

REPORT ON PROGRESS

The muon Smasher's guide

To cite this article: Hind Al Ali *et al* 2022 *Rep. Prog. Phys.* **85** 084201

View the [article online](#) for updates and enhancements.

You may also like

- [Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run](#)
B P Abbott, R Abbott, T D Abbott et al.

- [Model comparison from LIGO–Virgo data on GW170817's binary components and consequences for the merger remnant](#)
B P Abbott, R Abbott, T D Abbott et al.

- [A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals](#)
B P Abbott, R Abbott, T D Abbott et al.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Report on Progress

The muon Smasher's guide

Hind Al Ali¹, Nima Arkani-Hamed², Ian Banta¹, Sean Benevides¹,
Dario Buttazzo³, Tianji Cai¹, Junyi Cheng¹, Timothy Cohen⁴,
Nathaniel Craig^{1,*}, Majid Ekhterachian⁵, JiJi Fan⁶, Matthew Forslund⁷,
Isabel Garcia Garcia⁸, Samuel Homiller⁹, Seth Koren¹⁰,
Giacomo Koszegi¹, Zhen Liu^{5,11}, Qianshu Lu⁹, Kun-Feng Lyu¹²,
Alberto Mariotti¹³, Amara McCune¹, Patrick Meade⁷, Isobel Ojalvo¹⁴,
Umut Oktem¹, Diego Redigolo^{15,16}, Matthew Reece⁹, Filippo Sala¹⁷,
Raman Sundrum⁵, Dave Sutherland¹⁸, Andrea Tesi^{16,19}, Timothy Trott¹,
Chris Tully¹⁴, Lian-Tao Wang¹⁰ and Menghang Wang¹

¹ Department of Physics, University of California, Santa Barbara, CA 93106, United States of America

² School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, United States of America

³ INFN, Sezione di Pisa, Largo Bruno Pontecorvo 3, I-56127 Pisa, Italy

⁴ Institute for Fundamental Science, University of Oregon, Eugene, OR 97403, United States of America

⁵ Maryland Center for Fundamental Physics, University of Maryland, College Park, MD 20742, United States of America

⁶ Department of Physics, Brown University, Providence, RI 02912, United States of America

⁷ C. N. Yang Institute for Theoretical Physics, Stony Brook University, Stony Brook, NY 11794, United States of America

⁸ Kavli Institute for Theoretical Physics, University of California, Santa Barbara, CA 93106, United States of America

⁹ Department of Physics, Harvard University, Cambridge, MA 02138, United States of America

¹⁰ Department of Physics and Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, United States of America

¹¹ School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, United States of America

¹² Department of Physics, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong S.A.R., People's Republic of China

¹³ Theoretische Naturkunde and IIHE/ELEM, Vrije Universiteit Brussel, and International Solvay Institutes, Pleinlaan 2, B-1050 Brussels, Belgium

¹⁴ Princeton University, Princeton, NJ 08540, United States of America

¹⁵ Theoretical Physics Department, CERN, Geneva, Switzerland

¹⁶ INFN Sezione di Firenze, Via G. Sansone 1, I-50019 Sesto Fiorentino, Italy

¹⁷ LPTHE, CNRS & Sorbonne Université, 4 Place Jussieu, F-75252 Paris, France

¹⁸ INFN Sezione di Trieste, via Bonomea 265, 34136 Trieste, Italy

¹⁹ Department of Physics and Astronomy, University of Florence, Italy

E-mail: ncraig@ucsb.edu

Received 7 September 2021, revised 3 March 2022

Accepted for publication 12 April 2022

Published 5 July 2022



Abstract

We lay out a comprehensive physics case for a future high-energy muon collider, exploring a range of collision energies (from 1 to 100 TeV) and luminosities. We highlight the advantages of such a collider over proposed alternatives. We show how one can leverage both the point-like nature of the muons themselves as well as the cloud of electroweak radiation that

* Author to whom any correspondence should be addressed.

Corresponding editor: Dr Chiara Mariotti.

surrounds the beam to blur the dichotomy between energy and precision in the search for new physics. The physics case is buttressed by a range of studies with applications to electroweak symmetry breaking, dark matter, and the naturalness of the weak scale. Furthermore, we make sharp connections with complementary experiments that are probing new physics effects using electric dipole moments, flavor violation, and gravitational waves. An extensive appendix provides cross section predictions as a function of the center-of-mass energy for many canonical simplified models.

Keywords: muon, colliders, standard model, beyond the standard model

(Some figures may appear in colour only in the online journal)

Contents

1. Introduction	2	5.2.1. Lepton-flavor violating contact interactions	34
2. Muons vs protons	4	5.2.2. Direct probes of lepton-flavor violation in the MSSM	35
2.1. Muon annihilation	4	5.3. Gravitational waves	37
2.2. Vector boson fusion	6	6. Summary and future directions	39
2.3. Annihilation vs VBF	7	Acknowledgments	40
2.4. Signal vs background	7	Data availability statement	41
3. Muon colliders are gauge boson colliders	9	Appendix A. Simplified models	41
3.1. From the effective vector approximation to PDFs	9	A.1. Standard model	42
3.2. PDFs with broken electroweak symmetry	10	A.2. Supersymmetry	43
3.3. Impact of subleading logs	13	A.3. Vector-like quarks	43
3.4. Finite mass effects	14	A.4. Higgs portal	43
4. Physics	14	A.4.1. Missing mass analysis	45
4.1. Electroweak symmetry breaking	15	A.4.1.4. Missing mass combination	46
4.1.1. Higgs coupling sensitivity estimates from on-shell Higgs processes	15	A.4.2. Invariant mass analysis	46
4.1.2. Flavor and exotic couplings	16	A.5. Hidden valleys	51
4.1.2.1. $W+W \rightarrow tt$: a longitudinal scattering case study	16	A.6. Axion-like particles	51
4.1.3. The Higgs potential and the electroweak phase transition	17	References	52
4.1.4. Additional Higgs bosons	18	1. Introduction	
4.2. Dark matter	19	The discovery of the Higgs boson [1, 2] at the Large Hadron Collider (LHC) marks the end of one era and the dawn of another. The origin of mass has been explained, but in answering this question, the Higgs boson poses a host of others: is this the Higgs of the standard model? Is it the only Higgs, or one of many? Why is electroweak symmetry broken in the first place, and what sets the scale? How, if at all, is the origin of mass connected to the patterns of flavor, the nature of dark matter, or the abundance of matter over antimatter?	20
4.3. Naturalness	20	These questions make the call to explore shorter distances and higher energies as vibrant and clear as it has ever been.	21
4.3.1. The ‘big’ hierarchy problem: supersymmetry	20	Although the path forward is devoid of guaranteed discoveries, the journey thus far has always been more a matter of serendipity than inevitability. We build colliders not to confirm what we already know, but to explore what we do not.	22
4.3.1.1. Higgsinos, stops, and gluinos	21	In the wake of the Higgs boson’s discovery, the question is	23
4.3.1.2. The gravitino.	21		24
4.3.2. The ‘big’ hierarchy problem: compositeness	22		25
4.3.3. The ‘little’ hierarchy problem	22		26
5. Complementarity	23		27
5.1. EDMs	23		28
5.2. Flavor	24		29

not whether to build another collider, but which collider to build.

Over the course of the last decade, consensus has largely coalesced around linear or circular e^+e^- colliders [3–6] and circular pp colliders [7, 8], both of which constitute natural extensions of past and present machines. The strengths of these two approaches are largely complementary, with the precision of e^+e^- machines and the power of pp machines paving distinct paths toward the exploration of physics at shorter distances. Loosely speaking, the strength of the former is to reveal the fingerprints that new physics has left on the properties of the Higgs and other electroweak states, while the latter are positioned to produce the new physics directly. This has given rise to a familiar dichotomy between energy and precision as contrasting approaches to search for new physics.

Enter the muon. The potential advantages of high-energy muon colliders have long been recognized [9–15]. As a fundamental particle, the muon's full energy is available in a collision, with far cleaner final states relative to those produced by the dissociation of a composite particle like the proton. Its considerable mass suppresses the synchrotron radiation that effectively limits the energies of circular e^+e^- colliders, making both high energies and high luminosities achievable with a relatively small footprint. This raises the prospect that a muon collider could exceed the direct energy reach of the LHC, while achieving unprecedented precision measurements of Standard Model processes. The muon allows us to leverage the benefits of both energy and precision in a unified future collider program.

These advantages come at a cost: the colliding particles are no longer stable. The short lifetime of the muon imposes a series of technical challenges that must be overcome before such a collider can be realized. But progress towards this end has hastened considerably in recent years, spearheaded by the US muon accelerator program (MAP) [16–19], the muon ionization cooling experiment (MICE) [20–22], and the low emittance muon accelerator (LEMMa) concept [23]. Developments on the accelerator side have catalyzed experimental and theoretical activity, reflected by input to the European Particle Physics Strategy Update [24] and the proliferation of studies outlining aspects of the theory case for muon colliders at various energies. Recent contributions include studies of the electroweak boson PDF of the muon [25]; the production of new scalars [26–32] and diverse other states [33] in vector boson fusion (VBF); the discovery potential for minimal dark matter [34, 35]; the measurement of the Higgs self-couplings [36] and couplings to electroweak bosons [37]; the sensitivity to new physics encoded in irrelevant operators [38, 39]; and the coverage of potential BSM explanations for hints from the complementary experiments yielding the muon $g - 2$ [40–44], B meson [45], and K meson [46] anomalies.

In this paper, we present an aspirational theory case underlining the physics potential of a high-energy muon collider. We aim to identify energy and luminosity goals that would position such a collider as a natural successor to the LHC. Our

approach synthesizes some of the qualitative lessons of earlier studies (e.g. [28, 33, 34]), identifies entirely new physics objectives, and explores complementarity with forthcoming experiments across various frontiers. We summarize qualitative features of the most important production modes and characterize the electroweak gauge boson content of the initial state for both standard model and beyond-the-standard model final states; present a range of case studies demonstrating the muon collider's potential to shed light on electroweak symmetry breaking (EWSB), dark matter, and the naturalness of the weak scale; and sharpen connections to complementary experiments probing new physics through electric dipole moments (EDMs), flavor violation, and gravitational waves. Although our primary focus is on high-energy muon colliders, it bears emphasizing that many of the same physics considerations are applicable to other high-energy lepton colliders, and aspects of this work are relevant to the physics case for potential long-term upgrades of the ILC.

In the interest of identifying an optimal collider to succeed the LHC, we consider a variety of center-of-mass (c.m.) energies between 1 and 100 TeV enumerated in table 1, including energy benchmarks associated with various existing proposals. Wherever possible, we remain agnostic about the integrated luminosity attained at a given c.m. energy, preferring instead to determine the amount of integrated luminosity required to discover or constrain a particular point in parameter space. We will also provide forecasts, which for the sake of definiteness will make reference to two luminosity scalings enumerated in table 1. The first, an ‘optimistic’ scaling $\mathcal{L}_{\text{int}}^{\text{opt}}$, assumes integrated luminosity growing with s in order to compensate for the $1/s$ falloff in many interaction cross sections. This is chosen to make contact with recent muon collider physics studies above $\sqrt{s} = 10$ TeV, e.g. [33, 34, 37]. However, it bears emphasizing that the luminosity scaling in these studies reflects physics goals rather than realistic accelerator performance. The second luminosity benchmark, a ‘conservative’ scaling $\mathcal{L}_{\text{int}}^{\text{con}}$, follows the optimistic scaling up to $\sqrt{s} = 10$ TeV, after which it remains flat at 10 ab^{-1} for all subsequent energies. This is conservative only in relation to recent physics studies; we emphasize that the ‘optimistic’ and ‘conservative’ nomenclature is chosen to contrast the second luminosity assumption relative to the first, rather than to reflect definite plausibility. These benchmarks span current accelerator R & D targets up to $\sqrt{s} = 14$ TeV [47]. Above $\sqrt{s} = 14$ TeV, detailed accelerator studies do not currently exist, and attaining even the ‘conservative’ luminosity scaling at energies above $\sqrt{s} = 14$ TeV may prove infeasible. Among other considerations, neutrino flux hazard mitigation becomes increasingly difficult, as does beam focusing; current beam focusing targets at $\sqrt{s} = 10$ TeV are unlikely to be attainable at significantly higher energies. Although the ‘conservative’ luminosity scaling is compatible with an accelerator proposal [48] for a $\sqrt{s} = 100$ TeV muon collider leveraging a gamma factory to produce positrons for muon generation (in the spirit of the LEMMA concept [23]), conventional schemes are likely to yield considerably lower luminosity, and an integrated luminosity on the order of 10 fb^{-1}

Table 1. Energy and integrated luminosity benchmarks considered in this work. As detailed in the text, both the ‘optimistic’ and ‘conservative’ luminosity scalings $\mathcal{L}_{\text{int}}^{\text{opt}}$ and $\mathcal{L}_{\text{int}}^{\text{con}}$ are assumptions in this work for the sake of illustration. Detailed accelerator studies are required above $\sqrt{s} = 14$ TeV to develop robust projections, and are likely to result in more realistic luminosity scalings that fall off considerably by $\sqrt{s} = 100$ TeV.

\sqrt{s} (TeV)	1	3	6	10	14	30	50	100
$\mathcal{L}_{\text{int}}^{\text{opt}}$ (ab $^{-1}$)	0.2	1	4	10	20	90	250	1000
$\mathcal{L}_{\text{int}}^{\text{con}}$ (ab $^{-1}$)	0.2	1	4	10	10	10	10	10

at $\sqrt{s} = 100$ TeV may ultimately prove more realistic. Thus both the ‘optimistic’ and ‘conservative’ luminosity scalings in this paper should be treated as assumptions; detailed accelerator studies above $\sqrt{s} = 14$ TeV are required, and ultimately a realistic luminosity scaling may fall off above $\sqrt{s} = 14$ TeV and approach 10 fb^{-1} at $\sqrt{s} = 100$ TeV.

For the most part, the studies presented here involve rate measurements and an accounting of simple irreducible backgrounds. More detailed projections are necessarily subject to a host of experimental considerations, and future developments in accelerator, detector, and theory studies for a muon collider are closely intertwined. Muons at rest have a relatively short lifetime of $2.2 \mu\text{s}$ and, while the push to high momentum beams can extend the lab frame lifetime up to the order of seconds, the exponential decay of the muon produces an intense source of collinear off-momentum electrons. The electrons then interact with the beamline components, producing electromagnetic showers that result in a high flux of low-energy photons and soft neutrons; these are the primary source of background for a muon collider detector. The process of bending and ultimately focusing the beams to generate a high luminosity collision rate directs these off-momentum backgrounds into collimators upstream and also very close to the interaction point. The exact share of these backgrounds depends strongly on the machine lattice and the interaction point configuration. In all cases, the incorporation of shielding cones close to the interaction point has been identified as significant means of mitigating the effects of beam-induced background (BIB) inside the detector [47].

Detector studies performed in [48] show that the current approach to handle high detector backgrounds appears adequate to preserve the physics capabilities. These studies need to be extended and updated to incorporate recent breakthroughs in technology and higher center of mass energies. For backgrounds incident on the detector elements, the primary tool for separating collider events from BIB is the new generation of precision timing detectors, which leverage the large investment of effort going into HL-LHC upgrades [49, 50]. In the context of the present work, we take this as an encouraging indication that BIBs and associated reconstruction issues can be addressed, but emphasize that all of the studies herein represent estimates in need of detailed experimental study.

As this study will articulate, there is abundant motivation to build a future muon collider. The technological challenges do not appear to be insurmountable, and provide a wealth of

opportunity to develop new experimental techniques. The way the muon collider blurs the line between energy and precision opens the door to novel analysis approaches, while motivating new higher order calculations. And such a machine could readily furnish answers to many of the fundamental questions in particle physics.

The muon-smasher’s guide is organized as follows: in section 2, we sketch many of the main qualitative features of collisions at high-energy muon colliders, with an eye towards their advantages over pp colliders and the interplay between various production modes. In section 3, we turn to the physics of the initial state at high-energy muon colliders, characterizing the electroweak gauge boson content of high-energy muons and developing a pragmatic approach to capturing the most important effects. The broad physics case is developed in section 4, focusing mainly on the central themes (EWSB, dark matter, and naturalness) highlighted by the discovery of the Higgs. In section 5, we explore the complementarity of a muon collider with other experiments operating on compatible timelines, with a particular focus on EDMs, flavor violation, and gravitational waves. We summarize the central lessons of the study in section 6, underlining the energies and luminosities that would position a muon collider to address the questions posed by the Higgs discovery. We reserve a compendium of cross sections and the details of various analyses for appendix A.

2. Muons vs protons

High-energy muon colliders enjoy a host of advantages relative to their proton–proton counterparts, owing in part to the coexistence of scattering processes carrying nearly all of the collider energy (muon annihilation) with those carrying a smaller fraction (VBF). Before studying the physics potential of muon colliders in specific scenarios, we begin with a general exploration of the properties of muon annihilation and VBF, with an eye towards the comparison with proton–proton colliders.

2.1. Muon annihilation

The canonical class of scattering processes, familiar from lower-energy lepton colliders, is $\mu^+\mu^-$ annihilation. Well above the Z pole, the cross section falls off as $1/s$. On the one hand, this implies that the rate for producing both SM and BSM final states in $\mu^+\mu^-$ annihilation falls rapidly with collider energy, although it is worth emphasizing that particles with electroweak quantum numbers still enjoy attobarn-level

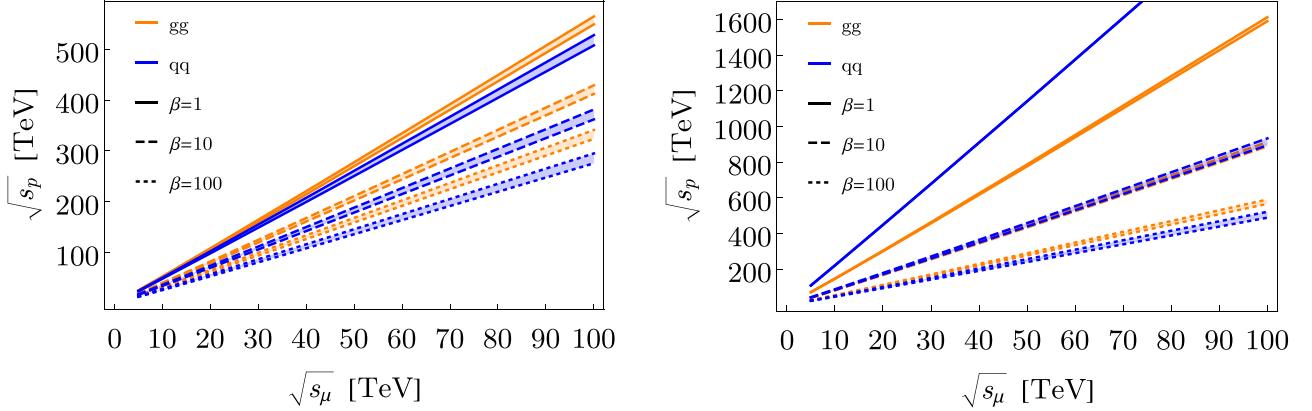


Figure 1. The c.m. energy $\sqrt{s_p}$ in TeV at a proton–proton collider versus $\sqrt{s_\mu}$ in TeV at a muon collider, which yield equivalent cross sections. Curves correspond to production via a gg (orange) or $q\bar{q}$ (blue) initial state at the proton–proton collider, while production at the muon collider is determined by $\mu^+\mu^-$. The partonic cross sections are related by $\beta \equiv [\hat{\sigma}]_p/[\hat{\sigma}]_\mu$. The bands correspond to two different choices of proton PDF sets, NNPDF3.0 LO (as in [33]) and CT18NNLO. The left (right) panel is for $2 \rightarrow 1$ ($2 \rightarrow 2$) scattering. Note that the comparison for $2 \rightarrow 1$ scattering generously assumes the muon collider is operating at the mass of the particle being produced.

cross sections at energies as high as $\sqrt{s} = 100$ TeV. On the other hand, these cross sections are relatively insensitive to the mass of the final state particles unless \sqrt{s} is close to the production threshold. Clearly, muon colliders have considerable discovery potential as long as the final state is sufficiently distinctive.

To quantitatively compare muon and proton colliders, we can compute the c.m. energies at which these two machines have equivalent cross sections [24, 33]. As we emphasize here, one of the great qualitative advantages of a muon collider over a pp collider is that the former generates interactions across all values of the partonic c.m. energy $\sqrt{\hat{s}}$, whereas the latter is dominated by $\sqrt{\hat{s}} \ll \sqrt{s}$ due to compositeness of the proton. In this section, we estimate that it can require between $\mathcal{O}(1-10)$ times more energy at a proton collider to achieve the equivalent production rate at a muon machine, see figure 1. Our focus is on annihilation processes, which are dominated by $x \sim 1$ at a muon collider, where x is the momentum fraction carried by the muon. This is in contrast to VBF-induced processes when the electroweak bosons radiated in the initial state become relevant, which typically have $x \ll 1$; we discuss qualitative features of VBF in this section, and defer a detailed study to section 3. The discussion in this section largely reprises the arguments given in [33].

To make a concrete comparison, we work in terms of generalized parton luminosities. We assume that the inclusive cross section for the final state F (with unspecified remnants X) arising from collisions of (possibly composite) particles A and B takes the form

$$\sigma(AB \rightarrow F + X) = \int_{\tau_0}^1 d\tau \sum_{ij} \frac{d\mathcal{L}_{ij}}{d\tau} \hat{\sigma}(ij \rightarrow F), \quad (1)$$

where hats denote partonic quantities, $\tau = \hat{s}/s$ in terms of the collider c.m. energy \sqrt{s} of the collider and partonic energy $\sqrt{\hat{s}}$, τ_0 is the production threshold, and the parton luminosity is

given by

$$\frac{d\mathcal{L}_{ij}}{d\tau}(\tau, \mu_f) = \frac{1}{1 + \delta_{ij}} \int_\tau^1 \frac{dx}{x} [f_i(x, \mu_f) f_j(\tau/x, \mu_f) + (i \leftrightarrow j)]. \quad (2)$$

Here the $f_i(x, \mu_f)$ are the parton distribution functions (PDFs) for parton i carrying a fraction x of the longitudinal momentum, at factorization scale μ_f , which we take to be $\mu_f = \sqrt{\hat{s}}/2$ when making figure 1.

First, we assume that the process results from a **2 → 1 collision**, i.e., $AB \rightarrow Y$ for a final state Y with mass $M = \sqrt{\hat{s}}$. In this case, the cross section σ_p at a proton–proton collider whose c.m. energy is $\sqrt{s_p}$ takes the form

$$\sigma_p(2 \rightarrow 1) = \int_{\tau_0}^1 d\tau \sum_{ij} \frac{d\mathcal{L}_{ij}}{d\tau} [\hat{\sigma}_{ij}]_p \delta \left(\tau - \frac{M^2}{s_p} \right). \quad (3)$$

At a muon collider whose c.m. energy is $\sqrt{s_\mu}$, the analogous production is dominated at threshold with $s_\mu = \hat{s} = M^2$ (the δ -function from the phase space measure is absorbed by the narrow width); the cross section may be simply approximated by the partonic one, $\sigma_\mu = [\hat{\sigma}]_\mu$. We then solve for relation between s_p and s_μ that yields equivalent cross sections:

$$\sigma_p = \sigma_\mu \implies \frac{[\hat{\sigma}]_p}{[\hat{\sigma}]_\mu} \sum_{ij} \frac{d\mathcal{L}_{ij}}{d\tau} \left(\frac{s_\mu}{s_p}, \frac{\sqrt{s_\mu}}{2} \right) \simeq 1, \quad (4)$$

where we are making the simplifying assumption that the partonic cross section is universal, $[\hat{\sigma}_{ij}]_p \simeq [\hat{\sigma}]_p$. This equation can be solved numerically for s_p in terms of s_μ for different assumptions about the relation between partonic cross sections:

$$\beta \equiv \frac{[\hat{\sigma}]_p}{[\hat{\sigma}]_\mu}. \quad (5)$$

For example, $\beta \simeq 10-100$ may be reasonable for a situation where a state is produced via QCD (electroweak) processes at the proton (muon) collider, keeping in mind not only the larger QCD coupling constant but also the many partonic channels available in pp collisions. We emphasize that

this comparison is biased in favor of muon colliders, as it assumes that the muon collider is fortuitously operating at the mass of the particle in question. In practice, the discovery of a hitherto-unknown state via resonant production at a muon collider is more likely to occur via VBF production, for which cross sections at $\hat{s} \sim s_\mu$ are much lower. In this situation, the strengths of a muon collider in $2 \rightarrow 1$ annihilation production would more likely come into play after discovery, where operation at $s_\mu \simeq M^2$ would allow for precision studies.

For $2 \rightarrow 2$ collisions, we assume the muon collider is optimized so that $\sqrt{s_\mu}$ is slightly above threshold, while at a proton–proton collider we take $[\hat{\sigma}]_p \propto 1/\hat{s}$, which is appropriate far above threshold. Then for the proton–proton case we can write

$$\sigma_p(2 \rightarrow 2) = \frac{1}{s_p} \int_{\tau_0}^1 d\tau \frac{1}{\tau} \sum_{ij} \frac{d\mathcal{L}_{ij}}{d\tau} [\hat{\sigma}_{ij} \hat{s}]_p, \quad (6)$$

while for the muon collider we have $\sigma_\mu = [\hat{\sigma} \hat{s}]_\mu / s_\mu$. In this case,

$$\sigma_p = \sigma_\mu \implies \frac{s_\mu}{s_p} \frac{[\hat{\sigma} \hat{s}]_p}{[\hat{\sigma} \hat{s}]_\mu} \int_{s_\mu/s_p}^1 d\tau \frac{1}{\tau} \frac{d\mathcal{L}_{ij}}{d\tau} \left(\tau, \frac{\sqrt{s_\mu}}{2} \right) \simeq 1, \quad (7)$$

which can again be solved numerically for s_p given various assumptions about the ratio of partonic cross sections.

The results of this exercise are shown in figure 1. For both $2 \rightarrow 1$ and $2 \rightarrow 2$ processes, the equivalent energy relationship is essentially linear, with only modest dependence on the choice of proton PDFs. A muon collider enjoys considerable advantages in $2 \rightarrow 1$ production, reaching the cross section of a 100 TeV pp collider by $\sqrt{s_\mu} \sim 20$ TeV assuming comparable partonic cross sections. The advantage is even sharper in $2 \rightarrow 2$ production, where a muon collider reaches the cross section of a 100 TeV pp collider between $\sqrt{s_\mu} \sim 5 - 7$ TeV for comparable partonic cross sections, depending on whether the process is $q\bar{q}$ - or gg -initiated. Even allowing for an enhancement of $\beta = 10$ at a pp collider (accounting for the difference between QCD and electroweak partonic cross sections), the $2 \rightarrow 2$ cross section of a 100 TeV pp collider is emulated by a muon collider operating at $\sqrt{s_\mu} \sim 12$ TeV.

As mentioned above, this comparison is necessarily favorable to muon colliders in the sense that it assumes the mass scale of new physics lies just below the c.m. energy of the collider, although radiative return can salvage some of the same conclusions at the cost of further suppressing the partonic cross section. Developing a comprehensive case for muon colliders still requires investigating production cross sections for a variety of new physics scenarios across a range of c.m. energies that go well above the production threshold. That is the motivation for the model dependent studies provided in sections 4 and 5 below.

2.2. Vector boson fusion

For collisions well above the production threshold, the virtual electroweak gauge boson content of high-energy muon beams becomes increasingly relevant, akin to the virtual gluon

content of high-energy proton beams. VBF becomes a correspondingly important channel for the production of SM and BSM particles alike, with cross sections typically scaling with c.m. energy as $\sim \log(s)$ far above threshold. The practical aspects of VBF at high-energy muon colliders have recently been studied in detail in [33], while the treatment of VBF in terms of electroweak PDFs has been initiated in [25]. Here we summarize the qualitative features of VBF production, reserving a detailed discussion of electroweak PDFs for section 3.

For the sake of illustration, representative electroweak PDFs and parton luminosities are shown in figure 2. The PDFs $f_i(x, Q)$ are shown as a function of the partonic momentum fraction x at the scale $Q = 10$ TeV for transverse polarizations of the photon and both transverse and longitudinal polarizations of the W and Z , while the parton luminosity functions $d\mathcal{L}_{ij}/d\tau(\tau, \mu_f)$ are shown as a function of $\sqrt{\tau} = \sqrt{\hat{s}/s}$ with factorization scale $\mu_f = \sqrt{\tau s}/2$ for $\sqrt{s} = 10$ TeV. Details of their derivation and scale dependence are presented in section 3.

These distributions illuminate many of the salient features of VBF production, modulo additional dependence on the partonic cross section for the process of interest. All of the electroweak boson PDFs peak at $x \sim 0$, a manifestation of the inherent infrared singularity. Photons constitute by far the largest component away from $x \sim 1$, as their logarithmic enhancement due to soft emission extends all the way down to m_μ , compared to the W and Z who are only non-zero for energies beginning at $\sim m_{W,Z}$. The relative size of Z PDFs relative to W PDFs reflects the familiar suppression of the Z coupling to muons. Longitudinal polarizations of the W and Z are suppressed relative to transverse polarizations due to the former’s modest coupling to muons, which is set by the muon mass.

Although these distributions provide a good qualitative sense of the various contributions to VBF processes, both the convolution of the parton luminosities with the partonic cross section and the imposition of realistic phase space cuts significantly affect the properties of VBF production cross sections. The sizable photon PDF reflects the abundance of soft photons at low scales, which do not contribute significantly to the production of particles with sizable transverse momentum. Even moderate phase space requirements on the final state, such as transverse momentum cuts, reduce the relative logarithmic enhancement enjoyed by the photon. Features of the partonic cross section for a given process can have a significant impact, most notably longitudinal enhancement. As a result, the dominant contribution to VBF production of particles carrying $SU(2)_L$ quantum numbers is often WW fusion, rather than $\gamma\gamma$ fusion, a conclusion borne out in the numerical results of both [33] and this study. For partonic processes with enhanced contributions from longitudinal polarizations, such as $W^+W^- \rightarrow t\bar{t}$ or $W^+W^- \rightarrow h$, this enhancement is often sufficient to overcome the relative suppression of the PDFs; this will be illustrated in a number of examples in section 4.1. As such, high-energy muon colliders are as much longitudinal gauge boson colliders as

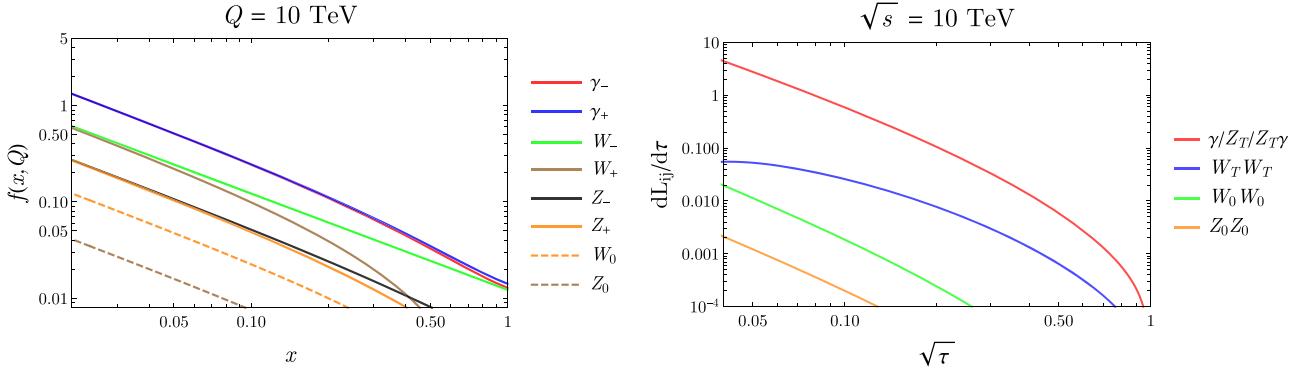


Figure 2. Left: the polarized PDF $f_i(x, Q)$ for electroweak gauge bosons separated by transverse and longitudinal helicities $\pm, 0$ as a function of partonic momentum fraction x at a scale $Q = 10$ TeV. Right: the parton luminosity functions $dL_{ij}/d\tau(\tau, \mu_f)$ as a function of $\sqrt{\tau} = \sqrt{\hat{s}}/s$ for $Q = \sqrt{s}/2$ at $\sqrt{s} = 10$ TeV.

they are transverse gauge boson colliders, subdominant PDFs notwithstanding.

2.3. Annihilation vs VBF

Ultimately, there is a rich interplay between annihilation and VBF production of both SM and BSM particles at high-energy muon colliders. For standard model processes well above threshold, the relative scaling as a function of collider energy \sqrt{s} is [33]

$$\frac{\sigma_{\text{VBF}}^{\text{SM}}}{\sigma_{\text{ann}}^{\text{SM}}} \propto \alpha_W^2 \frac{s}{m_V^2} \log^3 \frac{s}{m_V^2}, \quad (8)$$

where the triple logarithmic enhancement is due to a double collinear logarithm from the two electroweak PDFs and a single soft logarithm. The competition between the coupling suppression and energy growth leads to crossovers between standard model cross sections for annihilation and VBF production around energies of $\sqrt{s} \sim$ few TeV, with correspondingly higher crossover energies for higher-multiplicity final states. The scaling is analogous for production of BSM particles with a final-state mass scale m_X , for which the relative scaling well above threshold is [33]

$$\frac{\sigma_{\text{VBF}}^{\text{BSM}}}{\sigma_{\text{ann}}^{\text{BSM}}} \propto \alpha_W^2 \frac{s}{m_X^2} \log^2 \frac{s}{m_V^2} \log \frac{s}{m_X^2}, \quad (9)$$

where m_V is the mass scale of an intermediate state in the production process (often an electroweak vector boson, which in any case is assumed to satisfy $m_V \ll \sqrt{s}$). The collision energy \sqrt{s} at which annihilation and VBF cross sections for BSM final states cross over grows with the mass scale of the final state, but ultimately there are always collision energies at which VBF production wins for a fixed mass scale. This lends credence to the notion of high-energy muon colliders as gauge boson colliders, and highlights the importance of analyzing VBF production modes in characterizing the physics reach of these colliders.

Even so, it bears emphasizing that the above scaling assumes muon collisions are occurring well above production threshold. This is likely to be true for most standard model processes at a high-energy muon collider, but need not hold for BSM production. Although the underlying approximations

break down as $m_X \rightarrow \sqrt{s}$, the naive scaling in equation (9) indicates that annihilation production once again dominates in this limit. Despite the overall $1/s$ falloff in annihilation cross sections, even a handful of events with sufficiently distinctive final states near threshold may be sufficient for the discovery of new physics. We caution that projections based on VBF production modes alone would fail to capture these important cases.

The interplay between annihilation and VBF production is illustrated in figure 3, which shows the ratio of the annihilation cross section σ_{ann} and VBF cross section σ_{VBF} for two representative examples—a vector-like pair of $SU(2)$ singlet Weyl fermions with hypercharge ± 1 , and a vector-like pair of $SU(2)$ doublet Weyl fermions with hypercharge $\pm 1/2$ —as a function of the collider energy \sqrt{s} and the particle mass relative to threshold. In both cases, the annihilation cross section is computed analytically while the VBF cross section is computed by convolving partonic cross sections with the corresponding PDFs derived in section 3. For both the $SU(2)$ singlet and doublet, the crossover takes place once the fermion mass is above about 10% of $\sqrt{s}/2$. Ultimately the differences in the two cases are modest; although the WW fusion contribution to VBF is much larger for the doublet, the dominant contribution in both cases is ultimately from $\gamma\gamma$ fusion. For sufficiently distinctive final states, this is likely to favor production via annihilation as a discovery mode. Of course, the details depend on the relative sizes of signal and background, another aspect where muon colliders enjoy further advantages over their proton–proton counterparts, as we will now emphasize.

2.4. Signal vs background

A final generic advantage of a muon collider over a pp collider that we want to highlight has to do with the comparison of signal and background rates. Furthermore, this benefit is not restricted to production cross sections for high-mass states. This is not obvious at first glance; for low-mass states, including standard model particles, the rate advantage of pp colliders is considerable. For example, at $\sqrt{s} = 14$ TeV, the leading single Higgs production cross section is a

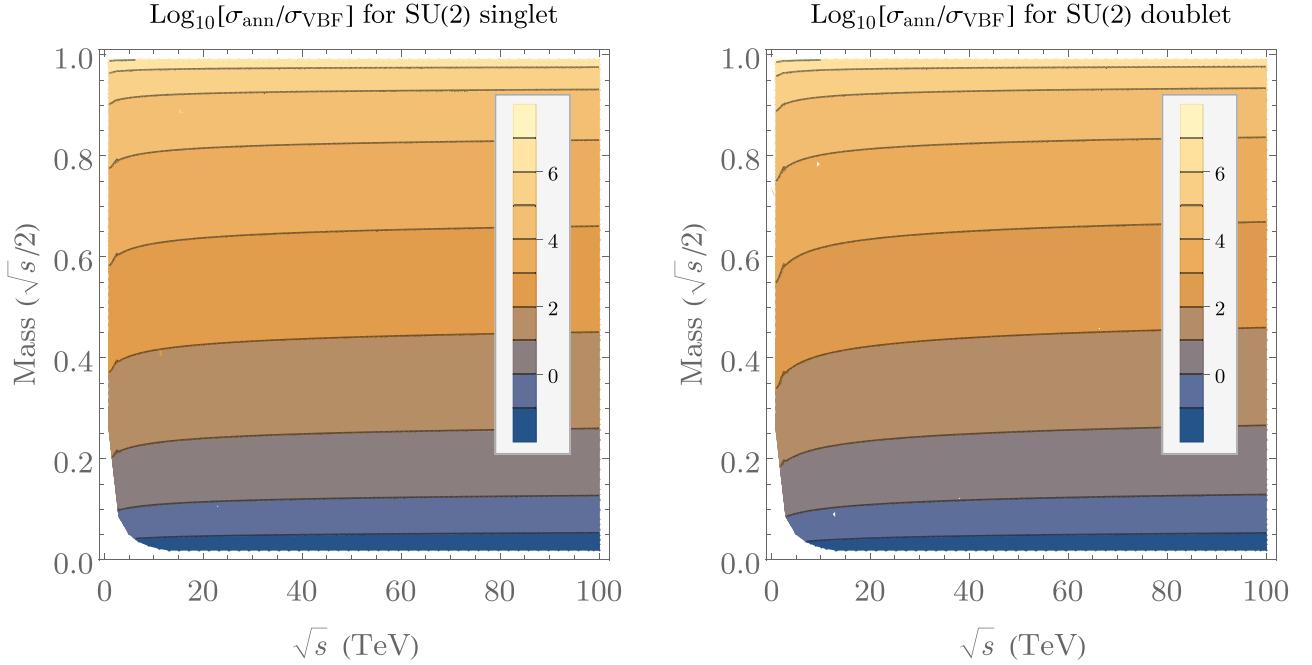


Figure 3. Left: log ratio of the annihilation cross section σ_{ann} and VBF cross section σ_{VBF} for a vector-like pair of $SU(2)$ singlet Weyl fermions with hypercharge ± 1 as a function of the collider energy \sqrt{s} and the particle mass relative to threshold. Right: the same ratio for a vector-like pair of $SU(2)$ doublet Weyl fermions with hypercharge $\pm 1/2$.

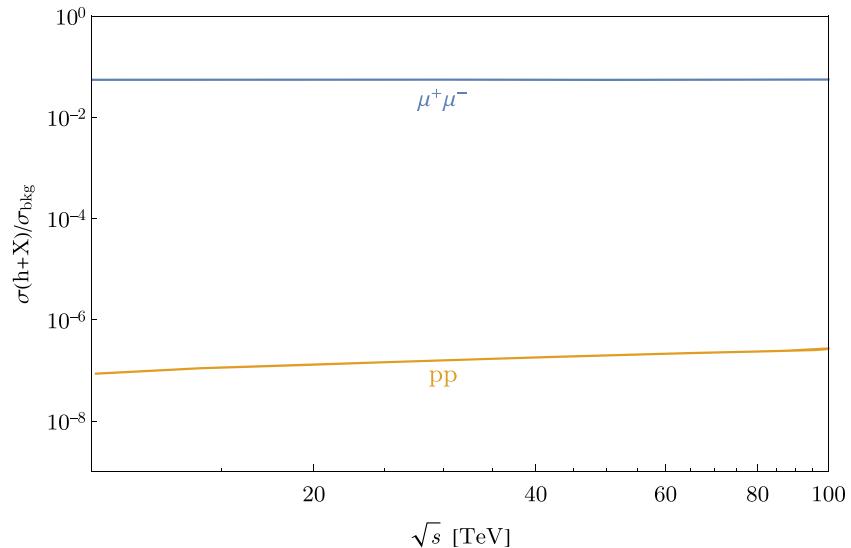


Figure 4. Higgs production cross section $\sigma(h+X)$ as a fraction of a representative ‘background’ cross section σ_{bkg} for $\mu^+\mu^-$ and pp colliders. For $\mu^+\mu^-$ colliders, we compute Higgs production using the LO cross section for $\mu^+\mu^- \rightarrow h + \nu\bar{\nu}$, while the ‘background’ cross section σ_{bkg} is taken to be the rate for single electroweak boson production, which is dominated by VBF production of W, Z, h, γ at these energies. For pp colliders we take the Higgs production cross section to be the N3LO cross section for $gg \rightarrow h$ [51] presented in [52], while the ‘background’ cross section σ_{bkg} is taken to be the $pp \rightarrow b\bar{b}$ cross section computed by MCFM [53].

factor of ~ 50 larger at a pp machine than its $\mu^+\mu^-$ counterpart. But ultimately, our ability to extract physics from the collider data sensitivity depends on the background rates, and here the advantage is decisively in favor of muon colliders.

