

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Buried in Background: Hunting for Physics Beyond the Standard Model

Permalink

<https://escholarship.org/uc/item/00q4n4gn>

Author

Koszegi, Giacomo

Publication Date

2023

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Buried in Background: Hunting for Physics Beyond the Standard Model

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

Doctor of Philosophy
in
Physics

by

Giacomo Koszegi

Committee in charge:

Professor Nathaniel Craig, Chair
Professor Xi Dong
Professor Jeffrey Richman

June 2023

The Dissertation of Giacomo Koszegi is approved.

Professor Xi Dong

Professor Jeffrey Richman

Professor Nathaniel Craig, Committee Chair

June 2023

Buried in Background: Hunting for Physics Beyond the Standard Model

Copyright © 2023

by

Giacomo Koszegi

To my parents

Acknowledgements

On the long list of those to whom I am indebted, I would first like to thank my advisor, Nathaniel Craig. His guidance, support, and encouragement over the past five years, as well as the independence he afforded me during that time, have been incalculably valuable to my growth as a physicist, teacher, and professional. I have learned an immense amount of physics, physical intuition, and pedagogy from Nathaniel, for which I am deeply grateful.

Of course, I owe thanks to all the wonderful collaborators I have had the pleasure of working with over the years, without whom this physics could not have been done. In particular, the continued mentorship I have received from Isabel Garcia Garcia has been a crucial and enriching feature of my graduate studies. And I especially want to thank Rudin Petrossian-Byrne, who patiently devoted countless hours (both in-person and late at night over Zoom) to carefully working through calculations with me and straightening out my many confusions; I am very fortunate to have had Rudin as a collaborator.

While this thesis represents the culmination of my research as a graduate student at UCSB, my academic path is inextricably linked to the education I received and the people I met at my undergraduate institution, Caltech. I must confess that, to this day, I am not quite certain why they admitted me in place of countless other worthy applicants (and I am not sure I ever will be). Nor can I adequately express how grateful I am that they *did*. To fully describe the profound impact that my time at Caltech has had would take far too many pages – suffice it to say that it has influenced nearly every part of who I am as a scientist.

I would like to thank my undergraduate advisors, Katherine Faber and Thomas Rosenbaum, who were kind and brave enough to let me do summer research projects in their labs. I owe a debt of gratitude to Tal Einav, whose prowess as a physics teach-

ing assistant I have sought to emulate. And I must thank my fellow undergraduates with whom I collaborated on innumerable problem sets through long nights and early mornings, day in and day out.

Above all else, I thank my family. I would not be here were it not for the love, encouragement, and unwavering support of my parents, Zoltan and Melinda, throughout my academic journey (and especially during the COVID-19 pandemic). And finally I wish to thank my younger siblings – Gianinna, Giovanni, and Giabrielle – for their love, curiosity, and enthusiasm. Words cannot convey my love and appreciation.

Curriculum Vitæ

Giacomo Koszegi

Education

- 2023 Ph.D. in Physics (Expected), University of California, Santa Barbara.
- 2021 M.A. in Physics, University of California, Santa Barbara.
- 2018 B.S. in Physics, California Institute of Technology.

Publications

“Reflections on Bubble Walls”

I. Garcia Garcia, G. Koszegi, R. Petrossian-Byrne
[arXiv:2212.10572 [hep-ph]] [1]

“The Muon Smasher’s Guide”

H. Al Ali, N. Arkani-Hamed, I. Banta, S. Benevedes, D. Buttazzo, T. Cai, J. Cheng, T. Cohen, N. Craig, M. Ekhterachian, J. Fan, M. Forsslund, I. Garcia Garcia, S. Homiller, S. Koren, G. Koszegi, Z. Liu, Q. Lu, K.-F. Lyu, A. Mariotti, A. McCune, P. Meade, I. Ojalvo, U. Oktem, D. Redigolo, M. Reece, F. Sala, R. Sundrum, D. Sutherland, A. Tesi, T. Trott, C. Tully, L.-T. Wang, M. Wang
Rep. Prog. Phys. 85 (2022) 8, 084201, [arXiv:2103.14043 [hep-ph]] [2]

“P Not PQ”

N. Craig, I. Garcia Garcia, G. Koszegi, A. McCune
JHEP 09 (2021) 130, [arXiv:2012.13416 [hep-ph]] [3]

“Deposition of Electrically Conductive Zirconium Monoxide via Plasma Spray-Physical Vapor Deposition”

B. J. Harder, B. Good, M. Schmitt, B. Kowalski, G. Koszegi, M. T. Johnson, K. T. Faber
J. Am. Ceram. Soc. 105 (2022) 5, 3568-3580 [4]

Abstract

Buried in Background: Hunting for Physics Beyond the Standard Model

by

Giacomo Koszegi

As experimental efforts to uncover the nature of physics beyond the Standard Model continue to push the boundaries of energy and sensitivity, theoretical predictions of well-motivated physics models will need to commensurately increase in precision so that we may reliably distinguish signal from background. A prime example may be found in the stochastic gravitational wave background that will soon be within reach of observatories and which could contain imprints from a variety of ultraviolet phenomena, including first order cosmological phase transitions and topological defects. We begin this thesis by discussing scenarios in which the gravitational wave spectrum due to a phase transition can be substantially altered by particle reflection off of relativistic bubble walls, an effect which has been largely ignored in the literature thus far. We then move on to discussing a particular class of parity-based solutions to the strong CP problem which also features a potential gravitational wave signal, this time due to domain wall topological defects. In addition, these models provide testable predictions for near-future colliders and tabletop experiments. Finally, we point out an exciting new computational avenue for discriminating between signal and background in particle collider data using machine learning coupled with a physically motivated metric on the space of collider events.

Permissions and Attributions

1. The content of Chapter 2 is the result of a collaboration with Isabel Garcia Garcia and Rudin Petrossian-Byrne. This work previously appeared in arXiv:2212.10572 [hep-ph].
2. The content of Chapter 3 is the result of a collaboration with Nathaniel Craig, Isabel Garcia Garcia, and Amara McCune. This work previously appeared in the Journal of High Energy Physics (JHEP **09** (2021) 130).
3. The content of Chapter 4 is the result of a collaboration with Tianji Cai, Junyi Cheng, Nathaniel Craig, and Andrew J. Larkoski. This work is ongoing.

