the form factor are small for very low momentum. Therefore, most applications of coherent scattering off nucleus should not be largely affected by the choice of the Helm one except close to the limit in eq. (1.11). We shall comment briefly about this again later in Part II.

Summary

Summarizing, the phenomenon of coherent scattering allows for a constructive interference between amplitudes of a composite target, enhancing the cross section in comparison to a single-component process. In the case of a nuclear target, if the coherence condition is satisfied, the enhancement factor will be approximately proportional to the mass number squared of the nucleus in question, massively boosting the likelihood of interaction in comparison to a single-nucleon process. For neutrino and dark matter experiments, which deal with very small couplings and therefore low event rate expectations, this opens up the possibility to probe a larger region of the available phase space accessible by the lower energy range processes.

We shall focus our attention on one such coherent scattering scenario, of a neutrino off a nucleus, known as $\text{CE}\nu\text{NS}$; this will be the our main tool throughout this text to explore the question of what may lie ahead in our search for Physics Beyond the Standard Model.

1.3 $\text{CE}\nu\text{NS}$

An acronym to ‘‘Coherent Elastic Neutrino-Nucleus Scattering’’, $\text{CE}\nu\text{NS}$ was first predicted by Friedmann in his 1974 paper [46], but eluded discovery for over four decades, with the first confirmed observation only happening in 2017, by the COHERENT Collaboration [47]. The reason is the immense difficulty involved in observing the phenomenon; detection requires to measure the energy of the recoiling nuclear target inside the detector, which is incredibly small. Following from the last section, the limited momentum window imposed by the coherence condition (1.11) plus assuming the non-relativistic limit implies

recoil energies of the order of

$$qR \ll 1 \implies E_R \simeq \frac{q^2}{2m_T} \lesssim \frac{1}{2m_T R^2} \xrightarrow[m_T \approx Am_p]{R \approx 2R_A \approx 2r_0 A^{1/3}} \frac{1}{8m_p r_0^2 A^{5/3}} \quad (1.17)$$

which, for a reasonably bulky nucleus with $A = 125$, and writing units explicitly, is

$$E_R \lesssim \frac{(hc)^2}{8m_p r_0^2 A^{5/3}} = \frac{(1.239841 \times 10^{-6})^2 \text{eV}^2 \cdot \text{m}^2}{8(0.938 \text{GeV})(1.44 \text{fm}^2)5^5} \approx 50 \text{ keV} \quad (1.18)$$

and generally of the order of $O(10 \text{ keV})$ for most viable choices of target. This requires extremely precise experimental control of uncertainties and extremely low detection thresholds. It is no wonder that it remained as a prediction for such a long time, with the necessary technology still far in the future at the time of its proposal; Freedman even remarked, in [46]:

Our suggestion may be an act of hubris, because the inevitable constraints of interaction rate, resolution, and background pose grave experimental difficulties for elastic neutrino-nucleus scattering.

With its discovery established [47, 48], the community is opening the doors to a precision phase of $\text{CE}\nu\text{NS}$, designing multiple approaches to measure it with increased statistics and higher accuracy. Currently, the interest in $\text{CE}\nu\text{NS}$ is manifest in mainly three fronts:

- Within the paradigm of the Standard Model, it allows to measure the sine-squared of the Weinberg/Weak mixing angle ($\sin^2 \theta_W$) at very low energies, as well as other experimentally challenging nuclear properties, such as the neutron skin [49];
- If the necessary parameters and emerging relevant properties can be determined independently (for instance, by the PREX-II collaboration [50]), it can serve as a powerful tool to probe New Physics scenarios, which we expect to further enhance the cross section in some sector of the parameter space, benefiting from the coherent enhancement to yield potentially measurable results.
- In the context of Dark Matter direct detection, it serves as a source of irreducible background due to the scattering of solar neutrinos, possibly limiting the discovery potential of the next generation of experiments as they push sensitivity limits down into this so called “neutrino floor” sector of the phase space. We refer to [51] for a description of this effect in a pure SM context, or [52–55] for the effect of NP.

As an added bonus, much of the formalism can also be carried to the process of coherent Dark Matter scattering[41], which needs to be taken into account properly for true exclusion limits in direct detection.

The rest of this work will be organized as follows: In the next chapter we shall formalize the concept of $\text{CE}\nu\text{NS}$, combining the definitions we presented thus far with appropriate mathematical descriptions. Next, we shall see how the results derived can be modified when considering Physics beyond the Standard Model, focusing on the case of light Z' models. We'll then close Part I by addressing the statistical formalism needed to extract information from experimental data. Finally, in Part II, we shall describe a few proposals for future $\text{CE}\nu\text{NS}$ experiments, and apply the tools we introduced to probe their sensitivity for the aforementioned class of models, in a compared analysis to current COHERENT results, among other bounds. Our objective will be to justify $\text{CE}\nu\text{NS}$ as a powerful tool to look at some interesting New Physics scenarios, as well as reaffirming the viability of the highlighted experimental proposals.

2 CE ν NS Formalism

In the previous chapter, we have introduced in broad strokes the notion of CE ν NS, building upon a need for tools to probe scenarios related to Dark Matter and Neutrino Physics. In this chapter, we shall formalize the description of CE ν NS from its fundamentals in the neutral current interaction between quark and neutrino, reaching a formula for a measurable event rate at dedicated experiments.

2.1 Process Kinematics




Figure 6 – Schematic view of the CE ν NS process, an incident neutrino scatters off a nuclear target T, which recoils as if it was a single entity

Before getting into the specifics of the amplitude and overall cross section computations, it's useful to take a look at the kinematics of the CE ν NS process. Following what we mentioned previously, the picture here is having the neutrino scatter off a nucleus in such a way that the latter recoils elastically as if it were a single body. Following the schematics of Figure 6, we'll define the following variables:

$$\Delta E(\nu) = E - E', \quad (2.1)$$

$$\vec{q} = \vec{p}' - \vec{p} = \vec{k} - \vec{k}', \quad (2.2)$$

where primed quantities are related to the final state. We'll define for this section the notation $\|\vec{x}\| \equiv x$ for the modulus of three-vectors, and write four-vectors with their Lorentz index explicit while representing associated scalar quantities, such as $x_\mu x^\mu \equiv x_\mu^2$, for clarity.

Back to the scattering scenario, imposing conservation of energy in the lab frame (where the nucleus is initially at rest) yields:

$$E + m_T = E' + E_T \quad (2.3)$$

and furthermore, since the process is elastic, the overall kinetic energy must be conserved as well. There is no change in the target nucleus. As such, we may define $E_T = m_T + E_R$,

where the last term is the recoil energy, and write it in term of our previously defined quantities as

$$E + m_T = E' + m_T + E_R \implies E_R = \Delta E(\nu). \quad (2.4)$$

As expected, the recoil energy is equal to the variation in the energy of the neutrino. Likewise, conservation of three-momentum constrains that the nucleus recoils with momentum \vec{q} in the lab frame. Assembling the four-vector of the recoiling nucleus in the same frame, we can also use an invariant to find an useful relation:

$$\|k'_\mu\|^2 = \|(m_T + E_R, \vec{q})\|^2 = (m_T + E_R)^2 - \|\vec{q}\|^2 = m_T^2 + 2m_T E_R + E_R^2 - q^2 \quad (2.5)$$

but the squared norm of the four-momentum on the LHS is just the square of the invariant mass of the nucleus, so

$$m_T^2 = m_T^2 + 2m_T E_R + E_R^2 - q^2 \implies q^2 = 2m_T E_R + E_R^2. \quad (2.6)$$

Since we are interested in the CE ν NS regime, it's interesting to check whether the typical scales of the problem, defined via eq. (1.11), are always in the non-relativistic regime for all choices of target nucleus. In natural units, we can consider that a target T with mass number A becomes relativistic when $q \simeq m_T \approx Am_p$, where m_p is the proton mass. Imposing the coherency condition, this implies

$$1 \gg qR \simeq m_T R \approx (Am_p)(2r_0 A^{1/3}) = 2m_p r_0 A^{4/3} \quad (2.7)$$

and, replacing the values for the proton mass and the proton charge radius r_0 on the last term, yields

$$2m_p r_0 A^{4/3} = \frac{2(0.938 \times 10^9 \text{eV})(1.2 \times 10^{-15} \text{m})A^{4/3}}{1.238 \times 10^{-6} \text{eV} \cdot \text{m}} = 1.818 A^{4/3} > 1. \quad (2.8)$$

which violates the coherency condition for any value of the mass number. Thus, we reach a contradiction, meaning that the CE ν NS regime defined by the coherency condition and the relativistic limit are incompatible, and cannot be simultaneously satisfied. Therefore, in general, we can assume the non-relativistic limit safely whenever working in the CE ν NS regime. In particular, this implies we can approximate the last term in eq. (2.6) as

$$q^2 \approx 2m_T E_R. \quad (2.9)$$

This is in principle only valid for the lab frame, however we can prove its validity for any inertial frame looking at the Lorentz-invariant Mandelstam variable t :

$$t = (k'_\mu - k_\mu)^2 = q_\mu^2 = (m_T + E_R - m_T, \vec{q})^2 = E_R^2 - q^2 = -2m_T E_R \quad (2.10)$$

where for the last equality we applied eq. (2.6).

So, for the relevant CE ν NS phase space, it will be useful to keep in mind that $-t = -q_\mu^2 = q^2 = 2m_T E_R$. We can also define some relations between the incident neutrino energy and the resulting recoil. First, let's check that the exchanged momentum follows

$$q^2 = |\vec{p}' - \vec{p}|^2 = |p'^2 - 2\vec{p}' \cdot \vec{p} + p^2| = |p'^2 - 2pp' \cos \theta + p^2| \quad (2.11)$$

but, since in this scenario the neutrinos are massless, $p_\mu^2 = E_\nu^2 - p^2 = 0 \implies p = E_\nu$, and the same applies for the primed quantities. Therefore

$$q^2 = |E_\nu^2 + E_\nu'^2 - 2E_\nu E_\nu' \cos \theta|. \quad (2.12)$$

Clearly, the maximum momentum exchange occurs with $\theta = \pi$, but this implies $\vec{p} = -\vec{p}'$, so $E_\nu' = p' = p = E_\nu$. We conclude that

$$q_{\max}^2 = E_\nu^2 + E_\nu' + 2E_\nu E_\nu' = 4E_\nu^2. \quad (2.13)$$

So, for a given neutrino energy, the maximum recoil energy the scattering can produce is set by replacing the previous expression into eq. (2.6):

$$\begin{aligned} 2m_T E_R^{\max} + (E_R^{\max})^2 &= q_{\max}^2 \implies 2m_T E_R^{\max} + (E_R^{\max})^2 = 4E_\nu^2 \\ \therefore E_R^{\max} &= \sqrt{m_T^2 + 4E_\nu^2} - m_T = m_T \left(\sqrt{1 + \frac{4E_\nu^2}{m_T^2}} - 1 \right) \\ &= m_T \left(1 + \frac{2E_\nu^2}{m_T^2} + O\left(\frac{E_\nu^4}{m_T^4}\right) - 1 \right) \approx \frac{2E_\nu^2}{m_T}. \end{aligned} \quad (2.14)$$

We can also ask, given a fixed recoil energy (eg. the detector threshold), what is the minimum energy of the incident neutrino to provoke an event equal or above it. We start from the same relation:

$$4(E_\nu^{\min})^2 = 2m_T E_R + E_R^2 \implies E_\nu^{\min} = \sqrt{\frac{m_T E_R}{2} + \frac{E_R^2}{4}} \approx \sqrt{\frac{m_T E_R}{2}}, \quad (2.15)$$

where we took the non-relativistic limit for the last step.

With this, we are ready to compute the cross section for the process, and the expected rate of events in an experiment.

2.2 Standard Model Cross Section for CE ν NS

The CE ν NS process, at its core, consists of a low-momentum Z-boson exchange between the incident neutrino and a quark in the target nucleus. However, keeping in mind the discussion in the previous chapter, we should be careful about which are the degrees of freedom we are actually probing in the relevant scale. Ultimately, we are far




Figure 7 – Feynman diagram for the CE ν NS process in the Fermi-like theory.

below the required energies to meaningfully probe the nucleon substructure, given that we are within the vicinity of the coherency condition, and nuclear effects should dominate. Furthermore, as evidenced in the deduction that led to eq. (2.8), we are also close to the non-relativistic limit.

All of the aforementioned facts feed into the intuition that a natural description of the interaction could be made assuming protons and neutrons as the degrees of freedom instead of quarks. We shall proceed by introducing a Fermi-like Lagrangian to generate the process in terms of nucleon states, and later show that the result obtained matches the computation made using the usual Standard Model Lagrangian. This will lead into a brief discussion on Effective Field Theory, which will be useful in later chapters when talking about our choices of Z' models of interest.

A Fermi-like Theory CE ν NS cross section

Let's start by writing down a trial Lagrangian, and discuss some of its features:

$$-\mathcal{L}_{\text{eff}} = G_F \sqrt{2} (\bar{\nu} \gamma^\mu P_L \nu) [\bar{N} \gamma_\mu (Q_N^V + Q_N^A \gamma^5) N], \quad (2.16)$$

where $P_L = \frac{1-\gamma_5}{2}$ is the left-handed projector, and $N = n, p$ denotes neutrons and protons.

We are being quite general, admitting both vector and axial contributions, and while the pre-factors and couplings have been intentionally chosen for later convenience, there's no real loss of generality. However, one thing that stands out is that this four-fermion operator is a term of dimension 6, which is not renormalizable. This demands

that G_F has dimensions of E^{-2} , and hints that this description is only valid up until some regime, breaking down at some cutoff. We'll set these points aside for now, and proceed with the computation of the amplitude via the diagram in Figure 7, with

$$\begin{aligned}\mathcal{A} &= iG_F\sqrt{2}\bar{u}(p')\gamma^\mu P_L u(p)\langle T'|\bar{N}\gamma_\mu(Q_N^V + Q_N^A\gamma^5)N|T\rangle \\ &\equiv iG_F\sqrt{2}L^\mu H_\mu\end{aligned}\quad (2.17)$$

where spinorial indices are being omitted for simplicity of notation, but are being properly accounted for. We have essentially a product of two terms: a lepton component and a hadronic component. The latter involves applying the nucleon operators in the asymptotic nuclear states, which as discussed in section 1.2, can be highly non-trivial. Still, we'll follow a simplified prescription from [41] in order to parametrize the explicit exchanged momentum dependence within form factors, which we will not calculate directly. In summary, we'll expand the target nuclear state as a sum over single-nucleon states with well-defined angular momenta, inserting a factor $\frac{2m_T}{2m_N} = \frac{m_T}{m_N}$ in order to account for the state normalization in the non-relativistic limit. We will omit the specific momentum/spin labels since they won't be relevant, and not worry about other labels of these states, since any dependence on extra degrees of freedom will be encoded in the structure components f_N^i , to be introduced later. Emphasizing the proper normalization of states, we have:

$$\begin{aligned}\mathcal{A} &= \sum_N iG_F\sqrt{2}\bar{u}(p')\gamma^\mu P_L u(p)\frac{m_T}{m_N}\langle N'|\bar{N}\gamma_\mu(Q_N^V + Q_N^A\gamma^5)N|N\rangle \\ &= iG_F\sqrt{2}\left[\bar{u}(p')\gamma^\mu P_L u(p)\right]\sum_N\frac{m_T}{m_N}\left[Q_N^V\langle N'|\bar{N}\gamma_\mu N|N\rangle + Q_N^A\langle N'|\bar{N}\gamma_\mu\gamma^5 N|N\rangle\right]\end{aligned}\quad (2.18)$$

First, we'll direct our attention to the hadronic term. We are summing over the nucleons making up the target and, as seen in the toy example in Sec. 1.2, ideally we want to obtain an enhancement factor proportional to the number of components. Our interaction, though, is not necessarily isospin-invariant; still, we could expect that it leads to Z and $(A - Z)$ factors for protons and neutrons, respectively, times some form factor. Going forward:

$$\begin{aligned}H_\mu &= \sum_N\frac{m_T}{m_N}\left[Q_N^V\langle N'|\bar{N}\gamma_\mu N|N\rangle + \langle N'|\bar{N}\gamma_\mu Q_N^A\gamma^5 N|N\rangle\right] \\ &= \sum_N\frac{m_T}{m_N}\left[Q_N^V f_N^V(q)\bar{u}^N(k')\gamma_\mu u^N(k) + Q_N^A f_N^A(q)\bar{u}^N(k')\gamma_\mu\gamma^5 u^N(k)\right].\end{aligned}\quad (2.19)$$

Here, we can apply a non-relativistic limit for the nucleon bilinears, justified by eq. (2.8). The process is straightforward, involving choosing a representation for the Dirac spinors and taking the $\gamma \rightarrow 1$ limit, where γ is the Lorentz factor. However, it is quite long and not particularly illuminating, so it will be omitted here. Instead, we'll refer to Sec. F of [41] for further details and just present the results. We'll write explicitly the spinorial

indices here for clarity, before dropping them again when returning to the amplitude calculation.

Assuming the chiral basis, the spinor can be written as:

$$u_N(\vec{p}, s) = \frac{1}{\sqrt{2(E_p + m_N)}} \begin{bmatrix} (p_\mu \sigma^\mu + m I_2) \xi^s \\ (p_\mu \bar{\sigma}^\mu + m I_2) \xi^s \end{bmatrix} \xrightarrow{NR} \frac{1}{\sqrt{4m}} \begin{bmatrix} (2m I_2 - \vec{p} \cdot \vec{\sigma}) \xi^s \\ (2m I_2 + \vec{p} \cdot \vec{\sigma}) \xi^s \end{bmatrix} \quad (2.20)$$

and the bilinears reduce to

$$\bar{u}'_N \gamma_\mu u_N \stackrel{NR}{=} \begin{pmatrix} 2m_N \delta_{s,s'} \\ \vec{K} \delta_{s,s'} - 2i\vec{q} \times \vec{s}_N \end{pmatrix}, \quad \bar{u}'_N \gamma_\mu \gamma^5 u_N \stackrel{NR}{=} \begin{pmatrix} 2\vec{K} \cdot \vec{s}_N \\ 4m_N \vec{s}_N \end{pmatrix}. \quad (2.21)$$

where we defined $\vec{K} = \vec{k} + \vec{k}'$ and $\vec{s}_N = \xi_s^\dagger \frac{\vec{\sigma}}{2} \xi_s$, and stressing that the terms in parenthesis above are four-vectors, not spinors. Before replacing these expressions in the hadronic term, however, it's useful to consider the impact of each of these components.

- \vec{s}_N is associated with the spin of the nucleon, and since we are summing over each of the nuclear components, the overall contribution of this term will depend on the number of unpaired nucleons in the orbitals accessible by the probe. This should be exactly zero for even-even nuclei, and close to zero for most viable targets, and thus we can safely ignore this term. This essentially kills the axial contribution, so we're left with the vector term.
- In the same vein, the non-relativistic approximation made depends on the momentum of the nucleus being dominated by its mass, so we can expect the zeroth-component of the vector current to dominate over the spatial components in this regime; we'll also safely drop the latter (ie. introduce a Kronecker delta $\delta_{i,j}$ with $j = 0$).

Returning to the full hadronic term, we're left with

$$\begin{aligned} H_\mu &= \sum_N Q_N^V \frac{m_T}{m_N} 2m_N f_N^V(q) \delta_{\mu 0} \\ &= \sum_{\{p\}} Q_p^V m_T 2f_p^V(q) \delta_{\mu 0} + \sum_{\{n\}} Q_n^V m_T 2f_n^V(q) \delta_{\mu 0}, \\ &= 2m_T \left(Z Q_p^V F_p^V(q) + (A - Z) Q_n^V F_n^V(q) \right) \delta_{\mu 0} \end{aligned} \quad (2.22)$$

where we defined the proton and nuclear form factors as the average of the structure components. Replacing it back into the amplitude:

$$\begin{aligned} \mathcal{A} &= iG_F \sqrt{2} \left[\bar{u}(p') \gamma^\mu P_L u(p) \right] 2m_T \left[Z Q_p^V F_p^V(q) + (A - Z) Q_n^V F_n^V(q) \right] \delta_{0\mu} \\ &= iG_F \sqrt{2} \left[\bar{u}(p') \gamma^\mu P_L u(p) \right] 2m_T Q_T \delta_{0\mu} \end{aligned} \quad (2.23)$$

with the definition $Q_T \equiv ZQ_p^V F_p^V(q) + (A-Z)Q_n^V F_n^V(q)$ for the effective charge of the target, and we've dropped the V superscript since there is no axial contribution to introduce ambiguity. Now, squaring the amplitude and summing over the spins:

$$\begin{aligned} |\overline{\mathcal{A}}|^2 &= 2G_F^2 4m_T^2 Q_T^2 \text{Tr} \left[\not{p}'' \gamma^\mu P_L \not{p} \gamma^\nu P_L \right] \delta_{\mu 0} \delta_{\nu 0} \\ &= 8G_F^2 m_T^2 Q_T^2 p'_\alpha p_\beta \text{Tr} \left[\gamma^\alpha \gamma^\mu P_L \gamma^\beta \gamma^\nu P_L \right] \delta_{\mu 0} \delta_{\nu 0}. \end{aligned} \quad (2.24)$$

We remark that we are summing, and not averaging over the initial states spins as usual. For the massless left-handed neutrinos, we have only a well-defined negative helicity state, and not a mixed density matrix, so there is no averaging required. Likewise, our sum over the nucleon states was constructed with the assumption that they can be described with well-defined angular momenta, and in particular the population of nuclear states imposes that they should be states with well-defined spin. Again, there is not an average.

Moving on, writing the projectors explicitly, and using linearity of the trace results in

$$\begin{aligned} \text{Tr} \left[\gamma^\alpha \gamma^\mu P_L \gamma^\beta \gamma^\nu P_L \right] &= \frac{1}{4} \left(\text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \right] - \text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \gamma^5 \right] \right. \\ &\quad \left. - \text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^5 \gamma^\beta \gamma^\nu \right] + \text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^5 \gamma^\beta \gamma^\nu \gamma^5 \right] \right) \end{aligned} \quad (2.25)$$

which we can solve using the trace identities

$$\text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \right] = 4 \left[g^{\alpha\mu} g^{\beta\nu} + g^{\alpha\nu} g^{\mu\beta} - g^{\alpha\beta} g^{\mu\nu} \right] \quad (2.26)$$

and

$$\text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \gamma^5 \right] = -4i\varepsilon^{\alpha\mu\beta\nu} \quad (2.27)$$

besides using the anticommutativity of γ^5 and its product with itself giving the identity to sum the first and fourth, and second and third terms. We obtain:

$$\begin{aligned} \text{Tr} \left[\gamma^\alpha \gamma^\mu P_L \gamma^\beta \gamma^\nu P_L \right] &= \frac{1}{2} \left(\text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \right] - \text{Tr} \left[\gamma^\alpha \gamma^\mu \gamma^\beta \gamma^\nu \gamma^5 \right] \right) \\ &= 2 \left(g^{\alpha\mu} g^{\beta\nu} + g^{\alpha\nu} g^{\mu\beta} - g^{\alpha\beta} g^{\mu\nu} + i\varepsilon^{\alpha\mu\beta\nu} \right) \end{aligned} \quad (2.28)$$

and inserting back into the expression for the amplitude,

$$\begin{aligned} |\overline{\mathcal{A}}|^2 &= 8G_F^2 m_T^2 Q_T^2 p'_\alpha p_\beta \left[2 \left(g^{\alpha\mu} g^{\beta\nu} + g^{\alpha\nu} g^{\mu\beta} - g^{\alpha\beta} g^{\mu\nu} + i\varepsilon^{\alpha\mu\beta\nu} \right) \right] \delta_{\mu 0} \delta_{\nu 0} \\ &= 16G_F^2 m_T^2 Q_T^2 \left(p'^0 p^0 + p'^0 p^0 - p' \cdot p + i p'_\alpha p_\beta \varepsilon^{\alpha\theta\beta\theta} \right)^0 \\ &= 16G_F^2 m_T^2 Q_T^2 \left(2E'E + \frac{q^2}{2} \right) \end{aligned} \quad (2.29)$$

where we used in the second line that $-q^2 = t = (p' - p)^2 = -2p' \cdot p$, following Sec. 2.1, since the neutrinos are treated as massless.

Using $E_R = E - E'$ and $q^2 \approx -|\vec{q}|^2$ (again, shown in Sec. 2.1), we obtain a final expression for the averaged amplitude:

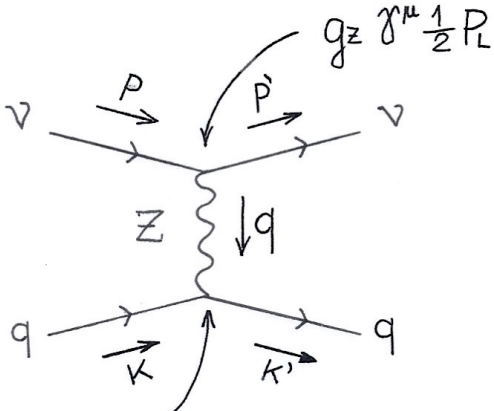
$$|\overline{\mathcal{A}}|^2 = 32G_F^2 m_T^2 Q_T^2 E^2 \left[1 - \frac{E_R}{E} - \frac{m_T E_R}{2E^2} \right]. \quad (2.30)$$

Using the two-body phase space in the lab frame, we can finally write down the differential cross section

$$\begin{aligned}
\frac{d\sigma}{dE_R} &= \frac{|\overline{\mathcal{A}}|^2}{4\sqrt{p \cdot k} 8\pi|\vec{p}|} \frac{1}{32G_F^2 m_T^2 Q_T^2 E^2 \left[1 - \frac{E_R}{E} - \frac{m_T E_R}{2E^2}\right]} \\
&= \frac{G_F^2 m_T}{\pi} \left[1 - \frac{E_R}{E} - \frac{m_T E_R}{2E^2}\right] Q_T^2 \\
&= \frac{G_F^2 m_T}{\pi} \left[1 - \frac{E_R}{E} - \frac{m_T E_R}{2E^2}\right] \left(Z Q_p F_p(q) + (A - Z) Q_n F_n(q)\right)^2. \quad (2.31)
\end{aligned}$$

While we have not defined the form factors nor the charges in this model, we can check that we have obtained something that indeed looks like the correct cross section, comparing to the literature [42]. While the former will not be computed explicitly here, we'll obtain some intuition on the latter by looking at the same computation in the Standard Model.