As illustrated in figure 4, the ratio between the single Higgs production cross section and a representative ‘background’ cross section at pp and $\mu^+\mu^-$ colliders operating at the same

\sqrt{s} differs by roughly six orders of magnitude and is relatively insensitive to \sqrt{s} . Here we have taken the representative ‘background’ cross section at a pp collider to be the $pp \rightarrow b\bar{b}$ cross section, a relevant measure of standard model backgrounds; the total or inelastic pp cross sections are orders of magnitude larger. The corresponding ‘background’ cross section at a $\mu^+\mu^-$ collider is taken to be the sum of inclusive single electroweak boson cross sections, representative

of the overall rate for hard standard model processes; the cross section for $b\bar{b}$ production with basic acceptance cuts is within the same order of magnitude [54]. Assuming it is possible to achieve comparable integrated luminosities at both experiments, it is clear that a precision Higgs program at a future muon collider provides many opportunities to study the Higgs sector in exquisite detail. Of course, these statements ultimately rest on details of detector performance and the ability to mitigate the beam induced muon collider backgrounds. Nonetheless, we see this as a clear sign that the muon collider will be an effective tool to probe both the precision and energy frontiers.

3. Muon colliders are gauge boson colliders

Before getting into the detailed physics case, this section will describe the physics of the initial state at a high energy muon collider. Naively, the advantage of a lepton collider is that the colliding beams are composed of elementary particles (so that the collision is relatively clean), which are in momentum eigenstates (so that the c.m. energy for each collision is known). This can be contrasted against proton colliders, where the beams are composed of composite states, so that the partonic c.m. energy varies from collision to collision. To make predictions in this case, one convolves the hard process of interest with universal PDFs. Additionally, the smashed protons leave a trail of debris in their wake, the so-called underlying event. As we will argue in this section, making predictions for a muon collider whose beams carry TeVs of energy has aspects in common with both better known types of machines: one must use PDFs, but the collision yields a small number of particles in the initial state that can be modeled reasonably well using perturbation theory.²⁰

At the theoretical level, the situation for a muon collider is simplified with respect to a proton collider since perturbative control can be maintained at every step of the calculation.²¹ For example, the boundary conditions for the proton PDFs are set at a scale where QCD is non-perturbative, implying that one must rely on inputs from experiment to numerically determine the proton PDFs. All of the complications that stem from this fact are avoided when studying muon PDFs. The muon colliders we discuss here have energy in the TeV to 100 TeV range, and so the masses of the weak gauge bosons can be treated as a small perturbation, i.e., it is typically reasonable to treat them as massless so that the PDF formalism applies; see section 3.4 for a brief discussion of finite mass effects. And since the electroweak gauge couplings are relatively small, working with leading order unresummed PDFs provides a reasonable approximation to the resummed result;

²⁰ Throughout this paper, we treat the muon beams as stable. Everything we say here is independent of this assumption, as long as our amazing accelerator colleagues can figure out how to provide us with a robust muon beam to play with.

²¹ Of course, some of these techniques are also relevant for past and proposed electron-position experiments, e.g., when predicting VBF initiated processes. However, the small mass of the electron effectively bounds the maximum energy for circular machines to be near the electroweak scale.

we will demonstrate the minimal impact of next-to-leading-log corrections in section 3.3 below. Interesting complications arise due to EWSB, but other than treating the mass versus gauge eigenbasis for the electroweak bosons consistently as we do below, these tend to have a small numerical effect on the cross section predictions. There are additionally subtleties associated with capturing the physics of the longitudinal gauge boson modes, and the interplay with the Goldstone equivalence theorem and unitarity; we will not comment on this further and will simply use the splitting functions in the ‘Goldstone equivalence gauge’ computed in [55]. Finally, while it is beyond the scope of this work, we note that one can also include the effects of QCD into the muon PDFs, as was recently described in [54].

In the rest of this section, we will first write down the formalism used to solve for the PDFs to leading logarithmic order using leading order splitting functions. This will provide us with a framework to explore the accuracy that can be achieved when taking different approximations. Our goal will be to demonstrate that the truncated leading log order (TLLO) PDFs provide a reasonable approximation to the more complete all log order results that result from integrating the DGLAP evolution equations. Given that the TLLO PDFs are easy to understand and can be expressed analytically, we advocate that these are all that are required to make predictions for a future muon collider in the energy range of interest here, unless high precision calculations are needed.

3.1. From the effective vector approximation to PDFs

The soft and collinear divergences inherent to theories of charged particles coupled to gauge bosons yield physical logarithmic enhancements that can spoil the convergence of perturbation theory. When considering colliding beams of charged particles, it is important to acknowledge our inherent inability to experimentally distinguish a single state in isolation from one that has emitted a nearly collinear or very soft additional particle. A framework for addressing this problem was first written down in 1934 by Weizsaecker [56] and Williams [57]; this is what is known as the ‘effective photon approximation’ or more generally the ‘effective vector approximation’ (EVA):

$$f_\gamma^{\text{EVA}}(x) \simeq \frac{\alpha}{2\pi} P_\gamma(x) \log \frac{E^2}{m_\mu^2}, \quad (10)$$

where the \log is the result of soft emissions, and the QED splitting function

$$P_\gamma(x) = \frac{1 + (1 - x)^2}{x} \quad (11)$$

can be derived by taking the collinear limit of a tree-level $1 \rightarrow 2$ process computed using perturbative QED. Note that in this approximation, E is the beam energy of the colliding charged particle while the emitted photon energy is given by $Q = xE$, where x is the momentum fraction carried by the photon.

Systematically improving this approximation requires developing the relevant DGLAP evolution equations, which

allow one to resum the \log that appears in equation (10), yielding the PDFs. It is perhaps under appreciated that the EVA and PDF approaches already differ at TLLO: the EVA is proportional to $\log E^2$, while the TLLO PDF is proportional to $\log Q^2$, where Q^2 is an unphysical renormalization scale, whose canonical value $Q^2 \sim (xE)^2$ is typically chosen to minimize higher order logarithms. With this choice of scale, the TLLO PDFs are

$$f_\gamma^{\text{PDF, TLLO}}(x) \simeq \frac{\alpha}{2\pi} P_\gamma(x) \log \frac{(xE)^2}{m_\mu^2}, \quad (12)$$

where the splitting function $P_\gamma(x)$ is still given by equation (11), and we are taking the gauge coupling to be fixed for simplicity. This is of course the logarithmic behavior one would find when computing at fixed order in perturbation theory, and additionally it follows from solving the DGLAP evolution equations for QED given in equation (15) to leading log order, as it must for self consistency. As we will see in what follows, simply solving the truncated DGLAP equation at the leading log order, as given in equation (12) for the full system of partons relevant at a muon collider provides a good approximation to the full solution to the DGLAP equations. Note that since it is trivial to implement, we do allow the gauge coupling to run (at one-loop order) when computing the leading log PDFs that are used to make some of the cross section predictions below.

3.2. PDFs with broken electroweak symmetry

Naively, one might expect that the PDFs for the massive electroweak gauge bosons can be derived by simply using the appropriate splitting functions, and replacing $m_\mu \rightarrow m_V$ inside the logarithm. This is the case for the W^\pm bosons, where the PDFs are given by²²

$$f_{W_T^-}^{\text{EVA}}(x) \simeq \frac{\alpha_2}{4\pi} P_{W_T^- \leftarrow \mu_L}(x) \log \frac{E^2}{m_W^2}, \quad (13)$$

with splitting functions

$$P_{W_T^- \leftarrow \mu_L}(x) = \frac{1 + (1 - x)^2}{x}, \quad (14)$$

which captures the splitting to W^- summed over both polarizations, assuming that the incoming muon beam has equal left-handed and right-handed helicity. Note that one must be careful to keep track of the helicity dependence, since the W^\pm couplings are chiral.

The computation of the Z -boson PDF is complicated by the fact that it mixes with the photon. Hence, one must first evolve the photon PDF from the scale $Q^2 = m_\mu^2 \rightarrow m_Z^2$, where the evolution equations change. Noting that the electromagnetic interactions conserve both C and P, we do not need to track the difference in helicities for this step of

²² We do note that there are some subtle questions about how to treat the longitudinal components, such that the Goldstone equivalence theorem is respected. However, this issue only appears at subleading order so we do not have to treat it carefully here; see [55] for a discussion.

the calculation. The DGLAP equations for the photon and muon are

$$\frac{d}{d \log Q^2} \begin{pmatrix} f_\gamma(x, Q^2) \\ f_\mu(x, Q^2) \end{pmatrix} = \begin{pmatrix} \mathcal{P}_{\gamma \leftarrow \gamma}(x) & \mathcal{P}_{\gamma \leftarrow \mu}(x) \\ \mathcal{P}_{\mu \leftarrow \gamma}(x) & \mathcal{P}_{\mu \leftarrow \mu}(x) \end{pmatrix} \otimes \begin{pmatrix} f_\gamma(x, Q^2) \\ f_\mu(x, Q^2) \end{pmatrix}, \quad (15)$$

where the convolution is defined in the standard way:

$$f(x) \otimes g(x) = \int_x^1 \frac{dz}{z} f(z) g\left(\frac{x}{z}\right). \quad (16)$$

Here we are treating the muon and electroweak gauge bosons as partons, but we ignore the contribution from the anti-muon since it does not appear at the order of interest. The explicit splitting functions are [55]

$$\mathcal{P}_{\gamma \leftarrow \mu}(x) = \frac{e^2}{8\pi^2} \frac{1 + (1 - x)^2}{x} \quad (17a)$$

$$\mathcal{P}_{\gamma \leftarrow \gamma}(x) = -\frac{2}{3} N_\gamma \delta(1 - x) \quad (17b)$$

$$\mathcal{P}_{\mu \leftarrow \mu}(x) = \frac{e^2}{8\pi^2} \left(\frac{1 + x^2}{(1 - x)_+} + \frac{3}{2} \delta(1 - x) \right), \quad (17c)$$

where N_γ counts the number of ways the photon can annihilate to quark and lepton pairs, and the plus function is defined in the standard way:

$$\int_x^1 dz \frac{f(z)}{(1 - z)_+} = \int_x^1 dz \frac{f(z) - f(1)}{(1 - z)}. \quad (18)$$

The boundary conditions are

$$f_\gamma(x, m_\mu^2) = 0 \quad \text{and} \quad f_\mu(x, m_\mu^2) = \delta(1 - x). \quad (19)$$

We solve these equations for f_γ and f_μ ; evaluating them at $Q^2 = m_Z^2$ provides the boundary conditions for the DGLAP evolution equations that include the Z boson. At this step, we must keep track of the different helicity dependence, so we assign half of the muon PDF to each helicity and then let them evolve independently. It is also critical to account for the difference between the gauge and mass eigenbases, since the interactions are diagonal in the former while physical processes are computed using the later. To this end, we need convert from $(\gamma, Z, Z\gamma)$ to (B, W_3, BW_3) using the transformation matrix

$$\begin{pmatrix} B \\ W_3 \\ BW_3 \end{pmatrix} = \begin{pmatrix} \cos^2 \theta_W & \sin^2 \theta_W & -\cos \theta_W \sin \theta_W \\ \sin^2 \theta_W & \cos^2 \theta_W & \cos \theta_W \sin \theta_W \\ 2 \cos \theta_W \sin \theta_W & -2 \cos \theta_W \sin \theta_W & \cos^2 \theta_W - \sin^2 \theta_W \end{pmatrix} \begin{pmatrix} \gamma \\ Z \\ Z\gamma \end{pmatrix}, \quad (20)$$

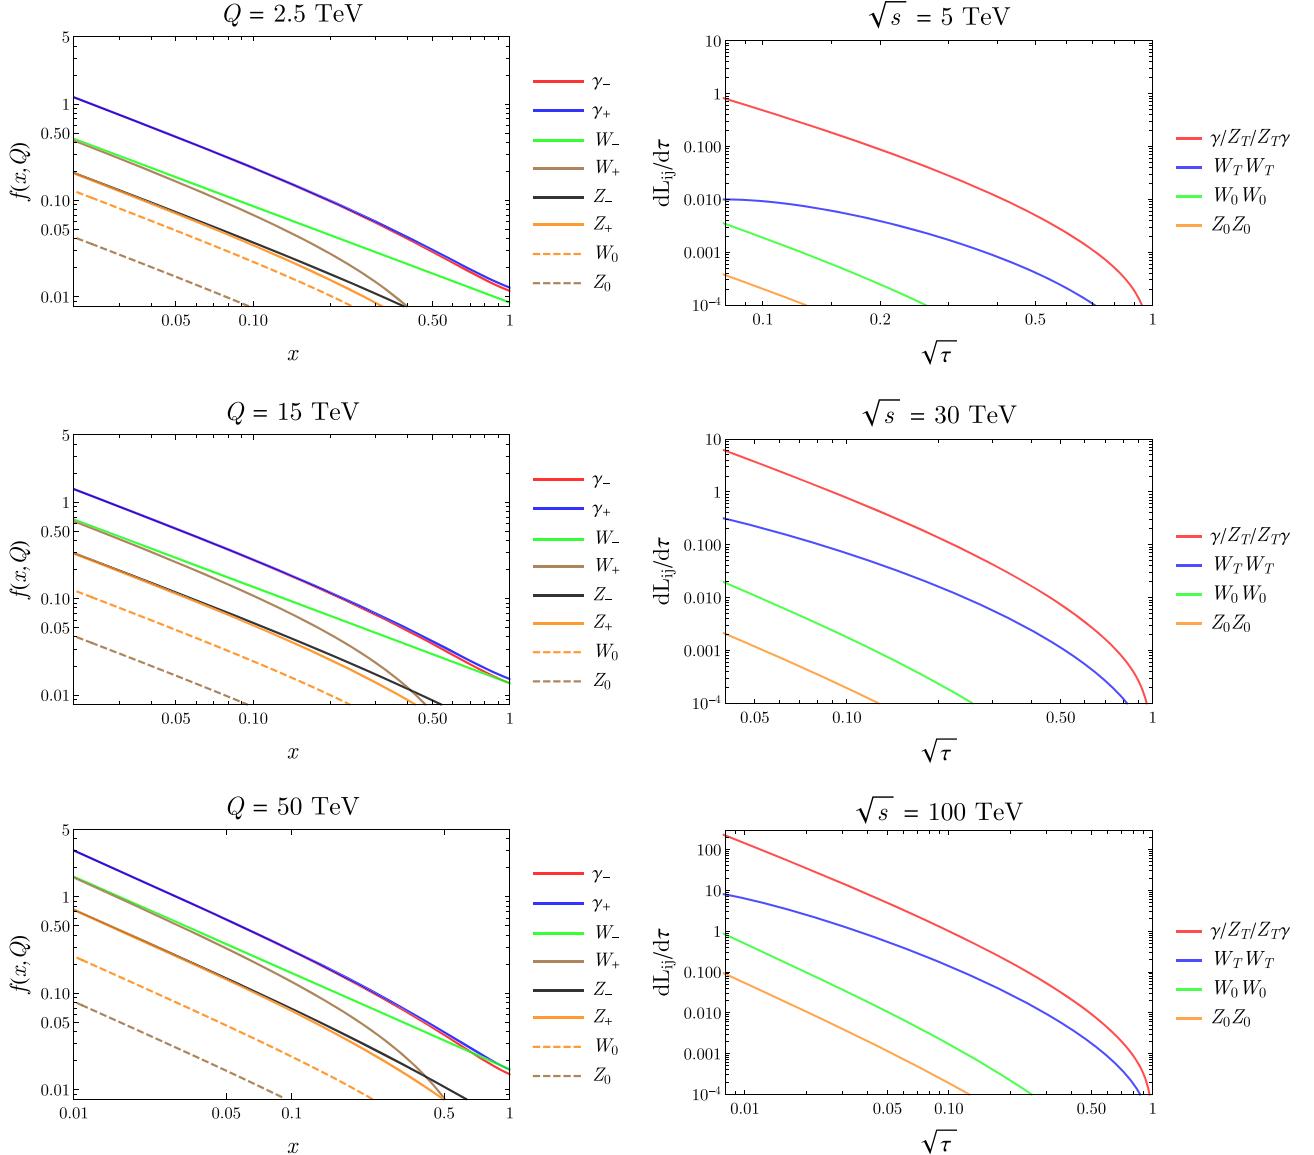


Figure 5. Left: the polarized PDF $f_i(x, Q)$ for the electroweak gauge bosons separated by helicity. Right: the parton luminosity functions $dL_{ij}/d\tau(\tau, Q = \sqrt{\tau s}/2)$ for gauge bosons separated by helicity.

where θ_W is the weak mixing angle, and the mixed $Z\gamma$ PDF accounts for possible interference effects among diagrams involving a $Z + \gamma$ initial state. For reference, the parton luminosity for Z and γ initial states is derived using the combination

$$\frac{dL}{d\tau} = f_Z(x, Q) \otimes f_\gamma(x, Q) + f_{Z\gamma}(x, Q) \otimes f_{Z\gamma}(x, Q). \quad (21)$$

The full DGLAP evolution equations are given in [55]. To provide an example, the DGLAP evolution equation for f_{B_-} is

$$\begin{aligned} \frac{d}{d \log Q^2} f_{B_-}(x, Q^2) &= \mathcal{P}_{B_- \leftarrow \mu_L}(x) \otimes f_{\mu_L}(x, Q^2) \\ &+ \mathcal{P}_{B_- \leftarrow \mu_R}(x) \otimes f_{\mu_R}(x, Q^2) \quad (22) \\ &+ \mathcal{P}_{B_- \leftarrow \nu_L}(x) \otimes f_{\nu_L}(x, Q^2) \\ &+ \mathcal{P}_{B_- \leftarrow B_-}(x) \otimes f_{B_-}(x, Q^2), \end{aligned}$$

where the splitting functions are given by²³

$$\mathcal{P}_{B_- \leftarrow \mu_L}(x) = \mathcal{P}_{B_- \leftarrow \nu_L}(x) = \frac{1}{8\pi^2} \left(\frac{-g_1}{2} \right)^2 \frac{1}{x} \quad (23a)$$

$$\mathcal{P}_{B_- \leftarrow \mu_R}(x) = \frac{1}{8\pi^2} (-g_1)^2 \frac{(1-x)^2}{x} \quad (23b)$$

$$\mathcal{P}_{B_- \leftarrow B_-}(x) = -\frac{g_1^2}{8\pi^2} \frac{415}{144} \delta(1-x). \quad (23c)$$

Although we do not write down the DGLAP evolution equation for f_{B_+} explicitly, there is a feature of the splitting

²³ Note that the splitting function $\mathcal{P}_{B_- \leftarrow B_-}(x)$ is determined by the decay rate of B_- . This expression includes the decay channel to a pair of massless top quarks, which is no longer a good approximation for small x . We have checked that this contribution to the B PDF is small, so for our purposes here we will simply use equation (23).

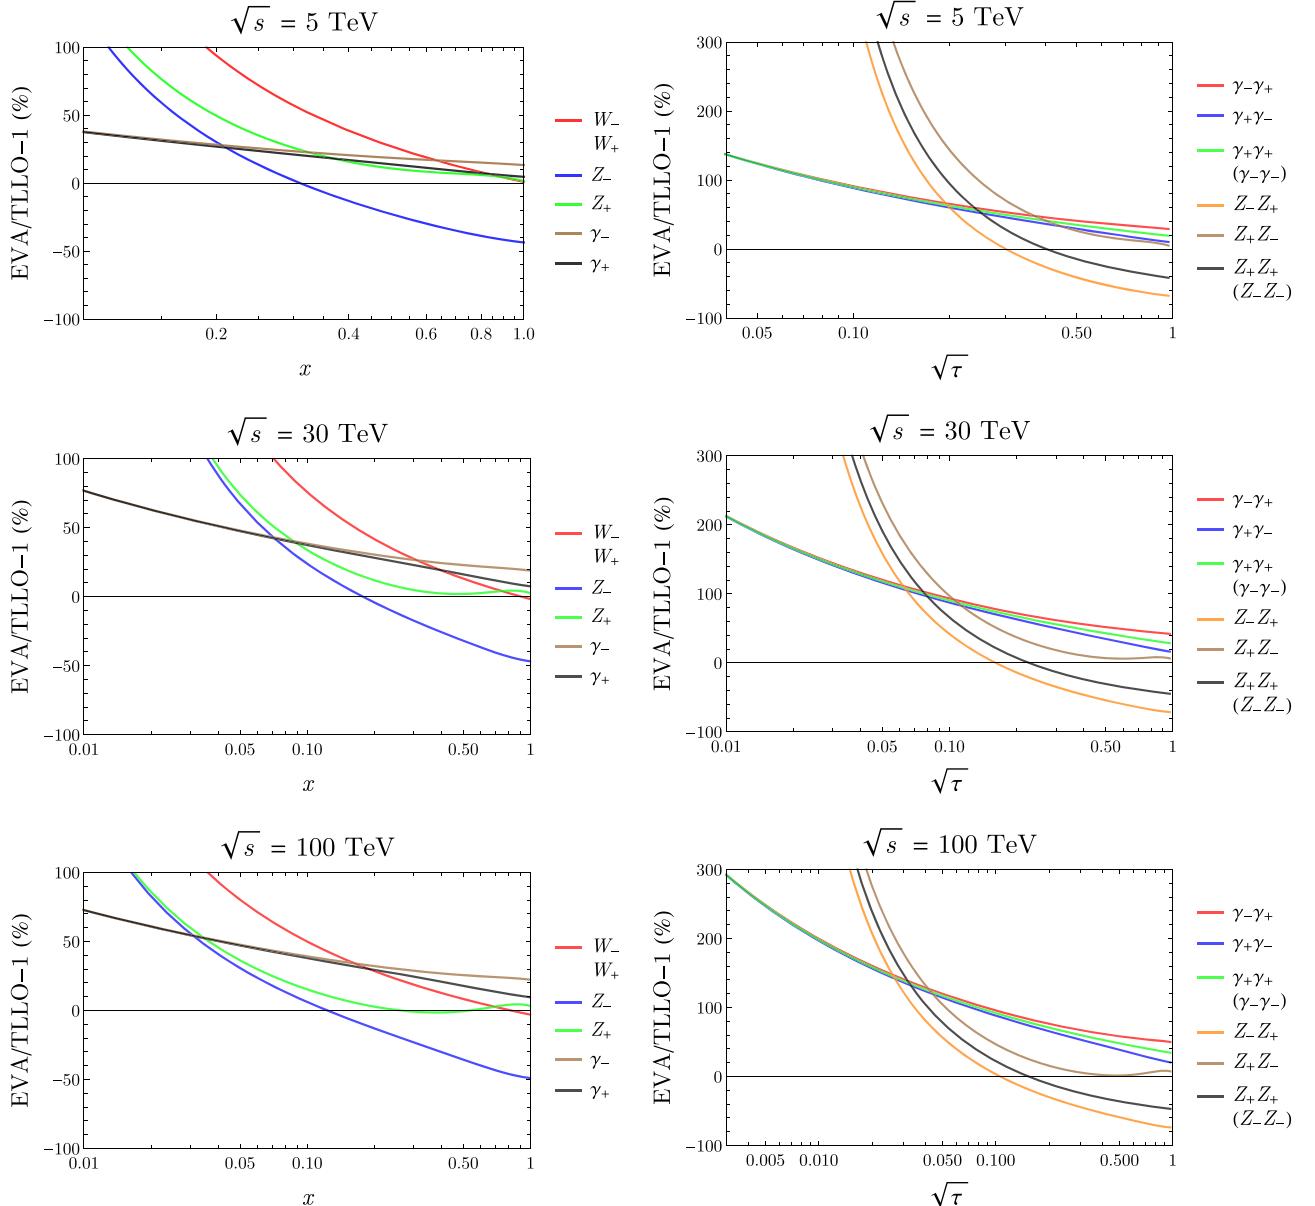


Figure 6. The percent difference between the EVA and TLLO PDFs for the gauge bosons separated by helicity. Left: a comparison of the PDFs taking $Q = x\sqrt{s}/2$ when evaluating the TLLO PDF. Right: a comparison of the parton luminosities taking $Q = \sqrt{\tau s}/2$.

functions that is worth noting. Due to CP invariance, the splitting function $\mathcal{P}_{B_{+} \leftarrow \mu_L} = \mathcal{P}_{B_{-} \leftarrow \mu_R}$, and so we use equation (23b) for both. Comparing equation (23a) with equation (23b), we see that the splitting functions for left-handed muon decaying to left-handed and right-handed gauge boson become equal in the $x \rightarrow 0$ limit. For contrast, at high x the latter approaches zero. This is due to the conservation of angular momentum. At high x , the probability to reverse the helicity goes to zero, which is one key reasons it is important to use the polarized PDFs for electroweak interactions.

Now that we have set up the detailed formalism, we simply solve the DGLAP evolution equations truncated to leading log order (which does not require performing any convolutions; see section 3.3 for more details), including one-loop running gauge couplings. The resulting approximated gauge

boson PDFs in the helicity basis are shown in figure 5. For these plots, we align the incoming muon (anti-muon) beam with the positive (negative) z -axis, and the positive z -axis with positive helicity. As we anticipated from the splitting functions, the right-handed gauge boson PDFs approach zero faster than the left-handed PDFs as $x \rightarrow 1$. We also note an interesting helicity-dependent effect for the photon PDF. If we had neglected the impact of rotating between the mass and gauge bases using equation (20), then clearly the photon PDF should not show any helicity dependence; see equation (15). However, after converting to the gauge basis, the neutral gauge bosons couple to left- and right-handed muons differently. This is the origin of the helicity dependence at large x for the photon PDF. Finally, we note that the PDF for the longitudinal gauge bosons is scale invariant, up to the minor

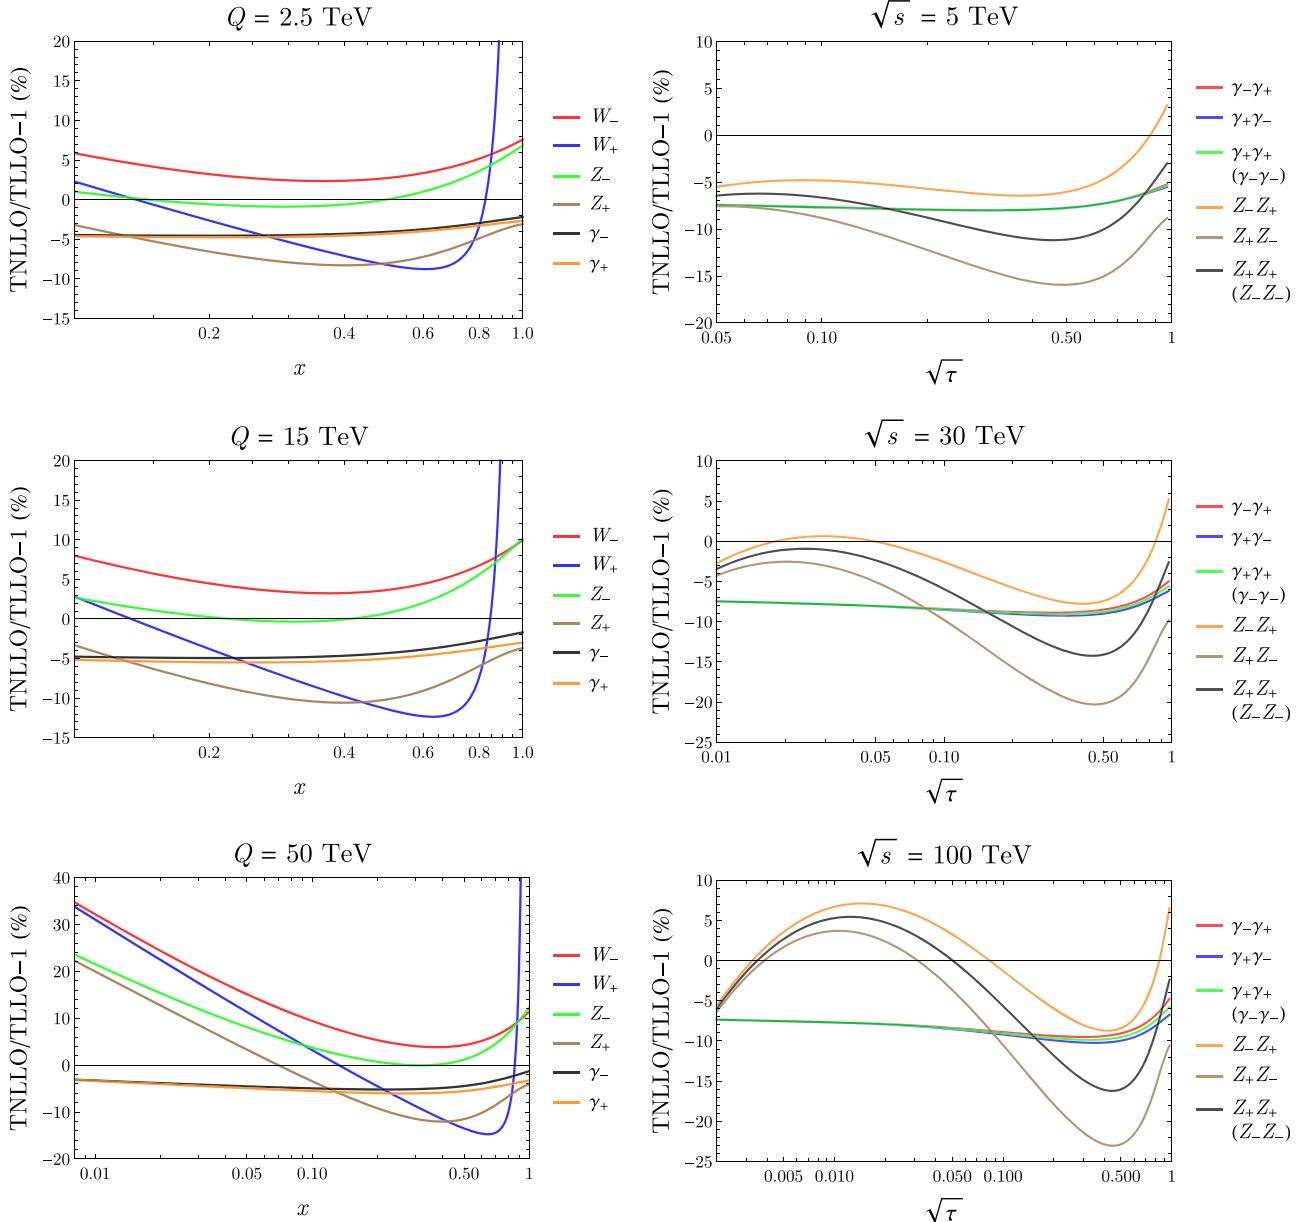


Figure 7. The percent difference between the TNLO to TLLO polarized PDF for different electroweak gauge bosons; see the legend. Left: a comparison of the PDFs for fixed Q . Right: a comparison of the parton luminosities taking $Q = \sqrt{s}/2$.

scale dependence from the running coupling. Naively, this is simply due to the fact that the polarization vector of the gauge boson is proportional to the momentum, and this implies that the associated splitting ends up being proportional to the fermion mass. This is therefore a small correction since the lepton mass is negligible at the energies of interest here. The leading order correction is a consequence of EW symmetry breaking, as we will explore in the following section. The naive power law behavior of these effects is suppressed in the so-called ‘ultra-collinear’ regime, see [55, 58] for a detailed discussion.²⁴

²⁴ We thank the referee for clarifying this point.

3.3. Impact of subleading logs

Now that we have computed the truncated leading order PDFs, we will briefly discuss the uncertainty associated with taking the truncated leading log approximation. Generically, the DGLAP evolution equations can be expressed as

$$\frac{d}{d \log Q^2} \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_n \end{pmatrix} = \begin{pmatrix} \mathcal{P}_{1 \leftarrow 1} & \dots & \mathcal{P}_{1 \leftarrow n} \\ \mathcal{P}_{2 \leftarrow 1} & \dots & \mathcal{P}_{2 \leftarrow n} \\ \vdots & & \vdots \\ \mathcal{P}_{n \leftarrow 1} & \dots & \mathcal{P}_{n \leftarrow n} \end{pmatrix} \otimes \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_n \end{pmatrix}, \quad (24)$$

where we have individual PDFs f_i for the polarization of each particle. Typically, to solve this matrix equation to all orders, one diagonalizes this matrix, transforms to Mellin space (where the convolution becomes a product), solves

the resulting differential equations, and then transforms and rotates back.

As we have already emphasized, the electroweak gauge couplings remain perturbative throughout the range of interest, and so unsurprisingly the TLLO solution provides a reasonable approximation. In figure 6, we compare the TLLO PDFs against the EVA. At low x , the EVA deviates significantly from the TLLO result. This behavior is easy to understand, since the former is proportional to $\log s/4m_Z^2$ while the latter is proportional to $\log Q^2/m_Z^2$. When $x \rightarrow 1$, they approach the same value up to small differences due to the running couplings (we evaluate the gauge couplings at the scale $\sqrt{s}/2$ for the EVA), and the polarization effect for the photon PDFs described above. We conclude that what scale appears within the logarithm is an important difference, and the EVA does *not* provide a good approximation.

Next, we turn to exploring the size of next-to-leading-log corrections. To this end, we will solve equation (24) iteratively to extract these subleading terms. Defining the n th logarithmic order as

$$f^{(n)}(x, Q^2) \sim (\alpha \log Q^2)^n, \quad (25)$$

we can obtain $f^{(n)}(x, Q^2)$ by inserting $f^{(n-1)}(x)$ into the right-hand-side of the DGLAP equation equation (24):

$$\frac{d}{d \log Q^2} f_i^{(n)}(x, Q^2) = \sum_{j=1}^n \mathcal{P}_{i \leftarrow j}(x) \otimes f_j^{(n-1)}(x, Q^2). \quad (26)$$

At zeroth order, the only non-zero PDF is

$$f_{\mu_L}^{(0)}(x, Q^2) = f_{\mu_R}^{(0)}(x, Q^2) = \delta(1-x), \quad (27)$$

which is simply the statement that the beams would be purely composed of muons in the absence of interactions.

In figure 7, we compare the TNLLO and TLLO PDFs and parton luminosity function. The discrepancy for the PDFs between the two levels of approximation are within $\sim 10\%$ ($\sim 40\%$) for $Q = 5$ TeV (50 TeV),²⁵ and this obviously implies that the parton luminosity is also under control. We take this as strong evidence that the TLLO PDFs are sufficient unless one is interested in making precision predictions at a level that is beyond the scope of this work.

3.4. Finite mass effects

Up until this point, we have treated all the gauge bosons as being massless (up to the fact that the longitudinal mode exists). This is a good approximation when the cross section is dominated by partonic collisions whose typical scale is significantly larger than the vector masses. However, if the process of interest has non-trivial support at low τ , finite mass

²⁵ This is true except for the W_+ PDF. For this case, the TLLO pdf goes to zero as $x \rightarrow 1$, which is the source of the divergent curve in the plot. This behavior is simply due to the fact that the probability to emit a single W_+ from an on-shell μ_L is zero. However, a non-zero contribution appears at NLL in the $x \rightarrow 1$ limit, since there can now be multiple emissions. Note that this issue arises in a region where the PDF is small, so that this effect has a negligible impact on observables.

effects could become important. To understand their numerical impact, we will briefly investigate how mass effects change the calculation of the TLLO PDF.

The dominant mass effect comes from the propagator. Note that in the massless case, the emitted vector boson has virtuality

$$q^2 = -\frac{k_T^2}{(1-z)}, \quad (28)$$

while for the massive vector boson, the off-shell propagator has virtuality

$$q^2 - m_V^2 = -\frac{k_T^2 + (1-z)m_V^2}{1-z}. \quad (29)$$

Hence, the splitting function is modified as

$$\mathcal{P}_{V \leftarrow \mu} \rightarrow \frac{k_T^4}{\tilde{k}_T^4} \mathcal{P}_{V \leftarrow \mu} \quad \text{with} \quad \tilde{k}_T^2 = k_T^2 + (1-z)m_V^2. \quad (30)$$

In order to estimate the impact on the PDFs simply, we neglect the running coupling, so the Q dependence for the massless W_T PDF is

$$\int_{m_W^2}^{Q^2} \frac{dk_T^2}{k_T^2} = \log \left(\frac{Q^2}{m_W^2} \right), \quad (31)$$

while for the massive W_T it is

$$\begin{aligned} \int_{m_W^2}^{Q^2} \frac{dk_T^2}{k_T^2} \frac{k_T^4}{(k_T^2 + (1-x)m_W^2)^2} &= \log \left(\frac{Q^2 + m_W^2(1-x)}{m_W^2(2-x)} \right) \\ &+ \frac{1}{2-x} - \frac{Q^2}{Q^2 + m_W^2(1-x)}. \end{aligned} \quad (32)$$

Clearly, the difference is only relevant near threshold with $x \ll 1$.

To get a quantitative sense of the impact of including the finite mass, figure 8 shows the ratio of the massive TLLO PDF (MTLLO) to the massless TLLO PDF. The left panel shows the ratio of $f(x, Q = x\sqrt{s}/2)$ while the right one displays the ratio of $d\mathcal{L}_{ij}/d\tau(\tau, \sqrt{s}/2)$ with $\hat{s} = \tau s$. To make the comparison to m_W more explicit, we have converted the x -axis of the right panel to \hat{s} . We see that the parton luminosities can be off by as much as 40%, even for \sqrt{s} above a TeV. We conclude that finite mass effects are more important than resumming logs. We leave the exploration of the impact of these effects for future work.²⁶

4. Physics

We turn next to the physics potential of a high-energy muon collider, focusing on some of the central themes—EWSB, naturalness, and dark matter (DM)—that have motivated new

²⁶ We also note that beyond TLLO order, there are additional important effects that result from maintaining momentum conservation when emitting multiple massive partons. It would also be interesting to investigate the impact this effect would have on the conclusions comparing TNLLO to TLLO PDFs, as we did in figure 7.

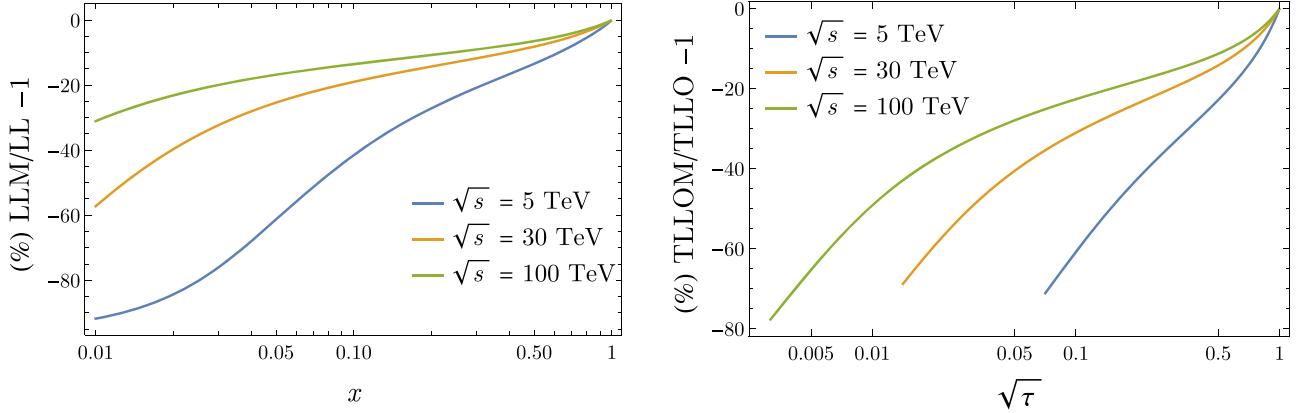


Figure 8. This figure compares the massive W_T PDF to massless W_T PDF to TLLO order. Left: a comparison of the PDFs taking $Q = x\sqrt{s}/2$. Right: a comparison of the parton luminosities $W_-^- W_+^+$ taking $Q = \sqrt{\tau s}/2$.

physics since the inception of the standard model. These considerations provide sharp goalposts for a future collider, indicating energies and luminosities that would enable such a collider to comprehensively explore the underpinnings of the standard model.

4.1. Electroweak symmetry breaking

The discovery of the Higgs completed the particle content of the SM. However, this discovery has also reinforced the puzzles associated with the Higgs field's role in the SM, generating as much confusion as clarity. The Higgs is the linchpin of the SM, responsible for all of the masses of elementary particles as well as flavor mixings, via EWSB. The majority of the SM parameters associated with the Higgs are not determined by gauge invariance, and their values must be measured. Moreover, the very fact that EWSB occurs via the Higgs is put into the SM by hand, in that we must specify the potential. Before we declare that the SM is complete, we must measure *all* of its parameters.

Unfortunately, the path to completing the goal of measuring all the SM parameters is often regarded as requiring two different colliders after the HL-LHC, due to the reliance on two qualitatively different types of observables. The first is to probe the Higgs couplings to other SM particles; we note that the light flavor Yukawa couplings have yet to be measured at all. The second is to explore the Higgs potential itself. To study the couplings of light flavors to the Higgs requires an extremely clean collider environment, which favors lepton colliders, such as the low-energy Higgs factories that have been proposed. Their advantage is clearly illustrated in figure 4, which shows that Higgs production is a relatively large fraction of the total cross section at lepton colliders, once its production via gauge bosons is kinematically allowed. Even at these colliders, there should be sufficient luminosity to probe the Yukawa couplings of the charm quark, while other light flavors pose a significant challenge; the capability to tag and measure light flavors is a subject of ongoing research [59]. However, lepton colliders offer at least the promise of measurements that would be overwhelmingly difficult at hadron colliders, where precision measurement of SM Higgs

branching fractions must overcome vast numbers of u , d , s , c and gluon background jets. Furthermore, a future lepton collider running at the Higgs mass pole could measure the s -channel resonance production to directly probe the lepton Yukawa coupling, precisely (at a muon collider) or with an upper limit of a few times the SM prediction (at an e^+e^- collider) [60–63].

While lepton colliders provide a cleaner environment than hadron colliders, their inherent drawback at the energies of the proposed ‘Higgs factories’ is the small absolute yield of Higgs particles compared to their hadron collider counterparts. For example, the Higgs factories proposed thus far would collect $\mathcal{O}(10^6)$ Higgses, whereas a 100 TeV pp collider would produce $\mathcal{O}(10^{10})$. This is exacerbated when taking into account multi-Higgs production, where a higher c.m. energy is needed. Only CLIC or the 1 TeV upgrade of an ILC would have sufficient energy for the multi-Higgs production processes to be useful to explore the Higgs potential. Therefore, a common view is that in addition to a Higgs factory, something akin to the FCC-hh is necessary to truly explore the Higgs potential.

