SEARCHING FOR NEW PHYSICS WITH FERMILAB'S PROTON BEAM

BY

DIANA FORBES

DISSERTATION

Submitted in partial fulfillment of the requirements
or the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2025

Urbana, Illinois

Doctoral Committee:

Associate Professor Jessie Shelton, Chair
Assistant Professor Yonatan Kahn, Director of Research
Professor George Gollin
Assistant Professor Helvi Witek

ABSTRACT

In this thesis, I summarize my work on two projects that use Fermilab's proton beam to search for new physics. The first project was in response to the anomalous muon magnetic moment announcement in 2021 [1]. This project proposed that we can use the SpinQuest experiment [2], with no additional installations needed, to look for a muonphillic scalar particle S whose coupling to muons could resolve the anomalous muon magnetic moment. In this setup, a $\sim$100 GeV proton beam would travel through a thick target material to produce a muon beam through scattering off of the target nuclei. This muon beam would then traverse the last length of the target, estimated to be $\sim$ 100 cm of material, to radiatively produce hypothetical scalar particles that would promptly decay into a muon pair. Hence, the signature would be an event with a pair of muons that has an invariant mass equal to the mass of the scalar particle. We simulated events for a choice of 3×10^{14} muons on target (MOT) with typical energies of $\sim$ 20 GeV, and, with a 15% invariant mass resolution, this strategy can probe the entire parameter space for which $\sim$ 200 MeV – GeV scalar particles resolve the muon $g-2$ anomaly. Ultimately, we estimated that we would only need about 6 years of beam time to achieve a 3σ discovery sensitivity.

For the second project, I collaborated with a group member named Rachel Nguyen to estimate the sensitivity of Fermilab's DUNE Near Detector site to detecting millicharged particles (MCPs) as well as heavy axion-like particles (ALPs). The majority of my work for this project focused on the MCP search. Regarding the MCP search, we utilized the setup of a proposed project named FerMINI [3] to look for a MCP would be generated from charged pion scattering- a production channel that has been so far overlooked. We would use the DUNE Near Detector complex at Fermilab, with the additional installation of scintillator arrays and photomultiplier tubes (PMTs), to detect MCPs through soft ionizations. Fermilab's proton beam would pass through a graphite target, producing charged pions that secondarily scatter in the target to produce a $\chi\bar{\chi}$ pair from a virtual photon. The signature for detecting a MCP would be three soft ionizations through the scintillator arrays that are collinear with the beam line and target, detectable through the collection of photoelectrons by PMTs. Using the framework of chiral perturbation theory, we found a small region of parameter space, around $1.5 < m_\chi < 2.8$ GeV, where the charged pion production channel dominates over previously-considered production mechanisms.

Regarding the ALP search, we proposed a search strategy that would use the DUNE Near

Detector complex and charged pions to detect heavy axion-like particles (ALPs) with low-energy couplings to gluons. Again, using the framework of chiral perturbation theory, we demonstrated regimes of parameter space where the charged pion production channel dominates over previously-considered production mechanisms for ALPs, thereby improving the sensitivity of DUNE to these new particles compared to previous studies.

ACKNOWLEDGEMENTS

During my time completing my PhD in physics at UIUC, I faced more challenges than I predicted. I was fortunate enough to have many people that made up my support network to get me through it, and I learned a lot about myself, academia, physics, and human nature in the process.

First, I'd like to thank my dad for many things. I appreciate the way that he raised me to be an independent thinker, passionate about seeking truth, kind to others, curious, analytical, and patient. I learn a lot from listening to him tell stories about how he grew up, what he's doing at work as a chemist, what his time in grad school was like, and challenges he used to face and that he still must overcome. I appreciate the short and efficient "dad texts" through the years to check in with me to make sure my finances and important stressors are in order. He speaks very little but he still has a way of inviting me into his inner world enough to tell me the important things I need to hear about. Without him, I would be disorganized, behind on taxes and bills, and I likely would never have had the resources to get a higher education.

I'd like to thank my mom for her continual compassion, support, and worry. She sends me food when I'm sick, mails me gifts and cards when I'm sad or she's feeling grateful, she makes sure I have what I need, she drives countless hours to me if she even suspects I'm in crisis, and she's always eager to hear from me on the phone.

I'd like to thank my friend and research group mate, Rachel Nguyen, for effectively taking me under her wing and giving me the tools and knowledge I needed to get through the program. She offered indispensable practical tricks to help with getting software to run and note-taking. She spent a great amount of time writing up and sending me personal research notes so that I could keep up with the material, clear up misunderstandings, and learn things that I wouldn't have even thought to ask about.

I'd like to thank my friends at this campus for spending time with me, comforting me when I was feeling anxious, taking me to birthday dinners, watching movies with me, allowing me to host dinners and pick what games we play, making art and music with me, and just generally being responsive to my needs, opinions, and experience of life. They introduced me to important and meaningful parts of the city and campus, they shared with me powerful life experiences and memories that gave me a wider understanding of human experience, and helped me develop more resilience.

I'd like to thank my research advisor, Yoni Kahn, for accepting me into his research group, for his practical advice, guidance with getting projects started, his patience, his kindness and generosity, and his countless hours helping me learn or getting me through difficult calculations that were important to my work. He also taught me valuable skills for sanity-checking my work, communicating with experimentalists, and making quick back-of-the-envelope estimates.

I'd like to thank Patrick Draper for talking about his experiences in academia, his honesty and openness in conversations, and his indispensable insights about managing anxiety. I'd also like to thank George Gollin for entertaining my independent electronics projects, being generous with his time and resources, showing kindness, and being a source of goodwill and happiness in the department. I'd like to think the additional members of my thesis committee, Helvi Witek and Jessie Shelton, for their goodwill, supporting my progress, and giving helpful feedback.

I'd like to thank the people in my research group who I leaned on most: Ian Harris, Michael Wentzel, Rachel Nguyen, and Hannah Day. They treated me with kindness, helped me learn about concepts that I had trouble with, and gave me a lot of helpful feedback that made me a stronger presenter and scientific communicator.

I'd like to thank one of my closest friends, Emma Weaver, for being a significant part of my life, during which she taught me to be more patient, forgiving, and accepting of peoples' unique experiences of life. She offers practical support and help with debugging errors that get me stuck for hours. She taught me a lot of coding best-practices that made me a more efficient and organized programmer. She is generous with her resources and offers a second home to come back to when I need it. But most of all, she's taught me how to not over-commit to obligations, to slow down, and to not avoid putting time into hobbies and relaxation when it's needed.

I'd like to thank my dad's brother, Jim Forbes, and his wife, Mary Ellen Forbes, for taking interest in my work, supporting me, and giving me helpful feedback during my time in grad school.

Finally, I'd like to acknowledge and thank my senior cat, Cinnamon, who I moved in with me from my childhood home during my first year at UIUC. She is the kindest and most expressive animal I have ever interacted with. She taught me to slow down, be present, and appreciate everything life has to offer. She is a daily reminder of the importance of taking care of things that are important to you, and how rewarding that can be. All of the components of my support network made me realize just how much we need each other, and how impossible it is to get through life and hardship alone.

Contents

CHAPTER 1: INTRODUCTION

The task of understanding how the world works and what it is made of is inspiring and arduous. Luckily, physicists around the world for generations have been working towards a coherent summary that describes just this. Called the Standard Model (SM), they have devised a theory that describes almost every type of particle interaction we can observe or imagine. The Standard Model describes that particles fall into one of two categories: fermions or bosons, distinguished by whether they have a spin of odd half integers (e.g. $\frac{1}{2}, \frac{3}{2}, \frac{5}{2}$, ...) or a spin of even half integers (e.g. $\frac{2}{2}, \frac{4}{2}, \frac{6}{2}$, ...). We think of fermions as what matter is made of, and bosons are what govern their motion. That is to say, all particles interact with each other or with themselves through the fundamental forces: electromagnetism, the weak force, the strong force, and the gravitational force. For all of these forces besides gravity, we have discovered, observed, and cataloged information about the bosons that mediate said force. Hence, the Standard Model includes photons that mediate electromagnetic interactions, $W^{\pm}$ and Z^0 bosons that mediate the weak force, and 8 different gluons that mediate the strong force [4].

Fermions can be thought of as belonging to two subcategories: quarks or leptons. There are 6 types of quarks and they join together under the strong force to make up larger composite particles called baryons (made of 3 quarks) and mesons (made of a quark and anti-quark). Examples of baryons include protons and neutrons, while an example of a meson is a pion. Leptons are fundamental particles that do not join together to make any larger structure object (excluding bound states like molecules/atoms). There are 6 known leptons: the electron, muon, tau particle, electron neutrino, muon neutrino, and the tau neutrino. According to the Standard Model, every particle is predicted to have a corresponding anti-particle of the same mass but opposite charge. All anti-particles have been observed such as anti-up, anti-down, anti-strange, anti-charge, anti-top, and anti-bottom quarks, in addition to the positron (anti-electron) and more.

Diving deeper, we note that more complex descriptions of how particles and their corresponding fields interact is described through quantum field theory; the foundational framework upon which the Standard Model is built. This machinery was able to predict the existence of the Higgs particle – a particle whose field allows matter to have mass – but it could not predict the Higgs particle's mass. A flurry of searches occurred between 1984 and 2010 to no avail. In 2011, experimentalists saw some possible results that indicated there may be a Higgs particle within the mass range of

120 - 145 GeV [5]. In 2012, data from 2011 and 2012 was combined to reveal that we find a statistically significant abundance of particles that the Higgs decays into, indicating that the Higgs particle exists and has a mass of about 126.5 GeV.

Although the SM is robust, scientific observations point towards the fact that it is still incomplete. For example, astronomers have observed several phenomena that tell us that dark matter (DM) – matter that does not emit light but causes gravitational effects –must exist. The earliest evidence for DM came from the observation that various luminous objects like gas clouds, stars, and galaxies, move faster than predicted by Newton's law of gravity that said they should move at speeds that drop with distance from the center of the galaxy r. Instead, astronomers measured that the luminous objects were moving at about a constant speeds as distance from the center of the galaxy increased for large r [6]. Hence, for this observation to make sense, galaxies with these galactic rotation curves must have much more matter in them than what we can see. DM has been observed many other ways including through gravitational lensing- a phenomenon where, according to general relativity, matter interacts with space to bend it in such a way that the trajectory of light gets bent. Observing the amount of gravitational lensing in galaxies has also provided strong evidence for DM [7]. The nature of DM has been an open question, but the general consensus at the moment is that no SM particle is a candidate for it. Searching for physics beyond the Standard Model gives us a chance to explain unresolved phenomena and understand the universe better.

In this thesis, I will summarize my work on several projects that explore different Beyond Standard Model (BSM) particles that could serve to solve outstanding problems with the Standard Model. All of these projects make use of a type of experiment called a beam dump, which I describe in detail in Section 1.1. In Section 1.2, I will describe an outstanding Standard Model issue related to the fact that the magnetic moment of the muon is different from Standard Model predictions, and then in Chapter 2, I will discuss my contributions towards a project that could resolve this discrepancy. In Section 1.3, I will describe a BSM particle called a millicharged particle (MCP), and then detail a search strategy for them that I worked on in Chapter 3. In Section Chapter 3, I also describe a search for one more BSM particle, an axion-like heavy particle (ALP), using a nearly identical experimental setup as the one for the MCP search. All of the work that makes up my thesis explores the shortcomings of our current models and offers alternative explanations that give experimentalists a clear path forward to have a chance at finding new physics. Note that this thesis will use natural units, i.e. $c = \hbar = 1$.

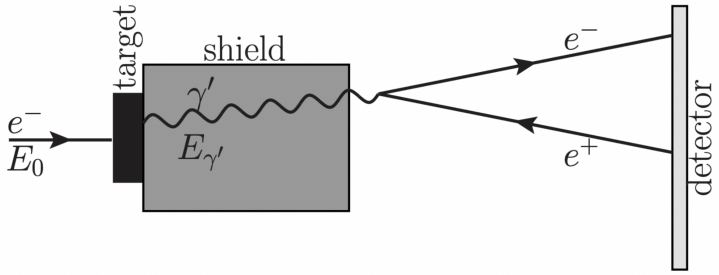

Figure 1. Schematic of dark photon production via electron beam dump through bremsstrahlung.

1.1 Beam Dump Searches

Beam dump[1] experiments are extremely useful in the search for new physics. These types of experiments use a high-energy proton or electron (and sometimes muon) beam, focused on a dense target to absorb hadronic cascades quickly, to potentially create new particles that can be detected outside of the target. Note that the new physics states that could be created from beam dumps can either decay "visibly" into other particles that are easily detected, or "invisibly" into particles that cannot be easily directly detected like neutrinos or dark matter. In the case of invisible decays, the signature for finding new physics would look like missing energy or momentum from the beam. Since electrons do not feel the strong force, their interactions are primarily electromagnetic, while proton beams have strong nuclear interactions and hence produce a wider variety of particles, in particular copious numbers of pions.

One of the earliest electron beam dump experiments occurred at the Stanford Linear Accelerator Center (SLAC) in 1982. This experiment, called E137, had a 20 GeV electron beam focused on an aluminum target, with shielding following, and then culminated in a 200 m decay region and calorimeter [9]. The team that proposed this experiment was hoping to have a chance at detecting neutral metastable particles like axion-like particles or photinos- a hypothetical spin-$1/2$ superpartner of the photon, predicted by supersymmetry, described in further detail in Sec. 1.2.2. Run 1 happened in January 1982, and the last run occured in November through December of that same year. They defined their signal roughly as seeing a single or double shower in the direction of the beam dump. Although they did not find any signal events, they were able to get some constraints on photino parameters, as well as the mass and lifetime of light axion-like bosons.

[1] Beam dump targets are longer than targets used for fixed target experiments, but I will use "beam dump" and "fixed target" interchangeably in this paper.

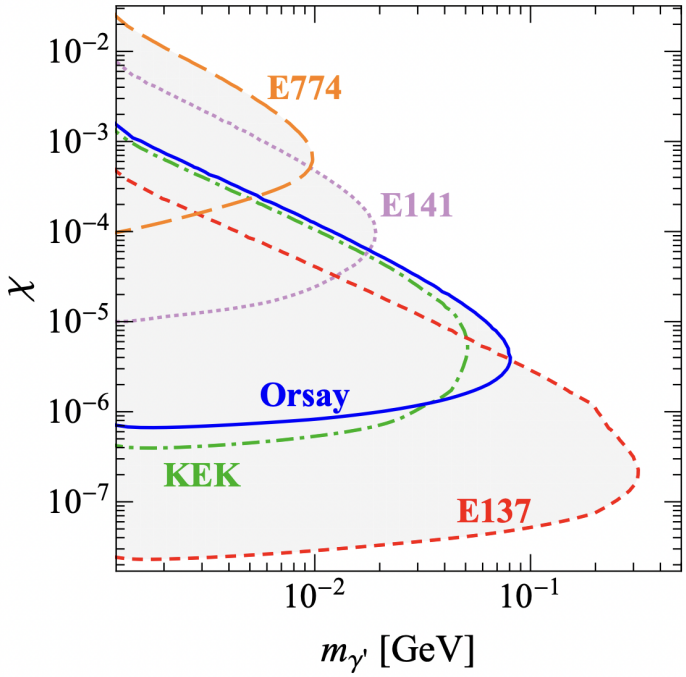

Figure 2. Here we see the 95% confidence level contours on the dark photon mass $m_{\gamma'}$ and kinetic mixing χ from early electron beam dump experiments, plotted in Ref. [8].

Other early electron beam dump experiments, mainly searching for light and neutral particles, include a search at KEK in 1986, E141 at SLAC in 1987 [10], a search for a light Higgs boson at the Laboratoire de l'Accélérateur Linéaire (LAL) in Orsay in 1989 [11], and an experiment called E774 at Fermilab in 1991 [12]. The constraints from these early experiments are shown in Fig. 2, pulled from Ref. [8]. It was later realized that all of these experiments are also probes for dark photons, a potential new particle that kinetically mix with SM photons and hence could mediate interactions between the SM and the BSM particles. These are described in further detail in Sec. 1.2.2. Fig. 1 shows a basic schematic of a beam dump search for dark photons, most often denoted by γ' or A'. This thesis will focus on additional BSM particles which can be produced at various beam dumps.

1.2 Muon $g-2$

1.2.1 Overview

A discrepancy between what we would expect to observe, according to the Standard Model, and what we find in experiments could point to a potential BSM particle to discover. In this section,

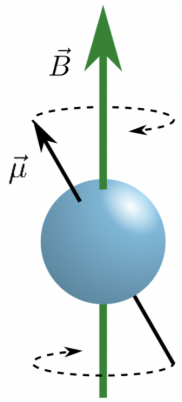

Figure 3. Precession of a particle's magnetic dipole moment (black) due to a magnetic field (green). Graphic adapted from Wikipedia.

we describe what motivates a suspicion that there could be a new particle that couples to muons enough to explain why experimentalists measure a slightly different magnetic moment $\vec{\mu}$ for a muon than theory predicts. A magnetic moment is a property of a particle that determines how that particle precesses in the presence of a magnetic field, depicted in Fig. 3. Specifically, the torque is given by

$$\vec{\tau} = \vec{\mu} \times \vec{B} = \left(g\frac{q}{2m}\right)\vec{S} \times \vec{B}, \tag{1}$$

where g is the dimensionless gyromagnetic ratio (or "g-factor"), $\vec{B}$ is the experimentally imposed magnetic field, and $\vec{S}$ is the particle's spin vector.

According to the Dirac equation, we should expect a muon's magnetic moment to be described by setting $g = 2$. However, quantum fluctuations from muon interactions with Standard Model particles, shown in Fig. 4, slightly increase this value by an amount captured by the anomalous magnetic moment $a_\mu = (g-2)/2$ [13]. The SM prediction of the anomalous magnetic moment is determined from the sum of all Standard Model contributions:

$$a_\mu^{SM} = a_\mu^{QED} + a_\mu^{EW} + a_\mu^{HVP} + a_\mu^{HLbL}, \tag{2}$$

where a_μ^{QED} is the contribution from quantum electrodynamics and a_μ^{EW} is the contribution from electroweak theory. The term a_μ^{HVP} represents the contribution from hadronic vacuum polarizations and the term a_μ^{HLbL} represents the contribution from hadronic light-by-light scattering. These

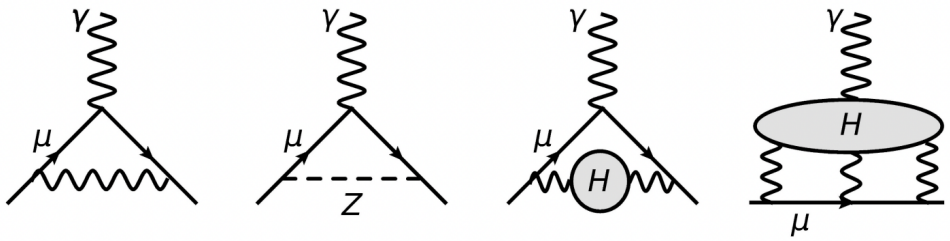

Figure 4. Here, we see the leading order contributions to the anomalous magnetic moment. The incoming and outgoing muons are drawn with solid lines with arrows and the external magnetic field is represented by a photon line coming in from the top. The diagrams shown from left to right are the one-loop QED diagram, the one-loop EW process involving Z-boson exchange, the leading-order hadronic vacuum polarization (HVP) diagram, and the hadronic light-by-light scattering (HLbL) contributions [13]. The grey region labeled "H" in the HVP diagram is a hadronic loop that demonstrates how electromagnetic interactions can interact with the vacuum to generate a quark-antiquark pair that can then annihilate to create photons. The HLbL diagram shows how photons can interact with each other through the exchange of virtual hadrons.

two terms involve the strong force and they dominate the uncertainty[2] in the overall a_μ^{SM} calculation. There are some competing methods for how to calculate the contributions from hadronic vacuum polarization to leading order (LO), next-to-leading order (NLO), and next-to-next-to leading order (NNLO). Similarly, there are also various methods for finding the contribution from hadronic light-by-light scattering. The various results from these methods are tabulated in Table 1, as well as their references [18]. Hence, there is some disagreement between theorists, but the theoretically computed anomalous magnetic moment a_μ^{sm} can range between $116\ 590\ 018.1 \times 10^{-11}$ and $116\ 590\ 019.8 \times 10^{-11}$, as depicted in Fig. 6, which is between a 5.1σ and 1.6σ deviation from experimental results, respectively [14].

Experimentalists have been measuring the anomalous magnetic moment of the muon since 1960, as shown is Table 2. Largely adapted from Ref. [19], we see the earliest experimental results for a_μ. Note that the precision over time was able to get better with improving technology, and Brookhaven National Laboratory (BNL), with its super–ferric storage ring magnet [38], was able to achieve a precision down to 0.72 ppm in 2001.

Fig. 5 shows the experimental setup at the 2001 BNL experiment. The Alternating Gradient Synchrotron (AGS) delivers protons to a target to produce pions that then decay into μ^+ particles. Through this process, the muons have their spins aligned with the direction of their flight axis. The muons are then injected into a storage ring with uniform magnetic field $\mathbf{B}$ throughout it. The muons

[2] Note that the uncertainty of a_μ is dominated by the hadronic contributions because at low energies, the strong force becomes non-perturbative, which complicates calculations. See references [14–18] for a summary.

Contribution	Value $\times 10^{11}$	Reference(s)
HVP LO (e^+e^-)	6931(40)	[19] [1]
HVP NLO (e^+e^-)	-98.3(7)	[20]
HVP NNLO (e^+e^-)	12.4(1)	[21]
HVP LO (lattice, udsc)	7116(184)	[22] [23] [24] [25]
HLbL (phenomenology)	92(19)	[26]
HLbL NLO (phenomenology)	2(1)	[27]
HLbL (lattice, uds)	79(35)	[28]
HLbL (phenomenology + lattice)	90(17)	[26]
QED	116 584 718.931(104)	[15] [16]
Electroweak	153.6(1.0)	[17]
HVP $(e^+e^-$, LO+NLO+NNLO)	6845(40)	[29]
HLbL (phenomenology + lattice + NLO)	92(18)	[26]
Total SM value	116 591 810(43)	[15]

Table 1. Contributions to a_μ^{sm} as calculated by different methods.

Experiment	Years	$a_\mu \times 10^{10}$	Precision	Reference
CERN I	1961	11450000(220000)	4300	[30]
CERN II	1962-68	11661600(3100)	270	[31]
CERN III	1974-76	11659100(110)	10	[32]
BNL	1997	11659251(150)	13	[33]
BNL	1998	11659191(59)	5	[34]
BNL	1999	11659202(15)	1.3	[35]
BNL	2000	11659204(9)	0.73	[36]
BNL	2001	11659214(9)	0.72	[37]
Average		11659208.0(6.3)	0.54	[37]

Table 2. Experimentally determined value for a_μ through the years, alongside the precision of the measurement.

are sent around the ring at a frequency of $w_c = \frac{eB}{m_\mu \gamma}$, and their spin axis precesses with angular frequency $w_s = \frac{eB}{m_\mu \gamma} + \frac{e}{m_\mu} a_\mu B$ for $\gamma = \frac{1}{\sqrt{1-v^2}}$. The difference between these two frequencies is captured by a quantity w_a which is related to the anomalous magnetic moment by $w_a = \frac{e}{m_\mu} a_\mu B$. Note that the precision of the measurement of a_μ is constrained most by the monitoring of the exact value of $\mathbf{B}$, determined by nuclear magnetic resonance (NMR) techniques [39]. Therefore, tracking the direction of the muon's spin vectors as they travel around the storage ring allows us to

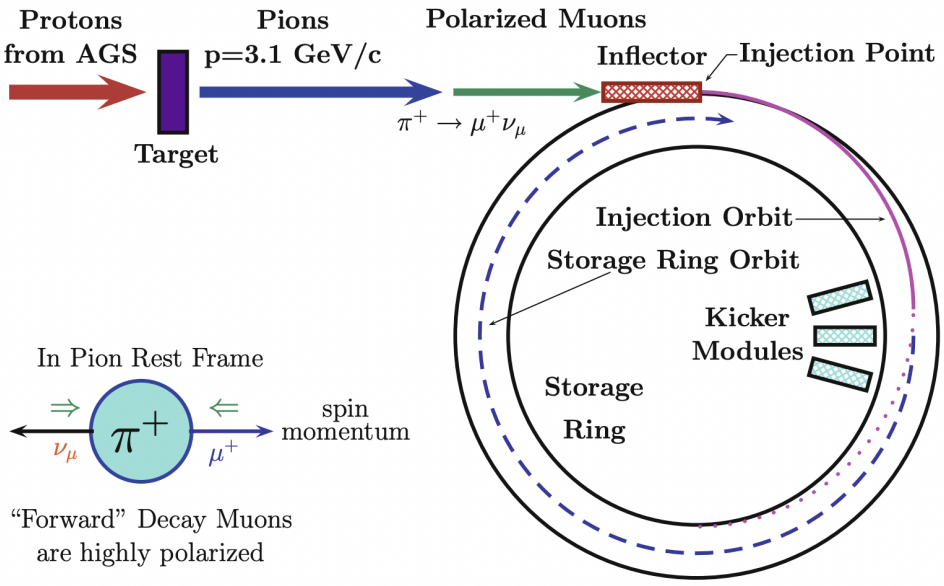

Figure 5. The experimental setup for the BNL 2001 determination of a_μ, pulled from Ref. [39].

calculate a precession frequency and hence a value for a_μ.

$$N_S = \sigma_{\text{accepted,sig}} \left(g_S{}^2\right) l \times \text{MOT} \tag{3}$$

$$N_B = \sigma_{\text{accepted,bg}}\, l \times \text{MOT} \tag{4}$$

BNL was able to take enough data between 2001 and 2006 that physicists could combine all the runs to conclude that the measured a_μ differed from the theoretically predicted value by about a 4.3σ deviation[3], which is a potential indicator of new physics but doesn't quite reach the 5σ benchmark to claim existence of a new particle [39]. In 2018, Fermilab was also able to take data in a similar experiment, but with various improvements such that they could have more precision with their measurements. Among the improvements was extensive use of 3 particle tracking simulation programs: the Geant4 simulation toolkit [40], and accelerator simulation toolkits BMAD [41] and COSY [42]. These tools helped them to get better estimates of various systematic uncertainties [13]. Analysis of the data took several years, and in 2021, the Muon G-2 Collaboration released an announcement that Fermilab's results confirmed the findings of the BNL measurements. The

[3] The claimed discrepancy is relative to choice of how a_μ^{sm} is calculated. This specific value corresponds to a choice of $a_\mu^{sm} \approx 116\,591\,783 \times 10^{-11}$.