Contents

Acknowledgements	v
Curriculum Vitae	vii
Abstract	viii
Permissions and Attributions	ix
1 Introduction	1
2 Reflections on Bubble Walls	6
2.1 Introduction	6
2.2 Setup	15
2.3 Reflection and transmission probabilities	18
2.3.1 Massive electromagnetism	19
2.3.2 A step wall	20
2.3.3 A smooth wall	23
2.4 Dark photon friction on bubble walls	27
2.4.1 Maximum Dynamic Pressure	29
2.4.2 Equilibrium γ -factor and energy budget	33
2.4.3 Fate of the dark radiation	36
2.5 Conclusions	41
Appendices	43
2.A Beyond the effective theory	43
2.A.1 Abelian Higgs UV-completion	43
2.A.2 Interaction rates	45
2.B Reflection probabilities – supplemental material	47
2.B.1 Longitudinal reflection from Goldstone Equivalence	47
2.B.2 Scattering on a δ' potential	49
2.B.3 The Born approximation	50
2.B.4 Numerical methods	52

3	P Not PQ	54
3.1	Introduction	54
3.2	P to solve strong CP	59
3.2.1	Parity as a solution to the strong CP problem	59
3.2.2	Scalar sector and fine-tuning	63
3.2.3	Fermion masses and a low parity-breaking scale	66
3.3	Dial P for Phenomenology	69
3.3.1	Collider bounds	70
3.3.2	Flavor constraints	74
3.4	Broken parity and the neutron EDM	79
3.4.1	Softly broken parity	80
3.4.2	Spontaneously broken parity and CP	83
3.4.3	Spontaneously broken parity alone	86
3.5	Strong CP and quantum gravity	88
3.5.1	Constraints from Planck-suppressed operators	88
3.5.2	Gravitational waves from the spontaneous breaking of parity	93
3.6	Conclusions	100
	Appendices	102
3.A	Mass eigenstates	102
3.A.1	Gauge and Higgs sectors	102
3.A.2	Fermion sector	103
3.B	Radiatively induced EDM	107
3.B.1	One-loop EDM	107
3.B.2	One-loop $\bar{\theta}$	112
3.C	Kaon mixing	115
4	The Power of Covariantized Phase Space	117
4.1	Introduction	117
4.2	Intrinsic geometry of N -body phase space	119
4.2.1	Phase space in the CM frame: a product manifold	120
4.2.2	Phase space in the CM frame: metric and explicit coordination	122
4.2.3	Phase space in the collinear limit	126
4.3	Implementation of the phase space metric for classification	130
4.3.1	Pre-processing of collider events	131
4.3.2	Processing the phase space manifold	135
4.3.3	Post-processing for classification	137
4.4	$t\bar{t}$ vs. QCD dijet event classification	138
4.5	Conclusions	146
	Appendices	148
4.A	Proof of minimal ordering for Euclidean metric	148

Chapter 1

Introduction

With the ultimate goal of discovering what lies beyond the Standard Model (BSM) of particle physics, recent years have seen the emergence of a plethora of upcoming experimental endeavors designed to achieve fantastic levels of precision, sensitivity, or energy. While the results of these experiments will certainly be valuable in their own right, it will be particularly interesting to interpret forthcoming data in the context of constraints on well-motivated models of BSM physics. Doing so reliably, though, requires comparable precision in the theoretical predictions against which we are comparing. For this reason, current efforts to refine the expected signals of BSM models at a variety of experiments – astrophysical, collider, dark matter, gravitational wave, neutrino, tabletop, etc. – are not only exciting for their near-term testability but crucial for the robustness of our conclusions. The hunt for BSM physics hinges on our ability to accurately separate signal from background.

This thesis represents a small step in that direction from theoretical and computational perspectives. We begin by focusing our efforts on gravitational wave observatories, the next generation of which promises measurements of increased sensitivity in a broader frequency range. In fact, upcoming experiments like LISA could very well be sensitive to

a stochastic background of gravitational waves [5, 6], whose existence is well-motivated. In addition to astrophysical events that have occurred throughout cosmological history, like inspirals and mergers, such a background would also have received contributions from more exotic phenomena like the Big Bang, inflation, topological defects, and (notably) first order cosmological phase transitions. LISA, in particular, will be well-suited to measuring the background coming from an electroweak phase transition [5, 6].

Stochastic gravitational waves are hence generally predicted by a wide array of particle physics models, but measuring them is only part of the challenge. There remains an *inverse problem*: given a background with some spectral shape, how do we narrow down the list of qualitatively different phenomena from which it could have arisen? For phase transitions, the background has a power-law shape whose precise form depends mainly on the strength, duration, and dynamics of the transition. While it is probably difficult to isolate a single microscopic model from a given background, it would be far more feasible to exclude large regions of parameter space based on qualitative differences such as these.

Alas, identifying even these qualitative differences can be highly nontrivial. First order phase transitions proceed through bubble nucleation, expansion, and then collision. One very important characteristic of such a phase transition is whether the bubble expansion *runs away* with constant acceleration, or if the bubbles reach some terminal velocity well before collision. This can alter the subsequent gravitational wave spectrum by affecting the duration and energy budget of the phase transition. Determining whether a phase transition will feature run-away bubbles requires a solid understanding of its dynamics. In general, a bubble expanding in vacuum will run away, but this is seldom an accurate picture. More often, bubbles expand against a surrounding plasma which can exert a frictional force that retards acceleration. If the friction (generally a function of bubble wall speed), grows large enough to match the pressure driving expansion, then

the bubbles will reach some terminal velocity.

In chapter 2 we present a new frictional effect based on particle reflection off of relativistic bubbles, which has typically been regarded as subleading in the existing literature. However, as we show, it is possible for reflections to change an otherwise run-away bubble into a terminal velocity bubble, underscoring the need for a more careful theoretical treatment of bubble wall dynamics in anticipation of highly sensitive experimental data from gravitational wave observatories.

Another avenue for tackling the stochastic background inverse problem is leveraging complementarity between *different* experiments (when possible) to more thoroughly test a model’s viability. The model we consider in chapter 2 also features a potential signal in the effective number of relativistic degrees of freedom, quantified by ΔN_{eff} , which upcoming experiments like CMB-Stage 4 could be sensitive to. Our frictional effect relies on the presence of either massive dark photons or $U(1)$ Goldstone bosons, including the QCD axion and axion-like particles (ALPs). We pay particular attention to the fact that such particles are well-motivated candidates for dark matter in many BSM models.

The QCD axion, however, was originally proposed with a very different motivation in mind – solving the strong CP problem. Experimentally, the strong CP problem arises from the fact that the electric dipole moment (EDM) of the neutron (comprised of electrically charged quarks) is constrained to be far smaller than we might expect *a priori*. Theoretically, this implies that the QCD interaction seems to respect CP symmetry to a very high degree, despite the fact that the electroweak interaction maximally violates it. This is one of several fine-tuning problems in the SM. One well-known way of solving it is by introducing a global $U(1)_{\text{PQ}}$ symmetry and using the associated Goldstone (the QCD axion) to dynamically relax the strong sector’s CP violating phase, $\bar{\theta}$, to zero. One of the major issues with this solution, the “axion quality problem”, is that the $U(1)_{\text{PQ}}$ symmetry must be protected from violating terms in the Lagrangian to an exceptional

level. This is particularly difficult to guarantee, given that quantum gravity is expected to generate irrelevant operators violating *all* global symmetries.