A high energy lepton collider, such as a muon collider with $\sqrt{s} \gtrsim \mathcal{O}(10)$ TeV, can completely change this narrative by offering both a clean environment for high-precision studies *and* the energy needed to produce new final states copiously. We will make this case through examples that follow. A first indication can be gleaned from estimates using the inclusive Higgs cross sections: at an $\mathcal{O}(10)$ TeV muon collider with $\mathcal{O}(10 \text{ ab})$ of luminosity, there will be an order of magnitude more Higgs bosons produced as compared to proposed Higgs factories. Additionally, there will be $\mathcal{O}(10^4)$ di-Higgs events, which are completely inaccessible at the low-energy Higgs factories. Although a 100 TeV hadron collider with $\mathcal{O}(10 \text{ ab})$ will produce $\mathcal{O}(10^7)$ di-Higgs events, there are severe backgrounds there that grow with collider energy [64]. Despite the higher yield at a 100 TeV hadron collider such as FCC-hh, the currently best estimated sensitivity to the triple-Higgs coupling of the SM [65] can be matched or exceeded with an $\mathcal{O}(10)$ TeV muon collider [37]. We perform some new studies in this section to sketch the potential for probing Higgs physics at high energy muon colliders.

Pursuing such a program of future Higgs measurements would not simply complete the SM, it *could likely* open the first window to physics beyond the SM. The Higgs boson is unique; it is the only apparently elementary scalar among all the particles observed in the Universe. Its distinct properties provide many compelling reasons to investigate it further. The Higgs provides the only source of flavor physics in the SM; the most relevant, invariant portal to other BSM sectors or dark matter; the unitarization of scattering amplitudes in the SM; a window on early Universe cosmology via the EW phase transition (EWPT), and potentially EW baryogenesis (EWBG); and, last but not least, the naturalness puzzle. We are strongly motivated to determine whether the Higgs is *solely* responsible for EWSB, and whether it is (partially) composite. In this section, we will discuss a muon collider's role in addressing these topic.²⁷ All of these questions can be attacked by measuring the Higgs's properties with sufficient precision. Many of them benefit from the large Higgs production rate and cleanliness of a high energy muon collider, as well as the dynamical range of c.m. energy that such a collider achieves by virtue of being a vector boson collider. An apt analogy for the path that started with finding the Higgs and continues by investigating it in sufficient detail is provided by cosmology. While the expansion of the Universe was known since Hubble, it was not until many decades later that the right observable was found and measured precisely enough that the accelerated expansion of the Universe was conclusively discovered. We are now just beginning to acquire experimental knowledge of the Higgs boson's properties, at a relatively coarse level. We must move toward the new era of precision Higgs physics, which, like precision cosmology, offers the hope of revolutionizing our understanding of the Universe.

4.1.1. Higgs coupling sensitivity estimates from on-shell Higgs processes. There are many new measurements of Higgs properties that are accessible via higher energies and cleaner environments, and we will explore a sampling of these in the subsequent sections. The desire to improve Higgs precision will drive one of the core programs for any future collider. Therefore, it is important to understand how precisely a high energy muon collider could measure Higgs properties on its own, as well as in combination with other colliders. The answer to this question depends both on the theoretical framework and experimental details, which leave an enormous range of possibilities that are beyond the scope of this paper to explore. In order to make a first quantitative estimate rather than simply stating that a large number of Higgs particles would be produced, we will make a number of simplifying assumptions. In this section, we focus on the processes in which on-shell Higgs bosons are produced. At a high energy muon collider, off-shell Higgs processes will in some cases offer an even more powerful probe of Higgs properties, a topic to which we will return in section 4.1.2.

²⁷ The Higgs potential also lies at the root of deep questions about the stability of the Universe, which we will leave for future investigations.

First, we will adopt the common κ fits for Higgs precision [66, 67]. This is not an endorsement of this methodology compared to any other, but a pragmatic choice for the sake of making comparisons, as all future collider proposals have an example of this type of fit (kappa-0 framework) [68]. The inputs to such a fit are the uncertainties on the cross section measurements in exclusive channels. These depend upon the signal cross section and physics backgrounds, as well as machine backgrounds, detector capabilities, and possible additional theoretical assumptions. The machine backgrounds and detector capabilities are particularly interesting in the context of high-energy muon colliders, as previously discussed. The BIB at muon-colliders serves both as a background to measurements and a driver of detector design. There is no optimized detector design available at all our benchmark c.m. energies, due to the fact that the BIB depends on the accelerator complex within roughly 25 m on each side of the interaction point. Therefore, we will simply choose our energy and luminosity benchmarks to be 10 TeV and 10 ab, and using the Muon collider detector card for the DELPHES fast simulation [69]. This choice of 'detector' does not serve as a final word, but allows us to begin exploring how the physics requirements interact with detector design. We do not include the effects of BIB beyond the anticipated reduction in geometric acceptance, as a realistic treatment is beyond the scope of this paper. Recent full-simulation studies provide encouraging signs that BIB may be brought under control, especially at higher energies [70, 71], but further study is warranted. Furthermore, for this toy study we *do not* include physics backgrounds.

While our assumptions may seem like too drastic of a simplification, there still is useful sensitivity information despite having made these naively non-conservative estimates. Our signal rates using DELPHES have a rather small acceptance, given that the detector card limits physics objects to $|\eta| < 2.5$ except for forward muons. There are a number of motivations for the detector card inspired from a hybrid of CLIC and FCC-hh for efficiencies and reconstruction [72], but these are not optimized for a particular physics target or energy. Additionally, the general acceptance roughly coincides with having BIB-suppressing tungsten nozzles [73, 74], with a 10° opening angle motivated by 1.5 TeV c.m. muon collider studies [74, 75]. The nozzle opening should be able to be reduced at higher energies since the radiation will be more forward [36] (and timing should also mitigate BIB effects). Physics backgrounds, of course, potentially matter a great deal more than BIB. However, they are significantly reduced at a lepton collider as shown in figure 4. Detailed sensitivity studies including physics backgrounds have been performed for 3 TeV lepton colliders for CLIC Higgs studies [76], and thus serve as a proof of principle (or potential floor) for our signal-driven sensitivities. There are a variety of studies that can and should be done in the future, but we hope this serves as a useful starting point by showing the effects of acceptance and efficiency via fast simulation. From the perspective of signal and BIB, we take this to be a conservative starting point.

Table 2. Signal rates and efficiencies for selected Higgs production channels at a 10 TeV muon collider using signal-only selection and the *DELPHES* muon collider fast simulation.

10 TeV @ 10 ab ⁻¹				
Production	Decay	Rate (fb)	$A \cdot \epsilon$ (%)	$\Delta\sigma/\sigma$ (%)
W-fusion	bb	490	7.4	0.17
	cc	24	1.4	1.7
	jj	72	37	0.19
	$\tau^+\tau^-$	53	6.5	0.54
	$WW^*(jj\ell\nu)$	53	21	0.30
	$WW^*(4j)$	86	4.9	0.49
	$ZZ^*(4\ell)$	0.1	6.6	12
	$ZZ^*(jj\ell^+\ell^-)$	2.1	8.9	2.3
	$ZZ^*(4j)$	11	4.6	1.4
	$\gamma\gamma$	1.9	33	1.3
Z-fusion	$Z(jj)\gamma$	0.9	27	2.0
	$\mu^+\mu^-$	0.2	37	0.37
W-fusion tth	bb	51	8.1	0.49
	$WW^*(4j)$	8.9	6.2	1.3
W-fusion tth	bb	0.06	12	12

Table 2 shows, for various channels, the results for cross sections, acceptance \times efficiency, and the measurement precision for our 10 TeV muon collider benchmark. The acceptances are based on minimal cuts. As a starting point, we choose these to approximate the one existing full-simulation study that focuses on $h \rightarrow b\bar{b}$ [71] at 1.5 TeV, which was also extrapolated to 10 TeV. For two-body final states of the Higgs, such as in [71], we make minimal cuts on reconstructed objects such as $p_T > 40$ GeV, with the VLC jet algorithm [77, 78] run for an exclusive number of jets with an $R = 0.5$ as implemented in *DELPHES*, and we use the tight-tagging working point for b -tagging. We find that we have good agreement with the very conservative study done in [71], and therefore it serves as another calibration point for extrapolating to the performance in other final states. For all other physics objects, we use the standard parameterization as found in the muon detector card, except for charm-tagging, which is not implemented. For the $c\bar{c}$ final state we apply a flat 20% tagging efficiency for each c -jet, as inspired by CLIC [79]. For greater than two-body final states, we reduce the p_T requirements to 20 GeV and reduce the tagging efficiencies to loose tags for b -tagging, such as in $t\bar{t}h$.

We stress again that the performance in the various channels shown in table 2 is in no way optimized and needs further study. However, with these putative sensitivities, we can perform a simple 10-parameter κ fit to compare this benchmark to other proposed colliders. Of course, the muon collider can be enhanced with complementary measurements from other colliders as well. Therefore we also perform the fit with the HL-LHC or a 250 GeV e^+e^- collider included. Here we took the CEPC input with full correlation matrix for different channels [6, 80] to represent the 250 GeV e^+e^- collider. The results and discussion on complementarity with other lepton collider Higgs factories would be similar. We present the results of these fits in table 3.

Table 3. Results of a 10-parameter fit to the Higgs couplings in the κ -framework, based on the attainable precision in each on-shell Higgs production and decay channel listed in table 2. Additionally, we include the effects of adding data sets projected from the HL-LHC and a 250 GeV e^+e^- Higgs factory. One should keep in mind that a muon collider will also strongly constrain Higgs properties via off-shell measurements, which are not included here.

Fit result (%)		
10 TeV muon collider	With HL-LHC	With HL-LHC + 250 GeV e^+e^-
κ_W	0.06	0.06
κ_Z	0.23	0.22
κ_g	0.15	0.15
κ_γ	0.64	0.57
$\kappa_{Z\gamma}$	1.0	1.0
κ_c	0.89	0.89
κ_t	6.0	2.8
κ_b	0.16	0.16
κ_μ	2.0	1.8
κ_τ	0.31	0.30
		0.27

It is impossible to directly compare our Higgs sensitivity measurements to other proposed colliders [68] given the signal-only nature of the results presented here. However, it is easy to see when looking at table 3 of [68] (which served as the input to the European Strategy Report) that the level of signal-only precision reached in our estimates is not orders of magnitude better than proposed Higgs factories. This can be traced back to the rather small $A \cdot \epsilon$ for our ‘detector’ choice. This motivates exploring detector design to see how much the effective acceptance can be increased while being able to keep BIB in check. Additionally, the cross sections we use are for unpolarized muons, and in fact, the single Higgs VBF cross section for unpolarized muons is not that different at 10 TeV from the polarized cross section for CLIC at 3 TeV. Therefore, it is important to understand the impact and potential for polarization of future high-energy muon colliders. Nevertheless, we emphasize that even in this first simplified study, a 10 TeV muon collider provides similar numbers to other colliders, and there is a great deal of room for improvement. Additionally, given the extra energy, the muon collider can achieve much better precision on Higgs self-interactions [37] using the same machine. Moreover, the real untapped potential for a high-energy muon collider comes from its ability to make novel measurements of off-shell Higgs couplings. This is a feature that is potentially shared in common with a 100 TeV pp collider, but is relatively unexplored as of yet. In the next section, we explore an example of this type of approach, which could measure the top Yukawa with a precision of $\mathcal{O}(1\%)$. With all of these caveats in mind, a high energy muon collider is an impressive Higgs factory as well as a discovery machine, and there are numerous interesting avenues for future work related to the Higgs.

4.1.2. Flavor and exotic couplings. Flavor physics in the SM only arises through the Higgs couplings, which determine both the mass pattern of the different generations and the mixings

that allow for flavor changing processes. Taking as motivation that flavor is one of the strangest aspects of the SM, there has been a rich history of testing the flavor structure of the SM indirectly using measurements from intensity frontier experiments. This program has resulted in stringent bounds on flavor changing processes, probing new physics scales that are naively well out of the direct reach of any future energy frontier experiment. Nevertheless, not all of the SM Yukawas have been measured yet, and large deviations in flavor diagonal Higgs couplings due to BSM physics are possible [81, 82] as well as smaller flavor-changing BSM Higgs couplings [83], depending on the particular flavors involved. Measuring the SM Yukawas may require more than an $\mathcal{O}(10)$ TeV muon collider. Any channel with a branching fraction similar to $\text{Br}(h \rightarrow \mu^+ \mu^-) \sim \mathcal{O}(10^{-4})$ will result in an absolute yield of 10^3 decays before backgrounds, acceptances, and efficiencies are accounted for. Nevertheless, if detectors are optimized, there is still the possibility to go after first generation couplings directly. If BSM deviations exist, even higher energy muon colliders will only have a greater physics potential. For example, current LHC data allows an enhancement of $\text{Br}(h \rightarrow d\bar{d})$ by $\mathcal{O}(10^6)$, which is well within the reach of a muon collider [82].

Beyond just measuring properties of the SM Higgs which should exist, we can use the Higgs as a potential window on unexpected new physics beyond the SM. The $H^\dagger H$ operator is the lowest dimension gauge and Lorentz invariant building block in the SM. Therefore, if there are new states beyond the SM that are lighter than the Higgs, it is quite likely that the Higgs will have some branching fraction to decay to them. These are known as *exotic Higgs decays* and have been a subject of significant recent study [84]. Given that a muon collider will have a clean environment with the additional benefit that it would produce a larger number of Higgses than the e^+e^- Higgs factories, it provides an excellent opportunity to investigate exotic Higgs decays further. Moreover, as at FCC-hh, given that the dynamic range of energies available for the Higgs grows as the lepton collider \sqrt{s} increases, there are a variety of other probes one can employ to test Higgs couplings beyond simply studying branching fractions of on-shell Higgs bosons.

4.1.2.1. $W^+W^- \rightarrow t\bar{t}$: a longitudinal scattering case study As an example of the power of having both precision and a dynamic range of energies available to measure Higgs couplings, we consider the classic example of the interplay between Higgs physics and perturbative unitarity. If the Higgs boson's couplings are not precisely those predicted by the standard model, then, in the absence of other new physics, scattering amplitudes of longitudinal gauge bosons will grow with energy and eventually violate perturbative unitarity bounds [85–90]. This allows high-energy colliders to probe new physics operators that involve the Higgs boson by studying scattering processes with external gauge bosons, rather than Higgses. This approach, very different from that taken at Higgs factories, has been dubbed ‘Higgs without Higgs’ [91].

A particularly interesting test case for this program at a muon collider is the measurement of the top Yukawa coupling. While an e^+e^- Higgs factory is especially well-suited to a high-precision measurement of the Higgs coupling to gauge bosons (especially through the Higgsstrahlung process), the top Yukawa coupling will be less constrained. If we assume that the top Yukawa deviates from its standard model value by a fraction δ_{BSM} , such that

$$y_t \mapsto y_t(1 + \delta_{\text{BSM}}), \quad (33)$$

then we expect the 95% confidence limit on δ_{BSM} after HL-LHC to be $|\delta_{\text{BSM}}| \lesssim 0.06$ (using figure 135 of [92]). This is a much weaker bound than the $\sim 10^{-3}$ precision of the hZZ coupling expected at a Higgs factory. Hence, we focus on the top Yukawa for a first case study of the potential Higgs coupling reach of a muon collider.

The scattering amplitude for top production via longitudinal W bosons when $\delta_{\text{BSM}} \neq 0$ scales as [90]

$$\mathcal{M}(W_L^+ W_L^- \rightarrow t\bar{t}) \simeq -\frac{m_t}{v^2} \delta_{\text{BSM}} \sqrt{\hat{s}}, \quad \text{with } \sqrt{\hat{s}} \gg m_t. \quad (34)$$

Taking into account only this growing term in the amplitude, we estimate that perturbative unitarity is violated at a scale $\Lambda_{\text{BSM}} \lesssim \frac{10 \text{ TeV}}{\delta_{\text{BSM}}}$. For small δ_{BSM} , this is well above the energy scale of a potential muon collider, so it is theoretically consistent to treat new physics in this sector via the parameter δ_{BSM} without specifying the UV completion.

In figure 9, we show the differential distribution $d\sigma/dm_{\bar{t}\bar{t}}$ at a 14 TeV muon collider, both for the case with unpolarized muon beams and the case with a fully left-handed μ^- and right-handed μ^+ . (FEYNARTS [93] and FEYNCALC [94–96] were used to perform these computations.) We see that the W^+W^- initial state is dominant in the case of polarized beams, increasing the possible sensitivity to the enhanced $W_L^+ W_L^- \rightarrow t\bar{t}$ process. In figure 10, we present the 2σ sensitivity of a muon collider to the parameter δ_{BSM} . The sensitivity is computed from the difference in BSM and SM predictions for the differential distribution $d\sigma/dm_{\bar{t}\bar{t}}$ integrated over a set of bins.²⁸ (This assumes that $t\bar{t}$ events can be detected with high efficiency, an assumption that should be checked with detector simulations in the future.) We find that percent-level deviations in the top Yukawa can be probed with luminosities $\sim 10 \text{ ab}^{-1}$ at a muon collider with polarized beams. The reach with unpolarized beams suffers (requiring roughly an order of magnitude more luminosity to achieve the same sensitivity). Notice that it is easier to probe negative values of δ_{BSM} , as these interfere constructively with the standard model process at smaller invariant masses $m_{\bar{t}\bar{t}}$.

Beyond allowing for precision measurements of Higgs couplings to SM particles, the strategy discussed in this section can

²⁸ We divide the energy range into 20 bins (with smaller bins at lower $m_{\bar{t}\bar{t}}$, where the cross section is larger) and find the value of δ_{BSM} for which the Poisson log likelihood difference $2\Delta \log L = 4$, where we compute this difference as $2\sum_{i \in \text{bins}} (n_i^{\text{SM}} - n_i^{\text{BSM}} + n_i^{\text{BSM}} \log(n_i^{\text{BSM}}/n_i^{\text{SM}}))$, where $n_i^{\text{(B)SM}}$ is the model's predicted mean for bin i and need not be an integer. This is a rough proxy for the sensitivity that one might obtain by doing pseudo-experiments. We have checked that the result is not very sensitive to the choice of binning.

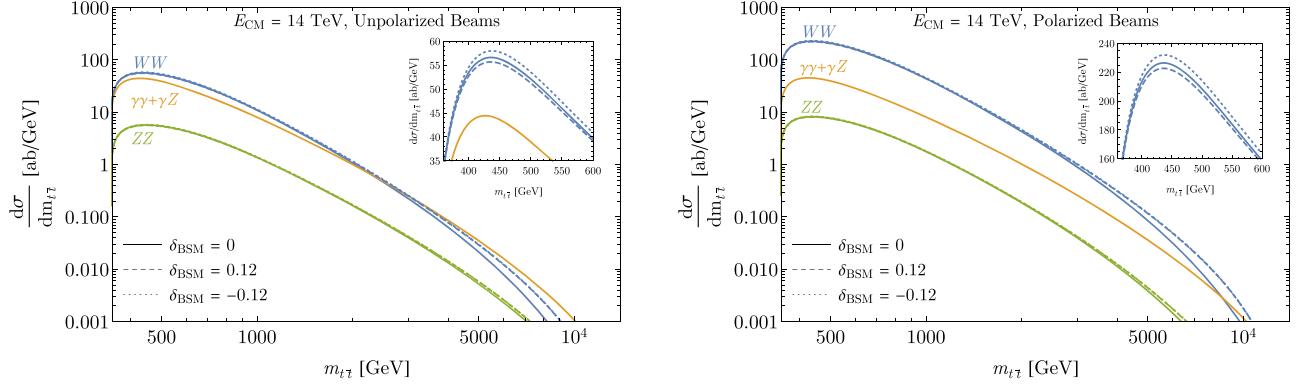


Figure 9. Differential cross section for $\mu^+ \mu^- \rightarrow t\bar{t} + X$ from different gauge boson fusion processes at a 14 TeV muon collider, with unpolarized beams (left) or fully polarized (left-handed μ^- and right-handed μ^+) beams (right). At high energies, a deviation from the standard model top Yukawa leads to a significant increase in the rates for the $W_L^+ W_L^- \rightarrow t\bar{t}$ process. At low energies (visible in the insets), it produces either destructive interference ($\delta_{\text{BSM}} > 0$) or constructive interference ($\delta_{\text{BSM}} < 0$).

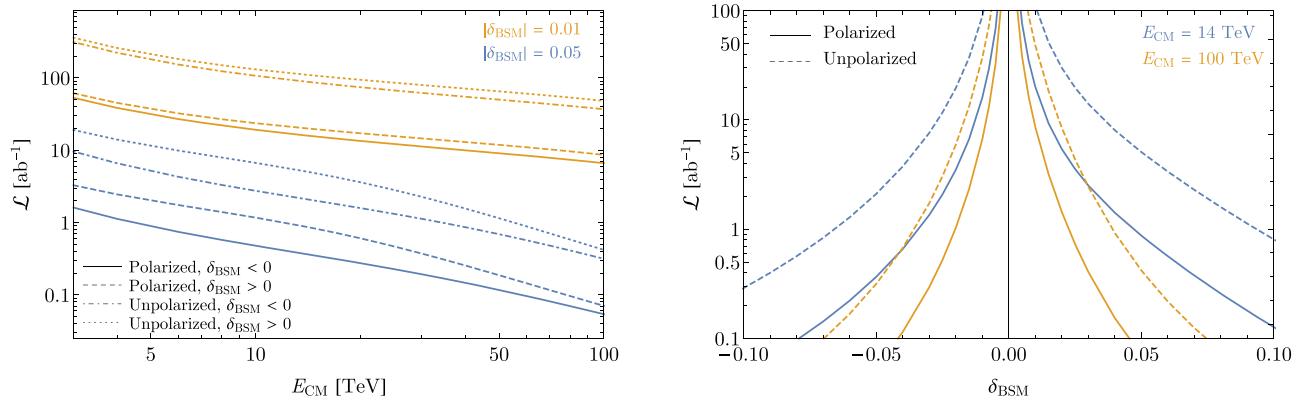


Figure 10. Luminosity needed to distinguish a modified top Yukawa coupling parameterized using δ_{BSM} from the standard model at 2σ confidence, using the differential rate $d\sigma/dm_{t\bar{t}}$ of the process $\mu^+ \mu^- \rightarrow t\bar{t} + X$. Left: the result is shown as a function of the c.m. energy for various benchmark values of δ_{BSM} . Right: the result is shown as a function of δ_{BSM} for $E_{\text{CM}} = 14$ and 100 TeV (right).

also be implemented to probe the existence of new degrees of freedom that, although kinematically accessible, are very rarely produced. In appendix A.4.2, we illustrate this point in the context of one of the most elusive BSM scenarios: a \mathbb{Z}_2 -symmetric SM-singlet that interacts with the SM only through the Higgs-portal. As we discuss there, a strategy based on exploiting the resulting kinematic features in the differential cross section for the process $\mu^+ \mu^- \rightarrow t\bar{t} + X$ may be competitive with the traditional missing-mass analysis that is the focus of appendix A.4.1.

4.1.3. The Higgs potential and the electroweak phase transition. One of the most intriguing aspects of EWSB is its role in the early Universe. Because we cannot directly observe the early Universe before the time of formation of the CMB other than through gravitational waves, we must make use of particle physics to draw inferences about what occurred. Many interesting and yet unmeasured epochs in cosmology are directly intertwined with EWSB. For example, the evolution of neutrinos in the Universe and the properties of the cosmic neutrino background depend crucially on the W and Z boson masses. The masses of SM particles arise from EW symmetry breaking, and so may have turned on

during the EWPT in a thermal history in which EW symmetry was restored at even earlier times and hotter temperatures. If the EWPT was strongly first order and other sources of CP violation exist—both of which require new physics beyond the SM—then EWBG could explain the matter/antimatter asymmetry in our Universe.

Since we cannot directly measure the Higgs potential at finite temperature, we are relegated to studying its zero temperature behavior, and possible couplings of the Higgs to other particles. Unfortunately, we also cannot access the Higgs potential away from its minimum at colliders, so we are left to study the shape of the potential locally through its derivatives, i.e., measuring the Higgs self couplings. This has motivated the intense study of what can be learned from di-Higgs production, since the previous Snowmass process in 2013. As mentioned earlier, a high energy muon collider has the ability to measure di-Higgs production and thus the triple Higgs coupling with precision similar to or better than a 100 TeV hadron collider. However, from the perspective of BSM physics, there is rarely *just* a shift in the triple Higgs coupling alone. This is clear from the EFT perspective, where there are multiple operators that can change the di-Higgs production rates. If one wants to focus

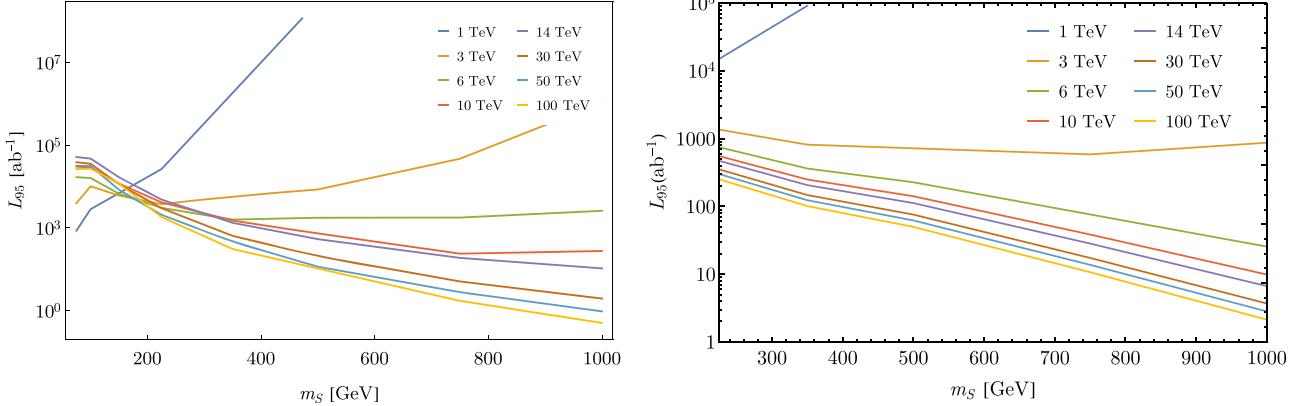


Figure 11. The integrated luminosity \mathcal{L} in units of ab^{-1} required to exclude a singlet scalar whose mass m_S is due entirely to EWSB at 95% CL, for muon colliders operating at various c.m. energies. Left: projected exclusion reach from missing energy searches using the naive combination of $SS + \gamma/h/Z$ final states. Right: projected exclusion reach from the invariant mass distribution of $t\bar{t}$ pairs produced in VBF. In both cases the integrated luminosity required for exclusion decreases at higher m_S because the singlet is assumed to acquire its entire mass from EWSB, in which case the signal cross section grows as a function of m_S .

solely on the h^6 operator, this can only be realized in singlet extensions of the SM.

While singlet extensions of the SM can come in a variety of forms, there is a particularly interesting model that serves as a ‘nightmare scenario’ [97], where the singlet is protected from mixing with the SM Higgs by a \mathbb{Z}_2 symmetry and is heavier than $m_h/2$ so that the Higgs cannot decay to this state. By studying this scenario, one can set a worst-case benchmark for how well a strong EWPT can be tested. Additionally, this \mathbb{Z}_2 singlet model can serve to benchmark the more recent investigations into EW symmetry non-restoration [98] and also as a proxy for neutral naturalness [99, 100].

There are at least two compelling channels in which to probe the ‘nightmare scenario’ at a high-energy muon collider, beyond indirect constraints that may be obtained from measurements of the triple Higgs coupling or Higgs couplings to other SM particles. One of these is the natural extension of the strategy pursued at the LHC and proposed for future proton–proton colliders: to use a search for missing energy in conjunction with one or more visible particles produced through ISR or associated production. In appendix A.4.1 we present a simplified analysis to assess the prospects of a high-energy muon collider in these final states, combining searches for missing mass in conjunction with an ISR photon, Higgs boson, or Z boson. Alternately, given that missing energy searches are limited by significant backgrounds even at a muon collider, in appendix A.4.2 we propose another novel search strategy that leverages the significant contribution of longitudinally polarized W and Z bosons to certain processes. In particular, radiative corrections from a \mathbb{Z}_2 singlet scalar with large coupling to the Higgs give rise to a feature as the invariant mass of the final state crosses over the singlet threshold of, e.g., $t\bar{t}$ pairs produced in VBF. This feature can be large enough to be distinguished from the standard model contribution provided sufficient control over theory systematics in the $t\bar{t}$ invariant mass distribution, and hence could offer comparable sensitivity to a direct search for the singlet in the missing energy final state. To our knowledge,

this is an entirely new way of searching for the ‘nightmare scenario,’ and one that leverages the unique strengths of a high-energy muon collider. The results of the two studies are summarized in figure 11, which shows the integrated luminosity required to exclude a singlet scalar obtaining all of its mass from EWSB for various collider energies. We conclude that a 10 TeV–30 TeV muon collider can easily compete with or exceed the reach of a 100 TeV collider for this compelling scenario [97, 101].

4.1.4. Additional Higgs bosons. As a final case study demonstrating the potential for a high-energy muon collider to illuminate the physics of EWSB, we consider the search for additional ‘Higgs bosons’ that acquire their standard model couplings by mixing with the Higgs. This is exemplified by one of the simplest extensions of the Higgs sector, a real scalar singlet with renormalizable couplings to the SM Higgs.²⁹ This encodes a large class of BSM theories which address the stability of the electroweak scale [99, 102], or relate the baryon asymmetry in the Universe today with the EWPT [103–107]. More generally, given the diversity of vector bosons and fermions in the SM, it is natural to ask if the scalar sector possesses similar depth.

A scalar singlet SM extension is a very useful benchmark to assess the capabilities of future colliders [28, 108], since it manifests itself in a two-fold way: indirectly, as modification of the Higgs decay rates, and directly, in single and double production channels. Both these effects are controlled by the same small set of parameters—notably the singlet mass and its mixing with the Higgs boson—allowing for an immediate comparison of the direct and indirect reach. As we shall see, the ability of a very high energy lepton collider to discover heavy resonances is crucial to overcoming the limitations of Higgs precision measurements, which are inevitably constrained by systematic uncertainties, and allows the exploration of entirely

²⁹ The sensitivity of muon colliders to extended Higgs sectors with electroweak doublets was recently studied in [31].

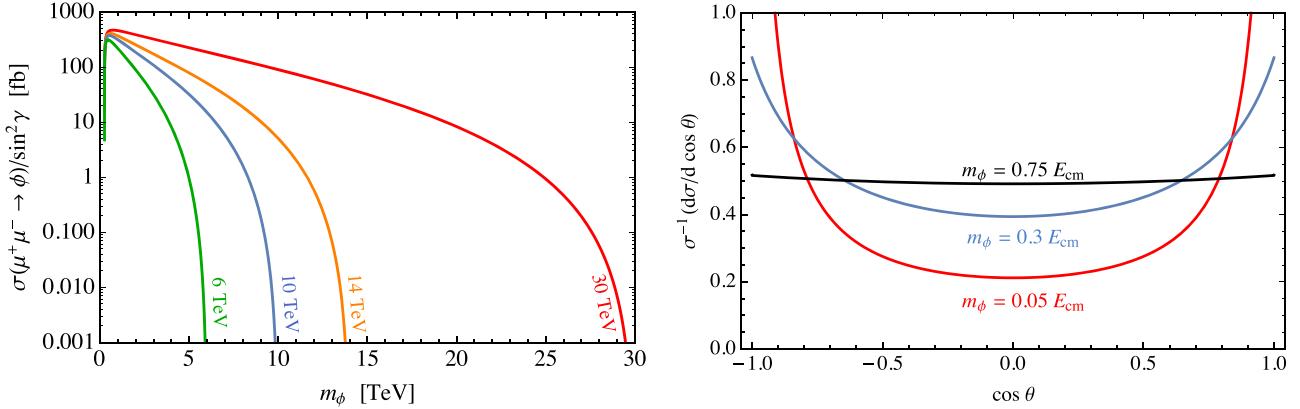


Figure 12. Left: total cross sections for $\ell^+\ell^- \rightarrow \phi\nu\bar{\nu}$ as functions of the mass m_ϕ , for different collider energies. Right: differential distributions of the decay products of $\phi \rightarrow hh$ in the scattering angle θ , for different values of m_ϕ ; one can see that lighter scalars are more boosted and decay in the forward and backward directions.

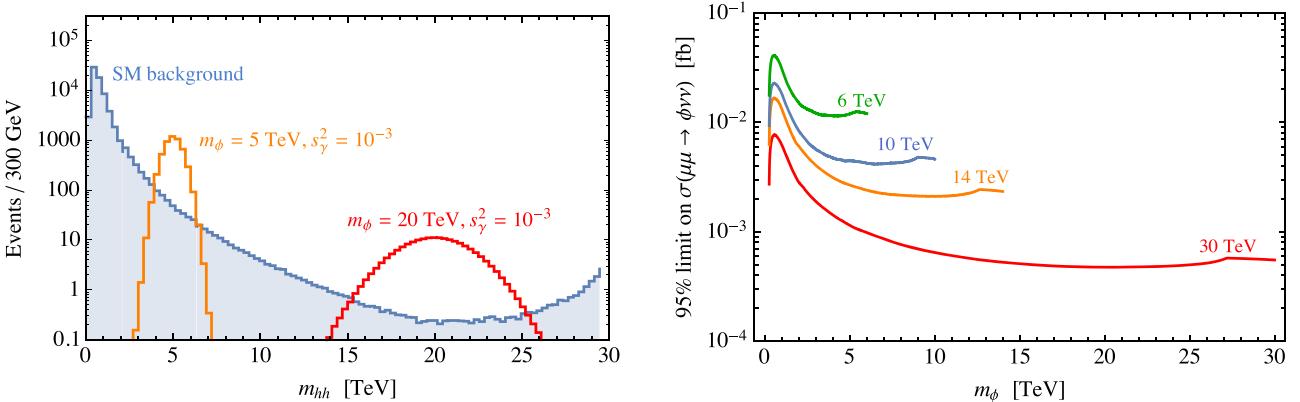


Figure 13. Left: distribution of double-Higgs events in the hh invariant mass at a muon collider with energy $E_{\text{cm}} = 30$ TeV. The SM background $\mu^+\mu^- \rightarrow hh\nu\bar{\nu}$ is shown in blue, and the resonant $\phi \rightarrow hh$ production (with $s_\gamma = 10^{-3}$) is superimposed for $m_\phi = 5$ TeV (orange) and $m_\phi = 20$ TeV (red). Right: expected 95% C.L. exclusion on the $\mu^+\mu^- \rightarrow \phi\nu\bar{\nu}$ fiducial cross section at the various muon collider benchmarks of table 1.

new territory involving weakly interacting new physics in the 10 TeV range.

The singlet phenomenology is dictated by the following Lagrangian

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{1}{2}(\partial_\mu S)^2 - V(S) - a_{\text{HS}} S |H|^2 - \lambda_{\text{HS}} S^2 |H|^2, \quad (35)$$

where two portal operators between the SM Higgs H and the singlet field S are possible at the renormalizable level. The mass eigenstates are identified via a rotation of an angle γ ,

$$h = h^0 \cos \gamma + S \sin \gamma, \quad \text{and} \quad \phi = S \cos \gamma - h^0 \sin \gamma, \quad (36)$$

where h^0 is the neutral component of the SM Higgs doublet, h is the SM-like state, and ϕ the new singlet-like scalar. Since S is a complete singlet, all interactions of ϕ with SM states other than the Higgs proceed through this mixing and are controlled by $\sin \gamma$. At the same time, all single Higgs boson couplings to SM fermions and vectors are rescaled by the same factor $\kappa_V = \kappa_f = \cos \gamma$.

The mixing angle can be read directly from the mass matrix

$$\gamma \simeq \frac{v(a_{\text{HS}} + \lambda_{\text{HS}} s)}{V''(s) + \lambda_{\text{HS}} v^2}, \quad (37)$$

under the assumption of small mixing angle, and where s is the VEV of the singlet. From the general formula above we can distinguish two cases: (i) if the singlet gets a VEV s due to its potential and $m_\phi \simeq g_* s$, where g_* is some coupling, then the mixing scales as $\gamma \simeq g_* v / m_\phi$; (ii) if the singlet gets a VEV only through its interaction with the SM Higgs then the mixing scales as $\gamma \simeq v a_{\text{HS}} / m_\phi^2$ and can be made arbitrarily small. For instance if we assume $a_{\text{HS}} = g_* v$, then $\gamma \simeq (g_* v / m_\phi)^2$ and the mixing decouples with one extra power of the ratio between the EW scale and the singlet mass compared to the case (i). We will be dealing mostly with the first class of models here, but the second scenario is useful to keep in mind.

Production modes and decay channels

At a high-energy muon collider, the dominant production mode for the scalar ϕ comes from VBF [28]. By exploiting

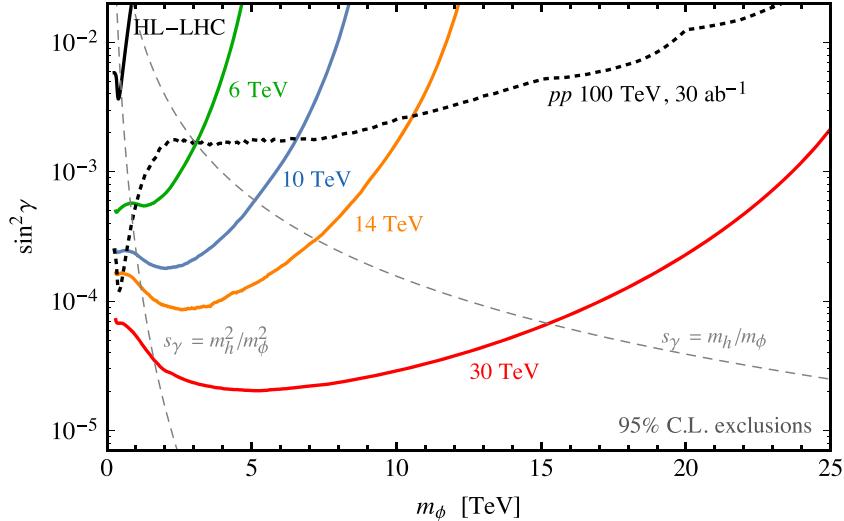


Figure 14. Exclusions on the mixing angle of a generic scalar singlet, $\sin^2\gamma = \kappa_V - 1$, as a function of the singlet mass m_ϕ for the various collider benchmarks (colored lines). The expected limits at HL-LHC (solid) and a FCC-hh (dashed) are shown as black lines for comparison. The thin dashed lines indicate the two possible scalings of the mixing angle with m_ϕ in realistic models with fixed coupling.

the scattering of equivalent Goldstone bosons we can compute both the single and double production analytically.

Single production proceeds via mixing with the Higgs. Its cross section is proportional to the mixing angle and is only logarithmically sensitive to the mass (at high energies), and it can be simply written as

$$\sigma_{\mu\mu \rightarrow \phi} \simeq \sin^2\gamma \frac{g^4}{256\pi^3} \frac{\log(s/m_\phi^2)}{v^2}. \quad (38)$$

A further dependence on the mass of the singlet-like state is hidden in the mixing $\sin^2\gamma$, as emphasized above. The dependence of the total cross section on the singlet mass is shown in the left panel of figure 12.

Double production, on the contrary, mainly depends on the quartic portal coupling and the mass of the singlet. The total cross section can again be computed analytically exploiting the scattering of equivalent Goldstone bosons, and reads

$$\sigma_{\mu\mu \rightarrow \phi\phi} \simeq \frac{g^4 |\lambda_{HS}|^2}{49 \cdot 152 \pi^5} \frac{\log(s/m_\phi^2)}{m_\phi^2}. \quad (39)$$

The decay channels of the singlet-like scalar are inherited by the mixing with the Higgs boson, and for $m_\phi \gg m_W$ they are related by an approximate $SO(4)$ (custodial) symmetry, which implies

$$\Gamma(h \rightarrow WW) = 2\Gamma(\phi \rightarrow ZZ) = 2\Gamma(\phi \rightarrow hh) = \sin^2\gamma \frac{m_\phi^3}{8\pi v^2}. \quad (40)$$

The width of the di-top decay channel is

$$\Gamma(\phi \rightarrow t\bar{t}) = \sin^2\gamma \frac{3y_t^2 m_\phi}{16\pi}, \quad (41)$$

and is subleading for all singlet masses. We note that invisible and/or displaced decays widths could be present if singlet interactions with extra dark sector states are allowed [109].

Sensitivity of a muon collider

We now estimate the reach of a muon collider for single scalar production. The main decay channel at a lepton collider is $\phi \rightarrow hh \rightarrow 4b$, which gives a rather clean signature with a large branching fraction. In the following, we assume $\text{Br}(\phi \rightarrow hh) = 25\%$, as predicted in the $m_\phi \gg m_h$ limit. Other relevant and complementary decay modes are $\phi \rightarrow W^+W^-, ZZ$, which are important when the $\text{Br}(\phi \rightarrow hh)$ is small.

We consider the search in the $hh \rightarrow 4b$ channel for a resonant hh pair over the SM $\ell^+\ell^- \rightarrow hh$ background. The latter is the main source of background, provided the Higgs bosons can be reconstructed with a sufficiently high accuracy. We do not include other sources of background, which mainly come from $\ell^+\ell^- \rightarrow Vh, VV$ (with $V = W, Z$), assuming they can be isolated from the signal, e.g., by including a window cut on the bb invariant masses. The SM hh background is simulated with MadGraph at parton level, without including Higgs decays or detector effects. We instead assume an overall signal selection efficiency of $\epsilon_{hh} = 30\%$, which is consistent with other analyses performed for high-energy lepton colliders [28, 39, 110, 111]. Furthermore, we impose an acceptance cut $p_T > 20 \text{ GeV}$, $\eta > 2$ on the Higgs bosons. Notice that these cuts are irrelevant for high invariant masses (relevant for heavy singlets), but become important for lower masses, where they cut-off the logarithmic enhancement of the cross section due to the forward singularity of VBF production.

