

Discovering New Light States at Neutrino Experiments

Rouven Essig,^{1,*} Roni Harnik,^{2,†} Jared Kaplan,^{1,‡} and Natalia Toro^{3,§}

¹*Theory Group, SLAC National Accelerator Laboratory, Menlo Park, CA 94025*

²*Theoretical Physics Department, Fermilab, Batavia, IL60510, USA*

³*Theory Group, Stanford University, Stanford, CA 94305*

(Dated: August 5, 2010)

Abstract

Experiments designed to measure neutrino oscillations also provide major opportunities for discovering very weakly coupled states. In order to produce neutrinos, experiments such as LSND collide thousands of Coulombs of protons into fixed targets, while MINOS and MiniBooNE also focus and then dump beams of muons. The neutrino detectors beyond these beam dumps are therefore an excellent arena in which to look for long-lived pseudoscalars or for vector bosons that kinetically mix with the photon. We show that these experiments have significant sensitivity beyond previous beam dumps, and are able to partially close the gap between laboratory experiments and supernovae constraints on pseudoscalars. Future upgrades to the NuMI beamline and Project X will lead to even greater opportunities for discovery. We also discuss thin target experiments with muon beams, such as those available in COMPASS, and show that they constitute a powerful probe for leptophilic PNGBs.

*Electronic address: rouven@stanford.edu

†Electronic address: roni@fnal.gov

‡Electronic address: jaredk@slac.stanford.edu

§Electronic address: ntoro@stanford.edu

Contents

I. Introduction	2
A. Models of Weakly Coupled Light Exotics	5
1. Pseudo-Goldstone Bosons	5
2. Kinetically Mixed Gauge Bosons	7
II. Constraints on PNGBs from Rare Meson Decays and Supernovas	8
A. Limits from Kaons and B-Meson Decays	8
B. Limits from Supernova SN 1987a	12
C. Limits from the anomalous muon magnetic moment	14
III. Constraints from LSND on PNGBs and Dark Gauge Bosons	14
A. Production of PNGBs at LSND and other Proton Beams	15
B. LSND Analyses Sensitive to PNGBs	17
C. LSND Limits on Light Vector Bosons	19
IV. Sensitivity of Modern Neutrino Experiments	20
A. Production from Proton-Nucleus Collisions	20
B. Production through Muon a -sstrahlung	22
V. Muon Fixed-Target Experiments with Thin Targets – COMPASS	26
VI. Conclusions	28
Acknowledgements	29
A. Pseudoscalar Production	30
References	32

I. INTRODUCTION

In the last decade, several experiments have explored neutrino masses and mixings, but these high-luminosity laboratories are also sensitive to rare production of new metastable particles. Neutrino beams such as the LAMPF Neutrino Source at Los Alamos and NuMI

and BooNE at Fermilab are produced through two basic stages. First, a high-intensity proton beam impinges on a target and produces a large number of pions (and other hadrons), which decay to muons and neutrinos. The muons are stopped in a thick layer of rock, while the neutrinos travel unimpeded through the rock to the detector. In fact, short-baseline neutrino detectors are situated behind the most intense proton *and muon* beam-dumps to date. Thus they are ideally configured to search for long-lived particles produced by rare proton-nucleus or muon-nucleus interactions.

Two classes of new physics scenarios naturally give rise to light, feebly coupled particles of this type. An approximate symmetry broken at a high mass scale F naturally gives rise to light pseudoscalars — pseudo-Nambu-Goldstone bosons (PNGBs, or “generic axions”) — with couplings of order m_X/F to Standard Model matter X . Alternately, a new “dark” $U(1)$ gauge boson can naturally have small kinetic mixing ϵ with the photon, giving rise to suppressed interactions with all electrically charged matter [1]. Either spin-1 or spin-0 bosons can be radiated in energetic-particle interactions with matter, with a very small rate proportional to the square of their weak coupling. The luminosities achieved in fixed-target experiments are such that thousands of these particles could be produced, and they offer a new window into weakly coupled sectors. Once produced, the lightest particles in a hidden sector could decay only through their weak couplings to Standard Model particles, and so would be quite long-lived and weakly interacting. While ordinary products of the collision are stopped in the shielding upstream of a neutrino detector, exotics could penetrate the shielding and decay within the detector, yielding a distinctive signal.

This note summarizes the several classes of exotic particles that naturally give rise to observable late-decay signals, their experimental signatures and typical kinematics. In any given model, the production cross section and lifetimes of these exotica are both determined by a single small coupling parameter and by the mass of the produced particle, so in particular, we present specific estimates for the sensitivity achievable with late-decay searches in MINOS/MINERvA, MiniBooNE, and LSND.

PNGB’s coupled to hadrons were searched for extensively in both proton and electron beam-dump experiments in the 1980s, most notably in CHARM [2] at CERN, E774 [3] at Fermilab, and the SLAC experiments E137 [4] and E141 [5]. Many of these limits have recently been re-interpreted ([6–8]) in the context of kinetically mixed gauge bosons and the associated scalar bosons that give them mass through the Higgs mechanism. Kinetically

mixed gauge bosons have been a subject of considerable recent interest [9–21] and discussions of other collider, accelerator, and direct and indirect astrophysical probes for them can be found in e.g. [22–39]. In particular, [7] discussed the sensitivity of neutrino experiments to hadronic production in the case of kinetic mixing and the potential importance of existing LSND data as a constraint on these models. Additional constraints on both classes of models from supernovae, rare decays, and radiative corrections have also been extensively discussed [6, 40–44].

Our aim in this note is to present a more complete analysis of the sensitivity of past, present, and future neutrino experiments to new weakly-coupled physics, with a particular focus on PNCB models. We hope such a unified summary will facilitate new analyses of neutrino-detector data to discover or constrain new weakly-coupled particles. We also discuss the potential reach for experiments with muon beams that strike a fixed thin target. In §II, we review constraints on generic pseudoscalars from other arenas, in particular supernova data, rare meson decays, and the anomalous muon magnetic moment. The range explored by neutrino experiments is complementary to all of these. In §III, we consider the implications of existing LSND analyses for both PNCB’s and dark gauge bosons. Due to the large number of protons dumped in LSND, we find that these analyses provide stronger constraints on PNCB’s than the CHARM experiment. Our results for dark gauge bosons are consistent with [7], but we have tried to clarify the experimental sensitivity. In §IV, we consider the sensitivity to PNCBs that could be achieved by analyses using modern neutrino beamlines, such as BooNE and NuMI, and their near detectors (the complementary analysis for kinetically mixed gauge bosons was presented in [7]). We consider both the standard production mode in proton-nucleus collisions and the production of very forward PNCBs off the stopping muons. The second mode is enhanced by the magnetic focusing of pions, so experiments using focused neutrino beams are uniquely sensitive to purely leptophilic PNCBs that do not couple to hadrons. For ordinary PNCBs coupled to both quarks and leptons, searches in MINOS/MINERvA and MiniBooNE would have sensitivity comparable to, or perhaps slightly better than, the CHARM beam-dump limit [2]. A future “Project X” could significantly extend this reach into new territory. For leptophilic PNCBs, MINOS/MINERvA has slightly better sensitivity than the constraint from E137. In §V, we discuss muon fixed-target experiments using thin targets. Such an experiment could be possible at the COMPASS experiment at CERN. We find that a COMPASS-like

setup can probe new territory in the parameter space of leptophilic PNGBs, closing a gap between muon $g - 2$ limits and those from neutrino experiments. We conclude in §VI, and an appendix discusses the details of PNGB production off muon beams.

A. Models of Weakly Coupled Light Exotics

1. Pseudo-Goldstone Bosons

Light pseudoscalars can arise as pseudo-Goldstone bosons in a large variety of well-motivated theories, such as multiple higgs doublet models, theories with an R-axion [45] (from spontaneous and explicit R-symmetry breaking in a supersymmetric theory), axion models [46], the Next-to-Minimal Supersymmetric Standard model [47] (NMSSM), and recent dark matter models with a scalar portal to the dark sector [48]. The most important point is that these particles are naturally light if there is an approximate shift symmetry, and that their interactions are proportional to the inverse of some symmetry breaking scale F . Using fermion equations of motion, the derivative coupling of a PNGB a to a fermion bilinear turns into the operator (in Weyl notation)

$$L \supset \frac{m_\chi}{F} a \chi \chi, \quad (1)$$

which is the coupling we assume for the leptons and/or quarks, χ , of the Standard Model. A phenomenologically interesting sub-class of PNGB models are those where the PNGB is leptophilic, i.e. it couples preferentially (or only) to leptons; this scenario could arise if the lepton sector has its own higgs mechanism separate from that of the quark sector. Leptophilic PNGB are essentially unconstrained by searches for rare meson decays and proton fixed target experiments, so experiments that have muon beams, such as MINOS/MINERvA, MiniBooNE, and COMPASS, can easily be the most sensitive probes of these particles.