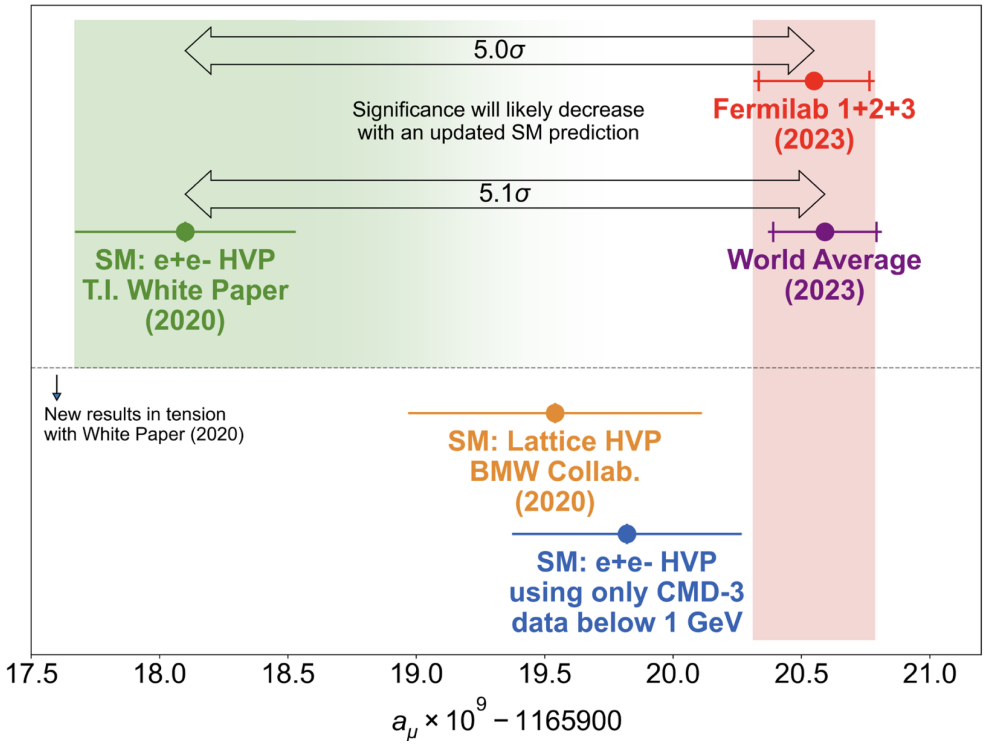

Figure 6. Here we see various Standard Model predictions of the anomalous muon magnetic moment a_μ plotted against the current experimentally determined world average, pulled from a Fermilab presentation [43]. Note that depending on the method to theoretically calculate a_μ, the discrepancy could be as large as $\sim 5\sigma$ or as small as $\sim 1.6\sigma$ [13].

measured anomalous muon magnetic moment is substantially different from the value that modern theory calculations predicts. In 2021 after run-1, the discrepancy was placed at about 4.2σ[4] [14]. Since that announcement, Fermilab has generated and analyzed additional data from several more runs, putting the current experimental average at $a_\mu^{exp} = (116\ 592\ 059 \pm 22) \times 10^{-11}$ [14]. Due to the disagreement of the calculated value of a_μ^{sm}, the community is hesitant to quote the exact discrepancy Δa_μ between experiment and theoretical prediction. The most recent experimental results alongside various theory predictions of a_μ are plotted in Fig. 6.

If Δa_μ can be explained by BSM states, we would use the following effective operator, in two-component fermion notation:

$$\mathcal{L}_{\text{eff}} = C_{\text{eff}} \frac{v}{\Lambda^2} (\mu_L \sigma^{\nu\rho} \mu^c) F_{\nu\rho} + h.c. \tag{5}$$

[4] This discrepancy uses $a_\mu^{\text{sm}} = 116\ 591\ 810 \times 10^{-11}$ obtained from the 2020 Whitepaper that a substantial portion of the community agrees with [18].

where C_{eff} is a dimensionless coefficient, $v = 246$ GeV is the Higgs vacuum expectation value (VEV), Λ is the BSM mass scale, μ_L and μ^c are the muon Weyl spinors, and $F_{\nu\rho}$ is the SM photon field strength [44]. Note that the BSM field, assumed to have a mass much larger than the muon mass, is integrated out. There are two cases where this contribution resolves $(g-2)_\mu$: one where the BSM states are singlets under the SM gauge group and one where the BSM states are charged under the electroweak (EW) gauge group [44, 45]. For the means of this summary and discussion, I focus on one-loop contributions to Δa_μ instead of more suppressed additional loop contributions.

1.2.2 Singlet Scenarios

Contributions to $(g-2)_\mu$ from singlet scenarios are suppressed by the SM muon Yukawa coupling $y_\mu \sim 10^{-3}$:

$$\Delta a_\mu \propto \frac{m_\mu y_\mu v}{(M_{\text{BSM}})^2} \tag{6}$$

where, as before, v is the Higgs vev and M_{BSM} is the mass of the BSM state. These singlet states can therefore be light, on the scale of $\lesssim$ GeV, and hence are amenable to beam dump searches [46, 47].

If Δa_μ is resolved by singlet BSM states, then Eq. 5 may be generated by new scalars or vectors. A vector interaction must come from a new $U(1)$ gauge extension to the SM:

$$\mathcal{L}_V \supset g_V V_\alpha \left(\mu_L^\dagger \bar{\sigma}^\alpha \mu_L + \mu^{L\dagger} \bar{\sigma}^\alpha \mu^c\right) + \frac{m_V^2}{2} V_\alpha V^\alpha. \tag{7}$$

One proposed vector solution to $g-2$ was to consider a simple extension of the SM to include a new vector singlet the corresponds to a hidden U(1) gauge group, called a massive dark photon. This particle would kinetically mix with SM photons and hence could mediate interactions between the SM and BSM particles. We would describe this model by the following Lagrangian:

$$\mathcal{L}'_{SM+U(1)} \supset \mathcal{L}_{SM} - \frac{1}{4} V_{\mu\nu}^2 + \frac{1}{2} \kappa V_{\mu\nu} F^Y_{\mu\nu} \tag{8}$$

where $V_{\mu\nu}$ is the field strength of $U(1)'$, $F^Y_{\mu\nu}$ is the field strength of $U(1)_{SM}$, and κ is the mixing angle [49]. Note that under this model, the dark photon couples to all charged fermions with equal strength due to the kinetic mixing between dark and visible photons. This model has since been ruled out due to data obtained from beam dump experiments (described in further detail in Sec.

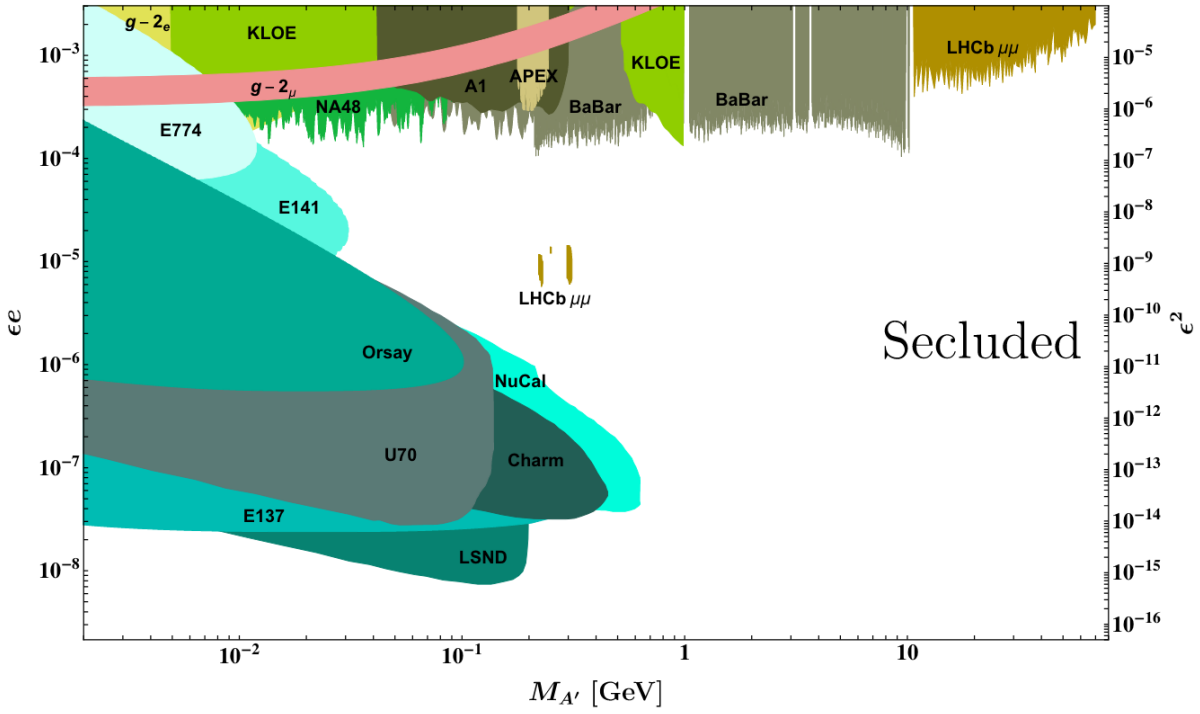

Figure 7. Here we see, from Ref. [48], constraints from past and current experiments on the dark photon mass $M_{A'}$ and kinetic mixing parameter ϵ.

1.1), B-factory searches, and rare meson decay searches, summarized in Fig. 7 [47, 48, 50].

Additional vector solutions have been exhaustively ruled out or have planned current or future experimental searches that will systematically test the remaining viable parameter space [47]. The only anomaly-free vector models where the BSM vector couples directly to muons are $U(1)_{B-L}$, $U(1)_{L_i-L_j}$, and $U(1)_{B-3L_i}$ for baryon number B, lepton number L, and lepton flavors L_i for $i = e, \mu, \tau$. Since these models all include the BSM vector coupling to electrons as well as muons, they have been ruled out by the same experiments that also ruled out the kinetically mixed dark photon solution [47, 48, 50]. The only remaining viable minimal vector explanation of $g-2$ is the gauged $U(1)_{L_\mu-L_\tau}$ model which will be fully tested soon by planned experiments [51–54].

For example, in 2020, a group at Fermilab suggested that NA62 look at kaon decays $K \to \mu\nu X$ where X is a gauged $L_\mu - L_\tau$ vector or a muonphilic scalar [55]. X could either decay into a pair of muons or it could decay invisibly. NA62 has not yet run this experiment. Similarly, another group in 2018 proposed a search called "M^3". This experiment would use Fermilab's 15 GeV muon beam, set on a fixed target, to generate a scalar S or vector V particle through bremsstrahlung that

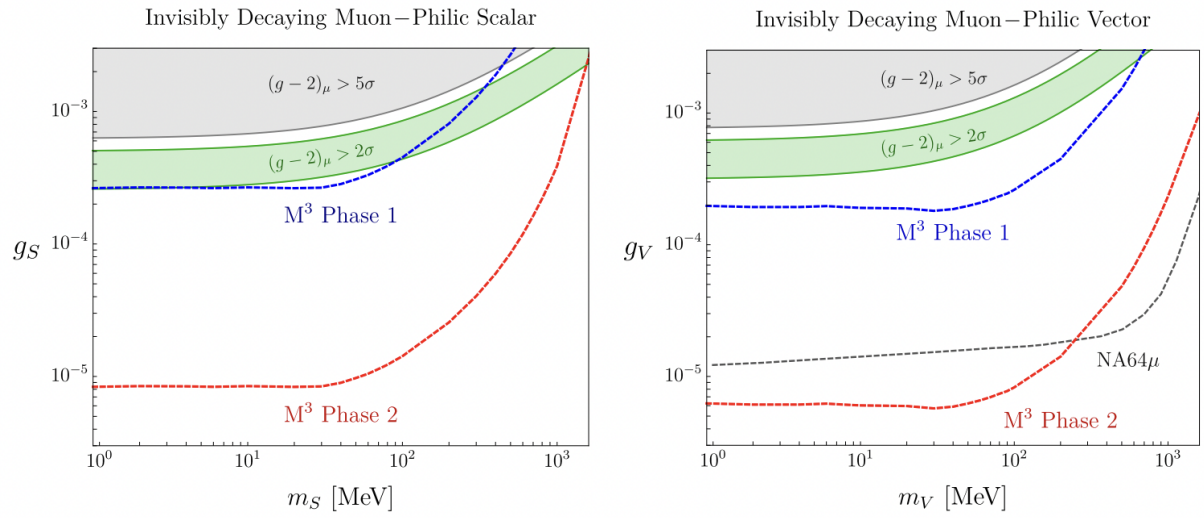

Figure 8. Here we see the projected sensitivity of two phases of the proposed M^3 experiment from Ref. [45]. Note that M^3 Phase 1 would use a flux of about 10^{10} MOT to probe the $(g-2)_\mu$ parameter space for vectors with $m_V \lesssim 500$ MeV and scalars with $m_S \lesssim 100$ MeV. M^3 Phase 2 would use a higher beam flux to probe additional parameter space that isn't relevant to $(g-2)_\mu$.

could invisibly decay into undetectable things like neutrinos [45]. For a process of $\mu^- N \to \mu^- N S$ or $\mu^- N \to \mu^- N V$, where S and V are the BSM states and N is the target nucleus, this proposed experiment would consider a signal event to be an event whose outgoing μ^- has a greater than 50% reduction in its incident energy. M^3 is important because a null result would exclude any BSM explanation of $(g-2)_\mu$ from invisibly-decaying muonphilic particles with masses below 100 MeV. M^3 has also not run yet. Finally, NA64μ, a missing-energy experiment at CERN, was run in 2023. This experiment used a 160 GeV muon beam from the M2 beamline at the CERN Super Proton Synchrotron, and it achieved a flux of 2×10^{10} muons on target [56]. They found no signal and were able to place a 90% exclusion limit on the gauged $L_\mu - L_\tau$ vector model parameter space, as seen in Fig. 9 alongside constraints from data and projected constraints from M^3.

For BSM scalar states coupling to the muon, they would be described by the Lagrangian

$$\mathcal{L}_\mathcal{S} \supset -(g_s S \mu_L \mu^c + h.c.) - \frac{1}{2} m_s^2 S^2 \tag{9}$$

where μ_L is the left-handed muon field and μ^c is the right-handed muon field. Note that the term that couples the S with muons is not gauge invariant under $SU(2)_L \times U(1)_Y$, so it must arise from

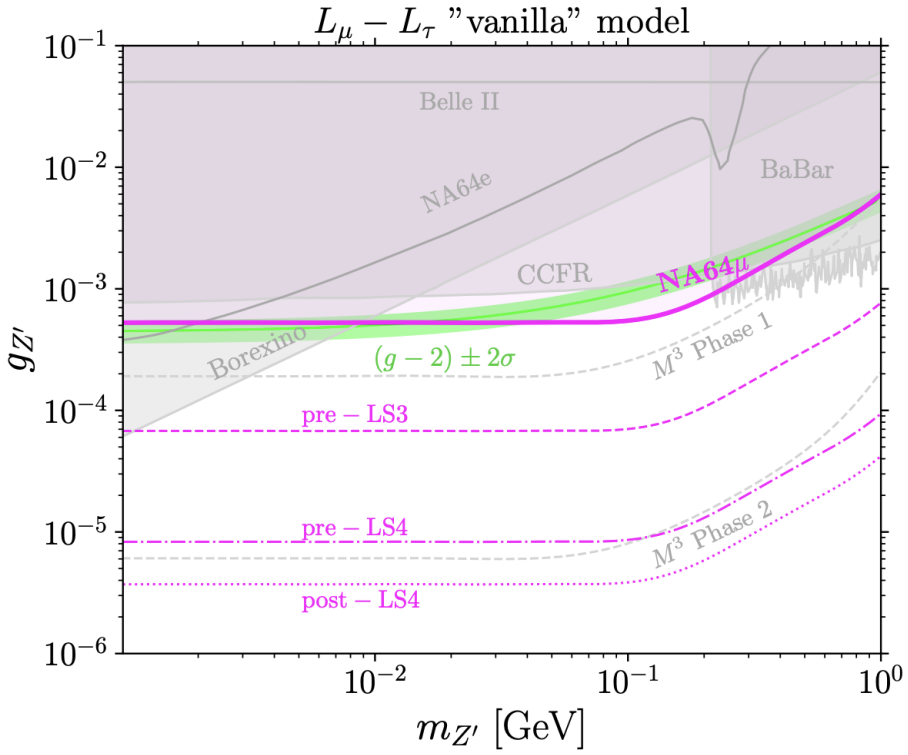

Figure 9. Here we see the 90% CL exclusion limits from NA64μ on the coupling $g_{Z'}$ as a function of the Z' mass, $m_{Z'}$, for the vanilla L_μ-L_τ model where the $Z' \to \bar{\nu}\nu$ dominates [56]. Note that LS3 stands for the CERN Long Shutdown 3- a period from 2026-2029 during which enhancements would be made to the Large Hadron Collider (LHC) to improve its performance. Similarly, LS4 stands for a shutdown that is planned for 2034-2036 during which more upgrades to the LHC would be installed [57]. According to Ref. [58], the expected MOT for each of these periods is 3×10^{11} for the pre-LS3 period, 2×10^{13} for the pre-LS4 period, and 10^{14} for the post-LS4 period.

a higher dimension operator involving the SM Higgs doublet H:

$$\mathcal{L} \supset -\frac{1}{\Lambda} S H^\dagger L \mu^c + h.c. \tag{10}$$

where Λ is an energy scale up to which this operator is valid. After electroweak symmetry breaking, this operator becomes $\mathcal{L}_S \supset -(g_s S \mu_L \mu^c + h.c.)$ for $g_s = \frac{v}{\Lambda}$ after electroweak symmetry breaking [59]. For $m_S < 2m_\mu$, the primary decay mode is $S \to \gamma\gamma$ which can be investigated by missing energy searches at colliders or visible decays at beam dump experiments [59]. For $m_S > 2m_\mu$, the primary decay mode is $S \to \mu^+\mu^-$, for which there have been many proposed searches [45, 60, 61]. Note that experimental data, for example data from the BABAR Collaboration [62], has ruled

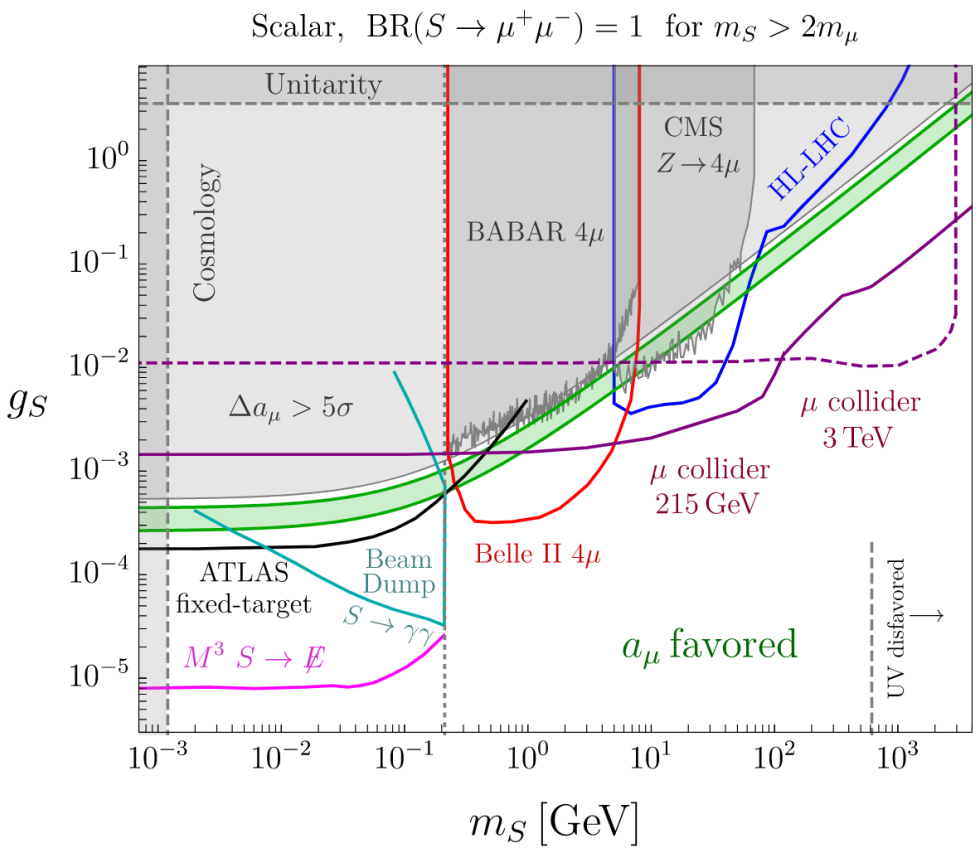

Figure 10. Pulled from Ref. [59], we see limits and projections on muonphilic singlets, assuming $BR(S \to \mu^+\mu^-) = 1$ and the scalar mass is greater than $2m_\mu$. We see existing experimental limits shaded in gray. We also see projected sensitivities from experiments M^3 [45], NA64μ [60], and ATLAS fixed-target experiments [61]. The Belle II 4μ and HL-LHC projected constraints are further discussed in Ref. [59].

out some parameter space for the muonphilic scalar model in addition to the gauged $L_\mu - L_\tau$ vector model. Various constraints from experimental data and projected limits for the muonphilic scalar model are shown in Fig. 10. For the majority of my thesis work that focused on the $g-2$ anomaly, I showed how to parasitically use a Fermilab-based existing experiment with minimal modifications to search for muonphilic scalars in the mass range $m_S > 2m_\mu$. This search would be complementary to the existing proposed searches to date. I describe my work on this further in Chapter 2.

1.2.3 Electroweak (EW) Scenarios

Contributions to $(g-2)_\mu$ from electroweak scenarios are not suppressed by the SM muon

Yukawa coupling:

$$\Delta a_\mu \propto \frac{m_\mu g_{\text{BSM}} v}{(M_{\text{BSM}})^2}. \tag{11}$$

Hence, EW scenario states are on a heavier mass scale than singlet states, and they are harder to search for in experiments [46, 47]. The models where the BSM states are charged under the EW gauge group have a mass scale for the BSM states that span a range of 100 GeV - 100 TeV [44, 46, 63–69]. Regarding explanations of new physics with large BSM masses, one explanation is a theory of minimal supersymmetry. The Minimal Supersymmetric Standard Model (MSSM) doubles the particle content of the SM by introducing a mapping between fermions and bosons such that each SM state X has a (yet undiscovered) supersymmetric state $\tilde{X}$ whose spin differs by half of a unit, which means that bosons have fermion superpartners, and fermions have boson superpartners. Superpartners of fermions get the prefix "s" and those of bosons get the suffix "-ino" [70]. For example, the neutrino, which is a fermion, has a corresponding superpartner called a "sneutrino" that is a boson. A Higgs boson has a superpartner called a Higgsino. This model also has two scalar Higgs doublets h_u and h_d, along with their superpartners $\tilde{h}_u$ and $\tilde{h}_d$. According to a MSSM model, the leading order contributions to a_μ would come from muon interactions with superpartners as shown in Fig. 11. For MSSM to resolve Δa_μ, the new physics particles must have a mass in the range of $100 < M_{SUSY} < 500$ GeV [71]. In addition to having the ability to explain the muon magnetic moment anomaly, MSSM is also attractive because it solves the Higgs hierarchy problem of the SM (i.e. SUSY would introduce a symmetry between fermions and bosons that can cancel large contributions to the Higgs mass, thus explaining why the Higgs mass is so much smaller than the Planck scale), and if superparticles are found, they would not only be candidates for dark matter, but their very existence could provide evidence for grand unification theories [71]. To date, no particles predicted by SUSY have been detected despite many efforts, and physicists have since moved away from SUSY as a viable minimal explanation of $(g-2)_\mu$. See Ref. [47] for a comprehensive summary of all electroweak (EW) scenario solutions.