The invariant mass distribution of the SM events is shown in blue in figure 13 (left), where two examples of signal are also shown. The limit on the cross section is obtained performing a cut-and-count experiment around the resonance peak, by requiring the di-Higgs invariant mass M_{hh} to lie within $\pm 15\%$ of the resonance m_ϕ . The reach on the cross section is obtained imposing

$$\frac{S}{\sqrt{S + B + \alpha^2 B^2}} = 2, \quad (42)$$

where the factor proportional to $\alpha = 3\%$ takes into account possible systematic uncertainties. The results for the various collider benchmarks are shown in figure 13 (right).

The limits on the cross sections are then translated into a reach on the mixing angle $\sin \gamma$ by comparing it with the signal cross section. The differential cross section as a function of the Higgs scattering angle θ_h for $\ell^+ \ell^- \rightarrow \phi \nu \nu \rightarrow hh \nu \nu$ can be computed analytically in the EVA (see section 3.1), and reads

$$\frac{d\sigma_{\ell^+ \ell^- \rightarrow hh \nu \nu}}{d \cos \theta_h} = \frac{m_W^4}{16\pi^3 v^6} \frac{1}{\sin^2 \theta_h} \left[\frac{m_\phi^2}{E_{\text{cm}}^2} - 1 + \frac{2m_\phi}{E_{\text{cm}} \sin \theta_h} \right. \\ \left. \times \left(\arctan \frac{E_{\text{cm}} \cot \frac{\theta_h}{2}}{m_\phi} - \arctan \frac{m_\phi \cot \frac{\theta_h}{2}}{E_{\text{cm}}} \right) \right]. \quad (43)$$

The angular distribution of hh events is plotted in the right panel of figure 12, for different values of the singlet mass. This expression is then integrated over the phase-space region defined by the acceptance cuts to find the number of signal events.³⁰ The result has been cross-checked by generating the signal events with MadGraph, after implementing the singlet model in FEYNRULES, finding perfect agreement. The reach in $\sin^2 \gamma$ is shown in figure 14 (left) as a function of the singlet mass. In the same figure we also show the reach of HL-LHC and FCC-hh for comparison. Notice that a muon collider with c.m. energy in the tens of TeV range could reach a sensitivity $s_\gamma^2 \lesssim 10^{-4}$ or lower, corresponding to deviations in Higgs couplings of $\mathcal{O}(10^{-4})$, which are beyond the capability of any present or other proposed future collider. This example highlights the great qualitative advantage of a high-energy muon collider in probing the Higgs sector: the longitudinal enhancement of VBF and relatively modest backgrounds give exquisite sensitivity to any additional scalars participating in EWSB.

4.2. Dark matter

The predominance of apparently non-baryonic matter in the Universe remains one of the few unambiguous indicators of physics beyond the standard model, and identifying the microscopic properties of dark matter is a central goal of multiple fields. Among the many candidates for particle dark matter, the weakly interacting massive particle (WIMP) paradigm has long been one of the most compelling. Within this paradigm, dark matter candidates arising as the lightest member of an electroweak (EW) multiplet form a particularly simple class of models [38, 112, 113]. The thermal relic abundance of such ‘minimal’ dark matter is fixed strictly in terms of the quantum numbers of the electroweak multiplet in question, picking out a high mass scale between 1–23 TeV for $SU(2)_L$ representations ranging from doublets to septuplets. This makes minimal dark matter a motivated but difficult scenario for colliders in

light of the high mass scale. Additionally, it is challenging from the detector point of view, because the typically small splittings of the EW multiplets suppress the amount of visible energy (and hence, missing momentum) in a typical event. Nevertheless, the abundant electroweak cross sections and relatively low irreducible backgrounds at a muon collider make it well positioned to search for minimal dark matter, to the point where a muon collider of sufficient energy could potentially render a decisive verdict on the scenario. In this section, we summarize the studies performed in [34], adapting their projections to the optimistic and conservative luminosity targets presented here.

Perhaps the best-known examples of minimal dark matter are the $SU(2)_L$ doublet and triplet, which can be mapped onto the higgsino and wino in supersymmetric theories. However, it is also interesting to consider multiplets with quantum numbers $(1, n, Y)$ under the SM gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$. For definiteness, we restrict our attention to fermions, whose only renormalizable SM interactions are with electroweak gauge bosons and whose mass arises from a vector-like mass parameter. The resulting mass degeneracy among members of the multiplet is split by EW loop corrections [112–116]. A number of considerations shape the motivated values of n and Y . For $n > 7$, the large electroweak charge induces a Landau pole in the standard model gauge couplings about one to two orders of magnitude above the mass of the EW multiplet [117]. As such, we will restrict our attention to $n \leq 7$. For a given n , a specific value of Y ensures that the lightest eigenstate of the EW multiplet is neutral and hence a suitable dark matter candidate.

For odd-dimensional multiplets $(1, n = 2T + 1, Y)$ with $T \in \mathbb{Z}^+$, $Y = 0$ ensures the electrically neutral member χ is always the lightest mass eigenstate in the multiplet, and fermions in these multiplets may be either Majorana or Dirac. Beyond the renormalizable level, irrelevant operators could allow the dark matter to decay. Such operators could be forbidden by introducing a small hypercharge $Y = \epsilon$ [118] such that the dark matter acquires an electric charge $Q = \epsilon$. However, other alternatives are possible, e.g., gauged discrete symmetries, which avoid the existence of global symmetries and ensure that the neutral particle is a good dark matter candidate. We will remain agnostic as to the particular mechanism ensuring the stability of the dark matter candidate, since it is largely irrelevant for the signals at a muon collider. In what follows, the notation $(1, n = 2T + 1, \epsilon)$ denotes a Dirac multiplet stabilized by additional considerations, while $(1, n = 2T + 1, 0)$ denotes a Majorana multiplet.

For even-dimensional multiplets, $Y = (n - 1)/2$ ensures the lightest mass eigenstate of the multiplet will be electrically neutral. Limits from direct detection already exclude all minimal cases with $Y \neq 0$, but even-dimensional multiplets may be rendered viable by introducing another state that mixes with the minimal multiplet after EWSB. This generates a small Majorana mass splitting between the neutral Dirac fermion pair [113], evading bounds from direct detection. For simplicity, in this case, we assume EW loop corrections still dominate

³⁰ Integrating equation (43) over $0 < \theta_h < \pi$ gives equation (38) in the limit $m_\phi \ll E_{\text{cm}}$.

Table 4. The fermionic minimal dark matter multiplets considered in this paper (and [34]), the mass target set by thermal abundance, and a brief summary of the 5σ discovery reach at a 30 TeV muon collider under the optimistic luminosity scaling in the four individual search channels described in the text. The 5σ discovery reach for muon colliders at $\sqrt{s} = 3, 6, 10, 14, 30, 100$ TeV under both conservative and optimistic luminosity assumptions is provided in the summary plots in figure 15. More details can be found in [34].

		5 σ discovery coverage (TeV)			
Model (color, n, Y)	Thermal target (TeV)	mono- γ	mono- μ	di- μ 's	disp. tracks
(1, 2, 1/2)	Dirac	1.1	—	2.8	—
(1, 3, 0)	Majorana	2.8	—	3.7	—
(1, 3, ϵ)	Dirac	2.0	0.9	4.6	—
(1, 5, 0)	Majorana	14	3.1	7.0	3.1
(1, 5, ϵ)	Dirac	6.6	6.9	7.8	4.2
(1, 7, 0)	Majorana	23	14	8.6	6.1
(1, 7, ϵ)	Dirac	16	13	9.2	7.4

the mass splitting between the neutral and charged members of the multiplet. This assumption is compatible with a relatively high scale of additional physics. For instance, when the additional state of mass M is heavy enough to be integrated out, the leading mass splitting is typically due to a dimension-5 operator which generates $\Delta m \propto v^2/M$. Requiring this splitting to be large enough to evade direct detection bounds but small enough to avoid altering the mass ordering from electroweak loop corrections implies $M \sim (10\text{--}1000)$ TeV. While this additional physics may itself be probed at a muon collider, this is more model-dependent and will not be further studied here. In what follows, we will focus on the case of the electroweak doublet ($n = 2$) as representative of the class of even-dimensional multiplets, with the above assumptions to reconcile limits from direct detection.

Of course, minimal dark matter candidates may also arise from real or complex scalar multiplets carrying electroweak quantum numbers. Scalars admit more renormalizable couplings to the standard model, most notably through Higgs portal operators of the schematic form $\chi\chi^\dagger HH^\dagger$. These couplings can induce significant tree-level mass splittings after electroweak symmetry is broken, introducing a high degree of model dependence. We leave this case for future study.

As noted above, the leading interactions between a minimal dark matter multiplet and the standard model are strictly controlled by the multiplet's EW quantum numbers. These interactions control the thermal relic abundance of cold dark matter resulting from freeze-out. Assuming that this is the sole source of the dark matter's abundance, matching current observations [119] determines the dark matter's ‘thermal target’ mass. The various fermionic dark matter candidates and their corresponding thermal target masses are enumerated in table 4. It bears emphasizing that the perturbative calculation of the thermal target mass is subject to large corrections from both Sommerfeld enhancement [120–122] and bound state effects [123, 124]. For the purposes of this discussion, we mainly use the thermal targets presented in [38], which themselves are primarily obtained by including Sommerfeld corrections to results in [118]. A notable exception is

the quintuplet Majorana fermion, for which bound state effects are significant; these lift the thermal target from 9 TeV to 14 TeV [124]. This is in contrast to the triplet Majorana fermion, for which bound state effects do not shift the thermal target relative to the Sommerfeld calculation [124]. For the septuplet Majorana fermion, we obtain an approximate target by using the fact that the degrees of freedom decrease by a factor of two relative to the Dirac fermion, pushing the thermal target higher by a factor of $\sqrt{2}$ relative to the Dirac case [118]. Needless to say, all thermal targets quoted here are subject to residual theoretical uncertainties. Experimental coverage of these targets is a compelling goal for a future collider program.

High-energy muon colliders are exceptionally well-positioned in this regard. There are a number of promising channels in which to search for minimal dark matter, including mono-photon, mono-muon, and VBF di-muon final states with an inclusive missing mass signature [34]. Alternately, the production of charged particles in the multiplet followed by decay into dark matter (and soft tracks) gives rise to a promising disappearing track signature, where the small mass splitting due to EW corrections translates into a macroscopic distance traversed by the charged particles before they decay. This channel's performance is subject to considerable uncertainties owing to the currently-unknown BIBs at a muon collider, but may significantly enhance the reach [35]. Here we summarize the performance of each channel, following [34].

- **Mono-muon:** in this channel, a charged particle of the EW multiplet is produced in association with a neutral one, leading to a single muon in the final state. This is a unique signal for a muon collider, with considerable discovery potential on account of a high signal-to-background ratio. Signal production is predominantly from the VBF processes $ZW, \gamma W \rightarrow \chi\chi$, although the process $\gamma Z \rightarrow \chi\chi$ also contributes to this channel when one of the initial state muons escapes detection following collinear emission of a γ or Z . Given the central role of VBF production for the signal, this channel is especially promising

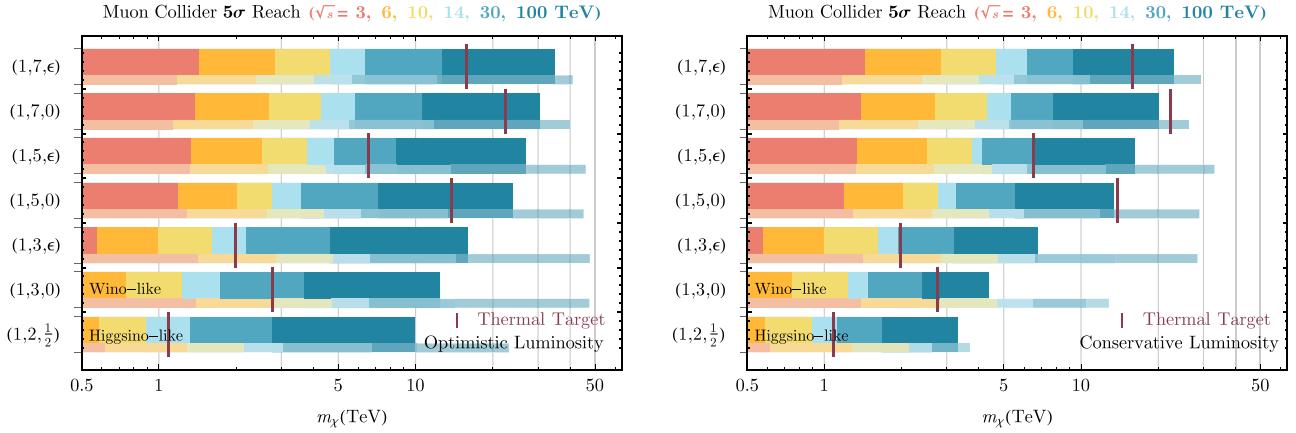


Figure 15. Summary of the 5σ discovery reach for minimal dark matter of muon colliders operating at various center of mass energies with the optimistic luminosity scaling (left) and conservative luminosity scaling (right). The wider bars indicate the combined reach from missing mass searches in the mono-photon, mono-muon, and VBF di-muon channels. The thinner bars indicate the estimated reach from the mono-photon plus one disappearing track search. The maroon vertical bars denote the thermal target for a given minimal DM multiplet. More details, including the detailed reaches for each channel and collider energies, can be found in [34]. We also note a recent dedicated study of disappearing tracks which projects a 10 TeV muon collider can cover 2.75 TeV (4.55 TeV) higgsinos (winos) [35].

for $m_\chi \ll \sqrt{s}/2$ and lower-dimensional EW multiplets, where the reach can exceed the mono-photon channel for $n \leq 3$. Higher-dimensional representations enjoy correspondingly higher rates and reach, but this channel cannot quite cover the (much higher) thermal target in these cases. The signal decreases with the dark matter mass as $1/m_\chi^4$, so that the reach in this channel does not extend all the way to the kinematic limit set by $\sqrt{s}/2$. The reach in mass scales approximately linearly with energy under our optimistic luminosity scaling, $\mathcal{L} \propto s$.

- **Mono-photon:** in this channel, the particles in the EW multiplet are pair produced in association with a photon from either initial or final state radiation. This channel is particularly effective for higher-dimensional multiplets due to the corresponding coupling enhancement, which scales as n^2 , and the reach exceeds that of the mono-muon channel for $n \geq 5$. It again does not quite reach the kinematic limit of $\sim \sqrt{s}/2$, despite coming close for $n = 7$. The primary challenge for the reach in this channel is the sizable irreducible mono-photon background, leading to a signal-to-background ratio on the order of $S/B < 10^{-2}$.

- **Di-muon:** in this channel, the particles in the EW multiple are pair-produced in neutral VBF from fusion of ZZ , $Z\gamma$, and $\gamma\gamma$, where the two final state muons are tagged. The acceptance of this channel is sensitive to the assumed muon angular acceptance; here, we have taken the muon acceptance to extend out to $|\eta_\mu| < 2.5$, leading to a significant reduction of the signal rate on account of the low likelihood of tagging two muons in the forward and backward regions. Even so, for higher-dimensional EW multiplets such as $n = 7$, this channel provides coverage competitive with the mono-photon channel while being more robust against systematics. Improving the reach in this channel motivates advanced detector design that would cover more of the forward regime, for instance,

covering $2.5 < |\eta_\mu| < 4.0$ (or potentially even out to $|\eta_\mu| < 8$ [72]).

- **Disappearing tracks:** in this channel, the charge ± 1 particle in the EW multiplet is pair produced and decays into the dark matter plus soft particles with a long lifetime due to the small radiative mass splitting. For the cases considered here, $c\tau$ ranges from 0.37 cm to 5.6 cm. If the charged particle hits several layers of the tracker before decaying, this results in the unique signature of a disappearing track, a potentially low-background process. However, making accurate projections for the reach of this channel at a muon collider is hampered by our current ignorance of tracker design and BIBs. Here we present an estimate based on the combination of a singlet displaced track plus another tagging object such as a photon (with the expectation that this will be required to suppress backgrounds), requiring tens of signal events for discovery. Requiring two displaced tracks would necessarily provide further background suppression, albeit with a significant loss of rate; this channel would be useful to study further in the event that backgrounds for the single-track channel prove to be prohibitive. Focusing on the single-track final state with an additional tagging object, the mono-photon channel with one disappearing track will have the largest signal rate, significantly extending the reach for all odd-dimensional cases. However, this channel fails to reach the kinematic threshold owing to the boost required for the charged particle to leave enough hits in the tracker before decaying.³¹ The triplet enjoys the greatest increase in sensitivity from this channel, coming close to the kinematic threshold, while for the doublet, this channel is stronger than the mono-muon channel.

³¹ In this case, further sensitivity may be obtained from using timing information [125]. However, the large out-of-time contribution from beam-induced backgrounds requires more detailed studies.

The 5σ discovery reach of muon colliders operating at various center of mass energies is summarized in figure 15 for the optimistic (conservative) integrated luminosity scaling scenarios defined in table 1. The sensitivity obtained by the combination of missing mass searches in the mono-photon, mono-muon, and VBF di-muon channels is shown separately from the sensitivity of the displaced track search. When combining the missing mass channels in the optimistic luminosity scaling scenario (left panel), the overall reach does not extend to the kinematic limit $m_\chi \sim \sqrt{s}/2$ (most notably for multiplets with $n \leq 3$) due to the low signal-to-background ratio. It is possible to cover (with 2σ confidence) the thermal targets of the doublet and Dirac triplet with a 10 TeV muon collider, while a 30 TeV option would suffice for the Majorana triplet. The thermal targets of Dirac (Majorana) quintuplet would be covered by muon colliders operating at 30 (100) TeV, while a 100 TeV collider would also cover the thermal target for the septuplet.

Rather than considering the reach of the benchmark collider energies, it is also interesting to note the minimum collider energy that would cover a given multiplet, assuming integrated luminosity scales with s . From this perspective, a Majorana triplet can be reached by a 20 TeV muon collider (still assuming integrated luminosity scales with s). A Majorana quintuplet can be covered by a 50 TeV muon collider, while a septuplet can be covered by a 70 TeV muon collider. The thermal targets of all the minimal multiplets considered here could be discovered at the 5σ level by a muon collider operating at 75 TeV.

Finally, we emphasize that the disappearing track signal has excellent potential, bringing the reach close to the kinematic threshold $m_\chi \sim \sqrt{s}/2$ on the basis of the current study [34]. For instance, a 10 TeV muon collider alone could reach the thermal target of both doublet and triplet cases with a disappearing track search, motivating further studies and careful consideration for detector design.

4.3. Naturalness

The hierarchy problem is a prime motivator that new physics should be accessible at colliders, as it strongly correlates the mass scale of additional degrees of freedom with those of the standard model. Precisely what degrees of freedom appear at scales indicated by the hierarchy problem is much less definite; in recent years it has become increasingly clear that there exists a plethora of solutions with wildly varying signatures. Nonetheless, the spectrum of solutions can be usefully divided into two categories: solutions of the ‘big’ hierarchy problem, namely those reaching from the weak scale all the way to the putative scale of quantum gravity, and solutions of the ‘little’ hierarchy problem, extending from the weak scale to the highest scales directly probed (thus far) by experiments.

There are two known solutions to the ‘big’ hierarchy problem: compositeness and supersymmetry. The lack of weak scale evidence for either solution suggests the existence of a mass gap between the Higgs and whatever physics resolves the hierarchy problem. Such a mass gap implies a significant

degree of fine tuning, somewhere between the percent and per mille level depending on the details of the UV completion. While it is entirely possible that the weak scale is finely tuned—after all, we do not actually know how nature computes fine tuning—a robust commitment to naturalness could suggest the existence of additional physics that bridges the gap between the weak scale and the appearance of supersymmetry or compositeness. Resolutions to this ‘little’ hierarchy problem need only span an order of magnitude in energy in order to reconcile the paucity of new physics at the weak scale with the expectations of naturalness. In contrast to the sparsity of qualitative solutions to the big hierarchy problem, there are innumerable solutions to the little hierarchy problem consistent with current data, ranging from dynamical mechanisms that relax the Higgs mass [126] to symmetry-based mechanisms that reside in ‘dark’ hidden sectors [99, 127]. Their signatures are equally diverse, often falling outside the scope of conventional collider signals.

Solutions to the hierarchy problem reduce the UV sensitivity of the Higgs mass parameter, making it possible to understand what sets the value of the weak scale and, ultimately, why electroweak symmetry is broken in the first place. To systematically test the naturalness of the weak scale, one would ideally like to pursue two lines of experimental inquiry: leveraging precision to directly test solutions to the ‘little’ hierarchy problem at or around the weak scale, and leveraging energy to reach the scale at which states associated with resolution of the ‘big’ hierarchy problem begin to appear. The great advantage of a high-energy muon collider is that, provided sufficient energy and luminosity, it may achieve both goals. On the one hand, the relatively low background rate and clean environment make it a promising tool for discovering new light states with weak (or no) standard model quantum numbers. On the other hand, the high c.m. energy gives it the reach to discover states well above the weak scale. In what follows, we illustrate this potential by considering aspects of muon collider sensitivity to representative solutions of the ‘big’ and ‘little’ hierarchy problems.

4.3.1. The ‘big’ hierarchy problem: supersymmetry. The many superpartners predicted by supersymmetry give rise to a host of experimental signatures; see e.g. [128] for an overview. Here we will focus on the three states most closely tied to the naturalness of the Higgs potential: the higgsino, stop, and gluino, which respectively contribute to the Higgs potential at tree level, one loop, and two loops; these are the calling cards of ‘natural supersymmetry’ [129–133]. In addition, we will explore a unique opportunity available to high-energy lepton colliders: probing low-scale supersymmetry breaking sectors through a direct search for the gravitino.

4.3.1.1. Higgsinos, stops, and gluinos The specific collider reach for a given superpartner requires detailed simulation as well as knowledge of beam and detector effects that is not available at this stage. Nonetheless, we can make some meaningful estimates that should hold up to $\mathcal{O}(1)$ factors. In general, the mass reach for superpartners with electroweak quantum numbers is closely correlated with the c.m. energy of the

collider. Although the s -channel production for these states falls as $1/s$, it remains $\gtrsim \mathcal{O}(\text{few tens})$ of attobarns at $\sqrt{s} = 30$ TeV and $\gtrsim \mathcal{O}(\text{few})$ attobarns at $\sqrt{s} = 100$ TeV independent of the mass \tilde{m} of the superpartner, provided $\tilde{m} \ll \sqrt{s}/2$. As $\tilde{m} \rightarrow \sqrt{s}/2$, the mass dependence of s -channel production typically enters with the velocity scaling $\propto \beta$ for fermionic superpartners and $\propto \beta^3$ for scalar superpartners. Even in the latter case, the cross section only falls by an order of magnitude once $\tilde{m} \sim 0.9 \times \sqrt{s}/2$. As long as the integrated luminosity at a collider with c.m. energy $\sqrt{s} \gtrsim 30$ TeV is at the level of inverse attobarns, there will be enough signal events to discover superpartners up to $\tilde{m} \lesssim \sqrt{s}/2$ provided backgrounds can be eliminated while retaining high signal efficiency.

The mass reach then becomes a question of the distinctiveness of the final state. For decays with large available phase space, given the relatively low irreducible backgrounds compared to hadron colliders, the final state can be made essentially background-free with high signal efficiency. Consider, for example, the decay $\tilde{t} \rightarrow t + \chi_1^0$ with $m_{\tilde{t}} \gg m_t + m_{\chi}$. The cross section for s -channel $t\bar{t}$ production is always within an order of magnitude of the $\tilde{t}\tilde{t}^*$ cross section until velocity suppression becomes significant. The cross section for the VBF mode $\mu^+ \mu^- \rightarrow t\bar{t} + \nu\bar{\nu}$ does not meaningfully exceed the s -channel $t\bar{t}$ cross section until $\sqrt{s} \sim 5$ TeV, at which point it grows relative to the s -channel cross section as $\sim s \log^3(s)$, exceeding the signal cross section by 3–4 orders of magnitude by $\sqrt{s} = 100$ TeV. Even at such high energies, this is far less daunting than the $t\bar{t}$ background to the same final state at the LHC, and sufficiently strong missing energy requirements and cuts on the distribution of visible particles in the final state should substantially reduce background with high signal efficiency. Thus, to first approximation, for distinctive final states we might assume the mass reach will be of order $\tilde{m} \sim 0.9 \times \sqrt{s}/2$ up to $\sqrt{s} = 100$ TeV.

Not all supersymmetric final states are so distinctive, however. For the higgsino, if standard model radiative corrections are the only source of mass splitting, then there is little phase space for missing energy without additional initial-state radiation. The mass reach then becomes more sensitive to backgrounds. Of course, if there is additional splitting in the higgsino multiplet, e.g., due to mixing with a partially decoupled bino or wino, then the final state rapidly becomes more distinctive and the estimate of the reach follows the same logic as above.

With this in mind, we estimate the reach for various superpartners, beginning with the higgsino. The mass of the higgsino is the most immediate measure of fine tuning in the Higgs potential, since supersymmetry relates the masses of the higgsino and Higgs doublets at tree level. In general the contribution of the higgsino to the tuning of the weak scale³² is

³² Quantified here by the Barbieri–Giudice measure [134] $\Delta_{M^2} = |\partial \ln m_h^2 / \partial \log M^2|$ in terms of m_h^2 . As always, we emphasize that individual fine-tuning measures should be taken with a grain of salt given our ignorance of nature’s prescription, and is only used here to provide a qualitative guide.

parametrically of the form

$$\Delta_{\tilde{h}} \simeq \frac{2m_{\tilde{h}}^2}{m_h^2}. \quad (44)$$

From this we can conclude that 10% tuning ($\Delta = 10$) corresponds to $m_{\tilde{h}} \sim 300$ GeV, while the percent and per mille levels are reached by $m_{\tilde{h}} \sim 900$ GeV and $m_{\tilde{h}} \sim 2.8$ TeV, respectively.

For a pure higgsino multiplet with only standard model radiative splittings, the final state is indistinct and a detailed study is required to forecast reach as a function of \sqrt{s} . Here we may rely on the dark matter analysis presented in section 4.2 for a pair of $SU(2)_L$ doublets with hypercharge $\pm 1/2$, which details the potential reach under both optimistic and conservative luminosity assumptions. Under the optimistic (conservative) luminosity assumptions, this analysis suggests that percent-level tuning can be probed at the 2σ level by a muon collider operating at $\sqrt{s} = 6$ TeV (10 TeV), while per mille-level tuning is accessible by a $\sqrt{s} = 30$ TeV (100 TeV) collider. The ultimate exclusion limit of a $\sqrt{s} = 100$ TeV collider under optimistic luminosity scaling approaches $m_{\tilde{h}} \sim 15$ TeV, corresponding to tuning at more than the *per myriad* level. But this is ultimately the worst-case scenario; even a modest splitting in the higgsino multiplet would make it possible to significantly suppress backgrounds beyond those considered here, and once sufficient phase space becomes available for on-shell decays within the higgsino multiplet, bounds should approach the $m_{\tilde{h}} \sim 0.9\sqrt{s}/2$ reach corresponding to distinctive final states. In this limit, percent tuning becomes accessible at $\sqrt{s} \gtrsim 2$ TeV, per mille tuning at $\sqrt{s} \gtrsim 6$ TeV, and per myriad tuning at $\sqrt{s} \gtrsim 20$ TeV.

Now we turn to the stop, whose contribution to the tuning of the weak scale at leading-logarithmic order is

$$\Delta_{\tilde{t}} \simeq \frac{3y_t^2}{4\pi^2} \frac{m_{\tilde{t}}^2}{m_h^2} \log \frac{\Lambda}{m_{\tilde{t}}}, \quad (45)$$

where Λ is a UV scale at which SUSY breaking is communicated; here we conservatively take $\Lambda = 10 \times m_{\tilde{t}}$. Here 10% tuning corresponds to $m_{\tilde{t}} \sim 950$ GeV, percent tuning to $m_{\tilde{t}} \sim 3$ TeV, and per mille tuning to $m_{\tilde{t}} \sim 9.5$ TeV.

Barring a high degree of degeneracy between the lightest stop and its decay products, the final state is quite distinctive and we can anticipate reach scaling as $m_{\tilde{t}} \sim 0.9\sqrt{s}/2$. To validate this expectation, we perform a parton-level analysis for two simplified stop-top-neutralino models: one in which the stop is an $SU(2)_L$ singlet denoted \tilde{t}_R , the other in which it is part of an $SU(2)_L$ doublet denoted \tilde{t}_L . In both models the neutralino χ is taken to be bino-like. To ensure the relevant production processes are not contaminated by contributions from heavy superpartners, we construct both models in FEYN-RULES using the t -channel dark matter framework [135] with couplings set to their supersymmetric values. We then simulate pair production of the stop in both muon annihilation and VBF in WHIZARD, followed by the decay of the stop to a top quark and neutralino. We generate samples for a range of stop masses up to $\sqrt{s}/2$, keeping the neutralino effectively massless with $m_{\chi} = 1$ GeV. We consider the leading background to be VBF

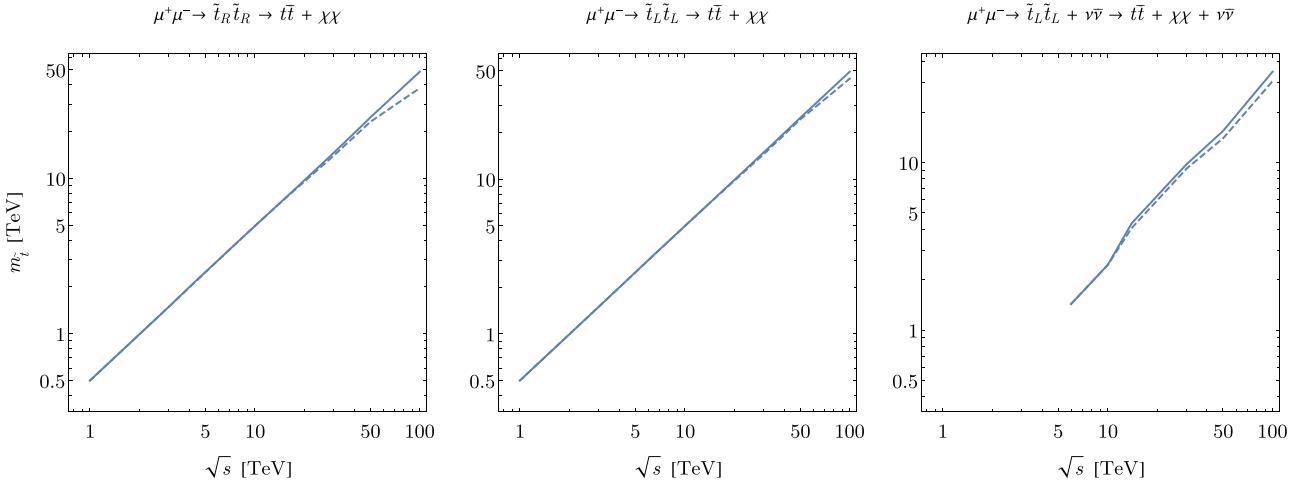


Figure 16. Exclusion reach on a stop-top-neutralino simplified model as a function of \sqrt{s} and m_t with $m_\chi = 1$ GeV, assuming ‘optimistic’ (solid) and ‘conservative’ (dashed) integrated luminosity scaling as detailed in the body of the text. Left: limit on pair production of an RH stop in muon annihilation. Center: limit on pair production of an LH stop in muon annihilation. Right: limit on pair production of an LH stop in VBF.

$t\bar{t}$ production, for which the W^+W^- fusion contribution yields a $t\bar{t} +$ missing energy final state closely mimicking the signal. For simplicity, we consider perfect reconstruction of the top quarks in the final state. We then separately consider two possible cuts to separate signal from background: (1) a cut on the M_C variable [136] constructed from the four-momenta of the t and \bar{t} ,

$$M_C = \sqrt{2(E_1 E_2 + \vec{p}_1 \cdot \vec{p}_2 + m_t^2)}, \quad (46)$$

for which the signal has an endpoint at $M_C^{\max} = m_t$ assuming a massless neutralino, and (2) a cut on the missing transverse momentum p_T . For each signal benchmark, we determine the M_C and p_T cuts that maximize S/\sqrt{B} , translating it into the exclusion reach shown in figure 16 corresponding to $S/\sqrt{B} = 2$ under either the optimistic or the conservative luminosity scaling. For the right-handed stop, the cross section for VBF pair production is generally too small to be a meaningful contribution to the total signal rate, while for the left-handed stop both muon annihilation and VBF production play an interesting role. In each case, the p_T cut alone is sufficient to provide sensitivity all the way up to $m_t \sim \sqrt{s}/2$, save for the largest values of \sqrt{s} where some degradation of the limit is observed under conservative luminosity assumptions due to the significant size of the $t\bar{t} + \nu\bar{\nu}$ background at these energies. Of course, this highly simplified analysis neglects a host of relevant effects associated with the decay and reconstruction of top quarks, energy resolution, forward reconstruction, and BIBs, but it supports the supposition that significant kinematic separation of signal and background is possible, and similar performance is likely achievable with an optimized selection using all kinematic information.

As such, we consider the stop mass reach of muon colliders operating at various values of \sqrt{s} assuming that the reach scales as $m_t \sim 0.9 \times \sqrt{s}/2$. As a starting point, the anticipated

2σ exclusion reach at the HL-LHC is $m_t \sim 1.7$ TeV, which suggests that a muon collider operating at $\sqrt{s} \gtrsim 4$ TeV is required to compete with LHC limits. Reaching percent-level tuning requires $\sqrt{s} \gtrsim 6$ TeV, while per mille requires $\sqrt{s} \gtrsim 20$ TeV and per myriad $\sqrt{s} \gtrsim 66$ TeV. Ultimately, the fine-tuning reach as a function of \sqrt{s} is comparable to the *worst-case* higgsino scenario with only standard model radiative splittings, the one loop suppression of the stop contribution to fine tuning competing with the more challenging backgrounds faced by the higgsino.

Apart from considerations of fine tuning, the stop mass is also central to the supersymmetric prediction of the observed Higgs mass. Famously, accommodating $m_h \sim 125$ GeV in the minimal supersymmetric extension of the standard model without significant mixing between the stop gauge eigenstates suggests that the stops lie around or above $\sim 5-10$ TeV; see, e.g., [137] for a review. A high-energy muon collider operating at $\sqrt{s} = 30$ TeV could cover the typical scale of the stop mass suggested by the observed Higgs mass at large values of $\tan\beta$, subject to further dependence on mixing angles and the remaining sparticle spectrum. At moderate values of $\tan\beta$, the stops can be heavier, $\sim 100 - 1000$ TeV in well-motivated models. However, in such scenarios the electroweakinos may be accessible (e.g., via the searches discussed in section 4.2), and the lightest sfermions, such as the right-handed stau, may be an order of magnitude lighter than the stops and could be directly accessible. Thus, the full range of searches for superparticles could cover a substantial portion of the parameter space motivated by the measured Higgs mass.

Finally, we turn to the gluino, which evades our reach estimates because it does not carry electroweak quantum numbers. Nonetheless, it may still be produced by a variety of higher-order processes, including gluon splitting from $q\bar{q}$ production. The gluino contribution to the tuning of the weak scale at

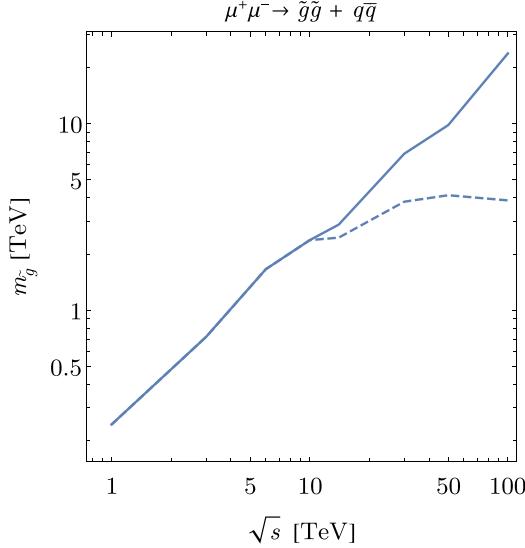


Figure 17. Gluino discovery reach from $\mu^+\mu^- \rightarrow \tilde{g}\tilde{g} + q\bar{q}$ as a function of \sqrt{s} , assuming ‘optimistic’ (solid) and ‘conservative’ (dashed) integrated luminosity scaling as detailed in the body of the text.

leading logarithmic order scales as

$$\Delta_{\tilde{g}} \simeq \frac{\alpha_s y_t^2}{\pi^3} \frac{m_{\tilde{g}}^2}{m_h^2} \log^2 \frac{\Lambda}{m_{\tilde{g}}}. \quad (47)$$

Again taking $\Lambda = 10 \times m_{\tilde{g}}$, 10% tuning corresponds to $m_{\tilde{g}} \sim 3$ TeV, percent tuning to $m_{\tilde{g}} \sim 9.5$ TeV, and per mille tuning to $m_{\tilde{g}} \sim 30$ TeV.

To estimate the production rate, we consider the process $\mu^+\mu^- \rightarrow \tilde{g}\tilde{g} + q\bar{q}$ with modest phase space cuts, as detailed in the appendix. Although the final state is quite distinctive—presumably zero background is achievable with high signal efficiency—the ultimate reach of a muon collider is quite sensitive to the integrated luminosity due to the relatively small cross section. In figure 17, we plot the discovery reach (assuming zero background, i.e., 5 signal events) for both optimistic and conservative luminosity assumptions. The optimistic luminosity assumption leads to a mass reach scaling with collider energy, while the conservative luminosity assumption leads to a relatively constant reach for $\sqrt{s} \gtrsim 10$ TeV.

As a starting point, the anticipated 2σ exclusion reach at the HL-LHC is $m_{\tilde{g}} \sim 3.2$ TeV, which suggests that the discovery reach of a muon collider could just reach the edge of LHC limits for $\sqrt{s} \gtrsim 10$ TeV and 10 ab^{-1} . The optimistic luminosity scaling paints a rosier picture, providing sensitivity to $m_{\tilde{g}} \sim 20$ TeV by $\sqrt{s} = 100$ TeV. Ultimately, the gluino reach lags behind the fine-tuning sensitivity of higgsino and stop reach, but nonetheless a muon collider is comparable to or better than the LHC depending on the c.m. energy and integrated luminosity.

Taken together, these projections suggest that a high-energy muon collider operating at $\sqrt{s} \sim 20 - 30$ TeV is capable of probing a natural supersymmetric explanation of the weak scale beyond the per mille level, with meaningful sensitivity to both of the states most important for understanding the scale

and origin of EWSB: the higgsinos and stops. Setting aside fine tuning, which may or may not be a sharp guide to the scale of new physics, a collider operating at $\sqrt{s} \sim 30$ TeV could reach the scale suggested by the Higgs mass prediction of the minimal supersymmetric standard model (MSSM), regardless of the stop sector mixing.

4.3.1.2. The gravitino. Although impressive in reach, the potential for a high-energy muon collider to probe standard model superpartners represents a continuation of the already-vast search program currently under way at the LHC. But a muon collider offers more than incremental progress in the search for supersymmetry; it would be perhaps the first collider with the potential to directly discover supersymmetry through its universal feature, the goldstone fermion of spontaneous supersymmetry breaking.

A universal prediction of spontaneous breaking of supersymmetry in the rigid limit is the presence of a massless Majorana fermion, the goldstino. When gravity effects are taken into account, the goldstino is eaten through super-Higgs mechanism by the spin 3/2 gravitino [138], which obtains a mass

$$m_{3/2} = \frac{F}{\sqrt{3}M_{\text{Pl}}}, \quad (48)$$

where \sqrt{F} is the scale of SUSY breaking and $M_{\text{Pl}} = 2.4 \times 10^{18}$ GeV. The scale of the superpartners is set schematically by the supersymmetry-breaking contributions

$$m_{\text{soft}} \sim \frac{g_{\text{eff}}^2}{16\pi^2} \sqrt{F} + m_{3/2}. \quad (49)$$

The first term in (49) is the gauge mediation contribution, with g_{eff} an effective coupling encoding suppressions or enhancements in the soft masses which are model dependent. For instance for the gaugino masses one typically has $g_{\text{eff}}^2 \sim g_{\text{SM}}^2 s_g N_{\text{mess}} \frac{\sqrt{F}}{M_{\text{mess}}}$, where $s_g \lesssim 1$ parameterizes possible gaugino screening effects. The second term in (49) is the gravity mediation contribution. We have assumed that a single supersymmetry-breaking scale controls both. Low-energy supersymmetry breaking models are those in which the gauge mediation contribution dominates, such that the gravitino is the lightest supersymmetric particle (LSP) in the spectrum and hence cosmologically stable unless R -parity is broken.

In low-energy supersymmetry breaking scenarios, the stable gravitino is dominantly produced in the early Universe through gluon-gluino scattering processes as computed in [139–142]. The gravitino yield can be written as

$$Y_{3/2}^{\text{UV}} = C_{\text{UV}} \frac{m_{\tilde{g}}^2 T_{\text{r.h.}}}{m_{3/2}^2 M_{\text{Pl}}}, \quad \text{with} \\ C_{\text{UV}} = \frac{45\sqrt{5}f_3}{8\pi^{13/2}g_*^{3/2}} \simeq 4 \times 10^{-5}, \quad (50)$$

where $m_{\tilde{g}}$ is the gluino soft mass and $T_{\text{r.h.}}$ is the reheating temperature. Requiring that the gravitino yield does not overclose the Universe typically results in strong constraints

on the maximal reheating temperature in these scenarios [143–146].