Standard Model CE ν NS cross section



$$g_Z \left[\left(\frac{\tau_L^3 - 2Q \sin^2 \theta_w}{2} \right) \gamma^\mu - \left(\frac{\tau_L^3}{2} \right) \gamma^\mu \gamma^5 \right]$$

Figure 8 – Feynman diagram for the CE ν NS process in the Standard Model

This time, we're adopting the usual Standard Model Lagrangian, with quark degrees of freedom; the relevant interaction term is

$$-\mathcal{L}_{\text{SM}}^{\text{NC}} \supset g_Z Z_\mu \left[\bar{f} \left(\tau_L^3 + Q \sin^2 \theta_w \right) \gamma^\mu P_L f - \bar{f} \left(-Q \sin^2 \theta_w \right) \gamma^\mu P_R f \right] \quad (2.32)$$

which we can also represent as

$$-\mathcal{L}_{\text{SM}}^{\text{NC}} \supset g_Z Z_\mu \left[\bar{f} \left(\frac{\tau_L^3 - 2Q \sin^2 \theta_w}{2} \right) \gamma^\mu f - \bar{f} \left(\frac{\tau_L^3}{2} \right) \gamma^\mu \gamma^5 f \right] \quad (2.33)$$

where τ_L^3 is the third generator of $SU(2)_L$, Q is the electric charge, θ_w is the weak mixing angle/Weinberg angle and we've written the interaction in an explicit V-A form.

Let's now write down the corresponding amplitude, via the diagram in Figure 8:

$$\begin{aligned} A &= (-i) \frac{g_Z^2}{2} \left[\bar{u}_L(p') \gamma^\mu u_L(p) \right] (-i) \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} (-i) \langle T' | \bar{q} \gamma^\nu \left(\frac{\tau_L^3 - 2s_W^2 Q_q}{2} \right) q | T \rangle \\ &\quad - (-i) \frac{g_Z^2}{2} \left[\bar{u}_L(p') \gamma^\mu u_L(p) \right] (-i) \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} (-i) \langle T' | \bar{q} \gamma^\nu \gamma^5 \left(\frac{\tau_L^3}{2} \right) q | T \rangle \\ &= i \frac{g_Z^2}{4} \left[\bar{u}_L(p') \gamma^\mu u_L(p) \right] \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} \left(\langle T' | \bar{q} \gamma^\nu \left(\tau_L^3 - 2s_W^2 Q_q \right) q | T \rangle \right. \\ &\quad \left. + \langle T' | \bar{q} \gamma^\nu \gamma^5 \left(\tau_L^3 \right) q | T \rangle \right) \\ &\equiv i \frac{g_Z^2}{4} L^\mu \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} (H_V^\nu + H_A^\nu) \end{aligned} \quad (2.34)$$

We have an axial and a vector term, each split into the neutrino current term, the Z-boson propagator and the hadronic matrix element. Let us focus on the latter first; following the same prescription as before, we'll expand the nuclear states into nucleon ones, and then to the same to obtain to get down to quarks. First, for the vector term,

$$\begin{aligned} H_V^\mu &= \langle T' | \bar{q} \gamma^\nu \left(\tau_L^3 - 2s_W^2 Q_q \right) q | T \rangle \\ &= \sum_N \frac{m_T}{m_N} f_N^V(q) \langle N' | \bar{q} \gamma^\nu \left(\tau_L^3 - 2s_W^2 Q_q \right) q | N \rangle \end{aligned} \quad (2.35)$$

Before proceeding to the quark matrix elements, we know the overall shape of the result given the possible Lorentz structure for the operator allowed by discrete symmetries (parity and charge conjugation)[41]; the current term can be decomposed as

$$\langle N' | \bar{q} \left(\tau_L^3 - 2s_W^2 Q_q \right) \gamma^\nu q | N \rangle = \sum_q Q_q^W \bar{u}(p_f) \left[\gamma^\mu F_{1,q}^{N,V}(q) + \sigma^{\mu\nu} \frac{i q_\nu}{2m_T} F_{2,q}^{N,V}(q) \right] u(p_i) \quad (2.36)$$

if the nucleon states are states with well-defined on-shell momenta $p_{i,f}$, where we defined the quark weak charge $Q_q^W \equiv \tau_L^3 - 2s_W^2 Q_q$, and F_1 and F_2 are respectively the Dirac and Pauli form factors, encoding the exchanged momentum dependency within each quark for each nucleon.

However, the second term is suppressed in the non-relativistic limit, again guaranteed by the coherency condition. Given the typical scale of the nucleon, which we can approximate by the order the proton charge radius, a regime satisfying the coherency

condition for the nucleus implies that for all purposes we are close to the zero-value of the form factors. At that limit, we are simply summing over contributions from the valence quarks of a given nucleon, so splitting eq. (2.36) we can define

$$\text{proton: } \sum_q Q_q^W \bar{u}(p_f) \gamma^\mu u(p_i) F_{1,q}^p(q) \stackrel{\text{NR}}{\equiv} (2Q_u^W + Q_d^W) \bar{u}(p_f) \gamma^\mu u(p_i) \equiv Q_p^W \bar{u}(p_f) \gamma^\mu u(p_i) \quad (2.37)$$

$$\text{neutron: } \sum_q Q_q^W \bar{u}(p_f) \gamma^\mu u(p_i) F_{1,q}^n(q) \stackrel{\text{NR}}{\equiv} (Q_u^W + 2Q_d^W) \bar{u}(p_f) \gamma^\mu u(p_i) \equiv Q_n^W \bar{u}(p_f) \gamma^\mu u(p_i), \quad (2.38)$$

and inserting these back into eq. (2.35), we obtain

$$H_V^\mu = \sum_N \frac{m_T}{m_N} Q_N^W f_N^V(q) \bar{u}(k') \gamma^\mu u(k) \quad (2.39)$$

which is of the same form as the vector term in eq. (2.19).

In a nutshell, we did exactly what we first assumed was reasonable in the previous model: reparametrize the amplitude in terms of the appropriate degrees of freedom. This time, however, we obtained a clear interpretation of the relation between the nucleon charges we had assigned and the fundamental quark couplings, which reduces to a simple sum over valence quarks in this low-momentum limit. We can't claim yet, comparing with eq. (2.19), that $Q_N = Q_N^W$ because of the pre-factors present in the overall amplitude; still, they should differ at most by a multiplicative constant, which we'll determine by the end.

The procedure for the axial hadronic term runs analogously, with the decomposition of the nucleon matrix element yielding [42]:

$$\langle N' | \bar{q} (\tau_L^3) \gamma^\nu \gamma^5 q | N \rangle = \sum_q \tau_L \bar{u}(p_f) \left[\gamma^\mu \gamma^5 G_{A,q}^N(q) - \frac{\gamma^5 q^\mu G_{P,q}^N(q) + \sigma^{\mu\nu} q_\nu \gamma^5 G_{T,q}^N(q)}{2m_T} \right] u(p_i) \quad (2.40)$$

where A, P, T stand for axial, pseudo-scalar and tensor components, respectively. Just like in the previous case, the non-relativistic limit allows us to discard the second term, resulting in

$$= \sum_q \tau_L \bar{u}(p_f) \gamma^\mu \gamma^5 G_{A,q}^N(q) u(p_i). \quad (2.41)$$

yielding something exactly analogue to eq. (2.36). We will skip the following steps, since we already know this contribution will be suppressed in the amplitude when taking the non-relativistic limit.

Finally, let's return to the amplitude in eq. (2.34), already dropping the axial term:

$$\mathcal{A} = i \frac{g_Z^2}{4} L^\mu \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} H_V^\nu \quad (2.42)$$

this time, let's focus on the propagator term. The mass of the Z boson is around 91 GeV; compared to that, we know from eq. (1.17) the typical recoil energy for the process, which is related to the exchanged four-momentum by $-q^2 = 2m_T E_R$. Therefore, the typical scale for q is of the order of $\sqrt{2A(0.938 \text{ GeV}) \times 50 \text{ keV}} \approx 10\sqrt{A} \text{ MeV}$, which is still far below the mass of the Z boson even for bulky stable nuclei with $A > 100$. At this scale, the Z cannot be produced on-shell, and its influence will be restricted to short distances. That means we can approximate $q^2 - m_Z^2 \simeq -m_Z^2$ and $\frac{q^2}{m_Z^2} \simeq 0$, which results in the following change to the propagator

$$\frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_Z^2}}{q^2 - m_Z^2} \xrightarrow{m_Z \gg \|q\|} -\frac{g_{\mu\nu}}{m_Z^2}. \quad (2.43)$$

and the amplitude now becomes (dropping again the V superscript since there's no axial term)

$$\mathcal{A} = i\frac{g_Z^2}{4} L^\mu \frac{g_{\mu\nu}}{m_Z^2} H^\nu = i\frac{g_Z^2}{4m_Z^2} L^\mu H_\mu. \quad (2.44)$$

This is exactly the amplitude from equation eq. (2.17), up to a pre-factor. In particular, the formulation and weak charge definitions would be equal if $G_F \sqrt{2} = \frac{g_Z^2}{m_Z^2}$.

While previously G_F was a general constant in the previous model's cross section, now we'll identify it with the Fermi Constant of the Standard Model, defined at tree-level as $G_F = \frac{1}{v^2 \sqrt{2}}$. We can then combine it with the definition of the Z mass:

$$G_F = \frac{1}{v^2 \sqrt{2}} \leftarrow m_Z = \frac{g_Z v}{2}$$

and thus, unsurprisingly, we obtain

$$\frac{g_Z^2}{4\sqrt{2}m_Z^2} = \frac{g_Z^2}{4\sqrt{2}\frac{g_Z^2 v^2}{4}} = \frac{1}{v^2 \sqrt{2}} = G_F \quad (2.45)$$

which confirms that indeed our whole derivation was consistent between the two models. The rest of the computation of the cross section occurs in following the same as before, since the amplitude and phase space are the same. Indeed, eq. (2.31) is the correct CE ν NS cross section [42], with

$$Q_n = Q_n^W = Q_u^W + 2Q_d^W \xrightarrow{Q_q = \tau_L^3 - 2Q \sin^2 \theta_W} = \frac{1}{2} - \frac{4 \sin^2 \theta_W}{3} - 1 + \frac{4 \sin^2 \theta_W}{3} = -\frac{1}{2} \quad (2.46)$$

$$Q_p = Q_p^W = 2Q_u^W + Q_d^W \rightarrow = 1 - \frac{8 \sin^2 \theta_W}{3} - \frac{1}{2} + \frac{2 \sin^2 \theta_W}{3} = \frac{1}{2} - 2 \sin^2 \theta_W \quad (2.47)$$

Before moving on to discuss the meaning of these matching results, we'll present some import considerations and common variations on eq. (2.31) found in the literature:

- It's common to drop the E_R/E term, since most experiments have a neutrino flux peaking at energies around 1 – 100 MeV, meaning that

$$\frac{E_R}{E} + \frac{m_T E_R}{2E^2} = \frac{E_R}{E} \left(1 + \frac{m_T}{2E}\right) \simeq \frac{E_R}{E} \frac{m_T}{2E} \quad (2.48)$$

for most of the range of the flux, which will generate the majority of events.

- There's some freedom in how we define the charges, up to a multiplicative constant, as we emphasized during the derivation of the amplitude; some texts (eg. [42]) prefer to work with $Q_N^W = 2\tau_L^2 - 4Q \sin^2 \theta_W$, which would result in a factor of 1/4 multiplying eq. (2.31), but both results are the same.
- The values for the weak charges in the Standard Model actually bring about an interesting cancellation - using the low-energy value of $\sin^2 \theta_W \simeq 0.2385$, the nucleon charges are

$$Q_n^W = -\frac{1}{2} = -0.5 \quad (2.49)$$

$$Q_p^W = \frac{1}{2} - 2 \times 0.2385 = 0.023 \quad (2.50)$$

so, in the Standard Model, the proton contribution is small compared to the neutron one, and CE ν NS cross section is approximately proportional to the neutron number $N = (A - Z)$.

- In principle, the form factors for protons and neutrons should be different, but a common approximation is to take $F_p(q) = F_n(q)$ as there is little available data for the neutron radius and distribution for most choices of nuclei, since electron scattering experiments are not effective to measure them. This can be properly computed and coded into other quantities (such as the neutron anomalous magnetic moment [42]) in a full computation of the form factor, but uncertainties should be subdominant for most of the range of experiments and thus it is usually more convenient to adopt the Helm parametrization or other approximated form.
- We still need to renormalize the theory before making predictions; radiative effects can be included in different ways depending on the renormalization scheme, so we'll follow the usual ρ -parametrization in the \overline{MS} scheme, referring to [44]. The nucleon weak charges run at low energy to the following values:

$$Q_V^p(\nu_e) = 0.0401, \quad (2.51)$$

$$Q_V^p(\nu_\mu) = 0.0318, \quad (2.52)$$

$$Q_V^n = -0.5094 \quad (2.53)$$

where the proton one acquires some flavor dependence from lepton loop corrections. None of the experiments we are working with will use tau neutrinos, so we'll just specify the values for electron and muon neutrinos.

Throughout this work, we'll work with the following form for the differential cross section, in order to compare sensitivity results with [44, 56]:

$$\frac{d\sigma}{dE_R} = \frac{G_F^2 m_T}{\pi} \left(1 - \frac{m_T E_R}{2E^2}\right) \left[Z Q_p F_p(q) + (A - Z) Q_n F_n(q) \right]^2, \quad (2.54)$$

$$Q_n = -\frac{1}{2}, Q_p = \frac{1}{2} - 2 \sin^2 \theta_W.$$

2.3 A Brief Interlude on Effective Field Theory

In the previous subsections, we have shown that the two descriptions actually yield the same cross section, once matching the couplings and charges. This is no coincidence, and points to a deeper principle of Physics, concerning separation of scales. In a general sense, most phenomena can be approximated with a local collection of contributions which does not need to include effects from far removed scales. For example, special relativity need not be invoked in scenarios with velocities much lower than the speed of light, in which Newtonian mechanics is a perfectly fine approximation.

In the case of the neutral current process, we are dealing with energies far below the mass of the Z-boson, such that it's far off-shell, and its dynamics does not play a large role. As such, its propagator can be expanded in a Maclaurin series as

$$\frac{1}{q^2 - m^2} = \frac{-1}{m^2} \frac{1}{1 - \frac{q^2}{m^2}} = \frac{-1}{m^2} \left[1 + O\left(\frac{q^4}{m^4}\right) \right] \quad (2.55)$$

as long as $m^2 > |q^2|$. Furthermore, should that difference be significant, truncating the series at the first term would already give a very precise approximation. As such, we diagrammatically represent the situation as a ‘‘collapse’’ of the propagator, represented in Figure 9, in which the boson exchange becomes manifestedly local, and at this scale the process reproduces a contact interaction (as seen by our formulation for the Fermi-like Lagrangian in eq. (2.16)).

The first thing to notice, however, is that this approximation is not valid at all scales; in particular, if one were to increase the energy of the process, it would become increasingly worse, and at the limit of its range of convergence $|q^2/m^2| < 1$, the series would diverge. This matches up with the fact that we'd be reaching the point where the Z-boson could be produced on-shell, and neglecting its dynamics would no longer make sense.

Introducing a bit of terminology, we say the Fermi-like theory is an *Effective (Field) Theory* description of the neutral current interaction. It is a low energy limit of the full theory in the Ultraviolet/UV (ie. high-energy) regime, which is the Electroweak theory, with a cutoff set at the mass of the Z-boson. The usage of ‘‘full’’ might be somewhat misleading, because the UV theory itself might also be an effective description for another




Figure 9 – Schematic view of the collapse of the propagator for an Effective interaction in terms of the Feynman diagrams

even more fundamental theory. One could even say, when properly formulated, all Physical theories are effective theories, in some sense.

In the opposite direction, in lower energy regimes, quantum chromodynamics becomes strongly coupled, and thus quarks become confined within nucleons. Redefining the dynamics in favor of baryons and mesons is also an application of Effective Field Theory, in which we are taking an Infrared/IR (ie., low energy) limit. There are many subtleties and advantages in applying properly the framework of Effective Field Theory to a problem, but we'll only focus on a few of them for our purposes. A more thorough breakdown of the topic and its formalism can be found in [57].

Following the discussion, the main feature of an EFT consists on, within its range of validity, giving a proper description of phenomena relying only on the relevant degrees of freedom at that given scale. In the effort to probe New Physics, this will serve to great effect as it will allow to parametrize our ignorance of Physics in the UV as perturbations to the Standard Model description, by means of effective higher-dimensional operators of Standard Model fields. We already did something similar in the previous example, introducing an effective four-fermion operator of dimension six to derive the CE ν NS cross section.

This will also serve us double-duty, as the mapping of the UV amplitudes to the Standard Model ones will introduce some redundancy, in the sense that multiple different theories can give rise to the same effective operator in the lower scale. Thus, constraining processes mediated by these via experiments will be a powerful and strongly model-independent tool to place bounds in a large class of UV models.

In summary, the Effective Field Theory approach will be applied to be able to neglect the effects of much heavier, non-dynamical states possibly required by our models of interest due to anomaly cancellations or other effects. Simultaneously, it will enable us to ready a relatively model-independent framework to systematically place limits on certain theories. Before moving on, there is still one definition that will be needed.

2.4 Measurable Quantity - Event Rate

Now that we have computed the CE ν NS cross section, it's necessary to connect it to a quantity to be measured in the experiments, which will be the event rate. Our standard experimental setup will consist of a neutrino beam colliding into a detector array, which will trigger upon the recoil of the nuclei inside its active volume above a certain threshold and return a signal. Running the apparatus over a given period, one can compare the number of measured events to the expected one from the Standard Model computation.

The event rate formula is very straightforward:

$$R = N_T \int_{E_R^{\min}}^{E_R^{\max}} \int_{E_\nu^{\min}}^{E_\nu^{\max}} dE_R dE_\nu \frac{d\sigma}{dE_R} \frac{d\phi}{dE_\nu}, \quad (2.56)$$

where we labeled the incoming neutrino energy $E \equiv E_\nu$ for clarity, and N_T is the number of target nuclei in the active volume, $\frac{d\phi}{dE_\nu}$ is the differential incident neutrino flux and $\frac{d\sigma}{dE_R}$ is the differential neutrino-nucleus cross section.

In particular, some experiments can measure the energy of the recoiling nucleus besides just triggering upon threshold sensitivity, or split contributions over another variable distribution; in this case, it might be convenient to introduce an event rate per bin. Furthermore, in cases with multiple targets or fluxes with multiple neutrino/antineutrino flavors, we should also expand it in a sum. Thus, we obtain the general rate-per-bin R_i as

$$R_i = \sum_T N_T \int_{E_R^{\min(i)}}^{E_R^{\max(i)}} \int_{E_\nu^{\min}}^{E_\nu^{\max}} dE_R dE_\nu \sum_{\{\nu\}} \frac{d\sigma(T, \nu)}{dE_R} \frac{d\phi(\nu)}{dE_\nu}. \quad (2.57)$$

Following Sec. 2.1, the recoil and neutrino energies are related by equations (2.15) and (2.14), and the explicit value for the limits and shape of the flux will be defined for each of the experiments of interest as required.

3 Light Z' Models

In this section, we shall define the class of models which extend the Standard Model scenario with which we will be working with. Furthermore, we shall explore how their introduction changes the results from the previous chapter, choosing a consistent description to relate the parameters of this new model with the observables measured in experiments.

3.1 Definition and Assumptions

Introducing new fields in the Standard Model must be done in a very careful manner. While it may seem straightforward to just add the necessary new terms into the Lagrangian, doing so improperly could upset the delicate gauge anomaly cancellation within its framework, leading to an inconsistent theory. We shall explore one of the simplest modifications available, consisting in adding a single new U(1) gauge symmetry, which past work has already proved that it can be done in a systematic way as so not to warp the Standard Model's anomaly cancellations[58, 59].

Therefore, with the knowledge that the theory can be made consistent, we'll extend the Standard Model by adding the following to the Lagrangian:

$$\mathcal{L} \supset \frac{-1}{4} F'^{\mu\nu} F'_{\mu\nu} + ig_{Z'} \sum_{i,j} Z'_\mu \bar{f}_i (Q_{i,j}^V + Q_{i,j}^A \gamma^5) \gamma^\mu f_j \quad (3.1)$$

where the first term is the kinetic term for the new boson Z' , $g_{Z'}$ is the coupling of this new boson to the fermions, f_i is a particular fermion and we're considering that there might be either vector or axial contributions. A correct combination of the charges $Q_{ij}^{V,A}$ guarantees that the new symmetry isn't anomalous, and from a model building perspective there is also the possibility to add new heavy fermionic fields not charged under the Standard Model group to guarantee the cancellations (though special care must still be taken with the gravity-gauge anomaly).

The previous example is something very simple and generic, which we could further specialize or expand by charging the new mediator under the Standard Model symmetries and/or introducing new charged fermions. We shall not go into these possibilities for now, and will address this possible need later, going back to the convenience of the Effective Field Theory formalism briefly discussed in Sec. 2.3.

We shall further restrict this broad category by making very few assumptions about the nature of this new mediator:

- First, it is not charged under any Standard Model symmetry;
- Second, it is massive;
- Third, it should couple to both quarks and neutrinos.

This first two conditions will define a class of U(1) extensions which can be called “Z' models”, due to its characteristics mimicking the usual Z boson in the electroweak sector as a neutral, massive mediator. If, furthermore, it also exhibits kinetic mixing with the Standard Model photon, it is instead called “Dark Photon”. The third point guarantees that effects of this extension can interfere with the Standard Model processes relevant to CE ν NS.

Closing our definition, while we do not make any assumptions on the mass of said mediator a priori, should it be lighter than the electroweak scale, we shall denominate the corresponding model as a “Light Z' Model”[60]. We will investigate the implications of that in the next section.

3.2 Neutrino Non-Standard Interactions

Closely related to the topic of Standard Model extension involving neutrinos, is the notion of Non-Standard Interactions [60]. As we've seen in chapter 2, we could apply an Effective Theory approach to parametrize the contribution of the heavy Z-boson inside the couplings of the theory. We can do the same for other new particles which might mediate a neutral current process, yielding an effective lagrangian in the similar in form as eq. (2.16). They are usually written as

$$-\mathcal{L}_{\text{NSI}} = 2\sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_{f,P} \epsilon_{f,P}^{i,j} \bar{f} \gamma_\mu P f \quad (3.2)$$

where i,j are flavor indices and f is a Standard Model fermion, with P being either P_R or P_L . Each of the terms $\epsilon_{f,P}^{i,j}$ is called *Non-Standard Interaction (NSI) Parameter*, and just like with the case of the Fermi-like theory, they are effective couplings encoding the contribution of new mediator states in the neutral current process of the Standard Model. For convenience, since we'll want to compare their contribution with the one from the Standard Model process, they are usually parametrized using the same scale G_F .

Another important feature of this framework is that it can even admit non-diagonal interactions, such that there is flavor mixing among other effects. We'll focus, however, on the case where $i = j$, such that there is no neutrino flavor mixing. These NSI parameters introduce a correction to the Standard Model amplitude, which can be included quite easily in our Fermi-like theory example. Before, it's convenient to rewrite eq. (3.2) in a

vector-axial form, since we already know the axial contribution can be neglected in the typical CE ν NS scale (see the discussion following eq. (2.21)). We obtain:

$$\begin{aligned}
-\mathcal{L}_{\text{NSI}} &= 2\sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_{f,P} \epsilon_{f,P}^{i,j} \bar{f} \gamma_\mu P f \\
&= 2\sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_f \left[\epsilon_{f,L}^{i,j} \bar{f} \gamma_\mu P_L f + \epsilon_{f,R}^{i,j} \bar{f} \gamma_\mu P_R f \right] \\
&= \sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_f \left[\epsilon_{f,L}^{i,j} \bar{f} \gamma_\mu (1 - \gamma^5) f + \epsilon_{f,R}^{i,j} \bar{f} \gamma_\mu (1 + \gamma^5) f \right] \\
&= \sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_f \left[(\epsilon_{f,L}^{i,j} + \epsilon_{f,R}^{i,j}) \bar{f} \gamma_\mu f - (\epsilon_{f,L}^{i,j} - \epsilon_{f,R}^{i,j}) \bar{f} \gamma_\mu \gamma^5 f \right] \\
&\equiv \sqrt{2}G_F \sum_{i,j} (\bar{\nu}_i \gamma^\mu P_L \nu_j) \sum_f \left[\epsilon_{f,V}^{i,j} \bar{f} \gamma_\mu f - \epsilon_{f,A}^{i,j} \bar{f} \gamma_\mu \gamma^5 f \right]
\end{aligned} \tag{3.3}$$

with the definitions

$$\epsilon_{f,V/A}^{i,j} \equiv \epsilon_{f,L}^{i,j} \pm \epsilon_{f,R}^{i,j}. \tag{3.4}$$

Now, we can compute the amplitude, following exactly the same diagram from Figure 7 and taking $i = j$ for the parameters above, but adding both the Standard Model and NSI amplitudes; we obtain, assuming no off-diagonal terms:

$$\begin{aligned}
\mathcal{A} &= \mathcal{A}_{\text{SM}} + \mathcal{A}_{\text{NSI}} \\
&= iG_F \sqrt{2} \bar{u}(p') \gamma^\mu P_L u(p) \langle T' | \bar{N} \gamma_\mu Q_N^V N | T \rangle \\
&\quad + iG_F \sqrt{2} \bar{u}(p') \gamma^\mu P_L u(p) \epsilon_{q,V} \langle T' | \bar{N} \gamma_\mu N | T \rangle.
\end{aligned} \tag{3.5}$$

The nuclear matrix element will follow the same derivation as in Sec. 2.2, and we'll assume a short-hand notation again, yielding:

$$\mathcal{A} = iG_F \sqrt{2} L^\mu \sum_N (Q_N^V + \epsilon_{N,V}) H_\mu, \tag{3.6}$$

where we defined $\epsilon_{p,V} \equiv 2\epsilon_{u,V} + \epsilon_{d,V}$ and $\epsilon_{n,V} \equiv \epsilon_{u,V} + 2\epsilon_{d,V}$ for protons and neutrons, respectively. The rest of the calculation follows completely analogously, by defining $\tilde{Q}_N^V \equiv Q_N^V + \epsilon_{N,V}$ and replacing the charges in eq. (2.54). Explicitly,

$$\begin{aligned}
\frac{d\sigma}{dE_R} &= \frac{G_F^2 m_T}{\pi} \left(1 - \frac{m_T E_R}{2E^2} \right) \left[Z Q_p F_p(q) + (A - Z) Q_n F_n(q) \right]^2, \\
Q_n &= -\frac{1}{2} + \epsilon_{u,V} + 2\epsilon_{d,V}, \\
Q_p &= \frac{1}{2} - 2\sin^2 \theta_W + 2\epsilon_{u,V} + \epsilon_{d,V}.
\end{aligned} \tag{3.7}$$

The impact of the contributions depends on the size of the NSI parameters with respect to the Standard Model charges, but we could expect in particular a sizeable effect for proton contribution, since it is ordinarily suppressed due to the value of the mixing

angle. In fact, many prior works have already applied COHERENT experiment datasets to constrain a number of these parameters to good precision [44, 60]. There is an important caveat though, in which this description, as essentially an EFT parametrization, is only valid up to the scale of which the typical energy of the process is small compared to the mass of the mediator involved in the UV description of the four-fermion operator.