The coupling of a PNGB with mass m_a to a fermion with mass m_ℓ induces a PNGB partial width

$$\Gamma_\ell = \frac{m_a}{8\pi} \left(\frac{m_\ell}{F} \right)^2 \sqrt{1 - (4m_\ell^2/m_a^2)} \quad (2)$$

and the total width is well approximated by $\Gamma_e + \Gamma_\mu$ for $m_a \lesssim 400$ MeV (for larger masses, hadronic decays can also become important but we use the leptonic widths for masses up

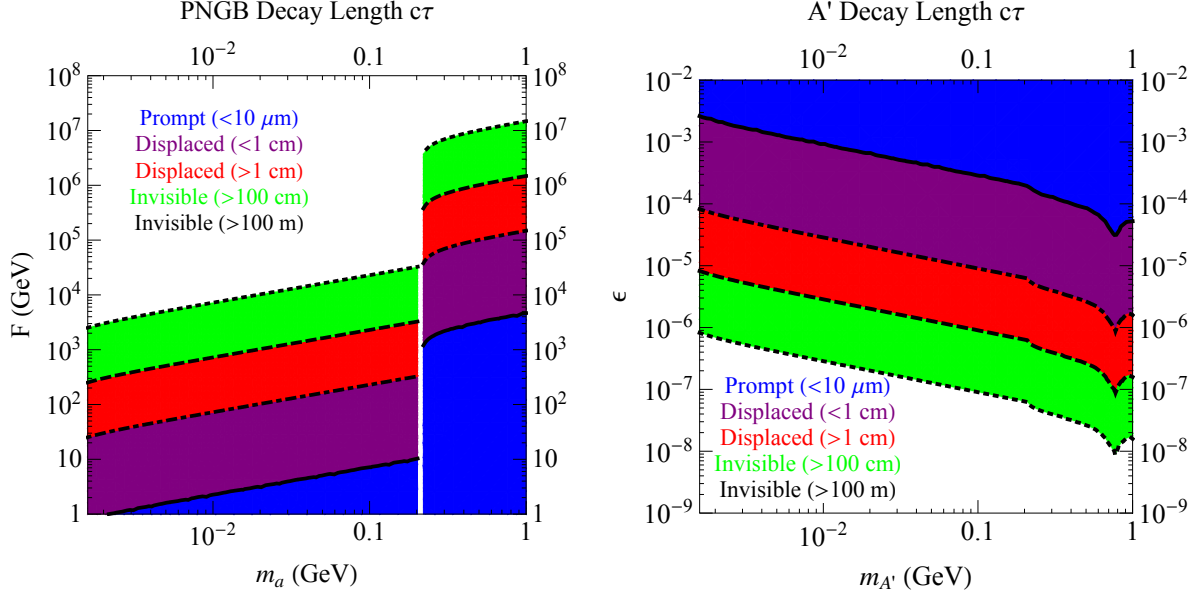


FIG. 1: **Left:** The rest-frame lifetime of a pseudo-Nambu-Goldstone boson (PNGB) as a function of its mass m_a and decay constant F . **Right:** The lifetime of a dark photon A' as a function of its mass $m_{A'}$ and ϵ , the strength of its mixing with the Standard Model hypercharge gauge boson. In both plots, the black lines correspond to different decay lengths ($c\tau$): $10 \mu\text{m}$ (solid), 1 cm (dot-dashed), 1 m (dashed), and 100 m (dotted). In the blue, purple, red, green, and white shaded regions the decays are prompt ($< 10 \mu\text{m}$), displaced with $< 1 \text{ cm}$, displaced with $> 1 \text{ cm}$, “invisible” with $> 100 \text{ cm}$, or “invisible” with $> 100 \text{ m}$, respectively. In fixed-target or beam dump experiments the particles typically get a large boost that increases their decay length by $E_{\text{beam}}/\text{mass}$. The feature in the left plot at $2m_\mu$ occurs since PNGB’s coupling to a Standard Model particle is proportional to that particle’s mass, and at this point decays to two muons are allowed. The dip in the right plot near 0.7 GeV is due to the ρ -resonance. The lifetime for both the PNGB and the A' is calculated assuming decays directly into Standard Model particles.

to 1 GeV). Thus, for example, proper lifetimes of 1 mm are obtained with $F \approx 70 \text{ GeV}$ ($m_a = 100 \text{ MeV}$) or $F \approx 20 \text{ TeV}$ ($m_a = 300 \text{ MeV}$). Fig. 1 (left plot) shows the decay length ($c\tau$) of a PNGB as a function of its mass m_a and decay constant F . Note that the decay length is very different above and below the muon threshold, due to the much stronger coupling to muons as compared to electrons. We see that for $F \lesssim 10^2 \text{ GeV}$, they decay promptly and colliders should be able to set the best constraints. For larger F , collider searches that look for displaced vertices or missing energy can still set limits, but searches in beam dump experiments (with a large shield) become relevant.

We ignore decays of PNGBs to two photons, since this is always subdominant in the mass range we consider in this paper.

2. *Kinetically Mixed Gauge Bosons*

Another class of light particles that has received significant interest in the last few years is a light, weakly coupled “sequestered sector” lying alongside the Standard Model. New states at the MeV–GeV scale are not in conflict with data, because the gauge and global symmetries of the Standard Model greatly restrict the couplings of ordinary matter to these states. It is, however, quite natural for the few allowed interactions to be suppressed by loop factors.

In particular, we consider a new MeV–GeV scale Abelian gauge boson A' coupled to electrically charged Standard Model particles ψ

$$\delta\mathcal{L} = \epsilon e A'_\mu \bar{\psi} \gamma^\mu \psi. \quad (3)$$

Such a coupling can generically originate from the kinetic mixing between the field strengths of the Standard Model hypercharge and a hidden sector $U(1)$ [1],

$$\delta\mathcal{L} = \frac{\epsilon_Y}{2} F'_{\mu\nu} F_Y^{\mu\nu}, \quad (4)$$

where $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the field strength of the A' gauge boson, and similarly $F_Y^{\mu\nu}$ is the hypercharge field strength. This mixing, assuming the hidden $U(1)$ is broken so that the A' is massive, is equivalent in low-energy interactions to assigning a hidden charge $\epsilon e q_i$ to Standard Model particles of electromagnetic charge q_i , where $\epsilon = \epsilon_Y / (\cos \theta_W)$ and θ_W is the Weinberg mixing angle (Eq. (3) if $q_\psi = 1$). Such a mixing can be generated in many ways, e.g. through loops of new heavy particles that couple to both the A' and Standard Model hypercharge. We refer the reader to e.g. [27] for more detailed discussions.

In Fig. 1 (right), we show the decay length of a vector boson A' as a function of its mass $m_{A'}$ and the parameter ϵ which sets the strength of its kinetic mixing with hypercharge. Assuming the A' decays into Standard Model particles rather than exotics, its lifetime is

$$\gamma c\tau \simeq \frac{3}{N_{\text{eff}} m_{A'} \alpha \epsilon^2} \simeq \frac{0.8 \text{cm}}{N_{\text{eff}}} \left(\frac{E_0}{10 \text{GeV}} \right) \left(\frac{10^{-4}}{\epsilon} \right)^2 \left(\frac{100 \text{MeV}}{m_{A'}} \right)^2, \quad (5)$$

where we have neglected phase-space corrections, and N_{eff} counts the number of available decay products. If the A' mixes kinetically with the photon, then $N_{\text{eff}} = 1$ for $m_{A'} < 2m_\mu$

Particle	Final state	Production mode	Momentum spectrum
PNGB	$\begin{cases} 2\mu & (m > 2m_\mu) \\ 2e, 2\gamma & (m < 2m_\mu) \end{cases}$	Proton-Nucleus (a/π^0 -mixing)	Same as π^0 spectrum
”	”	Muon a -strahlung	muon spectrum (avg. over material)
Dark gauge boson	$2e, 2\mu, \text{ hadronic modes}$	Proton-Nucleus (π^0 decay)	1/2 of typical final-state π^0 momentum
”	”	Muon-Nucleus (A' -sstrahlung)	muon spectrum (avg. over material)

TABLE I: Summary of Signal Properties for Light Exotics considered in this paper. The PNGB decay to two photons is never important for the mass range we consider, and can be ignored. More complicated signals are possible if the PNGB or A' can decay to other hidden sector particles before decaying to Standard Model particles, but we will not consider this possibility.

when only $A' \rightarrow e^+e^-$ decays are possible, and $2 + R(m_{A'})$ for $m_{A'} \geq 2m_\mu$, where $R = \frac{\sigma(e^+e^- \rightarrow \text{hadrons}; E=m_{A'})}{\sigma(e^+e^- \rightarrow \mu^+\mu^-; E=m_{A'})}$ [49].

We summarize the possible production and decay mechanisms for both PNGBs and A' 's in Table I. In this paper, we only consider the case for which the PNGB or A' decays directly to Standard Model matter. More complicated signals are possible if they can decay to other hidden sector particles before decaying to Standard Model particles.

II. CONSTRAINTS ON PNGBS FROM RARE MESON DECAYS AND SUPER-NOVAS

A. Limits from Kaons and B-Meson Decays

In this section and in Table II, we briefly summarize constraints from meson decays on PNGBs with sub-GeV masses. These constraints are controlled by two factors: the partial width for the rare meson decays into a PNGB, and the fraction of PNGB's that decay promptly enough to be included in the data samples (or, in the case of invisible-decay searches, the fraction that decay outside the detector). The combination of these searches is sensitive to PNGBs with $F \lesssim 10 - 100$ TeV over a wide range of masses. We call these limits “exclusions” below in the interest of brevity, but it should be emphasized that several of

Decay mode	B.R. limit	Mass range	Decay region
$B^+ \rightarrow K^+ + inv.$ [50]	3×10^{-6}	$0 < m_a < 2m_\mu^{(*)}$	$> 1.65\text{m}$
$K^+ \rightarrow \pi^+ + inv.$ [51]	$0.5 - 10 \times 10^{-10}$	$0 < m_a < 2m_\mu^{(*)}$	$> 1.3\text{m}$
$K^+ \rightarrow \pi^+ + X$ $X \rightarrow e^+e^-$ [52]	$1.5 - 5 \times 10^{-6}$	$10 \text{ MeV} < m_a < 120 \text{ MeV}$	$< 5\text{cm}^{(\dagger)}$
$B^+ \rightarrow K^+e^+e^-$ [53]	6.7×10^{-7}	$30 \text{ MeV} < m_a < 2m_\mu^{(*)}$	$< 100\mu\text{m}^{(\dagger)}$
$d\Gamma/dm^2(B^+ \rightarrow K^+\ell^+\ell^-)$ [54]	1.2×10^{-7}	$140 < m_a < 1440 \text{ MeV}$	$< .5\text{cm}$

TABLE II: Summary of constraints on PNCBs from various experiments (B.R. = Branching Ratio). For decay of meson X , a characteristic transverse energy $m_X/2$ is assumed in computing the fraction of decays either in or outside a cylinder whose radius is given in the “decay region” column.