An interesting allowed region is the so called *ultralight gravitino window* where $T_{\text{r.h.}}$ is required to be high enough to make the gravitino thermalize in the early Universe. Requiring $Y_{3/2}^{\text{UV}} > Y_{\text{eq}}$ with $Y_{\text{eq}} = 1.8 \times 10^{-3}$ gives

$$T_{\text{r.h.}} > \frac{45m_{3/2}^2 M_{\text{Pl}}}{m_{\tilde{g}}^2} \simeq 9 \text{ MeV} \left(\frac{m_{3/2}}{16 \text{ eV}} \right)^2 \left(\frac{2 \text{ TeV}}{m_{\tilde{g}}} \right)^2. \quad (51)$$

Since the gravitino freezes out when it is still relativistic, the matter power spectrum is going to be damped at small scales [147]. In [148] a combination of CMB data from WMAP and Lyman- α forest data was used to set an upper bound on the gravitino mass (and the SUSY-breaking scale)

$$m_{3/2} < 16 \text{ eV} \iff \sqrt{F} < 260 \text{ TeV}. \quad (52)$$

This bound will presumably be improved with current Planck data and even further with future cosmological surveys.

Given the upper bound on the SUSY-breaking scale, the gluino can be made to lie above the 2 TeV reach of the LHC provided

$$g_{\text{eff}}^2 > 1.2 \times \left(\frac{m_{\tilde{g}}}{2 \text{ TeV}} \right) \left(\frac{260 \text{ TeV}}{\sqrt{F}} \right), \quad (53)$$

following the parametrics in equation (49). We refer to [149, 150] for explicit models which realize a heavy superpartner spectrum with such a low SUSY-breaking scale. If we are agnostic about UV completions, the perturbativity of g_{eff} gives a lower bound on \sqrt{F} which is a loop factor below 260 TeV. The gravitino window can then be defined as

$$10^{-3} \text{ eV} < m_{3/2} < 16 \text{ eV} \iff 2.5 \text{ TeV} < \sqrt{F} < 260 \text{ TeV}, \quad (54)$$

where the lower bound is given by the current LHC bound on the gluinos plus perturbativity of the SM couplings, while the upper bound is from cosmological constraints on warm dark matter.

Within the light gravitino window defined in equation (54), we can safely assume that the superpartner spectrum is decoupled and the gravitino interactions with the SM are described by a universal EFT with couplings controlled by the supersymmetry breaking scale only [151–153]. Collider searches for direct gravitino production can hence provide a robust limit on the supersymmetry breaking scale, independent of the specific details of the superpartner spectrum. Traditionally, this has not been an emphasis of the supersymmetry search program at hadron colliders because direct searches for colored superpartners always exceed the sensitivity of direct gravitino searches. In contrast, at high-energy lepton colliders the reduced background and rapid growth of signal cross sections with \sqrt{s} makes this a competitive channel for the discovery of supersymmetry.

Here, we explore the sensitivity of a muon collider to gravitino pair production in the mono-photon final state via the production mode $\mu^+ \mu^- \rightarrow \tilde{G}\tilde{G}\gamma$. The analytic cross section

for this process in the gravitino EFT has been computed in [152]; it includes ISR photon topologies from contact interactions between leptons and gravitinos, as well as contact interactions involving leptons, gravitinos and a photon. In the limit of soft and collinear photon emission the cross section reads

$$\sigma(\mu^+ \mu^- \rightarrow \tilde{G}\tilde{G}\gamma) = \frac{\alpha s^3}{160\pi^2 F^4} \left[\frac{247}{60} + 2 \log \left(\frac{2E_{\text{min}}^\gamma}{\sqrt{s}} \right) \right] \times \log \left(\frac{1 - \cos \theta_{\text{min}}}{1 + \cos \theta_{\text{min}}} \right), \quad (55)$$

where \sqrt{s} is the c.m. energy, E_{min}^γ is the minimal energy of the photon, and θ_{min} is the minimal photon angle with respect to the beam direction. In the following numerical analysis we will employ the complete analytic formula, but we can use the above equation to get some reasonable estimates for the reach in terms of \sqrt{F} .

We assume minimal cuts on the photon ($E_{\text{min}}^\gamma > 50 \text{ GeV}$, $|\eta_\gamma| < 2.4$). The SM background at high \sqrt{s} is dominated by W -boson t -channel diagrams and by WW fusion (collectively denoted WW in the figure) as can be seen in figure 18, and it is roughly constant and equal to 2 pb for $\sqrt{s} \gtrsim 3 \text{ TeV}$. We further define the signal region by requiring the photon energy not to exceed the endpoint of the signal process, $E_{\text{max}}^\gamma = \sqrt{s}/2$. With these cuts, the 2σ sensitivity to \sqrt{F} using equation (55) is

$$\sqrt{F} \lesssim 61.7 \text{ TeV} \left(\frac{\mathcal{L}}{1000 \text{ ab}^{-1}} \right)^{1/16} \left(\frac{\sqrt{s}}{100 \text{ TeV}} \right)^{3/4} \times \left(4.8 + \log \left[\frac{\sqrt{s}}{100 \text{ TeV}} \right] \right)^{1/8}, \quad (56)$$

where we have only included statistical errors. The corresponding numerical result is displayed in figure 19. In the same figure, we show the decrease in the reach assuming a 1% systematic error on both signal and background. Given that the signal cross section grows as the 6th power of the c.m. energy, the energy increase is the most beneficial in increasing the sensitivity to high values of \sqrt{F} .

We conclude that a future high energy muon collider can almost certainly push up the lower bound of the ultralight gravitino window by one order of magnitude. An improvement of the cosmological bounds with respect to the ones derived from WMAP data could allow the light gravitino window to be completely closed in the future. More broadly, this exemplifies the ability of a high-energy muon collider to directly probe the mechanism of supersymmetry breaking.

4.3.2. The ‘big’ hierarchy problem: compositeness. We next turn our attention to the competing solution to the ‘big’ hierarchy problem, compositeness. Though compositeness may take many guises, its realizations all share a set of common features given the discovery of a light standard model-like Higgs with an apparent mass gap. The expectation is that fermionic top partners should exist at a scale $m_T \sim y_t f$ parametrically below the scale of compositeness $m_* = g_* f$ at which

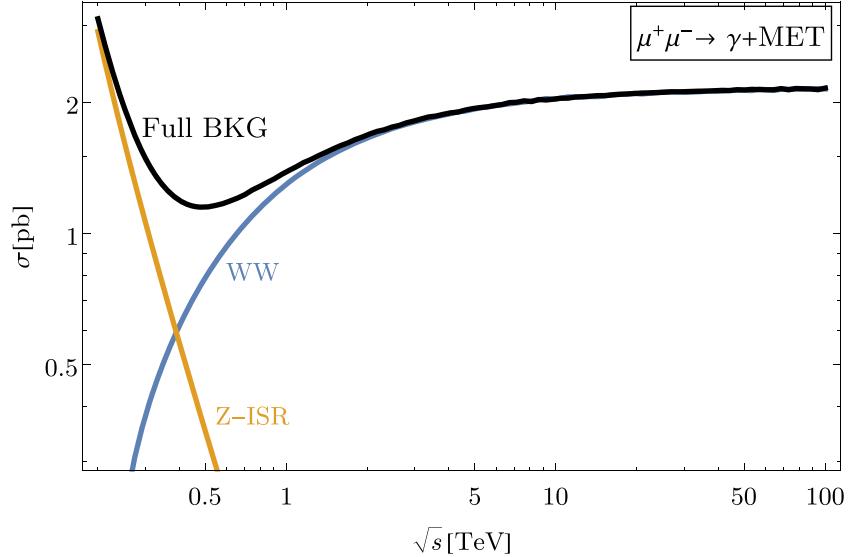


Figure 18. Background cross section for the $\gamma + \text{MET}$ final state as a function of the c.m. energy \sqrt{s} . The black line is the total background cross section in pb while the red and blue lines show the separate contributions of respectively $Z\gamma$ and WW (this denotes the combination of WW fusion and W -boson t -channel diagrams).

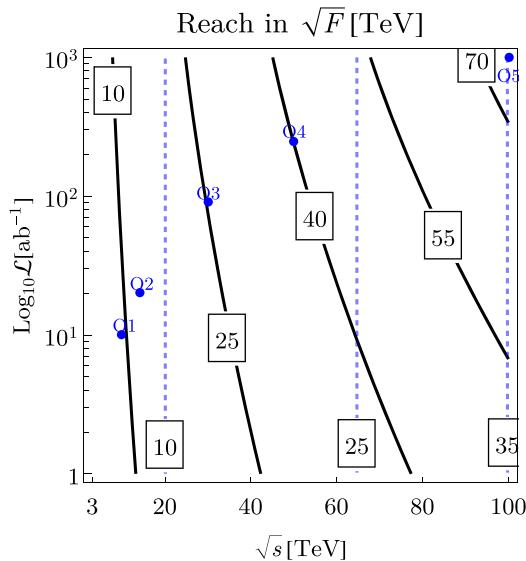


Figure 19. Estimated reach of the gravitino search interpreted as a constraint on the supersymmetry breaking scale \sqrt{F} . The solid black lines denote limit contours assuming only statistical errors: $S/\sqrt{S+B} = 1.96$. The dashed blue lines denote limit contours that include a systematic error: $S/\sqrt{S+B + \epsilon^2(S^2+B^2)} = 1.96$ with $\epsilon = 1\%$. The points O_1, \dots, O_5 indicate the optimistic benchmarks of table 1, starting from $\sqrt{s} = 10$ TeV and with increasing c.m. energies.

a host of resonances appear. See e.g. [154] for an excellent overview.

Perhaps the strongest tests of compositeness come from its distinctive imprints on the electroweak sector of the SM, including modification of Higgs couplings [155]. To the extent that the tuning associated with vacuum misalignment is generally on the order of $\Delta \sim f^2/v^2$ and the leading deviations from standard model predictions can be captured by irrelevant operators suppressed by f^2 or $(g_*/g)^2 f^2$, tests of such

deviations typically provide the most powerful constraints. Here, high-energy muon colliders enjoy a particular advantage, leveraging the energy growth of these deviations without the considerable backgrounds of pp counterparts. As demonstrated in [39], this allows a high-energy muon collider to access Higgs compositeness scales well above any other future collider project. For example, sensitivity to operators such as

$$\mathcal{O}_H = \frac{1}{2} (\partial_\mu |H|^2)^2 \quad \text{and} \\ \times \mathcal{O}_W = \frac{ig}{2} \left(H^\dagger \sigma^a \overset{\leftrightarrow}{D}^\mu H \right) D^\nu W_{\mu\nu}^a, \quad (57)$$

allows a muon collider operating at $\sqrt{s} = 10$ TeV to probe compositeness scales as high as $m_* \sim 45$ TeV with 10 ab^{-1} , while a collider operating at $\sqrt{s} = 30$ TeV can probe compositeness up to $m_* \sim 140$ TeV with 90 ab^{-1} [39]. The sensitivity at these two energies covers fine tuning associated with vacuum misalignment from the percent level to the per mille level.

To the extent that it is possible to decrease the tuning associated with vacuum alignment, sensitivity to fermionic top partners provides a complementary probe. Here the test of fine tuning is expected to be qualitatively similar to that of supersymmetric tops, with two relative advantages: for fermions the s -channel production cross section is only linearly suppressed by velocity near threshold, and in the absence of additional structure they typically decay directly to visible and highly distinct final states. Following the fine-tuning estimates in the supersymmetric case, reaching percent-level tuning requires $\sqrt{s} \gtrsim 6$ TeV, while per mille requires $\sqrt{s} \gtrsim 20$ TeV and per myriad $\sqrt{s} \gtrsim 66$ TeV. As in the case of supersymmetry, a high-energy muon collider operating at $\sqrt{s} \sim 20 - 30$ TeV would decisively test compositeness as a natural explanation for the weak scale.

4.3.3. The ‘little’ hierarchy problem. As we have seen, a high-energy muon collider operating at tens of TeV could provide satisfying coverage of known solutions to the ‘big’ hierarchy problem. It is no less suited to constraining or discovering solutions to the ‘little’ hierarchy problem, whose subtle signatures will remain largely untouched by the LHC. There are a plethora of such solutions which operate by relying on novel field theoretical [99, 100, 156–165] or cosmological [126, 166–169] ingredients. Here we will focus on solutions involving ‘neutral naturalness’ such as the twin Higgs [99] or hyperbolic Higgs [163], in which the partner particles are entirely neutral under the standard model. Such models feature (at least) four possible avenues to discovery at colliders: Higgs coupling deviations from mixing between scalars; direct production of SM singlet partner particles; displaced vertices from exotic Higgs decays; and direct production of the ‘radial mode’ associated with spontaneous breaking of the discrete symmetry.

In general, the Higgs coupling deviations are a subset of those expected in composite models, including most notably the oblique operator \mathcal{O}_H defined in equation (57). As with composite models, constraints on \mathcal{O}_H translate directly into bounds on f^2/v^2 , a typical measure of the fine tuning. The results of [39] suggest sensitivity to $f \sim 16v$ at $\sqrt{s} = 10$ TeV and $f \sim 46v$ at $\sqrt{s} = 30$ TeV, probing the tuning of the weak scale between the percent and per mille levels within the framework of the twin Higgs. The prospects for discovery via direct production of partner particles are considerably weaker. Although the precise reach depends on whether these partner particles are fermions (as in the twin Higgs) or bosons (as in the hyperbolic Higgs), the qualitative sensitivity can be inferred from the reach for the ‘nightmare scenario’ of a \mathbb{Z}_2 -symmetric singlet scalar summarized in section 4.1 and appendix A, extending at most to the hundreds of GeV.

A more exotic collider signature of models of neutral naturalness is the prediction of displaced decays of dark particles into standard model states [127]. The scalar mixing produces a non-zero branching ratio of Higgses into neutral partner states, which gives a portal for energy to be transferred to the partner sector at a collider. Some of the produced partner sector states are unstable to decaying through an off-shell Higgs into standard model states, bearing out ‘hidden valley’ type phenomenology [170, 171]. This produces spectacular signatures—for example vertices displaced from the beamline from which standard model jets appear—and this allows search strategies with very low background [125, 172]. The prospects for probing these signatures in neutral naturalness models at future lepton colliders have been studied in the context of Higgs factories [173, 174], and a branching ratio reach is projected which is competitive with or better than the LHC forecasts. Dedicated study for a higher-energy muon collider has yet to be performed.

The prospects for discovering the ‘radial mode’ are much stronger than for other partner particles. These can be obtained by interpreting the model independent results of figure 14 in the twin Higgs, where the mixing angle scales as

$\sin \gamma \sim g_* v / m_\phi$. Concretely, the Higgs potential in the simplest realization of the twin Higgs can be written as

$$V = \lambda_* \left(|H_A|^2 + |H_B|^2 - \frac{f_0^2}{2} \right)^2 + \kappa (|H_A|^4 + |H_B|^4) + \sigma_{\text{soft}} f^2 |H_A|^2. \quad (58)$$

The SM Higgs sector is extended by the addition of the twin Higgs H_B , a singlet under the SM and doublet under a mirror electroweak gauge group $SU(2)_B$. The twin Higgs is coupled to the SM Higgs H_A via a portal coupling λ_* . With only this quartic included, the potential linearly realizes an $SO(8)$ symmetry, spontaneously broken to $SO(7)$ at the scale f . The radiative stability of the construction is ensured by a \mathbb{Z}_2 symmetry between the SM and the mirror sector, which is softly broken by σ_{soft} to allow for $f > v$ and a viable phenomenology.³³ The $SO(8)$ -breaking quartic κ receives IR contributions from (mirror) top loops $\kappa \simeq (3y_t^4/8\pi^2) \log m_*/m_t$. Finally, $f_0 \simeq f$ up to corrections of order κ/λ_* .

After the spontaneous breaking of the $SO(8)$ symmetry, we are left with two real scalars in the spectrum: the SM-like Higgs h and the radial mode ϕ with masses $m_\phi^2 \simeq 2\lambda_* f^2$ and $m_h^2 \simeq 4\kappa v^2$ in the limit $\lambda_* \gg \kappa, \sigma_{\text{soft}}$. This shows that the Higgs mass is already of the correct size given the typical size of the irreducible contributions to κ . The requirement to reproduce the electroweak scale v and the Higgs mass m_h fixes 2 out of the 4 free parameters in equation (58). The remaining two are chosen to be (f, m_ϕ) while the mixing between the twin Higgs and the SM Higgs is predicted, and it scales as $\sin \gamma \simeq v/f$ in the $m_\phi \gg m_h$ limit.

In twin Higgs models, the fine tuning of the electroweak scale generically scales as

$$\frac{\delta m_h^2}{m_h^2} \simeq \frac{3y_t^4}{2\pi^2\lambda_*} \left(\frac{m_*}{m_h} \right)^2 \log \frac{m_*}{m_t}, \quad (59)$$

where m_* should be interpreted as the cutoff of the potential in equation (58), where colored states are expected to be necessary to stabilize the scale f from radiative corrections. Fixing the scale of the colored states m_* , the fine tuning of the electroweak scale in twin Higgs models is parametrically reduced with respect to that of conventional supersymmetry or composite scenarios by y_t^2/λ_* ; see, e.g., [175, 176]. The gain in fine tuning is limited by the perturbativity requirement on the linear potential in equation (58), which roughly requires $\lambda_* \lesssim 4\pi$, allowing colored states as heavy as 5 TeV with fine tuning on the order of 1%. The implications of a muon collider’s sensitivity to the radial mode are shown in figure 20, reinterpreting the reach illustrated in figure 14 in terms of the twin Higgs parameter space. Keeping the scale of the colored states fixed, the fine tuning of the electroweak scale goes like $1/\lambda_*$, making the sliver of parameter space with largest λ_* (or equivalently with largest m_ϕ/f) the most appealing.

³³ Here for simplicity, we set the \mathbb{Z}_2 -breaking quartic to zero. Having it non-zero will only impact the phenomenology at small m_ϕ , allowing the Higgs

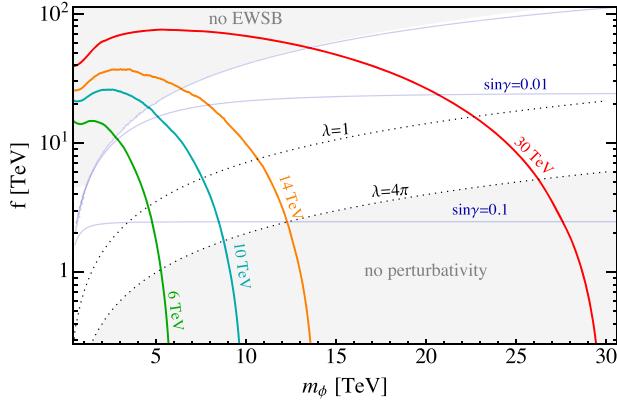


Figure 20. The same limits on a singlet scalar Higgs as shown in figure 14, interpreted in terms of the reach on the sigma-model scale f in the context of a twin Higgs model.

Apart from fine tuning considerations, figure 20 illustrates the correlation between direct searches and Higgs coupling deviations in the twin Higgs parameter space. Interestingly, a high energy muon collider of $\sqrt{s} = 14$ TeV will be able to fully probe the parameter space corresponding to a $\%$ deviation in Higgs couplings ($\sin \gamma \simeq 0.1$) by means of direct searches for the extra scalar. This may be taken as further motivation to reach such a high c.m. energy at a future collider. Conversely, moving into the region with small λ_* , the deviations in the Higgs couplings will be difficult to observe, but direct production at a muon collider could still cover a large portion of this parameter space.

5. Complementarity

Given the time required to achieve first collisions at a future collider, a number of planned and proposed experiments capable of extending our sensitivity to indirect signs of new physics could lead to orders of magnitude improvement of current limits. Among others, experiments searching for EDMs, anomalous flavor violation beyond the standard model, and stochastic gravitational wave backgrounds (SGWB) will probe scales ranging from tens to hundreds of TeV on relevant timescales.³⁴ Ultimately, a signature at any of these experiments would provide an indication of new physics at a scale amenable to further exploration. This motivates asking what energies and luminosities would be required for a muon collider to directly test the origin of indirect signals, providing another set of sharp goalposts.

5.1. EDMs

Electric dipole moments of elementary particles offer a nearly background-free probe of new physics beyond the standard

mass constraint to be satisfied at large f and small m_ϕ . See [28, 175] for a discussion.

³⁴ There are already hints that something interesting might be going with the muon, which could have exciting potential implications for a future muon collider, e.g., see recent studies regarding the muon $g - 2$ [40–44], B meson [45], and K meson [46] anomalies.

model, since all CP-violating SM effects are accompanied by flavor-violating spurions and give rise to extremely small EDMs. Recently, substantial progress has been made in experimental searches for EDMs using paramagnetic molecules. For instance, by studying the polar ThO molecule, the ACME collaboration has set a bound on the electron EDM of $|d_e| \leq 1.1 \times 10^{-29} e \text{ cm}$ at 90% confidence [177]. Rapid experimental progress towards the use of atoms and molecules, including novel approaches using polyatomic molecules, is likely to improve the electron EDM sensitivity to at least $10^{-32} e \text{ cm}$ in the coming decade (see [178, 179] and references therein). These approaches will also offer novel opportunities to probe hadronic EDMs.

The discovery of an EDM could definitively prove the existence of new physics beyond the SM. In a typical model, weakly-coupled k -loop physics, at a mass scale M with CP-violating phase δ_{CP} , could generate an electron EDM of the size

$$d_e \sim \sin(\delta_{\text{CP}}) \frac{e m_e}{M^2} \left(\frac{\alpha}{4\pi} \right)^k \simeq 10^{-32} e \text{ cm} \sin(\delta_{\text{CP}}) \\ \times \begin{cases} (1 \text{ PeV}/M)^2 & \text{for } k = 1 \\ (20 \text{ TeV}/M)^2 & \text{for } k = 2 \end{cases} \quad (60)$$

This is only a rough, order-of-magnitude guide, but it shows that an EDM experiment could provide the first discovery of the effects of new physics beyond the reach of current colliders. A one-loop EDM of this size could arise, for example, from sleptons with masses of order PeV and order one CP-violating phases [180, 181], or from lighter sleptons with smaller CP-violating phases. A two-loop EDM could arise, for instance, from chargino loops in split SUSY [182, 183]. More generally, *any* new particles interacting with electroweak gauge bosons and the Higgs can produce a two-loop EDM through Barr-Zee diagrams [184], see the left panel of figure 21.

An EDM discovery would tell us that new physics exists, but would tell us very little about the nature of the new physics. Colliders will have a crucial role to play, by providing more insight on the new particles responsible for the EDM and allowing us to determine which extension of the SM explains the effect. A one-loop EDM would be associated with new particles carrying electric charge which could be directly pair-produced at a muon collider if they are kinematically within reach, e.g., for sleptons we would search for $\mu^+ \mu^- \rightarrow \tilde{\ell}^+ \tilde{\ell}^-$. The range of possibilities at two loops is broader (e.g., [185, 186]), but the case of the Barr-Zee diagram offers a particularly appealing target. It involves new electroweak particles which could be pair produced directly. However, it *also* implies that these particles alter the interactions between Higgs and gauge bosons, as illustrated in figure 21. Since high-energy lepton colliders are electroweak gauge boson colliders, they offer a unique prospect to directly probe the same underlying electroweak interactions that generate the EDM, via precision studies of processes like $\gamma\gamma \rightarrow hh$. We leave a detailed assessment of the prospects for

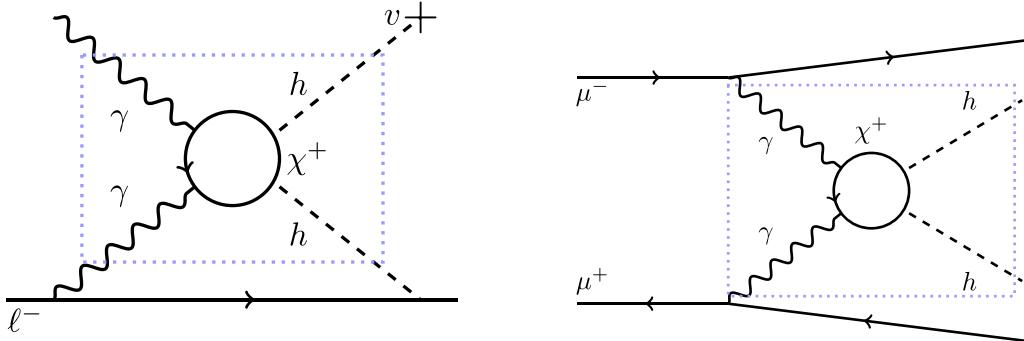


Figure 21. One possibility for testing the physics associated with a Barr-Zee type contribution to a lepton EDM at a future muon collider. Left: the two-loop Barr-Zee contribution to a lepton EDM. Right: a $\gamma\gamma \rightarrow hh$ process at a muon collider, sensitive to loops of charged particles that couple to the Higgs. The dotted blue box shows that both processes probe the same underlying physics.

EDM/collider complementarity along these lines for future work.

5.2. Flavor

One of the biggest puzzles in the SM is the pattern of fermion masses and mixings. Both the quark and lepton sectors have significant mass hierarchies, whereas the mixing matrices take a very different form in the two sectors. We expect that at high energies, where the flavor pattern of the SM is established, there may be much larger rates of flavor-changing processes than the SM predicts. This is a strong motivation for searching for flavor-violating processes at high-energy colliders.

Conversely, some of our most stringent bounds on physics beyond the SM come from low- or medium-energy precision tests of flavor-changing processes, e.g., $K - \bar{K}$ or $D - \bar{D}$ mixing [187]. Of particular interest, in the context of a muon collider, are precision tests of charged lepton flavor violation processes like $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$, $\tau \rightarrow 3\mu$, or μ -to- e conversion within atomic nuclei. As with EDMs, these processes are expected to be much more strongly constrained in the coming decades than they are at present, due to a number of currently operating or planned future experiments [188]. These experiments can indirectly probe physics at energies of 10s of TeV or even higher. A high-energy muon collider can probe the same physics (e.g., through direct searches for flavor-changing processes like $\mu^+ \mu^- \rightarrow \tau^\pm \mu^\mp$), or can help to elucidate the underlying mechanism of flavor violation by directly producing new particles with flavor-violating interactions, such as the mixed slepton production process $\mu^+ \mu^- \rightarrow \tilde{\ell}_i^+ \tilde{\ell}_j^-$. Below, we will give first estimates of the physics reach for both of these scenarios. We will see that there is a powerful complementarity between precision lepton flavor experiments and high-energy muon collider searches.

5.2.1. Lepton-flavor violating contact interactions. Belle has set a limit of $\text{Br}(\tau \rightarrow 3\mu) < 2.1 \times 10^{-8}$ [189], and this bound will be improved to 3.5×10^{-10} by Belle II [190]. These bounds can be contrasted with constraints on the $\mu \rightarrow 3e$ branching ratio, which is currently constrained to be

1.0×10^{-12} by SINDRUM [191], but will be dramatically improved, eventually to a sensitivity $\sim 1 \times 10^{-16}$, by the Mu3e experiment [188].

We parameterize the four-fermion operators relevant for the $\tau \rightarrow 3\mu$ decay via

$$\mathcal{L} \supset V_{LL}^{\tau 3\mu} (\bar{\mu} \gamma^\mu P_L \mu) (\bar{\tau} \gamma_\mu P_L \mu) + V_{LR}^{\tau 3\mu} (\bar{\mu} \gamma^\mu P_L \mu) (\bar{\tau} \gamma_\mu P_R \mu) + (L \leftrightarrow R) + \text{h.c.}, \quad (61)$$

with an equivalent set for the $\mu \rightarrow 3e$ decay. In what follows, we will assume all the $\tau 3\mu$ coefficients are equal:

$$V_{LL}^{\tau 3\mu} = V_{LR}^{\tau 3\mu} = V_{RL}^{\tau 3\mu} = V_{RR}^{\tau 3\mu} = \frac{c^{\tau 3\mu}}{\Lambda^2}, \quad (62)$$

where $c^{\tau 3\mu}$ is a dimensionless coefficient and Λ is to be interpreted as the scale of new physics. Setting the Wilson coefficients to unity, we see that the current Belle (future Belle II) constraint above translates into a bound of $\Lambda > 14.7$ TeV ($\Lambda > 40.9$ TeV), far below the corresponding constraints from SINDRUM (with $\tau 3\mu$ replaced by $\mu 3e$), which gives $\Lambda > 273$ TeV. Generically, though, one might expect new physics violating lepton flavor to manifest with hierarchies similar to the pattern of hierarchies in the SM lepton Yukawas, in which case the bounds from τ and μ might be more comparable. We will discuss this more in what follows.

At a muon collider, these same operators can be probed directly via $\mu^+ \mu^- \rightarrow \ell_i^+ \ell_j^-$. We will focus on $\mu\tau$ production, since it can be compared directly with the sensitivity from tau decays. Our analysis closely follows an analogous study at an $e^+ e^-$ collider in [192]. With only the four-fermion operator insertion, the rate for $\mu^+ \mu^- \rightarrow \mu\tau$ grows as s . Taking $c^{\tau 3\mu}/\Lambda^2 = 1/(50 \text{ TeV})^2$, the cross section grows from 0.059 fb at a 3 TeV muon collider to 66.1 fb at $\sqrt{s} = 100$ TeV.

Ignoring detector effects, there are two primary backgrounds for the flavor-violating process in the SM: $\tau^+ \tau^-$ production, where (at least) one of the taus decays to a muon and a neutrino and $\mu^+ \mu^- \rightarrow \mu\nu_\mu \tau\nu_\tau$ production via intermediate W bosons. These backgrounds can be significantly suppressed with simple cuts on the muon energy and the missing three-momentum. In particular, we demand that

Table 5. Number of signal and background events after kinematic cuts and estimating the loss of signal efficiency due to initial state radiation for 1 ab^{-1} of data and $c^{\tau^3\mu}/\Lambda^2 = 1/(50 \text{ TeV})^2$.

\sqrt{s} (TeV)	$N_{\text{post-cuts}}$		
	$\mu^+\mu^- \rightarrow \mu\tau$	$\mu^+\mu^- \rightarrow \mu\tau\nu_\mu\nu_\tau$	$\mu^+\mu^- \rightarrow \tau^+\tau^-$
0.125	0.0948	30.8	3.42×10^4
3	53.3	6.32×10^3	40.4
6	212	3.26×10^3	9.52
14	1.14×10^3	1.14×10^3	0.138
100	5.73×10^4	60.9	0.0312

the most energetic muon in the event has at least 90% of the beam energy and that the direction of the total missing three-momentum vector be at least 170° from the three-momentum of the most energetic muon. These cuts suppress the $\tau^+\tau^-$ background by a factor ~ 300 , and the W^+W^- background by a factor ~ 10 , depending slightly on the c.m. energy.

The signal process, on the other hand, passes these cuts with near perfect efficiency. Kinematically, there is no loss of signal events, and the only degradation is due to initial state radiation. This leads to a $\sim 10\%$ reduction at $\sqrt{s} = 3 \text{ TeV}$, slowly increasing to a 13% reduction at $\sqrt{s} = 100 \text{ TeV}$. The number of signal and background events after applying the kinematic cuts and initial state radiation degradation for 1 ab^{-1} of data with $c^{\tau^3\mu}/\Lambda^2 = 1/(50 \text{ TeV})^2$ are shown in table 5. We find a signal-to-background ratio ~ 1 to be achievable at a 14 TeV machine, and this dramatically rises to $\sim 10^3$ at a 100 TeV machine. The resulting bounds, assuming integrated luminosities of 1 ab^{-1} at 0.125, 3, 6, 14, and 100 TeV are shown in figure 22. It is clear that even a 3 TeV machine would be able to set a direct bound at the same level as the future Belle II sensitivity, and this constraint can be improved by up to ~ 2 orders of magnitude at a higher energy machine.

In figure 22, these results are compared to the constraints on the analogous four-fermion operator in the $\mu \rightarrow 3e$ decay with various ansatz regarding flavor violation. The diagonal lines show the expected relationship between the two Wilson coefficients assuming (i) flavor anarchy (all coefficients ~ 1), (ii) minimal leptonic flavor violation (MLFV) [193], (iii) the Wilson coefficients scale like the square root of the Yukawa couplings of the leptons involved in the flavor violation, and (iv) the Wilson coefficients scale like the product of the same Yukawa couplings.³⁵

While the muon decay sets the strongest limits assuming anarchical coefficients, we see that a 14 TeV muon collider could set a bound comparable to the current SINDRUM limit in the case of MLFV, and would be comparable to the stage-I Mu3e sensitivity if the coefficients scale like the square root of the Yukawa couplings. In the extreme case, where the Wilson coefficients behave like the product of the

³⁵ One should be cautious that if a flavorful ansatz is used, the inferred scale from the 100 TeV bounds may be low enough that an effective field theory description is no longer valid.

two Yukawas, even a 3 TeV muon collider would provide a bound complementary to the final Mu3e sensitivity, with higher energy machines improving this bound by orders of magnitude.

In addition to the $\tau^3\mu$ operators considered here, we expect roughly similar sensitivity to the $\mu^+\mu^- \rightarrow \mu^\pm e^\mp$ process, as well as to the processes such as $\mu^+\mu^- \rightarrow \tau^\pm e^\mp$ that violate lepton flavor by two units. Overall, we see that a muon collider would be capable of directly probing flavor-violating interactions that are quite complementary to future precision constraints.

5.2.2. Direct probes of lepton-flavor violation in the MSSM.

Charged lepton flavor violation in the MSSM arises as a result of the soft-breaking terms in the slepton mass matrix having non-diagonal entries in the basis where the SM lepton Yukawas are diagonal. In this case, the physical sleptons will be mixtures of different flavors, and their interactions with leptons and neutralinos/charginos will be flavor-violating. These flavor violating interactions lead to processes such as rare muon decays or muon-to-electron conversion at loop level, and thus, low energy experiments can indirectly probe these interactions with sensitivities extending beyond the TeV scale, depending on the flavor structure of the theory [180, 194]. A high-energy muon collider, on the other hand, would not only be capable of producing superpartners at high masses, but would also provide direct measurements of the lepton-flavor violating processes that would complement these low-energy probes and provide detailed insight into the mechanism of supersymmetry breaking.

For simplicity, we will consider a simplified scenario where the effects of all scalar superpartners except for \tilde{e}_R and $\tilde{\mu}_R$ decouple. In this case, the slepton mixing reduces to a 2×2 problem with slepton-mass squared matrix

$$\mathcal{M}_{\tilde{\ell},RR}^2 = \begin{pmatrix} \Delta_{RR,11} & \tilde{m}_{E,12}^2 \\ \tilde{m}_{E,12}^2 & \Delta_{RR,22} \end{pmatrix}, \quad (63)$$

where the diagonal terms are the sum of both soft-SUSY-breaking scalar masses (\tilde{m}_E^2) and D -terms as well as terms dictated by supersymmetry, and we have assumed the off-diagonal soft-breaking terms are CP conserving. This mass matrix can be diagonalized via a unitary matrix U_R to yield mass eigenstates $m_{\tilde{e}_1}^2, m_{\tilde{e}_2}^2$ with the mixing angle given by

$$\frac{1}{2} \sin(2\theta_R) = \frac{\tilde{m}_{E,12}^2}{m_{\tilde{e}_1}^2 - m_{\tilde{e}_2}^2}. \quad (64)$$

We will further consider the situation where the LSP is a pure bino with mass M_1 and assume the other neutralinos can be ignored. With $m_{\tilde{\ell}} > M_1$, the sleptons decay directly to a lepton and bino, and the latter will appear as missing momentum.

As a benchmark scenario, we consider the situation where the selectron and smuon are nearly degenerate in mass. This is well-motivated from gauge-mediated supersymmetry breaking scenarios, and also leads to a strong suppression of the lepton-flavor violation via a ‘super-GIM’ mechanism, allowing the superpartners to be relatively light. Such a scenario

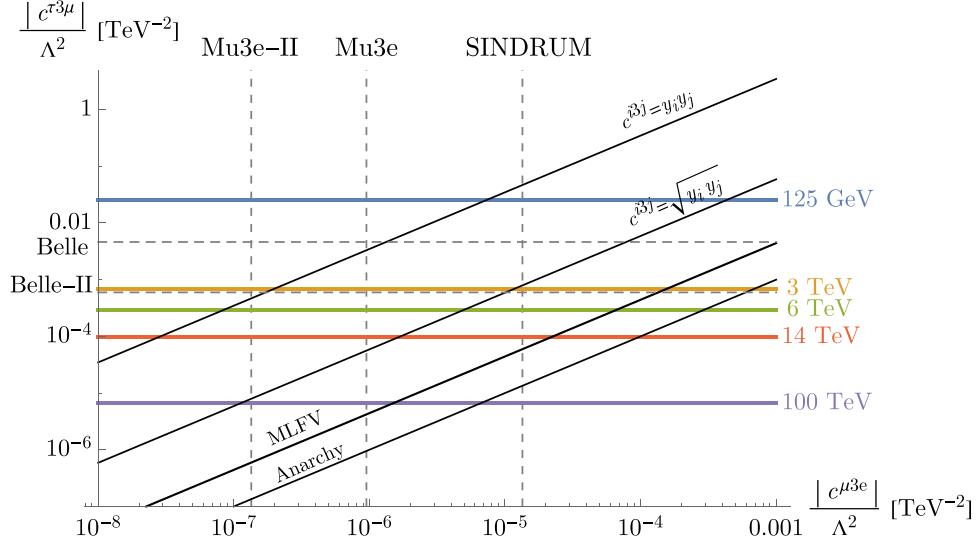


Figure 22. Summary of muon collider and precision constraints on flavor-violating three-body decays. The colored horizontal lines show the sensitivity to the $\tau^3\mu$ operator at various energies, all assuming 1 ab^{-1} of data. The dashed horizontal (vertical) lines show the current or expected sensitivity from $\tau \rightarrow 3\mu$ ($\mu \rightarrow 3e$) decays for comparison. The diagonal black lines show the expected relationship between different Wilson coefficients with various ansatz for the scaling of the flavor-violating operators (e.g., ‘Anarchy’ assumes that all Wilson coefficients are $\mathcal{O}(1)$).

was previously studied in the context of e^+e^- collisions in [195], but for relatively light superpartners. A high-energy muon collider would allow similar tests with a substantially more impressive mass reach.

In the limit of a small mass splitting, the parameter governing the amount of flavor violation is given by

$$(\delta_{RR})^{12} \equiv \frac{\tilde{m}_{E,12}^2}{\sqrt{\Delta_{RR,11}\Delta_{RR,22}}} \simeq \frac{1}{2} \frac{\Delta m^2}{\bar{m}^2} \sin(2\theta_R), \quad (65)$$

where we have introduced the average mass squared, \bar{m}^2 and mass-splitting $\Delta m^2 = m_{e_1}^2 - m_{e_2}^2$. Following [195], we can compute the probability that a gauge eigenstate $\tilde{\mu}$ decays into a final state with an electron—including the interference effects when $\bar{m}\Gamma$ and Δm^2 are of similar size—and thus find the cross section for the flavor violating process $\mu^+\mu^- \rightarrow \tilde{e}_{1,2}^+ \tilde{e}_{1,2}^- \rightarrow \mu^\pm e^\mp \chi_1^0 \chi_1^0$ as a function of the mass splitting and mixing angle. The results are shown in figure 23 for several choices of slepton and bino mass and at several different several of mass energies alongside the current bounds from $\mu \rightarrow e\gamma$ from MEG [196] and the expected sensitivities from several future experiments [197–199].

The primary background for these flavor-violating processes at a lepton collider is production of the different flavor final states and missing energy via intermediate W bosons. The total cross section for this background, including branching ratios, is 52 fb at a 6 TeV collider (15/0.6 fb at a 14/100 TeV machine), but the kinematics of this process are quite different from the slepton-pair production signal of interest. Moreover, for the flavor-violating scenarios at hand, it is likely that the relevant slepton and neutralino masses will have already been measured from the corresponding flavor-conserving processes (for details on how this can be done, see e.g., [200, 201]). With the slepton masses known, it is

then possible to fully reconstruct the two final state neutralino momenta by requiring that the neutralino and lepton momenta satisfy the slepton mass-shell constraint, along with conservation of energy and momentum. In general, these conditions will be impossible to satisfy for the background events, and indeed, we find in simulation that only $\sim 1/500$ background events can reconstruct the neutralino momenta while satisfying conservation of energy, while generally $\sim 98\%$ of the signal events reconstruct the momenta successfully. Further details of this analysis procedure are presented in [202].

In figure 23, we show the 5σ discovery reach for the three benchmark scenarios assuming 1 ab^{-1} (5 ab^{-1}) of data in red dashed (dotted) lines. We see that a 6 TeV muon collider would probe a great deal of parameter space complementary to the stage-II Mu2e bounds in the case of 1 TeV nearly-degenerate sleptons. More energetic colliders would be able to measure parameters that are beyond the reach of even the PRISM/PRIME future sensitivity.

Alternatively, focusing on the case of a relatively large mass splitting between the selectron and smuon, we can examine the reach as a function of δ_{RR}^{12} and the Bino mass M_1 . The results are shown in figure 24, where we again superimpose the current and future sensitivities from the low-energy experiments. The 1 and 5 ab^{-1} discovery reaches are again shown as dashed and dotted red lines, respectively.