This kind of scenario can correspond to the class of Light Z' mediators. If their mass is below the GeV scale, we can expect their dynamics to play a role and the NSI parametrization has to be modified. For a binned experiment, spread over a range of measured recoil energies, this can be particularly useful to break degeneracies between models and parameters - one can look for some characteristic dependence of the exchanged momentum carried by the mediator, which is directly related to the recoil energy as described in Sec. 2.1. Furthermore, as we'll see in the next sections, one can expect an enhancement of the new charge under lower exchanged momentum [60]. In order to analyze this scenario, and define our models of interest, we shall briefly introduce the concept of *Simplified Models*.

3.3 Simplified Model Approach

The aforementioned consideration about the range of validity of the parametrization could lead one to decide to instead use the “full” theory represented by the Lagrangian in eq. (3.1). However, one possible uncomfortable question to arise from that “simple” definition is: where would the mass of the new boson come from? The definition mentioned above does not contain a mass term for the Z' , and were it to follow the Standard Model scenario of Higgs Mechanism, one would need to introduce all sorts of new fields to produce either another Spontaneous Symmetry Breaking beyond the Electroweak one, or somehow justify this new boson evading detection while still gaining mass via its interaction with the Standard Model Higgs. In fact, many models containing Z' extensions are usually a byproduct of the breaking of a bigger group that includes the Standard Model's $SU(3)_C \times SU(2)_L \times U(1)_Y$ as a subgroup, perhaps introducing new scalars and fermions.

Since our objective is to study the influence of this model in CE ν NS experiments, it still makes sense to invoke the separation of scales and work in an effective field theory scenario: fields with masses above the momentum scale of the process we are working with will never be able to be produced on-shell, and thus their propagators will be suppressed at tree-level by a power of their mass. This also leads, as we discussed in Sec. 2.3, to some advantage from a phenomenological perspective, as multiple UV-complete theories can lead to the same low-energy behavior, and thus properly constraining the effective theory can help discard a larger number of theories at once. The question then is whether there is a way to reconcile this approach with the knowledge the mediator mass might be too

low to consider within the range of validity.

The solution is to incorporate another simple assumption: the UV theory itself might have some degree of separation of scales, such that while the mediator may be light, new fields such as extra scalars, right handed neutrinos and others are still very heavy. In particular, of the order of the Standard Model Z 's mass or higher. This will allow for a middle way, such that we can integrate out most of the fields, which have only a limited contribution. At the same time, we won't have to worry about origin of the mass of the mediator and other important but not relevant questions to the scenario at hand, while also keeping the dynamics of the light mediator. This is defined as a Simplified Model.

Summarizing, we still wish to apply the EFT parametrization to study the consequences of the new extension. Care must be taken, however, with the $U(1)$ mediator involved. Should its mass be of the order of the $CE\nu NS$ scale, neglecting the momentum part of its propagator could yield misleading results; since we are not making assumptions on what is its mass, we shall keep it in the low energy theory instead of integrating it out on some limit of the expansion parameter. At the same time, we will assume right-handed neutrinos, and other fermions included in the UV theory for consistency via anomaly cancellations, etc, to be sufficiently heavy to not interfere with our calculations at the $CE\nu NS$ scale.

3.4 Models of Interest and Revised Cross Section

Having made all the relevant definitions, we can now write an effective Lagrangian for the new interaction, containing the full propagator of the Z' , and match it to the amplitude computed before in this chapter to check how this will modify the NSI parametrization implied and match the charge definitions in the low energy model. This interaction term in this new Lagrangian reads:

$$\mathcal{L}_{\text{int } Z'} = g_{Z'} Z'_\mu \sum_f Q'_f \bar{f} \gamma^\mu f \quad (3.8)$$

where f are SM fermions, and we are assuming no tree-level mixing with the SM bosons. Again, given the Standard Model computation, this time we are not including axial contributions, since fundamentally they won't benefit from the same enhancement as a vector term.

The underlying process follows exactly the same diagram pictured previously in Figure 8, replacing Z by Z' and the relevant coupling/charges. The resulting amplitude then reads

$$\mathcal{A}_{Z'} = -i g_{Z'}^2 Q'_\nu \bar{u}(p') \gamma^\mu P_L u(p) \frac{g_{\mu\nu} - \frac{q_\mu q_\nu}{m_{Z'}^2}}{q^2 - m_{Z'}^2} \langle T' | Q'_q \bar{q} \gamma^\nu q | T \rangle. \quad (3.9)$$

Since $p'_\mu, p_\mu, q_\mu = p'_\mu - p_\mu$ are on-shell momenta, the Ward Identity guarantees that $q_\mu L^\mu = 0$, and thus we can safely drop the second term in the propagator without further consideration.

Using the usual shorthand notation from the previous examples, we may rewrite the amplitude as

$$\mathcal{A}_{Z'} = -i \sum_{N,q} g_{Z'}^2 Q'_\nu L^\mu \frac{g_{\mu\nu}}{q^2 - m_{Z'}^2} Q'_q H^\nu \quad (3.10)$$

and, scaling the charges by the Fermi constant:

$$\mathcal{A}_{Z'} = iG_F \sqrt{2} L^\mu \sum_{N,q} \left(\frac{-g_{Z'}^2 Q'_\nu Q'_q}{G_F \sqrt{2} (q^2 - m_{Z'}^2)} \right) H_\mu. \quad (3.11)$$

Comparing with eq. (3.6), this is exactly the same amplitude with the identification

$$\epsilon_{V,q} = \frac{-g_{Z'}^2 Q'_\nu Q'_q}{G_F \sqrt{2} (q^2 - m_{Z'}^2)} \simeq \frac{g_{Z'}^2 Q'_\nu Q'_q}{G_F \sqrt{2} (|\vec{q}|^2 + m_{Z'}^2)}, \quad (3.12)$$

where we used that $q^2 \simeq -|\vec{q}|^2$ in the last equality.

After fixing the charges to define a particular model, this expression has 2 free parameters: the new coupling $g_{Z'}$ and the mass of the new mediator $m_{Z'}$. Our objective will be to try to constrain the allowed values for these using experiments.

Based on this parametrization, we'll define two models of interest:

- **Universal Z' :** The new gauge boson couples to all fermions with the same strength, ie. $Q'_q = Q'_\nu = 1$ for all quarks and neutrino species. The Standard Model charges are modified as

$$Q_N^V \rightarrow \left(Q_N^V + \frac{3g_{Z'}^2}{\sqrt{2}G_F(|\vec{q}|^2 + m_{Z'}^2)} \right) \quad (3.13)$$

- **Z'_{B-L} :** The new gauge boson couples to the (anomaly-free) current of the gauged $U(1)_{B-L}$ group, where B and L are respectively the Baryon and Lepton numbers. As such, it corresponds to $Q'_q = 1/3$ and $Q'_\nu = -1$. The Standard Model charges are modified as

$$Q_N^V \rightarrow \left(Q_N^V - \frac{g_{Z'}^2}{\sqrt{2}G_F(|\vec{q}|^2 + m_{Z'}^2)} \right) \quad (3.14)$$

However, we shall also include an extra model of interest with a slightly different structure, in which it does not directly obeys the third condition we defined previously for Z' models, namely having non-zero coupling to quarks:

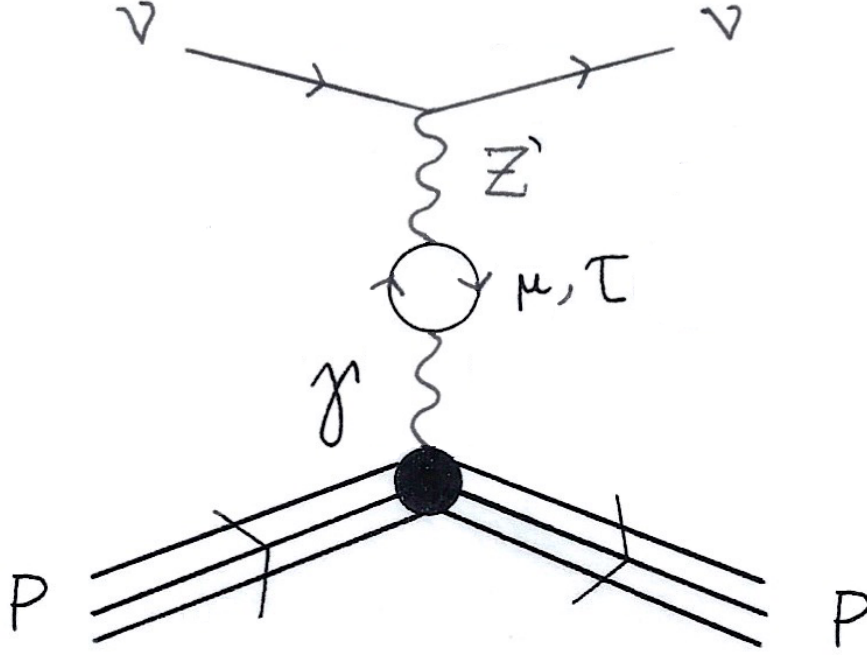


Figure 10 – Diagram for the loop-induced mixing between Z' and photon, giving rise to the proton neutral current charge modification in the $Z'_{\mu-\tau}$ model.

- $Z'_{\mu-\tau}$: In this case, the new gauge boson couples to the gauged $U(1)_{L\mu-L\tau}$ current, with $Q'_q = Q'_{\nu_e} = 0$ and $Q'_{\nu_\mu} = 1 = -Q'_{\nu_\tau}$. However, in this case, while one does not get a tree-level interaction with the quark sector, there is still an effective one. Instead of a tree-level exchange, here the process occurs via kinetic mixing between the photon and the Z' via a muon and tau loop, illustrated by the diagram in Figure 10. The quark-level couplings are modified as [61]:

$$\epsilon_{q,V} = + \frac{2 \alpha_{EM} g_{Z'}^2 Q_q}{\sqrt{2} \pi G_F (|\vec{q}|^2 + m_{Z'}^2)} \int_0^1 dx x(1-x) \log \left[\frac{m_\tau^2 + |\vec{q}|^2 x(1-x)}{m_\mu^2 + |\vec{q}|^2 x(1-x)} \right], \quad (3.15)$$

which can be approximated in the limit of $m_\tau, m_\mu \gg |\vec{q}|$ to

$$\epsilon_{q,V} = + \frac{\alpha_{EM} g_{Z'}^2 Q_q}{3\sqrt{2} \pi G_F (|\vec{q}|^2 + m_{Z'}^2)} \log \left(\frac{m_\tau^2}{m_\mu^2} \right), \quad (3.16)$$

where Q_q is the quarks' electric charge and α_{EM} is the usual electromagnetic fine-structure constant. While the tau mass is generally much above the typical CE ν NS scale, the same is not necessarily true for the muon mass. Still, as shown in [61], the full computation yields a negligible deviation for the relevant phase space. We shall include it fully when dealing with the ν BDX-DRIFT experiment in Part II to be careful, due to the higher energy of the incoming flux going beyond the region

considered in the reference, although the effect is also expected to be overall negligible for the sensitivity analysis that will follow. In conclusion, the Standard Model charges are modified as

$$Q_p^V \rightarrow \left[Q_p^V + \frac{\alpha_{EM} g_{Z'}}{3\sqrt{2}\pi G_F (|\vec{q}|^2 + m_{Z'}^2)} \log\left(\frac{m_\tau^2}{m_\mu^2}\right) \right], \quad (3.17)$$

$$Q_n^V \rightarrow Q_n^V. \quad (3.18)$$

Since the coupling occurs via a mixing with the photon, only the proton receives a contribution, since the neutron has no net electric charge to couple to at this scale.

Following what was mentioned in the end of Sec. 3.1, all of those modifications include a term which grows as $|\vec{q}|^2 \rightarrow 0$, thus naively enhancing the discrepancy to the Standard Model at lower exchanged momentum, if $m_{Z'}$ is of the order or below the typical scale of the process. This justifies CE ν NS as a very powerful tool to work with this sector of the landscape of light Z' models.

With the models of interest and their relevant parameters defined, the last piece of information we need in order to start analyzing the viability of these models and their sensitivity in future experiments is to define the testing framework itself. There must be a formalism in place to relate the underlying parameters, to the experimental datasets and infer the reliability of the results.

4 Hypothesis Testing and Wilks' Theorem

The formalism derived in the previous chapters can be readily applied to experiments to obtain predictions for the number of expected $CE\nu NS$ events. Although, as mentioned, most of the framework was put in place in the 70's, the first detection of this phenomena has only recently been made, by the COHERENT Collaboration [47]. There are currently, however, many different proposals for other $CE\nu NS$ -focused experiments, foreboding the doors opening for a $CE\nu NS$ "precision era" now that the discovery has been successfully achieved.

In this chapter, we shall proceed by first defining the general framework for hypothesis tests, which will be crucial later in order to probe the sensitivity of these experimental proposals to the parameters of our Light Z' Models of interest, defined in Sec 3.4.

4.1 Definitions and Formalism - Brief overview

Statistical Inference

Due to the intrinsic probabilistic nature of scattering events, one needs to define a proper way to connect an observed dataset to the underlying parameters of the model being tested, for which one needs the tools of Statistical Inference. That need can be promptly understood with very basic examples. The guiding principle is the same as for a simple coin-toss experiment: while unlikely, there is a non-zero chance of a fair coin to land on heads consecutively for any finite number of times. Just as well, were a coin-toss dataset to exhibit this behavior, it would still not be possible to discard it being fair with absolute certainty.

The role of statistical inference will be to define a systematic approach to use the available data to ascertain the relative likelihood of distinct hypotheses about the underlying distribution - after "n" consecutive heads results, we'd like to have a quantifiable way to express that the coin being fixed seems more likely than it being fair, and exactly how much more likely it is. While for a generic case it may be impossible to reach absolute certainty between two hypothesis (as one would have, after 10 tosses as heads one cannot guarantee the 11th wouldn't be tails without knowing the actual state of the coin), it is common to elect a threshold with which to claim exactly whether a hypothesis can be confirmed or excluded. For a fair coin, the probability of getting 100 consecutive heads is $\frac{1}{2^{100}} < 10^{-30}$, less than one in one *nonillion*, so intuitively one would never expect to reproduce such an unlikely event were this underlying hypothesis true, and for all effects could discard it without feeling anxious about making a mistake. In practice, thresholds

for Particle Physics discoveries are set at 5σ , whose precise meaning will be described later in this chapter.

Another feature illustrated in a coin-toss example is the dependence on the size of the dataset. As more tests are made, the cumulative outcomes that deviate from the mean become rarer, and in particular, there are fewer of them the farther they are from the true mean, in some sense. While there are subtleties than need to be addressed that go beyond this simple example, it's possible to accurately define this apparent empirical phenomenon imposing some simple assumptions on the underlying distribution.

For an experiment with a given set of outcomes with finite mean and finite non-zero variance, the probability of choosing a cumulative outcome of n samples such that its sample mean is within distance ϵ of the true mean tends to 1 as the n goes to infinity. This is known as the (Weak) **Law of Large Numbers**. Formally, if X denotes a random variable with n independently identically distributed samples X_1, X_2, \dots, X_n , with mean $E(X) = \mu$ and variance σ^2 , then

$$\lim_{n \rightarrow \infty} \Pr(|\bar{X}_n - \mu| < \epsilon) = 1, \forall \epsilon > 0 \quad (4.1)$$

where $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ denotes the sample average. In other words, the sample average *converges in probability* towards the expected value.

We can also invoke **Chebyshev's Inequality** to further bound how well the approximation may hold given that our sample has an underlying distribution with finite mean and finite non-zero variance; specifically, no more than $\frac{1}{k^2}$ samples can assume values k standard deviations or more away from the mean. Or,

$$\Pr(|X - \mu| \geq k\sigma) \leq \frac{1}{k^2}, \quad (4.2)$$

with X being a random variable, $\mu = E(X)$ the expected value, σ its standard deviation and k any real positive number.

We shall not demonstrate these theorems, as it is outside of the scope of this thesis, but our main takeaway involves the utility of what they can say about the goal of quantifying hypothesis tests. Together, they inform us that by taking the sample average from an experimental set of obtained results we should approach the true expected value. Furthermore, since the sample average itself is a random variable, we can systematically compare the probability of obtaining the result among different hypothesis about the mean value and variance.

That is to say, we can properly formalize the notion of *likelihood*: using a particular set ϵ of parameters yields an expected value $\mu(\epsilon)$ and variance $\sigma(\epsilon)^2$ to be compared to the observed sample mean \bar{X}_n , which defines a $k(\epsilon)$, linked to the probability of observing an outcome at least as extreme (ie. far from the mean μ). This set of parameters (or

equivalently, the hypothesis that states “ ε is the value of the parameters of the model M ”) will then be more likely than another set of parameters ε' if their associated k -value is smaller, ie. there's a larger probability of observing that set of results given the expected value computed from the theory for that set of parameters. For an independently and identically distributed random variable X , with a sample $\{x_n\}$ of outcomes, a likelihood function, therefore, will be a function

$$\mathbf{L}(\varepsilon) = f(\{x_n\}|\varepsilon) \quad (4.3)$$

describing the probability, within a given set of data, for it to have been generated by an underlying distribution with parameters ε . It's important to emphasize that this is a function of the parameters, not of the data set, which is not a variable.

While particularly illuminating to the reasoning and usefulness of inference tools, in practice we won't be using the bounds given by equation (4.2), in favor of stronger ones, and the procedure outlined will also be subtly different due to how we'll approach error control. We'll go back to our coin-toss example for a moment to illustrate this.

Definitions - Hypothesis and Errors

Let us perform a certain repetition of tosses in order to infer how likely it is that a coin is fair. We shall call “the coin is fair”, or equivalently, “The probability p of getting heads is 0.5” as the *Null-hypothesis*, often denoted as H_0 . The assertion $p = 0.1$ will be our *Alternative hypothesis* H_1 . The objective of the test will be to determine whether or not the outcome set rejects the Null-hypothesis within some predetermined threshold of likelihood in favor of our Alternative hypothesis. As mentioned before, for a general scenario one can never make the claim that “ H_0 is false” with absolute certainty, and analogously, that “ H_1 is true”, so even setting stringent thresholds there is always the possibility of errors. We shall classify them into two types:

- **Type I Error (false-positive):** This type of error occurs when a true Null-hypothesis ends up rejected. In the coin-toss example, this could happen by getting sufficiently more tails than heads within the sample, leading to the incorrect conclusion that the coin wasn't fair. It is closely associated with the significance level α , which is the threshold probability for observing a result at least as extreme for the sampling given that the Null-hypothesis is true.
- **Type II Error (false-negative):** This type of error occurs when a false Null-hypothesis cannot be rejected. For instance, if the coin-toss sample resulted in a proportion of heads and tails between the one from the null and alternate hypothesis, in such a way that the tolerance threshold wasn't enough to favor the Alternative hypothesis. The susceptibility to this type of error is said to determine the *Power of a test*, and it

can be reduced by larger sample sizes or less stringent thresholds for Null-hypothesis rejection.

In general, for the purposes of Particle Physics experiments, we'd like to minimize in particular the Type I errors, so as to avoid getting false experimental leads to direct theoretical developments in the wrong direction, leading to very stringent bounds for the significance α . Our focus, however, will be subtly different. We will study experimental proposals, to measure how much of the parameter space available for our models of interest can be excluded at a given significance if the Standard Model is the correct underlying theory. Therefore, for each point in the space of parameters, we shall conduct a hypothesis test with "Model M with input parameters ε " as the Null-hypothesis, against the Standard Model as the Alternative hypothesis. For that purpose, we will define a *test statistic*, whose distribution will allow us to properly define the significance of our test.

Test Statistic

Following [62], we shall define the test statistic following a Poisson likelihood. That is appropriate given that, following section 2.4, what we will ultimately be measuring in the experiments is the number of events observed within a limited interval of time. We'll start by defining the **Poisson Distribution**: A random variable X is said to be Poisson-distributed, with parameter $\lambda > 0$, if for each observed number of events k it has a probability mass function $f_{\text{Poiss}} = \Pr(X = k)$ respecting

$$f_{\text{Poiss}}(k, \lambda) = \Pr(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}. \quad (4.4)$$

Its expected value and variance are both equal to λ , which is the expected number of events in a particular period. One can also substitute $\lambda = rt$ to instead ask about the rate r in which an event happens within a time interval t . This type of distribution is valid for a model based on the assumption that events are occurring independently but at an average rate. This fits with our scenario of measuring the number of $\text{CE}\nu\text{NS}$ events in a detector over a specified time frame.

As stated previously, we can define a likelihood function as the probability of observing a determined set of outcomes for an identically and independently distributed random variable, if for the underlying distribution we assume a set ε of parameters; for a Poisson distribution, we then have:

$$\mathbf{L}(\varepsilon|x_1, x_2, \dots, x_n) = \Pr(X = x_1|\varepsilon) \times \Pr(X = x_2|\varepsilon) \times \dots \times \Pr(X = x_n|\varepsilon) \quad (4.5)$$

but each of the factors in the right hand side are just the Poisson probability mass function, so

$$\mathbf{L}(\varepsilon|x_1, x_2, \dots, x_n) = \prod_{i=1}^n f_{\text{Poiss}}(x_i, \lambda) = \prod_{i=1}^n \frac{\lambda^{x_i} e^{-\lambda}}{x_i!} = e^{-n\lambda} \frac{\lambda^{\sum_{i=1}^n x_i}}{\prod_{i=1}^n x_i!}. \quad (4.6)$$

For reasons we will show in the next sections, we will prefer to work with the **Log-Likelihood** instead:

$$-2 \log \mathbf{L}(\varepsilon | x_1, x_2, \dots, x_n) = 2 \left[n\lambda - \left(\sum_{i=1}^n x_i \right) \log \lambda + \log \left(\prod_{i=1}^n x_i \right) \right]. \quad (4.7)$$

So, now we have given one step ahead: we have found a way to quantify how likely it is to have some underlying parameter choice given a set of outcomes from our experiment: it is expressed by how big is the value of the respective likelihood function (or reciprocally how small, if working with $-2 \log(\mathbf{L})$). Furthermore, we can directly compare two hypothesis using it. Let \mathcal{P} be the phase space of parameters for a particular model, with $\mathcal{P}_0 \subset \mathcal{P}$ being the subset of this space containing the values satisfying a Null-hypothesis H_0 and $\mathcal{P}_1 \subset \mathcal{P}/\mathcal{P}_0$ the one containing values satisfying the Alternative hypothesis. Then, we can define the **Likelihood-ratio** Λ between them as

$$\Lambda = \frac{\max_{\varepsilon_0 \in \mathcal{P}_0} \mathbf{L}(\varepsilon_0)}{\max_{\varepsilon_1 \in \mathcal{P}_1} \mathbf{L}(\varepsilon_1)}, \quad (4.8)$$

with the test rejecting the Null-hypothesis at significance α if $\Lambda > c(\alpha)$, though we still have to introduce how to define this number c . Notice that we can also take the logarithm of the likelihood ratio:

$$\begin{aligned} -2 \log \Lambda &= -2 \log \left(\frac{\max_{\varepsilon_0 \in \mathcal{P}_0} \mathbf{L}(\varepsilon_0)}{\max_{\varepsilon_1 \in \mathcal{P}_1} \mathbf{L}(\varepsilon_1)} \right) \\ &= -2 \log \left(\max_{\varepsilon_0 \in \mathcal{P}_0} \mathbf{L}(\varepsilon_0) \right) + 2 \log \left(\max_{\varepsilon_1 \in \mathcal{P}_1} \mathbf{L}(\varepsilon_1) \right) \\ &= \min_{\varepsilon_0 \in \mathcal{P}_0} (-2 \log \mathbf{L}(\varepsilon_0)) - \min_{\varepsilon_1 \in \mathcal{P}_1} (-2 \log \mathbf{L}(\varepsilon_1)) \end{aligned} \quad (4.9)$$

which, for a Poisson-distributed experiment, will reduce to

$$-2 \log \Lambda = 2 \left[n\lambda_0 - \left(\sum_{i=1}^n x_i \right) \log \left(\frac{\lambda_0}{\lambda_1} \right) - n\lambda_1 \right] = 2n \left[\lambda_0 - \bar{x} \log \left(\frac{\lambda_0}{\lambda_1} \right) - \lambda_1 \right], \quad (4.10)$$

where \bar{x} is the sample average.

Wilks' Theorem

Finally, all that remains is finding a way to decide which threshold values are appropriate for rejecting a hypothesis for each given significance requirement (which, as defined above, is the probability of committing a type I error, ie. rejecting the Null-hypothesis even when it is true).

While there are multiple approaches to this problem, roughly grouped into “Bayesian” and “Frequentist” interpretations, we shall focus on the later given its prevalence in most of the literature [63]. Discussions about the merits of each can go into deeper philosophical

considerations about the inherent nature of probability, and thus are outside the scope of this work.

The test statistic was defined to formalize the notion of how extreme (or unlikely) one given set of outcomes can be for a set of parameters assumed by a hypothesis. Unsurprisingly, one of its factors depends on the sample average observed. However, were the experiment to be repeated a number of times, each could have sampled different values, even if they all follow the same underlying distribution. Thus, the test statistic itself should follow a related distribution. Indeed, in a Frequentist interpretation, we can define *Confidence Levels* for a given hypothesis test as intervals in the phase space of a parameter of interest in which the true value of that parameter under the Null-hypothesis is covered $(1 - \alpha)$ of the time when reproducing experiments. It is the reciprocal of the significance. If we can determine the distribution of the test statistic, setting a confidence level (or equivalently, significance) means simply computing the cumulative distribution function until a tabulated threshold.