($*$): Searches extend beyond $2m_\mu$ but only imply relevant limits on PNCB models below $2m_\mu$.

(\dagger): Length scales are guesses only, but overall exclusion is insensitive to cutoff because it overlaps with reach of invisible searches.

the results have been re-interpreted by non-experts (the authors) in a context very different from the original experimental design, from which significant inaccuracies in the boundaries could have resulted. We focus here on the most constraining searches; a more exhaustive set of limits are considered in [43], but our treatment of the decay regions and the PNCB lifetime differ from theirs.

Following [42], we consider two-higgs-doublet models as a generic framework for a new PNCB a that couples to the Standard Model (whereas [43] assumes NMSSM couplings). In this framework, rare meson decays to a lighter meson and a are mediated primarily by a top-quark loop, which receives contributions at all scales up to the highest energy at which the PNCB couples to the top quark (typically either the F scale or the electroweak symmetry-breaking scale). Therefore, the meson decay rates depend on the detailed UV dynamics through which the a particle couples to Standard Model fields, as discussed in [42]. We adopt the approximate formula

$$\Gamma(B^+ \rightarrow K^+a) \approx \frac{G_F^3 |V_{tb}^* V_{ts}|^2}{\sqrt{2} 2^{12} \pi^5} m_t^4 m_B^3 \left(\frac{2v^2}{F^2} \right) (\textit{kinematic}) [f_0(m_a)^2], \quad (6)$$

and likewise for $K^+ \rightarrow \pi^+a$, but with the product of CKM matrices $V_{tb}^* V_{ts}$ instead given by $V_{ts}^* V_{td}$. Here G_F is the Fermi coupling constant, $v = 174 \text{ GeV}$ is the electroweak Higgs vacuum expectation value, m_t is the top mass, and m_B the B -meson mass. This formula is obtained from (A.1) of [42] by assuming $\beta = 45^\circ$ in the two-higgs-doublet model, replacing $\sin \theta \rightarrow v/(2F)$, and setting the model-dependent combination $|X_1 + X_2| = 1$. This

combination depends on the physical Higgs mass but is typically in the range 1–10 so that our choice is conservative (although note that there are special choices of the Higgs mass where the sum of X 's crosses zero, in which case limits from meson decay are weaker). The kinematic factor is typically in the range 0.5-1 and is defined in (A.3) of [42], while the form factor $f_0 \approx 0.33$ [55]. The formula above with our parameter choices yields branching fractions about a factor of 10 smaller than those in [43].

The second factor that comes into play is the decay length of the PNGB. For visible meson decay searches, the PNGB must decay within a small distance (typically of order mm-cm) from the meson decay from which it originates, so that all of the tracks produced by the meson reconstruct a single vertex. Likewise, invisible-decay searches require that the PNGB decay *outside* the detector volume, with a typical size of about a meter. We have attempted to take these effects into account more precisely than [43], which uses a uniform length scale of 10 m for all limits.

The PNGB partial width to $\ell^+\ell^-$ is given in Eq. (2) and the lifetime is shown in Fig. 1. Naively, one would expect that PNGB's with characteristic lifetimes between 1 mm and 1 m might be poorly constrained by both visible-decay and invisible-decay searches. However, the fraction of events in which the PNGB decays within an atypically short distance from the meson vertex can still set significant constraints; moreover, the limits from experiments with different energy scales (m_K and m_B) overlap to fill this intermediate-lifetime region.

a. Invisible decay searches for $B^+ \rightarrow K^+\nu\bar{\nu}$ or $K^+ \rightarrow \pi^+X$, where X decays invisibly, peak in a specified mass range. Below the muon threshold, these are sensitive to the highest F 's probed in accelerator experiments.

- $B^+ \rightarrow K^+ + inv.$: A search at Belle [50] set an upper limit of 1.4×10^{-5} on the rate of $B^+ \rightarrow K^+\nu\bar{\nu}$. This search imposes an upper limit on the energy deposited in the electromagnetic calorimeter (ECL), which extends to 1.65 m from the beamline. Therefore, for this limit we consider only a decays outside this radius, assuming a typical transverse momentum of m_B/m_a . A tighter limit could likely be set by searching for a narrow invisible resonance in the B^+ decays.
- $K^+ \rightarrow \pi^+ + inv.$: The search for $K^+ \rightarrow \pi^+\nu\bar{\nu}$ at Brookhaven E787 [51] also set an explicit limit on the decay $K^+ \rightarrow \pi^+X$ where X is invisible. The branching fraction for this mode must be below about 5×10^{-11} for X masses below about 100 MeV, and

about 10^{-9} for X masses between 150 and 250 MeV. No limit is set between 100 to 150 MeV, where several $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ candidates were seen. The transverse size of the detector is roughly 1.3 m, and the kaons are produced at rest so that $m_K/2$ is a good approximation to the energy of the outgoing PNGB.

Of these two, the K^+ decay limit extends to higher F but the B^+ decay search is able to probe lower F because the PNGBs produced in B decays are more boosted and therefore longer-lived. There are additional constraints from CLEO and BaBar searches for $\Upsilon(1s)$ and $\Upsilon(3s)$ decays to γa [56–58], but this region is largely contained in the two identified above.

b. Visible decays Again, we focus here only on the most powerful visible-decay searches:

- $K^+ \rightarrow \pi^+ X$, $X \rightarrow e^+ e^-$: This resonant decay mode was searched for in the $K_{\mu 2}$ experiment at KEK [52], which excluded kaon branching fractions to below about 1.5×10^{-6} for PNGB masses between 10 and 80 MeV, and below about 5×10^{-6} for PNGB masses from 80 to 120 MeV. The experiment also set a limit above 140 MeV, but it is not as constraining as the B^+ -decay limits discussed below, so we do not include the constraint in this region for this mode. $K_{\mu 2}$ used stopped kaons, so we assume an initial energy of $m_K/2$. No vertex requirement is explicitly mentioned in the analysis, but common vertex requirements of order a few cm are frequently imposed in spectrometer analyses typical in these experiments; in any case, at much larger distances, the mass resolution would likely be degraded. We have used an *ad hoc* but conservative estimate of 1 cm to produce the limit in Fig. 2 but even if the vertex requirement is much tighter (as tight as 0.5 mm) the excluded region would overlap that of the $B^+ \rightarrow K^+ + inv.$ search.
- $B^+ \rightarrow K^+ \ell^+ \ell^-$: BaBar [53] and Belle [54, 59] have both measured the rate of the rare decay $B^+ \rightarrow K^+ \ell^+ \ell^-$ for $\ell = e, \mu$. The observed branching fractions, $(3\text{--}6) \times 10^{-6}$, are consistent with Standard Model predictions, and can be translated into rough (conservative) limits on $B^+ \rightarrow K^+ a$ by requiring that this exotic decay not exceed the total measured rate (we focus on the decays to K^\pm rather than K^* because the observed rates are lower). The BaBar measurement includes lower-mass electron pairs, down to 30 MeV (compared to 140 MeV at Belle). The most recent Belle analysis [54] bins events by invariant mass, so that we can obtain a tighter limit ($\approx 10^{-7}$) on

$B^+ \rightarrow K^+ a$ in the region of interest, 140-1440 MeV. In all cases we take the limit to be the central value plus 2σ . The BaBar measurement required that the $\ell^+\ell^-$ pair originate from the same vertex as the K^+ . To set a conservative limit we require that the a decay within $100\mu\text{m}$, the scale of BaBar’s vertex resolution. For the Belle analysis we require only that the PNGB decay within 0.5 cm, which would pass the requirement of [59].

- Similar but slightly weaker limits are obtained from measurements of $K \rightarrow \pi\ell^+\ell^-$, e.g. [60–62]. We refer the reader to the original results and to [43] for details.

It is worth emphasizing that the crude limits we have obtained are far weaker than the tightest limits that *could* be obtained by a directed analysis of BaBar or Belle data. Firstly, much tighter limits could be obtained by binning the 20–100 events in each sample more finely, and accounting for the detection efficiency as a function of mass. Further improvement could be obtained by including more displaced a decays, at the edge of the inner tracker (3 cm) or even beyond. *This direction is particularly worthy of exploration for the $B \rightarrow K\mu^+\mu^-$ mode, for which there are no complementary searches near the high- F boundary of the Belle-excluded region, and for which the muon system gives an additional handle for studying highly displaced decays.*

B. Limits from Supernova SN 1987a

For completeness we also include the constraints on PNGBs from SN 1987a [63]. We adapt the analysis of [64] to our setup, and obtain limits based on the assumption that PNGBs must not be the dominant mechanism of energy loss from the supernova. The temperature of the supernova core is conservatively estimated at $T \sim 30$ MeV [64], so PNGBs with mass significantly greater than these energies cannot be produced by the supernova, and are therefore unconstrained. The flux of PNGBs from the core of the supernova is approximately [64]

$$\frac{dN_a}{dE_a} \sim 10^{71} \left(\frac{1 \text{ GeV}}{F} \right)^2 e^{-E_a/T} \frac{1}{\text{GeV}} . \quad (7)$$

However, if these PNGBs decay or are re-absorbed, then they will not escape from the supernova and so they will not carry away any energy. Scattering and re-absorption dominate