Here we see the cross section varies only slightly as a function of M_1 , as the production is predominantly via s -channel photons and Z bosons. The muon collider reach, however, becomes much stronger at large bino masses. This can be understood as the signal having a very compressed phase space, due to the small difference between the slepton and bino masses. While the signal can still be reconstructed from knowing the initial and final visible four-momenta,

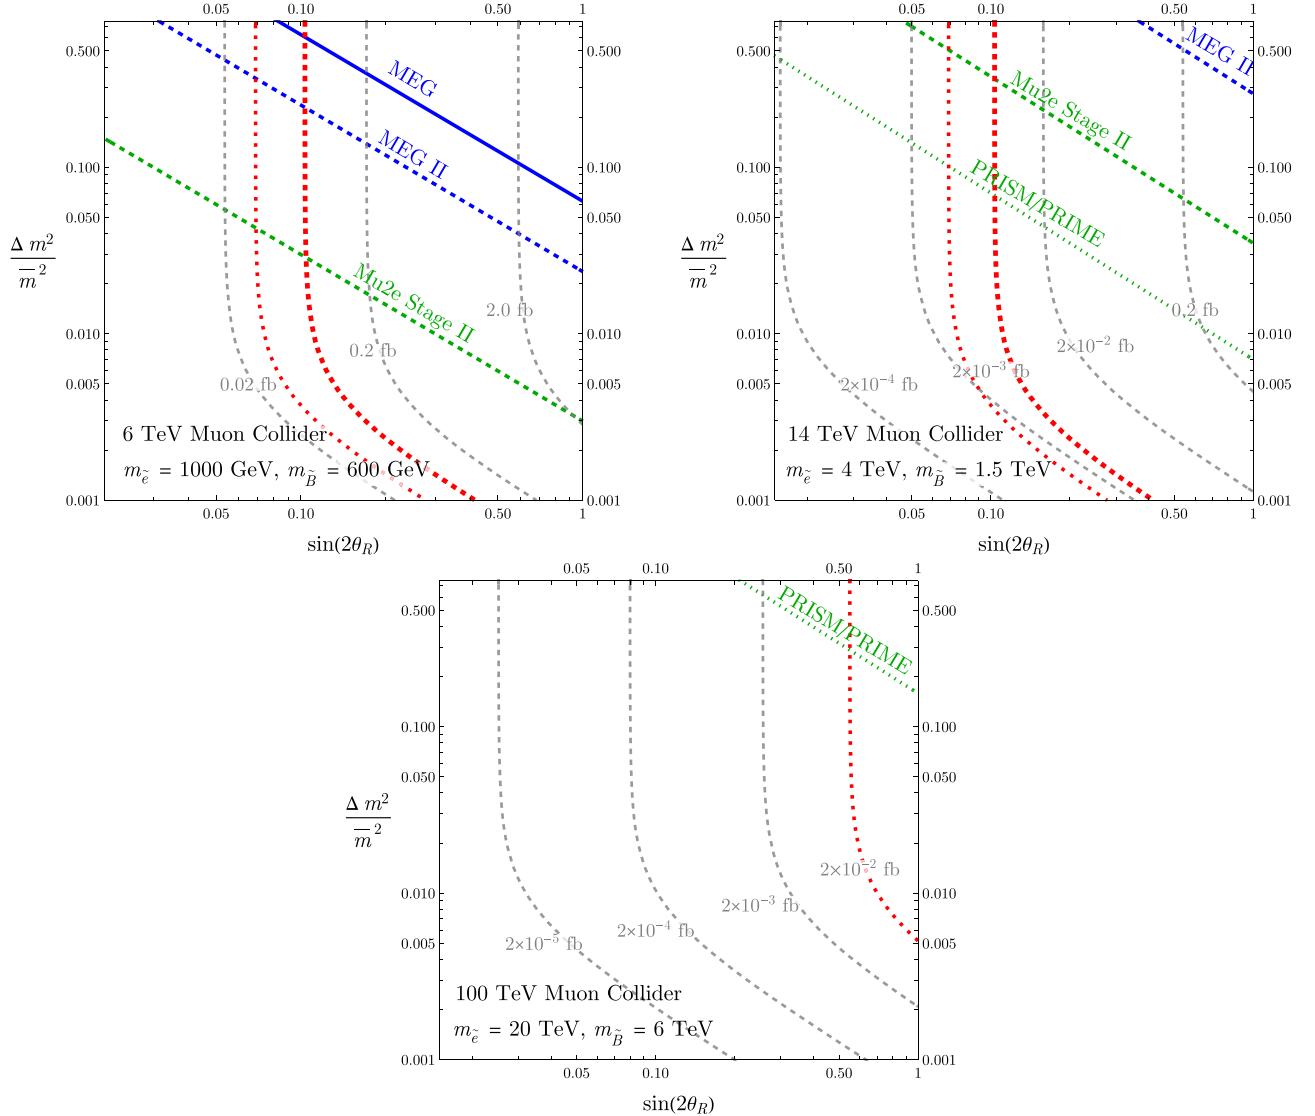


Figure 23. Contours (in dashed gray) of the cross section in the $\Delta m^2 / \bar{m}^2$ vs $\sin(2\theta_R)$ plane for the flavor violating process $\mu^+ \mu^- \rightarrow \mu^\pm e^\pm \chi_1^0 \chi_1^0$ for the nearly-degenerate slepton scenario described in the text. The bounds from current and future precision searches are overlaid in blue and green, along with red dashed and dotted curves showing the expected reach with 1 and 5 ab^{-1} of data. The different panels correspond to benchmark choices of c.m. energy, and superpartner masses.

the WW background generally does not fall into this small corner of phase space, and can thus be more efficiently vetoed.

We see again that a 6 TeV collider would cover much of the parameter space probed by MEG II, except perhaps at very small Bino masses where the signal more closely mimics the background, but would fall short of the eventual Mu2e Stage II sensitivity. A higher energy muon collider, however, could surpass the Mu2e sensitivity, and could measure flavor violating insertions as small as a few $\times 10^{-3}$ for sleptons at ~ 4 TeV with a muon collider at 14 TeV. A 100 TeV collider would be able to surpass even the most optimistic future μ -to- e conversion constraints for sleptons in the tens of TeV range. Overall, these benchmarks illustrate that muon colliders would have impressive capabilities for not only discovering superpartners, but measuring their flavor structure in detail.

5.3. Gravitational waves

The central importance and capabilities of a future collider must be seen within the context of the broader experimental efforts in particle physics. We have already discussed the connections between possible discoveries in precision flavor and CP violation experiments, direct and indirect dark matter detection, as well as searches for new light fields and dark forces. On the cosmological front, we are entering an era in which dramatic new forms of ‘fossil’ evidence for BSM physics may be found, within SGWB [203, 204] and within primordial non-Gaussianities in large scale structure (LSS) and high-redshift 21 cm 3D ‘maps’ [205]. While discoveries in any of these non-collider experiments would be spectacular, powerful new colliders would provide the ‘gold standard’ laboratory conditions to corroborate, connect, extend, and analyze their full significance.

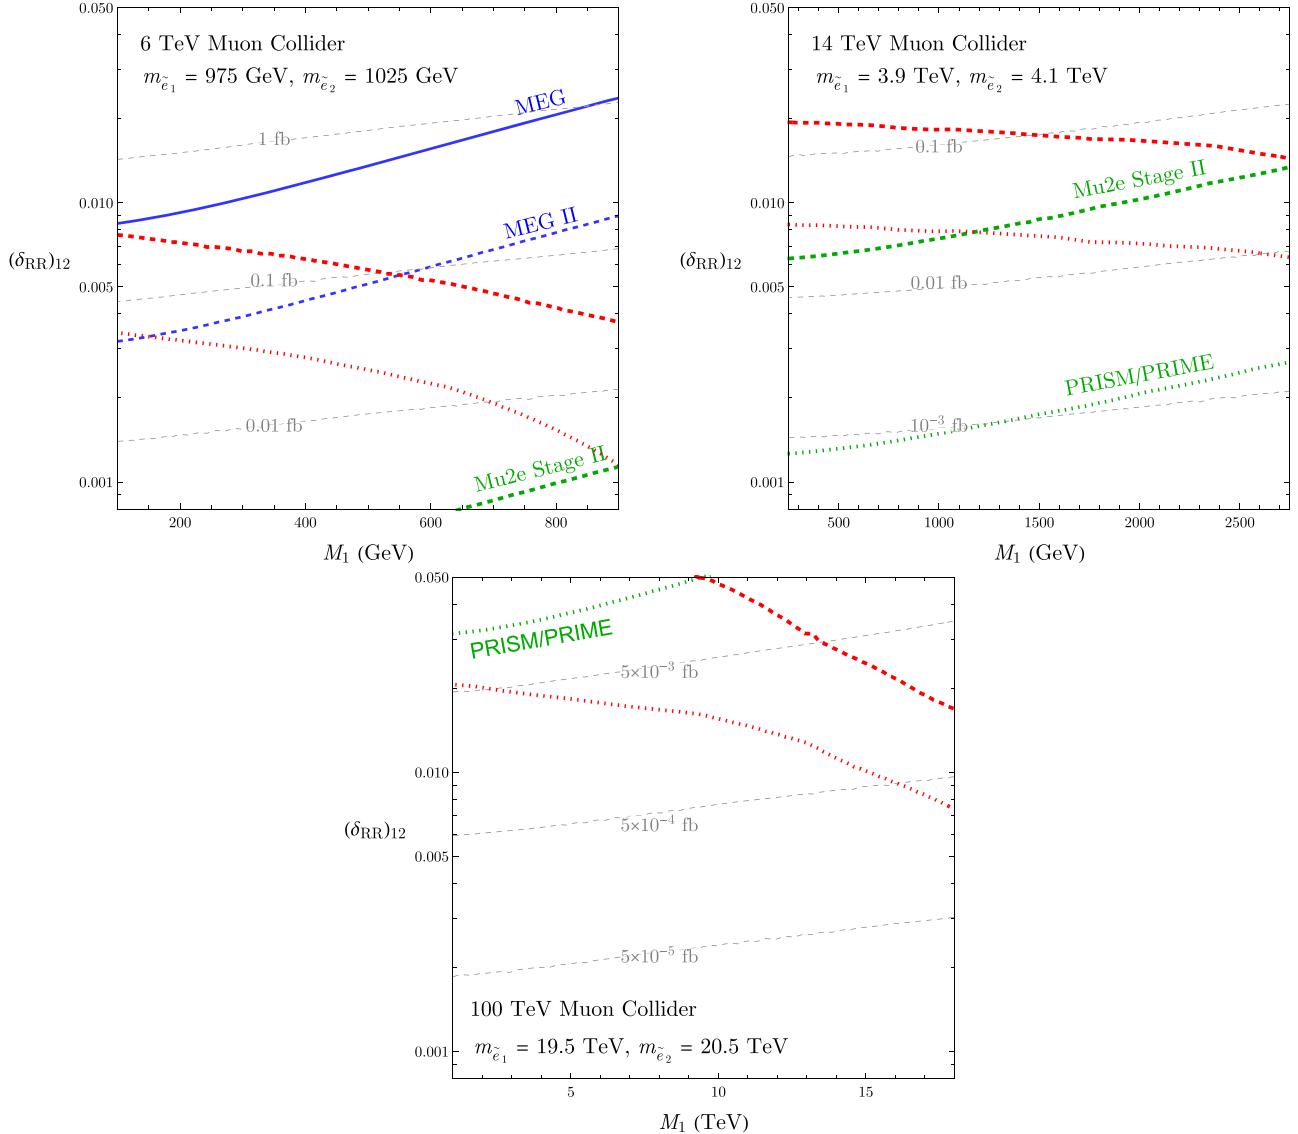


Figure 24. Similar to figure 23, but in the $(\delta_{RR})_{12}$ vs M_1 plane for a large mass splitting scenario.

First order cosmological phase transitions that could occur due to extensions of the SM or within dark sectors could be powerful sources of SGWB, while possibly providing the non-equilibrium conditions needed for generating matter asymmetries. The peak SGWB frequency, after redshifting from the time of production in the very early Universe, is given by

$$\omega \simeq 0.03 \text{ mHz} \frac{\beta}{H_{\text{PT}}} \frac{T_{\text{PT}}}{\text{TeV}}, \quad (66)$$

where T_{PT} is the temperature immediately after the phase transition, $1/\beta$ is essentially its duration, and H_{PT} is the Hubble expansion rate during this era [206, 207]. Typically, one expects $\beta \sim \mathcal{O}(10 - 100)H_{\text{PT}}$. Fortunately then, for T_{PT} in the BSM-motivated range, TeV–100 TeV, we can expect SGWB in roughly the mHz–Hz range accessible to proposed gravitational wave detectors such as LISA, BBO, and DECIGO. If a SGWB from a phase transition is detected, it would be critical to piece together the information in its frequency spectrum

with the complementary microphysics accessible within collider experiments to whatever extent possible. Quite plausibly, these elements relate to extensions of the SM Higgs sector. A high energy muon collider provides a balance of potential to probe the Higgs sector, to create very massive BSM states related to the phase transition, and given its clean environment to possibly produce and diagnose a small number of events resulting from the presence of a dark sector in case the SGWB originated there.

Let us illustrate the interplay between gravitational wave detection and a muon collider by one of the central questions of particle physics: is the Higgs boson truly elementary, or a composite of new strongly-coupled confined constituents? In the latter case, the compositeness would greatly mitigate the electroweak hierarchy problem (see [154, 208] for review). We would then expect to see excited composites at high-energy colliders as well as a deconfinement-to-confinement phase transition at gravitational wave detectors. This physics is usually discussed in its more theoretically tractable

AdS/CFT dual formulation as a 5D Randall–Sundrum I (RS1) model [209], translating to the collider phenomenology of Kaluza–Klein excitations and the transition between the black-brane horizon and IR-brane phases of RS1 [210–214]. For a very recent work refining phase transition dynamics and for references, see [215]. For example, if LISA were to detect a SGWB with peak frequency 0.5 mHz and gravitational power spectrum of 3×10^{-8} at the peak, in the RS1 framework this would correspond to a supercooled phase transition with $T_{\text{PT}} \simeq 1.3$ TeV, with Kaluza–Klein resonances in the (multi-)TeV regime accessible to muon colliders.³⁶ If say a 2 TeV ‘radion’ spin-0 excitation were discovered at a muon collider via VBF, its mass and cross-section would independently determine T_{PT} , giving a valuable cross-check of the underlying physics. Further, it would bound the heavy Kaluza–Klein spectrum to begin below 40 TeV. If say a spin-2 Kaluza–Klein graviton with mass 20 TeV is eventually discovered in associated production with the radion, the cross-section for this process (along with cross-sections for radion processes) together with the frequency and power spectrum from the SGWB data would allow us to extract the small ‘critical exponent’ responsible for generating the large electroweak hierarchy, in this example $\epsilon = 0.1$. We emphasize that this central parameter of the composite dynamics would be difficult to extract experimentally without combining collider and gravitational wave data. All the relevant measurements would be challenging, but worthy ambitions for future experiments.

Upcoming precision LSS and 21 cm surveys offer the potential to detect heavy particle production and propagation during inflation, imprinted on the non-Gaussian bispectrum in distinctive non-local effects (non-analytic in co-moving momenta). This field of ‘cosmological collider physics’ is sensitive to particle masses of order the inflationary Hubble scale or even somewhat higher [217, 218]. (For very recent work and references, see [219]). We do not as yet know the scales of inflation. If new particles are discovered in cosmological non-Gaussianities, they may lie far above the reach of terrestrial colliders, in which case they would give complementary information to what we learn from even a powerful muon collider. But there are two scenarios in which they could give us a (pre)view of collider-accessible physics: (i) if the inflationary Hubble scale is of order 100 TeV or less, then obviously ‘cosmological collider physics’ may directly be within reach of future terrestrial colliders; (ii) even if the inflationary Hubble scale is orders of magnitude above the TeV scale, there is a ‘heavy-lifting’ mechanism [220] whereby the particles seen in non-Gaussianities were given inflationary scale masses through strong curvature effects, but such effects are negligible today so that the particles may now be within terrestrial collider reach.

³⁶ Given a high degree of supercooling, say by a factor of 10^3 , we are assuming that bubble collisions dominantly source the SGWB signal and that these are captured by the ‘envelope approximation’. This approximation can be invalid in some regimes, but its replacements would not invalidate the general complementarity of the collider and gravitational wave experiments. See the discussion reviewed in [216].

6. Summary and future directions

The goal of this work is to paint the physics case of a high-energy muon collider with a broad brush, emphasizing the sense in which such a collider is positioned to answer the many questions posed or sharpened by the discovery of the Higgs. The broad outlines of this case are drawn by the physics of both muon annihilation and VBF, which in tandem provide compelling rates for standard model and beyond-the-standard model processes across a range of energies. Relative to recent work highlighting the significance of VBF, we have emphasized the value of muon annihilation as a discovery mode for sufficiently distinctive new physics. To characterize the rich physics of the initial state, we surveyed descriptions of the virtual electroweak gauge boson content of high-energy muons ranging from the EVA to electroweak PDFs. As a practical matter, we emphasized the sense in which the simple ‘leading log PDF’ is sufficient to capture most of the qualitative physics, and briefly explored the impact of finite W and Z masses on electroweak PDFs.

Turning to the physics case itself, we highlighted the potential of a muon collider to illuminate various aspects of EWSB, dark matter, and naturalness. The measurement of the top Yukawa coupling provided a particularly sharp test case for probing the physics of EWSB, in which deviations from the standard model prediction lead to significant changes in the rates for longitudinal VBF into $t\bar{t}$ pairs. This is sufficient for a muon collider at or above $\sqrt{s} = 14$ TeV to test the top Yukawa coupling beyond the expected sensitivity of the LHC. The increase in reach provided by polarized beams exemplifies their potential value. As a second test case relevant to the Higgs potential and the electroweak phase transition, we proposed two avenues for constraining the so-called ‘nightmare scenario’ in which the electroweak phase transition is strengthened by a light, \mathbb{Z}_2 -symmetric singlet scalar that only couples to the Higgs. In addition to the canonical search for pair production of the scalar in missing energy final states, we noted the possibility of again using longitudinal VBF into $t\bar{t}$ pairs, this time looking for a feature in the invariant mass distribution associated with the scalar threshold. In both channels, a collider operating above $\sqrt{s} = 10$ TeV can cover much of the motivated parameter space. We further considered direct searches for additional Higgs bosons, focusing on the case of a singlet scalar mixing with the Higgs. Here the abundant production of such scalars via VBF would allow a collider operating above $\sqrt{s} = 10$ TeV to probe mixing angles whose indirect imprints (in the form of Higgs coupling deviations) lie beyond the reach of any proposed collider, and a collider operating above $\sqrt{s} = 30$ TeV to far exceed the mass reach of a 100 TeV pp collider.

The implications of a muon collider for dark matter are exemplified by its coverage of ‘minimal dark matter’ models, in which the dark matter particle resides in an electroweak multiplet whose interactions with SM gauge bosons can generate the observed abundance. We highlighted two classes of search strategies for these multiplets, using either a missing mass or a disappearing track signature. Depending on the integrated luminosity, a high-energy muon collider operating at $\sqrt{s} = 10$

or 14 TeV can *discover* smaller electroweak representations (such as an $SU(2)$ doublet or triplet) at their thermal targets, while a collider operating at $\sqrt{s} = 30\text{--}100$ TeV can cover the thermal targets for higher electroweak representations as well.

We illustrated the potential of a muon collider to discover or constrain natural explanations of the weak scale by considering the reach for both solutions to the ‘big’ hierarchy problem (supersymmetry or compositeness) and the ‘little’ hierarchy problem (represented here by the twin Higgs). Within the context of supersymmetry, searches for the higgsino and the stop would allow a muon collider operating at $\sqrt{s} \gtrsim 20$ TeV to comprehensively cover tuning beyond the per mille level, and collisions at $\sqrt{s} = 30$ TeV would reach stop masses consistent with the observed Higgs mass even in the absence of significant stop mixing. Despite its lack of electroweak quantum numbers, even the gluino could be probed beyond LHC limits leveraging radiation off of $q\bar{q}$ pairs at $\sqrt{s} \gtrsim 10$ TeV. In addition to these conventional channels, a muon collider can provide the first direct collider test of the supersymmetry-breaking sector itself, achieving sensitivity to direct production of the gravitino consistent with LHC bounds on standard model superpartners, provided a low scale of supersymmetry breaking. The coverage of compositeness is comparable, with the addition of powerful indirect tests coming from the energy growth of irrelevant operators parameterizing mixing of the Higgs and composite states. Even solutions to the ‘little’ hierarchy problem without conventional LHC signatures, such as the twin Higgs, lie within reach of a muon collider. In addition to its considerable sensitivity to indirect effects, a high-energy muon collider would be able to directly access the radial mode of these models and decisively confirm or falsify them as a viable explanation of the weak scale.

A host of other experiments will indirectly probe physics as high as the PeV scale in the years preceding the first beams at a high-energy muon collider. To this end, we have considered the complementarity between a muon collider and potential signals of new physics in EDMs, flavor violation, and gravitational waves. Of particular interest are precision tests of charged lepton flavor violation in processes such as $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$, $\tau \rightarrow 3\mu$ and μ -to- e conversion. Here we have explored the detailed reach of a muon collider for both indirect sources of lepton flavor violation (such as flavor-violating four-fermion operators) and direct sources (such as flavor-violating slepton interactions in the MSSM), in both cases finding that muon colliders below $\sqrt{s} \sim 10$ TeV provide complementary sensitivity to experiments such as Mu2e and Mu3e, while more energetic colliders are capable of probing parameter space beyond the reach of future proposals.

Broadly speaking, we find that a muon collider operating at tens of TeV and tens of ab^{-1} is capable of surpassing the indirect reach of proposed e^+e^- Higgs factories and the direct reach of proposed 100 TeV pp colliders, covering a broad swathe of motivated physics beyond the standard model and probing explanations for potential signals in experiments across the many frontiers of particle physics. Perhaps the most

compelling gains stand to be made by colliders operating between $\sqrt{s} = 14\text{--}30$ TeV, for which most of the questions posed by the discovery of the Higgs boson may be decisively answered.

Needless to say, many of the projections made in this work are naive in light of the significant uncertainties and many unresolved challenges facing both accelerators and detectors. Nonetheless, we hope that they provide a qualitative guide to the energies and luminosities that would position a future muon collider as a comprehensive successor to the LHC, pinpointing a variety of directions that merit more careful study. Beyond characterizing the reach in conventional benchmarks for a future collider program, we have identified a number of opportunities uniquely suited to a muon collider, including direct tests of low-energy supersymmetry breaking. Key outstanding questions include the performance and prospects of searches involving missing energy, which are central to the coverage of dark matter, the electroweak phase transition, and supersymmetry breaking; the invariant mass resolution in heavy standard model final states, essential to making the most of the abundant opportunities provided by the fusion of longitudinal vector bosons; and the feasibility of fully instrumenting the forward region, which will shape the set of available observables and the composition of signal processes.

Much remains before us. But the muons are calling, and we must go.

Acknowledgments

We are grateful to Ken Van Tilburg for collaboration during the early stages of this work. We thank Kaustubh Agashe, Paolo Ciafaloni, Tao Han, Sergo Jindariani, Fabio Maltoni, and Dave Soper for useful conversations. The work of I Banta, T Cai, N Craig, G Koszegi, A McCune, and T Trott is supported by the US Department of Energy under the Grant DE-SC0011702. The work of D Buttazzo and A Tesi is supported by MIUR under Grant Number 2017L5W2PT. The work of T Cohen is supported by the US Department of Energy under Grant Number DE-SC0011640. The work of J Fan is supported by the US Department of Energy under Grant Number DE-SC-0010010. The work of I Garcia Garcia is funded by the Gordon and Betty Moore Foundation through Grant GBMF7392, and in part by the National Science Foundation under Grant No. NSF PHY-1748958. The work of S Homiller, Q Lu, and M Reece was supported in part by the DOE Grant DE-SC0013607 and by the Alfred P Sloan Foundation Grant No. G-2019-12504. The work of P Meade is supported in part by the National Science Foundation under Grant No. NSF PHY-1915093. The work of A Mariotti is supported in part by the SRP HEP-VUB and by the EOS-be.h n.30820817. D Sutherland has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 754496. The work of L-T Wang is supported by the DOE Grant DE-SC0013642. The work of S Koren is funded by a Mafalda and Reinhard Oehme Postdoctoral Fellowship. The work of M Ekhterachian and R Sundrum is supported by the NSF Grant PHY-1914731.

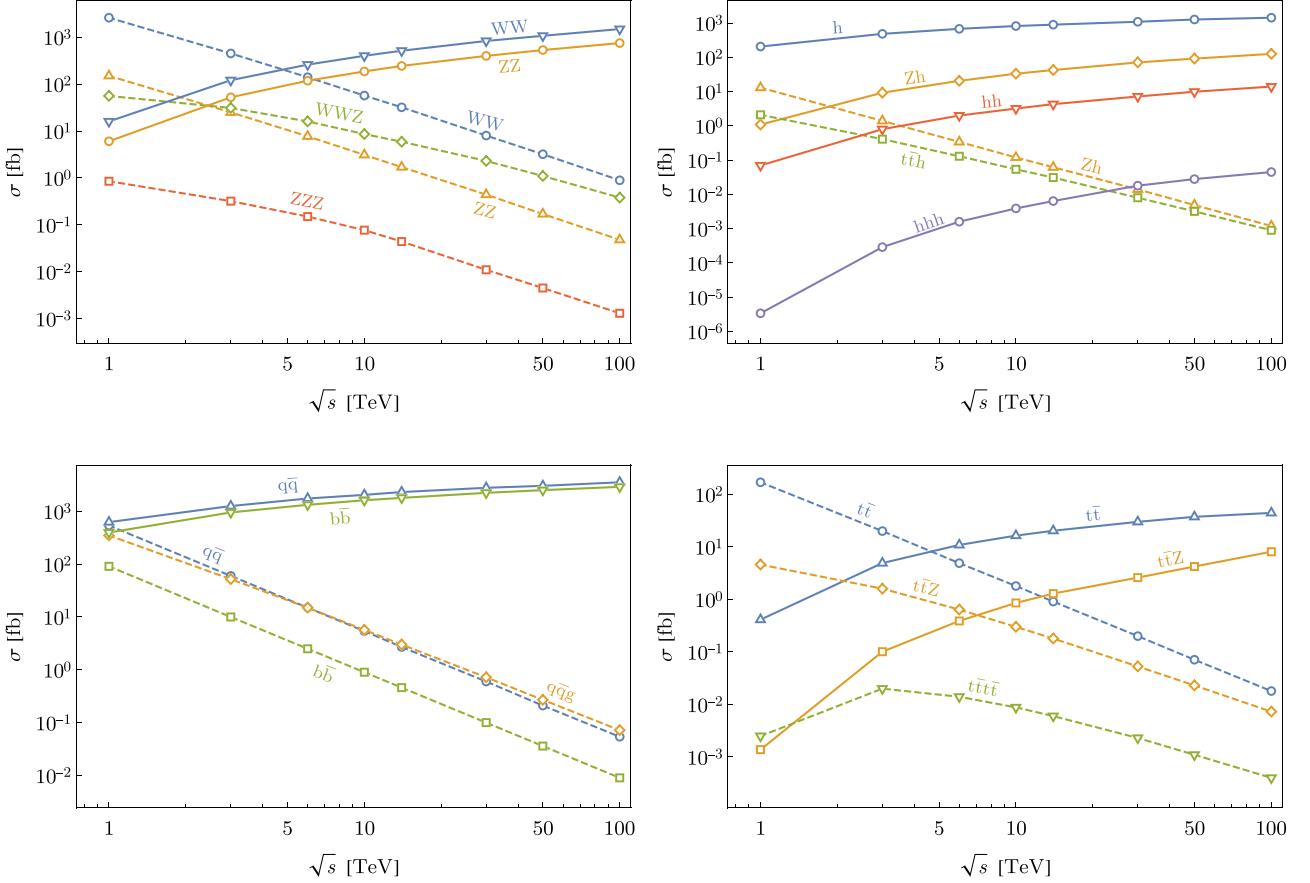


Figure 25. Cross sections for representative annihilation (dashed) and VBF (solid) SM processes as a function of \sqrt{s} .

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Appendix A. Simplified models

In order to benchmark the physics potential of muon colliders over a range of energies, we compute the rates for various processes using representative simplified models. Simplified models are defined in FEYNRULES using a combination of public and custom model files. Wherever possible, we simulate processes in both WHIZARD [221–223] and MADGRAPH5 [224]. We group the majority of processes into three categories: ‘annihilation’ when states are produced directly in $\mu^+\mu^-$ collisions; ‘neutral VBF’ when states are produced via VBF from combinations of electroweak vectors carrying zero total charge; and ‘charged VBF’ when states are produced via VBF from combinations of electroweak vectors carrying nonzero total charge. For the most part, we present results in annihilation and neutral VBF channels. For both charged and neutral VBF processes, we exclude contributions from on-shell W and Z bosons by imposing appropriate invariant mass cuts on the final state, as in [33]. In what follows, we typically do not display simulated VBF cross sections for mass points close to threshold given the strong

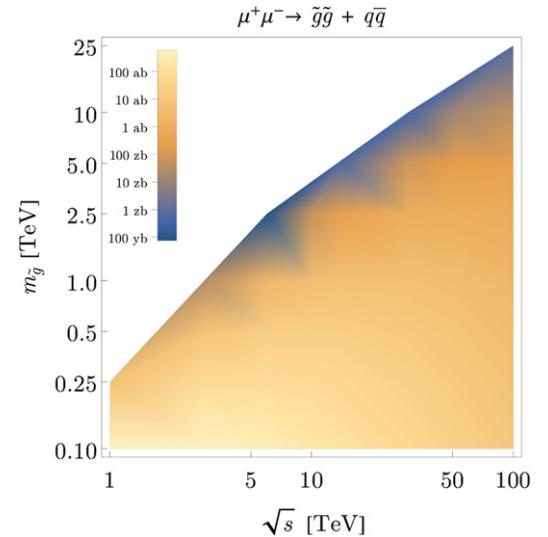


Figure 26. Cross section for $\tilde{g}\tilde{g}$ production resulting from the ‘bremsstrahlung’ of a gluon off of a $q\bar{q}$ pair as a function of $m_{\tilde{g}}$ and \sqrt{s} , computed using WHIZARD.

sensitivity to phase space cuts; as noted in the text, annihilation production dominates near threshold. Where relevant, we also consider QCD ‘bremsstrahlung’ processes, in which a gluon is radiated off a hard final state quark and then ‘splits’ into the strongly-interacting final state of interest; this is

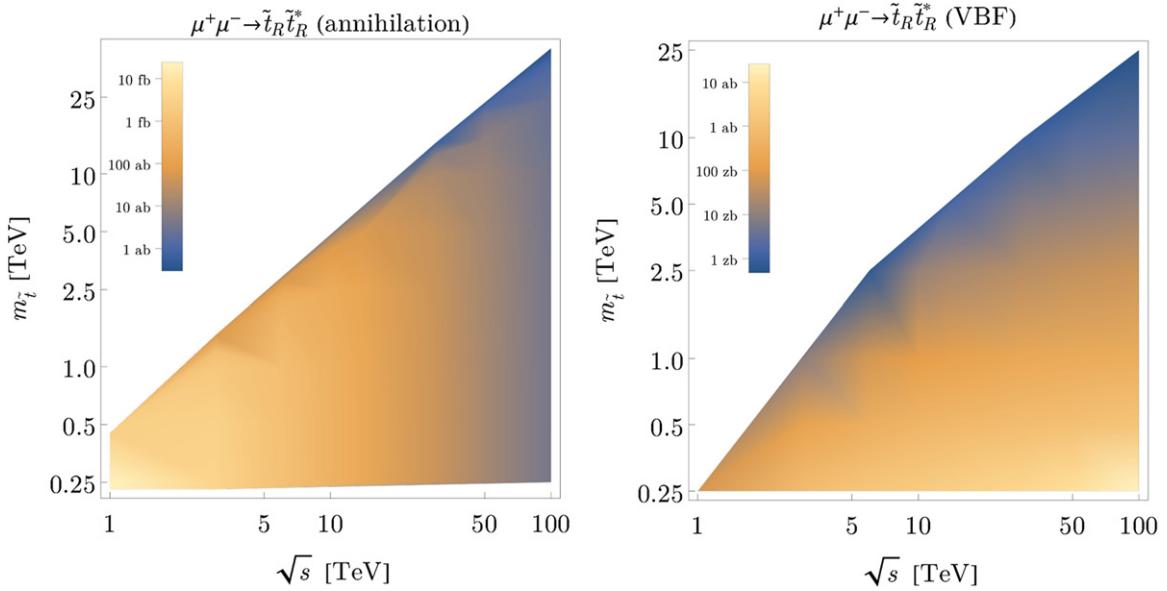


Figure 27. Cross section for $\tilde{t}_R \tilde{t}_R^*$ production via annihilation (left) and VBF (right) as a function of $m_{\tilde{t}}$ and \sqrt{s} , computed using WHIZARD.

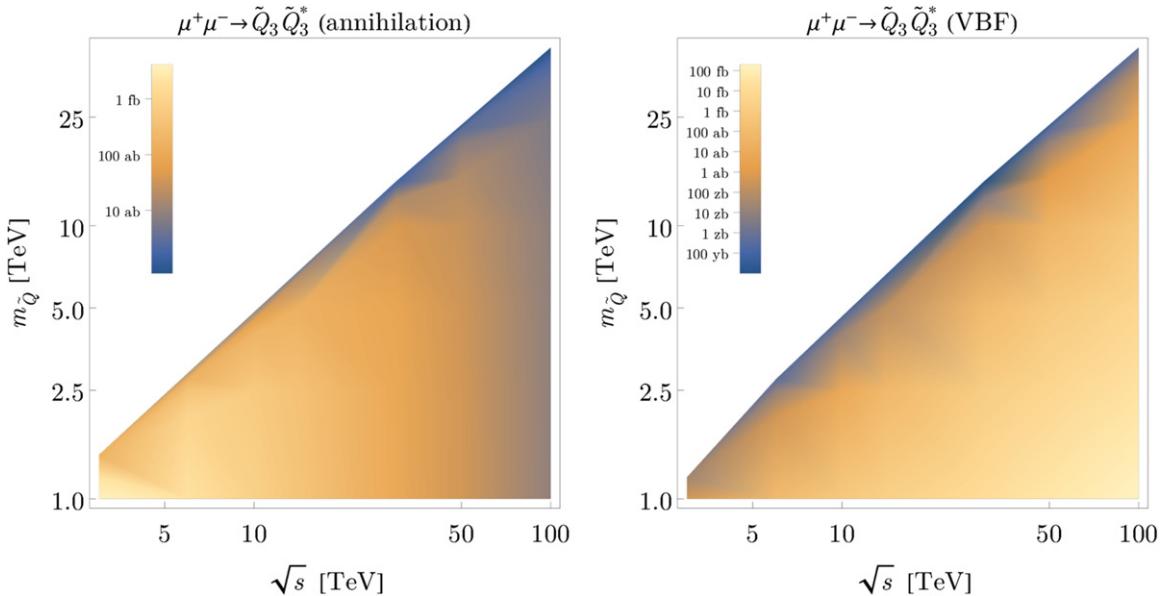


Figure 28. Cross section for $\tilde{Q}_3 \tilde{Q}_3^*$ production via annihilation (left) and VBF (right) as a function of $m_{\tilde{Q}}$ and \sqrt{s} , computed using WHIZARD.

particularly relevant for gluino production in supersymmetric models.

To validate the results of our simulations, we compute a selection of representative cross sections analytically using a combination of FEYNARTS, FORMCALC, and FEYNCALC. We compute select annihilation cross sections explicitly, while for select VBF cross sections we compute the ‘partonic’ cross sections explicitly and obtain total inclusive cross sections via numerical convolution with the LL luminosity functions detailed in section 3.

A.1. Standard model

We begin by exploring some SM rates as a function of \sqrt{s} . This provides some interesting benchmarks on its own, and also allows one to get a sense of the raw rate for backgrounds that are relevant when estimating the new physics potential of a muon collider.

For SM interactions, we treat the first and second generation fermions as massless and we use a diagonal CKM matrix for simplicity. We choose to work in terms of the

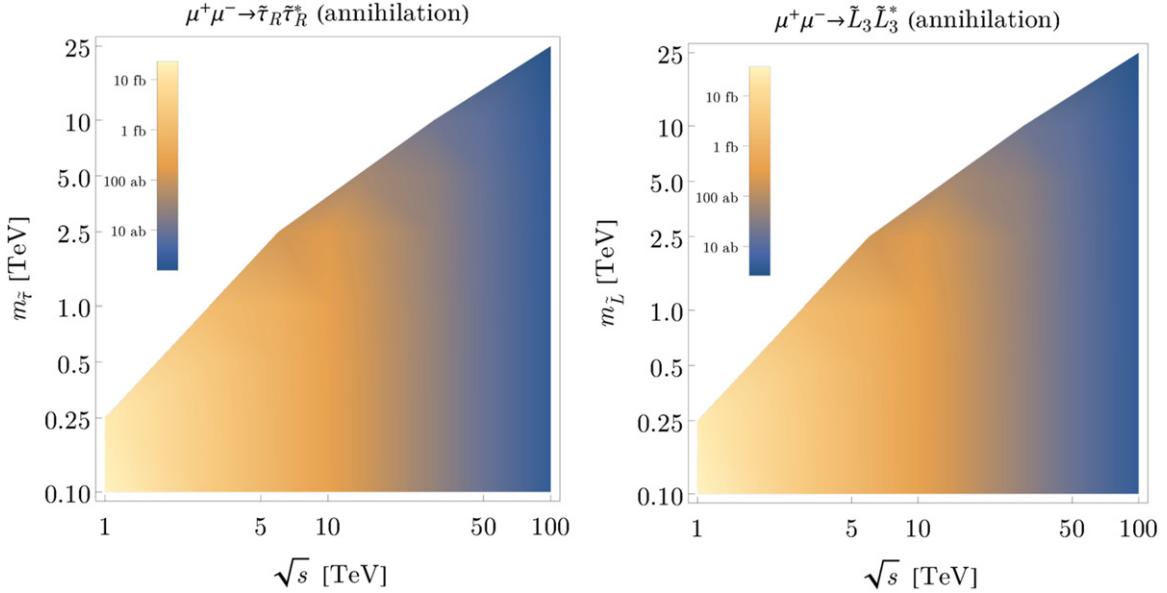


Figure 29. Cross section for $\tilde{\tau}_R \tilde{\tau}_R^*$ production (left) and $\tilde{L}_3 \tilde{L}_3^*$ production (right) via annihilation as a function of $m_{\tilde{\tau}}$ and \sqrt{s} , computed using WHIZARD.

measured value $\alpha(m_Z)^{-1} = 127.9$, at the cost of having $m_W = 79.8$ GeV, which is calculated using the tree-level relations in terms of G_F, α . We take $m_t = 173$ GeV and $m_h = 125$ GeV. Cross sections for representative processes computed using WHIZARD are illustrated in figure 25, and we find excellent agreement with MADGRAPH5. Unsurprisingly, as \sqrt{s} is increased the s -channel rates decrease as a power law, while the VBF processes grow logarithmically.

A.2. Supersymmetry

Perhaps the best studied simplified models are motivated by the MSSM. For supersymmetric processes, we use a variety of FEYNRULES models. For gluino pair production we use the default MSSM model in FEYNRULES, keeping the gluino light while decoupling all other sparticles. The cross section for $\tilde{g}\tilde{g}$ production resulting from the ‘bremsstrahlung’ of a gluon off of a $q\bar{q}$ pair is shown in figure 26. For stop pair production we use the t -channel dark matter FEYNRULES model [135] with couplings set to their supersymmetric values. Cross sections for the production of the $SU(2)$ singlet $\tilde{\tau}_R$ and $SU(2)$ doublet \tilde{Q}_3 via annihilation and VBF are shown in figures 27 and 28. For slepton pair production, we use the default MSSM model in FEYNRULES, alternately keeping either $\tilde{\tau}_R$ or \tilde{L}_3 light while decoupling all other sparticles. Cross sections for the production of the $SU(2)$ singlet $\tilde{\tau}_R$ and $SU(2)$ doublet \tilde{L}_3 via annihilation are shown in figure 29. For higgsino pair production, we use the default MSSM model in FEYNRULES, keeping the components χ_1^\pm and $\chi_{1,2}^0$ light while decoupling all other sparticles and taking care to keep the neutralino and chargino mixing matrices appropriately aligned. Cross sections for both chargino and neutralino pair production via annihilation and VBF are shown in figure 30. For wino pair production, we use the default MSSM model in FEYNRULES, keeping the components χ_1^\pm and χ_1^0 light while decoupling all other sparticles and

taking care to keep the neutralino and chargino mixing matrices appropriately aligned. Cross sections for both chargino and neutralino pair production via annihilation and VBF are shown in figure 31. Across all processes, we see the same trend as with the standard model processes: when both s -channel and VBF production is possible, the former dominates at low \sqrt{s} while the latter takes over for high \sqrt{s} . Another interesting point to note is that the squark production rate is significantly larger than the gluino rate, which implies that gluino production could be dramatically impacted when considering a more complete model that includes both gluinos and squarks.

A.3. Vector-like quarks

Next we turn to vector-like quarks, an important example that arises in scenarios ranging from global symmetry approaches to the hierarchy problem to parity solutions to the strong CP problem. We restrict our considerations to the gauge interactions of an $SU(3)$ triplet, $SU(2)$ singlet, hypercharge $+\frac{2}{3}$ Dirac fermion T , which we implement in FEYNRULES. Cross sections for $T\bar{T}$ production via annihilation and VBF are shown in figure 32, computed by convolving partonic amplitudes obtained via FEYNARTS and FORMCALC with the LL luminosity functions detailed in section 3.