On first view, it seems we have looped back into the same problem, since we do not know a priori the distribution of the test statistic. One choice of approach is to populate the outcome space by simulating pseudo-experiments, using Monte Carlo methods [63]; thus, choosing a parameter of interest and marginalizing over the rest yields a distribution of the test statistic over the values of the former, and the cumulative distribution can be sampled directly to determine the appropriate confidence interval. This is at the core of frequentist Feldman-Cousins methods. However, such a procedure can be extremely computationally intensive, specially when there are multiple parameters to consider.

Fortunately, under certain conditions, one does not need to rely on Monte Carlo due to the test statistic actually following asymptotically a known distribution, independent of the actual data. **Wilks' Theorem** states that, for a random variable distributed as some function $f(x|\varepsilon)$, which admits a maximum likelihood estimator $\bar{\varepsilon}$ under some regularity conditions (formally expressed in [64]), the test statistic $-2 \log \Lambda$ follows asymptotically a χ_k^2 distribution under the Null-hypothesis, where $k = \dim(\epsilon_1) - \dim(\epsilon_0)$ is the difference in the number of free parameters between the Alternate hypothesis and the Null-hypothesis. Specifically, it follows a χ_k^2 distribution up to factors of order $O(\frac{1}{\sqrt{n}})$, where n is the number of events. The proof of this theorem is outside of our scope, and can be found in [64].

Furthermore, following [63], there's a convenient way to express the Confidence Level Z of the test, by using units of standard deviation for the cumulative of a normal distribution; as a consequence of Wilks' Theorem it can be directly related to the test statistic as

$$Z(\varepsilon) = \sqrt{-2 \log \Lambda}. \quad (4.11)$$

The popular “ 5σ Limit” for a discovery in Particle Physics in practice means a confidence level of 99.99994%, or correspondingly, an α of 0.00006.

Experimental and Systematic Uncertainties

While we have successfully described a framework for hypothesis testing, so far we have not included anything concerning sources of errors in the measures themselves. Back to the example of the coin, we're now dealing with a case where there exists some inherent uncertainty about the state we're measuring. Maybe we're observing the coin being thrown in the far distance, such that heads and tails can only be distinguished with some non-perfect degree of certainty. This should also impact our description of the likelihood for the outcomes dataset in some way.

This will be included in our framework with the *nuisance parameters*, which are a set of parameters which are not related to the quantity of interest in the model we are testing, but that can impact the result. The common way to deal with those is to profile over them, choosing a combination to extremize the likelihood function with a given weight. Specifically, we shall make use of the **Pull Method**:

Let t be some χ^2 -distributed test statistic, in the absence of experimental/ systematic uncertainties, and $\vec{\xi}$ a set of nuisance parameters for the underlying distribution. Assuming that the errors related to these can be approximated by a normal distribution, with variance σ^2 for each corresponding nuisance parameter, we can define a new test statistic \tilde{t} as

$$\tilde{t} = t(\vec{\xi}) + \sum_i \left(\frac{\xi_i - 1}{\sigma_i} \right)^2. \quad (4.12)$$

Essentially, we're setting a scale for deviations around the mean by the variance related to the nuisance parameters, and then asking the same question that motivated our endeavor from the beginning: “how likely is it to observe this data distribution for some choice of parameters?”, and then combining it with the extremality of the prediction versus experimental outcomes.

Illustratively, we can think about a test of event rates in the Standard Model versus an experimental outcome. Assume there was some 10% overshooting of the theoretical prediction with respect to the observed number of events. However, the experiment has some uncertainties with measuring the number of events, because of flux normalizations, background considerations, etc. If the uncertainty is small, in this case which means reasonably below 10%, this kind of overshooting would be very unlikely to happen just due to detector effects, and thus the event dataset would be significant. On the other hand, if the uncertainty is large, this kind of deviation would be common, and the significance is reduced.

In the cases we are going to analyze, we'll be working mainly with two nuisance parameters: the normalization of the number of signal (ie. $\text{CE}\nu\text{NS}$) events (ξ_S) and the normalization of background events (ξ_B). We'll properly define them when proceeding with the sensitivity analysis in Chapter 8.

4.2 Testing Algorithm

With the formalism for the hypothesis test defined in the last section, we shall now outline exactly how this will be performed in practice for our context of future $\text{CE}\nu\text{NS}$ experiments. We shall proceed by steps, in a systematic way.

Our objective is to derive a profile separating the allowed and excluded region over the space of our model's parameters, which in this case are $(g_{Z'}, m_{Z'})$, at a fixed confidence level. In order to facilitate comparison with other bounds found in the literature, we will specify a 90% Confidence Level line as our target sensitivity, corresponding to a significance of 1.645σ . Given an experimental proposal E, and a model of interest Z'_i , we shall test the sensitivity of E with respect to Z'_i by following these steps:

1. Set a mass $m_{Z'} = M$ for the new mediator;
2. Compute the $\text{CE}\nu\text{NS}$ cross section for the Standard Model, and for the model Z'_i ;
3. Using the specifications of the experiment E, compute the event rate for the Standard Model and for the model Z'_i , following eq. (2.56), in each bin defined for E. We shall call the former $N_{SM,i}$ and the later $N_i(g_{Z'}, m_{Z'} = M)$;
4. Introduce systematic uncertainties via nuisance parameters ξ_S , for the overall normalization of N_i , and ξ_B , for the number of background events $N_{B,i}$, computed for E.
5. Since there is no data set for a planned experiment, we shall assume that it will follow the distribution implied by the Standard Model, in such a way that the sample average in each bin will be exactly $N_{SM,i} + N_{B,i}$;
6. For the fixed value of $m_{Z'} = M$, compute the test statistic following eqs. (4.9) and (4.10), assuming H_0 fixes $\varepsilon_0 = (g_{Z'} = g, m_{Z'} = M, \xi_S, \xi_B) \in \mathcal{P}_0$ and H_1 fixes $\varepsilon_1 = (g_{Z'} \neq g, m_{Z'} = M, \xi_S, \xi_B) \in \mathcal{P}_1$ such that we have parameters :

$$\begin{aligned} \chi^2(g_{Z'} = g; m_{Z'} = M) &= \sum_i \min_{\varepsilon_0 \in \mathcal{P}_0} (-2 \log \mathbf{L}_i(\varepsilon_0)) - \min_{\varepsilon_1 \in \mathcal{P}_1} (-2 \log \mathbf{L}_i(\varepsilon_1)) \\ &= \sum_i \min_{(\xi_S, \xi_B)} (-2 \log \mathbf{L}_i(\varepsilon_0)) - \min_{(g_{Z'}, \xi_S, \xi_B)} (-2 \log \mathbf{L}_i(\varepsilon_1)) \quad (4.13) \end{aligned}$$

where L_i is the Likelihood function for the i-th bin. But, following from the previous point, if the expected dataset just follows the Standard Model prediction, it's obvious

to expect that the minimum of the likelihood is reached for $g_{Z'} = 0, \xi_S = \xi_b = 1$, which is just the Standard Model result. Replacing accordingly:

$$\begin{aligned} \chi^2(g_{Z'} = g; m_{Z'} = M) = & 2 \min_{(\xi_S, \xi_B)} \sum_i \left[\xi_S N_i(g_{Z'} = g, m_{Z'} = M) + \xi_B N_{B,i} \right. \\ & - (N_{SM,i} + N_{B,i}) \log \left(\frac{\xi_S N_i(g_{Z'} = g, m_{Z'} = M) + \xi_B N_{B,i}}{N_{SM,i} + N_{B,i}} \right) \\ & \left. - (N_{SM,i} + N_{B,i}) \right] + \left(\frac{1 - \xi_S}{\sigma_S} \right)^2 + \left(\frac{1 - \xi_B}{\sigma_B} \right)^2. \end{aligned} \quad (4.14)$$

7. Scan over allowed values for $g_{Z'}$, minimizing $|\chi^2(g, M) - (1.645)^2|$ over $g - Z'$ to find the 90% C.L. contours, and reject the interval where $|\chi^2(g, M) - (1.645)^2| > 0$;
8. Scan over allowed values for $m_{Z'}$, returning to step 1 with a different choice for $m_{Z'}$.

The number of operations required is huge, and needs to be performed with the help of a computer, but are much less intensive than associated Monte Carlo methods. For this work, all steps have been implemented using Python and Wolfram Mathematica for verification.

Part II

Phenomenological Applications - CE ν NS Experiments

5 Benchmark Case - CENNS-10 Detector

Throughout Part I, we have carefully derived the formalism to compute predictions for $\text{CE}\nu\text{NS}$ experiments within the Standard Model and certain Z' extensions of interest, as well as introduced the framework with which to compare them to experimental data. Before applying it blindly to future prospects, it is convenient to test how well it may work for a well-studied benchmark case.

The COHERENT collaboration, at the time of this work, has announced three results of $\text{CE}\nu\text{NS}$ measurements [47, 48, 65], consisting in the first observation of the phenomena (2017) and a recent update on a Sodium-doped Cesium-Iodine scintillator (2021), and another recent observation in a liquid Argon, single phase scintillator detector (2020). We shall use the later as our example, due to the availability to good precision of Argon-40 nuclear properties, the simplicity of dealing with a single type of target nucleus, as well as having the advantage of an exact cancellation for the axial contribution for an even-even nucleus. Its correspondent data release is available at [66].

Our goals for this chapter are the following: extract the necessary quantities for the computation of the event rate via eq. (2.56) from the experiment's definitions and successfully compute the Standard Model event rate for the experiment following [44] and [48]. We shall also compute the expected number of events for our models of interest, described in 3.4. The corresponding sensitivity limits for this case will be presented later in Chapter 8, for the compared analysis with future proposals.

5.1 Detector properties

The COHERENT COH-AR-10 detector, formerly known as CENNS-10, is a 10 L Liquid Argon scintillator-based detector localized in the Neutrino Alley of the SNS facility, at the Oak Ridge National Laboratory, Fermilab. We shall refer to the detector exclusively as CENNS-10 henceforth, to match the presented references. In operation since 2015, it has reported the second observation of $\text{CE}\nu\text{NS}$ in 2020 [48]. Despite still suffering from large relative uncertainties from flux and quenching determination, it has achieved over 3σ significance in confirming $\text{CE}\nu\text{NS}$.

Brief overview of scintillation in Liquid Argon

A schematic view of the detector is presented in Figure 11. Incoming neutrinos produce a recoil event inside the fiducial volume, generating secondary scatters which will deposit energy in the liquid and trigger the formation of excited Argon dimers (ex-




Figure 11 – Schematic view of the CENNS-10 detector, including the fiducial LAr volume, overall detector vessel, vacuum chamber and shielding. Source: [48]

cimers)[67]. These consist of essentially a Rydberg-like state of a positive Ar_2^+ core with a bound electron; depending on the combined spin of the system, the ground state of the combination might lie in a singlet or triplet state, with mean lifetimes of 6ns and $1.5\mu s$, respectively. These excimers eventually decay back into two Argon atoms, emitting a photon in a characteristic 128nm wavelength (Vacuum Ultraviolet) as a result. The process occurs via two pathways, illustrated in Figure 12, with different yields for singlet and triplet states. The first involves direct excitation of an Argon atom to form an excimer, while the second relies on the recombination of a thermalized electron with an ionized dimer.

The liquid Argon medium is very transparent to its own scintillation light, due to the disparity between the typical separation scale in the liquid phase for free atoms and for the dimers[67]. It also has a very consistent yield of around 40.000 photons/MeV, which can propagate within the fiducial volume and reach the Photomultiplier Tubes (PMTs). There, the light is converted to an electric signal waveform, that is then digitalized. To maximize yield, the 128nm scintillation light is wavelength-shifted to a distribution peaked at 420nm [48]. This setup provides a nominal threshold nuclear recoil energy at around 20 keV.

The biggest advantage of this type of setup in the context of CE ν NS is due to its ability for Pulse-Shape Discrimination (PSD). One of the primary concerns for a coherent scattering proposal is to be able to distinguish Nuclear Recoil (NR) from Electron Recoil




Figure 12 – Representation of the two scintillation pathways of Argon; direct dimer excitation (above) and ionized dimer formation and recombination (below). Source: [68]

(ER) events. It has been verified [48] that each of these types of event populate singlet and triplet states in different proportions (namely, mostly singlet and mostly triplet, respectively), and thus looking at the proportion of prompt yield over total yield it is possible to discriminate between them by the resulting waveform. For CENNS-10, the PSD parameter F_{90} is defined as the integrated amplitude of the measured waveform in a 90ns windows over the total amplitude for the $6\mu\text{s}$ waveform, roughly tracking the singlet to triplet intensity ratio.

However, there are still other background concerns; the foremost is the incidence of Beam-related Neutrons (BRN) in the active detector volume, since they can scatter elastically off Argon to mimic relevant Nuclear Recoil events. The mitigation of this kind of background is implemented via shielding and by the choice of location [48]; an outer lead shield is responsible for most of the mitigation, with an inner water shield stopping Neutrino-induced Neutrons (NINs) which might be produced by the interaction of the neutrino beam with the outer shield. Other steady state, beam-unrelated backgrounds are studied by looking at the signals from off-beam times.

Quenching

So far, we have described the ideal operation of the detector, in which a recoil event occurs inside the active volume and the scintillation process starts with yield proportional

to the deposited energy. However, there are competing processes within the medium which do not generate scintillation light, with examples being energy dispersion through elastic collisions without excitation of the Argon, dimer-splitting interactions with impurities, or ionizing interactions between excited atoms. Together, these will constitute *quenching effects*, such that the underlying energy recoil transferred to the medium will be different than the one measured by the light yield. To distinguish between them, henceforth the former shall be discussed in units of *keVnr*, and the latter in *keVee*.

The conversion between the two units is called *quenching factor*, and its determination for CENNS-10 was done fitting global average values for Liquid Argon in the range of 0-125 keVnr, and maintained constant and equal to the edge values outside this range. The polynomial fit value and its uncertainties are [48]:

$$QF(\text{keVee}) = (0.246 \pm 0.006) + ((7.8 \pm 0.9) \times 10^{-4} \text{ keVnr}^{-1})E_R. \quad (5.1)$$

Experimental bins and most measured quantities will be given in keVee, and thus it will be crucial to properly account for the equivalent bin edges in keVnr using the quenching factor: the latter is the true exchanged energy in the primary scattering process, which enters the calculations for the differential cross section. The threshold energy previously mentioned, however, was already given in keVnr, since it corresponds to the underlying energy required to trigger the detector at a benchmark efficiency.

There are two choices of binning for CENNS-10 data, depending on background mitigation strategies at the level of analysis; in particular, we'll use Analysis A [48] due to the release of its dataset by the collaboration [66]. It consists of 12 bins 10 keVee wide, starting from zero, for a region of interest corresponding to the full acceptance region, even though the nuclear recoil threshold (defined at 60% efficiency) corresponds to roughly 4.2 keVee.

5.2 Neutrino Flux

With the detector settled, the next step involves determining the differential flux, which will enter as a parameter in eq. (2.56). The CENNS-10 detector was positioned within the Neutrino Alley at the SNS, illustrated in Figure 13, sharing a common neutrino source with other COHERENT collaboration detectors.

The neutrino flux is produced via a π DAR (Pion Decay-At-Rest) setup: a pulsed proton beam from the SNS collides into a liquid-Hg target, producing neutrons and an abundance of π^+ mesons, which lose energy inside the surrounding monolith and decay at rest primarily to an anti-muon μ^+ and a muon neutrino ν_μ . The anti-muon then later decays at rest through a three-body decay to a positron e^+ , an electron neutrino ν_e and a muon antineutrino $\bar{\nu}_\mu$. A schematic view is shown in Figure 14.




Figure 13 – The COHERENT Collaboration experiments in the Neutrino Alley, at the SNS. Source: [47].

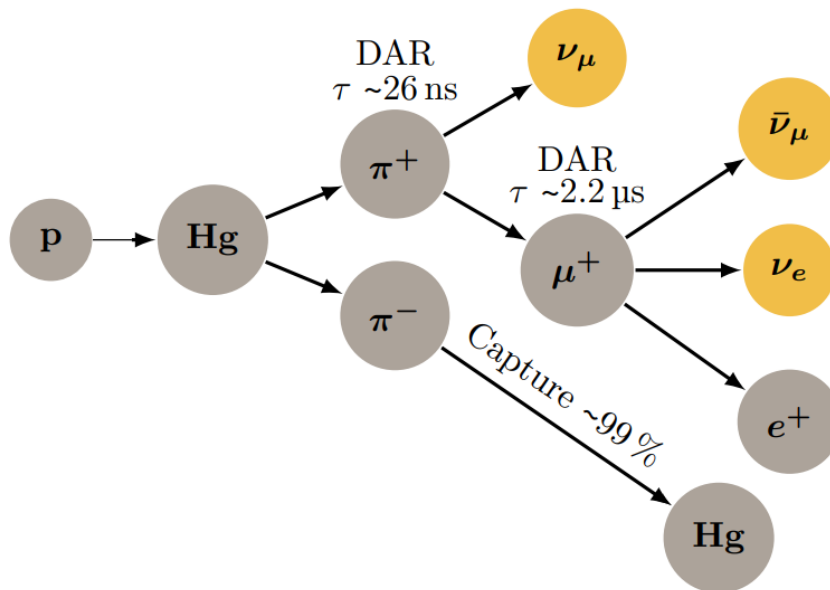


Figure 14 – Schematics of a pion decay at rest at the SNS. Source: [69].

Looking at the kinematics of the first process, in the LAB frame (in which the pion decays at rest):

$$\begin{aligned}
\pi^+ \rightarrow \mu^+ + \nu_\mu &\implies (m_\pi, 0) = (E_{\mu^+}, \vec{p}) + (E_{\nu_\mu}, -\vec{p}) \\
\stackrel{m_\nu \approx 0}{\implies} m_\pi^2 &= E_{\mu^+}^2 + 2E_{\mu^+}E_{\nu_\mu} + E_{\nu_\mu}^2 = m_{\mu^+}^2 + p^2 + 2p(m_\pi - p) + p^2 \\
&\implies 2p(m_\pi) = m_\pi^2 - m_\mu^2 \implies E_{\nu_\mu} = p = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} \quad (5.2)
\end{aligned}$$

one can see that the flux due to the final state neutrino constitutes a monoenergetic peak, sitting at $E_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} \simeq 29.8$ MeV. For the subsequent muon decay, this is no longer the case, as it decays through a three-body decay. In particular, if we consider the scenario where the outgoing positron has maximum momentum in the anti-muon rest frame,

$$\begin{aligned}
(m_\mu, 0) &= (E_e, \vec{p}) + (E_\nu, \vec{p}_1) + (E_{\bar{\nu}}, \vec{p}_2) \implies m_\mu - E_e = E_\nu + E_{\bar{\nu}} \\
&\implies m_\mu^2 = E_e^2 + 2E_e(E_\nu + E_{\bar{\nu}}) + (E_\nu + E_{\bar{\nu}})^2 \\
\stackrel{m_\nu \approx 0}{\implies} m_\mu^2 &= m_e^2 + p^2 + 2E_e(m_\mu - E_e) + p^2 \\
&\implies m_\mu^2 - m_e^2 = 2E_e^2 - 2m_e^2 + 2E_e(m_\mu - E_e) = 2E_e m_\mu - 2m_e^2 \\
\therefore E_e &= \frac{m_\mu^2 + m_e^2}{2m_\mu} \stackrel{m_\mu \gg m_e}{\longrightarrow} \frac{m_\mu}{2} \quad (5.3)
\end{aligned}$$

and for any other combination of momenta, the electron energy is smaller. Thus, we obtain a wider spectrum for all of the products, capped at $E_e, E_{\bar{\nu}_\mu}, E_{\bar{\nu}_e} < \frac{m_\mu}{2} \simeq 52.8$ MeV.

Due to the pulsing frequency of the SNS beam, compared to the muon lifetime, it is actually possible to distinguish the neutrinos from the prompt pion decay from the delayed muon decay, through their arrival time. This yields some degree of flavour distinction that allows for an easier characterization of the flux. Following the analysis of the kinematics of the processes, the total incoming neutrino flux is:

$$\frac{d\phi_{\text{Tot}}}{dE_\nu} = \frac{d\phi(\nu_\mu)}{dE_\nu} + \frac{d\phi(\nu_e)}{dE_\nu} + \frac{d\phi(\bar{\nu}_\mu)}{dE_\nu}, \quad (5.4)$$

$$\frac{d\phi(\nu_\mu)}{dE_\nu} = \eta \delta\left(E_\nu - \frac{m_\pi^2 - m_\mu^2}{2m_\pi}\right), \quad (5.5)$$

$$\frac{d\phi(\bar{\nu}_\mu)}{dE_\nu} = \eta \frac{64}{m_\mu} \left[\left(\frac{E_\nu}{m_\mu}\right)^2 \left(\frac{3}{4} - \frac{E_\nu}{m_\pi}\right) \right] \theta\left(E_\nu - \frac{m_\mu}{2}\right), \quad (5.6)$$

$$\frac{d\phi(\nu_e)}{dE_\nu} = \eta \frac{192}{m_\mu} \left[\left(\frac{E_\nu}{m_\mu}\right)^2 \left(\frac{1}{2} - \frac{E_\nu}{m_\pi}\right) \right] \theta\left(E_\nu - \frac{m_\mu}{2}\right), \quad (5.7)$$

where η is an overall normalization component we'll compute next, and with each component plotted in Figure 15 for better visualization.

In this setup, the neutrinos produced via decay are isotropically distributed with respect to the monolith target. Therefore, if r is the number of neutrinos of each flavor




Figure 15 – Components of the π DAR neutrino flux, given in arbitrary units. Source: [62].

produced by one proton-on-target, N_{POT} is the number of protons-on-target delivered for a given integrated exposure from the SNS beam, and L is the distance from the monolith to the detector, we can fix the normalization as

$$\eta = r \frac{N_{POT}}{4\pi L^2}. \quad (5.8)$$

For the CENNS-10 detector, following [44], we take $r = 0.09$, $N_{POT} = 13.7 \times 10^{22}$ and $L = 27.5\text{m}$. Since we are normalizing the flux by the integrated exposure, formally we're now talking about the number of events during the data collecting run of the detector; we'll still refer to it as the event rate to keep the definitions and formalism consistent throughout the text, but it should be understood as the number of events over the full exposure period.

5.3 Event Rate computation

Recalling equation eq. (2.57), the following quantities were necessary for the computation of the rate:

- **Differential Flux:** Provided in the last section by eq. (5.4) and the normalization we just presented;
- **Number of targets:** With a 10 L argon active volume at very high ($> 99\%$) purity, the number of target nuclei can be easily computed as:

$$N_T = N_A \frac{M_{\text{Det}}}{M_{\text{Ar}}^{\text{molar}}}, \quad (5.9)$$

where N_A is the Avogadro number, $M_{\text{Det}} = 24$ kg is the detector's mass, and $M_{\text{Ar}}^{\text{molar}} = 39.96$ g/mol is Argon-40's molar mass;

- **Recoil Energy bins:** There are 12 bins, 10 keVee wide. Starting from the second bin, the bin edges correspond, in keVnr, to:

$$\{E_R(i > 1)\} = \{36.4401, 67.0473, 93.959, 118.259, 145.56, 174.672, \\ 203.785, 232.897, 262.009, 291.121, 320.233, 349.345\}.$$

For the first bin ($i = 1$), its minimum energy depends on the threshold and acceptance function of the detector. In accordance to what was previously mentioned, in the following analysis we're including events below the nominal threshold at low efficiency, so that events are effectively cutoff at the acceptance limit of $E_R < 5$ keVnr [44].

- **Neutrino energy limits:** The minimum incident neutrino energy for a given recoil was already determined, via eq. (2.15), and it is naturally cutoff at $E_\nu^{\text{max}} = \frac{m_\mu}{2}$, as we've seen.

However, there is also another important factor to keep track of, which we have already alluded to but shall now present explicitly:

- **Efficiency:** While we formally wrote eq. (2.57) from a purely theoretical description, an important factor when connecting to the expected observed dataset at the experiment is the *Efficiency*. It corresponds to how effective the detector can be in triggering to events occurring inside its acceptance region. This should be convoluted with the event rate in order to get a real prediction for the event number observed. The efficiency per half-bin $A_{\text{eff}}(E_R)$ is included in the data release for the experiment [66].

Inserting everything in eq. (2.57), and including the efficiency, the total rate is therefore

$$R = \sum_{i=1}^{12} \frac{N_A M_{\text{Det}}}{M_{\text{Ar}}} \int_{E_R(i)}^{E_R(i+1)} dE_R A_{\text{eff}}(E_R) \int_{\sqrt{m_T E_R/2}}^{m_\mu/2} dE_\nu \sum_{\nu=\{\nu_\mu, \nu_e, \bar{\nu}_\mu\}} \frac{d\sigma(T, \nu)}{dE_R} \frac{d\phi(\nu)}{dE_\nu}. \quad (5.10)$$

The cross section is computed via eq. (2.31), and we used two different choices for the Form Factor: one calculation used the Helm parametrization described in eq. (1.15), while the other used the form factor described in [42], calculated from first principles following the discussion in 1.2, which includes semi-coherent responses as well. We have determined that the choice is completely negligible, accounting for a sub-% effect on the

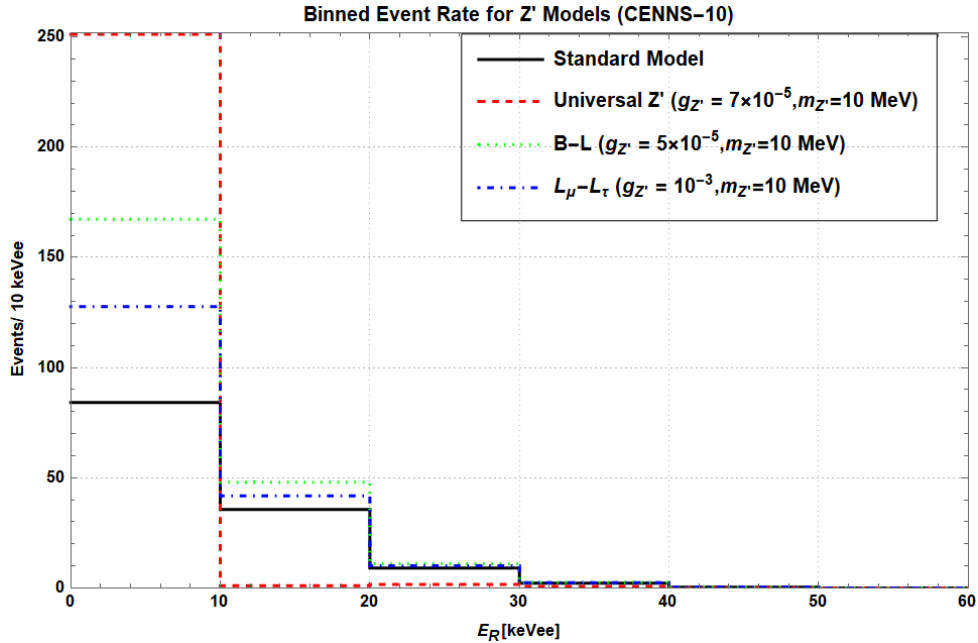


Figure 16 – Expected number of events for the models of interest compared to the Standard Model. Only the first 6 bins out of 12 are shown, with the remaining having a total 0 of expected events.

event rate per bin, at the relevant energies for CENNS-10. The resulting predicted event rates were:

$$R_i = \{84.21, 35.55, 9.07, 2.16, 0.32, 0., 0., 0., 0., 0., 0., 0.\} \quad (5.11)$$

with total $R = 131.32$, all in good agreement with [44, 56, 66]. Thus, we have successfully reproduced the CENNS-10 prediction for the Standard Model from the formalism constructed in Part I.