1.3 Millicharged particles

Millicharged particles (MCPs), are a fascinating candidate for BSM physics. The U(1) Standard Model gauge group allows particles to carry arbitrary small charges, i.e. $Q = \epsilon e$. MCPs are hence theoretical particles that could carry electric charges smaller than the smallest quark charges. They would interact weakly with Standard Model particles, but they would not couple to the strong or

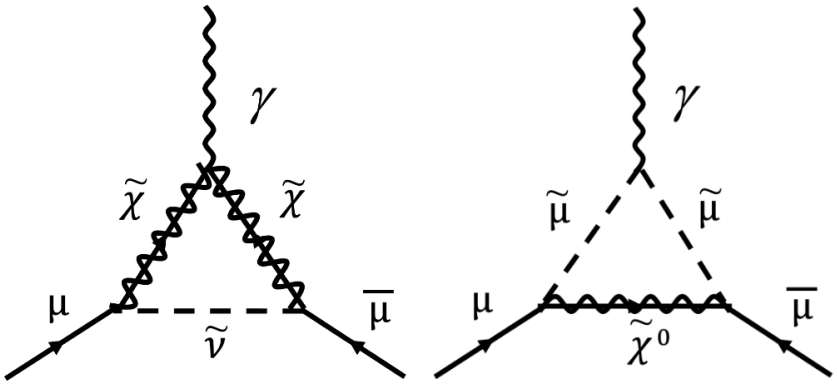

Figure 11. This sneutrino-chargino loop [left] and smuon-neutralino loop [right] contribute to the calculation of a_μ.

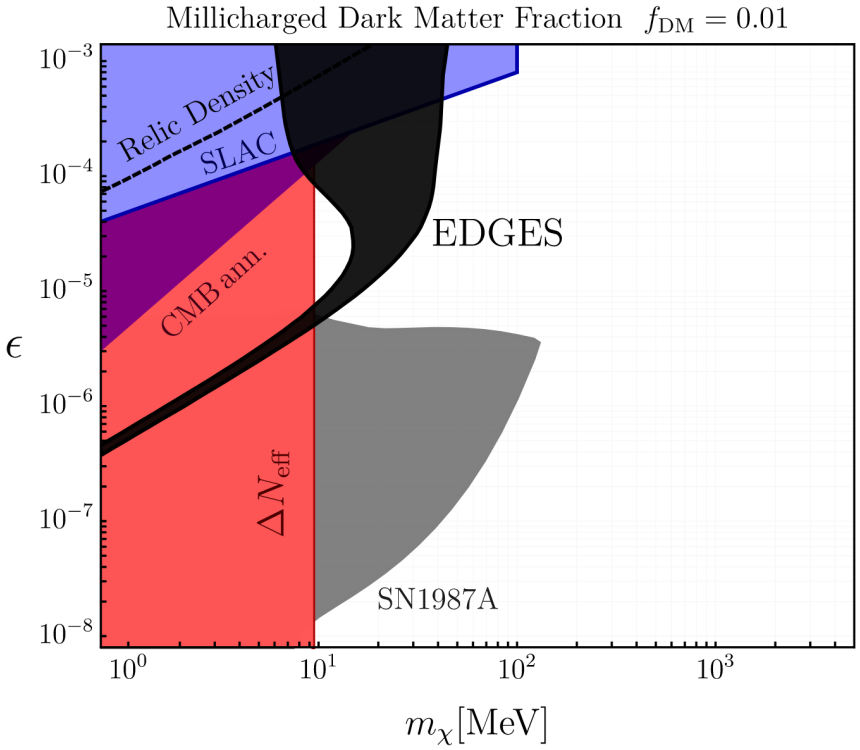

Figure 12. Pulled from Ref. [72], for a model where 1% of dark matter is made of MCP, we see the constraints on fermionic millicharged DM from Supernova 1987A (grey) [73]. This supernova was discovered in 1987 and data from it gives us information about MCP since the existence of weakly-coupled particles could change the neutrino emission from SN1987A by providing new channels to cool the star. We also see constraints from the search for millicharged particles at SLAC (blue), alongside other astrophysical constraints [74–77]. The solid black region represents the parameter space in which the fermionic millicharged DM could explain the anomalous additional signal that the EDGES Collaboration found.

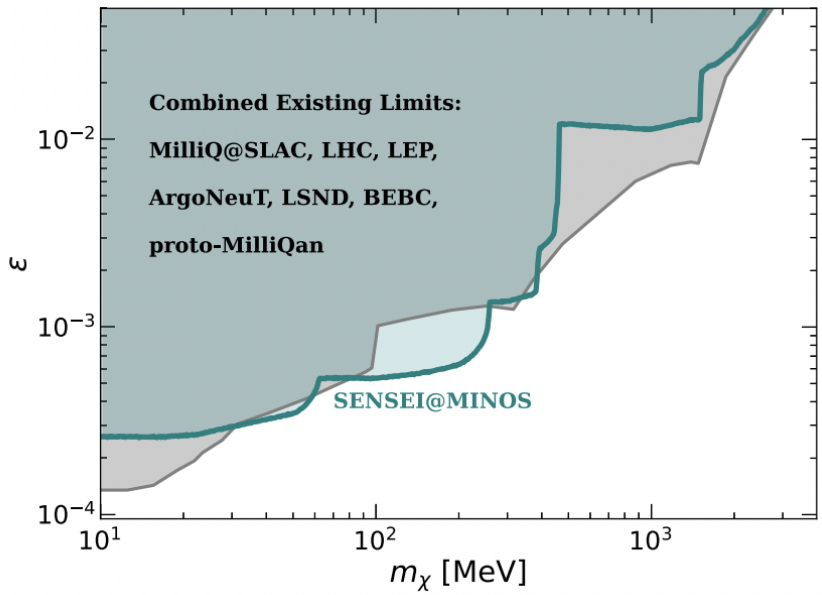

Figure 13. Combined existing limits on MCP parameter space combined with the existing limits from SENSEI at Fermilab's MINOS site [74, 82–84]. Note that the constraints from "proto-MilliQan" were obtained using the MilliQan demonstrator detector during LHC's Run 2.

weak force. These particles could potentially constitute a portion of the relic abundance of dark matter [72, 78]. Specifically, in 2018, the Experiment to Detect the Global Epoch of reionization Signature (EDGES) Collaboration found a signal in the absorption profile of the sky-averaged radio spectrum whose amplitude was larger than expected, at a confidence level of 3.8σ. One explanation for such a finding is that primordial gas in the early universe, around redshift $z \approx 17$, underwent additional cooling due to scattering from light DM [78–81]. Combining experimental and astrophysical constraints, it has been determined [72] that millicharged dark matter could only make up about $0.3 - 2\%$ of the total amount of dark matter in the universe. The parameter space has also been constrained to reveal that if this millicharged dark matter does exist, its mass must be in the range of $10 - 80$ MeV, and its charge ϵ must be between 10^{-6} and 10^{-4}, as depicted in Fig. 12.

The most sensitive laboratory-based probes of MCP are collider searches, electron fixed-target experiments, proton fixed-target experiments, and neutrino experiments [3]. The LHC performed its first MCP search during the LHC Run 2 between 2015 and 2018, using a center of mass energy $\sqrt{s} = 13$ TeV and an array of scintillators. This experiment used a detector called the "MilliQan Demonstrator", which served as the protoype of a larger detector called the "MilliQan Detector"

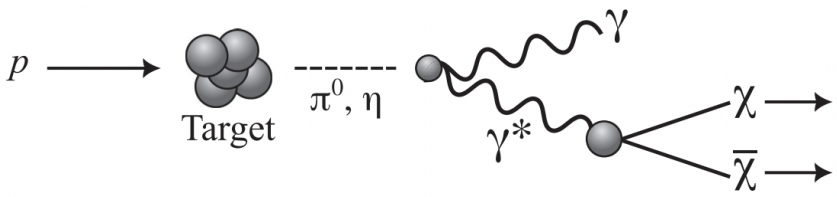

Figure 14. A schematic of MCP pair production from the decay of neutral mesons generated from proton-nucleon scattering [82].

[85]. This experiment excluded MCPs of masses $20 \leq m_\chi \leq 4700$ MeV with charges $0.006 \leq \epsilon \leq 0.3$ at a 95% confidence level. Plans to upgrade the MilliQan detector and collect data were planned to occur during the LHC Run 3, between 2022 and 2026. The upgrades are currently complete but data is still being taken. Note that the MilliQan detector inspired a proposal for an experiment named "FerMINI". Proposed in 2019, this experiment would search for MCPs at Fermilab's DUNE near detector hall using a scintillator array as well [3].

The overall leading MCP constraints are shown in Fig. 13. Note that the 2020 Fermilab-based experiment named SENSEI offers the best constraints for MCPs in the range $30 \leq m_\chi \leq 380$ MeV [82]. This experiment used a dark matter detector and neutrino beam to search for MCPs.

Production channels of MCPs that have been considered thus far in the literature are neutral meson decays, Drell-Yan production, and bremsstrahlung production. A schematic for neutral meson decay, adapted from Ref. [82], is shown in Fig. 14, where a virtual photon from meson decay generates two MCPs. Commonly investigated neutral meson decays include $\pi^0 \to \gamma\chi\bar{\chi}$, $\eta \to \gamma\chi\bar{\chi}$, $J/\psi \to \chi\bar{\chi}$, and $\Upsilon \to \chi\bar{\chi}$ [3]. Drell-Yan production involves the annihilation of quarks through a virtual photon to produce MCPs: $q\bar{q} \to \chi\bar{\chi}$. Finally, the last MCP production channel explored in literature is through bremsstrahlung: The primary beam particle scatters off of a heavy nucleus and emits a virtual photon which results in a MCP pair. However, the literature thus far has overlooked the production of MCPs through charged pion scattering, which motivated my second project described in Chapter 3. In this project, I adapted FerMINI's experimental setup [3] to investigate how much additional sensitivity to MCP detection we can obtain by considering additional signal from charged pions scattering off of target nuclei in beam dumps.

CHAPTER 2: NEW SCALAR FIELD TO RESOLVE $(g-2)_\mu$

The work in this section is based on "New Searches for Muonphilic Particles at Proton Beam Dump Spectrometers", published as an article in Physical Review D Volume 107, Article 116026 on 30 June 2023. The work was authored by myself, Christian Herwig, Yonatan Kahn, Gordan Krnjaic, Cristina Mantilla Suarez, Nhan Tran, and Andrew Whitbeck. Yonatan Kahn and I were affiliated with Department of Physics, University of Illinois at Urbana-Champaign. Christian Herwig, Cristina Mantilla Suarez, and Nhan Tran were affiliated with Fermi National Accelerator Laboratory. Gordan Krnjaic was affilliated with Fermi National Accelerator Laboratory, the Kavli Institute for Cosmological Physics at the University of Chicago, and Department of Astronomy and Astrophysics at the University of Chicago. Andrew Whitbeck was affiliated with Texas Tech University. The bulk of my work on this project involved simulating events with MadGraph [86], implementing the elastic and inelastic form factors (and finding a typo in the reference we were using), finding and implementing cuts that maximized our sensitivity, and collaborating with the group to generate the final reach plot.

In recent years, searches for muonphilic particles – new particles beyond the Standard Model that couple primarily to muons – have attracted considerable interest, inspiring novel strategies involving beam dumps [87], B-factories [88], missing energy and momentum experiments [45, 89], the Large Hadron Collider [90], and even future muon colliders [46] (see Ref. [59] and references therein for a survey of such techniques). In part, this popularity is related to the possible evidence for new physics from the Fermilab Muon $g-2$ collaboration, which has recently measured the anomalous magnetic moment of the muon [1, 91–93]. This new result is consistent with the earlier Brookhaven measurement [19] and the world average for $a_\mu \equiv \frac{1}{2}(g-2)_\mu$ now deviates from the Standard Model (SM) prediction [18] by

$$\Delta a_\mu = a_\mu^{\rm exp} - a_\mu^{\rm theory} = (251 \pm 59) \times 10^{-11} \,, \tag{12}$$

which constitutes a statistically significant 4.2σ discrepancy.[5]

If this discrepancy is due to new physics, there are necessarily new particles in nature that

[5] Previous lattice QCD extractions of the hadronic light-by-light [28, 94, 95] and hadronic vacuum polarization [96] contributions to a_μ are consistent with both the measured value and semi-analytical calculations based on R-ratio data. However, the BMW collaboration [97] has extracted a SM prediction of a_μ that is consistent with the measured value. This result is in tension with a_μ as determined by R-ratio methods and might be in tension with the SM electroweak fit [98–100], so future lattice calculations and improved R-ratio data will be necessary to conclude whether Eq. (12) is evidence of new physics.

couple to muons. These particles can be classified according to whether they are heavy with order-unity couplings to the muon (e.g. new weak-scale states charged under SM gauge interactions) or light and feebly coupled. If the former possibility is realized in nature, such states could be discovered with future high-energy collider searches [46, 47]. If instead, the particles in question are light and feebly coupled, the possibilities for new physics are much narrower: the particles must be muonphilic scalars or vectors, which are singlets under the SM gauge group [59]. It has been shown that, for most available decay channels, existing and planned intensity frontier experiments have sufficient sensitivity to discover these light new states if their couplings to the muon resolve Δa_μ [45, 59, 87, 89, 101–103]. However, there is a notable exception to this otherwise comprehensive coverage: scalars S with sub-GeV masses and prompt dimuon decays.

In this paper, we propose muon spectrometers at proton beam dumps as promising experiments to search for $S \to \mu^+\mu^-$ and potentially discover the new physics responsible for Δa_μ. While the parameter space of interest in this paper is framed around scalar particles that resolve $g-2$, the search strategy we present is general and can be adapted to search for any new particles that decay appreciably to dimuons. Proton beam dumps feature enormous luminosity and copious secondary production of muons through pion decay, evading some of the event rate limitations that cap the sensitivity of muon beam experiments such as M^3 [45] and NA64-μ [89]. The production of S from bremsstrahlung during nuclear scattering in the dump, $\mu^\pm N \to \mu^\pm N S$, is only weakly dependent on the muon energy, and thus a mono-energetic muon beam is not necessary. As we will show, for S couplings that resolve Δa_μ, the S decay is prompt and the signal is an invariant mass peak in opposite-sign muons emerging from a single vertex in the dump. The sensitivity is therefore driven by the invariant mass resolution, and we will argue that selecting events in the final portion of the dump, combined with a high-momentum-resolution spectrometer magnet, suffices to achieve the invariant mass resolution needed to observe a signal above the SM background from continuum QED production. We emphasize that our proposed search strategy involves no new hardware and can be parasitically implemented at any proton beam spectrometer experiment.

This paper is organized as follows. We introduce the scalar singlet model in Section 2.1 as a solution to Δa_μ, which can also be seen as a representative example of a muonphilic model. In Section 2.2, we introduce the basic experimental concept and discuss the generic requirements on the detector configuration. Then, in Section 2.3, we calculate the scalar singlet production rate as well as that of the leading irreducible background processes, and propose cuts which maximally

exploit the different kinematics of signal and background. In Sec. 2.4, we take the specific case of the SpinQuest experiment at Fermilab and detail key characteristics of the experiment, expected backgrounds, and the potential sensitivity to muonphilic scalars. Finally, we conclude and provide an outlook on the near-term experimental prospects in Section 2.5.

2.1 Scalar Singlet Model

We extend the SM by a scalar S which is a singlet under the SM gauge group, and which couples exclusively to muons through the Yukawa interaction

$$\mathcal{L} \supset g_S S \overline{\mu}\mu \,, \tag{13}$$

where g_S is a dimensionless coupling constant. This coupling induces a shift in Δa_μ at one-loop level which yields

$$\Delta a_\mu^S = \frac{g_S^2}{8\pi^2}\int_0^1 dz \frac{(1+z)(1-z)^2}{(1-z)^2 + z(m_S/m_\mu)^2} \approx 2\times 10^{-9}\left(\frac{g_S}{10^{-3}}\right)^2\left(\frac{700\,\mathrm{MeV}}{m_S}\right)^2 \,, \tag{14}$$

where m_S is the mass of S. The approximate equality holds in the limit $m_S \gg m_\mu$, and gives a sense for the typical size of couplings required. In the absence of other interactions, for $m_S > 2m_\mu$ the only tree-level decay channel is $S \to \mu^+\mu^-$, with a corresponding width of

$$\Gamma(S\to\mu^+\mu^-) = \frac{g_S^2 m_S}{8\pi}\left(1-\frac{4m_\mu^2}{m_S^2}\right)^{3/2} \,. \tag{15}$$

At one-loop level, there are also $S \to \gamma\gamma$ and $S \to \bar{\nu}\nu$ decay channels, but these are further suppressed by α^2 and G_F^2, respectively, so the dimuon channel has a branching fraction near unity.

As noted in Ref. [104, 105], the interaction in Eq. (13) is not gauge-invariant under the electroweak symmetry of the SM, but it may be generated from the dimension-5 operator

$$\mathcal{L}_{\mathrm{UV}} = \frac{1}{\Lambda} S H^\dagger L_\mu \mu^c + \mathrm{h.c.} \,, \tag{16}$$

where, in two-component Weyl fermion notation, L_μ is the second generation lepton doublet, μ^c is the right handed muon field, H is the SM Higgs doublet, and Λ is the mass scale at which new

particles have been integrated out. Matching Wilson coefficients yields $g_S = v/\Lambda$, so for $g_S = 10^{-3}$, the cutoff scale of the effective field theory defined by Eq. (13) is valid up to $\Lambda \approx 250$ TeV. Consequently, for the parameter space we consider in this paper ($m_S \lesssim 5$ GeV and muon beam energies below 50 GeV), the effective theory description is perfectly valid.

Previous work to constrain this model has focused on search strategies to probe the invisible decay channel if S couples to neutrinos or dark matter [45, 52, 106], as well as other loop-suppressed modes for $m_S < 2m_\mu$ which lead to long-lived S and displaced diphoton decays through one-loop processes [87]. By contrast, here we focus on the visible decay mode $S \to \mu^+\mu^-$, which we assume has a 100% branching fraction for $m_S > 2m_\mu$. Allowing for alternate invisible decay modes (which may imply the presence of new dark states) reproduces the phenomenology of previous invisible decay studies such as M^3 [45] and NA64-μ [106] and we will not consider these in this paper.

Finally, we note that the phenomenological search strategy introduced below may also be applied to new spin-1 vector particles whose couplings to the muon resolve Δa_μ. However, nearly all theoretically consistent models for visibly-decaying vectors in this mass range have been ruled out by laboratory searches [59, 107]. The only remaining anomaly-free $U(1)$ extension to the SM that can still resolve Δa_μ is gauged $L_\mu - L_\tau$ [107] for vector masses between $\sim$ 10-200 MeV, where the lower bound is set by cosmology [51] and the upper bound is set by the BABAR $e^+e^- \to 4\mu$ search [88]. Thus, for nearly all of the remaining viable parameter space in this model, the vector particle decays invisibly to neutrinos and is, therefore, testable with NA64-μ [108] and M^3 [45], which are optimized for missing momentum signatures. By contrast, scalar particles that resolve Δa_μ can still visibly decay to dimuons [59], so we focus on this scenario throughout our analysis.

2.2 Proton Beam Dump Spectrometer Concept

To search for muonphilic scalars as a possible explanation for Δa_μ, we require a large flux of muons on a target which will produce the scalars via bremsstrahlung, as shown in Fig. 15 (red boxed inset). From the decay width of the S to muons in Eq. (15), the lab-frame decay length is:

$$L \approx 8 \times 10^{-8}\,\mathrm{m} \left(\frac{E_S/m_S}{10}\right) \left(\frac{700\,\mathrm{MeV}}{m_S}\right) \left(\frac{10^{-3}}{g_S}\right)^2 , \tag{17}$$

where we have taken the $m_S \gg m_\mu$ limit. Except for a very small region of phase space just

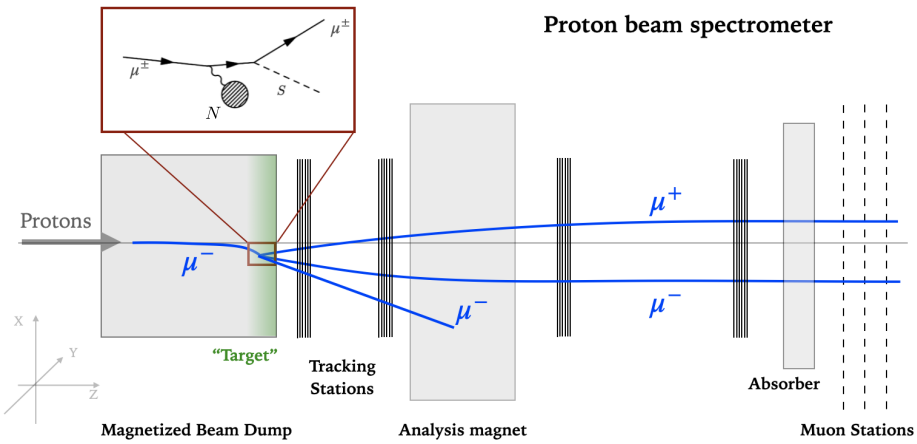

Figure 15. Proton beam dump spectrometer signature of prompt muonphilic scalars produced in the back of the beam dump, labeled as the "Target", and reconstructed by the downstream tracking stations. The spectrometer setup is inspired by the existing SpinQuest experiment, but we argue that the search strategy presented in this paper can work for other proton beam spectrometer configurations with a large flux of muons (see main text for details).

above the dimuon threshold, the couplings required to explain Δa_μ imply that the S must decay promptly. This remains true even if there are additional invisible decay modes, since those will only increase the total width and hence decrease the decay length.

Therefore, the target itself cannot be very dense or else the momentum resolution will be degraded by multiple scattering, so a large muon flux is important to compensate for this lower density. In this paper, we consider the SpinQuest spectrometer as an example of an experimental setup to search for such muonphilic scalars. A schematic inspired by SpinQuest is shown in Figure 15. The proton beam travels through some magnetized material producing a large fraction of $\mu^\pm$ with $\mathcal{O}(20\ \mathrm{GeV})$ energies, most of which originate from pion decays. These secondary muons are produced along the beam dump and those muons traversing the target region, which in Fig. 15 is denoted by "Target" in green, can produce the S during a nuclear scattering event and the outgoing daughter muons have enough momentum to exit the dump and be detected. The path of the beam muon is deflected by the magnetized dump, while the analysis magnet alters the trajectories of the three outgoing muons to measure their curvature and hence momenta. The signal is thus two or three muons originating from the same vertex in the dump (depending on whether the third muon has a high enough momentum to emerge from the tracking stations), with the invariant mass of an opposite-sign muon pair reconstructing the mass of the S. We note that Ref. [101] previously considered S production from secondary muons at SpinQuest but focused on the e^+e^- decay

mode, whereas we focus on the irreducible decay to $\mu^+\mu^-$.

This search strategy can be employed at any proton beam dump spectrometer with the following key features:

- a high-intensity proton beam with high repetition rate, from which a large flux of muons is produced. A high repetition rate, near-continuous wave beam is more valuable than a beam of similar current but lower duty factor (pulsed) in order to reduce combinatorial backgrounds;

- a beam dump that is many nuclear interaction lengths thick, to allow predominantly muons to exit through through the dump and greatly reduce hadronic backgrounds, with the last portion of the dump serving as the target;

- a beam of sufficient energy to produce secondary muons capable of traversing the entire dump, retaining sufficient momenta to both produce the S and to boost its decay products into the spectrometer's acceptance;

- a magnetized beam dump to spatially spread out positive and negatively charged beam muons to reduce the combinatorial backgrounds;

- a detector that can trigger on dimuon coincidences in the presence of a high muon flux, has good momentum and angular resolution such that the invariant mass resolution is mostly influenced by multiple scattering, and is granular enough to reject backgrounds.

As we will see below, the dominant irreducible background is muon pair production from QED, while an important reducible background is the combinatorial background from independent production of muons in the dump. Because the search strategy is a bump hunt in invariant mass, the relative size of the QED background will be driven both by the invariant mass resolution and the effective target length ℓ_T (the green region in Fig. 15). These requirements are in some tension because a larger ℓ_T improves the sensitivity (which scales as $\sqrt{\ell_T}$ assuming Poisson fluctuations on the background), but also leads to increased multiple scattering and a degraded mass resolution, as well as a larger combinatorial background. A full optimization of the sensitivity with respect to ℓ_T requires a concrete experimental design and is beyond the scope of this study, but as an example, we will take $\ell_T = 100$ cm and a 15% invariant mass resolution. This figure represents the

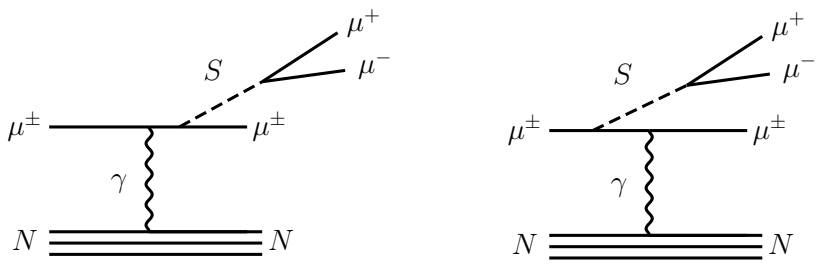

Figure 16. Feynman diagrams representing the dominant signal processes for S production in muon-nucleus scattering.

combined effects of an intrinsic 5% experimental resolution and multiple scattering and is further justified with simulations in Sec. 2.4.2.