In chapter 3 we consider a different class of solutions to the strong CP problem based on generalized parity. We show that the parity solutions under consideration greatly ameliorate the level of fine-tuning required to be consistent with existing experimental constraints and do not suffer from the need to forbid Planck suppressed operators as strongly as in the QCD axion case. The models we consider are by no means new; however, we do point out novel testable predictions that such models offer at a variety of upcoming experiments. Collider searches for heavy resonances, tabletop EDM measurements with even greater precision, and (depending on the mechanism of parity breaking) stochastic gravitational wave background observation all provide possible avenues for examining parity solutions to strong CP in the near future.

So far, we have discussed the theoretical refinements to various BSM model predictions that will be necessary to draw trustworthy inferences from data. In the case of particle colliders, the sheer volume of data produced also presents a computational challenge – how can we efficiently discriminate between signal events from some hypothetical model and background events from uninteresting sources?

The development of sophisticated machine learning (ML) algorithms offers an exciting new approach to this challenge. Even within the realm of ML, though, there are myriad options for the type of algorithm and the format of input data. For example, one could use an image classifier in which the input is some sort of visualization of a reconstructed collider event. While the goal is straightforward to understand (determining how similar two events *look*), the results would lack a clear physical interpretation. It would be preferable to use a simple algorithm coupled with a physically motivated representation of the data in the hopes of understanding *what* the machine learned and what the salient features are that allow one type of event to be distinguished from another.

Chapter 4 presents such a method for supervised event classification. We use an existing ML algorithm (support vector machine learning) based on a notion of distance between events – the input data takes the form of a matrix of pairwise distance values. The algorithm simply classifies events based on how close they are to one another, as measured by whatever metric was used to produce the distance matrix. The closer two events are, the more likely it is that they are of the same type (both signal or both background). The algorithm is “supervised” in the sense that it is trained on data for which it knows the correct answers before being asked to classify new data it has yet to encounter. The algorithm itself is thus relatively intuitive.

The physics enters through our choice of metric for the space of collider events. It turns out that there is actually a natural notion of distance between events when the final state particles can be treated as approximately massless, valid for collisions occurring at large energies in the center-of-momentum (CM) frame. In this case, the four-momenta of the final state particles reside in massless N -body phase space. It was recently shown that, in the CM frame, this phase space constitutes a $(3N - 4)$ -dimensional manifold that can be decomposed into the product of an $(N - 1)$ -dimensional simplex and a $(2N - 3)$ -dimensional sphere [7]. Combining the standard formulae for distances on a simplex and sphere in quadrature yields a metric for the overall manifold.

In addition to the CM frame, we also show in chapter 4 that this decomposition holds in the collinear limit, appropriate for high-energy jets. We then make use of this physically motivated metric to classify several types of simulated events and jets with a high degree of success. Our results demonstrate that this manifold decomposition encodes a fair amount of information about the underlying identity of each event/jet and opens up several topics for future study.

Chapter 2

Reflections on Bubble Walls

2.1 Introduction

Cosmological phase transitions in the early universe that proceed via nucleation of bubbles are a well-motivated possibility in minimal extensions of the Standard Model, as well as in more general scenarios featuring hidden sectors with their own dynamics. Such first order transitions result in the emission of gravitational radiation [8, 9, 10, 11, 12] that current and future observatories may be able to detect in the form of a stochastic background of gravitational waves (see e.g. [5, 6] for reviews). An observation of this kind would provide unambiguous evidence for the existence of degrees of freedom beyond the Standard Model. This has spurred significant interest in the gravitational wave signatures of hidden sectors in recent times [13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26].

Despite the abundance of particle physics models susceptible of undergoing an out-of-equilibrium transition, our ability to make use of the resulting gravitational wave signal to extract information about the relevant dynamics is extremely limited. The most revealing features concern the frequency peak of the stochastic background, as well as its spectral shape at high frequencies. The former determines the epoch at which

the transition takes place, whereas the latter contains information about the dominant source of gravitational radiation. For example, if most of the energy released during the transition goes into accelerating the bubble walls (as in vacuum [27]), these become relativistic and continue to expand at ever-increasing velocities. In this case, collisions of these “run-away” bubbles constitute the main source of gravitational waves, and the resulting signal falls off as f^{-1} at high frequencies [28]. Alternatively, pressure on the bubble walls due to particles in the thermal plasma may cause the expanding walls to reach a terminal speed. In this case, most of the latent heat gets damped instead into the thermal fluid, and it is its subsequent motion that provides the dominant source of gravitational radiation [29]. The high-frequency fall-off of the stochastic background is steeper, e.g. decreasing as f^{-4} for radiation sourced by sound waves [30, 31, 32, 33]. On the other hand, assuming radiation-domination at the time of gravitational wave production, causal propagation restricts the low frequency shape of the spectrum to grow as f^3 , independently of the dominant production channel [34].¹

Understanding the dynamics of expanding bubble walls in the early universe is clearly crucial to determine both quantitative and qualitative features of the resulting gravitational wave signal. But given the intrinsic degeneracy present in any stochastic background, it is equally important to explore alternative probes of the relevant dynamics. For example, the upcoming LISA experiment will be sensitive to phase transitions at electroweak to multi-TeV scale temperatures, probing the nature of the electroweak phase transition and potentially shedding light on the dynamics behind baryogenesis and electroweak symmetry breaking. In this range of energies, the complementarity between LISA and current and future colliders will no doubt be key in furthering our understanding of physics at and around the weak scale (see e.g. [39, 5, 6, 2] for reviews and

¹More generally, the low-frequency shape of the stochastic background depends on the equation of state [35, 36, 37] and on the existence of free-streaming particles [38, 36], and could thus provide non-trivial information about the universe at early times.

references). Beyond (and below) the electroweak scale, phase transitions within hidden sectors may occur at virtually any temperature, and the corresponding stochastic background could fall anywhere from the low frequency range of PTA observatories [40, 41] ($10^{-9} - 10^{-7}$ Hz), to the high frequencies probed by LIGO ($10 - 10^3$ Hz), and beyond. However, vacuum bubbles nucleated in hidden sector transitions may feature very different dynamics to those linked to the weak scale, and the relevant degrees of freedom may be inaccessible at laboratory experiments. Our work is motivated by the goal to more broadly understand the potential behavior of expanding bubbles in the early universe, as well as to identify alternative predictions that may accompany an observable stochastic background of gravitational waves.

In a first order phase transition, bubbles of true vacuum are nucleated at rest, and begin to expand fueled by the difference in free energy densities at either side of the interface. In vacuum, the velocity of the bubble walls evolves according to [27]

$$\frac{d|\vec{v}|}{dt} = \frac{1}{\gamma^3 R_0} \propto \frac{1}{\gamma^3} \frac{\Delta V}{\sigma}, \quad (2.1)$$

with $\gamma = 1/\sqrt{1 - \vec{v}^2}$ the usual Lorentz γ -factor, $R_0 \propto \sigma/\Delta V$ the critical bubble radius, and where ΔV and σ refer to the difference in vacuum energy densities and surface tension of the bubble wall. In reality, bubbles do not expand against a sea of false vacuum, but rather within a non-trivial environment that in the early universe must include the Standard Model plasma as well as, potentially, other ingredients such as dark matter. As the bubble wall speed grows, friction from the surrounding environment can exert a pressure on the interface that opposes the expansion of the bubble walls. Their evolution can still be written as in eq. (2.1), after the replacement [42, 43]:

$$\Delta V \rightarrow \Delta V - \mathcal{P}, \quad (2.2)$$

where \mathcal{P} refers to the frictional pressure resulting from the interaction between the interface and the surrounding medium, and is in general dependent on the speed of the bubble wall.