FIG. 2: **Left:** Constraints on pseudo-Nambu-Goldstone bosons as a function of their decay constants F and their mass m_a from various meson decays: $K^+ \rightarrow \text{anything} + e^+e^-$ (green), $K^+ \rightarrow \pi^+ + \text{invisible}$ (blue), $B^+ \rightarrow K^+\ell^+\ell^-$ (yellow) ($\ell = e, \mu$), and $B^+ \rightarrow K^+ + \text{invisible}$ (red). Constraints from $\Upsilon(1S)$ or $\Upsilon(3S) \rightarrow \gamma a \rightarrow \gamma + \text{invisible}$ and $K^+ \rightarrow \pi^+\ell^+\ell^-$ decays are weaker than those from $B^+ \rightarrow K^+ + \text{invisible}$ and $B^+ \rightarrow K^+\ell^+\ell^-$, respectively, and thus not shown. Details are in §II A. **Right:** Gray shaded background region is the combined exclusion region from meson decays in the left figure. In the green exclusion region, the proton beam dump experiment CHARM at CERN would have seen at least five events (this exclusion region agrees roughly with that in [2]) – see §IV. Here the PNGB is produced directly in the proton dump by a small mixing with the pion. For $m_a < 2m_\mu$, the PNGB decays to an electron pair, while in the “bubble” for $m_a > 2m_\mu$ the PNGB decays predominantly to a muon pair. The blue region is the limit from the supernova SN 1987a (see §IIB). The light red region is the constraint from the muon anomalous magnetic moment and fills the gap for low m_a and F left by the meson constraints (see §II C). The region excluded by the Fermilab E137 dump lies mostly within the CHARM excluded region and is not shown (it is instead shown in Fig. 5).

over PNGB decay for the relevant region of parameter space, giving a mean free path of

$$\lambda_{\text{mfp}} \sim 10 \text{ m} \left(\frac{F}{10^6 \text{ GeV}} \right)^2 \quad (8)$$

for PNGBs with $F < 10^8$ GeV. We see that for F significantly smaller than 10^6 GeV, the PNGB mean free path is much less than the estimated core size of 10 km, so for these smaller values of F , SN 1987a does not constrain the PNGB. The exclusion contours we have plotted in Fig. 2 (right) correspond to requiring that the PNGBs carry away less energy than neutrinos, meaning that the total integrated PNGB emission must be less than about 10^{53}

PNGBs, each with energy of order T .

C. Limits from the anomalous muon magnetic moment

PNGB's contribute to the anomalous magnetic moment of the muon, a_μ , at the loop level. For the mass range of interest in this paper ($\lesssim 1$ GeV), only the one-loop contribution is important, and it is given by (see e.g. [43, 65])

$$a_\mu^a = -\frac{1}{8\pi^2} \frac{m_\mu^2}{F^2} \int_0^1 dx \frac{x^3}{x^2 + (1-x)\frac{m_a^2}{m_\mu^2}}. \quad (9)$$

While the experimental measurement of a_μ is rather precise [66], the Standard Model prediction involves a hadronic contribution that must be estimated from experiments, which do not all agree. Using data from e^+e^- annihilation to hadrons, the theoretical value of a_μ is smaller than the measured value by $(316 \pm 79) \times 10^{-11}$ [67], which is a 4σ discrepancy. However, estimates from τ 's give a smaller disagreement, with [68] finding a difference of $(157 \pm 82) \times 10^{-11}$, which is a 1.9σ discrepancy.

Since the contribution from PNGB's is negative, a very conservative limit is obtained by using the 5σ lower bound in [68], i.e.

$$\begin{aligned} a_\mu^a &\geq (157 - 5 \times 82) \times 10^{-11}, \text{ i.e.} \\ a_\mu^a &\geq -253 \times 10^{-11}. \end{aligned} \quad (10)$$

This constraint is included in Fig. 2 (right) and 5.

III. CONSTRAINTS FROM LSND ON PNGBS AND DARK GAUGE BOSONS

The Liquid Scintillator Neutrino Detector (LSND) experiment ran at the Los Alamos Neutron Scattering Center (LANSCE) in the 1990's [69], and dumped $\mathcal{O}(10^{23})$ 800 MeV protons on a predominantly water-copper target. This produced $\sim 10^{22}$ pions, a very large number that allows LSND to be sensitive in principle to very weakly coupled PNGBs or gauge bosons. A detector of length 8.3 m and a diameter of 5.7 m was located 29.7 m away from the target, 12° off-axis [84], and filled with dilute liquid scintillator (there was no open

	E_{beam} (GeV)	N_p	X_t (m)	X_d (m)	$n_{\pi^0}\epsilon_{\text{geo}}$	\bar{E}_a (GeV)
CHARM [2]	400	2.4×10^{18}	480	515	0.12	25
LSND [69–71]	0.8	$\sim 10^{23}$	29.7	38	see text	0.3
MINOS / MINERvA [72, 73]	120	3.8×10^{20}	1050	1087	0.0006	20
MiniBooNE [74]	8.9	10^{21}	541	553	0.002	2.7

TABLE III: Shown are the total number of incident protons N_p , the distance from the target to the open decay region in front of the detector (i.e. the thickness of the shield) X_t , the distance from the target to the end of the detector X_d , the geometric acceptance times the number of pions per incident proton $n_{\pi^0}\epsilon_{\text{geo}}$, and the median PNGB energy E_a . These numbers were used to calculate the sensitivity of CHARM, MINOS/MINERvA (we always use the larger MINOS detector for estimates), and MiniBooNE to PNGBs produced directly in the proton dump. For LSND, we used a more involved procedure described in the text.

decay region in front of the detector). Although the LSND collaboration did not search for signals that could originate from decays of long-lived exotics, approximate limits can be extracted from two published LSND analyses.

We begin by reviewing the production of PNGB’s in proton-nucleus collisions (which will re-appear in our discussion of more recent neutrino experiments in §IV), then discuss the implications of two specific LSND analyses [70, 71] for PNGBs that decay to e^+e^- pairs or $\mu^+\mu^-$ pairs, respectively. The calculation involves considerable uncertainties and assumptions, and should be taken only as a sensitivity estimate — a dedicated analysis by the LSND collaboration is required to obtain a reliable limit. However, this analysis demonstrates that a dedicated analysis by the LSND collaboration would set some of the tightest constraints on PNGBs that couple to both quarks and leptons.

We next discuss the production of dark gauge bosons in pion decay, and the implications of the same two analyses for these models. Our results are consistent with [7], but we have considerably elaborated the discussion of the experimental sensitivity. The LSND sensitivity for this model overlaps closely with that of the SLAC electron beam-dump experiment E137 [4], as computed in [6].

A. Production of PNGBs at LSND and other Proton Beams

In this section, we will consider the production and experimental sensitivity to pseudoscalars from proton beam dumps. We will focus on the LSND experiment because it can

set the most stringent limits, although the MINOS and MiniBooNE experiments can also set interesting limits. We will discuss them more extensively in §IV.

Proton beam dumps can produce pseudoscalars directly through their mixing with pions. If the pseudoscalars couple to quarks they will interact via the operator

$$\frac{m_q}{F} a \bar{q} q \quad \Longrightarrow \quad c \frac{m_\pi^2 F_\pi}{F} a \pi^0 \quad (11)$$

where c is an $O(1)$ parameter that depends on the up and down quark masses and any coefficients in the pseudoscalar coupling to the quarks. Since the pseudoscalar mixes with the pion, for every pion that is produced through a QCD process there is a probability of approximately (for $c = 1$)

$$\left(\frac{F_\pi}{F} \right)^2 \quad (12)$$

of instead producing a PNGB.

We can estimate the production rate within the detector acceptance and the PNGB momentum distribution by using the measured rates for π^0 :

$$N_a = \left(\frac{F_\pi}{F} \right)^2 n_{\pi^0} N_p \epsilon_{\text{geo}}. \quad (13)$$

Here, n_{π^0} is the number of pions produced per incident proton, N_p is the total number of protons, and ϵ_{geo} is the geometric acceptance (the solid angle subtended by the detector at the target divided by the solid angle of the beam). In practice, only the product $n_{\pi^0} N_p \epsilon_{\text{geo}}$ is relevant but we have tabulated estimates for $n_{\pi^0} \epsilon_{\text{geo}}$ and N_p in Table III for various experiments, for the sake of comparison. To determine the number of observable e^+e^- pairs, we must also account for the probability of decaying in the detectable region (either inside the detector or in an open region upstream of the detector):

$$N_e = N_a \left(e^{-\frac{X_t}{\gamma c \tau}} - e^{-\frac{X_d}{\gamma c \tau}} \right) \simeq N_a \frac{X_d - X_t}{\gamma c \tau} e^{-\frac{L_0}{\gamma c \tau}}, \quad (14)$$

where X_t and X_d are the minimum and maximum decay lengths (roughly, X_t is the thickness of the shield and X_d is the distance from the target to the end of the detector), and in the second expression we have assumed that $\gamma c \tau \gg X_d - X_t$.




FIG. 3: Sensitivity of various neutrino experiments to pseudo-Nambu-Goldstone bosons as a function of their decay constants F and their mass m_a . The thick (thin) black solid line corresponds to 10 (1000) events in LSND, the thick (thin) dashed blue line corresponds to 3 (1000) events in MiniBooNE, and the thick (thin) dot-dashed red line corresponds to 3 (1000) events in MINOS/MINERvA (in each case the inner regions correspond to more events than indicated by the line). Here the PNGB is produced directly in the proton dump by a small mixing with the pion. For $m_a < 2m_\mu$, the PNGB decays to an electron pair, while in the “bubbles” for $m_a > 2m_\mu$ the PNGB decays predominantly to a muon pair. The gray shaded regions are the combined existing constraints from other beam dump experiments, meson decays, anomalous muon magnetic moment, and SN 1987a shown in the right plot of Fig. 2.

B. LSND Analyses Sensitive to PNGBs

We now focus on two LSND results and their implications for PNGB’s.