2.3 Signal and Background Rates, Kinematics, and Cuts

2.3.1 Signal and Irreducible Background

Our signal process is on-shell S production from muons scattering off a fixed target of nuclei N, $\mu^\pm N \to \mu^\pm N S$, followed by the prompt $S \to \mu^+\mu^-$ decay (Fig. 16). The experimental signature includes at least two opposite-sign muons in the final state originating from a single vertex which reconstruct the invariant mass of the S. The SM background is dominantly continuum muon pair production $\mu^\pm N \to \mu^\pm N \mu^+\mu^-$ through electromagnetic processes, including both the radiative (Fig. 17, left) and Bethe-Heitler trident (Fig. 17, right) diagrams [109, 110]. The background from off-shell Z production is negligible at the beam energies we consider. Photon-initiated background processes resulting from a hard photon bremsstrahlung, $\mu N \to \mu N \gamma$ followed by $\gamma N \to$ hadrons, are negligible since they can be mitigated with a mild di-muon selection and are further suppressed by $(m_e/m_\mu)^2$ for muon beams compared to the case of electron beams.

We simulated both signal and background processes with CalcHEP [111]. Although MadGraph [112] has become a standard tool for fixed-target experiments as well as collider experiments, it has known issues with collinear emission processes [87] which give large fluctuations in the computed cross section, while CalcHEP reliably gives sub-percent level Monte Carlo errors [113]. We implemented a custom form factor in CalcHEP for the nuclear target by scaling all

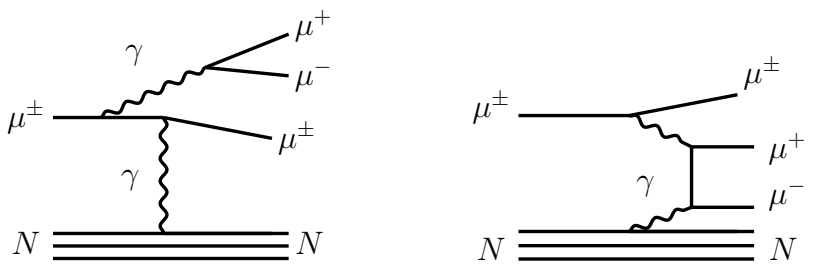

Figure 17. Feynman diagrams representing the dominant Standard Model background processes for dimuon pair production in muon-nucleus scattering. The radiative diagram (left) also has a contribution from final-state radiation, but radiation off the nucleus is suppressed since $m_N \gg m_\mu$.

the cross sections by a dipole form factor [114]:

$$G_2(t) = G_2^{\rm el}(t) + G_2^{\rm inel}(t) \ , \tag{18}$$

where the elastic contribution is

$$G_2^{\rm el}(t) = Z^2 \left(\frac{a^2 t}{1+a^2 t}\right)^2 \left(\frac{1}{1+t/d}\right)^2 , \tag{19}$$

and the inelastic contribution can be written as [6]

$$G_2^{\rm inel}(t) = Z \left(\frac{a'^2 t}{1+a'^2 t}\right)^2 W_2(t) \ , \quad W_2(t) = \left[\frac{1+\tau(\mu_p^2-1)}{(1+t/t_0)^4}\right] \ . \tag{20}$$

The parameters of this form factor model are

$$a = \frac{113 Z^{-1/3}}{m_e}, \quad a' = \frac{773 Z^{-2/3}}{m_e}, \quad d = 0.164\ \text{GeV}^2 A^{-2/3}, \quad t_0 = 0.71\ \text{GeV}^2 \ , \tag{21}$$

where $\mu_p = 2.79$, $\tau = t/(4m_p^2)$, Z and A are the atomic number and mass number of N, m_p and m_e are the proton and electron masses, and $t = -(p'_N - p_N)^2$ is the squared 4-momentum transfer to the nucleus (initial 4-momentum p_N, final momentum p'_N). For the signal events, in order to avoid any issues arising from the narrow width of the S, we generated on-shell S events $\mu^\pm N \to \mu^\pm N S$, decayed the S isotropically in its rest frame, and boosted back to the lab frame.

[6] Note that a typo in the literature has been propagated to various references from Ref. [114], now corrected here: the $W_2(t)$ function should not be squared.

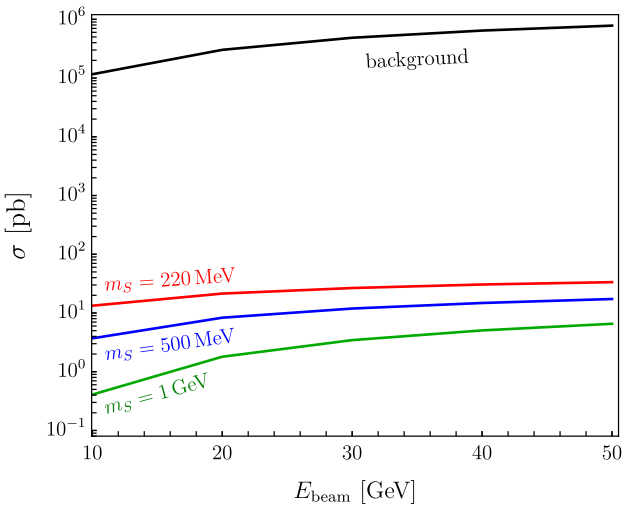

Figure 18. Cross sections as a function of beam energy for background and signal for various S masses, with g_S set to the minimum value consistent with the measured Δa_μ to within 2σ.

In order to confirm that the form factor was correctly implemented in CalcHEP, we reproduced the results of Fig. 8 of Ref. [113] for the process $\mu N \to \mu N S$ with N an aluminum nucleus. The number of signal events or background events $N_{S,B}$ is given by

$$N_{S,B} = \sigma_{S,B}^{(\mathrm{acc})}\, n_T\, \ell_T \times \mathrm{MOT}, \tag{22}$$

where $\sigma^{(\mathrm{acc})}$ is the accepted cross section given any cuts applied (described further below), n_T and ℓ_T are the number density of nuclei in the target and the effective target length, respectively, and MOT is the number of muons on target.

2.3.2 Kinematics and Cuts

As a starting point, based on the setup envisioned in Section 2.4 below, we will consider the example of a μ^- beam on an iron target with $Z = 26$ and $A = 56$, $n_T = 8.5 \times 10^{22}/\mathrm{cm}^3$, and $\ell_T = 100$ cm. Figure 18 shows, as a function of the beam energy, the cross section for $\mu^-\mu^-\mu^+$ events from the QED background process and signal process for different scalar masses, fixing the coupling g_S for each mass to be the minimum coupling required to resolve Δa_μ to within 2σ. The cross section is largely insensitive to the beam energy, but for all energies and masses, the QED background is several orders of magnitude larger, requiring additional cuts to render the signal visible. The most obvious such cut is an invariant mass cut; this turns out to be most efficient when

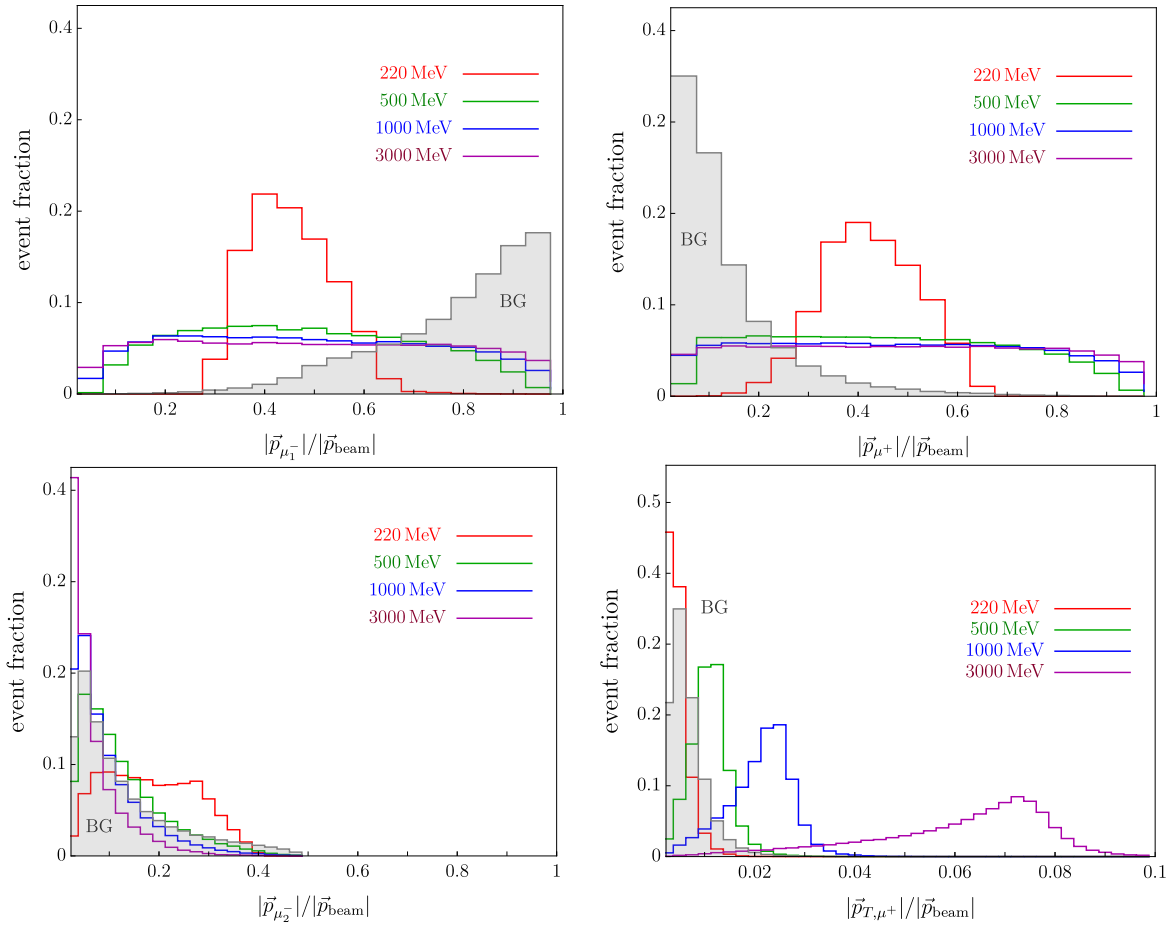

Figure 19. Kinematic distributions for $\mu^-\mu^-\mu^+$ events produced by SM QED background (shaded gray, labeled "BG") and signals of various scalar masses with a 20 GeV μ^- beam. Relative distributions (normalized to unity for illustration) are shown for the fraction of the beam momentum carried by the hardest final-state μ^-, labeled μ_1^- **(top left)**, μ^+ **(top right)**, and softer μ^-, labeled μ_2^- **(bottom left)**. The momentum of the leading μ^- transverse to the beam is also shown **(bottom left)**. For heavy S, the p_T and angular separation of the hardest muons peak at larger values than the background.

performed on the opposite-sign muon pair with the *hardest* μ^-, which we can justify by comparing the signal and background kinematics.

Figure 19 shows the kinematic distributions of both signal and background events for a 20 GeV μ^- beam, without applying any cuts. The principal difference between the signal and background processes is that the cross section for S production is peaked when S takes nearly all the beam energy [45, 115], at least for $m_S > m_\mu$ which is the case whenever S can decay visibly. As a result, the S is highly boosted, and thus the harder μ^- dominantly originates from S decay, with the momenta of the μ^+ and harder μ^- both peaked at $|\vec{p}_{\rm beam}|/2$. By contrast, the soft singularity of the massless photon in QED means that the μ^+ momentum distribution is peaked at zero, while

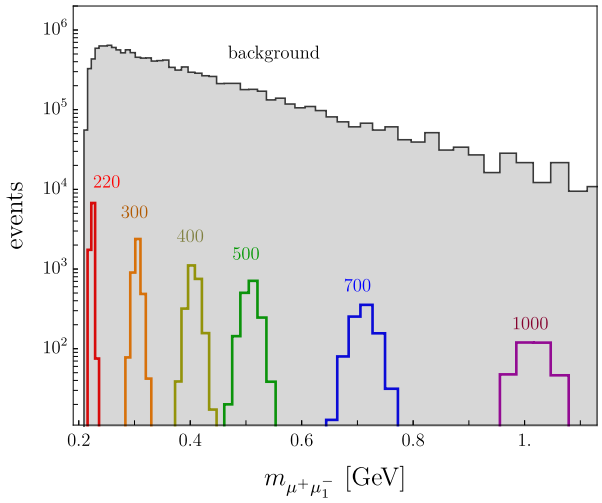

Figure 20. The $\mu^+\mu^-$ invariant mass distribution is shown for background and signal for several scalar masses (in MeV units), in events produced from a 20 GeV μ^- beam. A 5×10^{13} MoT sample (corresponding to approximately 1 year of running at SpinQuest, see Sec. 2.4 is shown with g_S set to the minimum values which resolve Δa_μ for each m_S value, as in Fig. 18. Here, the invariant mass is formed from the harder μ^- (labeled μ_1^-) in the event and 3-momenta are smeared to produce a 15% mass resolution. Each muon is required to satisfy $p_z > 5$ GeV.

the harder μ^- is typically the beam muon which still carries a large fraction of the original beam momentum.[7] Both of these features are apparent in the top panel of Fig. 19. Combined with the fact that the high boost afforded by fixed-target experiments makes almost all events purely in the forward direction, a cut on p_z for both the μ^+ and hardest μ^- allows an efficient separation of signal and background. We also note that for both signal and background, the recoil μ^- is very soft and often carries less than 25% of the beam momentum (bottom left). In practice, the soft μ^- may not even be observable if it curves away from the tracking stations, so absent any issues from combinatorial background due to the fact that the beam contains both μ^+ and μ^-, we can define the signal as two muons of opposite sign originating from the same vertex. We argue in Sec. 2.4 below that the combinatorial background is likely negligible compared to the irreducible background for $m_S \lesssim 1$ GeV. Finally, the bottom-right panel of Fig. 19 shows the transverse momentum distributions of the μ^+. The collinear singularity of QED (regulated by the muon mass) implies that the background peaks at small p_T, while for heavy S, the p_T is peaked at larger

[7] We note that the same distinctions between the QED and muonphilic particle kinematics hold for massive vectors, pseudoscalars, and axial vectors [116], and thus the same search strategies we propose should apply to those models too.

values.[8]

Fig. 20 shows the invariant mass spectrum of the μ^+ paired with the hardest μ^-, for both signal and background, after applying the p_z cut. As in Fig. 18, we have normalized the signal to the values of g_S consistent with Δa_μ. Keeping in mind the possible degradation of the mass resolution with multiple scattering in the target, we take as an example a 15% invariant mass cut, which means we require that the harder muon pair has an invariant mass within $m_S \pm r$, where $2r/m_S = 0.15$. In Fig. 20 we have applied a uniform Gaussian smearing to each of the 3-momentum components to qualitatively simulate a 15% mass resolution. For example, taking $m_S = 500$ MeV and implementing this invariant mass cut, the background is further reduced by a factor of $\sim$8 and we keep $\sim$90% of the signal (the remaining 10% are events where the softer μ^- originates from S decay). The requirements on achieving such an invariant mass resolution are discussed at length in Section 2.4.

Figure 21 (left) shows the effect on the sensitivity $N_S/\sqrt{N_B}$ from applying one or both cuts for a 20 GeV beam and 10^{15} MOT, using the same signal normalization that resolves Δa_μ. We show the effect of an invariant mass cut alone (green curve) to emphasize the fact that a standard "bump hunt" strategy is much less efficient for our particular signal process without applying the p_z cut. After applying both the invariant mass cut and the p_z cut, the signal efficiency is $\sim 83\%$ for low S masses near 220 MeV and $\sim 50\%$ for higher S masses near 1 GeV. Due to the clear separation of signal and background kinematics, the combined effect of both cuts can improve the sensitivity by an order of magnitude for $m_S \lesssim 1$ GeV, and even render the experiment free of irreducible QED background for larger masses. To illustrate the experimental requirements necessary to achieve full coverage of the $g-2$ parameter space for a given m_S, in Figure 21 (right) we plot the muon flux required to achieve approximate 3σ discovery sensitivity, $N_S/\sqrt{N_B} = 3$, as a function of m_S. With 15% invariant mass resolution and a 100 cm iron target region, we can probe the full parameter space that resolves the anomaly up to 1 GeV with $\simeq 3 \times 10^{14}$ MOT. For larger m_S at this luminosity, the QED background becomes subdominant to the combinatorial background in a realistic experimental implementation, as we will describe in Sec. 2.4.3 below.

2.4 Potential Reach at the SpinQuest Experiment

To make the discussion in Sec. 2.3 concrete, and to illustrate the experimental considerations

[8] A cut on p_T would only efficiently separate signal from background at masses $m_S \gtrsim 1$ GeV, where (as we illustrate below) the irreducible QED background is already subdominant to the expected combinatorial background after the invariant mass and p_z cuts.

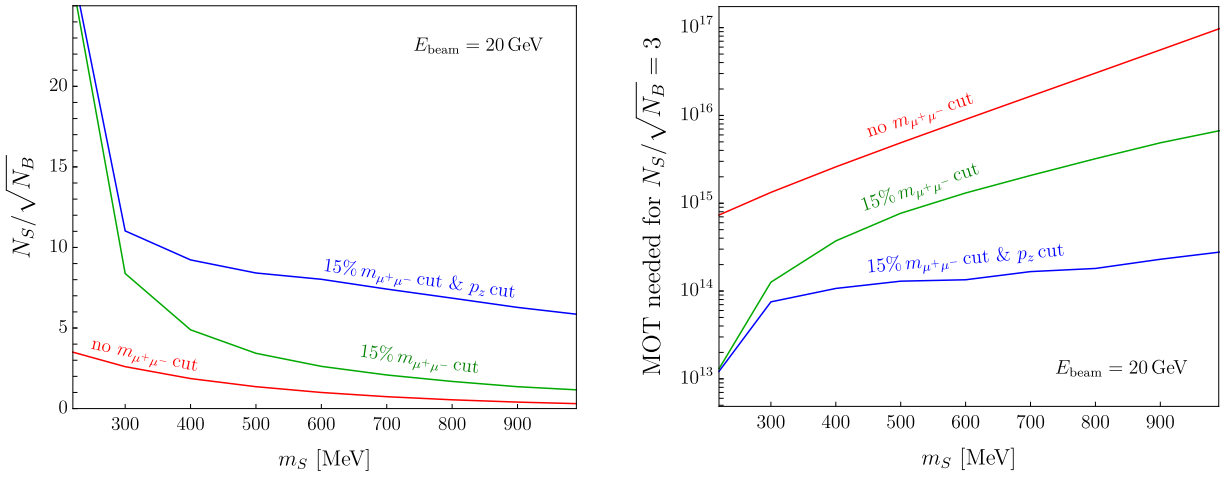

Figure 21. **Left:** Sensitivity as a function of scalar mass m_S for various cuts, setting g_S to the minimum value which resolves Δa_μ at each m_S and assuming 10^{15} MOT and $\ell_T = 100$ cm. **Right:** MOT required to reach $N_S/\sqrt{N_B} = 3$ as a function of m_S, for g_S as in the left panel. Both panels assume a muon beam energy of 20 GeV.

relevant for a practical implementation of our proposal, we imagine a setup similar to the SpinQuest experiment at Fermilab [117], where a 120 GeV proton beam impinges on a magnetized steel beam dump ($Z = 26$) many radiation lengths thick. Target interactions produce a spectrum of lower-energy secondary muons, which continue through the target and can produce S through bremsstrahlung during any scattering process, as described above. To identify the vertex corresponding to S decay, we consider only signal events which occur in the last portion of the 5 m beam dump, of length $\ell_T \simeq 100$ cm. As the simulated muon spectrum at proton beam dump experiments is still to be fully validated [103], we leave a detailed sensitivity calculation for future work and consider a monochromatic 20 GeV muon beam. We note, however, that both signal and background cross sections are largely independent of the beam energy as shown in Fig. 18, so treating the muons as monochromatic is a decent approximation up to experimental acceptance effects.

2.4.1 Muon Beam Properties

The signal sensitivity and estimated rate of backgrounds depend on the expected secondary muon momentum spectrum, following their traversal of the bulk of the beam dump. To facilitate these studies, we perform a toy simulation of muons passing through the dump, based on their initial spectra obtained from a full Geant4 simulation of the SpinQuest experimental setup [117]. Muons are randomly assigned to a position at the beginning of the dump based on the expected

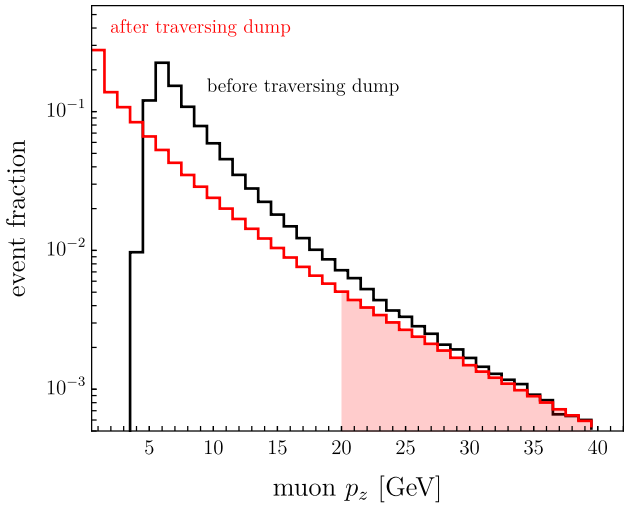

Figure 22. Secondary muon momentum spectrum from a primary 120 GeV proton beam. The part of the spectrum consisting of muons which exit the dump with $p_z > 20$ GeV, which we treat as our muon beam, is shaded in red.

beam profile and propagated to the end of the dump, accounting for (a) initial muon momenta transverse to the beam, (b) small and wide-angle scattering within the dump, and (c) the effect of a 1.9 T magnetic field. The momenta of muons before and after the dump are shown in Figure 22. The relative composition of muons is approximately 55% (45%) for positively (negatively) charged muons.

To estimate the time required to achieve a given number of MOT at SpinQuest, we can use the finding of Ref. [117] that 1.9 muons per RF bucket reach the first tracking station. Combining this with our simulation model, we find that 0.8 such muons have $p_z > 20$ GeV. Each 4 sec spill of 53 MHz buckets thus yields 2×10^8 MoT with energy above 20 GeV, or 3×10^{11} MoT in 24 hours of running. Assuming that the accelerator complex can deliver beam with an average duty factor of 50% over the full year, this amounts to 5×10^{13} MOT accumulated per year of operation in the nominal configuration, such that our target luminosity for m_S up to 1 GeV can be in principle achieved in 6 years of running.