Computations for the models of interest follow the exact same formula, replacing the Standard Model differential cross section to the one corrected by the contribution of the Z' boson. The results are presented in Figure 16, and also agree with [56].

With the benchmark case working properly, in the next chapters we'll use the formalism to compute predictions for some experimental proposals of interest, justifying why they were chosen, before a final comparative analysis with COHERENT's sensitivity results.

6 Experimental Proposal - ν BDX-DRIFT

The main challenge for $CE\nu$ NS experiments is twofold: one needs to get enough sensitivity to detect tiny recoil energies of the order of tens of keVs, while simultaneously being able to exclude the numerous background brought about by hadronic and electronic sources triggering on such a low threshold. This is also the case for Dark Matter direct detection experiments, leading to a large integration between the fields from the experimental point of view. With increasing worries about the financial and footprint burden of building new detectors and testing new technology, the act of repurposing decommissioned detectors is a natural approach. The fundamental question, of course, is whether it is possible to obtain the necessary data from this repurposed technology, and how it compares to current dedicated detectors. In this context, we present and analyze the first experimental proposal of interest, ν BDX-DRIFT [70], based on a Dark Matter detector concept.

6.1 Detector Properties

The ν BDX-DRIFT is based on the Beam Dump eXperiment - Directional Recoil Identification From Tracks (BDX-DRIFT), a long-researched experimental proposal consisting of the employment of negative ion drift Time Projection Chambers (TPCs) to achieve directional resolution of low threshold recoil events [71]. Although initially conceived as a Dark Matter detector, the commissioning of the DUNE project [72] led to a recent proposal [70] to repurpose its technology as a $CE\nu$ NS experiment, to be placed in the planned Near Detector Facility of the DUNE complex, at Fermilab.

A brief overview of DRIFT detectors

A BDX-DRIFT detector consists of a collection of $1\text{m} \times 1\text{m} \times 1\text{m}$ box units placed behind a beam dump. In the configuration discussed in Ref. [73], each detector unit is filled with a mixture of 40 Torr CS_2 + 1 Torr O_2 . Two readout planes with a central cathode between them are to be placed along the beam direction, configuring two drift volumes. Each readout plane would be composed of several anode wires connected to a gain element, to which the ionized gas molecules can release its electrons for normal avalanche to occur. A simplified cut-view of a detector unit is shown in Figure 17. The distance between wires is related to the minimum track resolution length and sets a limit on the lower threshold resolution.

In contrast to the scintillation mechanism discussed in the previous chapter for




Figure 17 – Sketch of the BDX-DRIFT-1m detector, where z denotes the beam axis; it was originally planned to detect Dark Matter (χ) produced via a beam dump experiment. Source: [73].

the CENNS-10 detector, the working principle of a drift chamber relies on measuring the deposited energy via ionization tracks left by scattering events. These charged tracks are transported through the gas medium at constant speed due to an external electric field applied to the detector, until reaching the readout. There, the carried charge is deposited and yields a signal. The dispersion of these tracks due to natural diffusion and thermalization is one of the obstacles to setting proper detector precision, usually requiring an applied extra magnetic field to maintain cohesion, which limits the scalability of the setup [71].

While typical drift chambers rely on the ionization tracks being transported by electrons moving through the gas to the readout region, BDX-DRIFT's differential lies in using negative ion drift instead. This is more effective in preserving track shape, due to the ions remaining near thermal equilibrium with the gas up to relatively high drift fields, suppressing diffusion to the thermal limit [74]. In order to achieve this, a gas with strong electronegative nature is required, to bind the electrons tightly enough to maintain binding while subjected to the external electric field. CS_2 is the usual choice for the main

detector component. The low pressure configuration of the setup is also helpful to push recoil energy thresholds by limiting diffusion during drift.

A second electronegative gas is also used in order to get the production of additional ions which drift at a different speed compared to the main detector component. This allows for a measurement along the beam direction of the distance between the initial events and readout planes, as the arrival time difference is unique for each point along that axis. This is all done in order to ensure a careful 3D fiducialization of the active detector volume, allowing for a high degree of background rejection, in particular with respect to non-beam related sources. This concept has indeed proven successful in providing very clean results for previous runs of CS_2+O_2 -based DRIFT detectors in Dark Matter searches [75].

Another crucial characteristic of the setup is the ability to distinguish between nuclear and electron recoil events. This is imposed by restricting a threshold ionization current on readout; the principle is that for the same deposited energy, electrons leave much longer tracks due to their small mass (via the Bethe-Block Equation), in such a way that both the ionization density is reduced and the arrival time for the full tracks is much longer. Thus, imposing short reconstructions at readout at certain threshold compared to the typical arrival time filters out electron recoil events efficiently [76].

Application of DRIFT to $\text{CE}\nu\text{NS}$

The ν BDX-DRIFT proposal [70] takes advantage of the BDX-DRIFT detector setup described above to achieve good background mitigation and low detection thresholds. In a $\text{CE}\nu\text{NS}$ event, most of the recoil events in the typical CS_2 configuration of the detector would happen off Sulfur nuclei (with recoil energies up to tens of keVs). The resulting tracks would then end in a direction lying, at most, at 1 degree from the beam line, allowing for careful exclusion of most types of background-induced signals which would take place outside this region. However, beam-related backgrounds, such as neutrino-induced neutrons or inelastic neutrino processes, could still potentially induce scattering events similar to $\text{CE}\nu\text{NS}$ events. Their contribution has been computed in [70] and shown to be largely subdominant for ν BDX-DRIFT, therefore not presenting a problem for the detector's capabilities.

While previous deployments of DRIFT detectors worked mostly with a fixed 40 (+1) Torr configuration, it has been suggested in [70] that other possibilities could maximize the $\text{CE}\nu\text{NS}$ yield. Other choices of gas have also been considered and tested as viable alternatives to CS_2 (eg. CF_4 in [74]), and could be considered instead.

The effect of these different configurations on the expected event number is two-fold: it modifies the number of targets and changes the recoil energy threshold. Considering

a detector with a fixed active volume, working at room temperature, and taking a perfect gas approximation, the number of targets N_T is determined by

$$N_T = \rho(P) \frac{N_A}{m_{\text{molar}}} V_{\text{det}}, \quad \rho = 5.5 \times 10^{-5} \times \left(\frac{m_{\text{molar}}}{\text{g/mol}} \right) \left(\frac{P}{\text{Torr}} \right) \frac{\text{kg}}{\text{m}^3}, \quad (6.1)$$

where N_A is the Avogadro number, V_{det} is the detector volume, m_{molar} is the molar mass of the gas in the detector and P is the pressure.

Turning to the dependence of E_R^{min} on the gas pressure and chemical composition, we have

$$E_R^{\text{min}} = E_{\text{th}}(\text{Nuc}_i) = f_i \left(\frac{P}{40 \text{ Torr}} \right) \text{ keV}, \quad (6.2)$$

with $f_i = \{7.5, 20\}$ for $\text{Nuc}_i = \{C, S\}$, respectively [75]. Eq. (6.2) takes into account that the lower threshold for the recoil energy depends on the length of the tracks that can be resolved when they reach the readout, with lighter nuclei producing longer tracks and thus having a lower threshold. In addition, using the Bethe-Bloch equation, the energy loss scales with the density of the medium, limiting track length for a given recoil energy.

Since we are working in the ideal gas approximation, the density dependence can be converted into a pressure dependence using Eq. (6.1). Comparing Eq. (6.1) and (6.2) we see that larger pressures lead to a larger number of targets, but at the same time increase the minimum recoil energy, discarding lower energy events. Therefore, there is a trade-off that must be considered in searching for the ideal pressure configuration given the chemical composition of the gas inside the DRIFT chamber. It's also interesting to notice that this implies that the relative contribution of each nuclear species to the total rate can be drastically altered, due to the distinct energy regions probing different sectors of the form factors, which drops off differently for each nucleus. This can be seen explicitly in Figure 18, where we plot both form factors over an arbitrary recoil energy range.

Event discrimination and Upper limits

We have established how the working principles of the DRIFT detectors can be applied to CE ν NS, and how its setup can influence the expected event rate. One crucial detail needs to be emphasized, though, before moving on with the computation of the event rate. For CENNS-10, it was clear how to reconstruct the recoil energy in keVee for events measured, allowing for proper multibinned analysis over the recoil energy distribution of events. This is not so trivial for the case of ν BDX-DRIFT. This time, there are two relevant nuclear species inside the active volume, namely Carbon and Sulfur (Oxygen is neglected since its presence is mostly in trace amounts).

Previously, electron and nuclear recoil events could be discriminated looking at the inherent difference in track length and topology between them, but in the case of Carbon and Sulfur their mass difference is too small for a separation as clear-cut. In fact, for most

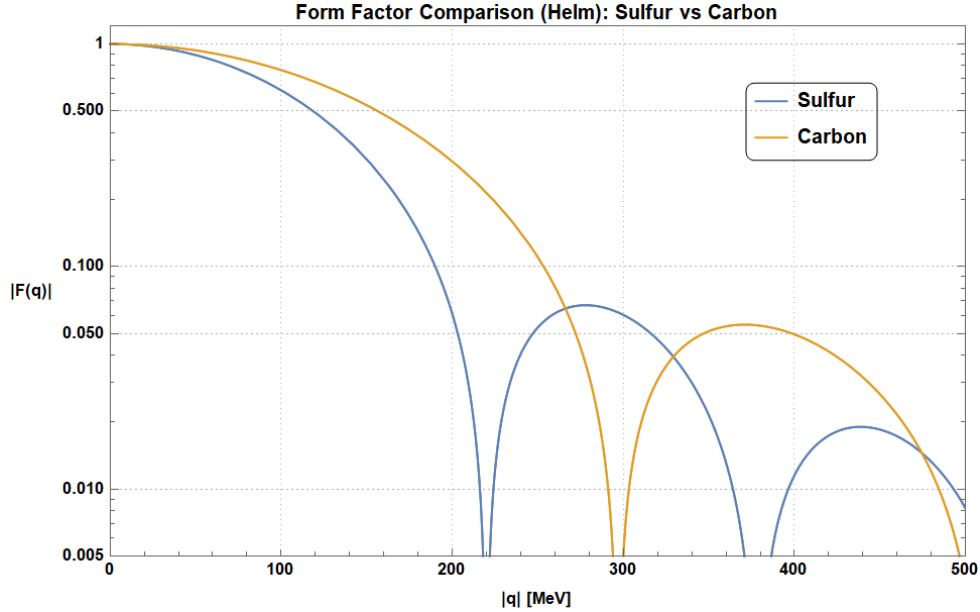


Figure 18 – Comparison between the Sulfur and the Carbon Form Factors, assuming the Helm Parametrization. Values for the nuclear mass and radii are taken from [45].

of the region of interest, there is an overlap - one can reconstruct the recoil energy given a track distribution only if the recoiling element is known for most of the phase space. While previous work [76] has shown some degree of distinction being possible by measuring other quantities besides track length, such as number of ion pairs produced, we have not found explicit demonstration of this being applied in a data-taking setting. Faced with this, we choose to move forward conservatively, assuming that no such distinction is possible in the overlapping region. As such, we'll limit our analysis to a single bin, since the definition of the recoil energy for a measured track is dubious above $E_{\text{th}}(S)$ and below the minimum between $E_R^{\text{max}}(C)$, $E_R^{\text{max}}(S)$. The original ν BDX-DRIFT proposal, likewise, uses a single-binned analysis for their results [70]. Since we are only interested in the overall event count in the whole region of interest, there is also no need to quantify quenching for the detector.

6.2 Neutrino Flux and Region of Interest

For the CENNS-10 case, the flux had a closed analytic form, with well-defined kinematics from the π DAR setup at the SNS. This will not be the case for ν BDX-DRIFT, where the incident neutrino flux is given by the LBNF Beam Dump spectrum, described in the DUNE Experiments's Technical Report [72], and reproduced here in Figure 19. Given the main goal of the complex as a laboratory for oscillation studies, the setup is optimized for the flux to be well-approximated as a pure ν_μ beam as it reaches the Near Detector Facility. In order to allow for the probing of CP Violation effects, a pure $\bar{\nu}_\mu$

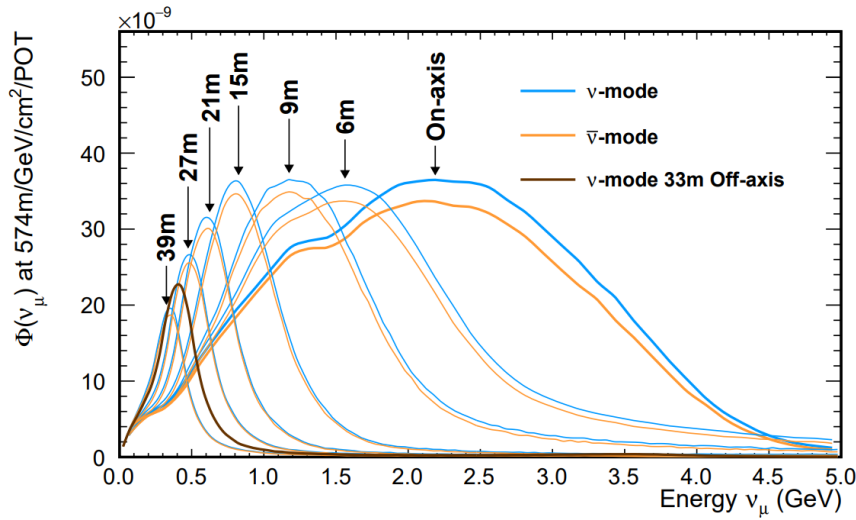


Figure 19 – Predicted Neutrino flux at the Near Detector Facility of the DUNE complex, for several positions with respect to the beam axis. Source: [72].

setting is also available.

The Near Detector facility allows for different configurations with respect to the beam axis; the inherent trade-off one must consider is that going off-axis will diminish the incident flux, but also mitigate beam-related backgrounds. For our case, the computation of the latter has already been mentioned as largely subdominant for the whole available spectrum [70], and thus only the on-axis configuration will ultimately be considered, in ν mode.

The energy and intensity of the DUNE beam are both bigger than that of the SNS beam, peaking between 2 and 3 GeV and with a tail extending all the way up to 5 GeV. In principle, this should already be considerably above the ideal energy region for $\text{CE}\nu\text{NS}$ phenomena, but we can again invoke the beam-related background computation in [70] to guarantee there is no significant loss of precision in using the whole available flux, instead of trying to impose cuts. This motivates the choice of the maximum incident neutrino energy as $E_\nu^{\text{max}} = 5$ GeV. In fact, this is particularly convenient because, under our conservative assumption of indistinguishability between Carbon and Sulfur events, there is no simple way to introduce cuts in recoil energy. Using the whole flux will solve this by setting a natural cutoff corresponding to the kinematic limit of the cross section, related to the maximum between each species' maximum recoil energy.

6.3 Event Rate

Summarizing the main points of the previous sections, we have fixed the main quantities needed for the computation of the event rate prediction:

- **Differential Flux:** Provided by the DUNE Technical Report [72], and presented in Figure 19;
- **Number of Targets:** Fixed by equation (6.1) for a given pressure configuration, assuming an ideal gas approximation. For this proposal, we assume an active volume of $10m^3$, corresponding to ten functional units connected along the beam axis;
- **Recoil Energy bins:** We'll assume a conservative single-binned analysis, with E_R^{\min} given by eq. (6.2) for each nuclear species, and extending through all of the allowed phase space, set for $E_r^{\max} = 2(E_\nu^{\max})^2/m_N \simeq 3192$ keV, set as the recoil limit for Carbon, which is chosen so as to not discard any Sulfur or Carbon event;
- **Neutrino Energy limits:** The minimum neutrino energy is again given by (2.15), by using the smaller recoil energy threshold between the two nuclear species as defined in the previous item. The maximum neutrino energy is set by the DUNE flux, as $E_\nu^{\max} = 5$ GeV, as mentioned before;
- **Efficiency:** Since the experiment is a future proposal, there hasn't been realistic studies of its efficiency. We follow [70] in assuming 100% detector efficiency.

The total event rate will then be a weighted sum of the individual rates due to each nuclear species, with coefficients determined by stoichiometry. Since Oxygen is only present in trace amounts compared to Carbon and Sulfur, we have

$$R_{\text{Tot}}(P) = R(C, P) + 2R(S, P), \quad (6.3)$$

where $R(X, P)$ corresponds to the event rate in eq. (2.56) for a given choice of pressure configuration, which determines $N_T(P)$ following (6.1) and E_R^{\min} via (6.2). Here, we employ the Helm Form Factor again for the computation for convenience, but we remark that at the higher energies of the DUNE flux, we could probe the region toward the tail-end of the form factor, where semi-coherent responses become more significant. In fact, [70] does a brief study of associated uncertainties among different phenomenological parametrizations for the form factor, and there indeed seems to be some discrepancy. We will present our analysis as a benchmark case, and leave the determination of the associated nuclear uncertainties for future work.

The number of Standard Model events as a function of the pressure inside the detector is show in Figure 20, for a science run of 7 years. We present each nucleus' contribution separately, in addition to their weighted sum; this could prove useful in future work if the capability for track identification of the detector were to be established.

Benchmark values for the models introduced in 3.4 are illustrated in Figure 21, under the same conditions.




Figure 20 – Expected number of $\text{CE}\nu\text{NS}$ events in $\nu\text{BDX-DRIFT}$ as a function of the pressure configuration (in Torr) for a data-taking period of 7 years. The vertical grey line shows the pressure value where the Carbon contribution overtakes the Sulfur one. The number of events is computed in the SM and, for illustrative purposes, we do not impose any cut on the maximum recoil energy.

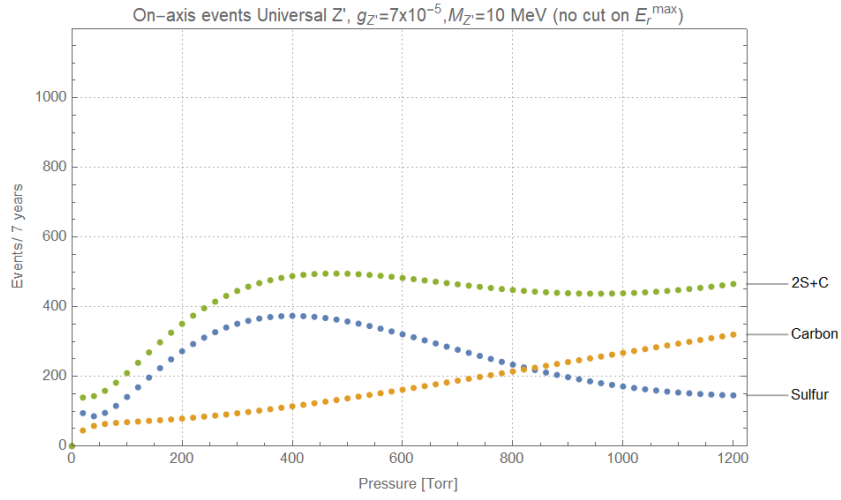

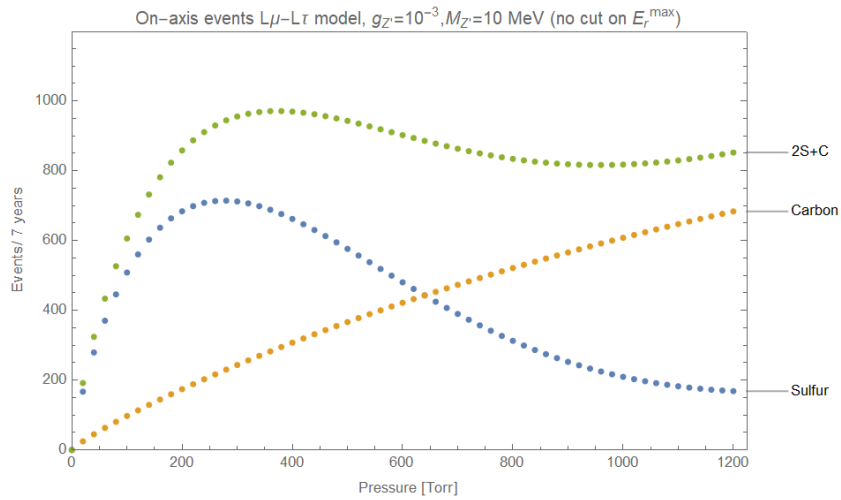
(a) ν BDX-DRIFT event rate predicted for the Z'_{Uni} model(b) ν BDX-DRIFT event rate predicted for the Z'_{B-L} model(c) ν BDX-DRIFT event rate predicted for the $Z'_{\mu-\tau}$ model

Figure 21 – Expected number of CE ν NS events in ν BDX-DRIFT as a function of the pressure configuration (in Torr) for a data-taking period of 7 years, for benchmark choices of $g_{Z'}$, $m_{Z'}$ in each of the models of interest.

7 Experimental Proposal - ESS-CE ν NS Detectors

Last chapter introduced an interesting idea using the “know-how” from Dark Matter detection designs to setup a somewhat unorthodox proposal for a CE ν NS experiment, taking advantage of the massive supportive structure of DUNE. In this chapter, we’ll present an ensemble of proposals closer to the concept of the COHERENT Collaboration, also trying to show the advantages of inserting them in the context of a new facility.

The European Spallation Source (ESS) is a multi-disciplinary research facility complex currently under construction on the outskirts of Lund, Sweden. It is planned to house the world’s most powerful pulsed neutron source, combining the world’s most powerful superconducting proton linac with an advanced hydrogen moderator [62]. With primary infrastructure construction finished as of December 2021, instrument deployment and calibration is ongoing until the planned start of its scientific programme in 2023.

While the primary research goals of the facilities are centered around Neutron Physics, its setup with a proton beam colliding into a dense target inside a monolith provides conditions almost identical to those of the Oak Ridge Laboratory’s SNS. Besides the neutron flux generated by spallation, this construction should also generate a π DAR neutrino flux which could be used for CE ν NS experiments in a realization close to the one adopted by the COHERENT Collaboration. This has led to myriad of proposals [77] by a part of the specialized community about ways to harness this new neutrino source for breakthrough studies. Before getting into the specifics of some of these proposals, it’s useful to look compare the capabilities of the ESS with the ones offered by the SNS.

7.1 ESS vs SNS

There are three main differences between the ESS and SNS: energy output, target composition and pulse timing. Since the neutrino fluxes for both originate from pion and muon decays-at-rest, the spectrum is essentially fixed up to a normalization, following the shape we already presented in eq. (5.4) and Figure 15.

This normalization is given by the total exposure (in POT) over a data-taking period multiplied by the neutrino yield per proton-on-target, and distributed isotropically from the target monolith to the detector position, following eq. (5.8). The exposure (in POT) can be computed by a simple formula:

$$N_{\text{POT}} = \frac{W}{E_p} \times T, \quad (7.1)$$

where W is the power of the beam, E_p is the average proton energy, and T is the accumulated time of beam delivery (ie. total time with the beam turned “on”). The SNS, at the time of the CENNS-10 operation, had a nominal power of 1.4 MW, delivering 1 GeV protons on target [48]. For its full operation capacity, the ESS is expected to reach nominal power of 5 MW, with 2 GeV protons [62]. With both operating for 5000 hours of on-beam time per year, this implies an increase in exposure by a factor of 2.5 for the ESS in comparison to the SNS. Naturally, this leads to an intensification of the flux, since its normalization is proportional to N_{POT} , and to a higher number of expected CE ν NS events over the same time period.

Another quantity influencing the normalization of the flux is r , the neutrino yield per proton-on-target. This is directly related to the number of pions produced for each collision of protons from the beam pulse with the target; different materials can result in different yields. There’s also a dependence on the incoming proton energy: higher energies facilitate double-pion production per interaction [78], raising the yield. Following [62], dedicated simulations performed show a conservative improvement of a factor of 3.75 for the neutrino yield favoring the ESS at full-capacity, with a tungsten target, compared to the Hg target at SNS with beam parameters $W = 1$ MW, $E_p = 1$ GeV.

The pulse time of the SNS is a crucial feature for flavor discrimination via prompt/delayed components of the flux. For the CENNS-10 configuration, 360ns Full-Width at Half-Maximum (FWHM) pulses were delivered on target at a 60Hz rate. This sits between the orders of the mean lifetimes of the pion ($\tau(\pi^+) \simeq 26$ ns) and of the muon ($\tau(\mu^+) \simeq 2200$ ns). The result is that pions produced from the late protons reaching the target in the tail of the pulse should still decay before the bulk of muon decays - therefore, the mono-energetic peak of ν_μ from the pion decay at rest should arrive earlier at the detector, dominating over the other neutrino components from muon decay for an early, narrow time window. This will not be the case for the ESS, with a design pulse timing of 2.8 ms, at 14 Hz, thus making it unsuitable for flavor discrimination. Some degree of flavor separation should still be possible, however, based on recoil energy distribution of events [62].

Another consequence of the long pulse timing at lower frequency is a substantial increase of the duty factor, corresponding to the ratio of beam on/off timings. This is inversely proportional to background mitigation, since it leaves a larger window open for background events due to surrounding environment effects. This is specially significant for the control of cosmic ray background in ground-level and shallow-depth experiments in general. Despite the duty factor increase, the ESS should still have a better signal-to-background figure of merit, due to the flux intensification [62]. This scenario does not take into account the influence of beam-related backgrounds, though.