PNGBs decaying to e^+e^- : The analysis in [70] used $\sim 0.92 \times 10^{23}$ protons on target and looked for $\nu_\mu \rightarrow \nu_e$ oscillations using ν_μ from π^+ decay in flight [85]. The ν_e are detected via the inclusive charged-current reaction $\nu_e + {}^{12}\text{C} \rightarrow e^- + X$. This analysis focused on identifying electrons in the energy range 60 MeV to 200 MeV. Various cuts were used in the analysis with an energy-dependent efficiency that is always near 10%. Clearly this analysis

should be sensitive to PNGBs decaying to electrons inside the LSND detector, although it is impossible to accurately estimate the efficiency without a dedicated analysis. One difference between our signal and the study in [70] is that in our case both an electron and a positron are produced in the detector, as opposed to just a single electron. However, it is impossible to distinguish e^+e^- events from a single electron (or single photon) event, and it has been suggested [75] that we should assume that the total e^+e^- pair energy (i.e. the PNGB energy) would have been measured as the energy of a “single-electron”. For simplicity, we will assume that the detection efficiency for an e^+e^- pair is roughly the same as for the single electron analysis, i.e. $\epsilon_{\text{eff},1} \sim 0.1$.

We have to estimate the energy distribution and number of PNGBs incident on the LSND detector. This should be roughly equivalent to the number and energy distribution of pions, which has been simulated with Monte Carlo by the LSND collaboration [75]. Specifically, we will model the PNGB kinetic energy distribution using the predicted pion kinetic energy distribution from [75], and rescaling the total rate by F_π^2/F^2 from the total number of pions incident on the detector, $(8.6 \pm 2.1) \times 10^{14} \text{ cm}^{-2}$. This should be an excellent approximation for $m_a \approx m_\pi$; it may be subject to additional $\mathcal{O}(1)$ uncertainties for $m_a \ll m_\pi$ and especially for $m_a \gg m_\pi$, which we neglect. In addition to the reconstruction efficiency, we must also account for the fraction of PNGBs with kinetic energy between $60 - m_a$ and $200 - m_a$, which for small m_a and the assumed distribution is approximately 25%. We note that the mean of the pion kinetic energy distribution is at ~ 275 MeV with a root-mean-square spread of ~ 130 MeV, so an analysis including higher-energy electrons would be significantly more sensitive.

We show the number of e^+e^- events obtained from PNGB decays in the F versus m_a parameter space in Fig. 3, where the solid black thick and thin lines for $m_a < 2m_\mu$ show 10 and 1000 e^+e^- events, respectively. We have assumed that the number of PNGB decays inside the detector is given by Eq. (14) integrated over an PNGB energy from $60 \text{ MeV} - m_a$ to $200 \text{ MeV} - m_a$, and multiplied by the efficiency $\epsilon_{\text{eff},1}$. The analysis of [70] (see their Fig. 29) indicates that 10 PNGB events in a 20 MeV energy bin below 200 MeV would have been easily noticed. We note that the sensitivity could have been increased by increasing the energy threshold for electrons beyond 200 MeV.

PNGBs decaying to $\mu^+\mu^-$: We now turn our attention to the analysis in [71], which would have been sensitive to $\mu^+\mu^-$ events. The analysis in [71] considered $\sim 1.8 \times 10^{23}$ protons on

target and searched for $\pi^0 \rightarrow \nu_\mu \bar{\nu}_\mu$. The neutrinos interact with the nuclei in the detector to produce muons through the reactions $\nu_\mu + {}^{12}\text{C} \rightarrow \mu^- + p + X$ and $\bar{\nu}_\mu + {}^{12}\text{C} \rightarrow \mu^+ + n + X$. This analysis focused on identifying high-energy, muon-like beam excess events in the energy range 160 MeV to 600 MeV electron equivalent. The muon was required to decay inside the detector. Various selection cuts were used in the analysis [71] with an overall efficiency of ~ 0.15 for identifying neutrino interactions in the detector. This analysis should clearly have been sensitive to PNGBs decaying to $\mu^+\mu^-$ pairs inside the LSND detector, although it is impossible to accurately calculate the efficiency without a dedicated analysis. We simply estimate the efficiency for $\mu^+\mu^-$ pairs to be reconstructed as single-muon events to be $\epsilon_{\text{eff},2} \sim 0.1$, a little bit less than the efficiency for the original analysis. We estimate the PNGB kinetic energy distribution as described above. To obtain the total number of pions incident on the detector, we rescale the value we used above for the $a \rightarrow e^+e^-$ analysis to $1.7 \times 10^{15} \text{ cm}^{-2}$, since this analysis uses a factor of ~ 2 more protons. Integrating over an PNGB energy from $160 \text{ MeV} - m_a + 2m_\mu$ to $600 \text{ MeV} - m_a + 2m_\mu$, we show the number of $\mu^+\mu^-$ events obtained from PNGB decays in the F versus m_a parameter space in Fig. 3, where the solid black thick and thin lines for the region $m_a > 2m_\mu$ show 10 and 1000 $\mu^+\mu^-$ events, respectively. The analysis in [71] indicates that 10 PNGB events in a given energy bin would have been noticed, and so we believe that this is a reasonable number of events with which to set a tentative limit.

C. LSND Limits on Light Vector Bosons

In [7] it was shown that LSND may be sensitive to a new light vector boson A' that mixes with $U(1)$ hypercharge; however, the precise sensitivity was unclear due to uncertainties in the LSND experiment. We believe that we have clarified these uncertainties using information from [70, 71, 75].

With approximately 1.8×10^{23} protons on target at 800 MeV, the LSND experiment produced about 1.4×10^{22} π^0 's. The vast majority of these π^0 's decay to $\gamma\gamma$, so with a kinetic mixing parameter ϵ we expect a branching fraction of $2\epsilon^2$ for $\pi^0 \rightarrow \gamma A'$. We will only

consider the simplest scenario, where the A' decays to e^+e^- with a lifetime

$$\tau_{A'} = \frac{3}{\alpha\epsilon^2 m_{A'} \sqrt{1 - \frac{4m_e^2}{m_{A'}^2} \left(1 + \frac{2m_e^2}{m_{A'}^2}\right)}} \quad (15)$$

that is also set by the kinetic mixing. Due to the huge number of pions produced at LSND, the A' may be visible even for very small ϵ .

To set a reliable limit, we use the results of the $\pi^0 \rightarrow \nu\bar{\nu}$ search discussed previously [71]. They performed a Monte Carlo simulation to determine the flux of neutrinos in the LSND detector from this decay and plotted it in their Fig. 1 [71]. In our case we want to consider $\pi^0 \rightarrow \gamma A'$, but in the limit that $m_{A'} \ll m_\pi^0$ the neutrino results should give a very good approximation to the distribution of A' particles; beyond this limit we will continue to use the neutrino results with the caveat that we expect an $O(1)$ uncertainty in the rates. Given this distribution of A' particles, we computed the number of A' 's that would have decayed inside the LSND detector using the lifetime formula above, and the $2\epsilon^2$ branching ratio of $\pi^0 \rightarrow \gamma A'$.

The only reliable data that can exclude an A' decaying to e^+e^- inside the LSND detector is the analysis of [70], which we discussed above in order to determine the LSND sensitivity to PNGBs. From the study of π^0 backgrounds in [70] and from discussions with LSND collaborators [75], we expect that an A' decay to e^+e^- would have been interpreted as a single-electron event in the LSND detector. The data on these types of events has only been published for energy depositions below 200 MeV, and the analysis in question had an overall reconstruction efficiency of about 10%. Thus we use the Monte Carlo data [70, 75] below this energy to set our tentative limits, which are displayed in Fig. 4. Our limits are based on the conservative assumption that a signal of 50 or more events after the 10% efficiency cut should have been visible.

IV. SENSITIVITY OF MODERN NEUTRINO EXPERIMENTS

A. Production from Proton-Nucleus Collisions

We have discussed this production mode already in the context of LSND, in §III A. A similar analysis was also carried out by the CHARM experiment at CERN [2], a proton




FIG. 4: The sensitivity of various beam dump, collider, and astrophysical probes of light vector bosons that kinetically mix with the standard model hypercharge, as a function of the kinetic mixing parameter ϵ and the vector boson mass $m_{A'}$ (from [6]). The thick (thin) solid black line corresponds to 10 (1000) signal events in LSND. The region enclosed by the thick line can be viewed as a very rough exclusion limit from the LSND experiment. A re-analysis of the LSND data by the LSND collaboration could further extend the sensitivity of that experiment [7]. Further details are described in the text.

beam dump with a 400 GeV beam, a 35 m decay region, and a 3×3 meter detector. Similar searches could be done using the MINOS or MINERvA detectors on the NuMI beamline, and with the MiniBooNE detector. We summarize the parameters for these experiments in Table III.

In all cases we estimate the production of PGNBs by rescaling the pion production rate (determined either experimentally or from Monte Carlo) by a factor of $(F_\pi/F)^2$. However, the production rate of very weakly coupled PGNBs could actually be significantly greater than this scaled rate because PGNBs produced within the target will virtually always escape from the target, whereas many of the pions produced by proton interactions might not make it out of the beam dump. This issue may merit further exploration.

For the CHARM experiment, we take the experimental parameters from [2]. Our calcu-

lation of the PNGB exclusion region is shown in Fig. 2 (right) and roughly agrees with the results in [2].

The MiniBooNE estimate is shown as a thick (thin) blue dashed line in Fig. 3, corresponding to 3 (1000) events. It is obtained as follows. We take the π^0 production cross section to be the average of the π^+ and π^- production cross sections found in [74]. These cross sections are given as a function of angle and momentum, and one can show that the fraction of π^0 's produced in the target that point towards the 12 m detector at a distance of 541 m away is roughly 0.002. In addition, [74] gives the average number of π^+ and π^- produced per “particle-producing reaction”, which we take to be per incident proton on target – we average these to find about 0.89 π^0 's per incident proton.

The MINOS/MINERvA estimate is shown as a thick (thin) red dot-dashed line in Fig. 3, again corresponding to 3 (1000) events. This estimate for MINOS was obtained by using data on their π^+ momentum distributions from Fig. 9.3 of [72]. Based on the assumption that the momentum distribution for a general PNGB would be similar, we estimate from that figure that a fraction 0.0035 of PNGBs produced by the NuMI proton dump would point towards the MINOS detector. Data from [72] shows that there were roughly 0.18 pions produced per proton on target, which allows us to estimate the total number of PNGBs produced.