2.4.2 Reconstruction efficiencies and mass resolution

The capability to efficiently trigger and accurately reconstruct di-muons will drive the sensitivity of muon spectrometer experiment to scalar decays. The forward region of SpinQuest is instrumented with an emphasis on detecting muons from the decay of a high mass virtual photon or meson produced through target interactions. Consequently, we consider only muons with

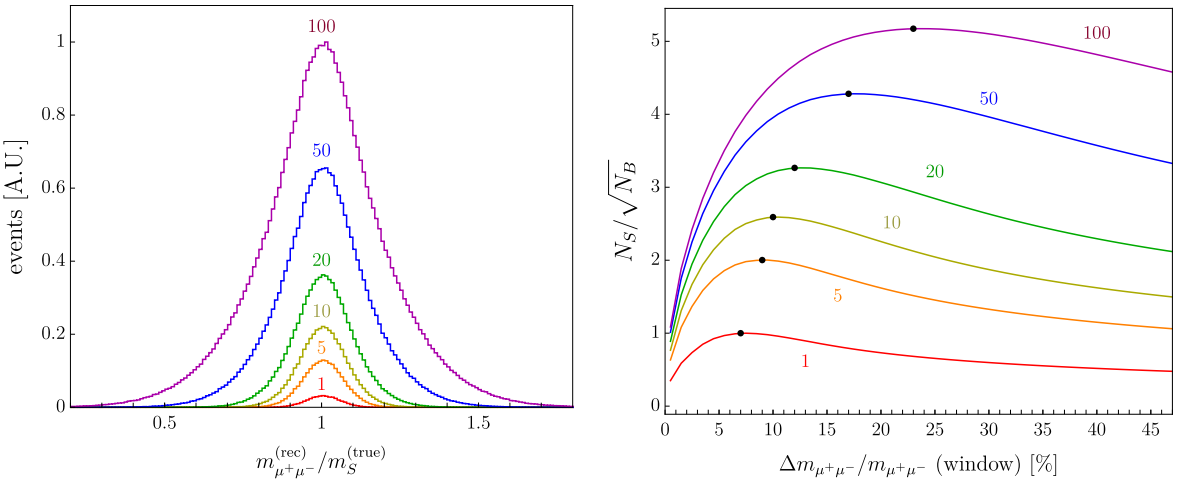

Figure 23. **Left:** Reconstructed di-muon mass is shown for events with $m_S = 1\,\mathrm{GeV}$ produced in the last portion of the dump, taking into account both multiple scattering effects and a 5% experimental resolution. Mass distributions are shown for several target thicknesses in units of X_0, with larger decay volumes increasing yields but degrading the mass resolution. Note that our fiducial value of $\ell_T = 100\,\mathrm{cm}$ corresponds to about $57X_0$ in steel. **Right:** The corresponding signal sensitivities ($N_S/\sqrt{N_B}$) as a function of the mass window for each of the same thicknesses. The optimal window size is indicated with a black dot. A flat background in the neighborhood of the scalar mass is assumed.

momenta above 5 GeV in the following analysis, assuming the trajectories of particles with lower momenta cannot be accurately measured. While imperfect acceptance can lead to a small additional inefficiency, the effect has a complex dependence on the exact geometry of the detector and details of the reconstruction algorithms, which we do not attempt to reproduce in this work. We note that despite the small loss in acceptance, a minimum requirement on the muon momentum removes background processes far more efficiently than signal, due to the soft sub-leading muon expected from radiative processes, as seen in Figures 19 and 21.

Another key driver of the search sensitivity is the achievable di-muon mass resolution, which should be small to reduce the level of backgrounds, particularly for small scalar masses. We assume a 5% mass intrinsic experimental resolution, in accordance with the studies of Ref. [117] based on decays between the dump and tracking stations. However, because the scalar decays promptly after being produced within the last fraction of the dump, this mass resolution will generally be further degraded by multiple scattering. Using the simulation described above, we propagated this effect to the mass resolution of $S \to \mu^+\mu^-$ pairs as a function of the target length traversed. This leads to a trade-off, illustrated in Figure 23, where scalars produced earlier in the dump can be selected to enhance the overall signal rate, at the price of poorer mass resolution. In this study,

we take an effective target thickness of $\ell_T = 100$ cm as a representative choice, corresponding to an approximate 15% total mass resolution including multiple-scattering effects. Finally, we need to be able to determine if the candidate dimuon mass pair originated from the target region. The displaced vertex resolution in the z axis determined from simulation studies in Ref. [117] is approximately 10 cm. This would be degraded by multiple scattering effects in the last $\sim$1 m of the dump but would still be sufficient to determine if the dimuon pair is in the target region with reasonable efficiency.

2.4.3 Combinatorial Background

In addition to the irreducible QED background discussed in Sec. 2.3, an additional experimental background can arise through combinations of independently-produced muons in the dump that conspire to produce an apparent di-muon vertex. This effect, unlike the irreducible radiative and trident backgrounds, is sensitive to details of the experimental setup such as vertex resolution, which we take as $\sigma_{x,y} = 1$ cm in line with SpinQuest's projected capabilities [117].

The chance of finding a $\mu^-\mu^+$ vertex from uncorrelated muons is generally low because the magnetic field separates the oppositely-charged particles produced early in the dump. Here the "wrong-sign" muon can only form a potential $\mu^-\mu^+$ vertex if the combined effects of initial beam momenta and scattering in the dump are enough to compensate for the strong magnetic field bending its path in the opposing direction. Figure 24 (left) shows the distribution of μ^+ and μ^- positions in the bending plane at the end of the dump. As $\mu^+ S \to \mu^+\mu^-\mu^+$ signal events would result from the production of an S off of a μ^+ late in the dump, signal vertices will generally be concentrated at large x (signal efficiency is $\epsilon_{\text{sig}} = 95\%$ for $x_{\text{dump}} > 10$ cm), and vice-versa for $\mu^- S$ events. Our simulation predicts a contamination rate of wrong-sign muons of 3×10^{-4} after selecting events displaced by 10 cm in the bending plane.

The rate of combinatorial $\mu^+\mu^-$ pairs can be computed from the rate of wrong-sign scatters into the signal-enriched region together with the expected position spread and vertex resolution. The experimental resolution motivates a requirement that the reconstructed coordinates transverse to the beams agree to within 2 cm, which occurs randomly in only 0.6% of background events. All together, this toy simulation predicts a rate of 5×10^{-6} candidate pairs of 5 GeV muons per each time sample. The distribution of masses from these background events are shown in Figure 24 (right) for a sample of 5×10^{13} MoT. Considering a 15% window in di-muon mass, we expect

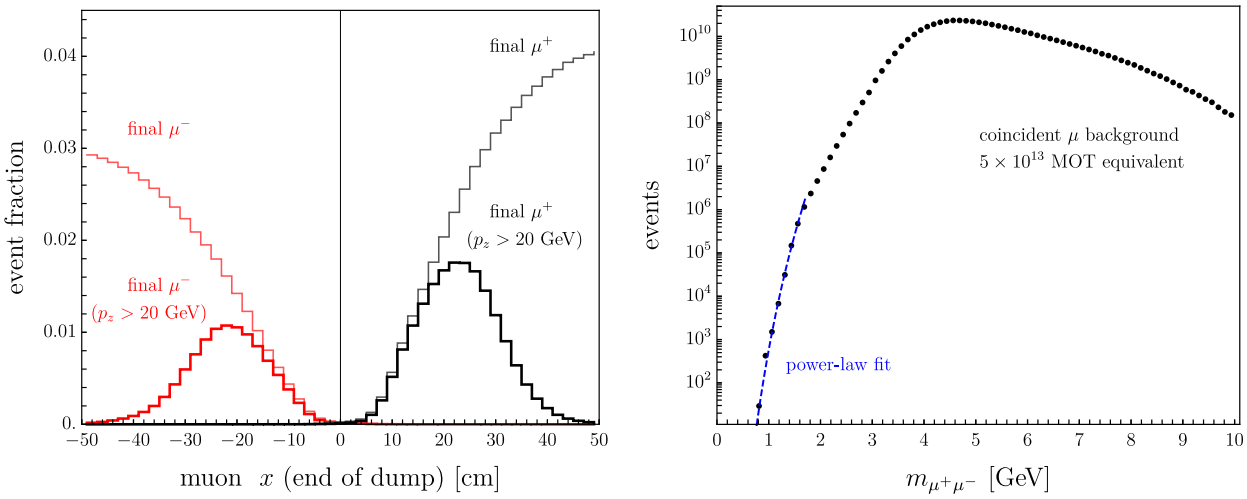

Figure 24. **Left:** The position of positively and negatively charged muons in the bending plane, after traversing the beam dump. **Right:** The $m_{\mu^-\mu^+}$ mass distribution for $\mu^-\mu^+$ pairs passing a common vertex requirement.

this sample to contain about 100 combinatorial background events at 1 GeV. This background grows rapidly with mass (a power law fit yields a dependence $\sim m^{15}_{\mu^+\mu^-}$), with 7×10^6 events for $m_S = 2\,\text{GeV}$, illustrating the need for novel experimental strategies to probe the high-mass region. Less than one combinatorial background event is predicted below 800 MeV for this luminosity.

2.4.4 Estimated Sensitivity at SpinQuest

Fig. 25 shows our estimated sensitivity to the $g-2$ parameter space at a SpinQuest-like experiment, including the invariant mass resolution, target length, and combinatorial effects discussed above. Our limit on g_S is defined[9] by solving $N_S/\sqrt{N_B} = 3$ for g_S at each mass value m_S. As anticipated, with 3×10^{14} MOT, we can cover the full Δa_μ parameter space up to 2σ for $m_S < 1$ GeV. We estimated the combinatorial background by scaling the results of Fig. 24, but the extremely steep scaling with mass results in a sharp degradation of the sensitivity of our proposed experiment at around 1 GeV. We have verified that the sensitivity is largely unchanged for different beam energies; the same muon flux with 30 GeV or 40 GeV muons can probe the same parameter space, which suggests that our estimates are likely robust to the precise secondary muon spectrum at proton beam dump experiments.

[9] While systematic uncertainties are neglected in this analysis, they are expected to have negligible impact due to the robust background estimate from $m_{\mu^+\mu^-}$ side-band fits. Background sculpting due to kinematic thresholds near $m_S \sim 2m_\mu$ may require careful treatment, which we will leave for future work.

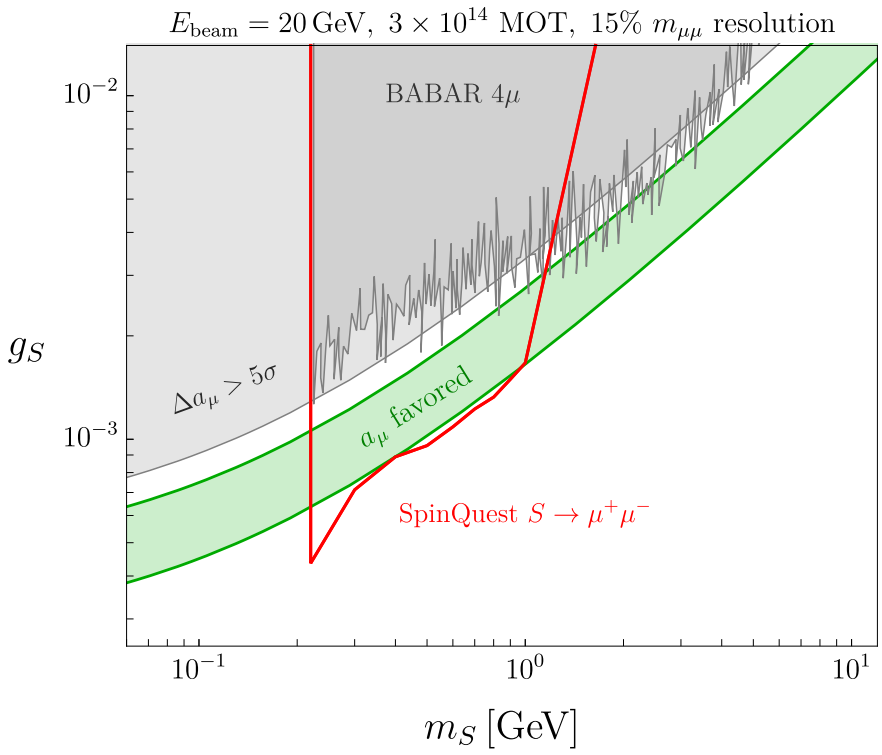

Figure 25. Sensitivity to the $(g-2)_\mu$ parameter space for a SpinQuest-like spectrometer at a proton beam dump experiments. We assume the same fiducial parameters as discussed in Sec. 2.3, namely 20 GeV beam energy, $\ell_T = 100$ cm, 15% invariant mass resolution, and $p_T > 5$ GeV on the hardest $\mu^+\mu^-$ pair. We include both irreducible QED background and reducible combinatorial background. A muon flux of 3×10^{14} MOT, corresponding to approximately 6 years of running at nominal SpinQuest luminosity, can fully cover the preferred region for Δa_μ to 2σ for $m_S < 1$ GeV.

2.5 Conclusions and Outlook

In this paper, we have introduced a new search strategy for muonphilic particles using proton beam dump spectrometers. As primary protons impinge on a fixed target, they produce a beam of secondary muons whose interactions with the target material can produce new muonphilic particles. If these new states decay visibly to dimuons, the daughter particles emerge from the target region and their combined invariant mass can be reconstructed using a downstream tracking station. This search strategy can be parasitically executed at the Fermilab SpinQuest experiment and can achieve unprecedented sensitivity to new muonphilic particles with 3×10^{14} muons on target and appropriate analysis cuts. Note that this experiment has been built and has started collecting data as of Summer 2024.

While the search strategy we outlines here is general for any muonphilic particles, our discus-

sion has been framed around particles that resolve the longstanding muon $g-2$ anomaly, which is arguably the longest-standing disagreement between SM predictions and experimental measurements. Assuming the theoretical prediction of $g-2$ remains unchanged with the inclusion of recent lattice QCD results, all possible beyond-the-SM solutions should be tested comprehensively. Here we have found that a proton beam dump spectrometer can cover parameter space in a highly complementary region to missing-momentum experiments such as M^3 [45], which can probe $m_S < 2m_\mu$, and NA64-μ [89], which can fully probe the parameter space for an $L_\mu - L_\tau$ gauge boson. Future B-factories such as Belle-II can also cover the visible scalar decay parameter space [59], but the full luminosity may be a decade away. In the intervening years, our analysis shows that proton beam dump experiments such as SpinQuest can potentially discover the new physics responsible for Δa_μ at masses below 1 GeV in a reasonable $\sim$6 years of running, and likewise has sensitivity to other muonphilic particles in this mass range.

More detailed analyses by experimental collaborations are required to produced a more refined sensitivity projection, including in-situ measurements of the SpinQuest muon spectrum and several detector effects such as detector reconstruction efficiency and trigger efficiency. Additional and more sophisticated analysis techniques include multivariate kinematic selections and an optimization of the beam dump target region. Nonetheless, even with such considerations, the sensitivity of the SpinQuest experiment can benefit even more significantly from beamline considerations and detector improvements. Increasing the duty factor of the SpinQuest experiment, which currently takes data for 4 s out of 1 minute, could increase the expected MOT in a year by approximately an order of magnitude – of course at the expense of other experiments. Furthermore, the RF bucket-to-bucket intensity is limited by the efficiency of the slow extraction of the beam from the main injector and the occupancy of the detector. By improving the proton beam intensity and using more modern highly-granular tracking detectors, especially at the front of the spectrometer, the sensitivity of the experiment could also be drastically improved.

CHAPTER 3: SEARCHES FOR EXOTIC PARTICLES AT THE DUNE NEAR DETECTOR

The work in this section is based on "Exotic Particles at the DUNE Near Detector from Charged Pion Scattering", published as an article in Physical Review D Volume 110, Article 095029 on 26 November 2024. The work was authored by myself, Yonatan Kahn, and Rachel Nguyen. During the writing and publication of this paper, the three of us were affiliated with the Department of Physics at the University of Illinois at Urbana-Champaign, as well as the Illinois Center for Advanced Studies of the Universe here at UIUC. In this paper, we analyzed the sensitivity of the DUNE Near Detector complex to millicharged particles (MCPs) and heavy axion-like particles (ALPs) with low-energy couplings to gluons. Using the framework of chiral perturbation theory, we demonstrated regimes of parameter space where the charged pion production channel dominates over previously-considered production mechanisms for both MCPs and ALPs, thereby improving the sensitivity of DUNE to these new particles compared to previous studies. The bulk of my work that contributed to this paper was focused on the MCP search. I simulated the events with MadGraph, I talked to the authors of a reference we followed closely to ensure we understood the search strategy and definition of the signal well, I implemented various cuts by hand, and I calculated and plotted the MCP search sensitivity.]

In the search for light, weakly-coupled extensions to the Standard Model, fixed-target experiments have proved to be an extremely useful approach (see Ref. [118] and references therein). Trading energy for luminosity compared to colliding-beam experiments, these $\mathcal{O}(10-100)$ GeV proton and electron beams can deliver Avogadro's number of beam particles on target, allowing us to search for a wide variety of low-mass particles beyond the Standard Model (BSM) with small production cross sections. In particular, proton beams produce copious amounts of secondary particles, including neutral mesons and muons, which can generate BSM particles through rare decays, mixing, or bremsstrahlung. Recent studies [3, 119–121] have demonstrated that as part of the upcoming DUNE program [122–125], the neutral meson channel yields excellent sensitivity to millicharged particles (MCPs) (denoted χ) or heavy axion-like particles (ALPs) coupled to QCD (denoted a).

However, previous analyses have overlooked an additional production mechanism for these new particles, namely charged pion scattering. At DUNE, the 120 GeV proton beam creates $\sim 6.5\,\pi^{\pm}$ for every proton on target (POT) compared to $\sim 3.5\,\pi^0$ per POT [101]. As a charged pion traverses

the target, it can scatter off the target nuclei and produce BSM particles through bremsstrahlung-like processes, which may be described in the framework of chiral perturbation theory as long as the BSM particles have masses well below $4\pi f_\pi \sim 1.2\,\mathrm{GeV}$, where $f_\pi = 93\,\mathrm{MeV}$ is the pion decay constant. This production channel was first considered for dark photon production at DarkQuest [126], where it was shown that charged pion bremsstrahlung generated comparable event rates to the Drell-Yan production channel for GeV-scale dark photons.

In this paper, we follow the logic of Ref. [126] to search for MCPs and heavy ALPs at DUNE through secondary charged pion scattering. The MCP production channel is very similar to the previously-studied dark photon production, involving an off-shell virtual photon generating a $\chi\bar{\chi}$ pair instead of an on-shell dark photon decaying to visible states, and we will show that charged pion scattering does indeed yield modest gains in sensitivity for GeV-scale χ. For heavy axions, QCD-coupled axion production from charged pions dominates the meson mixing production channel except in narrow resonance windows around the neutral meson masses, and is comparable to the expected rate from kaon decays [121]. At high axion masses, charged pion scattering provides additional sensitivity at large couplings all the way up to the regime of validity of chiral perturbation theory.

This paper is organized as follows. In Sec. 3.1, we briefly describe the experimental design of the DUNE Near Detector (ND) Complex located at Fermilab, including additional detector components proposed for detecting MCPs. In Sec. 3.1.1, we give an overview of our simulation pipeline that we use to model the production and detection of BSM particles from charged pion scattering. In Sec. 3.2, we describe our MCP model and calculate the additional sensitivity gained at DUNE ND at high MCP masses from charged-pion scattering, compared to the production modes previously considered in Refs. [3, 119]. In Sec. 3.3, we do the same for the heavy axion model of Ref. [120]. We conclude in Sec. 3.4.

3.1 DUNE as a BSM Detector

DUNE is primarily designed to produce neutrinos, but its 120 GeV proton beam can also be used as a source for BSM particles, which can be detected at the DUNE Near Detector (ND) Complex. The LBNF-DUNE beam is a 120 GeV proton beam incident on a graphite target, with an expected luminosity of 1.47×10^{21} protons on target (POT) each year, and a total of 1.47×10^{22}

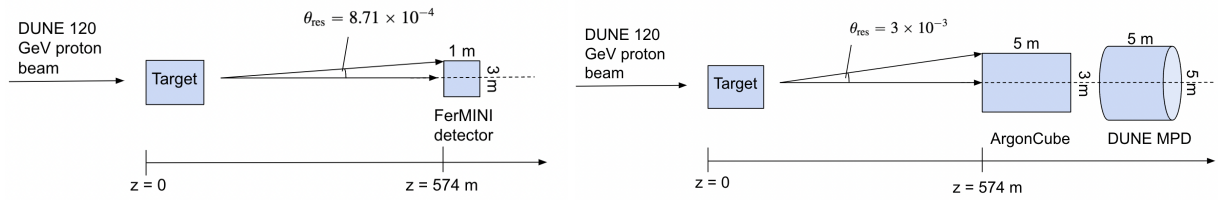

Figure 26. A simple schematic of the experimental setup for MCP (left) and ALP (right) searches. Note that the DUNE ND Complex contains ArgonCube and the Multi-Purpose Detector, which are the detectors used for our ALP search. Additionally, for the sake of our MCP search, one could imagine installing scintillator arrays and PMTs inside the Dune ND Complex inspired by the FerMINI proposal [3], described in more detail in Sec. 3.2.2.

POT over a 10-year run [122–125, 127]. The ND Complex is located 574 meters down the beam line from the target, and will contain two argon detectors: ArgonCube, a 50 ton liquid argon time projection chamber (LArTPC), and the Multi-Purpose Detector (MPD), a gaseous argon time projection chamber designed to detect muons and other particles that are not stopped in ArgonCube. ArgonCube is 7 m in width, 3 m in height, and 5 m in length along the beamline. The MPD is situated directly downstream from ArgonCube and has a cylindrical volume that is 5 m in diameter and 5 m in height. Both ArgonCube and the MPD are sensitive to diphoton and hadron signatures by reconstructing the invariant mass and direction of the event, and these will be the primary decay modes of the heavy axions we consider.

To search for MCPs, one may also imagine adding scintillator arrays in the Near Detector Complex, as was proposed in [3]. To make a direct comparison with Ref. [3], we consider the same setup: three stacks of scintillator arrays in a 1 m $\times$ 1 m detection area, coupled to photomultiplier tubes (PMTs) to detect charged particles. The distinctive signature of a MCP would be soft ionizations with yield below that expected from a minimum-ionizing particle of charge $|Q| = 1$.

In Fig. 26, we show a schematic of the DUNE ND complex and the detectors of interest. Relevant background processes for each of our BSM signatures will be discussed in Secs. 3.2 and 3.3.

3.1.1 Simulation Details

In order to calculate the production and detection of new particles at DUNE ND, we perform a Monte Carlo simulation. In this section, we outline the general simulation setup that generates BSM events from charged pion-nucleon scattering and estimates the detection efficiency within

the detectors described above. In Secs. 3.2 and 3.3, we will specify to MCPs and heavy axions, respectively.

3.1.2 Charged Pion Production

When the 120 GeV DUNE proton beam collides with the target, it produces a multitude of SM mesons, including π, η, η' as well as heavier mesons. While mesons are produced during several nuclear interaction lengths, in order to provide a conservative estimate of the reach and to make a direct comparison with previous studies [3, 101, 119, 120, 126], we only consider the mesons produced within the first proton interaction length of the target, which amounts to assuming that every incident proton scatters once.

To estimate the SM meson spectrum, we adapt the `Pythia 8.2` simulation of a 120 GeV proton incident on a proton target from [101], rescaled to match the nucleon composition of the DUNE graphite target. The number of mesons produced per 120 GeV proton is $N_{\pi^0} \sim 3.5$, $N_{\pi^\pm} \sim 6.5$, $N_\eta \sim 0.40$, and $N_{\eta'} \sim 0.04$.[10] The total number of charged pions can be written as an integral over the energy spectrum of charged pions produced,

$$N_\pi = \int dE_\pi \frac{dN_\pi}{dE_\pi}. \tag{23}$$

For simplicity, we assume $N_{\pi^+} = N_{\pi^-} = N_{\pi^\pm}/2$. The energy distribution of the charged pions is shown in Figure 27. The majority of charged pions are produced with low energy, $E_\pi \sim 1 - 10$ GeV. However, the high-energy tail of the distribution, combined with the large quantity of charged pions produced in the target (of order 10^{23} over the lifetime of DUNE) can produce a detectable number of BSM particles that are boosted enough to reach the DUNE ND.

3.1.3 BSM Particle Flux

Exotic particles may be created from pion-nucleon scattering, $\pi^\pm N \to \pi^\pm N' X$, where we will later consider $X = a$ or $\chi\bar{\chi}$.[11] For sufficiently small momentum transfers, we can describe the scattering process using chiral perturbation theory. In this approximation, the pions would scatter incoherently off the nucleons contained within the target's nuclei. Typically, staying within the regime of validity of chiral perturbation theory limits the mass of the BSM particles to $\lesssim 1$ GeV,

[10] Treating the target as free protons is clearly a very rough approximation, however it allows us to make a direct comparison with previous studies in the literature [3, 101]. Ref. [3] (also Ref. [120]) found $N_{\pi^0} \sim 4.5$ (2.9), $N_\eta \sim 0.5$ (0.33) for the DUNE target, which only differ by $\sim 20\%$ from our values.

[11] Axions can also be produced through the subdominant pion conversion $\pi N \to aN$.

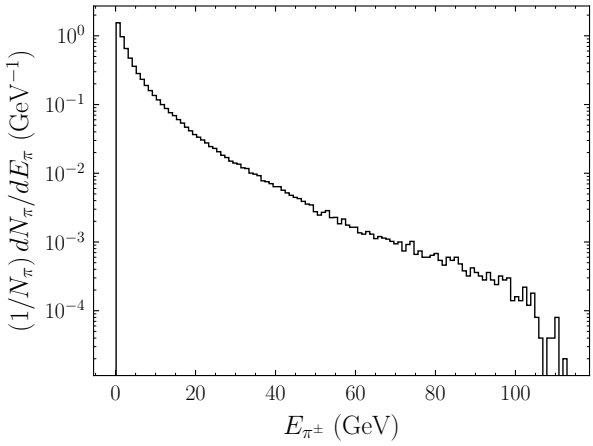

Figure 27. The energy distribution of charged pions from a `Pythia 8.2` simulation of a 120 GeV proton hitting a proton target. The pions were created within the first thin slice of the target. We rescale this distribution to match the nucleon content of the DUNE graphite target. This distribution was adapted from the analysis of [101].

though for MCPs in particular, the production rate remains significant up to about 5 GeV. A full treatment of this process would require matching onto the perturbative QCD regime, which is beyond the scope of this work but has significant synergy with the DUNE neutrino cross section program [128].