7.2 Proposed Detector Setups

Despite the caveats due to pulse characteristics, the higher intensity of the neutrino flux proportioned at the ESS will be favorable to collect statistics at a projected faster pace than COHERENT's detectors at the SNS [62]. With the intention of the community to have the field moving past the discovery era into a precision one, this feature is crucial step in that direction. As such, several proposals of upgraded detection techniques to be employed at a CE ν NS program at the ESS have been put forward to try to make full use of its capabilities for Neutrino Physics. We shall very briefly go over some of them, referring to [62] for more complete descriptions and further references, and [77] for a grouping of other ESS-related CE ν NS proposals.

- **Cryogenic Scintillator (ESS-CsI)**

This proposal intends to use a liquid nitrogen-cooled (77k) undoped CsI scintillator as a low threshold detector array. It is similar to the one used for the first detection of CE ν NS by the COHERENT Collaboration [48], which was a Sodium-doped CsI[Na] detector at room temperature. The main advantage of the former over the latter would be its higher light output and an expectation of the measurement of quenching effects down to a sub-keV scale, boosting precision and allowing for higher statistics.

The suggested design for the ESS involves a small array of four $5 \times 5 \times 50$ cm³ crystals, for a total mass of 22.5 Kg, each read out by two LAAPDs (Large Area Avalanche Photodiode). With conservative estimates for the quenching factor, the threshold recoil energy could reach values around 1 keV_{nr}.

- **Charged-Coupled Device (ESS-Si)**

Silicon-based charged-coupled devices (CCDs) are a popular instrument for rare event searches, due to their capability for low threshold energies with reduced read-out noise. These have already been used in reactor-based CE ν NS initiatives (eg. [79]), as well as Dark Matter searches (eg. [80]) in recent years with good success. This proposal involves a Skipper CCD detector array, which is characterized by its capability to reserve the charge in a pixel for repeated non-destructive sampling, yielding in the ideal case a sub-electron readout noise for a large number of re-samplings. This boosts background rejection capabilities and massively lowers the threshold even in comparison to ordinary CCDs.

The suggested design for the ESS mimics the one proposed for DAMIC@SNOLAB [81], with a 1 kg skipper CCD and sampling reaching a threshold recoil energy of 0.16 keV_{nr}.

- **High-pressure Gaseous TPC (ESS-Xe)**

Another popular choice of low-threshold detector commonly used for Dark Matter searches, high-pressure noble gas TPCs present excellent energy resolution compared to other proposals even at very low energies. The pure gas setup is preferred in favor of dual phase proposals in this context to mitigate charge trapping and delayed release in the interface. These can typically be very limiting for low energy sensitivity detectors at shallow depths due to cosmic ray backgrounds, even with proper shielding.

While a variety of noble gases can be used, Xenon is the usual choice due to its high atomic weight, well-studied quenching and the possibility for analysis with different triggering requirements, either by primary scintillation (S1) only data or by its ratio to electroluminescent amplification of ionization (S2).

The suggested design for the ESS is inspired by the detector of the NEXT experiment [82], with a 20 kg active volume of Xenon with projected 0.9 keVnr threshold recoil energy.

- **p-type point contact HPGe (ESS-Ge)**

P-type Point Contact (PPC) Germanium detectors are semiconductor detectors which benefit from an enhanced capability of background rejection via pulse shape discrimination. These are specially efficient in distinguishing localized single-scatter events from multiple-site energy depositions, which are typical for Compton-like scattering backgrounds from cosmic rays and other sources. Combined with the possibility of sub-keV threshold and low noise ratio, this kind of detector is very well suited for rare, low energy searches, and already employed in other CE ν NS and neutrinoless double beta decay initiatives [83].

The suggested design for the ESS consists of two Germanium PPCs, totaling 7 kg of active mass, equipped with shield with double active veto to tag neutrino-induced neutrons produced from the outer shield. The projected threshold recoil energy for the detector is of 0.6 keVnr.

- **Scintillating Bubble Chamber (ESS-Ar)**

The CENNS-10 detector already demonstrated the powerful capability of scintillating liquid Argon chambers as a CE ν NS detector, but their threshold energy and low energy resolution are limited by a background derived from beta emission of cosmogenically produced Ar³⁹ inside the detector volume. This issue is mitigated with the use of liquid argon in a bubble chamber setup, which is extremely efficient in filtering out this kind of electron-recoil background, resulting in improved threshold and resolution [62]. Combined with the intrinsic discrimination of electron and nuclear recoils proportioned outside of the Ar³⁹ emission region via scintillation

timings, described previously for CENNS-10, this yields the best degree of NR/ER discrimination amongst all of the detector proposals.

As a trade-off, however, energy reconstruction is not viable for this case, limiting the analysis to a single-binned total event rate count, as was the case for the conservative approach taken for the ν BDX-DRIFT proposal. While we will not delve into this possibility, an excellent control of the energy thresholds could enable a two-step analysis with different threshold choices, allowing the experiment to probe possible momentum dependencies of New Physics scenarios in specific windows.

The suggested design for the ESS consists of an active volume with 10 kg, and projected threshold recoil energy of 0.1 keVnr.

- **Standard Bubble Chamber (ESS-C₃F₈)**

The last proposal concerns a C₃F₈ bubble chamber, inspired by the PICO Experiment [84]. Like for the Argon bubble chamber discussed above, it boasts high degree of discrimination for nuclear and electron recoil events, though it is limited to total event rate analysis and presents a higher threshold compared to the other proposals. However, a vast amount of data and further studies of related backgrounds, detector capabilities and overall operation is available due to past and current research for the PICO collaboration, presenting a very realistic test study.

The suggested design for the ESS consists of an active volume with 10 kg, and 2 keVnr threshold recoil energy.

7.3 Neutrino Flux and Region of Interest

As aforementioned, the neutrino flux at the ESS will follow a traditional π DAR distribution, given by eq. (5.4), with normalization given by (5.8). While the exact positioning of the detectors will depend on available space within the beam dump alley of the facility, we shall consider a distance $L = 20$ m for all proposals. At full capacity, the exposure is 3×10^{23} POT/year, and a full science run comprised of 3 years of data taking is assumed. The higher energy compared to the SNS yields a rate of neutrinos per POT, per flavor of $r = 0.3$.

Again, the whole available range of the flux will be considered, setting $E_\nu^{\max} \simeq 52.8$ MeV, and the maximum recoil energy for each nuclear species given by (2.14). Except for the bubble chamber designs, which will only work with a total event rate analysis, we shall define the bins for the proposals according to the energy resolution of each detector at threshold, via a Gaussian energy smearing such that

$$\sigma(E_R) = \sigma_0 \sqrt{\frac{E_R}{E_R^{\text{th}}}}, \quad (7.2)$$

where σ is the energy resolution and σ_0 the resolution at $E_R = E_R^{\text{th}}$. Each bin is then defined by

$$2\sigma \left(\frac{E_R^r(i) + E_R^l(i)}{2} \right) = E_R^r(i) - E_R^l(i), \quad (7.3)$$

meaning that for the i -th bin with left and right edges at, respectively, $(E_R^l(i), E_R^r(i))$, its width is equal to twice the energy resolution at its center. To construct the bins, we define $E_R^l(1) = E_R^{\text{th}}$, for each detector.

7.4 Event Rate

Summarizing the main points of the previous sections, we have again fixed the main quantities needed for the computation of the event rate prediction:

- **Differential Flux:** Provided by eq. (5.4), with normalization given by eq. (5.8) with $r = 0.3$, $N_{\text{POT}} = 6 \times 10^{23}$ POT for a 3-year run and $L = 20$ m for all proposals;
- **Number of Targets:** Fixed for each detector by equation (5.9), where M^{molar} is given by the molecular molar mass for detectors with multiple nuclear species.
- **Recoil Energy bins:** For the bubble chambers, a single-binned analysis will be assumed, with E_R^{min} given by each detector's threshold recoil energy, and extending through all of the kinematically allowed region, set by eq. (2.14). For the other experiments, the bins are given by eq. (7.3), setting $E_R^l(1) = E_R^{\text{th}}$;
- **Neutrino Energy limits:** The minimum neutrino energy is again given by (2.15), by using the recoil energy threshold for each detector. The maximum neutrino energy is set by the kinematically allowed region, encompassing the whole flux, with $E_\nu^{\text{max}} = \frac{m_\mu}{2}$;
- **Efficiency:** Following [62], the efficiency will be assumed as a step-function, with $A_{\text{eff}(E_R)} = 0$ for $E_R < E_R^{\text{th}}$ and $A_{\text{eff}(E_R)} = 0.8$ for $E_R \geq E_R^{\text{th}}$, for each detector.

The total rate per bin for a detector will again be a weighted sum over the nuclear species present, with coefficients given by simple stoichiometry. The rate per bin for a given detector is given by

$$R(i) = \sum_X C_X \frac{N_A M_{\text{Det}}}{M_{\text{molar}}} \int_{E_R(i)}^{E_R(i+1)} dE_R A_{\text{eff}}(E_R) \int_{E_\nu^{\text{min}}}^{E_\nu^{\text{max}}} dE_\nu \sum_{\nu=\{\nu_\mu, \nu_e, \bar{\nu}_\mu\}} \frac{d\sigma(T, \nu)}{dE_R} \frac{d\phi(\nu)}{dE_\nu}. \quad (7.4)$$

with C_X the stoichiometry coefficient for a nuclear species X, and all the relevant quantities are given for this species as previously defined.

All the relevant information for each detector is summarized in Table 1. We also adopt a background prediction based on [62], already including the duty factor reduction for the ESS.

Detector Technology	Target	Mass [kg]	E_r^{\min} [keV _{nr}]	$\frac{\Delta E}{E}$ E_{th}	(%)	E_r^{\max} [keV _{nr}]	background [day ⁻¹]
Cryogenic scintillator	CsI	22.5	1	30	46.1	406	
Charge-coupled device	Si	1	0.16	60	212.9	8.5	
High-pressure gaseous TPC	Xe	20	0.9	40	45.6	357.6	
p-type point contact HPGe	Ge	7	0.6	15	78.9	329	
Scintillating bubble chamber	Ar	10	0.1	40	150.0	4×10^{-2}	
Standard bubble chamber	C ₃ F ₈	10	2	40	329.6	4×10^{-2}	

Table 1 – Summary of properties for the detectors possibly employed at the ESS: target, mass, recoil energy threshold, width for the smearing and maximum recoil energy, steady-state background. The steady-state background includes the 4×10^{-2} reduction by the ESS duty factor.

In order to compute the event rates, we have assumed again the Helm Form Factor via eq. (1.15). Following the verification done against the alternative form factor of [42] in the context of the SNS’ π DAR flux, this should not influence results at the energy range we are considering. Due to limitations on the precision and availability of data on the neutron radius and its distribution within each particular nuclear species, we shall work with a $F_p(q) = F_n(q)$ approximation, which shouldn’t introduce large errors in this energy range. The relevant isotopes and respective radii necessary for the computation of the form factor are presented in Table 2, with masses taken from [45].

	¹³³ Cs	¹²⁷ I	²⁸ Si	¹³² Xe	⁷² Ge	⁴⁰ Ar	¹² C	¹⁹ F
R_p [fm]	4.80	4.75	3.12	4.79	4.06	3.43	2.47	2.90

Table 2 – Isotopes and respective nuclear radius values, taken from [45].

The resulting total rates for the Standard Model prediction on each of the detectors are presented in Table 3, in agreement with the predictions from [62]. We shall not present benchmark computations for the Z' prime models in favor of focusing on the sensitivity results in the next chapter, due to the large number of detectors and models considered.

	ESS-CsI	ESS-Si	ESS-Xe	ESS-Ge	ESS-Ar	Ess-C ₃ F ₈
SM Events/year	8591	81	8033	1603	1393	540

Table 3 – Standard Model prediction for the total number of events per year for all of the detector proposals.

8 Sensitivity Analysis

After careful construction of the necessary theoretical framework in Part I, in Part II we have so far applied that formalism to a well-studied benchmark case for verification, and followed it up with descriptions for new experimental proposals, pointing out how to extract the needed information for event rate computation from them. This has been all done in preparation for the analysis in this chapter. We shall combine everything we've described and discussed in the rest of this work in order to measure how effective these new experiments may be in constraining the parameters from the models defined in Sec. 3.4. Specifically, we want to answer: how much of the parameter space of these models may be excluded at a 90% Confidence Level relative to the Standard Model prediction?

We start by reintroducing our definition for the binned χ^2 statistic, previously defined in Chapter 4:

$$\begin{aligned} \chi^2(g_{Z'}, m_{Z'}) = & 2 \min_{(\xi_S, \xi_B)} \sum_i \left[\xi_S N_i(g_{Z'}, m_{Z'}) + \xi_B N_{B,i} \right. \\ & \left. - (N_{SM,i} + N_{B,i}) \left(1 + \log \left(\frac{\xi_S N_i(g_{Z'}, m_{Z'}) + \xi_B N_{B,i}}{N_{SM,i} + N_{B,i}} \right) \right) \right] \\ & + \left(\frac{1 - \xi_S}{\sigma_S} \right)^2 + \left(\frac{1 - \xi_B}{\sigma_B} \right)^2. \end{aligned} \quad (8.1)$$

Here, $N_i(g_{Z'}, m_{Z'})$ is the CE ν NS event rate for the i -th bin computed for one of the Z' models for a given value of coupling and mediator mass; $N_{B,i}$ is the number of background events for the i -th bin; $N_{SM,i}$ is the number of CE ν NS events predicted by the Standard Model for the i -th bin; ξ_S, ξ_B are nuisance parameters respectively related to the normalizations of the signal event rate (which is CE ν NS) and background rate, with corresponding uncertainties σ_S and σ_B .

Assuming that our sample is large enough so that Wilks' Theorem applies, we can build the desired confidence interval at 90% C.L. by determining the boundary via eq. 4.11, such that $\sqrt{\chi^2} = 1.645$. This assumption should be reasonable given the verification of its validity for the CENNS-10 detector results in [48] and studies of diagonal NSI parameters with CsI in [47], according to the work presented in [85].

ν BDX-DRIFT Background and Uncertainties

For the ν BDX-DRIFT proposal, we use a single bin constraining only the total event rate. While the beam-related background due to Neutrino-induced Neutrons is negligible, the degree of steady-state and other beam-unrelated backgrounds is hard to determine precisely, and should vary a lot depending on the effectiveness of shielding

techniques employed. While results from the COHERENT collaboration [47, 48, 65] prove optimistic in that front, we proceed with the most conservative case outlined in [70]. This consists in setting $N_{B,i} = 0.25 \times N_{SM,i}$. In other words, we are adding the background as a 25% fraction on top of the expected signal event rate in each bin.

The uncertainty σ_S , related to normalization of the signal event rate, is dominated mostly by the uncertainty in the determination of the neutrino flux and the contribution from nuclear effects appearing due to the simplified form factor parametrization, specially in the higher energy tail. We estimate uncertainties of 10% for both, adding in quadrature to obtain $\sigma_S \simeq 0.14$ for this analysis. Due to the capabilities of direction reconstruction for events, we expect systematics related to the background events to be small, defining $\sigma_B = 0.01$.

ESS Backgrounds and Uncertainties

For the ESS proposals, we shall closely follow the values described in [62] for the background distribution, presented here in Table 3. Bins are defined following eq. (7.3) for each of the detector proposals with the exception of the Argon and C_3F_8 bubble chambers, which will use a single-binned analysis following their capabilities described in Chapter 7. Again following [62], we shall adopt signal and background normalization uncertainties of $\sigma_S = 0.1$, $\sigma_B = 0.01$ for all detectors. The former is based on an estimate of improvement for the quenching factor uncertainty with respect to [47], which was the dominant contribution for the first run of the CsI[Na] detector; this has indeed been concretized in the most recent results for the detector [65]. Likewise, the latter assumes a similar degree of uncertainty in background determination from on/off beam analysis in comparison to COHERENT's previous results.

Using the aforementioned values, we apply the testing algorithm outlined in Sec. 4.2, iterating over the three models of Sec. 3.4 for each of the detectors. The obtained sensitivity curves are presented next.

8.1 Sensitivity for the Universal Z' Model

The results for the Universal Z' Model are shown in Figure 22, showing the 90% C.L. limits obtained for ν BDX-DRIFT (left) in two distinct pressure configurations, and for the various ESS-based detectors (right).

In the left panel, the orange sensitivity curve is obtained by simulating the ν BDX-DRIFT detector filled with CS_2 at $P = 411$ Torr, determined in [70] as the configuration which maximized the signal rate for the Standard Model. However, the sensitivity is dependent on the difference and on the ratio between the Standard Model and Universal Z' predictions, and not just on the absolute value of the two. It's natural, therefore, to

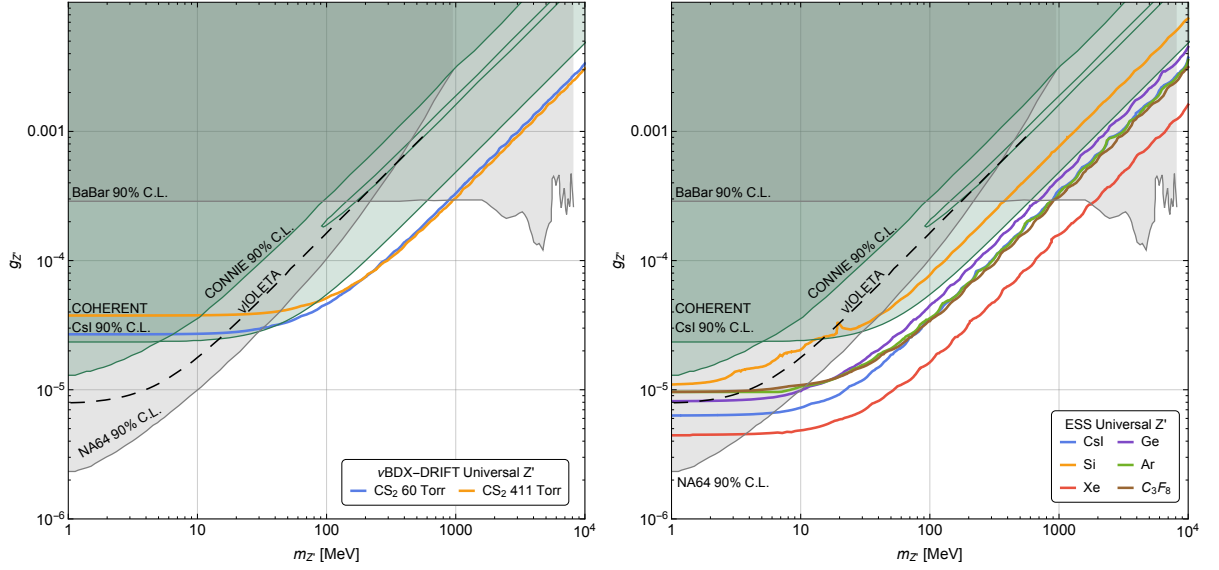


Figure 22 – Future sensitivity at 90% C.L. in the $m_{Z'}$ – $g_{Z'}$ plane for the universal Z' model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green regions are excluded at 90% C.L. by CONNIE [79] and COHERENT [56] assuming dominant Z' decays to SM states. The dashed black curve shows the sensitivity of ν IOLETA [86]. The gray shaded areas are excluded at 90% C.L. by BaBar [87] and NA64 [88, 89] assuming that the Z' decays dominantly in invisible dark sector states.

wonder whether there might be an alternative configuration which yields better sensitivity, at the cost of a reduced number of overall events.

The blue curve is obtained for an alternative configuration choosing $P = 60$ Torr, and shows a more stringent result compared to the previous one in the $m_{Z'} \lesssim 100$ MeV region. Still, it is very slightly dominated for the higher mass regions. There were two main motivations for this choice of alternative pressure value:

- Pressure configurations around 40 Torr have already been extensively studied and applied in the context of Dark Matter searches using previous detectors from the DRIFT line, being a convenient choice from a practical point of view [73];
- Following the determination of the Z' contributions in the Simplified Model approach and the related discussion in Sec. 3.4, the size of the modifications to the Standard Model charge increase for processes with lower exchanged momentum, which is related to the recoil energy of the event (2.9):

$$|\epsilon_N| \propto \frac{1}{m_{Z'}^2 + |\vec{q}|^2} \simeq \frac{1}{m_{Z'}^2 + 2m_T E_R}. \quad (8.2)$$

Here, the mass of the mediator acts as a regulator of sorts, since the factor becomes essentially constant when $m_{Z'} \gg |\vec{q}|$. Inversely, if the mass of the mediator is suffi-

ciently small, there exists a range of recoil energies where there will be a significant effect following the variation of the exchanged momentum. For a conservative example, using $|\vec{q}| \approx 50$ MeV as the maximum momentum established for the coherent condition back in eq. (1.11), and assuming a mediator mass of $m_{Z'} = 100$ MeV, this corresponds to a nearly 25% variation for $|\epsilon_N|$ within the whole kinematically available range.

We've remarked back in Sec. 6.1 that there is a trade-off to be considered involving the pressure configuration value, since it is proportional to the number of targets N_T , but increasing it also raises the threshold recoil energy for the detector, pushing more low-energy events out of the rate prediction. It is not trivial to determine analytically how much these effects offset each other. Nevertheless, what we just established is that the second effect could in principle introduce a larger variation, due to extending a range where the Z' contribution is maximal via its unique q -dependence. Since the sensitivity, as aforementioned, is related to the difference between the predicted rates, it's a reasonable assumption to expect that lower pressure configurations could exhibit better sensitivity for lighter mediator masses (in particular, $\lesssim O(100\text{MeV})$).

Once the bound for the $P = 411$ Torr configuration was fixed, we randomly selected a few points over the curve $f_{90\%} = (g_{Z'}, m_{Z'})$. We then computed again the significance for those values of the parameters with randomly chosen alternative pressure P' configurations, as a qualitative quick check of how close to optimal $P = 411$ Torr is for the sensitivity determination. If they resulted in $\sqrt{\chi^2(g_{Z'}, m_{Z'}, P')} > 1.645$ (the 90% C.L. threshold), shifting configurations would result in gain of sensitivity for that point. As a general trend, lower pressures relative to the orange curve yielded better sensitivity for the region of lower mediator mass, a possibility outlined in the second point above. Taking into account the practical aspects outlined in the first point, a representative value of $P = 60$ Torr was chosen for a full rerun of the analysis, resulting in the blue curve in the left panel of Figure 22.

The right panel shows the the obtained sensitivity for the six detector proposals for the ESS under consideration, identified by their target material as CsI (blue), Ge (purple), Si (orange), Ar (green), Xe (red) and C_3F_8 (brown), following the identification in Table 1. Stronger sensitivity is obtained by the detectors with higher atomic mass number (CsI and Xe) as would be naively expected due to the coherent enhancement. C_3F_8 however is an exception to this, due to possessing the highest threshold among all the detectors.

Blind Regions

For both panels, there is a thin diagonal strip inside the region delimited by COHERENT-CsI which is actually not excluded at 90% C.L., as shown in [56]. It starts

around $m_{Z'} \approx 100$ MeV, extending continuously in the direction of heavier mediator masses. We'll name this as a *blind region* for the detector. In order to understand its meaning, let's look again at $|\epsilon_N|$, which is the Z' contribution shown in eq. (3.13).

While the previous discussion highlighted the effect of the $|\vec{q}|^2$ dependence becoming more significant for lighter mediators, one can also consider what happens in the case of heavier Z' masses. Following the discussion in Sec. 2.3, as $m_{Z'} \rightarrow m_Z$ the q -dependence becomes suppressed, and we can again neglect the dynamics of the mediator in the Simplified Model approach in favor of a Fermi-like EFT. Ultimately, it corresponds to adopting the NSI parametrization as discussed in Sec. 3.2. Since we're taking $F_p(q) \simeq F_n(q)$, in this limit the nucleon charges factorize from the energy dependency, and the Rate will be proportional to

$$R \propto \left[Z \left(\frac{1}{2} - 2 \sin^2 \theta_W \right) - \frac{A-Z}{2} + A \frac{3g_{Z'}^2}{\sqrt{2}G_F m_{Z'}^2} \right]^2 \equiv \tilde{Q}^2. \quad (8.3)$$

This rate is clearly invariant under a transformation $\tilde{Q} \rightarrow -\tilde{Q}$. Neglecting the proton contribution from the Standard Model due to the value of the mixing angle, this implies that we have an exact degeneracy between the Universal Z' model and the Standard Model rates if the NSI parameter $\epsilon \equiv g_{Z'}^2/m_{Z'}^2$ satisfies

$$\left[-\frac{A-Z}{2} + A \frac{3\epsilon}{\sqrt{2}G_F} \right] = \frac{A-Z}{2} \implies \epsilon = \frac{G_F \sqrt{2}}{3} \frac{A-Z}{A}. \quad (8.4)$$

This explains why the blind region consists of a diagonal strip in the $g_{Z'} \times m_{Z'}$ space, as it is equivalent to a line with $g_{Z'}/m_{Z'}$ constant, so the value of ϵ that yields the degeneracy is preserved throughout its extension.

However, this is not exact. While not neglecting the proton contribution will still yield an exact (although, with a less elegant expression) degeneracy, introducing the proper q -dependence lifts it since now the charge can't be carried through the E_R integration, such that the rate is now instead proportional to

$$R \propto \int dE_R \left[-\frac{A-Z}{2} + A \frac{3g_{Z'}^2}{\sqrt{2}G_F(m_{Z'}^2 + 2m_T E_R)} \right]^2, \quad (8.5)$$

remembering that $|\vec{q}|^2 = 2m_T E_R$, following eq. (2.9). Now, fixing a point $(g_{Z'}, m_{Z'})$, we can *still* solve the exact degeneracy condition via (8.4) simply adding to the denominator the $|\vec{q}|^2$ term. Although, since we will now be integrating over E_R , we are effectively summing over an interval where every value yields a non-degenerate differential rate, except for a single point; in general, this should not yield a degenerate total rate.