Searches in either MiniBooNE and MINOS/MINERvA would have significant overlap with the existing CHARM result, perhaps slightly extending the region probed. However, the proposed “Project X” upgrades to NuMI, which would increase the beam intensity by a factor of ~ 5 -10 and possibly include a larger near detector could extend sensitivity to larger F .

B. Production through Muon a -sstrahlung

A more unique search opportunity at MINOS/MINERvA and MiniBooNE arises from their magnetic focusing of charged pions, which in turn focuses the neutrinos and muons from their decay towards the detectors. As a result of this focusing, approximately one in a thousand protons on target produces a muon that points toward the neutrino detector. As these muons stop in the rock upstream of the detector, they can radiate very forward PNGB's in a process that is analogous to ordinary bremsstrahlung. Unlike photon bremsstrahlung, however, the typical PNGB produced by this process carries a large fraction of the incident

	E_{beam} (GeV)	N_ℓ	\bar{E}_ℓ (GeV)	X_t (m)	X_d (m)	W (m)	E_{thresh} (GeV)
E137 [4]	20.0	$N_e = 1.2 \times 10^{20}$	20	180	380	3	2.0
MINOS/MINERvA [76]	120.0	$N_\mu = 2.7 \times 10^{17}$	20	240	270	3	1.0
MiniBooNE [77]	8.9	$N_\mu = 2.0 \times 10^{17}$	1.3	450	462	12	0.1

TABLE IV: Shown are the proton beam energy E_{beam} , the total number of incident electrons (for E137) or muons (for MINOS/MINERvA and MiniBooNE) N_ℓ and their average energy \bar{E}_ℓ , the distance from the start of the electron or muon dump to the open decay region (if any) in front of the detector or to the detector itself (i.e. the thickness of the shield) X_t , the distance from the start of the muon dump to the end of the detector X_d , the diameter of the detector W , and the threshold energy E_{thresh} to detect an electron or muon that originates from an PNGB decay. These numbers were used to calculate the sensitivity of E137 [4], MINOS/MINERvA [76] (we always use the larger MINOS detector for estimates), and MiniBooNE [77] to PNGBs produced by bremsstrahlung off an electron beam (in the case of E137) or off a muon beam that is produced in a proton dump.

muon's energy. The production rate and kinematics can be reliably calculated using the Weizsäcker-Williams approximation [78], where the nuclei in the fixed target provide an effective photon beam. We have relegated the details to Appendix A, so in this section we will simply give the results along with their physical motivation.

We will be interested in a wide range of values for m_a ; with this in mind we note that production is dominated by emission angles

$$\theta_a \lesssim \max\left(\frac{m_a}{E_0}, \frac{m_\mu}{E_0}\right) \quad (16)$$

for $E_a \sim E_0$. This is very useful for estimating the angular acceptance, which we compute by taking the ratio of the solid angle subtended by the detector to the solid angle within θ_a of the beam direction. We also considered acceptance issues associated with the PNGB decay, but because we require the PNGB to decay within the MiniBooNE detector or in the relatively small open decay region in front of the MINOS/MINERvA detector, this is not important, and we ignore it.

The cross section is always peaked near $E_a \approx E_0$, even when $m_a \ll m_\mu$. This is in stark contrast to the analogous formula for photon bremsstrahlung, where the rate is proportional to the inverse of the photon energy. The difference is due to the contrasting soft emission behavior of gauge bosons and goldstone bosons – the former have soft singularities, while emission of the latter vanishes in the soft limit due to their derivative couplings. The




FIG. 5: Sensitivity and constraints of various experiments to leptophilic pseudo-Nambu-Goldstone bosons as a function of their decay constants F and their mass m_a . Here the PNGB is produced by bremsstrahlung off an incident muon or electron beam. Thick (thin) lines show rough sensitivity regions and correspond to 3 (1000) displaced e^+e^- pairs in MINOS/MINERvA (red dot-dashed lines), MiniBooNE (blue dashed lines), and in a thin target experiment using the COMPASS muon beam (green dotted line, see §V). The thick (thin) dotted black lines correspond to $S/\sqrt{B} = 3$ (10), where S (B) are the number of prompt $\mu^+\mu^-$ signal (radiative background) events in COMPASS (see §V) (we have ignored the Bethe-Heitler background and the finite acceptance, so these lines should not be viewed as real significance lines but only as very rough estimates of what could be probed). Inside the gray shaded region, E137 would have seen at least one event – since they saw none, this region gives their approximate constraint. Details are described in the text. The light red region is the constraint from the muon anomalous magnetic moment (see §II C).

peaking toward high energy fraction E_a/E_0 is further enhanced by a larger phase space when $m_a \gg m_\mu$. The total cross section for pseudoscalar production from a muon beam has the parametric form

$$\sigma \approx \frac{m_\mu^2}{F^2} \frac{2\alpha^2}{\max(m_\mu^2, m_a^2)} \quad (17)$$

Note that the formula in the case of an electron beam dump would be identical except with

$m_\mu \rightarrow m_e$, so we see that muons are a much more efficient source of PNGBs as long as $m_a \gg m_e$.

To get a rough idea of our experimental reach, it is useful to have approximate formulas for the total production rates. For a thin target, the yield is

$$N_a \sim N_\mu \frac{m_\mu^2}{\alpha F^2} \frac{T_e m_e^2}{\max(m_\mu^2, m_a^2)} \quad (18)$$

where T_e is the number of electron radiation lengths of material. In the case of a thick target, where the beam of muons is completely dumped, it is more difficult to give a simple parametric formula for the pseudoscalar yield because muons stop due to minimum ionization interactions, as opposed to bremsstrahlung. However the equation above with $T_e \sim 100$ provides a rough estimate of the pseudoscalar yield for MINOS and MiniBooNE.

Fig. 5 shows the number of e^+e^- pairs in MINOS/MINERvA (red dot-dashed lines) and MiniBooNE (blue dashed lines) as a function of the PNGB decay constant F and the PNGB mass m_a . The thick (thin) lines correspond to 3 events (1000 events). The gray shaded region corresponds to the approximate constraint from E137: inside the region, more than 1 e^+e^- event would have been seen (note that we calculated this region using the procedure in [6], but changing the couplings from an A' to those relevant for PNGBs). Our estimate for the E137 constrained region agrees with that in [4]. We see that MINOS/MINERvA can extend the E137 region, although the MiniBooNE region is contained within the E137 region since many PNGBs are produced in the dump with a large enough angle causing them to miss the detector (this is because the MiniBooNE muon beam has a lower energy and their detector is further away compared to MINOS/MINERvA). Note that these lines have been calculated with the full formula as detailed in the appendix. The NOVA experiment and “Project X” upgrades to the NuMI beamline will have a factor of 5–10 more protons and so will significantly extend the reach. The light red region is the constraint from the muon anomalous magnetic moment, while the other lines will be discussed in §V.

V. MUON FIXED-TARGET EXPERIMENTS WITH THIN TARGETS – COMPASS

In this note, we have argued that new light states may be produced by bremsstrahlung off a dumped muon beam. Here we will briefly comment on the potential of using a fixed target setup to constrain weakly coupled light states. Compared to a dump experiment, the target will be much thinner, and we can thus search for particles with a much shorter lifetime.

Electron fixed target experiments have been used to probe new light states [6, 26, 37, 79], but muon fixed-target setups are different in several respects. The intensity of muon beams is obviously lower, but this may be partly compensated for by a much thicker target — a muon beam can easily traverse a meter of material leaving a relatively quiet off-beam environment. Furthermore, muon beams will have an obvious advantage over electrons in PNGB searches, as their couplings are proportional to particle mass. For example, it is instructive to consider the COMPASS experiment at CERN [80], where a 160 GeV muon beam strikes a low- Z polarized target (they have also used higher- Z , but much thinner, targets). A detailed description of their polarized target is given in [80], and we approximate it as a 130 cm long target consisting of Lithium, with a packing factor of 0.5 (i.e. a column density of 34.7 g/cm², or 0.42 radiation lengths). In total, we estimate that this experiment has collided about $\sim 10^{15}$ muons. Here we will illustrate the rough reach of this COMPASS-like setup assuming the Lithium target. We have checked that a higher Z target, such as Tungsten, of a similar thickness would probe a somewhat larger parameter space, including high- F regions above the muon threshold. We will not show this in any figures, focusing instead on the data set already collected by COMPASS.

There are two regions of parameter space to consider, one in which F is low and the PNGB decays promptly, and one in which F is high and the PNGB decays with a displaced vertex. In the case of a prompt decay, one must search for a peak in the di-muon invariant mass on top of a sizable standard model background (decays to electrons may not be searched for this way because their interactions in the target considerably degrade the mass resolution). The signal to background ratio (where the background comes from di-lepton production via

an off-shell photon) for such a search is roughly [6]

$$\frac{S}{B} \sim \frac{m_\mu^2}{F^2} \frac{3}{4\alpha^2} \frac{m_a}{\delta m}, \quad (19)$$

where δm is the resolution-limited mass window used in the search. Based on the angular and momentum resolutions for COMPASS [81], we infer a mass resolution $\sigma_m = 11$ MeV, making a window width $\delta m \approx 2.5\sigma_m = 27$ MeV appropriate for estimating sensitivity. Additional two-photon diagrams (which we refer to as the ‘‘Bethe-Heitler’’ background) and their interference have been neglected; these actually dominate over the ‘‘radiative’’ dilepton production assumed in 19, but can be reduced to an order-1 fraction of the background by kinematic selection.