The total cross section for the scattering process in the lab frame is

$$\sigma = \frac{1}{4E_\pi M_n} \int d\Pi \left(Z\langle |M^{(p)}|^2 \rangle + (A-Z)\langle |M^{(n)}|^2 \rangle \right), \tag{24}$$

where E_π is the energy of the incoming pion, M_n is the nucleon mass, $d\Pi$ is the 3- or 4-body phase space measure, $Z = 6$ is the atomic number of the graphite target, $A = 12$ is the mass number, and $\langle |M^{(p,n)}|^2 \rangle$ are the spin-averaged matrix elements for $\pi^\pm$ scattering off protons or neutrons, respectively.[12] For both axion and MCP production, we perform a Monte Carlo integral of the cross section using `MadGraph5` [130] to generate events.For both models, we compute the relevant matrix elements using chiral perturbation theory, where the Lagrangian for the MCP model is given by Eqs. (33)–(34), and for the axion model by Eq. (39); further details of the calculations for each model are given in Sec. 3.2 and Sec. 3.3 below. To remain within the regime of validity

[12] There is a subtlety in treating the target mass in the cross section calculation. Here, we have treated the pions as scattering off individual nucleons, but the full center-of-mass energy (and hence the kinematic threshold) of the scattering process involves the nucleus mass. To treat the target mass properly, we should apply a target mass correction [129], but in this work we have chosen to take the conservative approach (with the lowest possible kinematic threshold) and use the nucleon mass only.

of leading-order chiral perturbation theory, we implement momentum cutoffs,

$$t_n \equiv -(p'_n - p_n)^2 \leq (4\pi f_\pi)^2 \tag{25}$$

$$t_\pi \equiv -(p'_\pi - p_\pi)^2 \leq (4\pi f_\pi)^2 \tag{26}$$

as phase space cuts, where p_n (p_π) and p'_n (p'_π) are the initial and final nucleon (pion) 4-vectors, respectively. Because MCP production proceeds through a virtual photon, low-mass MCPs will be dominantly produced in the region of small photon virtuality, where the soft singularity of QED makes the Monte Carlo sampling unreliable. Consequently, we will focus only on MCPs above about 400 MeV, and we will use the closely-related study of dark photon production in [126] to argue that there are limited sensitivity gains for lighter MCPs.

The cross section depends on the mass m_X of the BSM particle X, its couplings to the SM generically denoted by κ, and the energy of the incoming pion E_π. We bin the charged pion energies into $i = 1, ..., N_0$ bins, denoted $E_{\pi,i}$, and denote the number of charged pions in each energy bin as $N_{\pi,i}$. We then generate the cross section, $\sigma(E_{\pi,i}, m, \kappa)$, in `MadGraph5` for each energy bin, mass, and coupling value we consider. With the discretized charged pion energy distribution and cross section, we can calculate the number of BSM particles that can be produced at DUNE. Conservatively, we will restrict production to take place within the first interaction length of the pion, which is $l_\pi = 53.30$ cm for a graphite target. Since charged pions typically travel several interaction lengths before decaying and/or capturing on nuclei, additional production is possible deeper in the target, and is likely significant for lower-mass BSM particles; we intend to return to this point in a future in-depth study. For a given mass and coupling, the number of BSM particles produced within one interaction length of the dump is

$$N_X(m_X, \kappa) = L \times \sum_{i=1}^{N_0} N_{\pi,i}\, \sigma_\pi(E_{\pi,i}, m_X, \kappa). \tag{27}$$

Here the luminosity is defined as

$$L = l_\pi n_T \text{POT}, \tag{28}$$

where $n_T \approx 1.0 \times 10^{23}/\text{cm}^{-3}$ is the target number density.

To determine the flux of these particles at DUNE ND, we only consider the particles that are within the angular resolution of the detector of interest at the DUNE ND Hall. We estimate this by

generating $N_{\rm MC} = 10000$ Monte Carlo events for each BSM process and keeping only events that enter the detector geometry. The number of accepted BSM particles is then

$$N_X(m,\kappa)_{\rm acc.} = \frac{L}{N_{\rm MC}} \sum_{\rm events \in geom.} \sum_{i=1}^{N_0} N_{\pi,i}\, \sigma_\pi(E_{\pi,i}, m, \kappa). \tag{29}$$

3.1.4 BSM Particle Detection

Lastly, we must determine if the BSM particle produces a signal event in the detectors. If the BSM particle does not need to decay in order to be detected, as is the case for MCPs, the detection efficiency is simply a function of the detector properties, and we will explore this scenario in Sec. 3.2 below.

For heavy axions, detection relies on observing the decay products. We will only consider the long-lived case, where the BSM particle decays within the detectors at the DUNE ND Hall. The decay region consists of the length of LArTPC and MPD, which is 10 m long. We assume that if scintillators to detect MCPs are installed, these will not greatly alter this geometry. To calculate the number of signal events, we must determine the probability that the particle decays within the DUNE ND decay region,

$$P_{\rm decay} = e^{-L_{\rm det}/(\gamma c\tau)} \times (1 - e^{-L_{\rm dec}/(\gamma c\tau)}), \tag{30}$$

where $L_{\rm det}$ is the distance to the detectors, $L_{\rm dec}$ is the decay region of the detectors, γ is the boost factor of the BSM particle, and τ is its lifetime. The number of signal events is then

$$N_S = P_{\rm decay} \times {\rm BR}(X \to Y) \times N_{X,{\rm acc.}}, \tag{31}$$

where $N_{X,{\rm acc.}}$ is given in Eq. (29) and ${\rm BR}(X \to Y)$ is the branching ratio of the BSM particle X decaying to the SM particles Y which can be detected.

3.2 Millicharged Particles

To make contact with previous studies of MCPs [3, 119], we consider a model where a fermionic MCP χ with mass m_χ has electric charge $Q = e\epsilon$, with $e = \sqrt{4\pi\alpha}$ the QED gauge coupling and

$\epsilon \ll 1$. The MCP therefore couples to the SM photon A_μ through the Lagrangian

$$\mathcal{L}_{\chi\bar{\chi}} \supset \bar{\chi}\left(i\gamma^\mu\partial_\mu - m_\chi\right)\chi + Q\bar{\chi}\gamma^\mu\chi A_\mu + \mathcal{L}_{\pi N} + \mathcal{L}_{\pi NA}, \tag{32}$$

where $\mathcal{L}_{\pi N}$ and $\mathcal{L}_{\pi NA}$ are pion-nucleon and pion-nucleon-photon Lagrangians described by chiral perturbation theory. Following Refs. [126, 131–135], pion-nucleon interactions can be described by

$$\begin{aligned}
\mathcal{L}_{\pi N} \supset\ & \frac{g_A}{2f_\pi}\Big(\bar{p}\gamma^\mu\gamma^5 p\partial_\mu\pi^0 - \bar{n}\gamma^\mu\gamma^5 n\partial_\mu\pi^0 \\
& + \sqrt{2}\bar{p}\gamma^\mu\gamma^5 n\partial_\mu\pi^+ + \sqrt{2}\bar{n}\gamma^\mu\gamma^5 p\partial_\mu\pi^-\Big) \\
& + \frac{i}{4f_\pi^2}\Big(\bar{p}\gamma^\mu p(\pi^+\partial_\mu\pi^- - \pi^-\partial_\mu\pi^+) \\
& + \bar{n}\gamma^\mu n(\pi^-\partial_\mu\pi^+ - \pi^+\partial_\mu\pi^-) \\
& + \sqrt{2}\bar{n}\gamma^\mu p(\pi^-\partial_\mu\pi^0 - \pi^0\partial_\mu\pi^-) \\
& + \sqrt{2}\bar{p}\gamma^\mu n(\pi^0\partial_\mu\pi^+ - \pi^+\partial_\mu\pi^0)\Big)
\end{aligned} \tag{33}$$

where p and n are the proton and neutron spinors, π^0 and $\pi^\pm$ are the pion fields, $f_\pi = 93$ MeV is the pion decay constant, and $g_A = 1.27$ is the axial coupling determined from neutron beta decay [136]. Similarly, the pion-nucleon-photon interactions are described by

$$\begin{aligned}
\mathcal{L}_{\pi NA} \supset\ & ieA_\mu\Big(\pi^-\partial^\mu\pi^+ - \pi^+\partial^\mu\pi^-\Big) \\
& + e^2 A_\mu A^\mu\pi^+\pi^- + e\epsilon A_\mu\bar{p}\gamma^\mu p \\
& + \frac{ieg_A}{\sqrt{2}f_\pi}A_\mu\Big(\bar{n}\gamma^\mu\gamma^5 p\pi^- - \bar{p}\gamma^\mu\gamma^5 n\pi^+\Big) \\
& + \frac{e}{2f_\pi^2}A_\mu\Big(\bar{n}\gamma^\mu n\pi^+\pi^- - \bar{p}\gamma^\mu p\pi^+\pi^- \\
& + \frac{1}{\sqrt{2}}\bar{n}\gamma^\mu p\pi^0\pi^- + \frac{1}{\sqrt{2}}\bar{p}\gamma^\mu n\pi^0\pi^+\Big).
\end{aligned} \tag{34}$$

Our signal process is the scattering process $\pi^\pm N \to \pi^\pm N\chi\bar{\chi}$, where the $\chi\bar{\chi}$ pair is produced from a virtual photon radiated by the charged pion or the proton. The leading-order Feynman diagrams for MCP production from $\pi^- p$ scattering are shown in Fig. 28. Note that in order to preserve the Ward identity, we must sum over all diagrams and cannot claim that a subset are parametrically

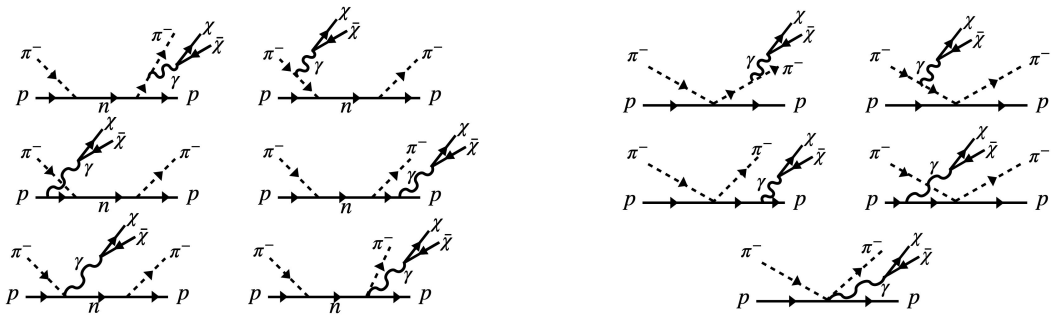

Figure 28. Leading-order Feynman diagrams for MCP production through π^--proton scattering. A similar set of t-channel diagrams exists for π^+-proton scattering. Similarly, corresponding s- and t-channel diagrams also exist for π^--neutron scattering.

suppressed compared to the rest; the usual intuition that bremsstrahlung is suppressed for radiation by heavy particles does not hold in the non-renormalizable chiral Lagrangian [126]. A similar set of diagrams related by crossing and isospin symmetry exist for $\pi^- n$, $\pi^+ p$, and $\pi^+ n$ scattering. To obtain the full cross section, we must sum incoherently over all four channels.

3.2.1 MCP Flux

Given the large multiplicity of Feynman diagrams contributing to $\pi N \to \pi N \chi\bar{\chi}$ and the 4-body final state, we use `MadGraph` to calculate the cross section with momentum cutoffs given by Eqs. (25)–(26) that ensures our cross sections are reliable within the regime of validity of chiral perturbation theory. As shown in Fig. 29, the cross section is only weakly dependent on incident pion energy above the kinematic threshold for $\chi\bar{\chi}$ production.

The flux of MCP at the DUNE ND is given by

$$N_\chi = L N_{\pi^\pm} \left(\sigma_{\pi^+} + \sigma_{\pi^-}\right). \tag{35}$$

To model how many MCPs can reach the detector, we implemented an angular cut on MCP trajectories as depicted in Fig. 26. For an expected detector area of $1\ \mathrm{m}^2$ at a distance of 574 m from the target, the angular acceptance is $\theta_{\mathrm{acc.}} = 8.71 \times 10^{-4}$ rad. The expected flux of MCPs as a function of MCP mass is shown in Fig. 30. For $m_\chi \gtrsim 1$ GeV, we see that the MCP flux from charged pion production exceeds that of Drell-Yan production all the way up to the kinematic threshold of the 120 GeV beam. Interestingly, this is true even for the conservative assumption of the target mass equal to the bare nucleon mass, while for dark photon production in Ref. [126], the relative

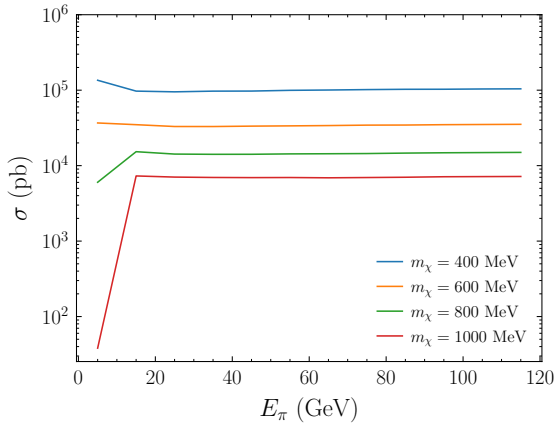

Figure 29. The total inclusive cross section for the scattering process $\pi N \to \pi N \chi \bar{\chi}$, summed over all charged pion and nucleon channels as in Eq. (24), as a function of incoming pion energy E_π for various MCP masses m_χ. The apparent increase in the cross section at low E_π for $m_\chi = 400$ MeV is an artifact of the soft singularity of QED in `MadGraph`, which is why we restrict our analysis to $m_\chi > 400$ MeV.

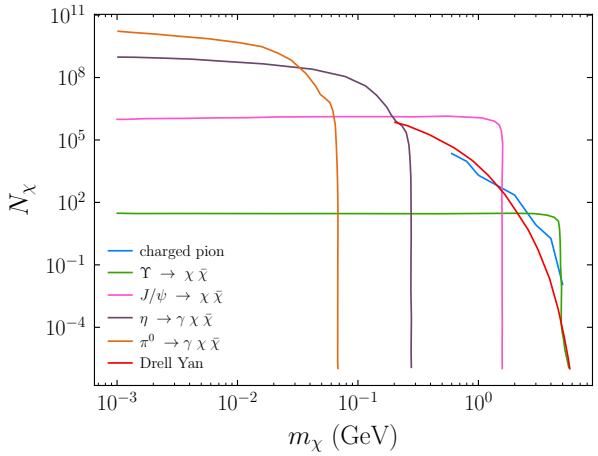

Figure 30. Number of MCP N_χ incident at the DUNE ND, for 10^{21} POT and $\epsilon = 10^{-3}$. Charged pion scattering exceeds Drell-Yan scattering for $m_\chi \gtrsim 1$ GeV up to highest masses (~ 5 GeV) allowed by the chiral cutoff, and is the dominant production mode between 2–3 GeV.

benefit of charged pion scattering is somewhat smaller. As alluded to earlier, `MadGraph` becomes unreliable for small MCP masses as the virtuality of the photon becomes small (this can be seen in the spurious increase of the cross section for $m_\chi = 400$ MeV as $E_\pi \to 0$ in Fig. 29), so we do not show the flux for $m_\chi < 400$ MeV. That said, based on the results of Ref. [126] for similar kinematics, we strongly expect that the flux will be subdominant to meson decays in this mass

range and will not yield any additional gains in sensitivity.

3.2.2 MCP Detector Signatures and Backgrounds

The proposed MCP detector design is based on that of MilliQan [137, 138]. The signature for detecting a MCP is three soft ionizations through the scintillator arrays that are collinear with the beam line and target, detectable through the collection of photoelectrons by PMTs. Following Ref. [3], we also require that all three soft ionizations occur within a small 15 nanosecond window to further reduce background.

There are two types of background that would look like additional soft ionizations throughout the scintillators: detector-related background and beam-related background. To make a direct comparison with Ref. [3], we make identical assumptions about the backgrounds, yielding an estimate of about 600 events per year; we refer the reader to Ref. [3] for more details about the background processes.

3.2.3 MCP Sensitivity at DUNE ND

In the FerMINI design [3, 139], a minimum-ionizing particle with charge $|Q| = 1$ would produce approximately 10^6 photoelectrons (PE) in a 1 m plastic scintillator with a density of $\rho_{\text{scint.}} = 1$ g/cm^3. Thus MCPs with charge $Q = e\epsilon$ would produce PEs that scale as ϵ^2. We can then quantify the probability of a MCP producing at least one PE in all three scintillator stacks as

$$P(\epsilon) = \left(1 - e^{-\overline{N}_{\text{PE}}(\epsilon)}\right)^3, \tag{36}$$

where $\overline{N}_{\text{PE}}$ denotes the average number of photoelectrons produced from soft ionization. We use the approximation of $\overline{N}_{\text{PE}}(\epsilon) = \left(\frac{\epsilon}{\zeta}\right)^2$ where $\zeta = 2 \times 10^{-3}$. This value of ζ corresponds to a Saint-Gobain BC-408 plastic scintillator [3, 138]. Given these choices, such an experiment could be sensitive to ϵ of order 10^{-4}.

With the probability of detecting a PE in the scintillators, we can now calculate the number of MCP signal events using Eqs. (35) and (36),

$$N_S = P(\epsilon)\, N_\chi(m_\chi, \epsilon). \tag{37}$$

Assuming that the scintillators can detect values as low as $\epsilon = 10^{-4}$, we compute the projected 95%

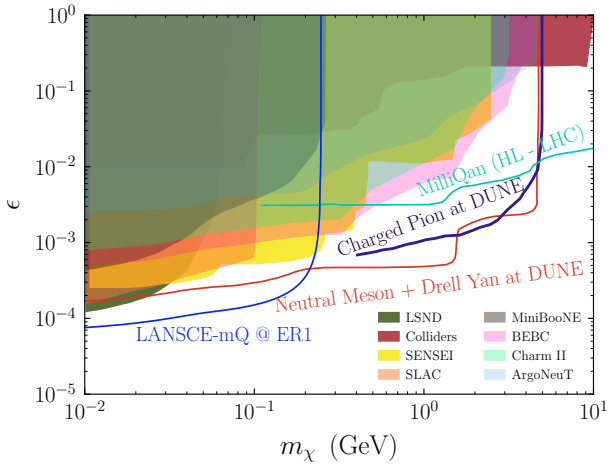

Figure 31. 95% CL sensitivity to MCP detection at DUNE ND with 10^{21} POT (about one year of beam time). The charged pion production channel is shown in purple, and the red-orange curve labeled "Neutral Meson + Drell Yan" is our reproduction of the projected sensitivity from Ref. [3]. The milliQan HL-LHC projected sensitivity [137, 138] is plotted in teal, and the LANSCE-mQ at ER1 [141] projected sensitivity, using 5.9×10^{22} POT (about 3 years of beam time), is plotted in blue. Shaded regions are existing limits from SLAC ([74], peach), SENSEI ([142], yellow), BEBC ([143], light pink), Charm II ([144], light green), LSND ([83], olive green), MiniBooNE ([83], gray), ArgoNeuT ([84], light blue), and colliders ([145], burgundy).

CL limit, i.e. $N_S/\sqrt{N_B} = 2$ where N_B is the number of background events, for the DUNE ND setup to MCP from charged pion scattering. The results are shown in Fig. 31 along with constraints and projections from other beam-dump, collider, and direct detection experiments. We also show the projected limit from neutral meson decay and Drell-Yan scattering at DUNE, following Fig. 1 of Ref. [3] for the total number of MCPs that reach the target, and then applying Eq. (36) in order to model the scintillator response. We find that we reproduce the sensitivity in Ref. [3] to within about a factor of 2, with any differences likely due to slightly different treatments of the angular acceptance [140], which lets us directly assess the increase in sensitivity from charged pion scattering. Overall, the charged pion production channel (which is irreducible in this model) can offer additional sensitivity to detecting MCPs for masses $1.5\ \text{GeV} \lesssim m_\chi \lesssim 3\ \text{GeV}$.

3.3 Heavy Axion-like Particles

The DUNE ND is an excellent experiment to look for long-lived exotic particles at MeV-GeV mass scales. One such particle that we can search for at DUNE ND is a heavy axion coupled to gluons at low energies. This particle may be a generic ALP, but in some models with additional gauge groups, such a heavy axion may in fact be the QCD axion which solves the strong CP

problem [146]. In this section, we will be agnostic as to the UV interpretation of such a particle, and simply use this model as a benchmark to make a direct comparison with the model of Ref. [120]. As we have done with MCPs in Sec. 3.2, we will demonstrate the additional sensitivity gained from charged pion production, which is an irreducible production channel in this model.

As in Ref. [120], we consider a low-energy ALP model where at scales $\mu = \mathcal{O}(1)$ GeV the ALP is only coupled to gluons,

$$\mathcal{L}_a = \frac{\alpha_s}{8\pi f_a} a\, G^a_{\mu\nu} \tilde{G}^{a\mu\nu}. \tag{38}$$

Here, f_a is the axion decay constant, α_s is the strong fine structure constant, and $G^a_{\mu\nu}$ is the gluon field strength tensor. Following [147–149], the axion couplings to nucleons and pions in the chiral Lagrangian can be derived by performing a chiral rotation to shift the axion field into the quark mass matrix. This rotation is arbitrary, but in the conventional parameterization which eliminates mass mixing between the axion and the π^0, we obtain

$$\begin{aligned}
\mathcal{L}_{a\pi N} =& \frac{g_A}{2f_\pi}\Big(\partial_\mu \pi^0 (\bar{p}\gamma^\mu\gamma_5 p - \bar{n}\gamma^\mu\gamma_5 n) \\
&+ \sqrt{2}\partial_\mu \pi^+ \bar{p}\gamma^\mu\gamma_5 n + \sqrt{2}\partial_\mu \pi^- \bar{n}\gamma^\mu\gamma_5 p\Big) \\
&+ \frac{\partial_\mu a}{2f_a}\Big(C_{ap}\bar{p}\gamma^\mu\gamma_5 p + C_{an}\bar{n}\gamma^\mu\gamma_5 n \\
&+ \frac{C_{a\pi N}}{f_\pi}(i\pi^+\bar{p}\gamma^\mu n - i\bar{n}\gamma^\mu p)\Big) \\
+ \frac{\partial_\mu a}{2f_a}\frac{C_{a\pi}}{f_\pi}&(\pi^0\pi^+\partial^\mu\pi^- + \pi^0\pi^-\partial^\mu\pi^+ - 2\pi^+\pi^-\partial^\mu\pi^0),
\end{aligned} \tag{39}$$

where the various axion couplings are given as

$$\begin{aligned}
C_{ap} - C_{an} &= g_A\left(\frac{m_u - m_d}{m_u + m_d}\right), \\
C_{ap} + C_{an} &= -g_0, \\
C_{a\pi N} &= \frac{C_{ap} - C_{an}}{\sqrt{2}g_A}, \\
C_{a\pi} &= \frac{2(C_{ap} - C_{an})}{3g_A}.
\end{aligned} \tag{40}$$

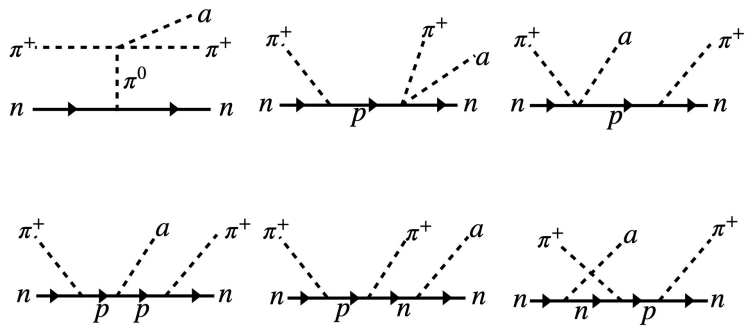

Figure 32. Leading-order diagrams for axion production from π^+-neutron scattering. A similar set of t-channel diagrams exists for π^+-proton scattering, and corresponding s- and t-channel diagrams also exist for π^- scattering.

As in Sec. 3.2, $f_\pi = 93$ MeV is the pion decay constant, $g_A = 1.27$ is the nucleon axial coupling, and $g_0 = 0.52$ is another axial coupling extracted from lattice QCD simulations and evaluated at an $\overline{\text{MS}}$ scale $\mu = 2$ GeV [150]. Note that the chiral rotation which generates these couplings will also generate an axion-photon coupling; in the spirit of this "gluon dominance" model, we will ignore production from this operator and only use it to permit decays $a \to \gamma\gamma$ for $m_a < 3m_\pi$ (see Sec. below). Furthermore, we note that axion-neutral meson kinetic mixing will also be present in the chiral Lagrangian. In order to compare the parametric dependence of scattering processes to previously-considered mixing processes, we are neglecting diagrams which involve kinetic mixing insertions, and we leave to future work a full analysis which simultaneously treats mixing and scattering in order to capture any potential interference effects.