A.4. Higgs portal

Next, we turn to one of the models that has implications for baryogenesis, dark matter, and models of neutral naturalness. Furthermore, this model is simultaneously very simple to state while being very difficult to discover. Specifically, we extend the SM using a \mathbb{Z}_2 symmetric singlet S that couples via the

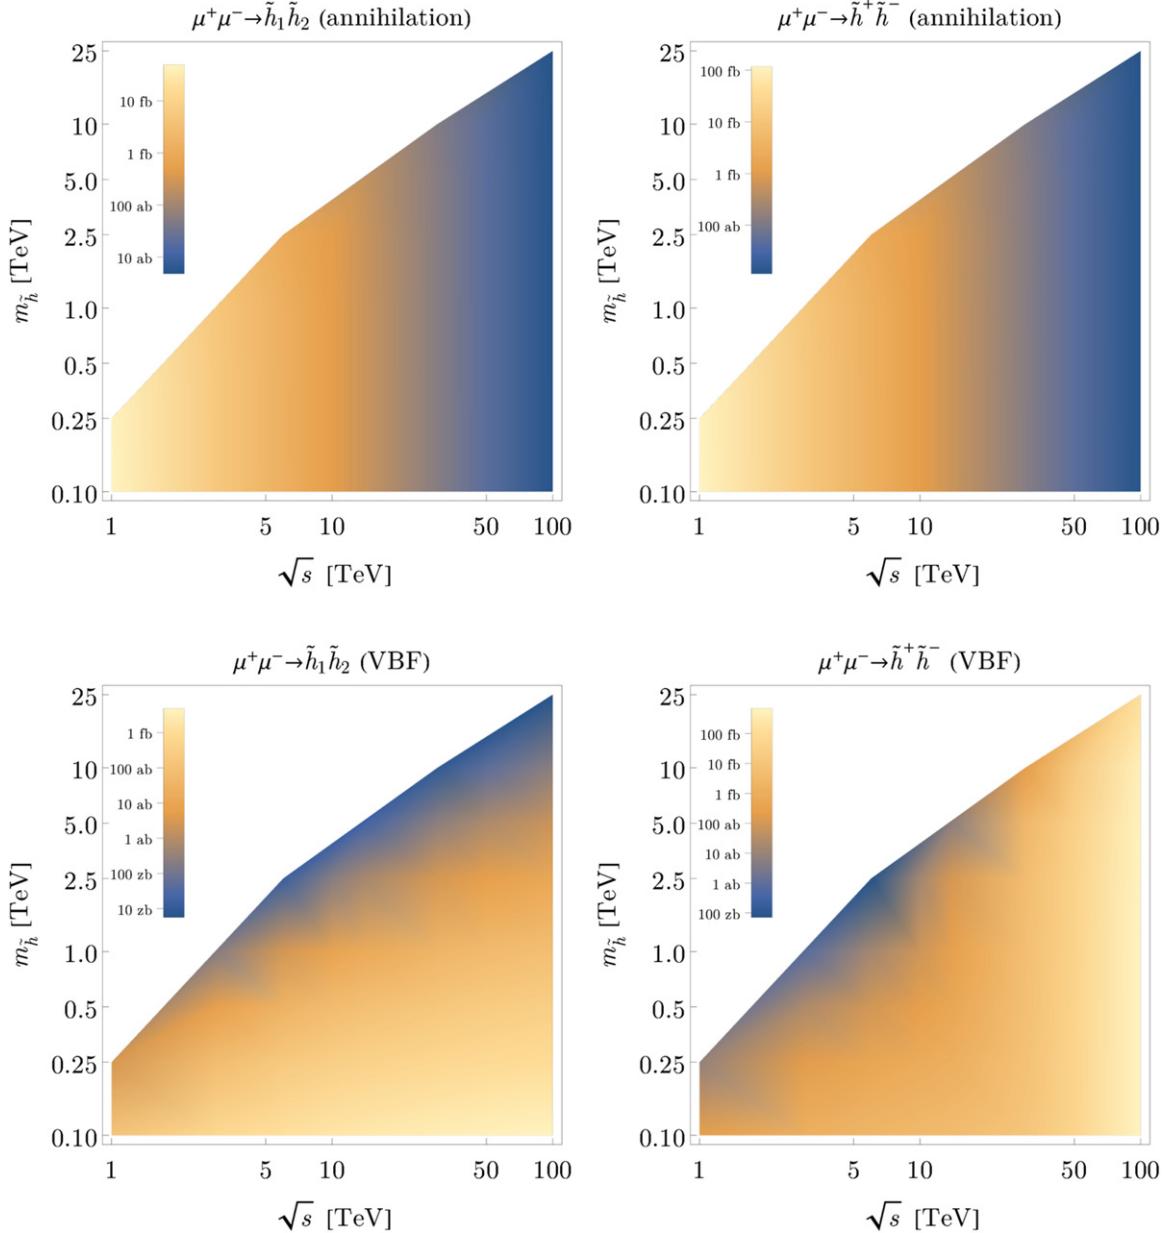


Figure 30. Cross section for $\tilde{h}_1^0\tilde{h}_2^0$ and $\tilde{h}^+\tilde{h}^-$ production as a function of $m_{\tilde{h}}$ and \sqrt{s} via annihilation (top) and VBF (bottom), computed using WHIZARD.

Higgs portal:

$$\mathcal{L} = \frac{1}{2}(\partial S)^2 - \frac{1}{2}M_S^2 S^2 - \lambda_{HS} S^2 |H|^2. \quad (67)$$

We omit a possible singlet quartic, which does not influence the phenomenology, and assume M_S, λ_{HS} are such that S is stabilized at the origin and the \mathbb{Z}_2 symmetry is unbroken. The physical mass of the singlet in the broken phase is

$$m_S^2 = M_S^2 + \lambda_{HS} v^2. \quad (68)$$

As written, the singlet is absolutely stable and may constitute a dark matter candidate. It also may be rendered unstable by explicit \mathbb{Z}_2 breaking, leading to prompt or long-lived signatures.

The differential cross section for s -channel production of SS in association with a Z boson is [225]

$$\begin{aligned} \frac{d\sigma}{dt du}(\mu^+\mu^- \rightarrow ZSS) = & \frac{\left[\left(-\frac{1}{2} + s_W^2 \right)^2 + s_W^4 \right] \lambda_{HS}^2 v^2 m_Z^2}{16(2\pi)^3 s} \\ & \times \left(\frac{e}{s_W c_W} \right)^4 \left[\frac{s + (t - m_Z^2)(u - m_Z^2)/m_Z^2}{(s - m_Z^2)^2 + \Gamma_Z^2 m_Z^2} \right] \\ & \times \left[\frac{1}{s + t + u - m_Z^2 - m_h^2} \right]^2 \\ & \times \left[1 - \frac{4m_S^2}{s + t + u - m_Z^2} \right]^{1/2}, \end{aligned} \quad (69)$$

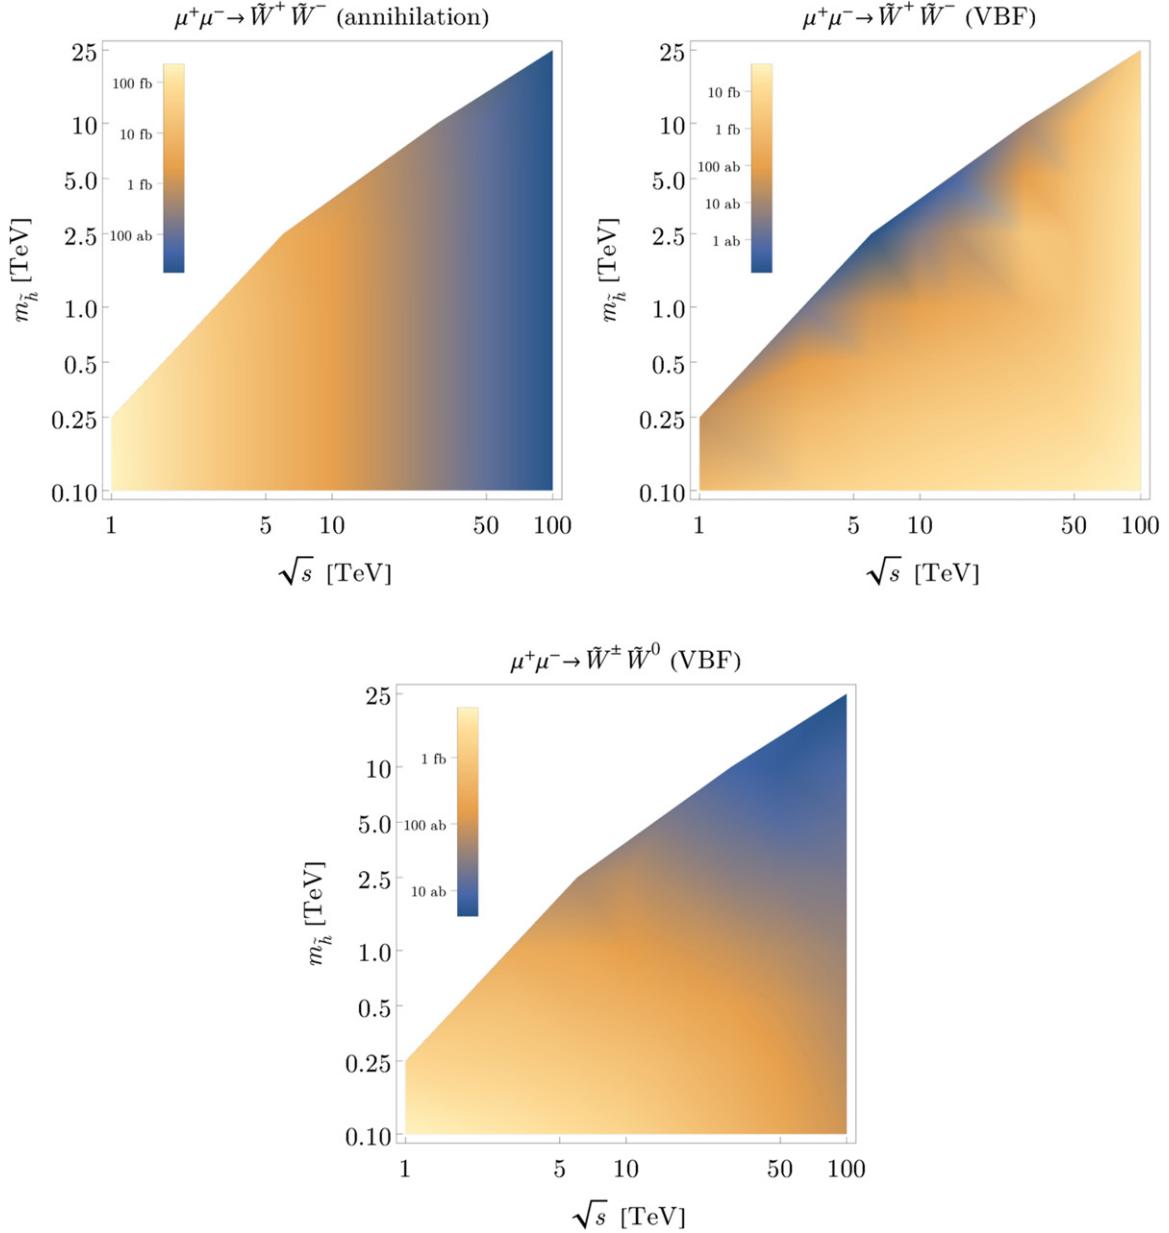


Figure 31. Cross section for $\tilde{W}^+ \tilde{W}^-$ production via annihilation (top left) and VBF (top right), $\tilde{W}^\pm \tilde{W}^0$ production via VBF (bottom) as a function of $m_{\tilde{W}}$ and \sqrt{s} , computed using **MADGRAPH5**.

where the Mandelstam variables are defined in the standard way for $2 \rightarrow 2$ processes such that $s + t + u = m_Z^2 + Q^2$, where Q^2 is the square of the invariant mass of the two singlets in the final state. At large s , this leads to a cross section falling as $1/s^2$. This motivates considering a variety of VBF-like processes, including $W^+ W^-$ fusion into SS and HSS final states, as well as $WZ/W\gamma$ fusion into the WSS final state. We implement the model in **FEYNRULES** and compute cross sections primarily in **WHIZARD**, finding good agreement with both **MADGRAPH5** and the result of convolving the partonic cross sections with our LL PDFs. Cross sections for all four processes as a function of m_S and \sqrt{s} are shown in figure 33.

A.4.1. Missing mass analysis. Given the central importance of the singlet scalar Higgs portal in BSM scenarios ranging from electroweak baryogenesis to neutral naturalness, we carry out a series of simplified analyses to benchmark the reach of dedicated searches for S , assuming that S is detector stable. These analyses cover three different production modes: $W^+ W^-$ fusion into an SS pair in association with an ISR photon; $W^+ W^-$ fusion into an SS pair in association with a Higgs boson; and s -channel production of an SS pair in association with a Z boson. In each case and for each benchmark value of \sqrt{s} , we simulate $100k$ background events and $10k$ signal events for each of $m_S = 75, 100, 150, 225, 350, 500, 750, 1000$ GeV in **WHIZARD**.

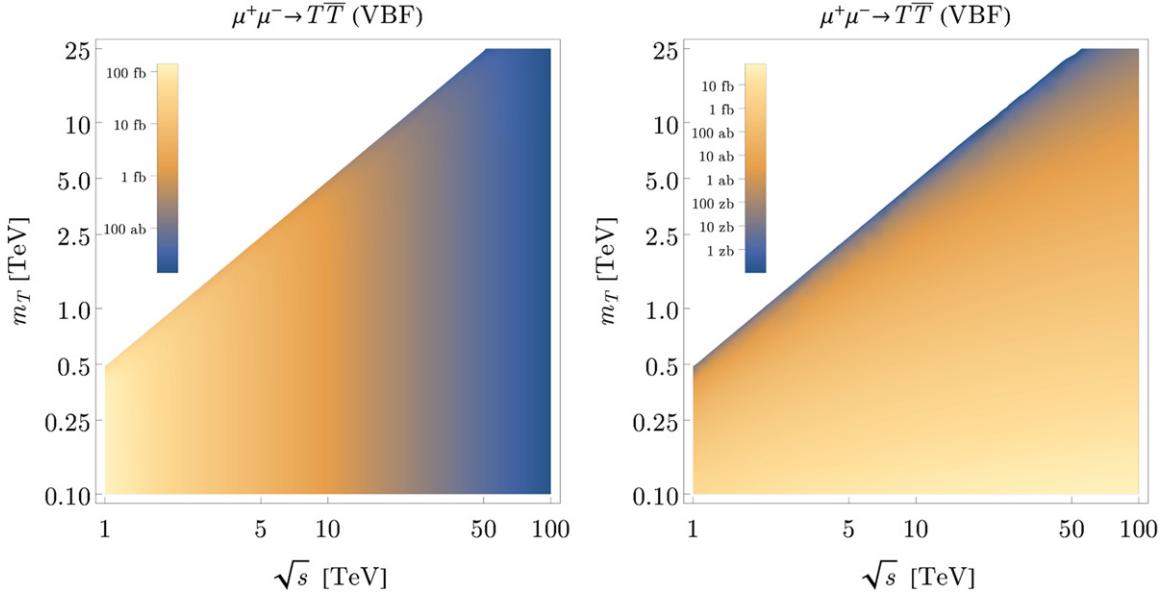


Figure 32. Cross section for $T\bar{T}$ production via annihilation (left) and VBF (right) as a function of m_T and \sqrt{s} , computed by convolving partonic amplitudes obtained via FEYNARTS and FORMCALC with the LL luminosity functions detailed in section 3.

SS+ γ analysis. For this signal, we consider the production of SS via W^+W^- fusion accompanied by an ISR photon, $\mu^+\mu^- \rightarrow SS + \nu\bar{\nu} + \gamma$, where S is treated as invisible. The primary SM background for this final state is $\mu^+\mu^- \rightarrow \nu\bar{\nu} + \gamma$, i.e., production of a single photon via W^+W^- fusion. We require that the ISR photon falls within the detector acceptance, here defined as $10^\circ < \theta < 170^\circ$, and employ a ‘missing mass’ strategy in which the photon energy is required to lie between

$$50 \text{ GeV} < E_\gamma < \frac{s - 4m_S^2}{2\sqrt{s}}. \quad (70)$$

The background spectrum is harder than the signal spectrum, so the lower bound on the photon energy serves primarily to avoid singular regions of phase space. The upper bound on the photon energy reflects the lower bound on the missing mass distribution.

The integrated luminosity required to exclude a singlet scalar at 95% CL using this analysis is shown in figure 34.

SS+ h analysis. For this signal, we consider the production of hSS via W^+W^- fusion, $\mu^+\mu^- \rightarrow SS + h + \nu\bar{\nu}$, where S is treated as invisible. The primary SM background for this final state is $\mu^+\mu^- \rightarrow \nu\bar{\nu} + h$, i.e., production of a single Higgs via W^+W^- fusion. We require the Higgs boson falls within the detector acceptance, here defined as $10^\circ < \theta < 170^\circ$, and employ a ‘missing mass’ strategy in which the Higgs energy is required to lie between

$$E_{\min} < E_h < \frac{s - 4m_S^2 + m_h^2}{2\sqrt{s}}. \quad (71)$$

The signal spectrum is harder than the background spectrum, so the lower bound E_{\min} on the Higgs energy is chosen for each value of m_S to maximize S/\sqrt{B} . The upper bound on the Higgs energy reflects the lower bound on the missing mass distribution.

The integrated luminosity required to exclude a singlet scalar at 95% CL using this analysis is shown in figure 35.

SS+ Z analysis. For this signal, we consider s -channel production of $ZSS\mu^+\mu^- \rightarrow SS + Z$, where S is treated as invisible. The primary SM background for this final state is $\mu^+\mu^- \rightarrow Z + \nu\bar{\nu}$, which has contributions from both W^+W^- fusion and $\mu^+\mu^-$ annihilation. We require that the Z boson falls within the detector acceptance, here defined as $10^\circ < \theta < 170^\circ$, and employ a ‘missing mass’ strategy in which the Z energy is required to lie between

$$E_{\min} < E_Z < \frac{s - 4m_S^2 + m_Z^2}{2\sqrt{s}}. \quad (72)$$

The signal spectrum is harder than the background spectrum, so the lower bound E_{\min} on the Z energy is chosen for each value of m_S to maximize S/\sqrt{B} . The upper bound on the Z energy reflects the lower bound on the missing mass distribution. At higher energies, the optimal value of E_{\min} leads to background acceptance of less than 10^{-4} for lower values of m_S , in which case we generate additional exclusive background samples to improve Monte Carlo statistical error.

The integrated luminosity required to exclude a singlet scalar at 95% CL using this analysis is shown in figure 36.

A.4.1.4. Missing mass combination To benchmark the combined sensitivity of these three searches, we consider the naive combination that results from adding their significance in quadrature. The integrated luminosity required to exclude a singlet scalar at 95% CL using the combination of analyses is shown in figure 37.

A.4.2. Invariant mass analysis. An alternative strategy to probe the Higgs portal singlet scenario is to search for the

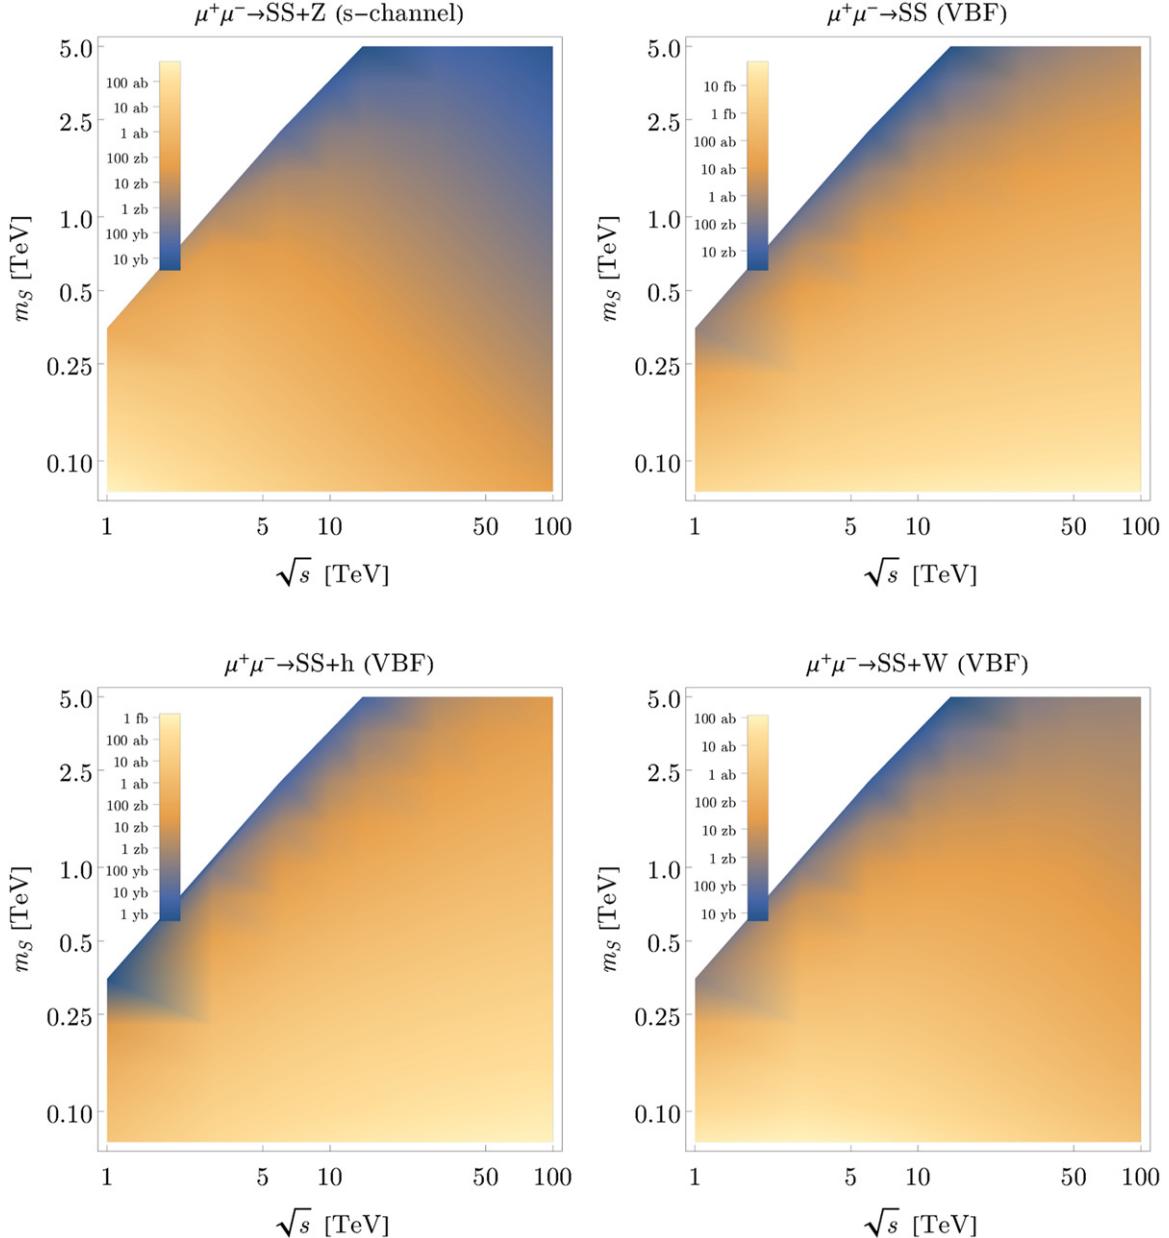


Figure 33. Cross section for various Higgs portal production modes as a function of m_S and \sqrt{s} computed using **WHIZARD**, including s -channel $SS + Z$ production (upper left), VBF SS production (upper right), VBF $SS + h$ production (lower left), and VBF $SS + W$ production (lower right). As a benchmark, we take $\lambda_{HS} = 1$ here.

virtual effects of S through its influence on the Higgs sector, rather than looking for this state on-shell as described in appendix A.4.1. At one-loop order, the simplified model of equation (67) includes modifications in both the Higgs two-point function, as well as the couplings between the Higgs and SM degrees of freedom. As already noted in section 4.1, this kind of tampering in the Higgs sector will affect the behavior of scattering amplitudes involving longitudinally polarized gauge bosons, given the role of the Higgs in maintaining perturbative unitarity of the SM. In this section, we illustrate how the resulting features in the differential cross section for the process $\mu^+\mu^- \rightarrow \bar{t}t + X$ may be leveraged to probe this model.

The one-loop correction to the Higgs two-point function in the model of equation (67) includes both mass and wavefunction renormalization effects, and it is given by

$$\Sigma_2(p^2) = \frac{(\lambda_{HS}v)^2}{16\pi^2} \int_0^1 dx \log \left(\frac{m_S^2 - p^2 x(1-x)}{m_S^2 - m_h^2 x(1-x)} \right) - (p^2 - m_h^2) \delta Z, \quad (73)$$

where

$$\delta Z = -\frac{(\lambda_{HS}v)^2}{16\pi^2} \int_0^1 dx \frac{x(1-x)}{m_S^2 - m_h^2 x(1-x)}, \quad (74)$$

and we have imposed renormalization conditions such that the Higgs propagator has a single pole at $p^2 = m_h^2$ with residue i .

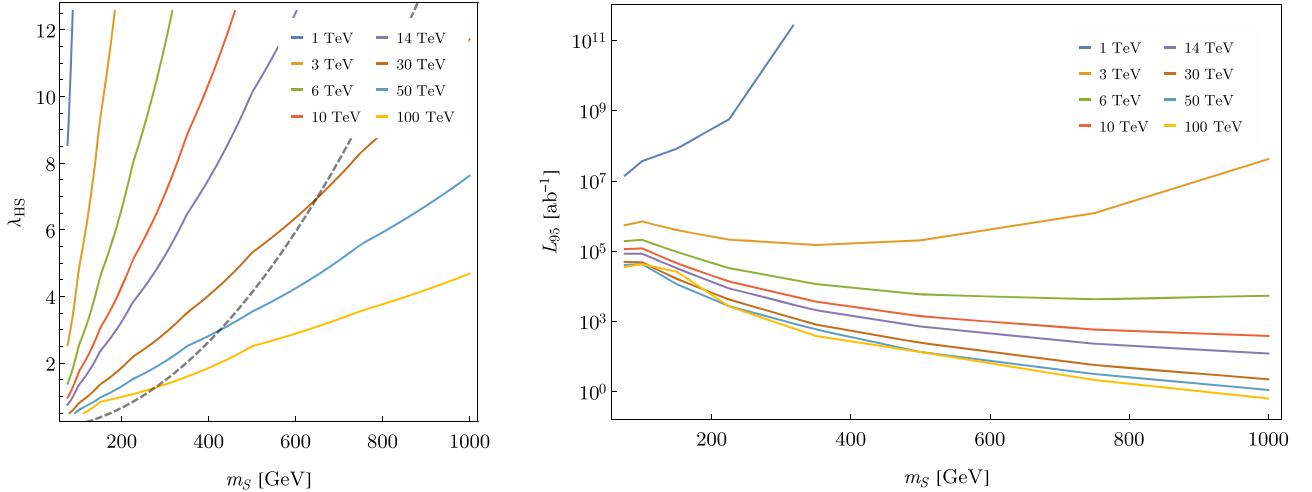


Figure 34. Left: projected 95% CL exclusion limits as a function of m_S and λ_{HS} from the $SS + \gamma$ final state for various energy and luminosity benchmarks described in the text. The dashed line corresponds to a singlet scalar acquiring mass entirely from EWSB. Right: the integrated luminosity \mathcal{L} in units of ab^{-1} required to exclude a singlet scalar whose mass m_S is due entirely to EWSB at 95% CL in the $SS + \gamma$ final state.

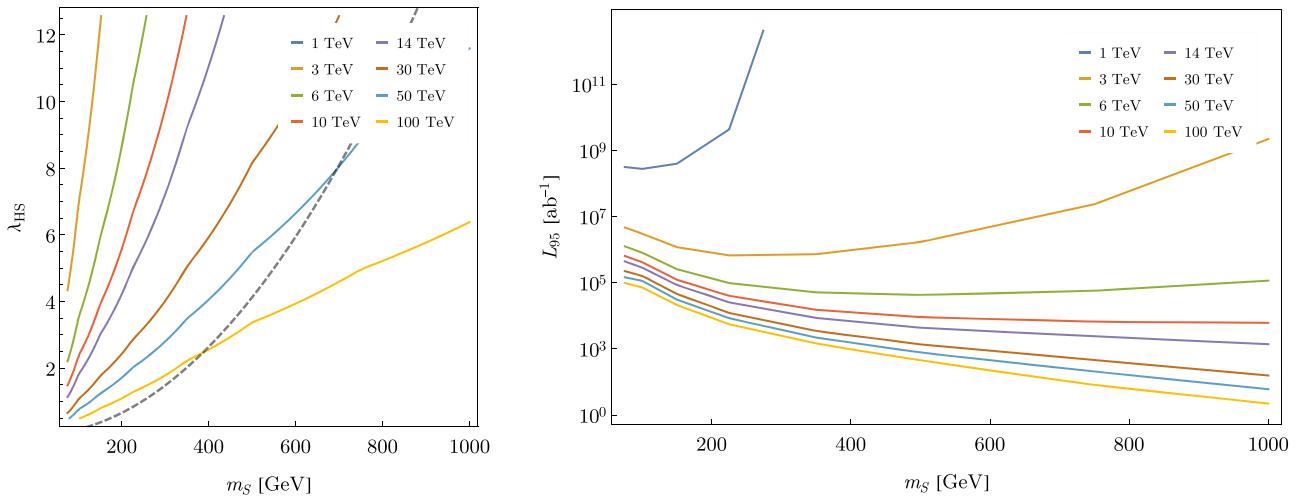


Figure 35. Left: projected 95% CL exclusion limits as a function of m_S and λ_{HS} from the $SS + h$ final state for various energy and luminosity benchmarks described in the text. The dashed line corresponds to a singlet scalar acquiring mass entirely from EWSB. Right: the integrated luminosity \mathcal{L} in units of ab^{-1} required to exclude a singlet scalar whose mass m_S is due entirely to EWSB at 95% CL in the $SS + h$ final state.

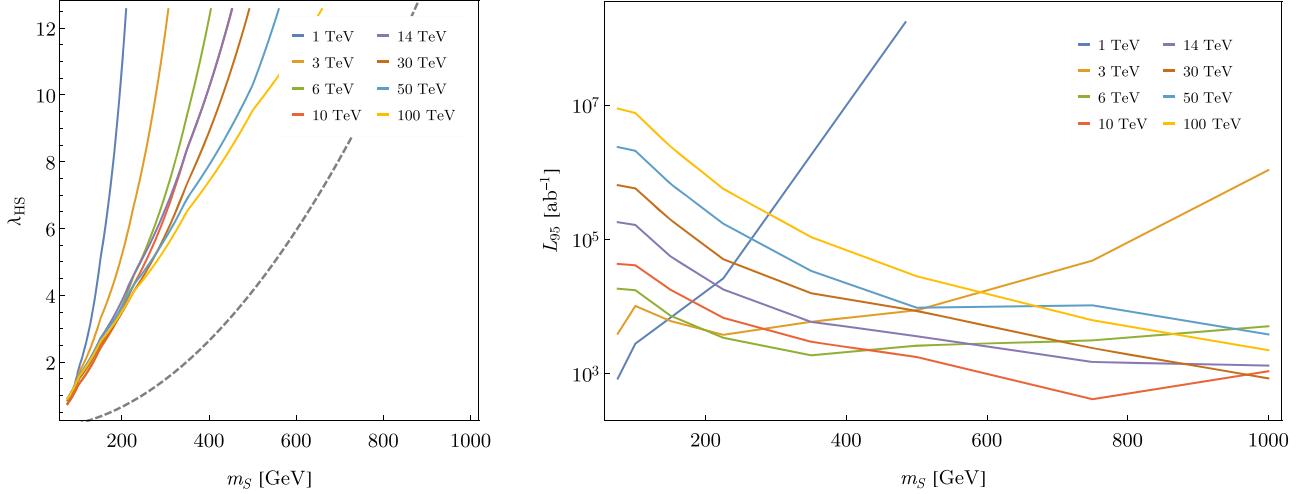


Figure 36. Left: projected 95% CL exclusion limits as a function of m_S and λ_{HS} from the $SS + Z$ final state for various energy and luminosity benchmarks described in the text. The dashed line corresponds to a singlet scalar acquiring mass entirely from electroweak symmetry breaking. Right: the integrated luminosity \mathcal{L} in units of ab^{-1} required to exclude a singlet scalar whose mass m_S is due entirely to electroweak symmetry breaking at 95% CL in the $SS + Z$ final state.

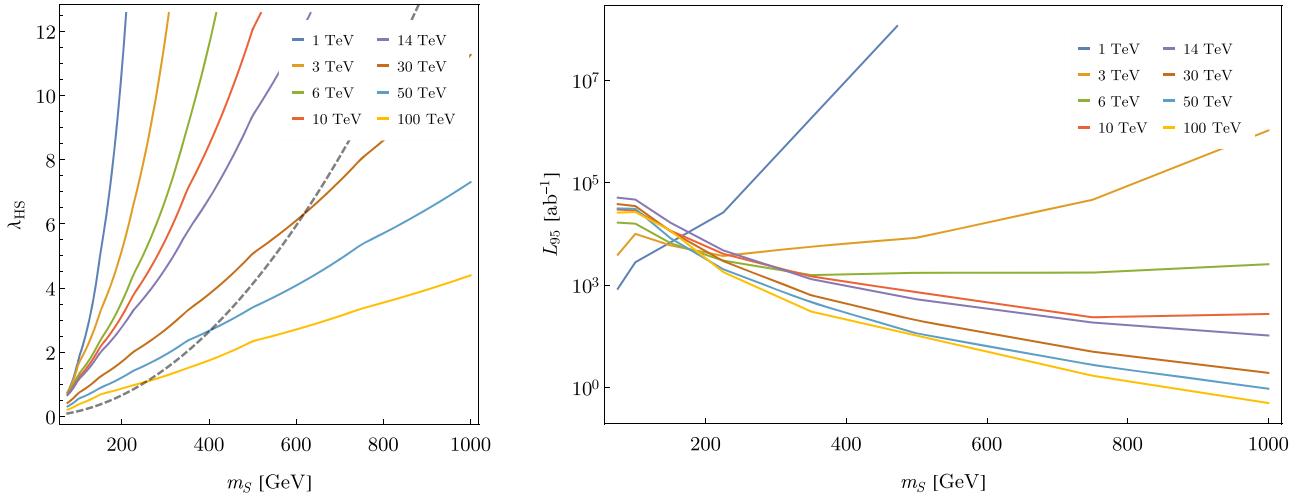


Figure 37. Left: projected 95% CL exclusion limits as a function of m_S and λ_{HS} from the naive combination of $SS + \gamma/h/Z$ final states for various energy and luminosity benchmarks described in the text. The dashed line corresponds to a singlet scalar acquiring mass entirely from electroweak symmetry breaking. Right: the integrated luminosity \mathcal{L} in units of ab^{-1} required to exclude a singlet scalar whose mass m_S is due entirely to electroweak symmetry breaking at 95% CL from the naive combination of $SS + \gamma/h/Z$ final states.

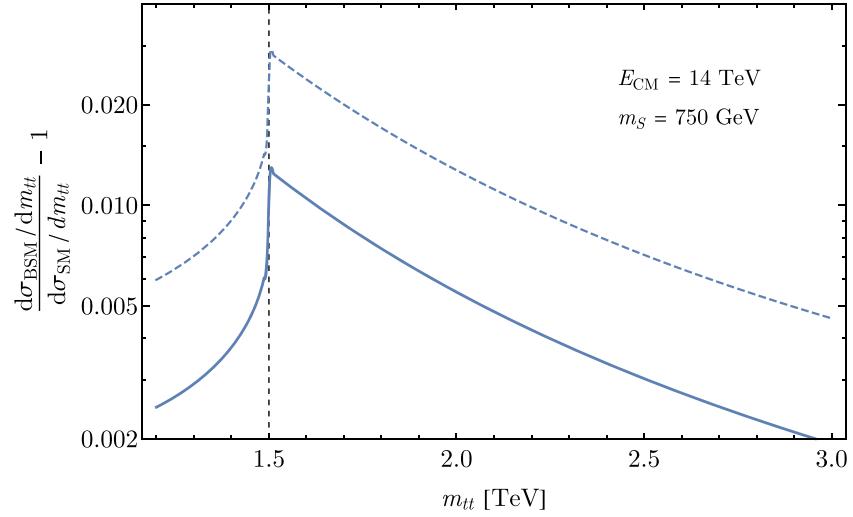


Figure 38. Fractional deviation in the differential cross section for the process $\mu^+\mu^- \rightarrow \bar{t}t + X$ as a result of the radiative corrections present in the model of equation (67) (see equations (73) and (74) and surrounding discussion). The c.m.-energy of the colliding muons is taken to be 14 TeV, and the mass of the singlet $m_S = 750$ GeV, all of which arises from electroweak symmetry breaking. The solid line includes both charged and neutral intermediate VBF processes, whereas the dashed line includes only contributions from W^+W^- fusion, as would be appropriate if background events could be eliminated by identifying the outgoing muons. The vertical dashed line corresponds to $m_{\bar{t}t} = 2m_S = 1.5$ TeV, at which the peak of the kinematic feature in the differential cross section takes place.

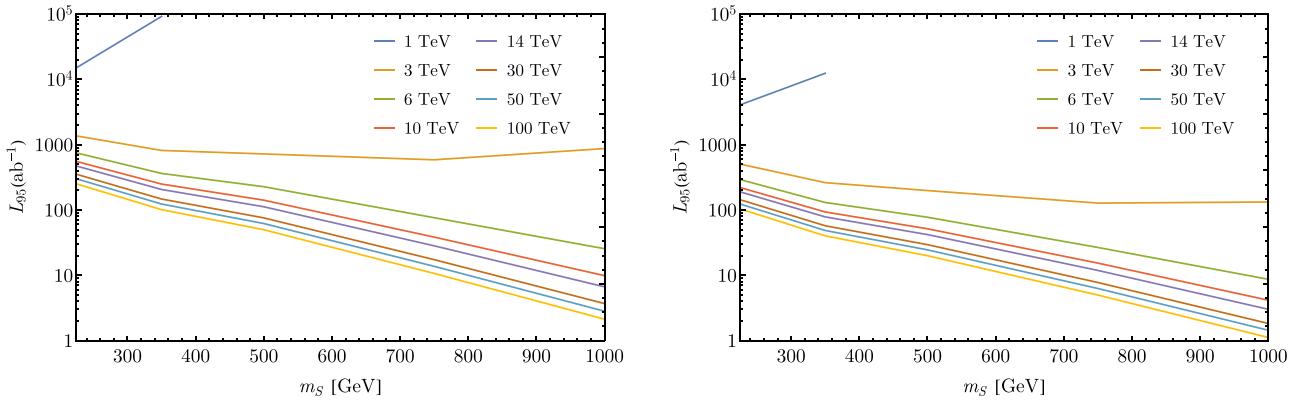


Figure 39. Estimate of the integrated luminosity required to exclude a Higgs portal singlet S whose mass arises entirely from electroweak symmetry breaking, for various c.m. energies. Left: $\bar{t}t$ production through intermediate neutral and charged VBF is treated as background. Right: only $\bar{t}t$ production through charged VBF is treated as background, under the assumption that the outgoing muons accompanying neutral VBF can be identified.

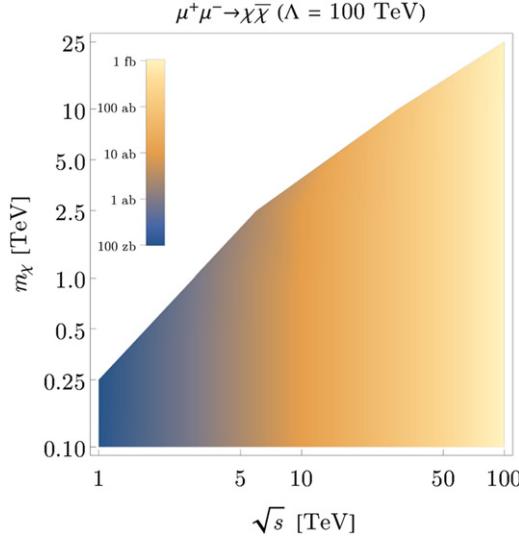


Figure 40. Cross section for $\chi\bar{\chi}$ production in the hidden valley simplified model as a function of m_χ and \sqrt{s} for $\Lambda = 100$ TeV and $c = 1$, computed using MADGRAPH5.

Wavefunction renormalization further affects SM Higgs couplings, which are modified by a universal factor of $(1 + \delta_*)$, with $\delta_* \simeq \delta Z/2$.

At the partonic level, the processes most sensitive to these modifications in the Higgs sector are those involving longitudinally polarized W and Z bosons in the initial state, with $W_L^+ W_L^- \rightarrow t\bar{t}$ being the most relevant given the enhanced W content of the muons. This results in a distinctive kinematic feature in the differential cross section for the process $\mu^+ \mu^- \rightarrow t\bar{t} + X$ that peaks at a scale $m_{t\bar{t}} \simeq 2m_S$, where $m_{t\bar{t}}$ is the invariant mass of the $t\bar{t}$ pair. For illustration, figure 38 shows the fractional deviation in $d\sigma/dm_{t\bar{t}}$ with respect to the SM for a collision c.m. energy $E_{\text{CM}} = 14$ TeV, and a singlet mass $m_S = 750$ GeV that results entirely from EWSB. The maximum size of the deviation can be of order $\sim 1\%$. Our calculation includes the one-loop effects described in the previous paragraph, but is otherwise performed at tree-level. This approximation is justified as long as λ_{HS} is larger than any of the couplings of the SM, which is always the case when the singlet mass arises entirely from its coupling to the Higgs and provided $m_S > m_t$.

An estimate of the integrated luminosity required to exclude a singlet whose mass is due entirely to EWSB at 95% CL using this analysis is shown in figure 39. The number of signal and background events have been computed by integrating the differential cross section over a mass window spanning the range between 90% and 150% of $2m_S$. The corresponding

S/\sqrt{B} ratio has been computed for $m_S = 225, 350, 500, 750$, and 1000 GeV and linearly interpolated for intermediate masses.³⁷

Although a more sophisticated analysis would be required to draw robust conclusions, the results in figure 39 suggest that this search strategy could improve on the individual channels described in appendix A.4.1, and may be competitive with their combination. The analysis could be further improved by considerations related to the angular distribution of the outgoing $t\bar{t}$ pair, or through the use of polarized beams. More importantly, our discussion illustrates how muon colliders offer qualitatively new possibilities to search for new physics, by taking advantage of both their high energy reach as well as their underlying identity as gauge boson colliders.

A.5. Hidden valleys

To benchmark the sensitivity of muon colliders to hidden valley scenarios [170, 171], we consider a particularly minimal realization in which the vector current of a SM-neutral Dirac fermion χ is coupled to the muon vector current of the standard model via a dimension-6 contact term:

$$\mathcal{L} = i\bar{\chi}D\chi - m_\chi\bar{\chi}\chi + \frac{c}{\Lambda^2}(\bar{\mu}\gamma^\mu\mu)(\bar{\chi}\gamma_\mu\chi). \quad (75)$$

We implement this model in FEYNRULES and compute the cross section using MADGRAPH5. The cross section for $\chi\bar{\chi}$ production as a function of m_χ and \sqrt{s} is shown in figure 40 for $\Lambda = 100$ TeV and $c = 1$.