Two qualitatively different scenarios are thus possible. If $\Delta V \gg \mathcal{P}$ during the entire evolution of the bubble walls, these effectively behave as if they were in vacuum and most of the energy released as the bubbles grow goes into accelerating the expanding interface. The walls then “run away” – that is, they continue to expand with ever increasing velocities. Alternatively, if \mathcal{P} grows large enough so as to neutralize the difference in vacuum energies, $\Delta V = \mathcal{P}$, the bubble walls reach an equilibrium regime of constant speed. Once in equilibrium, the fraction of the total energy that becomes localized on the interface quickly becomes tiny. Instead, most of the energy gets damped into whatever sector of the surrounding environment is responsible for halting the acceleration of the bubble walls. Calculating the pressure experienced by bubble walls as they expand is a classic problem [44, 45, 46, 42, 47, 43] that has received renewed attention in recent times [48, 29, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60].

Within a thermal plasma, particles with phase-dependent mass create a pressure on the expanding walls that asymptotes to a constant $\mathcal{P}_\infty \sim \Delta m^2 T^2$ in the ultra-relativistic limit ², independently of the type of particle [48]. An additional source of friction may be present if the spectrum contains gauge bosons, of the form $\mathcal{P}_\infty \sim \gamma g^2 \Delta m_v T^3$, with g the relevant gauge coupling and Δm_v the change in the mass of the gauge boson at either side of the bubble wall [50, 51, 56]. This effect has its origin in the transition radiation emitted by charged particles as they cross the wall and its γ -dependence can easily render it the most significant source of friction on fast expanding bubbles. Indeed, if the electroweak phase transition were first order, transition radiation would likely cause

²The ultra-relativistic limit refers to the kinematic regime where the energy of the incident particles in the rest frame of the bubble wall is the largest energy scale. Alternatively, this corresponds to the limit $\gamma \rightarrow \infty$ for the γ -factor of the bubble wall.

the bubble walls to reach an equilibrium γ -factor as low as $\gamma_{\text{eq}} = \mathcal{O}(10)$ [50].

A crucial aspect of all sources of friction known so far is that the corresponding pressure is a monotonically increasing function of the wall speed. As put forward in [48, 50], this allows for a simple criterion to determine whether bubble walls during a cosmological phase transition become run-away, by comparing the pressure in the ultra-relativistic limit, \mathcal{P}_∞ , to the difference in vacuum energy density across the wall:

$$\text{Run-away criterion:} \quad \Delta V > \mathcal{P}_\infty \quad [48, 50] . \quad (2.3)$$

In this chapter, we discuss a new physical effect that can qualitatively alter the dynamics of bubble walls during a cosmological phase transition. Namely, the existence of a transient relativistic regime characterized by an approximately constant reflection probability of longitudinal massive vectors off an expanding interface. Effectively, the wall behaves temporarily like an imperfect mirror that reflects a fraction of longitudinal – but not transverse – modes. Two conditions need to be satisfied for this regime to be accessible: (*i*) that the expansion of the bubble walls takes place against a population of massive vectors whose mass changes across the interface; and (*ii*) that the expanding walls are sufficiently “thin”. By thin, we mean that the wall thickness (in the rest frame of the bubble wall) be much smaller than the Compton wavelength of the massive vector, i.e.

$$L \ll m^{-1} . \quad (2.4)$$

In this case, the regime of constant longitudinal reflection corresponds to Lorentz γ -factors in the range

$$1 \ll \gamma \ll (Lm)^{-1} , \quad (2.5)$$

ending when γ is so large that the Lorentz-contracted Compton wavelength of the dark

photon becomes smaller than the wall thickness. We will refer to this kinematic regime as the region of “inter-relativistic” motion. Equation (2.4) ensures that $(Lm)^{-1} \gg 1$, and that this regime is indeed accessible during the evolution of the expanding bubbles. Once γ becomes $\gg (Lm)^{-1}$, reflection probabilities for all polarizations die off exponentially – a well-known feature of the ultra-relativistic limit.

Most notably, the effect described above leads to an additional source of friction on expanding bubble walls. Unlike previously known cases, the corresponding pressure features a characteristic non-monotonic dependence on the relevant γ -factor, reaching a maximum at $\gamma \sim (Lm)^{-1}$ before turning-off at larger values. In (superficial) analogy with the behavior of spacecraft shortly after launch, we will refer to this pressure peak as Maximum Dynamic Pressure. Its existence can make it much harder for bubble walls to become run-away than previously believed, and we will show that eq. (2.3) can be qualitatively misleading in phase transitions where the bubble walls expand against an existing population of phase-dependent massive dark photons.

This chapter is dedicated to deriving the claims made in the previous two paragraphs, as well as illustrating some of their phenomenological consequences. We will refer to the temperature of the Standard Model plasma at the epoch of the phase transition as T_* , and denote the phase transition strength via the usual dimensionless quantity [5, 6]:

$$\alpha \equiv \frac{\Delta V}{\rho_{\text{SM}}(T_*)} = \frac{\Delta V}{\frac{\pi^2}{30} g_*(T_*) T_*^4} . \quad (2.6)$$

We will focus on bubble walls that expand against a population of cold and non-interacting dark photons. Despite its simplicity, this system will be relevant in some physically interesting cases, such as when the dark photons furnish the dark matter [61, 62, 63, 64, 65, 66] – a well-motivated benchmark that we often refer to throughout this work.³ When the

³Other work exploring the interplay between cosmological phase transitions and dark matter includes, e.g. [67, 68].

bubble walls reach an equilibrium regime as a result of longitudinal reflections, the fraction of the total energy that goes into accelerating the bubble walls becomes increasingly small. Instead, most of the available energy goes into making the reflected dark photons relativistic, turning them into dark radiation. If the dark radiation remains relativistic until late times, an observable contribution to ΔN_{eff} is possible. In particular, current bounds on ΔN_{eff} could probe phase transitions with strength $\alpha \gtrsim 10^{-1}$, whereas CMB S-4 measurements could be sensitive to scenarios down to $\alpha \sim 10^{-2}$ for all relevant frequencies.