Nevertheless, for the Universal Z' Model, it's a fact that there is a non-trivial region where its rate approaches the Standard Model one due to a sign reversal of the effective charge as the new coupling grows. The general position and shape depends on

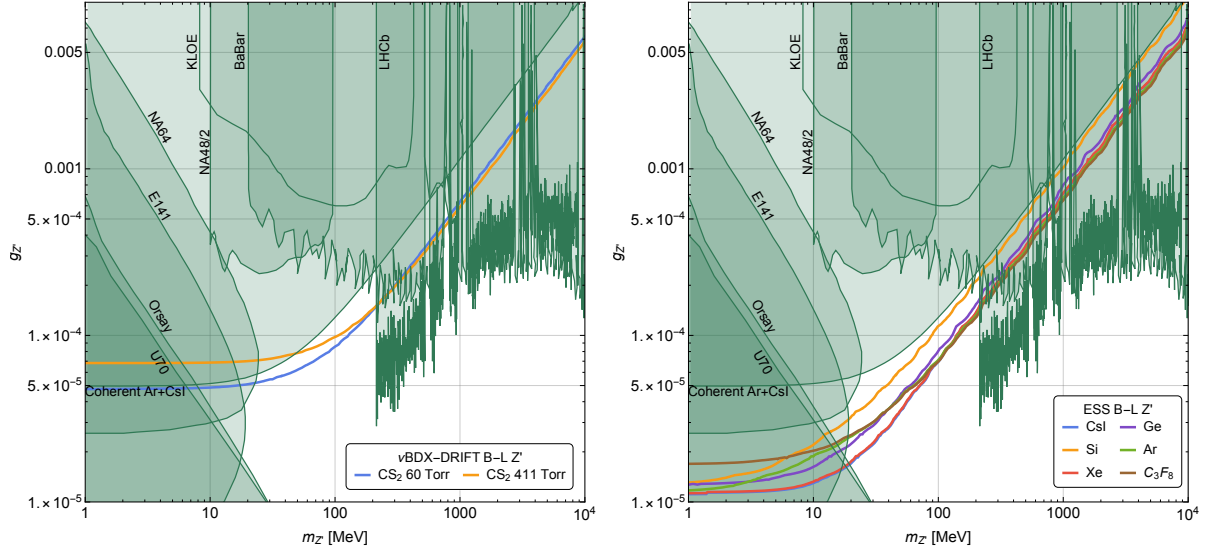


Figure 23 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $B - L$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green regions are excluded at 90% C.L. by NA64 [95], E141 [96], Orsay [97], U70 [98], COHERENT [56], NA48/2 [99], KLOE [100, 101], BaBar [102] and LHCb [103] assuming dominant Z' decays to SM states.

the detector’s composition, with a rough proportionality to $\frac{A-Z}{A}$ as the mediator mass grows. This will be a good motivator for the deployment of multiple $\text{CE}\nu\text{NS}$ detectors with different target compositions, so as to cover each others’ blind regions.

Due to our procedure for constructing these limits, mapping the blind region in detail is very intensive computationally. As such, we performed limited benchmark tests which showed a very thin line, already within the region excluded by COHERENT for both ν BDX-DRIFT configurations and most ESS detectors. As such, we do not present it in Figure 22. The same applies to the $L_\mu - L_\tau$ model, which will be discussed later.

Back to the plots, for both panels, the dark green shaded regions are already excluded at 90% C.L. by CONNIE [79] and COHERENT [56]. The black dashed curve furthermore shows the sensitivity that another proposed experiment, the Neutrino Interaction Observation with a Low Energy Threshold Array (ν IOLETA), can reach for the universal Z' model [86, 90–94]. These searches assume that the Z' decays most of the time to SM states. If the Z' decays predominantly to invisible dark sector particles, existing searches from the BaBar [87] and the NA64 [88, 89] experiments apply and exclude the gray region.

8.2 Sensitivity for the $B - L$ model

Figure 23 presents the 90% C.L. sensitivity obtained for the $B - L$ model for each of the experimental proposals. The dark green shaded area is excluded by existing experimental searches assuming dominant Z' decays into SM particles, from fixed target experiments (NA64 [95], E141 [96], Orsay [97] and U70 [98]), COHERENT [56], KLOE [100, 101] and NA48/2 [99]. Furthermore, as shown in Ref. [104], it is possible to reinterpret the BaBar [102] and LHCb [103] bounds on dark photons and adapt them to the $B - L$ case.

The panel on the left of Figure 23 shows the sensitivity for the ν BDX-DRIFT detector with the same configuration as used for the Universal Z' analysis, CS_2 at 60 Torr (plotted in blue) and 411 Torr (plotted in orange). While in the previous case there was only a very marginal improvement over COHERENT-CsI detector's bounds, for the B-L model there is a more significant improvement roughly over the range of $30\text{MeV} < m_{Z'} < 200\text{MeV}$, in particular for the 60 Torr case. This region sits between the most stringent limits from accelerator and fixed targets experiments, and ν BDX-DRIFT would help to improve the coverage between them

The 90% C.L. sensitivity curves for the experiments at the ESS are shown in the right panel. These detectors, and in particular the one based on Xe, will be able to probe a good part of the parameter space that currently is not excluded by COHERENT, collider or fixed target experiments. It also encompasses the entirety of the region sensitive to ν BDX-DRIFT.

On the other hand, we show in Figure 24 how the future ν BDX-DRIFT and ESS sensitivity compare to current bounds assuming the Z' to dominantly decay into invisible particles belonging to a dark sector. The gray region is currently bounded by BaBar [87] and NA64 [88, 89] searches.¹ Both sets of detectors have the potential to improve the limits in the range $100 < m_{Z'}/\text{MeV} < 500$ and above 8 GeV, where the BaBar experiment abruptly loses sensitivity and the bounds from LEP [106] are too weak.

Differently from the Universal Z' case, we do not expect blind regions for the B-L model. This can be understood directly by looking at the degeneracy condition in eq. (8.4) after replacing the B-L contribution:

$$\left[-\frac{A-Z}{2} - A\frac{\epsilon}{\sqrt{2}G_F} \right] = \frac{A-Z}{2} \implies \epsilon = -G_F\sqrt{2}\frac{A-Z}{A}. \quad (8.6)$$

Since $A > Z$, this implies $\epsilon < 0$. However, we defined $\epsilon = g_{Z'}^2/m_{Z'}^2$ previously, where both parameters are real numbers, and therefore ϵ must be non-negative, such that the degeneracy condition is not satisfied. Furthermore, in the previous case, the blind region

¹ The $B - L$ model can also be probed in neutrino-electron scattering at DUNE [105]. Since, however, the future sensitivity is not derived using CE ν NS, we do not show this region in our plot.




Figure 24 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $B - L$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The gray shaded areas are excluded at 90% C.L. by BaBar [87] and NA64 [88, 89] assuming that the Z' decays dominantly in invisible dark sector states.

was effectively generated by a change of sign of the effective charge, which is not present here: both the Standard Model and Z' contributions have the same (negative) sign, and the event rate is monotonic with respect to $g_{Z'}$.

8.3 Sensitivity for the $L_\mu - L_\tau$ model

Figure 25 shows the 90% C.L. sensitivity obtained by the ν BDX-DRIFT detector at the LBNF (left) and several hypothetical experiments at the ESS (right) for the $L_\mu - L_\tau$ model. The sensitivities are compared with the 95% C.L. excluded regions (dark green) obtained in [111] exploiting the neutrino trident cross section measured by the CCFR collaboration [112], by the SM Z bosons decaying into four leptons searches at the ATLAS [108, 109] and CMS [110] experiments (which can be reinterpreted assuming that the SM Z boson decays into a Z' and two muons), and by the BaBar search for $e^+e^- \rightarrow Z\mu^+\mu^-$, where the Z' decays into muons [107]. Finally, we show also the reinterpretation by [114] of the Borexino limits [113] (dark green) and the 2σ region needed to explain the anomalous magnetic moment of the muon (red region) [7, 115, 116].

While the ν BDX-DRIFT detector will not be sensitive to regions not already excluded by current existing searches, the various experiments proposed for the ESS have the potential to probe a large sector of unexplored region, in the range $3 < m_{Z'}/\text{MeV} < 70$. Additionally, all of them should be able to reach the so called “ $(g - 2)_\mu$ ” band in the low mass region. It is defined as the combination of coupling and masses needed for to explain




Figure 25 – Future sensitivity at 90% C.L. in the $m_{Z'} - g_{Z'}$ plane for the $L_\mu - L_\tau$ model for the ν BDX-DRIFT detector (left) exploiting CS_2 at two different pressures, and the various detectors described in the text at the ESS (right). The dark green shaded areas are excluded at 95% C.L. by BaBar [107], ATLAS [108, 109], CMS [110], CCFR [111, 112] and at 90% C.L. by Borexino [113, 114] assuming that the Z' decays to muons. In the red region the model explains at 2σ the anomalous magnetic moment of the muon [7, 115, 116]. The yellow band shows the region that can explain the Xenon1T excess [117] in some specifically extended $L_\mu - L_\tau$ model [118], while the black diamonds refer to a model [119] that explain the cosmic neutrino spectrum features observed by IceCube [120].

the observed anomalous magnetic moment of the muon [7] in the context of this model. In particular, they may be able to completely exclude the solution for $(g - 2)_\mu$ for $m_{Z'} < 20$ MeV.

Furthermore, it is interesting to note that in specifically extended models [118] there is a some overlap between the region preferred by the excess of low energy electrons observed by the Xenon1T experiment [117] (yellow band) and the $(g - 2)_\mu$ anomaly (red region). Again, both are contemplated within the range of the ESS detectors, which could probe combinations of $g_{Z'}$ and $m_{Z'}$ that simultaneously solve both. Finally, reference points of a specific model [119] explaining peculiar features observed in the cosmic neutrino spectrum by the IceCube collaboration [120] (also shown in Fig. 25) are also within reach, posing the possibility of tying the three anomalies at once with a single observation effort.

9 Conclusion

This work was first motivated through a need to search for new directions for Physics beyond the Standard Model. While it has proven time and again its consistency and sturdiness faced with experimental tests, the puzzle of the origin of neutrino masses has been unquestionably verified to be outside of its scope. Thus, we turned to neutrino-related phenomena as a natural choice to try to explore signs of New Physics.

Coherent Scattering also arises as a natural focus when considering the close relation to Dark Matter searches, whose properties also consist another well-established problem beyond the current Standard Model. The relatively small interaction rates of neutrinos and Dark Matter also push for the idea of any effect which may boost the cross section of related processes, connecting to the idea of $CE\nu NS$.

From that point, we derived the formalism to describe this type of interaction in the Standard Model, also noting how a convenient reparametrization of the problem in terms of relevant degrees of freedom could yield a characterization consistent with the fundamental description in terms of quark and neutrinos interacting via a neutral current process. Thus, we also introduced the notion of effective field theory.

The next logical step was to determine exactly which types of models could be efficiently probed by $CE\nu NS$ -dedicated experiments, relying on the description we built. $U(1)$ extensions to the Standard Model proved to be a simple, yet quite encompassing class which could both be obtained in a consistent way. They were also still generic enough such that one could match their realizations unto more complete theories in the UV description.

The selection of light Z' realizations inside this wide class was also well-motivated by the characteristic interaction and scales of $CE\nu NS$, and justified the introduction of a simplified model approach. We chose three models to study in detail, based on the surrounding interest in the literature: (i) a universal Z' model in which the light spin-1 particle couples to all the SM fermions with universal strength; (ii) a model in which the Z' couples to the anomaly free $B - L$ current; (iii) a model in which the Z' couples to the anomaly free $L_\mu - L_\tau$ current.

With the formalism for the full description of the process both in the Standard Model and the light Z' models of interest in place, the only remaining need was for a framework to connect predictions and experiments, and allow for the extraction of the relevant parameters of the theory. Thus, we invoked the tools of statistical inference, properly defining a systematic way to verify findings.

All these efforts laid the ground for the application to a real-world, well-studied

case, in order to confirm the validity of our approach. The COHERENT Liquid Argon detector was the obvious choice, given the very complete literature surrounding it, as well as the large collection of available data.

Once our methods were verified, we selected two different experimental proposals to establish their sensitivity with respect to the light Z' models of interest. We proceeded to show how to extract all the necessary quantities for the computation of event rates, and finally obtained the projected sensitivity, analysed in a comparative manner to other established and expected bounds. The whole work has been connected by a logical thread, returning to the initial need to probe what lies beyond the Standard Model.

Our results were presented in Figures 22 – 25. We have shown that the ν BDX-DRIFT proposal presents a weaker sensitivity to the unexplored parameter space in comparison to the ESS proposals. Yet, it could still prove useful in the context of the community as a whole, adding another valuable experimental point to measure the A and Z dependency of the $\text{CE}\nu\text{NS}$ cross section, as it is the only experiment to our knowledge with Sulfur as an active target.

Other dedicated studies of nuclear uncertainties related to the form factor using the unique high-energy flux from DUNE could also be incredibly useful, and are left as a possible follow up to this work. The possibility also remains open for upgrades of the design, introducing reliable track identification, to massively boost its capabilities for future studies.

In contrast, the ESS proposals have all generally demonstrated aptitude at exploring a larger portion of untested parameter space. This is particularly true for the $L_\mu - L_\tau$ model, where our results show that the detector exploiting Xe and CsI will be able to test the region $3 \lesssim m_{Z'}/\text{MeV} \lesssim 60$, where the Z' is able to explain the anomalous muon magnetic moment measurements [7, 115, 116] and the Xenon1T excess [117, 118]. Moreover, all the detectors at the ESS will be able to test a specific model [119] that can explain peculiar features observed in the cosmic neutrino spectrum by IceCube [120].

All the tools and formalism derived in this work are also readily applicable to other $\text{CE}\nu\text{NS}$ experiments, as well as Dark Matter direct detection efforts, which might be pursued in future work.

Bibliography

- [1] “The quantum theory of the emission and absorption of radiation”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 114.767 (Mar. 1927), pp. 243–265. DOI: [10.1098/rspa.1927.0039](https://doi.org/10.1098/rspa.1927.0039). URL: <https://doi.org/10.1098/rspa.1927.0039>.
- [2] R.P. Feynman and A. Zee. *QED: The Strange Theory of Light and Matter*. Alix G. Mautner memorial lectures. Princeton University Press, 2006. ISBN: 9780691125756. URL: <https://books.google.com.br/books?id=Uv-uxB0sRKEC>.
- [3] Serguei Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [hep-ex].
- [4] Georges Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214](https://arxiv.org/abs/1207.7214) [hep-ex].
- [5] Mu-Chun Chen. “TASI 2006 Lectures on Leptogenesis”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Exploring New Frontiers Using Colliders and Neutrinos*. Mar. 2007, pp. 123–176. arXiv: [hep-ph/0703087](https://arxiv.org/abs/hep-ph/0703087).
- [6] Anson Hook. “TASI Lectures on the Strong CP Problem and Axions”. In: *PoS TASI2018* (2019), p. 004. arXiv: [1812.02669](https://arxiv.org/abs/1812.02669) [hep-ph].
- [7] B. Abi et al. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm”. In: *Phys. Rev. Lett.* 126.14 (2021), p. 141801. DOI: [10.1103/PhysRevLett.126.141801](https://doi.org/10.1103/PhysRevLett.126.141801). arXiv: [2104.03281](https://arxiv.org/abs/2104.03281) [hep-ex].
- [8] Elena Graverini. “Flavour anomalies: a review”. In: *Journal of Physics: Conference Series* 1137 (Jan. 2019), p. 012025. ISSN: 1742-6596. DOI: [10.1088/1742-6596/1137/1/012025](https://doi.org/10.1088/1742-6596/1137/1/012025). URL: <http://dx.doi.org/10.1088/1742-6596/1137/1/012025>.
- [9] Roel Aaij et al. “Test of lepton universality in beauty-quark decays”. In: (Mar. 2021). arXiv: [2103.11769](https://arxiv.org/abs/2103.11769) [hep-ex].
- [10] Sz. Borsanyi et al. “Leading hadronic contribution to the muon magnetic moment from lattice QCD”. In: *Nature* 593.7857 (2021), pp. 51–55. DOI: [10.1038/s41586-021-03418-1](https://doi.org/10.1038/s41586-021-03418-1). arXiv: [2002.12347](https://arxiv.org/abs/2002.12347) [hep-lat].
- [11] Stephen P. Martin. “A Supersymmetry primer”. In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. DOI: [10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001). arXiv: [hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).

- [12] Keith A. Ulmer. “Supersymmetry: Experimental Status”. In: *3rd Large Hadron Collider Physics Conference*. 2015, pp. 180–191. arXiv: [1601.03774](https://arxiv.org/abs/1601.03774) [hep-ex].
- [13] F. Giordano. “SUSY searches at the LHC Run2”. In: *Nuovo Cim. C* 40.1 (2017), p. 2. DOI: [10.1393/ncc/i2017-17002-1](https://doi.org/10.1393/ncc/i2017-17002-1).
- [14] Takaaki Kajita. *Nobel Prize Lecture - Discovery of Atmospheric Neutrino Oscillations*. <https://nobelprize.org/uploads/2018/06/kajita-lecture.pdf>. Accessed: 2022-02-01. 2015.
- [15] Arthur B. McDonald. *Nobel Prize Lecture - The Sudbury Neutrino Observatory: Observation of Flavor Change for Solar Neutrinos*. <https://www.nobelprize.org/uploads/2018/06/mcdonald-lecture.pdf>. Accessed: 2022-02-01. 2015.
- [16] Leonard S. Kisslinger. “Review of neutrino oscillations with sterile and active neutrinos”. In: *International Journal of Modern Physics A* 31.23 (Aug. 2016), p. 1630037. ISSN: 1793-656X. DOI: [10.1142/S0217751X16300374](https://doi.org/10.1142/S0217751X16300374). URL: <http://dx.doi.org/10.1142/S0217751X16300374>.
- [17] Andre de Gouvea. “TASI lectures on neutrino physics”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics in $D \geq 4$* . Nov. 2004, pp. 197–258. arXiv: [hep-ph/0411274](https://arxiv.org/abs/hep-ph/0411274).
- [18] Fritz Zwicky. “On the Masses of Nebulae and of Clusters of Nebulae”. In: *The Astrophysical Journal* 86 (1937), p. 217.
- [19] V.C. Rubin, N. Thonnard, and Jr. Ford W.K. “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/”. In: *Astrophys. J.* 238 (1980), p. 471. DOI: [10.1086/158003](https://doi.org/10.1086/158003).
- [20] Maxim Markevitch et al. “Direct constraints on the dark matter self-interaction cross section from the merging galaxy cluster 1E 0657–56”. In: *The Astrophysical Journal* 606.2 (2004), p. 819.
- [21] Pablo Villanueva-Domingo, Olga Mena, and Sergio Palomares-Ruiz. “A Brief Review on Primordial Black Holes as Dark Matter”. In: *Frontiers in Astronomy and Space Sciences* 8 (May 2021). ISSN: 2296-987X. DOI: [10.3389/fspas.2021.681084](https://doi.org/10.3389/fspas.2021.681084). URL: <http://dx.doi.org/10.3389/fspas.2021.681084>.
- [22] KIM GRIEST. “The Search for the Dark Matter: WIMPs and MACHOs”. In: *Annals of the New York Academy of Sciences* 688.1 (June 1993), pp. 390–407. ISSN: 0077-8923. DOI: [10.1111/j.1749-6632.1993.tb43912.x](https://doi.org/10.1111/j.1749-6632.1993.tb43912.x). URL: <http://dx.doi.org/10.1111/j.1749-6632.1993.tb43912.x>.
- [23] R. H. Sanders. “A historical perspective on modified Newtonian dynamics”. In: *Can. J. Phys.* 93.2 (2015), pp. 126–138. DOI: [10.1139/cjp-2014-0206](https://doi.org/10.1139/cjp-2014-0206). arXiv: [1404.0531](https://arxiv.org/abs/1404.0531) [physics.hist-ph].

- [24] Albert Petrov. “Introduction to Modified Gravity”. In: *SpringerBriefs in Physics* (2020). ISSN: 2191-5431. DOI: [10.1007/978-3-030-52862-1](https://doi.org/10.1007/978-3-030-52862-1). URL: <http://dx.doi.org/10.1007/978-3-030-52862-1>.
- [25] Nabila Aghanim et al. “Planck 2018 results-VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (2020), A6.
- [26] Wayne Hu and Scott Dodelson. “Cosmic Microwave Background Anisotropies”. In: *Annual Review of Astronomy and Astrophysics* 40.1 (Sept. 2002), pp. 171–216. ISSN: 1545-4282. DOI: [10.1146/annurev.astro.40.060401.093926](https://doi.org/10.1146/annurev.astro.40.060401.093926). URL: <http://dx.doi.org/10.1146/annurev.astro.40.060401.093926>.
- [27] Joel R. Primack. *Dark Matter and Structure Formation in the Universe*. 1997. arXiv: [astro-ph/9707285](https://arxiv.org/abs/astro-ph/9707285) [[astro-ph](#)].
- [28] Tongyan Lin. *TASI lectures on dark matter models and direct detection*. 2019. arXiv: [1904.07915](https://arxiv.org/abs/1904.07915) [[hep-ph](#)].
- [29] Yonit Hochberg et al. “Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles”. In: *Phys. Rev. Lett.* 113 (2014), p. 171301. DOI: [10.1103/PhysRevLett.113.171301](https://doi.org/10.1103/PhysRevLett.113.171301). arXiv: [1402.5143](https://arxiv.org/abs/1402.5143) [[hep-ph](#)].
- [30] A. Arbey and F. Mahmoudi. “Dark matter and the early Universe: a review”. In: *Prog. Part. Nucl. Phys.* 119 (2021), p. 103865. DOI: [10.1016/j.pnpnp.2021.103865](https://doi.org/10.1016/j.pnpnp.2021.103865). arXiv: [2104.11488](https://arxiv.org/abs/2104.11488) [[hep-ph](#)].
- [31] E. Armengaud et al. “Searching for low-mass dark matter particles with a massive Ge bolometer operated above-ground”. In: *Phys. Rev. D* 99.8 (2019), p. 082003. DOI: [10.1103/PhysRevD.99.082003](https://doi.org/10.1103/PhysRevD.99.082003). arXiv: [1901.03588](https://arxiv.org/abs/1901.03588) [[astro-ph.GA](#)].
- [32] K. Abe et al. “A direct dark matter search in XMASS-I”. In: *Phys. Lett. B* 789 (2019), pp. 45–53. DOI: [10.1016/j.physletb.2018.10.070](https://doi.org/10.1016/j.physletb.2018.10.070). arXiv: [1804.02180](https://arxiv.org/abs/1804.02180) [[astro-ph.CO](#)].
- [33] D. Yu. Akimov et al. “WIMP-nucleon cross-section results from the second science run of ZEPLIN-III”. In: *Phys. Lett. B* 709 (2012), pp. 14–20. DOI: [10.1016/j.physletb.2012.01.064](https://doi.org/10.1016/j.physletb.2012.01.064). arXiv: [1110.4769](https://arxiv.org/abs/1110.4769) [[astro-ph.CO](#)].
- [34] E. Aprile et al. “Light Dark Matter Search with Ionization Signals in XENON1T”. In: *Phys. Rev. Lett.* 123.25 (2019), p. 251801. DOI: [10.1103/PhysRevLett.123.251801](https://doi.org/10.1103/PhysRevLett.123.251801). arXiv: [1907.11485](https://arxiv.org/abs/1907.11485) [[hep-ex](#)].
- [35] A. H. Abdelhameed et al. “First results from the CRESST-III low-mass dark matter program”. In: *Phys. Rev. D* 100.10 (2019), p. 102002. DOI: [10.1103/PhysRevD.100.102002](https://doi.org/10.1103/PhysRevD.100.102002). arXiv: [1904.00498](https://arxiv.org/abs/1904.00498) [[astro-ph.CO](#)].

- [36] Xiangyi Cui et al. “Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment”. In: *Phys. Rev. Lett.* 119.18 (2017), p. 181302. DOI: [10.1103/PhysRevLett.119.181302](https://doi.org/10.1103/PhysRevLett.119.181302). arXiv: [1708.06917](https://arxiv.org/abs/1708.06917) [astro-ph.CO].
- [37] I. Alkhatib et al. “Light Dark Matter Search with a High-Resolution Athermal Phonon Detector Operated Above Ground”. In: *Phys. Rev. Lett.* 127 (2021), p. 061801. DOI: [10.1103/PhysRevLett.127.061801](https://doi.org/10.1103/PhysRevLett.127.061801). arXiv: [2007.14289](https://arxiv.org/abs/2007.14289) [hep-ex].
- [38] F. Ruppin et al. “Complementarity of dark matter detectors in light of the neutrino background”. In: *Phys. Rev. D* 90.8 (2014), p. 083510. DOI: [10.1103/PhysRevD.90.083510](https://doi.org/10.1103/PhysRevD.90.083510). arXiv: [1408.3581](https://arxiv.org/abs/1408.3581) [hep-ph].
- [39] Daniel Z. Freedman, David N. Schramm, and David L. Tubbs. “The Weak Neutral Current and Its Effects in Stellar Collapse”. In: *Ann. Rev. Nucl. Part. Sci.* 27 (1977), pp. 167–207. DOI: [10.1146/annurev.ns.27.120177.001123](https://doi.org/10.1146/annurev.ns.27.120177.001123).
- [40] A. Liam Fitzpatrick et al. “The effective field theory of dark matter direct detection”. In: *Journal of Cosmology and Astroparticle Physics* 2013.02 (Feb. 2013), pp. 004–004. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2013/02/004](https://doi.org/10.1088/1475-7516/2013/02/004). URL: <http://dx.doi.org/10.1088/1475-7516/2013/02/004>.
- [41] Eugenio Del Nobile. “Appendiciario – A hands-on manual on the theory of direct Dark Matter detection”. In: (Apr. 2021). arXiv: [2104.12785](https://arxiv.org/abs/2104.12785) [hep-ph].
- [42] Martin Hoferichter, Javier Menéndez, and Achim Schwenk. “Coherent elastic neutrino-nucleus scattering: EFT analysis and nuclear responses”. In: *Physical Review D* 102.7 (Oct. 2020). ISSN: 2470-0029. DOI: [10.1103/physrevd.102.074018](https://doi.org/10.1103/physrevd.102.074018). URL: <http://dx.doi.org/10.1103/PhysRevD.102.074018>.
- [43] Richard H. Helm. “Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei”. In: *Phys. Rev.* 104 (5 Dec. 1956), pp. 1466–1475. DOI: [10.1103/PhysRev.104.1466](https://doi.org/10.1103/PhysRev.104.1466). URL: <https://link.aps.org/doi/10.1103/PhysRev.104.1466>.
- [44] M. Cadeddu et al. “Physics results from the first COHERENT observation of coherent elastic neutrino-nucleus scattering in argon and their combination with cesium-iodide data”. In: *Phys. Rev. D* 102.1 (2020), p. 015030. DOI: [10.1103/PhysRevD.102.015030](https://doi.org/10.1103/PhysRevD.102.015030). arXiv: [2005.01645](https://arxiv.org/abs/2005.01645) [hep-ph].
- [45] I. Angeli and K. P. Marinova. “Table of experimental nuclear ground state charge radii: An update”. In: *Atom. Data Nucl. Data Tabl.* 99.1 (2013), pp. 69–95. DOI: [10.1016/j.adt.2011.12.006](https://doi.org/10.1016/j.adt.2011.12.006).
- [46] Daniel Z. Freedman. “Coherent Neutrino Nucleus Scattering as a Probe of the Weak Neutral Current”. In: *Phys. Rev. D* 9 (1974), pp. 1389–1392. DOI: [10.1103/PhysRevD.9.1389](https://doi.org/10.1103/PhysRevD.9.1389).