A.6. Axion-like particles

As our final example, we consider a simplified model in which an axion-like particle a couples to electroweak field strengths. We use the ALPSeFT FEYNRULES model file documented in [226], taking the Lagrangian to be

$$\mathcal{L} \supset \frac{1}{2}(\partial a)^2 - \frac{1}{2}m_a^2a^2 - \frac{c_{\tilde{B}}}{f_a}aB_{\mu\nu}\tilde{B}^{\mu\nu} - \frac{c_{\tilde{W}}}{f_a}aW_{\mu\nu}^a\tilde{W}^{a,\mu\nu}. \quad (76)$$

Cross sections for VBF production of a as a function of m_a and \sqrt{s} are shown in figure 41 for $f_a = 100$ TeV and $(c_{\tilde{B}}, c_{\tilde{W}}) = (1, 0)$ [left] and $(0, 1)$ [right].

³⁷ For a center-of-mass energy of the incoming muons $E_{\text{CM}} = 1$ TeV, we only estimate the required luminosity for $m_S = 225$ and 350 GeV, and linearly interpolate between the two. For the larger values of m_S , it is not possible for the deviation in the differential cross section to reach its peak at $m_{t\bar{t}} = 2m_S$, as the singlet never becomes on-shell, and therefore the method discussed here is not applicable.

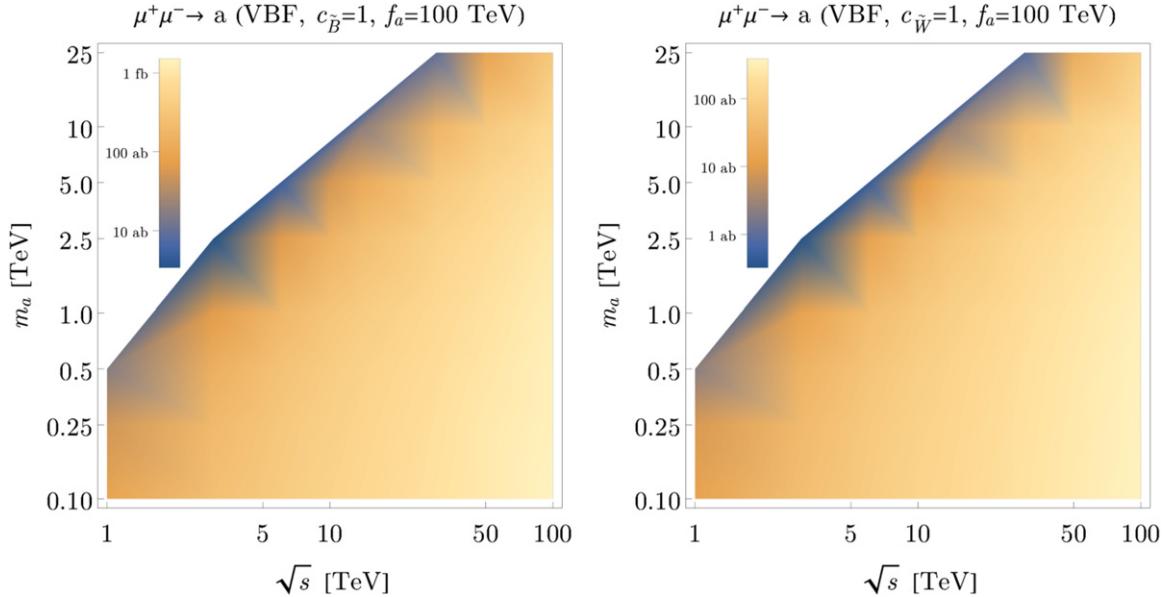


Figure 41. Cross section for VBF production of a in the axion-like particle simplified model as a function of m_a and \sqrt{s} for $f_a = 100$ TeV, computed using **MADGRAPH5**. Left: $c_{\tilde{B}} = 1$, $c_{\tilde{W}} = 0$. Right: $c_{\tilde{B}} = 0$, $c_{\tilde{W}} = 1$.

ORCID iDs

Tianji Cai <https://orcid.org/0000-0002-3235-9486>
 Nathaniel Craig <https://orcid.org/0000-0003-2214-8447>
 Matthew Forslund <https://orcid.org/0000-0003-0144-9339>
 Samuel Homiller <https://orcid.org/0000-0002-0063-6856>

References

- [1] ATLAS Collaboration Aad G *et al* 2012 Observation of a new particle in the search for the standard model Higgs boson with the ATLAS detector at the LHC *Phys. Lett. B* **716** 1–29
- [2] CMS Collaboration Chatrchyan S *et al* 2012 Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC *Phys. Lett. B* **716** 30–61
- [3] Bambade P *et al* 2019 The international linear collider: a global project (arXiv:1903.01629 [hep-ex])
- [4] CLICdp, CLIC Collaboration Charles T *et al* 2018 The compact linear collider (CLIC)—2018 summary report *CERN Yellow Rep. Monogr.* **1802** 1–98
- [5] FCC Collaboration Abada A *et al* 2019 FCC-ee: the lepton collider: future circular collider conceptual design report volume 2 *Eur. Phys. J. ST* **228** 261–623
- [6] CEPC Study Group Collaboration Dong M *et al* 2018 CEPC conceptual design report: volume 2—physics & detector (arXiv:1811.10545 [hep-ex])
- [7] FCC Collaboration Abada A *et al* 2019 FCC-hh: the hadron collider: future circular collider conceptual design report volume 3 *Eur. Phys. J. ST* **228** 755–1107
- [8] Ahmad M *et al* 2015 CEPC-SPPC preliminary conceptual design report: 1. Physics and detector *IHEP-CEPC-DR-2015-01, IHEP-TH-2015-01, IHEP-EP-2015-01*
- [9] Budker G I 1969 Accelerators and colliding beams *Conf. Proc. C* **690827** 33–9
- [10] Parkhomchuk V V and Skrinsky A N 1996 Ionization cooling: physics and applications *AIP Conf. Proc.* **352** 7–9
- [11] Neuffer D 1983 Principles and applications of muon cooling *Part. Accel.* **14** 75–90
- [12] Neuffer D 1987 Multi-TeV muon colliders *AIP Conf. Proc.* **156** 201–8
- [13] Barger V, Berger M S, Gunion J F and Han T 1995 s -channel Higgs boson production at a muon–muon collider *Phys. Rev. Lett.* **75** 1462–5
- [14] Barger V, Berger M S, Gunion J F and Han T 1997 Higgs boson physics in the s -channel at $\mu + \mu$ -colliders *Phys. Rep.* **286** 1–51
- [15] Ankenbrandt C M *et al* 1999 Status of muon collider research and development and future plans *Phys. Rev. ST Accel. Beams* **2** 081001
- [16] Delahaye J-P *et al* 2013 Enabling intensity and energy frontier science with a muon accelerator facility in the US: a white paper submitted to the 2013 US community summer study of the division of particles and fields of the American physical Society *Community Summer Study 2013: Snowmass on the Mississippi* vol 8 (arXiv:1308.0494 [physics.acc-ph])
- [17] Delahaye J-P *et al* 2014 A staged muon accelerator facility for neutrino and collider physics *5th Int. Particle Accelerator Conf., WEZA02* vol 6 (arXiv:1502.01647 [physics.acc-ph])
- [18] Ryne R D *et al* 2015 Design concepts for muon-based accelerators *6th Int. Particle Accelerator Conf.*
- [19] Long K, Lucchesi D, Palmer M, Pastrone N, Schulte D and Shiltsev V 2020 Muon colliders to expand frontiers of particle physics (arXiv:2007.15684 [physics.acc-ph])
- [20] MICE Collaboration Mohayai T A 2018 First demonstration of ionization cooling in MICE *9th Int. Particle Accelerator Conf.*
- [21] MICE Collaboration Blackmore V 2018 Recent results from the study of emittance evolution in MICE *9th Int. Particle Accelerator Conf.*
- [22] MICE Collaboration Bogomilov M *et al* 2020 Demonstration of cooling by the muon ionization cooling experiment *Nature* **578** 53–9
- [23] Antonelli M, Boscolo M, Di Nardo R and Raimondi P 2016 Novel proposal for a low emittance muon beam using positron beam on target *Nucl. Instrum. Methods Phys. Res. A* **807** 101–7
- [24] Delahaye J P, Diemoz M, Long K, Mansoulié B, Pastrone N, Rivkin L, Schulte D, Skrinsky A and Wulzer A 2019 Muon colliders (arXiv:1901.06150 [physics.acc-ph])

[25] Han T, Ma Y and Xie K 2021 High energy leptonic collisions and electroweak parton distribution functions *Phys. Rev. D* **103** 031301

[26] Eichten E and Martin A 2014 The muon collider as a *H/A* factory *Phys. Lett. B* **728** 125–30

[27] Chakrabarty N, Han T, Liu Z and Mukhopadhyaya B 2015 Radiative return for heavy Higgs boson at a muon collider *Phys. Rev. D* **91** 015008

[28] Buttazzo D, Redigolo D, Sala F and Tesi A 2018 Fusing vectors into scalars at high energy lepton colliders *J. High Energy Phys.* **JHEP11(2018)144**

[29] Bandyopadhyay P and Costantini A 2021 Obscure Higgs boson at colliders *Phys. Rev. D* **103** 015025

[30] Kalinowski J, Robens T, Sokolowska D and Żarnecki A F 2021 IDM benchmarks for the LHC and future colliders *Symmetry* **13** 991

[31] Han T, Li S, Su S, Su W and Wu Y 2021 Heavy Higgs bosons in 2HDM at a muon collider (arXiv:2102.08386 [hep-ph])

[32] Liu W and Xie K-P 2021 Probing electroweak phase transition with multi-TeV muon colliders and gravitational waves (arXiv:2101.10469 [hep-ph])

[33] Costantini A, De Lillo F, Maltoni F, Mantani L, Mattelaer O, Ruiz R and Zhao X 2020 Vector boson fusion at multi-TeV muon colliders *J. High Energy Phys.* **JHEP09(2020)080**

[34] Han T, Liu Z, Wang L-T and Wang X 2020 WIMPs at high energy muon colliders (arXiv:2009.11287 [hep-ph])

[35] Capdevilla R, Meloni F, Simoniello R and Zurita J 2021 Hunting wino and higgsino dark matter at the muon collider with disappearing tracks (arXiv:2102.11292 [hep-ph])

[36] Chiesa M, Maltoni F, Mantani L, Mele B, Piccinini F and Zhao X 2020 Measuring the quartic Higgs self-coupling at a multi-TeV muon collider *J. High Energy Phys.* **JHEP09(2020)098**

[37] Han T, Liu D, Low I and Wang X 2021 Electroweak couplings of the Higgs boson at a multi-TeV muon collider *Phys. Rev. D* **103** 013002

[38] Di Luzio L, Gröber R and Panico G 2019 Probing new electroweak states via precision measurements at the LHC and future colliders *J. High Energy Phys.* **JHEP01(2019)011**

[39] Buttazzo D, Franceschini R and Wulzer A 2020 Two paths towards precision at a very high energy lepton collider (arXiv:2012.11555 [hep-ph])

[40] Capdevilla R, Curtin D, Kahn Y and Krnjaic G 2020 A guaranteed discovery at future muon colliders (arXiv:2006.16277 [hep-ph])

[41] Buttazzo D and Paradisi P 2020 Probing the muon $g-2$ anomaly at a muon collider (arXiv:2012.02769 [hep-ph])

[42] Capdevilla R, Curtin D, Kahn Y and Krnjaic G 2021 A no-lose theorem for discovering the new physics of $(g-2)\mu$ at muon colliders (arXiv:2101.10334 [hep-ph])

[43] Chen N, Wang B and Yao C-Y 2021 The collider tests of a leptophilic scalar for the anomalous magnetic moments (arXiv:2102.05619 [hep-ph])

[44] Yin W and Yamaguchi M 2020 Muon $g-2$ at multi-TeV muon collider (arXiv:2012.03928 [hep-ph])

[45] Huang G-Y, Queiroz F S and Rodejohann W 2021 Gauged $L_\mu - L_\tau$ at a muon collider (arXiv:2101.04956 [hep-ph])

[46] Huang G-Y, Jana S, Queiroz F S and Rodejohann W 2021 Probing the $R_{K^{(*)}}$ anomaly at a muon collider (arXiv:2103.01617 [hep-ph])

[47] Bartosik N *et al* 2020 Detector and physics performance at a muon collider *JINST* **15** P05001

[48] Mokhov N V and Striganov S I 2012 Detector background at muon colliders (arXiv:1204.6721 [physics.ins-det])

[49] CMS C 2019 A MIP timing detector for the CMS phase-2 upgrade *Technical Report CERN-LHCC-2019-003. CMS-TDR-020* CERN Geneva

[50] ATLAS Collaboration Collaboration 2018 Technical proposal: a high-granularity timing detector for the ATLAS phase-II upgrade *Technical Report CERN-LHCC-2018-023. LHCC-P-012* CERN Geneva

[51] Anastasiou C, Duhr C, Dulat F, Furlan E, Gehrmann T, Herzog F, Lazopoulos A and Mistlberger B 2016 High precision determination of the gluon fusion Higgs boson cross-section at the LHC *J. High Energy Phys.* **JHEP05(2016)058**

[52] LHC Higgs Cross Section Working Group Collaboration de Florian D *et al* 2016 Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector (arXiv:1610.07922 [hep-ph])

[53] Boughezal R, Campbell J M, Ellis R K, Focke C, Giele W, Liu X, Petriello F and Williams C 2017 Color-singlet production at NNLO in MCFM *Eur. Phys. J. C* **77** 7

[54] Han T, Ma Y and Xie K 2021 Quark and gluon contents of a lepton at high energies (arXiv:2103.09844 [hep-ph])

[55] Chen J, Han T and Tweedie B 2017 Electroweak splitting functions and high energy showering *J. High Energy Phys.* **JHEP11(2017)093**

[56] von Weizsäcker C 1934 Radiation emitted in collisions of very fast electrons *Z. Phys.* **88** 612–25

[57] Williams E J 1934 Nature of the high energy particles of penetrating radiation and status of ionization and radiation formulae *Phys. Rev.* **45** 729–30

[58] Dawson S 1985 The effective W approximation *Nucl. Phys. B* **249** 42–60

[59] Gao J 2018 Probing light-quark Yukawa couplings via hadronic event shapes at lepton colliders *J. High Energy Phys.* **JHEP01(2018)038**

[60] Altmannshofer W, Brod J and Schmaltz M 2015 Experimental constraints on the coupling of the Higgs boson to electrons *J. High Energy Phys.* **JHEP05(2015)125**

[61] Greco M, Han T and Liu Z 2016 ISR effects for resonant Higgs production at future lepton colliders *Phys. Lett. B* **763** 409–15

[62] Valdivia Garcia M A, Faus-Golfe A and Zimmermann F 2016 Towards a monochromatization scheme for direct Higgs production at FCC-ee (<https://cds.cern.ch/record/2159683>)

[63] d'Enterria D 2017 Higgs physics at the future circular collider *PoS ICHEP2016* 434

[64] Homiller S and Meade P 2019 Measurement of the triple Higgs coupling at a HE-LHC *J. High Energy Phys.* **JHEP03(2019)055**

[65] Mangano M L, Ortona G and Selvaggi M 2020 Measuring the Higgs self-coupling via Higgs-pair production at a 100 TeV p–p collider *Eur. Phys. J. C* **80** 1030

[66] LHC Higgs Cross Section Working Group Collaboration *et al* 2012 LHC HXSWG interim recommendations to explore the coupling structure of a Higgs-like particle (arXiv:1209.0040 [hep-ph])

[67] LHC Higgs Cross Section Working Group Collaboration Andersen J R *et al* 2013 Handbook of LHC Higgs cross sections: 3. Higgs properties (arXiv:1307.1347 [hep-ph])

[68] de Blas J *et al* 2020 Higgs boson studies at future particle colliders *J. High Energy Phys.* **JHEP01(2020)139**

[69] DELPHES 3 Collaboration de Favereau J, Delaere C, Demin P, Giannanco A, Lemaître V, Mertens A and Selvaggi M 2014 DELPHES 3, a modular framework for fast simulation of a generic collider experiment *J. High Energy Phys.* **JHEP02(2014)057**

[70] Bartosik N *et al* 2019 Preliminary report on the study of beam-induced background effects at a muon collider (arXiv:1905.03725 [hep-ex])

[71] Bartosik N *et al* 2020 Detector and physics performance at a muon collider *JINST* **15** P05001

[72] Selvaggi M 2020 Talk at mdi studies meeting of the muon collider collaboration (https://indico.cern.ch/event/957299/contributions/4023467/attachments/2106044/3541874/delphes_card_mucol_mdi%20.pdf 21-09-2020)

[73] Foster G W and Mokhov N V 1996 Backgrounds and detector performance at a 2×2 TeV $\mu^+ \mu^-$ collider *AIP Conf. Proc.* **352** 178–90

[74] Mokhov N V, Alexahin Y I, Kashikhin V V, Striganov S I and Zlobin A V 2011 Muon collider interaction region and machine-detector interface design *Conf. Proc. C* **110328** 82–4

[75] Mokhov N V and Striganov S I 2012 Detector backgrounds at muon colliders *Phys. Proc.* **37** 2015–22

[76] Abramowicz H *et al* 2017 Higgs physics at the CLIC electron–positron linear collider *Eur. Phys. J. C* **77** 475

[77] Boronat M, Fuster J, García I, Ros E and Vos M 2015 A robust jet reconstruction algorithm for high-energy lepton colliders *Phys. Lett. B* **750** 95–9

[78] Boronat M, Fuster J, García I, Roloff P, Simoniello R and Vos M 2018 Jet reconstruction at high-energy electron–positron colliders *Eur. Phys. J. C* **78** 144

[79] Linssen L, Miyamoto A, Stanitzki M and Weerts H (ed) 2012 Physics and detectors at CLIC: CLIC conceptual design report *CERN Yellow Report CERN-2012-003* (arXiv:1202.5940)

[80] An F *et al* 2019 Precision Higgs physics at the CEPC *Chin. Phys. C* **43** 043002

[81] Egana-Ugrinovic D, Homiller S and Meade P 2019 Aligned and spontaneous flavor violation *Phys. Rev. Lett.* **123** 031802

[82] Egana-Ugrinovic D, Homiller S and Meade P 2019 Higgs bosons with large couplings to light quarks *Phys. Rev. D* **100** 115041

[83] Altmannshofer W, Eby J, Gori S, Lotito M, Martone M and Tuckler D 2016 Collider signatures of flavorful Higgs bosons *Phys. Rev. D* **94** 115032

[84] Curtin D *et al* 2014 Exotic decays of the 125 GeV Higgs boson *Phys. Rev. D* **90** 075004

[85] Dicus D A and Mathur V S 1973 Upper bounds on the values of masses in unified gauge theories *Phys. Rev. D* **7** 3111–4

[86] Lee B W, Quigg C and Thacker H B 1977 Strength of weak interactions at very high energies and the Higgs boson mass *Phys. Rev. Lett.* **38** 883–5

[87] Lee B W, Quigg C and Thacker H B 1977 Weak interactions at very high energies: the role of the Higgs-boson mass *Phys. Rev. D* **16** 1519

[88] Veltman M 1977 Second threshold in weak interactions *Acta Phys. Pol. B* **8** 475

[89] Chanowitz M S and Gaillard M K 1985 The TeV physics of strongly interacting W 's and Z 's *Nucl. Phys. B* **261** 379–431

[90] Appelquist T and Chanowitz M S 1987 Unitarity bound on the scale of fermion mass generation *Phys. Rev. Lett.* **59** 2405
Appelquist T and Chanowitz M S 1988 *Phys. Rev. Lett.* **60** 1589 Erratum

[91] Henning B, Lombardo D, Riembau M and Riva F 2019 Measuring Higgs couplings without Higgs bosons *Phys. Rev. Lett.* **123** 181801

[92] Cepeda M *et al* 2019 Report from working group 2: Higgs physics at the HL-LHC and HE-LHC (arXiv:1902.00134 [hep-ph]) vol 7 221–584

[93] Hahn T 2001 Generating Feynman diagrams and amplitudes with FeynArts3 *Comput. Phys. Commun.* **140** 418–31

[94] Mertig R, Böhm M and Denner A 1991 Feyn calc—computer-algebraic calculation of Feynman amplitudes *Comput. Phys. Commun.* **64** 345–59

[95] Shtabovenko V, Mertig R and Orellana F 2016 New developments in FeynCalc9.0 *Comput. Phys. Commun.* **207** 432–44

[96] Shtabovenko V, Mertig R and Orellana F 2020 FeynCalc 9.3: new features and improvements *Comput. Phys. Commun.* **256** 107478

[97] Curtin D, Meade P and Yu C-T 2014 Testing electroweak baryogenesis with future colliders *J. High Energy Phys. JHEP11(2014)127*

[98] Meade P and Ramani H 2019 Unrestored electroweak symmetry *Phys. Rev. Lett.* **122** 041802

[99] Chacko Z, Goh H-S and Harnik R 2006 Natural electroweak breaking from a mirror symmetry *Phys. Rev. Lett.* **96** 231802

[100] Craig N, Knapen S and Longhi P 2015 Neutral naturalness from orbifold Higgs models *Phys. Rev. Lett.* **114** 061803

[101] Craig N, Lou H K, McCullough M and Thalapillil A 2016 The Higgs portal above threshold *J. High Energy Phys. JHEP02(2016)127*

[102] Ellwanger U, Hugonie C and Teixeira A M 2010 The next-to-minimal supersymmetric standard model *Phys. Rep.* **496** 1–77

[103] Profumo S, Ramsey-Musolf M J and Shaughnessy G 2007 Singlet Higgs phenomenology and the electroweak phase transition *J. High Energy Phys. JHEP08(2007)010*

[104] Espinosa J R, Konstandin T and Riva F 2012 Strong electroweak phase transitions in the standard model with a singlet *Nucl. Phys. B* **854** 592–630

[105] Morrissey D E and Ramsey-Musolf M J 2012 Electroweak baryogenesis *New J. Phys.* **14** 125003

[106] Kotwal A V, Ramsey-Musolf M J, No J M and Winslow P 2016 Singlet-catalyzed electroweak phase transitions in the 100 TeV frontier *Phys. Rev. D* **94** 035022

[107] Kurup G and Perelstein M 2017 Dynamics of electroweak phase transition in singlet-scalar extension of the standard model *Phys. Rev. D* **96** 015036

[108] Buttazzo D, Sala F and Tesi A 2015 Singlet-like Higgs bosons at present and future colliders *J. High Energy Phys. JHEP11(2015)158*

[109] Alipour-Fard S, Craig N, Gori S, Koren S and Redigolo D 2020 The second Higgs at the lifetime frontier *J. High Energy Phys. JHEP07(2020)029*

[110] Franceschini R *et al* 2018 The CLIC potential for new physics *CERN Yellow Rep. Monogr.* **3/2018**

[111] CLICdp Collaboration Roloff P, Schnoor U, Simoniello R and Xu B 2020 Double Higgs boson production and Higgs self-coupling extraction at CLIC *Eur. Phys. J. C* **80** 1010

[112] Cirelli M, Fornengo N and Strumia A 2006 Minimal dark matter *Nucl. Phys. B* **753** 178–94

[113] Cirelli M and Strumia A 2009 Minimal dark matter: model and results *New J. Phys.* **11** 105005

[114] Thomas S and Wells J D 1998 Phenomenology of massive vectorlike doublet leptons *Phys. Rev. Lett.* **81** 34–7

[115] Buckley M R, Randall L and Shuve B 2011 LHC searches for non-chiral weakly charged multiplets *J. High Energy Phys. JHEP05(2011)097*

[116] Ibe M, Matsumoto S and Sato R 2013 Mass splitting between charged and neutral winos at two-loop level *Phys. Lett. B* **721** 252–60

[117] Di Luzio L, Gröber R, Kamenik J F and Nardecchia M 2015 Accidental matter at the LHC *J. High Energy Phys. JHEP07(2015)074*

[118] Nobile E D, Nardecchia M and Panci P 2016 Millicharge or decay: a critical take on minimal dark matter *J. Cosmol. Astropart. Phys. JCAP04(2016)048*

[119] Planck Collaboration Aghanim N *et al* 2020 Planck 2018 results: VI. Cosmological parameters *Astron. Astrophys.* **641** A6

[120] Belotsky K M, Khlopov M, Legonkov S and Shibaev K 2005 Effects of new long-range interaction: recombination of relic heavy neutrinos and antineutrinos *Grav. Cosmol.* **11** 27–33

[121] Hisano J, Matsumot S, Nagai M, Saito O and Senami M 2007 Non-perturbative effect on thermal relic abundance of dark matter *Phys. Lett. B* **646** 34–8

[122] Cirelli M, Strumia A and Tamburini M 2007 Cosmology and astrophysics of minimal dark matter *Nucl. Phys. B* **787** 152–75

[123] An H, Wise M B and Zhang Y 2016 Effects of bound states on dark matter annihilation *Phys. Rev. D* **93** 115020

[124] Mitridate A, Redi M, Smirnov J and Strumia A 2017 Cosmological implications of dark matter bound states *J. Cosmol. Astropart. Phys.* **JCAP05(2017)006**

[125] Liu J, Liu Z and Wang L-T 2019 Enhancing long-lived particles searches at the LHC with precision timing information *Phys. Rev. Lett.* **122** 131801

[126] Graham P W, Kaplan D E and Rajendran S 2015 Cosmological relaxation of the electroweak scale *Phys. Rev. Lett.* **115** 221801

[127] Craig N, Katz A, Strassler M and Sundrum R 2015 Naturalness in the dark at the LHC *J. High Energy Phys.* **JHEP07(2015)105**

[128] Craig N 2013 The state of supersymmetry after run I of the LHC, in beyond the standard model after the first run of the LHC (arXiv:1309.0528 [hep-ph])

[129] Dimopoulos S and Giudice G F 1995 Naturalness constraints in supersymmetric theories with non-universal soft terms *Phys. Lett. B* **357** 573–8

[130] Pomarol A and Tommasini D 1996 Horizontal symmetries for the supersymmetric flavor problem *Nucl. Phys. B* **466** 3–24

[131] Cohen A G, Kaplan D B and Nelson A E 1996 The More minimal supersymmetric standard model *Phys. Lett. B* **388** 588–98

[132] Brust C, Katz A, Lawrence S and Sundrum R 2012 SUSY, the third generation and the LHC *J. High Energy Phys.* **JHEP03(2012)103**

[133] Papucci M, Ruderman J T and Weiler A 2012 Natural SUSY endures *J. High Energy Phys.* **JHEP09(2012)035**

[134] Barbieri R and Giudice G F 1988 Upper bounds on supersymmetric particle masses *Nucl. Phys. B* **306** 63–76

[135] Arina C, Fuks B and Mantani L 2020 A universal framework for t -channel dark matter models *Eur. Phys. J. C* **80** 409

[136] Tovey D R 2008 On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders *J. High Energy Phys.* **JHEP04(2008)034**

[137] Draper P and Rzehak H 2016 A review of Higgs mass calculations in supersymmetric models *Phys. Rep.* **619** 1–24

[138] Deser S and Zumino B 1977 Broken supersymmetry and supergravity *Phys. Rev. Lett.* **38** 1433–6

[139] Bolz M, Brandenburg A and Buchmüller W 2001 Thermal production of gravitinos *Nucl. Phys. B* **606** 518–44

Bolz M, Brandenburg A and Buchmüller W 2008 *Nucl. Phys. B* **790** 336–7 Erratum

[140] Pradler J and Steffen F D 2007 Thermal gravitino production and collider tests of leptogenesis *Phys. Rev. D* **75** 023509

[141] Pradler J and Steffen F D 2007 Constraints on the reheating temperature in gravitino dark matter scenarios *Phys. Lett. B* **648** 224–35

[142] Rychkov V S and Strumia A 2007 Thermal production of gravitinos *Phys. Rev. D* **75** 075011

[143] Moroi T, Murayama H and Yamaguchi M 1993 Cosmological constraints on the light stable gravitino *Phys. Lett. B* **303** 289–94

[144] Kawasaki M and Moroi T 1995 Gravitino production in the inflationary universe and the effects on Big-Bang Nucleosynthesis *Prog. Theor. Phys.* **93** 879–99

[145] Moroi T 1995 Effects of the gravitino on the inflationary universe *PhD Thesis* (arXiv:hep-ph/9503210)

[146] Hall L J, Ruderman J T and Volansky T 2015 A cosmological upper bound on superpartner masses *J. High Energy Phys.* **JHEP02(2015)094**

[147] Pierpaoli E, Borgani S, Masiero A and Yamaguchi M 1998 Formation of cosmic structures in a light gravitino-dominated universe *Phys. Rev. D* **57** 2089–100

[148] Viel M, Lesgourgues J, Haehnelt M G, Matarrese S and Riotto A 2005 Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman-alpha forest *Phys. Rev. D* **71** 063534

[149] Hook A and Murayama H 2015 Low-energy supersymmetry breaking without the gravitino problem *Phys. Rev. D* **92** 015004

[150] Hook A, McGehee R and Murayama H 2018 Cosmologically viable low-energy supersymmetry breaking *Phys. Rev. D* **98** 115036

[151] Brignole A, Feruglio F and Zwirner F 1997 On the effective interactions of a light gravitino with matter fermions *J. High Energy Phys.* **JHEP11(1997)001**

[152] Brignole A, Feruglio F. and Zwirner F 1998 Signals of a superlight gravitino at e^+e^- colliders when the other superparticles are heavy *Nucl. Phys. B* **516** 13–28

Brignole A, Feruglio F. and Zwirner F 1999 *Nucl. Phys. B* **555** 653–5 Erratum

[153] Brignole A, Feruglio F, Mangano M L and Zwirner F 1998 Signals of a superlight gravitino at hadron colliders when the other superparticles are heavy *Nucl. Phys. B* **526** 136–52

Brignole A, Feruglio F, Mangano M L and Zwirner F 2000 *Nucl. Phys. B* **582** 759–61 Erratum

[154] Panico G and Wulzer A 2016 *The Composite Nambu-Goldstone Higgs* vol 913 (Berlin: Springer)

[155] Giudice G F, Grojean C, Pomarol A and Rattazzi R 2007 The strongly-interacting light Higgs *J. High Energy Phys.* **JHEP06(2007)045**

[156] Arkani-Hamed N, Cohen A G and Georgi H 2001 Electroweak symmetry breaking from dimensional deconstruction *Phys. Lett. B* **513** 232–40

[157] Arkani-Hamed N, Gregoire T, Wacker J G and Cohen A G 2002 Phenomenology of electroweak symmetry breaking from theory space *J. High Energy Phys.* **JHEP08(2002)020**

[158] Arkani-Hamed N, Cohen A G, Katz E, Nelson A E, Gregoire T and Wacker J G 2002 The minimal moose for a little Higgs *J. High Energy Phys.* **JHEP08(2002)021**

[159] Arkani-Hamed N, Cohen A G, Katz E and Nelson A E 2002 The littlest Higgs *J. High Energy Phys.* **JHEP07(2002)034**

[160] Fox P J, Nelson A E and Weiner N 2002 Dirac gaugino masses and supersoft supersymmetry breaking *J. High Energy Phys.* **JHEP08(2002)035**

[161] Burdman G, Chacko Z, Goh H-S and Harnik R 2007 Folded supersymmetry and the LEP paradox *J. High Energy Phys.* **JHEP02(2007)009**

[162] Craig N, Knapen S and Longhi P 2015 The orbifold Higgs *J. High Energy Phys.* **JHEP03(2015)106**

[163] Cohen T, Craig N, Giudice G F and Mccullough M 2018 The hyperbolic Higgs *J. High Energy Phys.* **JHEP05(2018)091**

[164] Cheung C and Saraswat P 2018 Mass hierarchy and vacuum energy (arXiv:1811.12390 [hep-ph])

[165] Cohen T, Craig N, Koren S, Mccullough M and Tooby-Smith J 2020 *Phys. Rev. Lett.* **125** 151801

[166] Arkani-Hamed N, Cohen T, D’Agnolo R T, Hook A, Kim H D and Pinner D 2016 Solving the hierarchy problem at reheat-ing with a large number of degrees of freedom *Phys. Rev. Lett.* **117** 251801

[167] Hook A and Marques-Tavares G 2016 Relaxation from particle production *J. High Energy Phys.* **JHEP12(2016)101**

[168] Geller M, Hochberg Y and Kuflik E 2019 Inflating to the weak scale *Phys. Rev. Lett.* **122** 191802

[169] Csáki C, D’Agnolo R T, Geller M and Ismail A 2021 Crunching Dilaton, hidden naturalness *Phys. Rev. Lett.* **126** 091801

[170] Strassler M J and Zurek K M 2007 Echoes of a hidden valley at hadron colliders *Phys. Lett. B* **651** 374–9

[171] Han T, Si Z, Zurek K M and Strassler M J 2008 Phenomenology of hidden valleys at hadron colliders *J. High Energy Phys.* **JHEP07(2008)008**

[172] Liu Z, Wang L-T and Zhang H 2017 Exotic decays of the 125 GeV Higgs boson at future e^+e^- colliders *Chin. Phys. C* **41** 063102

[173] Alipour-Fard S, Craig N, Jiang M and Koren S 2019 Long live the Higgs factory: Higgs decays to long-lived particles at future lepton colliders *Chin. Phys. C* **43** 053101

[174] Wang Z S and Wang K 2020 Physics with far detectors at future lepton colliders *Phys. Rev. D* **101** 075046

[175] Katz A, Mariotti A, Pokorski S, Redigolo D and Ziegler R 2017 SUSY meets her twin *J. High Energy Phys. JHEP01(2017)142*

[176] Contino R, Greco D, Mahbubani R, Rattazzi R and Torre R 2017 Precision tests and fine tuning in twin Higgs models *Phys. Rev. D* **96** 095036

[177] ACME Collaboration Andreev V *et al* 2018 Improved limit on the electric dipole moment of the electron *Nature* **562** 355–60

[178] Cairncross W B and Ye J 2019 Atoms and molecules in the search for time-reversal symmetry violation *Nat. Rev. Phys.* **1** 510–21

[179] Hutzler N R 2020 Polyatomic molecules as quantum sensors for fundamental physics *Quantum Sci. Technol.* **5** 044011

[180] Altmannshofer W, Harnik R and Zupan J 2013 Low energy probes of PeV scale sfermions *J. High Energy Phys. JHEP11(2013)202*

[181] McKeen D, Pospelov M and Ritz A 2013 Electric dipole moment signatures of PeV-scale superpartners *Phys. Rev. D* **87** 113002

[182] Arkani-Hamed N, Dimopoulos S, Giudice G F and Romanino A 2005 Aspects of split supersymmetry *Nucl. Phys. B* **709** 3–46

[183] Giudice G F and Romanino A 2006 Electric dipole moments in split supersymmetry *Phys. Lett. B* **634** 307–14

[184] Barr S M and Zee A 1990 Electric dipole moment of the electron and of the neutron *Phys. Rev. Lett.* **65** 21–4
Barr S M and Zee A 1990 *Phys. Rev. Lett.* **65** 2920 Erratum

[185] Cesarotti C, Lu Q, Nakai Y, Parikh A and Reece M 2019 Interpreting the electron EDM constraint *J. High Energy Phys. JHEP05(2019)059*

[186] Panico G, Pomarol A and Riembau M 2019 EFT approach to the electron electric dipole moment at the two-loop level *J. High Energy Phys. JHEP04(2019)090*

[187] Isidori G, Nir Y and Perez G 2010 Flavor physics constraints for physics beyond the standard model *Annu. Rev. Nucl. Part. Sci.* **60** 355

[188] Baldini A *et al* 2018 A submission to the 2020 update of the European Strategy for Particle Physics on behalf of the COMET, MEG, Mu2e and Mu3e collaborations (arXiv:1812.06540 [hep-ex])

[189] Hayasaka K *et al* 2010 Search for lepton flavor violating tau decays into three leptons with 719 million produced $\tau^+\tau^-$ pairs *Phys. Lett. B* **687** 139–43

[190] Belle-II Collaboration Altmannshofer W *et al* 2019 The Belle II physics book *PTEP* **123C01**
(Belle-II Collaboration) Altmannshofer W *et al* 2020 *PTEP* **029201** Erratum

[191] SINDRUM Collaboration Bellgardt U *et al* 1988 Search for the decay $\mu^+ \rightarrow e^+e^+e^-$ *Nucl. Phys. B* **299** 1–6

[192] Murakami B and Tait T M 2015 Searching for lepton flavor violation at a future high energy e^+e^- collider *Phys. Rev. D* **91** 015002

[193] Cirigliano V, Grinstein B, Isidori G and Wise M B 2005 Minimal flavor violation in the lepton sector *Nucl. Phys. B* **728** 121–34

[194] Ellis S A R and Pierce A 2016 Impact of future lepton flavor violation measurements in the minimal supersymmetric standard model *Phys. Rev. D* **94** 015014

[195] Arkani-Hamed N, Cheng H-C, Feng J L and Hall L J 1996 Probing lepton flavor violation at future colliders *Phys. Rev. Lett.* **77** 1937–40

[196] MEG Collaboration Baldini A *et al* 2016 Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+\gamma$ with the full dataset of the MEG experiment *Eur. Phys. J. C* **76** 434

[197] MEG II Collaboration Baldini A *et al* 2018 The design of the MEG II experiment *Eur. Phys. J. C* **78** 380

[198] Mu2e Collaboration Bartoszek L *et al* *Mu2e Technical Design Report*

[199] Mu2e Collaboration Abusalma F *et al* 2018 Expression of interest for evolution of the Mu2e experiment (arXiv:1802.02599 [physics.ins-det])

[200] ECFA/DESY LC Physics Working Group Collaboration Aguilar-Saavedra J *et al* 2001 TESLA: the superconducting electron positron linear collider with an integrated x-ray laser laboratory. Technical design report: Part 3. Physics at an e^+e^- linear collider (arXiv:hep-ph/0106315)

[201] Freitas A, Martyn H-U, Nauenberg U and Zerwas P 2004 Sleptons: masses, mixings, couplings *Int. Conf. Linear Colliders (LCWS 04)* vol 9 pp 939–46

[202] Homiller S, Lu Q and Reece M unpublished

[203] Caprini C and Figueroa D G 2018 Cosmological backgrounds of gravitational waves *Class. Quantum Grav.* **35** 163001

[204] Christensen N 2019 Stochastic gravitational wave backgrounds *Rep. Prog. Phys.* **82** 016903

[205] Meerburg P D *et al* 2019 Primordial non-Gaussianity (arXiv:1903.04409 [astro-ph.CO])

[206] Kamionkowski M, Kosowsky A and Turner M S 1994 Gravitational radiation from first-order phase transitions *Phys. Rev. D* **49** 2837–51

[207] Huber S J and Konstandin T 2008 Gravitational wave production by collisions: more bubbles *J. Cosmol. Astropart. Phys. JCAP09(2008)022*

[208] Bellazzini B, Csaki C and Serra J 2014 Composite Higgses *Eur. Phys. J. C* **74** 2766

[209] Randall L and Sundrum R 1999 Large mass hierarchy from a small extra dimension *Phys. Rev. Lett.* **83** 3370–3

[210] Creminelli P, Nicolis A and Rattazzi R 2002 Holography and the electroweak phase transition *J. High Energy Phys. JHEP03(2002)051*

[211] Randall L and Servant G 2007 Gravitational waves from warped spacetime *J. High Energy Phys. JHEP05(2007)054*

[212] Nardini G, Quirós M and Wulzer A 2007 A confining strong first-order electroweak phase transition *J. High Energy Phys. JHEP09(2007)077*

[213] Konstandin T, Nardini G and Quiros M 2010 Gravitational backreaction effects on the holographic phase transition *Phys. Rev. D* **82** 083513

[214] Konstandin T and Servant G 2011 Cosmological consequences of nearly conformal dynamics at the TeV scale *J. Cosmol. Astropart. Phys. JCAP12(2011)009*

[215] Agashe K, Du P, Ekhterachian M, Kumar S and Sundrum R 2021 Phase transitions from the fifth dimension *J. High Energy Phys. JHEP02(2021)051*

[216] Caprini C *et al* 2020 Detecting gravitational waves from cosmological phase transitions with LISA: an update *J. Cosmol. Astropart. Phys. JCAP03(2020)024*

[217] Chen X and Wang Y 2010 Quasi-single field inflation and non-Gaussianities *J. Cosmol. Astropart. Phys. JCAP04(2010)027*

[218] Arkani-Hamed N and Maldacena J 2015 Cosmological collider physics (arXiv:1503.08043 [hep-th])

[219] Bodas A, Kumar S and Sundrum R 2020 The scalar chemical potential in cosmological collider physics *J. High Energy Phys. JHEP21(2020)079*

[220] Kumar S and Sundrum R 2018 Heavy-lifting of gauge theories by cosmic inflation *J. High Energy Phys.* **JHEP05(2018)011**

[221] Moretti M, Ohl T and Reuter J 2001 O'Mega: an optimizing matrix element generator (arXiv:[hep-ph/0102195](https://arxiv.org/abs/hep-ph/0102195))

[222] Kilian W, Ohl T and Reuter J 2011 WHIZARD-simulating multi-particle processes at LHC and ILC *Eur. Phys. J. C* **71** 1742

[223] Christensen N D, Duhr C, Fuks B, Reuter J and Speckner C 2012 Introducing an interface between FeynRules and WHIZARD *Eur. Phys. J. C* **72** 1990

[224] Alwall J *et al* 2014 The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations *J. High Energy Phys.* **JHEP07(2014)079**

[225] Burgess C P, Pospelov M and ter Veldhuis T 2001 The minimal model of nonbaryonic dark matter: a singlet scalar *Nucl. Phys. B* **619** 709–28

[226] Brivio I, Gavela M B, Merlo L, Mimasu K, No J M, del Rey R and Sanz V 2017 ALPs effective field theory and collider signatures *Eur. Phys. J. C* **77** 572