Extensions of the Standard Model featuring massive vectors are popular both because of their minimality as well as their potential to furnish the dark matter, with a variety of production mechanisms spanning a wide mass range [63, 64, 65, 66]. Before we move on, let us summarize why the existence of *massive* dark photons whose mass changes in the course of our cosmological history is not only a well-motivated possibility, but may be an unavoidable feature in a wide class of models. At the renormalizable level, the physical system that we focus on is described by a Lagrangian of the form

$$\mathcal{L} \supset \frac{1}{2} (\partial_\mu \phi)^2 - V(\phi) - \frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \frac{1}{2} m^2 V_\mu V^\mu, \quad (2.7)$$

where ϕ is a real scalar field and V_μ is the massive dark photon. By assumption, the potential for ϕ features two non-degenerate vacua such that the scalar sector undergoes a phase transition in the early universe. WLOG we take the false and the true vacuum to lie at $\langle \phi \rangle = 0$ and $\langle \phi \rangle = v$. $V(\phi)$ may be a finite temperature effective potential, or it may be a zero-temperature potential in which case the transition proceeds via quantum tunneling. If the scalar sector is in thermal equilibrium with the Standard Model then $T_* \sim v$, whereas $T_* \ll v$ is possible if the transition is ‘super-cooled’, or if ϕ belongs in a hidden sector that is decoupled from the thermal plasma. At nucleation, the thickness

of the bubble walls is determined by the features of the scalar potential, and typically $L \sim (\sqrt{\lambda} v)^{-1}$, with λ the typical quartic coupling in $V(\phi)$. The size of the dimensionless combination $T_* L$ will have a quantitative effect on our results, as we will discuss when relevant.

At the level of eq. (2.7), the scalar and dark photon sectors are fully decoupled. Beyond the renormalizable terms of eq. (2.7), this need not remain true. For example, the following operator

$$\mathcal{L} \supset \frac{\kappa}{2} \phi^2 V^\mu V_\mu, \quad (2.8)$$

leads to an additional contribution to the vector mass in the true vacuum, of the form $\Delta m^2 = \kappa v^2$. As the scalar vev ‘turns on’ in the early universe the dark photon mass will shift accordingly. Of course, the scalar field may be complex instead of real, which can be trivially accommodated by writing $|\phi|^2$ instead. Indeed, an effective interaction $\propto |H|^2 V^\mu V_\mu$, with H the Standard Model Higgs doublet, would lead to a shift in the dark photon mass before and after the electroweak phase transition. We emphasize that eq. (2.8) cannot be forbidden on the basis of symmetry, and its manifest lack of gauge redundancy is a moot point given that the theory under consideration already contains a mass for V^μ . Legalistically, one might object to eq. (2.8) on the basis that it differs from a Stückelberg mass, $\frac{1}{2} m^2 V^\mu V_\mu$, in that the former is not a renormalizable interaction and demands UV-completion. However, upholding the laws of effective field theory, we have no choice but to overrule this objection. Accepting the existence of a finite cutoff in our description of nature, a theory with a massive dark photon and scalar degrees of freedom will in general feature effective interactions as in eq. (2.8).

Indeed, a non-zero κ in eq. (2.8) sets an upper bound on the scale of UV completion. Provided that the term in eq. (2.8) only accounts for a subleading contribution to the overall mass of the dark photon in the true vacuum ($\Delta m^2 \ll m^2$), the upper bound on

the UV cutoff can be conveniently written as [69]

$$\Lambda \lesssim \frac{4\pi v}{\sqrt{\Delta m^2/m^2}}. \quad (2.9)$$

Here, we indeed focus on cases where the change in the dark photon mass is tiny. This allows for a separation between the scale of the phase transition, v , and the UV cutoff, allowing us to neglect the effect of heavy degrees of freedom on the dynamics of the phase transition and focus instead on the consequences of non-zero Δm^2 on the evolution of the bubble walls. As we will see, even when $\Delta m^2/m^2 \ll 1$, the implications for the evolution of cosmological vacuum bubbles can be significant.⁴

An important comment before we proceed. As made clear in the preceding paragraph, the content of this chapter is only relevant in the presence of vector bosons that are *massive* at either side of the bubble wall. Our results therefore do not affect the pressure created by *massless* gauge bosons that gain a mass as they cross the wall, and so we have nothing to add to e.g. the pressure created by W and Z bosons during a first order electroweak phase transition.

The rest of this chapter is organized as follows. We begin in section 2.2 with a summary of the physical setup that we will focus on. In section 2.3, we present our calculation of the reflection probability for phase-dependent massive vectors. Figure 2.2 illustrates the main result of this section: the existence of a transient relativistic regime characterized by an approximately constant reflection probability of longitudinally polarized dark photons. Section 2.4 focuses on fleshing out the consequences of our results for the evolution of bubble walls in the early universe, including the existence of a Maximum Dynamic Pressure in section 2.4.1, as well as a self-consistent determination of the

⁴Although most of our subsequent discussion will proceed within the effective theory defined in eqs. (2.7) and (2.8), we discuss in section 2.A.1 how this effective description can arise from a UV-complete model.

equilibrium γ -factor when this pressure is large enough to halt the acceleration of the bubble walls in section 2.4.2. The fate of the reflected dark photons depends sensitively on a variety of considerations, most notably on whether the sector undergoing the phase transition is hot or cold, as we discuss in section 2.4.3. We summarize our conclusions in section 2.5.

2.2 Setup

The rest of this chapter is dedicated to calculating the pressure on an expanding bubble wall due to an existing population of phase-dependent massive dark photons and to discussing the implications of our results. With this goal, we consider an expanding planar interface, representing a portion of a sufficiently large bubble wall, moving with local velocity \vec{v} and corresponding γ -factor $\gamma \equiv 1/\sqrt{1 - \vec{v}^2}$. The wall is not expanding in vacuum, but rather against a population of cold and non-interacting massive vector bosons with number density n_V . WLOG, we take the velocity of the bubble wall in the rest frame of the dark photons to be along the negative z -axis, $\vec{v} = -|\vec{v}|\hat{z}$. At leading order, to compute the pressure on the wall we need the momentum transfer from particles that either reflect or transmit across the interface. Our assumption that the dark photon sector is non-interacting allows us to consider the individual interactions of particles with the wall, although this approach is more generally valid whenever the relevant mean free path exceeds the wall thickness.

Like previous work focused on computing pressure on expanding bubbles, we find it convenient to work in the rest frame of the interface. In this frame, the scalar order parameter characterizing the transition varies only as a function of the spatial coordinates, which in our convention will be along the z -axis. A dark photon wind moving from the false to the true vacuum hits the bubble wall with velocity $-\vec{v} = |\vec{v}|\hat{z}$. All particles hit

the interface at normal incidence and there are no particles traveling in the opposite direction – both observations follow from our assumption that the dark photon sector is cold ⁵. This setup is depicted in fig. 2.1.




Figure 2.1: In the rest frame of the bubble wall, a dark photon wind hits the interface from the region of false vacuum. The energy of the incident dark photons, $\omega = \gamma m$, remains conserved in the interaction with the wall, whereas momentum along the z -direction changes resulting in a net momentum transfer to the interface. L refers to the wall thickness, i.e. the typical length scale over which the order parameter varies. In this chapter, we focus on the case where the vector mass changes by a small amount at either side of the bubble wall, i.e. $\Delta m^2 \ll m^2 \simeq \tilde{m}^2$. Much of our work will be concerned with the dynamics of the inter-relativistic kinematic regime of eq. (2.5) where $m \ll \omega \ll L^{-1}$.