We can combine Eq. (19) with the expected number of signal events using the formulas in Appendix A to get a very rough estimate for the sensitivity of such a search. These estimates omit several factors, all of which degrade sensitivity: the finite detector acceptance, the additional Bethe-Heitler background, and the finite acceptance of kinematic selection needed to veto the dominant Bethe-Heitler background. Assuming 10^{15} muons on target and a mass window $\delta m = 27$ MeV [81], we show the ‘‘ 3σ ’’ and ‘‘ 10σ ’’ estimated reach in Fig. 5 (the neglected factors are expected to reduce the sensitivity each point in parameter space by a factor of 2–4). Such a search in existing COMPASS data (or that of a similar experiment) would be sensitive to leptophilic PNGBs with F near the weak scale and masses between the dimuon threshold and a few GeV. Higher F ’s can be probed by increasing the target thickness in radiation lengths, for example by using a high- Z target of comparable length.

Another parameter range that can be studied by muon fixed-target experiments is that where the PNGB decays are significantly displaced. If the vertex can be reconstructed to be downstream of the target region, beyond the tails of Standard Model backgrounds, a few events could be enough to claim a discovery. The transverse vertex resolution of COMPASS is about 0.1 mm [80], so that the vertex resolution in the direction of the beam is about $\sigma_z = (0.1 \text{ mm}) \times E_{\text{beam}}/m_a$. In Fig. 5, we show the area of parameter space that yields 3 or 1000 events in the region between $5\sigma_z$ and 10 meters behind the 130 cm Lithium target (for this estimate we assume that the PNGB is produced in the middle of the target). Note that the vertex resolution is worse for lower-mass PNGBs, so that for $m_a \lesssim 10 \text{ MeV}$, we

have $5\sigma_z \gtrsim 10$ m, and there is no good search region (which is why the plot cuts off for low m_a). An analysis of the data in COMPASS, or a similar experiment of this type, can cover new ground with respect to E137 or the neutrino experiments at lower F , and can close the window between the latter searches and the muon $g - 2$ limit. With a long higher- Z target such as Tungsten, more PNGBs would be produced so that some region of parameter space at $F \sim 10^4$ GeV above the muon threshold may also give rise to observably large rates of displaced decays.

In summary, the very rough estimates above suggest that a fixed target experiment with a focused muon beam may be able to probe unconstrained regions in PNGB parameter space. In fact, the existing data set of the COMPASS experiment may already be able to set some interesting new limits for leptophilic PNGBs. It may also be worthwhile to consider using more diffuse muon “beams” such as those in neutrino factories in a thin target setup due to their higher intensity. We leave this for future thought.

VI. CONCLUSIONS

We have explored the sensitivity of neutrino oscillation experiments to three types of new light states – vector bosons that kinetically mix with the photon, pseudoscalars that couple to quarks and leptons, and pseudoscalars that couple preferentially to leptons. The first two are strongly constrained by rare decays, fixed target experiments, and supernova 1987a (all of which we have reviewed), whereas there are fewer tests of the third class.

The sensitivity of the LSND experiment to vector bosons was discussed in [7], but the details of the LSND detector and analyses were not considered. We have shown that the analyses of [70, 71] would have been sensitive to vector bosons with mass below $2m_\mu$ and a large range of coupling strengths. This sensitivity has significant overlap with the E137 experiment [4], but LSND does probe a new region [at very weak coupling, and in any case the LSND results serve as an important cross-check. Because LSND dumped a larger number of protons, other neutrino experiments are not as sensitive to the light vector boson scenario.

We also considered the sensitivity of various neutrino experiments to pseudo-scalars that couple to quarks, so that they can be produced in proton beam dumps. LSND remains the most sensitive experiment for pseudoscalars with mass below $2m_\mu$, nearly closing the gap between fixed target experiments and supernova constraints. However, other experiments

are more sensitive to heavier pseudoscalars. The MiniBooNE and MINOS experiments are currently competitive with the best limits on these particles. However, the estimates that we derive for these experiments may be too conservative because we estimate the rate of PGNB production by scaling the pion production rate, which may be an underestimate if many pions get stuck in the beam dump. This issue merits further study.

The MINOS and MiniBooNE experiments produce neutrinos from focused muon beams; the requisite muon beam dumps provide a unique opportunity to search for pseudoscalars that couple preferentially to leptons, since we expect these particles to couple far more to muons than to electrons. The MINOS experiment is sensitive to leptophilic pseudoscalars with decay constants almost an order of magnitude greater than any previous experiment (in particular E137), while MiniBooNE can probe a region that is contained within that of E137.

A thin target experiment with a muon beam, such as that available in COMPASS, offers a unique probe for leptophilic PGNBs. An analysis using the existing COMPASS data set and looking for either e^+e^- pairs originating from displaced vertices behind their (Lithium) target or for a spike in the $\mu^+\mu^-$ invariant mass spectrum of muons coming from the target should be sensitive to new regions of parameter space. A similar experiment using a higher- Z target would have even more sensitivity.

Upgrades to the NuMi beamline followed by the proposed “Project X” experiment will explore new parameter space for both standard and leptophilic pseudoscalars. It is our hope that in the future these experiments will perform dedicated analyses to explore and constrain new weakly coupled low-mass particles.

Acknowledgements

We thank Philip Schuster for collaboration in early stages of this work. We also thank James Bjorken, Joe Lykken, Chris Polly, Geoff Mills, Simona Murgia, Michael Peskin, Ronald Ransome, Brian Rebel, Byron Roe, David Schmitz, Tomer Volansky, Jay Wacker, Hywel White, and GERALYN Zeller for very useful discussions. We especially want to thank William Louie for extensive discussions and correspondence about the LSND detector and publications. RE and JK are supported by the US DOE under contract number DE-AC02-76SF00515. Fermilab is operated by Fermi Research Alliance, LLC, under Contract DE-

AC02-07CH11359 with the United States Department of Energy. We acknowledge the hospitality of the Aspen Center for Physics where part of this work was done.

Appendix A: Pseudoscalar Production

We are interested in the pseudoscalar production rate from muons braking in a fixed target. This process is analogous to ordinary bremsstrahlung, and it can be reliably calculated using the Weizsäcker-Williams approximation [78], where the nuclei in the fixed target provide an effective photon beam. When the incoming muon has energy E_0 , the differential cross section to produce a pseudoscalar of mass m_a with energy $E_a = xE_0$ is

$$\frac{d\sigma}{dx d\cos\theta_a} = \frac{m_\mu^2 2\alpha^2 E_0 x}{F^2 U^2} \left[x^2 - \frac{2m_a^2 x(1-x)}{U} + \frac{2m_a^2}{U} (m_a^2(1-x)^2 + m_\mu^2 x^2(1-x)) \right] \chi \quad (\text{A1})$$

where

$$U = E_0^2 \theta_a^2 x + m_\mu^2 x + m_a^2 \frac{1-x}{x} \quad (\text{A2})$$

is the virtuality of the intermediate muon in initial state bremsstrahlung, and χ is a form factor that can be found in [78]. We will be interested in a wide range of values for m_a ; with this in mind we see that production is dominated by

$$\theta_a \lesssim \max\left(\frac{m_a}{E_0}, \frac{m_\mu}{E_0}\right) \quad (\text{A3})$$

for $x \sim 1$. For angles larger than this, the differential cross section falls off rapidly, as $1/\theta_a^4$, so this angular scale sets the width of the PNGB beam for the purposes of angular acceptance. We compute the angular acceptance by taking the ratio of the solid angle subtended by the detector to the solid angle within θ_a of the beam direction.

Integrating Eq. (A1) over the angle θ_a , we find

$$\frac{d\sigma}{dx} = \frac{m_\mu^2 2\alpha^2}{F^2 m_\mu^2} x \left[\frac{1 + \frac{2}{3}f}{(1+f)^2} \chi_1(Z) + \left(\frac{1}{3f^2}(1+f) \log(1+f) - \frac{1+4f+2f^2}{3f(1+f)^2} \right) \chi_2(Z) \right] \quad (\text{A4})$$

where

$$f = \frac{m_a^2(1-x)}{m_\mu^2 x^2} \quad (\text{A5})$$

and we have now included the form factors

$$\chi_1(Z) = Z^2 \ln(184 Z^{-1/3}) + Z \ln(1194 Z^{-2/3}) \quad (\text{A6})$$

$$\chi_2(Z) = Z^2 + Z. \quad (\text{A7})$$

It is important to note that the cross section is always peaked near $x \sim 1$, even when $m_a \ll m_\mu$. This is in stark contrast to the analogous formula for bremsstrahlung, which is proportional to $1/x$. The difference is due to the different soft emission behavior of gauge bosons and goldstone bosons – the former have soft singularities, while emission of the latter vanishes in the soft limit. The cross section is dominated for $x \sim 1$ with the parametric form

$$\sigma \approx \frac{m_\mu^2}{F^2} \frac{2\alpha^2}{\max(m_\mu^2, m_a^2)} \quad (\text{A8})$$

The $\max(m_\mu^2, m_a^2)$ factor comes from the presence of the function $f \sim m_a^2/m_\mu^2$ in the denominator for larger m_a/m_μ . The formula in the case of electrons would be identical with $m_\mu \rightarrow m_e$, so we see that as claimed, muons are a much more efficient source of PNGBs as long as $m_a \gg m_e$.

To use these formulae we must account for the way that the muon slows in a beam dump. The number of PNGBs produced per incident muon is

$$\frac{dY}{dx} = \frac{N_A X_e}{A} \int_{E_a}^{E_0} dE_1 \int_0^{T_e} dt_e I_\mu(E_0, E_1, t_e) \frac{d\sigma}{dx'} \quad (\text{A9})$$

where E_0 is the energy of the original incident muons, E_a is the energy of the produced PNGBs, $x' = E_a/E_1$, $x = E_a/E_0$, N_A is Avogadro's number, X_e is the unit (electron) radiation length in g/cm^2 , A is the atomic number of the material, and T_e is the total number of (electron) radiation lengths in the target or beam dump. The function $I_\mu(E_0, E_1, t_e)$ is the distribution of muon energies after the muons have traversed t_e radiation lengths.