- [47] D. Akimov et al. “Observation of Coherent Elastic Neutrino-Nucleus Scattering”. In: *Science* 357.6356 (2017), pp. 1123–1126. DOI: [10.1126/science.aao0990](https://doi.org/10.1126/science.aao0990). arXiv: [1708.01294](https://arxiv.org/abs/1708.01294) [nucl-ex].
- [48] D. Akimov et al. “COHERENT Collaboration data release from the first detection of coherent elastic neutrino-nucleus scattering on argon”. In: (June 2020). DOI: [10.5281/zenodo.3903810](https://doi.org/10.5281/zenodo.3903810). arXiv: [2006.12659](https://arxiv.org/abs/2006.12659) [nucl-ex].
- [49] M. Cadeddu et al. “Physics results from the first COHERENT observation of coherent elastic neutrino-nucleus scattering in argon and their combination with cesium-iodide data”. In: *Physical Review D* 102.1 (July 2020). ISSN: 2470-0029. DOI: [10.1103/physrevd.102.015030](https://doi.org/10.1103/physrevd.102.015030). URL: <http://dx.doi.org/10.1103/PhysRevD.102.015030>.
- [50] D. Adhikari et al. “Accurate Determination of the Neutron Skin Thickness of ^{208}Pb through Parity-Violation in Electron Scattering”. In: *Phys. Rev. Lett.* 126.17 (2021), p. 172502. DOI: [10.1103/PhysRevLett.126.172502](https://doi.org/10.1103/PhysRevLett.126.172502). arXiv: [2102.10767](https://arxiv.org/abs/2102.10767) [nucl-ex].
- [51] Louis E. Strigari. “Neutrino Coherent Scattering Rates at Direct Dark Matter Detectors”. In: *New J. Phys.* 11 (2009), p. 105011. DOI: [10.1088/1367-2630/11/10/105011](https://doi.org/10.1088/1367-2630/11/10/105011). arXiv: [0903.3630](https://arxiv.org/abs/0903.3630) [astro-ph.CO].
- [52] Roni Harnik, Joachim Kopp, and Pedro A. N. Machado. “Exploring ν Signals in Dark Matter Detectors”. In: *JCAP* 07 (2012), p. 026. DOI: [10.1088/1475-7516/2012/07/026](https://doi.org/10.1088/1475-7516/2012/07/026). arXiv: [1202.6073](https://arxiv.org/abs/1202.6073) [hep-ph].
- [53] Enrico Bertuzzo et al. “Dark Matter and Exotic Neutrino Interactions in Direct Detection Searches”. In: *JHEP* 04 (2017), p. 073. DOI: [10.1007/JHEP04\(2017\)073](https://doi.org/10.1007/JHEP04(2017)073). arXiv: [1701.07443](https://arxiv.org/abs/1701.07443) [hep-ph].
- [54] M. C. Gonzalez-Garcia et al. “Neutrino Discovery Limit of Dark Matter Direct Detection Experiments in the Presence of Non-Standard Interactions”. In: *JHEP* 07 (2018), p. 019. DOI: [10.1007/JHEP07\(2018\)019](https://doi.org/10.1007/JHEP07(2018)019). arXiv: [1803.03650](https://arxiv.org/abs/1803.03650) [hep-ph].
- [55] C. Boehm et al. “How high is the neutrino floor?” In: *JCAP* 01 (2019), p. 043. DOI: [10.1088/1475-7516/2019/01/043](https://doi.org/10.1088/1475-7516/2019/01/043). arXiv: [1809.06385](https://arxiv.org/abs/1809.06385) [hep-ph].
- [56] M. Cadeddu et al. “Constraints on light vector mediators through coherent elastic neutrino nucleus scattering data from COHERENT”. In: *JHEP* 01 (2021), p. 116. DOI: [10.1007/JHEP01\(2021\)116](https://doi.org/10.1007/JHEP01(2021)116). arXiv: [2008.05022](https://arxiv.org/abs/2008.05022) [hep-ph].
- [57] Riccardo Penco. “An Introduction to Effective Field Theories”. In: (June 2020). arXiv: [2006.16285](https://arxiv.org/abs/2006.16285) [hep-th].

- [58] Davi B. Costa, Bogdan A. Dobrescu, and Patrick J. Fox. “General Solution to the U(1) Anomaly Equations”. In: *Physical Review Letters* 123.15 (Oct. 2019). ISSN: 1079-7114. DOI: [10.1103/physrevlett.123.151601](https://doi.org/10.1103/physrevlett.123.151601). URL: <http://dx.doi.org/10.1103/PhysRevLett.123.151601>.
- [59] Davi B. Costa. “Anomaly-free U(1)_m extensions of the Standard Model”. In: *Physical Review D* 102.11 (Dec. 2020). ISSN: 2470-0029. DOI: [10.1103/physrevd.102.115006](https://doi.org/10.1103/physrevd.102.115006). URL: <http://dx.doi.org/10.1103/PhysRevD.102.115006>.
- [60] *Neutrino Non-Standard Interactions: A Status Report*. Vol. 2. 2019, p. 001. DOI: [10.21468/SciPostPhysProc.2.001](https://doi.org/10.21468/SciPostPhysProc.2.001). arXiv: [1907.00991](https://arxiv.org/abs/1907.00991) [hep-ph].
- [61] M. Atzori Corona et al. “Probing light mediators and $(g - 2)_\mu$ through detection of coherent elastic neutrino nucleus scattering at COHERENT”. In: (Feb. 2022). arXiv: [2202.11002](https://arxiv.org/abs/2202.11002) [hep-ph].
- [62] D. Baxter et al. “Coherent Elastic Neutrino-Nucleus Scattering at the European Spallation Source”. In: *JHEP* 02 (2020), p. 123. DOI: [10.1007/JHEP02\(2020\)123](https://doi.org/10.1007/JHEP02(2020)123). arXiv: [1911.00762](https://arxiv.org/abs/1911.00762) [physics.ins-det].
- [63] Kyle Cranmer. *Practical Statistics for the LHC*. 2015. arXiv: [1503.07622](https://arxiv.org/abs/1503.07622).
- [64] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *Annals Math. Statist.* 9.1 (1938), pp. 60–62. DOI: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360).
- [65] D. Akimov et al. “Measurement of the Coherent Elastic Neutrino-Nucleus Scattering Cross Section on CsI by COHERENT”. In: (Oct. 2021). arXiv: [2110.07730](https://arxiv.org/abs/2110.07730) [hep-ex].
- [66] D. Akimov et al. *COHERENT Collaboration data release from the first detection of coherent elastic neutrino-nucleus scattering on argon*. Version 1.0. Zenodo, June 2020. DOI: [10.5281/zenodo.3903810](https://doi.org/10.5281/zenodo.3903810). URL: <https://doi.org/10.5281/zenodo.3903810>.
- [67] Ben Jones. *Introduction to Scintillation Light in Liquid Argon*. Presented at the Liquid Argon TPC R&D Workshop (LArTPC14). July 2014. URL: <https://indico.fnal.gov/event/8381/contributions/107078/attachments/69965/83895/LArTPCWorkshopScintLight.pdf>.
- [68] C Amsler et al. “Luminescence quenching of the triplet excimer state by air traces in gaseous argon”. In: *Journal of Instrumentation* 3.02 (Feb. 2008), P02001–P02001. ISSN: 1748-0221. DOI: [10.1088/1748-0221/3/02/p02001](https://doi.org/10.1088/1748-0221/3/02/p02001). URL: <http://dx.doi.org/10.1088/1748-0221/3/02/P02001>.

- [69] Jacob Zetlemoyer. *Results of a CEvNS Search with the CENNS-10 Liquid Argon Detector*. Presented at the 2019 Magnificent CEvNS Workshop. Nov. 2019. URL: <https://indico.cern.ch/event/844613/contributions/3607503/attachments/1942028/3220533/JCZCENNS10M7s2019EngRunFinalv2.pdf>.
- [70] D. Aristizabal Sierra et al. “Coherent elastic neutrino-nucleus scattering with the ν BDX-DRIFT directional detector at next generation neutrino facilities”. In: *Phys. Rev. D* 104.3 (2021), p. 033004. DOI: [10.1103/PhysRevD.104.033004](https://doi.org/10.1103/PhysRevD.104.033004). arXiv: [2103.10857](https://arxiv.org/abs/2103.10857) [hep-ph].
- [71] D. P. Snowden-Ifft, C. J. Martoff, and J. M. Burwell. “Low pressure negative ion time projection chamber for dark matter search”. In: *Phys. Rev. D* 61 (10 Apr. 2000), p. 101301. DOI: [10.1103/PhysRevD.61.101301](https://doi.org/10.1103/PhysRevD.61.101301). URL: <https://link.aps.org/doi/10.1103/PhysRevD.61.101301>.
- [72] Babak Abi et al. “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics”. In: (Feb. 2020). arXiv: [2002.03005](https://arxiv.org/abs/2002.03005) [hep-ex].
- [73] D. P. Snowden-Ifft et al. “Directional light-WIMP time-projection-chamber detector for electron beam-dump experiments”. In: *Physical Review D* 99.6 (Mar. 2019). ISSN: 2470-0029. DOI: [10.1103/physrevd.99.061301](https://doi.org/10.1103/physrevd.99.061301). URL: <http://dx.doi.org/10.1103/PhysRevD.99.061301>.
- [74] D. P. Snowden-Ifft and J.-L. Gauvreau. “High precision measurements of carbon disulfide negative ion mobility and diffusion”. In: *Review of Scientific Instruments* 84.5 (May 2013), p. 053304. ISSN: 1089-7623. DOI: [10.1063/1.4803004](https://doi.org/10.1063/1.4803004). URL: <http://dx.doi.org/10.1063/1.4803004>.
- [75] J. B. R. Battat et al. “First background-free limit from a directional dark matter experiment: results from a fully fiducialised DRIFT detector”. In: *Phys. Dark Univ.* 9-10 (2015), pp. 1–7. DOI: [10.1016/j.dark.2015.06.001](https://doi.org/10.1016/j.dark.2015.06.001). arXiv: [1410.7821](https://arxiv.org/abs/1410.7821) [hep-ex].
- [76] J.B.R. Battat et al. “Low threshold results and limits from the DRIFT directional dark matter detector”. In: *Astroparticle Physics* 91 (May 2017), pp. 65–74. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2017.03.007](https://doi.org/10.1016/j.astropartphys.2017.03.007). URL: <http://dx.doi.org/10.1016/j.astropartphys.2017.03.007>.
- [77] Ivan Esteban. *European Spallation Source: a future for Coherent Neutrino Nucleus Scattering*. 2021. arXiv: [2105.04669](https://arxiv.org/abs/2105.04669) [hep-ph].
- [78] R. L. Burman and P. Plischke. “Neutrino fluxes from a high-intensity spallation neutron facility”. In: *Nucl. Instrum. Meth. A* 398 (1997), pp. 147–156. DOI: [10.1016/S0168-9002\(97\)00821-8](https://doi.org/10.1016/S0168-9002(97)00821-8).

- [79] Alexis Aguilar-Arevalo et al. “Search for light mediators in the low-energy data of the CONNIE reactor neutrino experiment”. In: *JHEP* 04 (2020), p. 054. DOI: [10.1007/JHEP04\(2020\)054](https://doi.org/10.1007/JHEP04(2020)054). arXiv: [1910.04951](https://arxiv.org/abs/1910.04951) [hep-ex].
- [80] Liron Barak et al. “SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper CCD”. In: *Physical Review Letters* 125.17 (Oct. 2020). ISSN: 1079-7114. DOI: [10.1103/physrevlett.125.171802](https://doi.org/10.1103/physrevlett.125.171802). URL: <http://dx.doi.org/10.1103/PhysRevLett.125.171802>.
- [81] A. Aguilar-Arevalo et al. “Search for low-mass WIMPs in a 0.6 kg day exposure of the DAMIC experiment at SNOLAB”. In: *Phys. Rev. D* 94.8 (2016), p. 082006. DOI: [10.1103/PhysRevD.94.082006](https://doi.org/10.1103/PhysRevD.94.082006). arXiv: [1607.07410](https://arxiv.org/abs/1607.07410) [astro-ph.CO].
- [82] J. J. Gomez-Cadenas. “Status and prospects of the NEXT experiment for neutrinoless double beta decay searches”. In: *54th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. 2019, pp. 201–206. arXiv: [1906.01743](https://arxiv.org/abs/1906.01743) [hep-ex].
- [83] Frank Edzards et al. “Surface Characterization of P-Type Point Contact Germanium Detectors”. In: *Particles* 4.4 (Oct. 2021), pp. 489–511. ISSN: 2571-712X. DOI: [10.3390/particles4040036](https://doi.org/10.3390/particles4040036). URL: <http://dx.doi.org/10.3390/particles4040036>.
- [84] C. Amole et al. “Dark Matter Search Results from the Complete Exposure of the PICO-60 C₃F₈ Bubble Chamber”. In: *Phys. Rev. D* 100.2 (2019), p. 022001. DOI: [10.1103/PhysRevD.100.022001](https://doi.org/10.1103/PhysRevD.100.022001). arXiv: [1902.04031](https://arxiv.org/abs/1902.04031) [astro-ph.CO].
- [85] Peter B. Denton and Julia Gehrlein. “A statistical analysis of the COHERENT data and applications to new physics”. In: *Journal of High Energy Physics* 2021.4 (Apr. 2021). ISSN: 1029-8479. DOI: [10.1007/jhep04\(2021\)266](https://doi.org/10.1007/jhep04(2021)266). URL: [http://dx.doi.org/10.1007/JHEP04\(2021\)266](http://dx.doi.org/10.1007/JHEP04(2021)266).
- [86] G. Fernandez-Moroni et al. “The physics potential of a reactor neutrino experiment with Skipper-CCDs: Searching for new physics with light mediators”. In: (Aug. 2021). arXiv: [2108.07310](https://arxiv.org/abs/2108.07310) [hep-ph].
- [87] J. P. Lees et al. “Search for Invisible Decays of a Dark Photon Produced in e^+e^- Collisions at BaBar”. In: *Phys. Rev. Lett.* 119.13 (2017), p. 131804. DOI: [10.1103/PhysRevLett.119.131804](https://doi.org/10.1103/PhysRevLett.119.131804). arXiv: [1702.03327](https://arxiv.org/abs/1702.03327) [hep-ex].
- [88] D. Banerjee et al. “Search for vector mediator of Dark Matter production in invisible decay mode”. In: *Phys. Rev. D* 97.7 (2018), p. 072002. DOI: [10.1103/PhysRevD.97.072002](https://doi.org/10.1103/PhysRevD.97.072002). arXiv: [1710.00971](https://arxiv.org/abs/1710.00971) [hep-ex].
- [89] D. Banerjee et al. “Dark matter search in missing energy events with NA64”. In: *Phys. Rev. Lett.* 123.12 (2019), p. 121801. DOI: [10.1103/PhysRevLett.123.121801](https://doi.org/10.1103/PhysRevLett.123.121801). arXiv: [1906.00176](https://arxiv.org/abs/1906.00176) [hep-ex].

- [90] *The vIOLETA collaboration website*. [Link](https://www.violetaexperiment.com/). URL: <https://www.violetaexperiment.com/>.
- [91] Dario Rodrigues et al. *vIOLETA: Neutrino Interaction Observation with a Low Energy Threshold Array*. [Link](https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-521.pdf). July 2020. URL: <https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-521.pdf>.
- [92] Emiliano Pozzi et al. *Short baseline neutrino program in Argentina*. [Link](https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-523.pdf). July 2020. URL: <https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-523.pdf>.
- [93] Ivan Jesus Martinez-Soler et al. *A first study of the physics potential of a reactor neutrino experiment with Skipper-CCDs*. [Link](https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-508.pdf). July 2020. URL: <https://nusoft.fnal.gov/nova/nu2020postersession/pdf/posterPDF-508.pdf>.
- [94] Guillermo Fernandez-Moroni et al. “The physics potential of a reactor neutrino experiment with Skipper CCDs: Measuring the weak mixing angle”. In: *JHEP* 03 (2021), p. 186. DOI: [10.1007/JHEP03\(2021\)186](https://doi.org/10.1007/JHEP03(2021)186). arXiv: [2009.10741](https://arxiv.org/abs/2009.10741) [hep-ph].
- [95] D. Banerjee et al. “Improved limits on a hypothetical X(16.7) boson and a dark photon decaying into e^+e^- pairs”. In: *Phys. Rev. D* 101.7 (2020), p. 071101. DOI: [10.1103/PhysRevD.101.071101](https://doi.org/10.1103/PhysRevD.101.071101). arXiv: [1912.11389](https://arxiv.org/abs/1912.11389) [hep-ex].
- [96] E. M. Riordan et al. “Search for short-lived axions in an electron-beam-dump experiment”. In: *Phys. Rev. Lett.* 59 (7 Aug. 1987), pp. 755–758. DOI: [10.1103/PhysRevLett.59.755](https://doi.org/10.1103/PhysRevLett.59.755). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.59.755>.
- [97] M. Davier and H. Nguyen Ngoc. “An Unambiguous Search for a Light Higgs Boson”. In: *Phys. Lett. B* 229 (1989), pp. 150–155. DOI: [10.1016/0370-2693\(89\)90174-3](https://doi.org/10.1016/0370-2693(89)90174-3).
- [98] Johannes Blumlein and Jurgen Brunner. “New Exclusion Limits for Dark Gauge Forces from Beam-Dump Data”. In: *Phys. Lett. B* 701 (2011), pp. 155–159. DOI: [10.1016/j.physletb.2011.05.046](https://doi.org/10.1016/j.physletb.2011.05.046). arXiv: [1104.2747](https://arxiv.org/abs/1104.2747) [hep-ex].
- [99] J. R. Batley et al. “Search for the dark photon in π^0 decays”. In: *Phys. Lett. B* 746 (2015), pp. 178–185. DOI: [10.1016/j.physletb.2015.04.068](https://doi.org/10.1016/j.physletb.2015.04.068). arXiv: [1504.00607](https://arxiv.org/abs/1504.00607) [hep-ex].
- [100] F. Archilli et al. “Search for a vector gauge boson in ϕ meson decays with the KLOE detector”. In: *Phys. Lett. B* 706 (2012), pp. 251–255. DOI: [10.1016/j.physletb.2011.11.033](https://doi.org/10.1016/j.physletb.2011.11.033). arXiv: [1110.0411](https://arxiv.org/abs/1110.0411) [hep-ex].
- [101] D. Babusci et al. “Limit on the production of a light vector gauge boson in phi meson decays with the KLOE detector”. In: *Phys. Lett. B* 720 (2013), pp. 111–115. DOI: [10.1016/j.physletb.2013.01.067](https://doi.org/10.1016/j.physletb.2013.01.067). arXiv: [1210.3927](https://arxiv.org/abs/1210.3927) [hep-ex].

- [102] J. P. Lees et al. “Search for a Dark Photon in e^+e^- Collisions at BaBar”. In: *Phys. Rev. Lett.* 113.20 (2014), p. 201801. DOI: [10.1103/PhysRevLett.113.201801](https://doi.org/10.1103/PhysRevLett.113.201801). arXiv: [1406.2980](https://arxiv.org/abs/1406.2980) [hep-ex].
- [103] Roel Aaij et al. “Search for $A' \rightarrow \mu^+\mu^-$ Decays”. In: *Phys. Rev. Lett.* 124.4 (2020), p. 041801. DOI: [10.1103/PhysRevLett.124.041801](https://doi.org/10.1103/PhysRevLett.124.041801). arXiv: [1910.06926](https://arxiv.org/abs/1910.06926) [hep-ex].
- [104] Philip Ilten et al. “Serendipity in dark photon searches”. In: *JHEP* 06 (2018), p. 004. DOI: [10.1007/JHEP06\(2018\)004](https://doi.org/10.1007/JHEP06(2018)004). arXiv: [1801.04847](https://arxiv.org/abs/1801.04847) [hep-ph].
- [105] Kaustav Chakraborty et al. “Constraining general U(1) interactions from neutrino-electron scattering measurements at DUNE near detector”. In: (Nov. 2021). arXiv: [2111.08767](https://arxiv.org/abs/2111.08767) [hep-ph].
- [106] Patrick J. Fox et al. “LEP Shines Light on Dark Matter”. In: *Phys. Rev. D* 84 (2011), p. 014028. DOI: [10.1103/PhysRevD.84.014028](https://doi.org/10.1103/PhysRevD.84.014028). arXiv: [1103.0240](https://arxiv.org/abs/1103.0240) [hep-ph].
- [107] J. P. Lees et al. “Search for a muonic dark force at BABAR”. In: *Phys. Rev. D* 94.1 (2016), p. 011102. DOI: [10.1103/PhysRevD.94.011102](https://doi.org/10.1103/PhysRevD.94.011102). arXiv: [1606.03501](https://arxiv.org/abs/1606.03501) [hep-ex].
- [108] Wolfgang Altmannshofer et al. “Explaining dark matter and B decay anomalies with an $L_\mu - L_\tau$ model”. In: *JHEP* 12 (2016), p. 106. DOI: [10.1007/JHEP12\(2016\)106](https://doi.org/10.1007/JHEP12(2016)106). arXiv: [1609.04026](https://arxiv.org/abs/1609.04026) [hep-ph].
- [109] Georges Aad et al. “Measurements of Four-Lepton Production at the Z Resonance in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with ATLAS”. In: *Phys. Rev. Lett.* 112.23 (2014), p. 231806. DOI: [10.1103/PhysRevLett.112.231806](https://doi.org/10.1103/PhysRevLett.112.231806). arXiv: [1403.5657](https://arxiv.org/abs/1403.5657) [hep-ex].
- [110] Albert M Sirunyan et al. “Search for an $L_\mu - L_\tau$ gauge boson using $Z \rightarrow 4\mu$ events in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Lett. B* 792 (2019), pp. 345–368. DOI: [10.1016/j.physletb.2019.01.072](https://doi.org/10.1016/j.physletb.2019.01.072). arXiv: [1808.03684](https://arxiv.org/abs/1808.03684) [hep-ex].
- [111] Wolfgang Altmannshofer et al. “Neutrino Trident Production: A Powerful Probe of New Physics with Neutrino Beams”. In: *Phys. Rev. Lett.* 113 (2014), p. 091801. DOI: [10.1103/PhysRevLett.113.091801](https://doi.org/10.1103/PhysRevLett.113.091801). arXiv: [1406.2332](https://arxiv.org/abs/1406.2332) [hep-ph].
- [112] S. R. Mishra et al. “Neutrino tridents and W-Z interference”. In: *Phys. Rev. Lett.* 66 (24 June 1991), pp. 3117–3120. DOI: [10.1103/PhysRevLett.66.3117](https://doi.org/10.1103/PhysRevLett.66.3117). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.66.3117>.
- [113] G. Bellini et al. “Precision measurement of the ^7Be solar neutrino interaction rate in Borexino”. In: *Phys. Rev. Lett.* 107 (2011), p. 141302. DOI: [10.1103/PhysRevLett.107.141302](https://doi.org/10.1103/PhysRevLett.107.141302). arXiv: [1104.1816](https://arxiv.org/abs/1104.1816) [hep-ex].

- [114] Sergei Gninenko and Dmitry Gorbunov. “Refining constraints from Borexino measurements on a light Z' -boson coupled to $L\mu$ - $L\tau$ current”. In: *Phys. Lett. B* 823 (2021), p. 136739. DOI: [10.1016/j.physletb.2021.136739](https://doi.org/10.1016/j.physletb.2021.136739). arXiv: [2007.16098](https://arxiv.org/abs/2007.16098) [hep-ph].
- [115] G. W. Bennett et al. “Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL”. In: *Phys. Rev. D* 73 (2006), p. 072003. DOI: [10.1103/PhysRevD.73.072003](https://doi.org/10.1103/PhysRevD.73.072003). arXiv: [hep-ex/0602035](https://arxiv.org/abs/hep-ex/0602035).
- [116] T. Aoyama et al. “The anomalous magnetic moment of the muon in the Standard Model”. In: *Phys. Rept.* 887 (2020), pp. 1–166. DOI: [10.1016/j.physrep.2020.07.006](https://doi.org/10.1016/j.physrep.2020.07.006). arXiv: [2006.04822](https://arxiv.org/abs/2006.04822) [hep-ph].
- [117] E. Aprile et al. “Excess electronic recoil events in XENON1T”. In: *Phys. Rev. D* 102.7 (2020), p. 072004. DOI: [10.1103/PhysRevD.102.072004](https://doi.org/10.1103/PhysRevD.102.072004). arXiv: [2006.09721](https://arxiv.org/abs/2006.09721) [hep-ex].
- [118] Debasish Borah et al. “Muon ($g - 2$) and XENON1T excess with boosted dark matter in $L\mu - L\tau$ model”. In: *Phys. Lett. B* 820 (2021), p. 136577. DOI: [10.1016/j.physletb.2021.136577](https://doi.org/10.1016/j.physletb.2021.136577). arXiv: [2104.05656](https://arxiv.org/abs/2104.05656) [hep-ph].
- [119] Takeshi Araki et al. “MeV scale leptonic force for cosmic neutrino spectrum and muon anomalous magnetic moment”. In: *Phys. Rev. D* 93.1 (2016), p. 013014. DOI: [10.1103/PhysRevD.93.013014](https://doi.org/10.1103/PhysRevD.93.013014). arXiv: [1508.07471](https://arxiv.org/abs/1508.07471) [hep-ph].
- [120] M. G. Aartsen et al. “A combined maximum-likelihood analysis of the high-energy astrophysical neutrino flux measured with IceCube”. In: *Astrophys. J.* 809.1 (2015), p. 98. DOI: [10.1088/0004-637X/809/1/98](https://doi.org/10.1088/0004-637X/809/1/98). arXiv: [1507.03991](https://arxiv.org/abs/1507.03991) [astro-ph.HE].