By assumption, the dark photon mass depends on the order parameter characterizing the phase transition. As anticipated in the Introduction, and further discussed in section 2.A.1, realistic situations often feature a dependence of the form $m_V^2(z) = m^2 + \kappa v(z)^2$, with v the relevant order parameter and κ a constant determined by the underlying model. However, for most of our discussion, it will be sufficient to just assume that the vector mass-squared parameter is z -dependent, with asymptotic values m^2 and \tilde{m}^2 for large negative and positive z , as indicated in fig. 2.1. We define the difference in squared masses as $\Delta m^2 \equiv \tilde{m}^2 - m^2 > 0$ and we will always assume that $\Delta m^2 \ll m^2 \simeq \tilde{m}^2$,

⁵Quantitatively, the assumption of normal incidence requires that, in the wall frame, the normal component of the dark photon momentum is much greater than the components along the plane of the bubble wall, i.e. $|\vec{k}_\perp| \gg |\vec{k}_\parallel|$. For a gas of dark photons at temperature $T_{\text{dp}} \ll m$ and a relativistic wall, $|\vec{k}_\perp| \sim \gamma m \gg |\vec{k}_\parallel| \sim \sqrt{T_{\text{dp}} m}$ provided $\gamma \gg \sqrt{T_{\text{dp}}/m}$, which is satisfied trivially. It should be straightforward to generalize our results beyond these assumptions.

in keeping with the discussion around eq. (2.9).

Since the wall background is independent of t , the energy of the incident particles, $\omega = \gamma m$, remains conserved in the interaction with the wall, whereas momentum in the z -direction changes in a way that results in a net momentum transfer to the interface. In total, the pressure exerted on the wall as a result of this dark photon wind can be written as

$$\mathcal{P} = \gamma |\vec{v}| n_V \times \frac{1}{3} \sum_{\lambda} (R_{\lambda} \Delta k_R + T_{\lambda} \Delta k_T) . \quad (2.10)$$

The factor $\gamma |\vec{v}| n_V$ corresponds to the flux of incoming massive vectors hitting the wall (in the wall frame), whereas the quantity in parentheses represents the average momentum transfer to the wall from an incoming particle with fixed polarization λ , with R_{λ} and T_{λ} the corresponding reflection and transmission probabilities. The momentum transfer from reflected and transmitted particles is given by $\Delta k_R = 2k_z$ and $\Delta k_T = k_z - \tilde{k}_z$, with k_z and \tilde{k}_z the asymptotic transverse momenta at either side of the bubble wall, as defined in fig. 2.1. The sum over λ includes the three physical polarizations of a massive vector, and the lack of mixing between different polarizations in eq. (2.10) follows as a result of normal incidence.

For example, in the regime of ultra-relativistic motion – when ω is the largest energy scale in the problem, $\omega \gg m, L^{-1}$ – reflection and transmission probabilities must asymptote to zero and unity respectively, and eq. (2.10) takes the form,

$$\mathcal{P}_{\infty} = \lim_{\gamma \rightarrow \infty} \mathcal{P} \simeq \gamma n_V \Delta k_T = \rho_V \frac{\Delta m^2}{2m^2} , \quad (2.11)$$

where we have used $\Delta k_T \simeq \Delta m^2 / (2\gamma m)$, and $\rho_V = mn_V$ is the energy density of dark photons. If the massive vector population outside of the bubble wall was instead described in terms of a fully thermal distribution, with $T \gg m$, we would have $\omega \sim \gamma T$ and

$n_V \sim T^3$, and therefore $\mathcal{P}_\infty \sim T^2 \Delta m^2$ as is well known [48]. As seen in eq. (2.11), the frictional pressure reaches a constant (independent of the wall velocity) in this limit.

A comment is in order before we move on. Comparing the asymptotic pressure in eq. (2.11) to the difference in vacuum energy densities, one finds

$$\frac{\mathcal{P}_\infty}{\Delta V} \simeq \frac{\Delta m^2}{2m^2} \frac{\rho_V}{\Delta V} \sim \frac{1}{\alpha} \frac{\Delta m^2}{m^2} \frac{\rho_V(T_*)}{\rho_{\text{dm}}(T_*)} \frac{\rho_{\text{dm}}(T_*)}{\rho_{\text{SM}}(T_*)}, \quad (2.12)$$

where α is defined in eq. (2.6) and in the last step we have made it explicit that all the relevant energy densities are to be evaluated at the phase transition epoch. The ratio ρ_V/ρ_{dm} in eq. (2.12) is 1 if the dark photons account for all of the dark matter, whereas α and $\Delta m^2/m^2$ are $\ll 1$ as discussed previously. The ratio of dark matter to Standard Model radiation in the early universe is $\rho_{\text{dm}}/\rho_{\text{SM}} \lll 1$, making eq. (2.12) correspondingly tiny. As per the run-away criterion of eq. (2.3), one would be tempted to conclude that a small change in the dark matter mass would have a negligible effect on the evolution of cosmological bubble walls. Our results will show that – in general – this expectation can be mistaken.

To proceed, we must obtain reflection and transmission probabilities, as a function of γ , for massive vectors interacting with the non-trivial background of the bubble wall. This problem reduces to solving the equations of motion for massive electromagnetism with a spatially varying photon mass, as we discuss next.

2.3 Reflection and transmission probabilities

We begin in section 2.3.1 with a brief discussion regarding massive electromagnetism with a spatially varying vector mass. In section 2.3.2, we obtain reflection and transmission probabilities in one the of the few non-trivial cases where an analytic solution is

accessible: a step function change in the mass of the dark photon. This corresponds to the limit of vanishing wall thickness, and it provides an accurate description of the system in the regime $\omega \ll L^{-1}$. Considering a step wall allows us to illustrate one of our main results: that the reflection probability for longitudinal modes asymptotes to a constant in the inter-relativistic regime of eq. (2.5). In section 2.3.3 we turn to the realistic situation of finite width, and show that the existence of a relativistic regime of near-constant longitudinal reflection is a generic feature of thin walls with finite thickness.