It is essential to remember that muons stop primarily through minimum ionizing interactions, and not through radiation, so the number of radiation lengths is not directly related to the muon energy loss. We therefore use electron radiation lengths in our computations, so that t_e and T_e can be much greater than 1 without completely depleting the energy of the muon beam. We estimate the function I_μ using the relevant material properties (if the intervening material is earth, we use Silicon).

To compute the number of electron or muon pairs from PNGB decays in a detector, we simply integrate Eq. (A9) times the probability that an PNGB with this energy decays in (or in some cases in front of) the relevant detector (see Eq. (14)).

-
- [1] B. Holdom, Phys. Lett. **B166**, 196 (1986).
 - [2] F. Bergsma et al. (CHARM), Phys. Lett. **B157**, 458 (1985).
 - [3] A. Bross et al., Phys. Rev. Lett. **67**, 2942 (1991).
 - [4] J. D. Bjorken et al., Phys. Rev. **D38**, 3375 (1988).
 - [5] E. M. Riordan et al., Phys. Rev. Lett. **59**, 755 (1987).
 - [6] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. **D80**, 075018 (2009), 0906.0580.
 - [7] B. Batell, M. Pospelov, and A. Ritz, Phys. Rev. **D80**, 095024 (2009), 0906.5614.
 - [8] P. Schuster, N. Toro, and I. Yavin (2009), 0910.1602.
 - [9] D. P. Finkbeiner and N. Weiner, Phys. Rev. **D76**, 083519 (2007), astro-ph/0702587.
 - [10] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner, Phys. Rev. **D79**, 015014 (2009), 0810.0713.
 - [11] M. Pospelov and A. Ritz, Phys. Lett. **B671**, 391 (2009), 0810.1502.
 - [12] I. Cholis, D. P. Finkbeiner, L. Goodenough, and N. Weiner (2008), 0810.5344.
 - [13] I. Cholis, G. Dobler, D. P. Finkbeiner, L. Goodenough, and N. Weiner (2008), 0811.3641.
 - [14] J. T. Ruderman and T. Volansky, JHEP **02**, 024 (2010), 0908.1570.
 - [15] K. R. Dienes, C. F. Kolda, and J. March-Russell, Nucl. Phys. **B492**, 104 (1997), hep-ph/9610479.
 - [16] C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin (2009), 0902.3246.
 - [17] A. Katz and R. Sundrum (2009), 0902.3271.
 - [18] D. E. Morrissey, D. Poland, and K. M. Zurek (2009), 0904.2567.
 - [19] Y. Cui, D. E. Morrissey, D. Poland, and L. Randall, JHEP **05**, 076 (2009), 0901.0557.
 - [20] S. A. Abel, M. D. Goodsell, J. Jaeckel, V. V. Khoze, and A. Ringwald, JHEP **07**, 124 (2008), 0803.1449.
 - [21] A. Ringwald (2008), 0810.3106.
 - [22] B. Batell, M. Pospelov, and A. Ritz (2009), 0903.0363.
 - [23] R. Essig, P. Schuster, and N. Toro (2009), 0903.3941.

- [24] F. Bossi (2009), 0904.3815.
- [25] P.-f. Yin, J. Liu, and S.-h. Zhu (2009), 0904.4644.
- [26] M. Freytsis, G. Ovanesyan, and J. Thaler (2009), 0909.2862.
- [27] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, JHEP **04**, 014 (2009), 0901.0283.
- [28] C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin (2009), 0909.0290.
- [29] V. M. Abazov et al. (D0), Phys. Rev. Lett. **103**, 081802 (2009), 0905.1478.
- [30] B. Batell, M. Pospelov, A. Ritz, and Y. Shang (2009), 0910.1567.
- [31] P. Schuster, N. Toro, N. Weiner, and I. Yavin (2009), 0910.1839.
- [32] P. Meade, S. Nussinov, M. Papucci, and T. Volansky (2009), 0910.4160.
- [33] P.-f. Yin and S.-h. Zhu (2009), 0911.3338.
- [34] R. Essig, N. Sehgal, and L. E. Strigari (2009), 0902.4750.
- [35] S. Galli, F. Iocco, G. Bertone, and A. Melchiorri, Phys. Rev. **D80**, 023505 (2009), 0905.0003.
- [36] T. R. Slatyer, N. Padmanabhan, and D. P. Finkbeiner, Phys. Rev. **D80**, 043526 (2009), 0906.1197.
- [37] R. Essig, P. Schuster, N. Toro, and B. Wojtsekhowski (2010), 1001.2557.
- [38] R. Essig, N. Sehgal, L. E. Strigari, M. Geha, and J. D. Simon (2010), 1007.4199.
- [39] M. Pospelov and J. Pradler (2010), 1006.4172.
- [40] M. Pospelov (2008), 0811.1030.
- [41] B. Batell, M. Pospelov, and A. Ritz (2009), 0911.4938.
- [42] M. Freytsis, Z. Ligeti, and J. Thaler, Phys. Rev. **D81**, 034001 (2010), 0911.5355.
- [43] S. Andreas, O. Lebedev, S. Ramos-Sanchez, and A. Ringwald (2010), 1005.3978.
- [44] H.-S. Goh and M. Ibe, JHEP **03**, 049 (2009), 0810.5773.
- [45] A. E. Nelson and N. Seiberg, Nucl. Phys. **B416**, 46 (1994), hep-ph/9309299.
- [46] C. Hagmann, H. Murayama, G. G. Raffelt, L. J. Rosenberg, and K. v. Bibber (2008).
- [47] R. Dermisek and J. F. Gunion, Phys. Rev. Lett. **95**, 041801 (2005), hep-ph/0502105.
- [48] Y. Nomura and J. Thaler (2008), 0810.5397.
- [49] C. Amsler et al. (Particle Data Group), Phys. Lett. **B667**, 1 (2008).
- [50] K. F. Chen et al. (BELLE), Phys. Rev. Lett. **99**, 221802 (2007), 0707.0138.
- [51] S. S. Adler et al. (E787), Phys. Rev. **D70**, 037102 (2004), hep-ex/0403034.
- [52] T. Yamazaki et al., Phys. Rev. Lett. **52**, 1089 (1984).

- [53] B. Aubert et al. (BaBar) (2005), hep-ex/0507005.
- [54] I. Adachi et al. (Belle) (2008), 0810.0335.
- [55] P. Ball and R. Zwicky, Phys. Rev. **D71**, 014015 (2005), hep-ph/0406232.
- [56] R. Balest et al. (CLEO), Phys. Rev. **D51**, 2053 (1995).
- [57] W. Love et al. (CLEO), Phys. Rev. Lett. **101**, 151802 (2008), 0807.1427.
- [58] P. del Amo Sanchez et al. (BaBar) (2010), 1007.4646.
- [59] A. Ishikawa et al. (Belle), Phys. Rev. Lett. **91**, 261601 (2003), hep-ex/0308044.
- [60] R. Appel et al. (E865), Phys. Rev. Lett. **83**, 4482 (1999), hep-ex/9907045.
- [61] E. Goudzovski, PoS **KAON09**, 009 (2009), 0908.3860.
- [62] E. Marinova (2010), 1006.2338.
- [63] G. G. Raffelt, Phys. Rept. **198**, 1 (1990).
- [64] A. Burrows, M. S. Turner, and R. P. Brinkmann, Phys. Rev. **D39**, 1020 (1989).
- [65] F. Domingo and U. Ellwanger, JHEP **07**, 079 (2008), 0806.0733.
- [66] G. W. Bennett et al. (Muon G-2), Phys. Rev. **D73**, 072003 (2006), hep-ex/0602035.
- [67] T. Teubner, K. Hagiwara, R. Liao, A. D. Martin, and D. Nomura (2010), 1001.5401.
- [68] M. Davier et al., Eur. Phys. J. **C66**, 127 (2010), 0906.5443.
- [69] C. Athanassopoulos et al. (LSND), Nucl. Instrum. Meth. **A388**, 149 (1997), nucl-ex/9605002.
- [70] C. Athanassopoulos et al. (LSND), Phys. Rev. **C58**, 2489 (1998), nucl-ex/9706006.
- [71] L. B. Auerbach et al. (LSND), Phys. Rev. Lett. **92**, 091801 (2004), hep-ex/0310060.
- [72] A. V. Lebedev (2007), FERMILAB-THESIS-2007-76.
- [73] MINERvA (2006), FERMILAB-DESIGN-2006-01.
- [74] A. A. Aguilar-Arevalo et al. (MiniBooNE), Phys. Rev. **D79**, 072002 (2009), 0806.1449.
- [75] Private Communication with William C. Louie.
- [76] K. Anderson et al. (1998), FERMILAB-DESIGN-1998-01.
- [77] I. Stancu et al. (MiniBooNE) (2001), FERMILAB-TM-2207.
- [78] Y.-S. Tsai, Phys. Rev. **D34**, 1326 (1986).
- [79] M. Reece and L.-T. Wang (2009), 0904.1743.
- [80] P. Abbon et al. (COMPASS), Nucl. Instrum. Meth. **A577**, 455 (2007), hep-ex/0703049.
- [81] E. S. Ageev et al. (COMPASS), Eur. Phys. J. **C41**, 469 (2005), hep-ex/0503033.
- [82] B. Armbruster et al. (KARMEN), Phys. Rev. **D65**, 112001 (2002), hep-ex/0203021.
- [83] A. Aguilar et al. (LSND), Phys. Rev. **D64**, 112007 (2001), hep-ex/0104049.

- [84] The KARMEN experiment [82] had similar features and also dumped a large number of protons, above 9000 C . However the KARMEN detector was 90° off axis and is thus not useful for our purposes.
- [85] This analysis should not be confused with [83] which looked for $\nu_\mu \rightarrow \nu_e$ oscillations from muon decays at rest and saw a the famous LSND excess. This analysis is not suitable for us because the search included a 2.2 MeV photon associated with the capture of the free neutron produced in the charged current reaction.