3.3.1 ALP Flux

Axions may be emitted from pion-nucleus scattering through $\pi^\pm N \to \pi^\pm N' a$. Below squared 4-momentum transfers of $(4\pi f_\pi)^2$, we can describe the scattering using chiral perturbation theory, using the Lagrangian in (39). The pions would incoherently scatter off the nucleons contained within the target nuclei. The leading-order amplitudes Feynman diagrams for $\pi^+ n$ scattering are shown in Fig. 32, with similar diagrams for the other charge and isospin channels. We implement this model in `Madgraph` according to the methods in Sec 3.1.2, and show the cross section as a function of incident pion energy in Fig. 33. The derivative couplings in Eq. (39) yield a cross section growing faster with E_π than the MCP cross section studied in Sec. 3.2. In addition, we

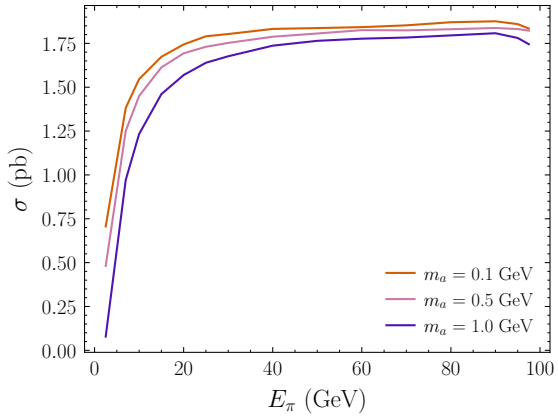

Figure 33. The cross section for the scattering process $\pi N \to \pi N a$ as a function of incoming pion energy E_π for range of axion masses and $f_a = 10^4$ GeV. The cross section scales with increasing pion energy due to the derivative couplings in the chiral Lagrangian. The cross section also has a weak dependence on m_a, which is only significant at low pion energies.

note that our cross section has very weak m_a dependence until the kinematic threshold imposed by the chiral momentum cutoff.

Since the axion has the same quantum numbers as the π^0, axions may also be produced through conversion by replacing the π^0 in any isospin-exchange process, e.g. $\pi^+ n \to a p$. However, it turns out that the production cross section for $2 \to 3$ scattering is an order of magnitude larger than $2 \to 2$ scattering over all masses and couplings we consider. This is perhaps surprising given that one would naively expect the $2 \to 3$ process to be phase-space suppressed. However, charge conservation only allows a π^+ to scatter off a neutron and the π^- to scatter off a proton in $2 \to 2$ scattering. In contrast, in the $2 \to 3$ process, a π^+ could scatter off both protons and neutrons, taking advantage of the full incoherent nuclear scattering. In addition, the $2 \to 3$ process contains six Feynman diagrams compared to three diagrams for $2 \to 2$, and the combinatoric addition of these diagrams seems to overcome the phase space suppression from an additional final-state particles. Furthermore, while the axion could also be produced from scattering initiated by neutral pions as well as charged pions, the neutral pion cross section is suppressed due to the particular values of the coupling constants $C_{ap}, C_{an},$ and $C_{a\pi N}$ in the SM.[13] We explore this somewhat counterintuitive feature of the axion cross section in Appendix .

[13] Due to the arbitrariness of the chiral rotation which generates the axion-pion couplings, and given that our choice eliminates axion-pion mass mixing, it is possible that the relative size of the π^0 cross section is basis-dependent, but the physical inclusive cross section is basis-independent. We thank Joshua Berger for discussions on this point.

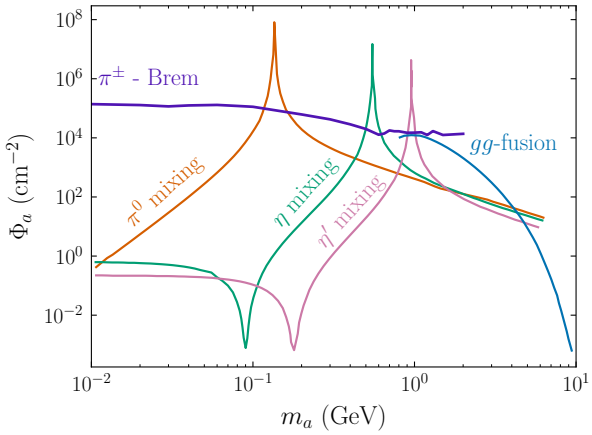

Figure 34. The expected flux of axions at ArgonCube and MPD, located at the DUNE ND Hall. We compare the axion production from charged pion bremsstrahlung (blue) to the production from π^0 decay (orange), η decay (green), η' decay (pink), and gluon-gluon fusion (purple), all assuming $f_a = 10^4$ GeV. For $m_a \lesssim 4\pi f_\pi$, axion production from charged pion scattering dominates the flux, with the exception of where m_a resonates with the neutral meson masses.

We can now calculate the flux of axions at the DUNE ND, using

$$\Phi_a = \frac{L}{A} \times (N_{\pi^+}\sigma_{\pi^+} + N_{\pi^-}\sigma_{\pi^-}), \tag{41}$$

where A is the cross-sectional area of the detector. As with MCPs, we will just compute the number of axions produced within the first pion interaction length in the target, and as such, we will not consider the effect of subsequent target layers on the pion energy distribution.

Fig. 34 shows the flux of axions produced that reach ArgonCube and MPD at the DUNE ND Hall. We compare the flux from charged pion bremsstrahlung to the flux from neutral meson decay and gluon-gluon fusion. With the exception of the resonance peaks, the flux from charged pion bremsstrahlung is approximately $\mathcal{O}(10 - 10^6)$ larger than the flux from neutral meson decay and gluon-gluon fusion depending on the axion mass. This increase in flux will increase the sensitivity of the DUNE ND to axion detection for m_a away from the neutral meson masses.

We can understand this enhancement through a dimensional analysis estimate. The number of axions produced through neutral pion mixing is [120]

$$N_{\pi^0 \to a} \sim N_{\rm POT} N_{\pi^0} \left(\frac{1}{6} \frac{f_\pi}{f_a} \frac{m_a^2}{m_a^2 - m_\pi^2} \right)^2 . \tag{42}$$

For the production of axions through charged pion scattering, each diagram in Fig. 32 has two

pion lines contributing a factor of $1/f_\pi$ each, an axion contributing $1/f_a$, and derivative couplings which contribute factors of the overall energy scale E_π. The $2 \to 3$ cross section therefore scales as

$$\sigma(\pi N \to \pi N a) \sim \frac{1}{16\pi} \frac{E_\pi^4}{f_a^2 f_\pi^4}, \tag{43}$$

where we have included a 3-body phase space factor $\frac{1}{16\pi}$ and neglected the order-1 factors g_0 and g_A as well as any combinatorial factors from the multiplicity of Feynman diagrams. The total number of axions will scale as $\sigma \times l_\pi n_T N_{\pi^\pm}$ POT, so the ratio of production channels assuming $m_a \ll m_\pi$ and $N_{\pi^\pm} \approx 2N_{\pi^0}$ is

$$\begin{aligned} \frac{N_{\pi N \to \pi N a}}{N_{\pi^0 \to a}} &\sim \frac{9}{2\pi} \frac{E_\pi^4}{m_a^4} \frac{l_\pi n_T m_\pi^4}{f_\pi^6} \\ &\sim 1.5 \times 10^8 \left(\frac{E_\pi}{1\ \text{GeV}}\right)^4 \left(\frac{10\ \text{MeV}}{m_a}\right)^4. \end{aligned} \tag{44}$$

This is reasonably close to the $\sim 10^6$ enhancement seen in Fig. 34, which accounts for all of the order-1 coupling factors, the true pion energy distribution, and the detector's geometric acceptance.

3.3.2 ALP Detector Signatures and Backgrounds

The axions that are created within the angular acceptance of ArgonCube and MPD at DUNE ND can be detected by their decay into SM particles. For the axion masses we are interested in, the axion can decay into photons or hadrons. Even in the gluon-dominated model, axions are able to decay to photons through mixing with neutral mesons (indeed, this was the primary production mechanism of Ref. [120]). For $m_a < 3m_\pi$, the axion will primarily decay into two photons, $a \to \gamma\gamma$. The decay width is given by [151]

$$\Gamma_{a\to\gamma\gamma} = \frac{\alpha_{\text{EM}}}{256\pi^3} \frac{m_a^3}{f_a^2} c_\gamma^2, \tag{45}$$

where

$$\begin{aligned} c_\gamma = &-1.92 + \frac{1}{3} \frac{m_a^2}{m_a^2 - m_\pi^2} + \frac{8}{9} \frac{m_a^2 - \frac{4}{9} m_\pi^2}{m_a^2 - m_\eta^2} \\ &+ \frac{7}{9} \frac{m_a^2 - \frac{16}{9} m_\pi^2}{m_a^2 - m_{\eta'}^2}. \end{aligned} \tag{46}$$

For $3m_\pi < m_a \lesssim 1$ GeV, the axion decays to primarily to hadrons through $a \to 3\pi$ and $a \to \pi\pi\gamma$ [120, 151].

We consider the case where the axion decays within the liquid and gaseous argon detectors at DUNE ND, between 574 and 584 m from the target, with an angular acceptance of ~ 3 mrad (see Fig. 26). Approximately 1% of the total number of axions produced are detected due to the small angular acceptance. These axions must also be highly boosted in order to travel 574 m and be within the angular resolution of the detectors. As a result, the decay products are also highly boosted compared to background events [120]. Much of the background in both detectors comes from neutrino-related scattering processes in the argon tanks [120, 127]. Both the diphoton and hadronic backgrounds would have low energies and would result in a roughly isotropic signature. Since our axion signal is highly boosted, we can easily distinguish signal from background events [120].

3.3.3 ALP Sensitivity at DUNE ND

Combining the expected signal from $a \to \gamma\gamma$ and $a \to$ hadrons, we can determine the axion parameter space to which DUNE ND is sensitive. We calculate the number of signal events according to the methods described in Sec. 3.1.4, using the axion decay length $c\tau_a$ from Refs. [120, 152]. The number of axion signal events is then

$$N_S = P_{\text{decay}} \times \text{BR}(a \to X) \times N_a, \tag{47}$$

where $N_a = A\Phi_a$ with Φ_a in (41), and $\text{BR}(a \to X)$ is the branching ratio of the axion decaying to either photons or hadrons [151, 152].

In Figure 35, we show the sensitivity of DUNE ND to 3 or more signal events in 10 years of data collection. The increased production rate from charged pion scattering away from meson resonances improves the projected sensitivity of DUNE in both the long-lifetime (small $1/f_a$) and short-lifetime (large $1/f_a$) regimes. Because we only calculated N_a up to the chiral cutoff $t_n, t_\pi < (4\pi f_\pi)^2$, we do not find additional sensitivity beyond the largest m_a considered in Ref. [120]. For axions below the kaon mass, the sensitivity from charged pion scattering alone is within an order-1 factor of axions produced from kaon decay, recently argued in Ref. [121] to be the dominant production channel in this mass range. Including charged pion scattering beyond the first interaction length of the dump could yield additional sensitivity, possibly making these two

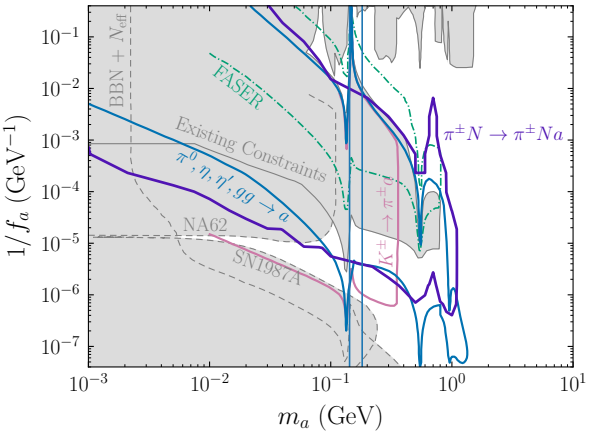

Figure 35. The sensitivity of DUNE ND to 3 or more heavy ALP signal events over 10 years of data collection from charged pion scattering (purple) along with other projections and constraints. We show projections for DUNE ND (solid) along with FASER [153] (green dot-dashed). The projection for an axion signal from neutral meson and gluon-gluon fusion production are shown in blue [120]. The pink curve, adapted from Ref. [121] and rescaled to the same event rate as the other curves, shows diphoton signal events from $K^\pm \to \pi^\pm a$. In shaded gray we show constraints from NA62 [154] for invisible kaon decays, SN1987A [155, 156], and cosmology [157]. In the shaded region labeled "Existing Constraints" we include bounds from invisible kaon decays from E787 and E949 [155], electron beam dump [158–160], CHARM [161], visible kaon decays [162], B meson decays [151], and LHC di-jet searches [163].

production modes comparable.

3.4 Conclusion

In this chapter, I discussed our exploration of the sensitivity of the DUNE ND to heavy axions and millicharged particles produced through charged pion scattering. We found that the expected flux of both ALPs and MCP produced from charged pion scattering was greater than previously-considered production modes meson for some masses in the MeV–GeV range. With this increased flux, we were able to claim a conservative increase in sensitivity of DUNE ND to both ALPs and MCPs. The increase in sensitivity to MCP was modest, but combined with the additional production mechanisms can modestly increase the sensitivity at masses above 2 GeV. The increase in sensitivity to ALPs is similar, and combined with the strong projected constraints from kaon decays in Ref. [121], can likely fully probe remaining parameter space for $m_a \lesssim 250\,\mathrm{MeV}$ between the recent NA62 limits and the SN1987A constraint.

For both sensitivity calculations, we only considered the BSM particle production from a scattering event within the first pion interaction length. However, the energy of the pion is not completely depleted after this event but is only reduced by a factor of $1/e$ on average. The sensitivity to

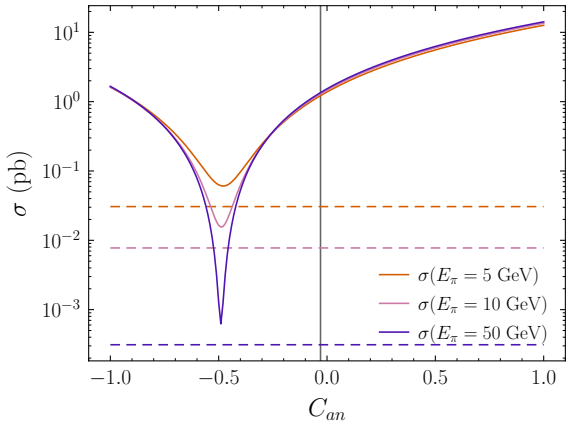

Figure 36. The cross sections $\sigma(\pi^- p \to an)$ (solid) and $\sigma(\pi^0 p \to ap)$ (dashed) as a function of C_{an} values. We used the parameters $m_a = 0.5$ GeV, $f_a = 10^4$ PeV, and $C_{ap} = -0.49$. The vertical gray line indicates our model value $C_{an} = -0.03$. We have plotted the cross section for three incoming pion energies $E_\pi = 5$ GeV (orange), $E_\pi = 10$ GeV (pink), and $E_\pi = 50$ GeV (purple). We see that $\sigma(\pi^- p \to an)$ approaches $2 \times \sigma(\pi^0 p \to ap)$ as C_{an} approaches C_{ap}. For all other values of C_{an}, $\sigma(\pi^- p \to an)$ exceeds $\sigma(\pi^0 p \to ap)$ by at least a factor of 2.

ALPs and MCPs could very likely be improved by incorporating multiple pion scatterings through the target. In addition, there is a substantial flux of pions that survive the ~ 200 m decay region after being focused toward the DUNE detectors, which could rescatter in the rock between the target and ND complex. Axions produced from these charged pion scattering events would only need to travel half the distance to the DUNE ND, thereby increasing the flux of short-lifetime ALPS at large $1/f_a$ which would not have otherwise survived to the detectors. We leave this extension to future work.

Finally, it is reassuring that our chiral perturbation theory calculation closely matches the parton-level Drell-Yan or gluon fusion calculations for BSM particle masses around the chiral cutoff, but there are likely additional gains in sensitivity to both MCPs and axions from a complete treatment of the high-momentum-transfer regime which interpolates between the chiral Lagrangian and perturbative QCD. Such a treatment would involve, for example, quasi-elastic scattering processes and target mass corrections, and for ALPs would consider the scattering and mixing processes in the same theoretical framework. This study would be highly synergistic with the DUNE neutrino physics program where the intermediate-energy regime is crucial for understanding neutrino-nucleus cross sections [128]. Further improvements in the theoretical framework for GeV-scale neutrino cross sections at DUNE would therefore translate directly into improved new physics sensitivity at DUNE, and we look forward to such collaborations involving lattice QCD,

BSM phenomenology, and neutrino physics to maximize the science potential of DUNE.

CHAPTER 4: CONCLUSION

This thesis has focused on searching for new physics by proposing new beam dump searches or modifications to existing beam dump searches. In section , we proposed a search strategy for a scalar particle S whose coupling to the muon could resolve the $(g-2)_\mu$ anomaly. The search strategy would take advantage of the S particle's prompt decay into a pair of muons in order to find it. Specifically, we suggested hitting a low density target with a high-flux muon beam to produce the S particle through bremsstrahlung, which would then decay into a pair of muons. Hence, the signal would be an outgoing μ^- and μ^+ whose invariant mass is equal to the mass of the scalar particle. We suggested accepting only events that passed a 15% invariant mass cut, described in section 2.3.2, as well as a cut requiring that the momentum of the higher energy μ^- and corresponding μ^+, measured in the forward direction along the beam line, were both greater than 5 GeV. These cuts would allow the experiment to be sensitive to the relevant parameter space that would resolve the $(g-2)_\mu$ anomaly. With a flux of 3×10^{14} muons on target, the search would take about 6 years of running at SpinQuest at Fermilab. The SpinQuest collaborators at Fermilab began taking data during the summer of 2024 and we will know within the next 5 years whether they found any signal that would indicate BSM physics. If such evidence is found, physicists would not only have discovered a new type of particle, but as a community, we would be one step closer to resolving existing problems with the Standard Model.

In section 3.1, we described the equipment at Fermilab that we would use for a MCP search. Following the setup of an experiment called FerMINI, we would use DUNE's 120 GeV proton beam and graphite target, alongside installation of PMTs and scintillator arrays, to detect MCPs. We noticed that no existing literature discussed what additional reach you would get if you considered the additional MCPs generated from charged pion scattering, so this was the focus of our work on this project. The protons would scatter off of nuclei through the graphite target to generate charged pions that would then secondarily scatter to produce a $\chi\bar{\chi}$ pair from a virtual photon. After carefully implementing an angular resolution cut to take the detector geometry into consideration, as well as throwing out events that did not satisfy the chiral cutoff described in section 3.1.3, we found that charged pion scattering offers additional sensitivity to detecting MCPs for masses $1.5\ \mathrm{GeV} \lesssim m_\chi \lesssim 3\ \mathrm{GeV}$. Lastly, we proposed a search for heavy axion-like particles (ALPs) using charged pion scattering at DUNE's Near Detector Complex. We would use their proton beam

to generate charged pions that would then scatter to generate ALPs, detectable through their decays into SM particles. Using chiral perturbation theory, we found that the increased production rate from charged pion scattering improves the projected sensitivity of DUNE in both the long-lifetime (small $1/f_a$) and short-lifetime (large $1/f_a$) regimes. Our proposed MCP and ALP searches pave the way for experimentalists to use existing detectors with minimal modifications to have a chance to discover new physics. As technology improves, detectors will have increasing sensitivity to smaller couplings which could open the door for further discovery.

REFERENCES

[1] B. Abi *et al.* (Muon g-2), Phys. Rev. Lett. **126**, 141801 (2021), arXiv:2104.03281 [hep-ex].

[2] A. Apyan, B. Batell, A. Berlin, N. Blinov, C. Chaharom, S. Cuadra, Z. Demiragli, A. Duran, Y. Feng, I. P. Fernando, S. Gori, P. Harris, D. Hoang, D. Keller, E. Kowalczyk, M. Leys, K. Liu, M. Liu, W. Lorenzon, P. Maksimovic, C. M. Suarez, H. Marukyan, A. Mitra, Y. Miyachi, P. McCormack, E. A. Moreno, Y. C. Morales, N. Paladino, M. Rai, S. Rotella, L. Saunders, S. Sawada, C. Smith, D. Sperka, R. Tesarek, N. Tran, Y.-D. Tsai, Z. Wan, and M. Wynne, "Darkquest: A dark sector upgrade to spinquest at the 120 gev fermilab main injector," (2022), arXiv:2203.08322 [hep-ex].

[3] K. J. Kelly and Y.-D. Tsai, Phys. Rev. D **100**, 015043 (2019), arXiv:1812.03998 [hep-ph].

[4] J. J. Bevelacqua, Pub Med **99,5**, 613 (2010).

[5] S. L. Wu, Modern Physics Letters A **29**, 1330027 (2014).

[6] K. Olive *et al.* (Particle Data Group), Chin. Phys. C **38**, 1676 (2014).

[7] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, W. Forman, C. Jones, S. Murray, and W. Tucker, The Astrophysical Journal **606**, 819 (2004).

[8] S. Andreas, C. Niebuhr, and A. Ringwald, Physical Review D **86** (2012), 10.1103/physrevd.86.095019.

[9] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, L. Mo, P. Rassmann, C. Church, and T. Nunamaker, in *4th Moriond Workshop: Massive Neutrinos in Particle Astrophysics* (1984) pp. 227–242.

[10] E. M. Riordan, M. W. Krasny, K. Lang, P. de Barbaro, A. Bodek, S. Dasu, N. Varelas, X. Wang, R. Arnold, D. Benton, P. Bosted, L. Clogher, A. Lung, S. Rock, Z. Szalata, B. W. Filippone, R. C. Walker, J. D. Bjorken, M. Crisler, A. Para, J. Lambert, J. Button-Shafer, B. Debebe, M. Frodyma, R. S. Hicks, G. A. Peterson, and R. Gearhart, Phys. Rev. Lett. **59**, 755 (1987).

[11] M. Davier and H. Nguyen Ngoc, Physics Letters B **229**, 150 (1989).

[12] A. Bross, M. Crisler, S. Pordes, J. Volk, S. Errede, and J. Wrbanek, Phys. Rev. Lett. **67**, 2942 (1991).

[13] A. Keshavarzi, K. S. Khaw, and T. Yoshioka, Nuclear Physics B **975**, 115675 (2022).

[14] D. P. Aguillard *et al.* (The Muon $g-2$ Collaboration), Phys. Rev. Lett. **131**, 161802 (2023).

[15] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio, Phys. Rev. Lett. **109**, 111808 (2012),

arXiv:1205.5370 [hep-ph].

[16] T. Aoyama, T. Kinoshita, and M. Nio, Atoms **7**, 28 (2019).

[17] A. Czarnecki, W. J. Marciano, and A. Vainshtein, Phys. Rev. D **67**, 073006 (2003), [Erratum: Phys.Rev.D 73, 119901 (2006)], arXiv:hep-ph/0212229.

[18] T. Aoyama *et al.*, Phys. Rept. **887**, 1 (2020), arXiv:2006.04822 [hep-ph].

[19] G. W. Bennett *et al.* (Muon g-2), Phys. Rev. D **73**, 072003 (2006), arXiv:hep-ex/0602035.

[20] A. Keshavarzi, D. Nomura, and T. Teubner, Phys. Rev. **D101**, 014029 (2020), arXiv:1911.00367 [hep-ph].

[21] A. Kurz, T. Liu, P. Marquard, and M. Steinhauser, Phys. Lett. **B734**, 144 (2014), arXiv:1403.6400 [hep-ph].

[22] B. Chakraborty *et al.* (Fermilab Lattice, LATTICE-HPQCD, MILC), Phys. Rev. Lett. **120**, 152001 (2018), arXiv:1710.11212 [hep-lat].

[23] S. Borsanyi *et al.* (Budapest-Marseille-Wuppertal), Phys. Rev. Lett. **121**, 022002 (2018), arXiv:1711.04980 [hep-lat].

[24] T. Blum, P. A. Boyle, V. Gülpers, T. Izubuchi, L. Jin, C. Jung, A. Jüttner, C. Lehner, A. Portelli, and J. T. Tsang (RBC, UKQCD), Phys. Rev. Lett. **121**, 022003 (2018), arXiv:1801.07224 [hep-lat].

[25] D. Giusti, V. Lubicz, G. Martinelli, F. Sanfilippo, and S. Simula (ETM), Phys. Rev. **D99**, 114502 (2019), arXiv:1901.10462 [hep-lat].

[26] K. Melnikov and A. Vainshtein, Phys. Rev. D **70**, 113006 (2004), arXiv:hep-ph/0312226.

[27] G. Colangelo, M. Hoferichter, A. Nyffeler, M. Passera, and P. Stoffer, Phys. Lett. **B735**, 90 (2014), arXiv:1403.7512 [hep-ph].

[28] T. Blum, N. Christ, M. Hayakawa, T. Izubuchi, L. Jin, C. Jung, and C. Lehner, Phys. Rev. Lett. **124**, 132002 (2020), arXiv:1911.08123 [hep-lat].