2.3.1 Massive electromagnetism

In the absence of charged sources, Maxwell's equations in the presence of a varying vector mass read

$$\partial_\mu F^{\mu\nu} + m_V^2(x)V^\nu = 0 , \quad (2.13)$$

which, in turn, imply the consistency condition

$$\partial_\mu (m_V^2(x)V^\mu) = 0 . \quad (2.14)$$

As is well known, when the vector mass is constant, $m_V^2(x) = m^2$, eq. (2.14) reduces to the familiar requirement that $\partial_\mu V^\mu = 0$, and Maxwell's equations admit plane-wave solutions of the form

$$V^\mu(x) = v^\mu e^{-ik \cdot x} \quad (\text{constant vector mass}) , \quad (2.15)$$

for all k^μ such that $k^2 = m^2$, and the v^μ are momentum-dependent complex coefficients satisfying $k_\mu v^\mu = 0$.⁶ As usual, the v^μ can be written as a sum over unit-normalized

⁶The field V^μ is of course real-valued. It is implicitly assumed that only the real part of eq. (2.15) has physical significance. The same applies to all other expressions of this form in the remainder of this chapter.

polarization vectors, one for each of the three physical degrees of freedom of a massive spin-1 particle. For $\vec{k} = \pm|\vec{k}|\hat{z}$, a convenient basis is given by

$$\varepsilon_x^\mu = (0, 1, 0, 0) , \quad \varepsilon_y^\mu = (0, 0, 1, 0) \quad \text{and} \quad \varepsilon_l^\mu = \left(\frac{|\vec{k}|}{m}, 0, 0, \pm \frac{\omega}{m} \right) . \quad (2.16)$$

When the photon mass has a non-trivial profile, analytic solutions to eqs. (2.13) and (2.14) only exist in some special cases. Assuming the vector mass-squared features no time dependence, as appropriate in the rest frame of the bubble wall, eq. (2.14) can be written as

$$\partial_\mu V^\mu = - \left(\vec{\nabla} \log m_V^2(\vec{x}) \right) \cdot \vec{V} . \quad (2.17)$$

If $\vec{k} \parallel \vec{\nabla} m_V^2$, as in the case of normal incidence, eq. (2.17) is non-vanishing for the longitudinal component, whereas $\partial_\mu V^\mu = 0$ for the transverse modes. As summarized around fig. 2.1, this is indeed the setup that we focus on in this work.

2.3.2 A step wall

We will first consider a step-function change in the vector mass:

$$m_V^2(z) = m^2 + \Delta m_V^2(z) \quad \text{with} \quad \Delta m_V^2(z) = \Delta m^2 \Theta(z) , \quad (2.18)$$

where $\Delta m^2 = \tilde{m}^2 - m^2 > 0$ and $\Theta(z)$ is the Heaviside step function. With this mass profile, the equations of motion feature plane-wave solutions on either side of the wall

localized at $z = 0$. These can be written as

$$V_{\perp}^{\mu}(t, z) = e^{-i\omega t} \begin{cases} (0, 1, 1, 0)e^{ik_z z} + r_{\perp}(0, 1, 1, 0)e^{-ik_z z} & z < 0 \\ t_{\perp}(0, 1, 1, 0)e^{i\tilde{k}_z z} & z > 0 \end{cases} \quad (2.19)$$

for the transverse modes,⁷ and

$$V_l^{\mu}(t, z) = e^{-i\omega t} \begin{cases} \left(\frac{k_z}{m}, 0, 0, \frac{\omega}{m}\right) e^{ik_z z} + r_l \left(\frac{k_z}{m}, 0, 0, -\frac{\omega}{m}\right) e^{-ik_z z} & z < 0 \\ t_l \left(\frac{\tilde{k}_z}{m}, 0, 0, \frac{\omega}{m}\right) e^{i\tilde{k}_z z} & z > 0 \end{cases} \quad (2.20)$$

for the longitudinal component, with $k_z = \sqrt{\omega^2 - m^2}$ and $\tilde{k}_z = \sqrt{\omega^2 - \tilde{m}^2}$ as defined in fig. 2.1. The overall normalization of eqs. (2.19) and (2.20) is arbitrary, and with this choice the reflection and transmission probabilities are given by

$$R_{\alpha} = |r_{\alpha}|^2 \quad \text{and} \quad T_{\alpha} = \frac{\tilde{k}_z}{k_z} |t_{\alpha}|^2, \quad (2.21)$$

with $\alpha = \perp, l$ for transverse and longitudinal modes respectively.

We can now obtain analytic solutions for both reflection and transmission probabilities by integrating the equations of motion across the interface. Let us discuss the transverse modes first. As discussed below eq. (2.17), the transverse modes satisfy $\partial_{\mu} V_{\perp}^{\mu} = 0$, and eq. (2.13) reads $(\square + m_V^2(z))V_{\perp}^{\mu} = 0$. We then have:

$$\lim_{\epsilon \rightarrow 0} \int_{-\epsilon}^{+\epsilon} dz (\square + m_V^2(z)) V_{\perp}^{\mu} = 0 \quad \Rightarrow \quad \partial_z V_{\perp}^{\mu} \text{ is continuous at } z = 0. \quad (2.22)$$

Combined with the requirement that V_{\perp}^{μ} itself remains continuous, we can solve for r_{\perp}

⁷We have taken advantage of rotational invariance in the $x - y$ plane to set $t_x = t_y \equiv t_{\perp}$ and $r_x = r_y \equiv r_{\perp}$.

and t_\perp . In particular, the reflection probability is given by

$$R_\perp = |r_\perp|^2 = \left| \frac{k_z - \tilde{k}_z}{k_z + \tilde{k}_z} \right|^2 \xrightarrow{\omega \gg m, \tilde{m}} \left(\frac{\Delta m^2}{4\omega^2} \right)^2 = \gamma^{-4} \left(\frac{\Delta m^2}{4m^2} \right)^2. \quad (2.23)$$

Unsurprisingly, R_\perp falls off rapidly in the regime of relativistic motion.

The behavior of the longitudinal mode is starkly different. We can obtain a first matching condition by integrating eq. (2.14) across the wall:

$$\lim_{\epsilon \rightarrow 0} \int_{-\epsilon}^{+\epsilon} dz \partial_\mu (m_V^2(z) V_l^\mu) = 0 \quad \Rightarrow \quad m_V^2(z) V_l^3 \text{ is continuous at } z = 0. \quad (2.24)$$

Integrating eq. (2.13) provides the second condition we need. Since $\partial_\mu V_l^\mu$ no longer vanishes, expanding eq. (2.13) we now have $(\square + m_V^2(z)) V_l^\mu - \partial^\mu (\partial_\nu V_l^\nu) = 0$. The matching condition arising from the $\mu = 0$ equation is degenerate with eq. (2.24). Instead, focusing on $\mu = 3$, we find

$$\lim_{\epsilon \rightarrow 0} \int_{-\epsilon}^{+\epsilon} dz ((\square + m_V^2(z)) V_l^3 + \partial_z (\partial_\nu V_l^\nu)) = 0 \quad \Rightarrow \quad \partial_t V_l^0 \text{ is continuous at } z = 0. \quad (2.25)$$

Given that the time dependence of V^μ is of the form $e^{-i\omega t}$ for all z , the previous requirement is equivalent to demanding that V_l^0 itself remains continuous. The corresponding reflection probability now reads

$$R_l = |r_l|^2 = \left| \frac{\tilde{m}^2 k_z - m^2 \tilde{k}_z}{\tilde{m}^2 k_z + m^2 \tilde{k}_z} \right|^2 \xrightarrow{\omega \gg m, \tilde{m}} \left(\frac{\tilde{m}^2 - m^2}{\tilde{m}^2 + m^2} \right)^2 \simeq \left(\frac{\Delta m^2}{2m^2} \right)^2, \quad (2.26)$$

where the last step assumes $\Delta m^2 \ll m^2 \simeq \tilde{m}^2$. As advertised in the Introduction, in the regime of relativistic motion the longitudinal reflection probability approaches a constant,

independent of γ .^{8 9}

2.3.3 A smooth wall

In a realistic situation where the wall thickness is finite and the vector mass varies smoothly, the analytic results of the previous subsection only provide an accurate approximation to the reflection and transmission probabilities in the regime $\omega \ll L^{-1}$. In what follows, we perform a more general analysis of the case of finite width. We parametrize the dark photon mass as in eq. (2.18), except now

$$\Delta m_V^2(z) = \Delta m^2 \Theta_L(z) , \quad (2.27)$$

where $\Theta_L(z)$ is no longer a step function, but rather some smooth function that approaches 0 and 1 for large negative and positive z respectively, and with an appropriate step-function limit as $L \rightarrow 0$. We will discuss a specific choice of mass profile at the end of this subsection, but will otherwise keep things general.