[29] M. Davier, A. Hoecker, B. Malaescu, and Z. Zhang, Eur. Phys. J. C **77**, 827 (2017), arXiv:1706.09436 [hep-ph].

[30] G. Charpak, F. J. M. Farley, R. L. Garwin, T. Muller, J. C. Sens, V. L. Telegdi, and A. Zichichi, Phys. Rev. Lett. **6**, 128 (1961).

[31] J. Bailey *et al.* (CERN Muon Storage Ring), Phys. Lett. B **67**, 225 (1977).

[32] M. Deile *et al.*, Nuclear Physics B - Proceedings Supplements **116**, 215 (2003), proceedings of the 6th International Symposium on Radiative Corrections and the 6th Zeuthen Workshop on Elementary

Particle Theory.

[33] R. M. Carey *et al.*, Phys. Rev. Lett. **82**, 1632 (1999).

[34] H. N. Brown *et al.* (Muon (g-2) Collaboration), Phys. Rev. D **62**, 091101 (2000).

[35] H. N. Brown *et al.* (Muon (g – 2) Collaboration), Phys. Rev. Lett. **86**, 2227 (2001).

[36] G. W. Bennett *et al.* (Muon (g-2) Collaboration), Phys. Rev. Lett. **89**, 101804 (2002).

[37] G. W. Bennett *et al.* (Muon (g-2) Collaboration), Phys. Rev. Lett. **92**, 161802 (2004).

[38] G. W. Bennett *et al.* (Muon g-2 Collaboration), Phys. Rev. D **73**, 072003 (2006).

[39] F. Jegerlehner, *The anomalous magnetic moment of the muon*, Vol. 222 (Springer, 2008).

[40] S. Agostinelli *et al.* (GEANT4), Nucl. Instrum. Meth. A **506**, 250 (2003).

[41] D. Sagan, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **558**, 356 (2006), proceedings of the 8th International Computational Accelerator Physics Conference.

[42] K. Makino and M. Berz, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **558**, 346 (2006), proceedings of the 8th International Computational Accelerator Physics Conference.

[43] J. Mott (Fermilab Wine and Cheese Seminar, 2023).

[44] R. Capdevilla, D. Curtin, Y. Kahn, and G. Krnjaic, Journal of High Energy Physics **2022** (2022), 10.1007/jhep04(2022)129.

[45] Y. Kahn, G. Krnjaic, N. Tran, and A. Whitbeck, JHEP **09**, 153 (2018), arXiv:1804.03144 [hep-ph].

[46] R. Capdevilla, D. Curtin, Y. Kahn, and G. Krnjaic, Phys. Rev. D **103**, 075028 (2021), arXiv:2006.16277 [hep-ph].

[47] R. Capdevilla, D. Curtin, Y. Kahn, and G. Krnjaic, Phys. Rev. D **105**, 015028 (2022), arXiv:2101.10334 [hep-ph].

[48] M. Bauer, P. Foldenauer, and J. Jaeckel, Journal of High Energy Physics **2018** (2018), 10.1007/jhep07(2018)094.

[49] M. Pospelov, Physical Review D **80** (2009), 10.1103/physrevd.80.095002.

[50] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Journal of High Energy Physics **2018** (2018), 10.1007/jhep06(2018)004.

[51] M. Escudero, D. Hooper, G. Krnjaic, and M. Pierre, JHEP **03**, 071 (2019), arXiv:1901.02010 [hep-ph].

[52] G. Krnjaic, G. Marques-Tavares, D. Redigolo, and K. Tobioka, Phys. Rev. Lett. **124**, 041802 (2020), arXiv:1902.07715 [hep-ph].

[53] W. Altmannshofer, S. Gori, M. Pospelov, and I. Yavin, Physical Review Letters **113** (2014), 10.1103/physrevlett.113.091801.

[54] P. Ballett, M. Hostert, S. Pascoli, Y. F. Perez-Gonzalez, Z. Tabrizi, and R. Z. Funchal, Phys. Rev. D **100**, 055012 (2019).

[55] G. Krnjaic, G. Marques-Tavares, D. Redigolo, and K. Tobioka, Physical Review Letters **124** (2020), 10.1103/physrevlett.124.041802.

[56] Y. M. Andreev, D. Banerjee, B. Banto Oberhauser, J. Bernhard, P. Bisio, N. Charitonidis, P. Crivelli, E. Depero, A. V. Dermenev, S. V. Donskov, R. R. Dusaev, T. Enik, V. N. Frolov, A. Gardikiotis, S. V. Gertsenberger, S. Girod, S. N. Gninenko, M. Hösgen, R. Joosten, V. A. Kachanov, Y. Kambar, A. E. Karneyeu, E. A. Kasianova, G. Kekelidze, B. Ketzer, D. V. Kirpichnikov, M. M. Kirsanov, V. N. Kolosov, V. A. Kramarenko, L. V. Kravchuk, N. V. Krasnikov, S. V. Kuleshov, V. E. Lyubovitskij, V. Lysan, V. A. Matveev, R. Mena Fredes, R. G. Mena Yanssen, L. Molina Bueno, M. Mongillo, D. V. Peshekhonov, V. A. Polyakov, B. Radics, K. M. Salamatin, V. D. Samoylenko, D. A. Shchukin, O. Soto, H. Sieber, V. O. Tikhomirov, I. V. Tlisova, A. N. Toropin, M. Tuzi, M. B. Veit, P. V. Volkov, V. Y. Volkov, I. V. Voronchikhin, J. Zamora-Saá, and A. S. Zhevlakov, Phys. Rev. Lett. **132**, 211803 (2024).

[57] "LHC long term schedule — lhc-commissioning.web.cern.ch," https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm, [Accessed 04-11-2024].

[58] Y. M. Andreev, D. Banerjee, B. B. Oberhauser, J. Bernhard, P. Bisio, N. Charitonidis, P. Crivelli, E. Depero, A. V. Dermenev, S. V. Donskov, R. R. Dusaev, T. Enik, V. N. Frolov, A. Gardikiotis, S. V. Gertsenberger, S. Girod, S. N. Gninenko, M. Hösgen, R. Joosten, V. A. Kachanov, Y. Kambar, A. E. Karneyeu, E. A. Kasianova, G. Kekelidze, B. Ketzer, D. V. Kirpichnikov, M. M. Kirsanov, V. N. Kolosov, V. A. Kramarenko, L. V. Kravchuk, N. V. Krasnikov, S. V. Kuleshov, V. E. Lyubovitskij, V. Lysan, V. A. Matveev, R. M. Fredes, R. G. M. Yanssen, L. M. Bueno, M. Mongillo, D. V. Peshekhonov, V. A. Polyakov, B. Radics, K. M. Salamatin, V. D. Samoylenko, D. A. Shchukin, O. Soto, H. Sieber, V. O. Tikhomirov, I. V. Tlisova, A. N. Toropin, M. Tuzi, B. M. Veit, P. V. Volkov, V. Y. Volkov, I. V. Voronchikhin, J. Zamora-Saá, and A. S. Zhevlakov, "Shedding light on dark

sectors with high-energy muons at the na64 experiment at the cern sps," (2024), arXiv:2409.10128 [hep-ex].

[59] R. Capdevilla, D. Curtin, Y. Kahn, and G. Krnjaic, JHEP **04**, 129 (2022), arXiv:2112.08377 [hep-ph].

[60] S. Gninenko, D. Kirpichnikov, and N. Krasnikov, Physical Review D **100** (2019), 10.1103/physrevd.100.035003.

[61] I. Galon, E. Kajamovitz, D. Shih, Y. Soreq, and S. Tarem, Physical Review D **101** (2020), 10.1103/physrevd.101.011701.

[62] B. Collaboration, Physical Review D **94** (2016), 10.1103/physrevd.94.011102.

[63] P. Athron, C. Balázs, D. H. J. Jacob, W. Kotlarski, D. Stöckinger, and H. Stöckinger-Kim, Journal of High Energy Physics **2021** (2021), 10.1007/jhep09(2021)080.

[64] R. Dermíšek and A. Raval, Physical Review D **88** (2013), 10.1103/physrevd.88.013017.

[65] A. Freitas, J. Lykken, S. Kell, and S. Westhoff, Journal of High Energy Physics **2014** (2014), 10.1007/jhep05(2014)145.

[66] F. S. Queiroz and W. Shepherd, Physical Review D **89** (2014), 10.1103/physrevd.89.095024.

[67] L. Calibbi, R. Ziegler, and J. Zupan, Journal of High Energy Physics **2018** (2018), 10.1007/jhep07(2018)046.

[68] R. Dermisek, K. Hermanek, and N. McGinnis, Physical Review D **104** (2021), 10.1103/physrevd.104.l091301.

[69] N. Arkani-Hamed and K. Harigaya, "Naturalness and the muon magnetic moment," (2021), arXiv:2106.01373 [hep-ph].

[70] F. Jegerlehner and A. Nyffeler, Physics Reports **477**, 1–110 (2009).

[71] A. Czarnecki and W. J. Marciano, Nuclear Physics B - Proceedings Supplements **76**, 245–252 (1999).

[72] A. Berlin, D. Hooper, G. Krnjaic, and S. D. McDermott, Phys. Rev. Lett. **121**, 011102 (2018), arXiv:1803.02804 [hep-ph].

[73] J. H. Chang, R. Essig, and S. D. McDermott, Journal of High Energy Physics **2018** (2018), 10.1007/jhep09(2018)051.

[74] A. A. Prinz, *The Search for millicharged particles at SLAC*, Other thesis (2001).

[75] C. Bœhm, M. J. Dolan, and C. McCabe, Journal of Cosmology and Astroparticle Physics **2013**, 041–041 (2013).

[76] S. D. McDermott, H.-B. Yu, and K. M. Zurek, Physical Review D **83** (2011), 10.1103/phys-

revd.83.063509.

[77] T. R. Slatyer, Physical Review D **93** (2016), 10.1103/physrevd.93.023527.

[78] J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, Nature News (2018).

[79] R. Barkana, Nature News (2018), https://doi.org/10.1038/nature25791.

[80] A. Fialkov, R. Barkana, and A. Cohen, Physical Review Letters **121** (2018), 10.1103/physrevlett.121.011101.

[81] H. Tashiro, K. Kadota, and J. Silk, Physical Review D **90** (2014), 10.1103/physrevd.90.083522.

[82] L. Barak, I. M. Bloch, A. M. Botti, M. Cababie, G. Cancelo, B. A. Cervantes-Vergara, L. Chaplinsky, M. Crisler, A. Drlica-Wagner, R. Essig, J. Estrada, E. Etzion, G. F. Moroni, S. E. Holland, Y. Korn, I. Lawson, S. Luoma, S. Munagavalasa, A. Orly, S. E. Perez, D. Rodrigues, N. A. Saffold, S. Scorza, A. Singal, M. S. Haro, L. Stefanazzi, K. Stifter, J. Tiffenberg, S. Uemura, T. Volansky, T.-T. Yu, R. Harnik, Z. Liu, and R. Plestid (SENSEI Collaboration), Phys. Rev. Lett. **133**, 071801 (2024).

[83] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, Phys. Rev. Lett. **122**, 071801 (2019), arXiv:1806.03310 [hep-ph].

[84] R. Acciarri *et al.* (ArgoNeuT), Phys. Rev. Lett. **124**, 131801 (2020), arXiv:1911.07996 [hep-ex].

[85] A. Ball *et al.*, Physical Review D **102** (2020), 10.1103/physrevd.102.032002.

[86] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn, D. L. Rainwater, and T. Stelzer, JHEP **09**, 028 (2007), arXiv:0706.2334 [hep-ph].

[87] C.-Y. Chen, M. Pospelov, and Y.-M. Zhong, Phys. Rev. D **95**, 115005 (2017), arXiv:1701.07437 [hep-ph].

[88] J. P. Lees *et al.* (BaBar), Phys. Rev. D **94**, 011102 (2016), arXiv:1606.03501 [hep-ex].

[89] H. Sieber, D. Banerjee, P. Crivelli, E. Depero, S. N. Gninenko, D. V. Kirpichnikov, M. M. Kirsanov, V. Poliakov, and L. Molina Bueno, Phys. Rev. D **105**, 052006 (2022), arXiv:2110.15111 [hep-ex].

[90] I. Galon, E. Kajamovitz, D. Shih, Y. Soreq, and S. Tarem, Phys. Rev. D **101**, 011701 (2020), arXiv:1906.09272 [hep-ph].

[91] T. Albahri *et al.* (Muon g-2), Phys. Rev. D **103**, 072002 (2021), arXiv:2104.03247 [hep-ex].

[92] T. Albahri *et al.* (Muon g-2), Phys. Rev. A **103**, 042208 (2021), arXiv:2104.03201 [hep-ex].

[93] T. Albahri *et al.* (Muon g-2), Phys. Rev. Accel. Beams **24**, 044002 (2021), arXiv:2104.03240 [physics.acc-ph].

[94] A. Gérardin, H. B. Meyer, and A. Nyffeler, Phys. Rev. D **100**, 034520 (2019), arXiv:1903.09471 [hep-lat].

[95] E.-H. Chao, R. J. Hudspith, A. Gérardin, J. R. Green, H. B. Meyer, and K. Ottnad, Eur. Phys. J. C **81**, 651 (2021), arXiv:2104.02632 [hep-lat].

[96] B. Chakraborty *et al.* (Fermilab Lattice, LATTICE-HPQCD, MILC), Phys. Rev. Lett. **120**, 152001 (2018), arXiv:1710.11212 [hep-lat].

[97] S. Borsanyi *et al.*, Nature **593**, 51 (2021), arXiv:2002.12347 [hep-lat].

[98] M. Passera, W. J. Marciano, and A. Sirlin, Phys. Rev. D **78**, 013009 (2008), arXiv:0804.1142 [hep-ph].

[99] A. Crivellin, M. Hoferichter, C. A. Manzari, and M. Montull, Phys. Rev. Lett. **125**, 091801 (2020), arXiv:2003.04886 [hep-ph].

[100] A. Keshavarzi, W. J. Marciano, M. Passera, and A. Sirlin, Phys. Rev. D **102**, 033002 (2020), arXiv:2006.12666 [hep-ph].

[101] A. Berlin, S. Gori, P. Schuster, and N. Toro, Phys. Rev. D **98**, 035011 (2018), arXiv:1804.00661 [hep-ph].

[102] R. Coy and X.-J. Xu, JHEP **10**, 189 (2021), arXiv:2108.05147 [hep-ph].

[103] C. Rella, B. Döbrich, and T.-T. Yu, Phys. Rev. D **106**, 035023 (2022), arXiv:2205.09870 [hep-ph].

[104] B. Batell, A. Freitas, A. Ismail, and D. Mckeen, Phys. Rev. D **98**, 055026 (2018), arXiv:1712.10022 [hep-ph].

[105] B. Batell, A. Freitas, A. Ismail, D. McKeen, and M. Rai, Phys. Rev. D **104**, 115032 (2021), arXiv:2107.08059 [hep-ph].

[106] S. N. Gninenko, N. V. Krasnikov, M. M. Kirsanov, and D. V. Kirpichnikov, (2016), arXiv:1604.08432 [hep-ph].

[107] M. Bauer, P. Foldenauer, and J. Jaeckel, JHEP **07**, 094 (2018), arXiv:1803.05466 [hep-ph].

[108] S. N. Gninenko, D. V. Kirpichnikov, M. M. Kirsanov, and N. V. Krasnikov, Phys. Lett. **B796**, 117 (2019), arXiv:1903.07899 [hep-ph].

[109] V. Ganapathi and J. Smith, Phys. Rev. D **23**, 75 (1981).

[110] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. D **80**, 075018 (2009), arXiv:0906.0580 [hep-ph].

[111] A. Belyaev, N. D. Christensen, and A. Pukhov, Comput. Phys. Commun. **184**, 1729 (2013),

arXiv:1207.6082 [hep-ph].

[112] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, JHEP **06**, 128 (2011), arXiv:1106.0522 [hep-ph].

[113] L. Marsicano, M. Battaglieri, A. Celentano, R. De Vita, and Y.-M. Zhong, Phys. Rev. D **98**, 115022 (2018), arXiv:1812.03829 [hep-ex].

[114] R. Essig, P. Schuster, N. Toro, and B. Wojtsekhowski, JHEP **02**, 009 (2011), arXiv:1001.2557 [hep-ph].

[115] Y.-S. Liu, D. McKeen, and G. A. Miller, Phys. Rev. D **95**, 036010 (2017), arXiv:1609.06781 [hep-ph].

[116] Y.-S. Liu and G. A. Miller, Phys. Rev. D **96**, 016004 (2017), arXiv:1705.01633 [hep-ph].

[117] A. Apyan *et al.*, in *2022 Snowmass Summer Study* (2022) arXiv:2203.08322 [hep-ex].

[118] S. Gori *et al.*, (2022), arXiv:2209.04671 [hep-ph].

[119] R. Harnik, Z. Liu, and O. Palamara, JHEP **07**, 170 (2019), arXiv:1902.03246 [hep-ph].

[120] K. J. Kelly, S. Kumar, and Z. Liu, Phys. Rev. D **103**, 095002 (2021), arXiv:2011.05995 [hep-ph].

[121] J. Berger and G. Putnam, (2024), arXiv:2405.18480 [hep-ph].

[122] R. Acciarri *et al.* (DUNE), (2015), arXiv:1512.06148 [physics.ins-det].

[123] J. Strait *et al.* (DUNE), (2016), arXiv:1601.05823 [physics.ins-det].

[124] R. Acciarri *et al.* (DUNE), (2016), arXiv:1601.05471 [physics.ins-det].

[125] R. Acciarri *et al.* (DUNE), (2016), arXiv:1601.02984 [physics.ins-det].

[126] D. Curtin, Y. Kahn, and R. Nguyen, Phys. Rev. D **108**, 095039 (2023), arXiv:2305.19309 [hep-ph].

[127] J. M. Berryman, A. de Gouvea, P. J. Fox, B. J. Kayser, K. J. Kelly, and J. L. Raaf, JHEP **02**, 174 (2020), arXiv:1912.07622 [hep-ph].

[128] L. A. Ruso *et al.*, (2022), arXiv:2203.09030 [hep-ph].

[129] R. Ruiz *et al.*, Prog. Part. Nucl. Phys. **136**, 104096 (2024), arXiv:2301.07715 [hep-ph].

[130] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079 (2014), arXiv:1405.0301 [hep-ph].

[131] S. Scherer, Adv. Nucl. Phys. **27**, 277 (2003), arXiv:hep-ph/0210398.

[132] C. S. Shin and S. Yun, (2022), arXiv:2211.15677 [hep-ph].

[133] P. J. Ellis and H.-B. Tang, Phys. Rev. C **57**, 3356 (1998), arXiv:hep-ph/9709354.

[134] V. Bernard, N. Kaiser, and U. G. Meissner, Nucl. Phys. B **457**, 147 (1995), arXiv:hep-ph/9507418.

[135] V. Bernard, N. Kaiser, and U.-G. Meissner, Nucl. Phys. A **619**, 261 (1997), arXiv:hep-ph/9703218.

[136] A. Czarnecki, W. J. Marciano, and A. Sirlin, Phys. Rev. Lett. **120**, 202002 (2018), arXiv:1802.01804 [hep-ph].

[137] A. Haas, C. S. Hill, E. Izaguirre, and I. Yavin, Phys. Lett. B **746**, 117 (2015), arXiv:1410.6816 [hep-ph].

[138] A. Ball *et al.*, (2016), arXiv:1607.04669 [physics.ins-det].

[139] C. Patrignani *et al.* (Particle Data Group), Chin. Phys. C **40**, 100001 (2016).

[140] K. Kelly, private communication.

[141] Y.-D. Tsai, I. Hwang, R. Schmitz, M. Citron, K. Gunthoti, J. Steenis, H. Jeong, H. Moon, J. H. Yoo, and M. X. Liu, (2024), arXiv:2407.07142 [hep-ph].

[142] O. Abramoff *et al.* (SENSEI), Phys. Rev. Lett. **122**, 161801 (2019), arXiv:1901.10478 [hep-ex].

[143] G. Marocco and S. Sarkar, SciPost Phys. **10**, 043 (2021), arXiv:2011.08153 [hep-ph].

[144] K. De Winter *et al.* (CHARM-II), Nucl. Instrum. Meth. A **278**, 670 (1989).

[145] S. Davidson, S. Hannestad, and G. Raffelt, JHEP **05**, 003 (2000), arXiv:hep-ph/0001179.

[146] A. Hook, S. Kumar, Z. Liu, and R. Sundrum, Phys. Rev. Lett. **124**, 221801 (2020), arXiv:1911.12364 [hep-ph].

[147] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, Phys. Rept. **870**, 1 (2020), arXiv:2003.01100 [hep-ph].

[148] S. Chang and K. Choi, Phys. Lett. B **316**, 51 (1993), arXiv:hep-ph/9306216.

[149] K. Choi, H. J. Kim, H. Seong, and C. S. Shin, JHEP **02**, 143 (2022), arXiv:2110.01972 [hep-ph].

[150] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro, JHEP **01**, 034 (2016), arXiv:1511.02867 [hep-ph].

[151] D. Aloni, Y. Soreq, and M. Williams, Phys. Rev. Lett. **123**, 031803 (2019), arXiv:1811.03474 [hep-ph].

[152] N. Blinov, E. Kowalczyk, and M. Wynne, JHEP **02**, 036 (2022), arXiv:2112.09814 [hep-ph].

[153] A. Ariga *et al.* (FASER), Phys. Rev. D **99**, 095011 (2019), arXiv:1811.12522 [hep-ph].

[154] E. Cortina Gil *et al.* (NA62), JHEP **06**, 093 (2021), arXiv:2103.15389 [hep-ex].

[155] F. Ertas and F. Kahlhoefer, JHEP **07**, 050 (2020), arXiv:2004.01193 [hep-ph].

[156] J. H. Chang, R. Essig, and S. D. McDermott, JHEP **09**, 051 (2018), arXiv:1803.00993 [hep-ph].

[157] P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, JCAP **05**, 009 (2020), arXiv:2002.08370 [hep-

ph].

[158] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, JHEP **02**, 018 (2016), arXiv:1512.03069 [hep-ph].

[159] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, JHEP **12**, 094 (2017), [Erratum: JHEP 03, 190 (2021)], arXiv:1709.00009 [hep-ph].

[160] D. Banerjee *et al.* (NA64), Phys. Rev. Lett. **125**, 081801 (2020), arXiv:2005.02710 [hep-ex].

[161] F. Bergsma *et al.* (CHARM), Phys. Lett. B **157**, 458 (1985).

[162] S. Gori, G. Perez, and K. Tobioka, JHEP **08**, 110 (2020), arXiv:2005.05170 [hep-ph].

[163] A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka, Phys. Lett. B **783**, 13 (2018), arXiv:1710.01743 [hep-ph].

APPENDIX: AXION CROSS SECTION PECULIARITIES

In this Appendix, we investigate some unusual features of the pion-nucleon-axion cross section. First, we consider the relative sizes of the charged and neutral pion cross sections. In our calculation of $\sigma(\pi^0 N \to \pi^0 N a)$ in `Madgraph`, we found that the flux of axions was approximately an order of magnitude less than that of $\pi^{\pm} N$-scattering in Fig. 34. It appears that this is due to the hierarchy of the low-energy coupling constants C_{ap} and C_{an}; we can see this feature in the simpler $2 \to 2$ cross sections $\pi^0 n \to an$, $\pi^0 p \to ap$, $\pi^+ n \to ap$, $\pi^+ p \to an$.

In the gluon dominance model, $C_{ap} = -0.49$ and $C_{an} = -0.03$. To isolate the dependence on the proton and neutron couplings, we compare $\sigma(\pi^0 p \to ap)$, which only depends on C_{ap}, to $\sigma(\pi^- p \to an)$, which depends on both C_{ap} and C_{an}. These two processes have similar Feynman diagrams, apart from an additional contact diagram for the $\pi^- p$ scattering. Since $C_{a\pi N} = 0$ when $C_{ap} = C_{an}$, we expect $\sigma(\pi^- p \to an) = 2 \times \sigma(\pi^0 p \to ap)$ when the couplings are the same. Away from this limit, any large variation between the charged and neutral pion cross sections must be due to interference effects.

In Figure 36, we plot the cross sections $\sigma(\pi^- p \to an)$ and $\sigma(\pi^0 p \to ap)$ for a fixed $C_{ap} = -0.49$ while varying $C_{an} \in [-1, 1]$. We chose to vary C_{an} values because the neutral pion scattering does not depend on C_{an}. We see the expected behavior as we approach $C_{ap} = C_{an}$, but away from this limit, $\sigma(\pi^- p \to an)$ can be orders of magnitude larger than $\sigma(\pi^0 p \to ap)$. At the physical value $C_{an} = -0.03$, the ratio of cross sections is:

$$\frac{\sigma(\pi^- p \to an)}{\sigma(\pi^0 p \to ap)} = 2. \tag{48}$$

We therefore conclude that the large hierarchy which we also observe in $2 \to 3$ scattering is likely due to interference effects.