Let us start by building some intuition behind the radically different behavior exhibited by the transverse and longitudinal components discussed in section 2.3.2 by taking a closer look at the corresponding field equations. It will be helpful to factor out the time

⁸At the level of a Stückelberg theory, a discontinuity in the number of degrees of freedom (dof) occurs in the limit $m \rightarrow 0$, with 2 vs 3 physical polarizations at either side of the interface. The fact that $R_l \rightarrow 1$ as $m \rightarrow 0$ in eq. (2.26) reflects the observation that the longitudinal mode would be unphysical in the region $z < 0$, and therefore must not propagate into the region $z > 0$. In an Abelian Higgs UV-completion, the number of dof of course stays continuous, with the would-be longitudinal accounted for by the appropriate linear combination of the two real dof of the complex Higgs. The physical dof are two transverse modes and the suitable combination of the real scalars, and therefore the treatment presented in this section is not applicable in the massless regime. Scattering with $m = 0$ has been studied long ago, e.g. [70]. As emphasized in the Introduction, and in section 2.A.1, we instead focus on cases where (in the Abelian Higgs language) the Higgs vev is ‘on’ on both sides of the wall.

⁹See section 2.B.1 for an alternative derivation of eq. (2.26) making use the of Goldstone nature of longitudinals at high energies.

dependence of V^μ , as follows

$$V^\mu(t, z) = v^\mu(z)e^{-i\omega t} , \quad (2.28)$$

and instead focus on the behavior of the $v^\mu(z)$. As we discussed previously, the equations of motion for the transverse modes read $(\square + m_V^2(z))V_\perp^\mu = 0$. In terms of the v^μ , this can be written as

$$(\partial_z^2 + k_z^2) v_\perp^\mu(z) = \Delta m_V^2(z) v_\perp^\mu(z) , \quad (2.29)$$

with $k_z^2 = \omega^2 - m^2$. The transverse components obey a Schrödinger-like equation for a particle moving in one dimension with “energy” k_z^2 in the presence of a potential $U_\perp(z) = \Delta m_V^2(z)$. In the $L \rightarrow 0$ limit, our problem reduces to one-dimensional quantum mechanical scattering on a step potential, with energies above the step. Indeed, eq. (2.23) is just the reflection probability for this classic problem.

On the other hand, the Schrödinger equation governing the behavior of the longitudinal component reads

$$(\partial_z^2 + k_z^2) \lambda(z) = U_l(z)\lambda(z) \quad \text{with} \quad U_l(z) = \Delta m_V^2(z) + \frac{3}{4} \left(\frac{\partial_z m_V^2(z)}{m_V^2(z)} \right)^2 - \frac{1}{2} \frac{\partial_z^2 m_V^2(z)}{m_V^2(z)} , \quad (2.30)$$

where we have defined

$$\lambda(z) \equiv \frac{m_V(z)}{\omega} v_l^3(z) \quad \text{such that} \quad \lambda(z) \rightarrow \begin{cases} e^{ik_z z} + r_l e^{-ik_z z} & \text{for } z \ll -L \\ t_l e^{i\tilde{k}_z z} & \text{for } z \gg +L \end{cases} . \quad (2.31)$$

The effective scattering potential for the longitudinal component, $U_l(z)$, depends on the

length scale over which the mass changes not just through the choice of mass profile but through its derivatives – much unlike its transverse counterparts. It is this crucial difference that leads to the strikingly different behavior of longitudinal and transverse modes.

In the relativistic limit, $k_z \simeq \omega \gg m$, we might neglect the first term in $U_l(z)$. The derivative terms in the effective potential are localized in a region of thickness $\sim L$, and the second-derivative term is dominant whenever $\Delta m^2/m^2$ is tiny. We then have

$$U_l(z) \simeq -\frac{1}{2} \left(\frac{\partial_z^2 m_V^2(z)}{m_V^2(z)} \right) \simeq -\frac{\Delta m^2}{2m^2} \Theta_L''(z) . \quad (2.32)$$

The step-function limit discussed in the previous subsection corresponds to $\Theta_L''(z) \xrightarrow{L \rightarrow 0} \delta'(z)$, and eq. (2.30) reduces to the Schrödinger equation in a potential $U_l(z) = -\frac{\Delta m^2}{2m^2} \delta'(z)$. Solving this equation subject to the appropriate boundary conditions, one indeed recovers the result of eq. (2.26) to leading order in the ratio $\Delta m^2/m^2$, as we summarize in section 2.B.2.

Since the longitudinal reflection coefficient in the relativistic regime is $R_l \ll 1$ whenever the change in the dark photon mass is tiny, we can go further and leverage the one-dimensional Born approximation familiar from quantum mechanics to obtain an analytic expression for the reflection probability in the case of a smoothly varying mass. As we summarize in section 2.B.3, the Born approximation to the longitudinal reflection coefficient can be written in closed form in terms of the effective scattering potential in eq. (2.30) as

$$R_{l,\text{Born}} = \frac{1}{4k_z^2} \left| \int_{-\infty}^{\infty} dz e^{2ik_z z} U_l(z) \right|^2 . \quad (2.33)$$

Plugging eq. (2.32) into eq. (2.33) and integrating by parts, we find

$$R_{l, \text{Born}} \simeq \frac{1}{4k_z^2} \left(\frac{\Delta m^2}{2m^2} \right)^2 \left| \left[e^{2ik_z z} \Theta'_L(z) \right]_{-\infty}^{+\infty} - 2ik_z \int_{-\infty}^{\infty} dz e^{2ik_z z} \Theta'_L(z) \right|^2 \quad (2.34)$$

$$\simeq \left(\frac{\Delta m^2}{2m^2} \right)^2 \left| \int_{-L}^L dz (1 + 2ik_z z + \mathcal{O}(k_z^2 z^2)) \Theta'_L(z) \right|^2 \quad (2.35)$$

$$= \left(\frac{\Delta m^2}{2m^2} \right)^2 (1 + \mathcal{O}(k_z^2 L^2)) . \quad (2.36)$$

To get the second line, we used that Θ'_L quickly vanishes for $|z| \gtrsim L$, and we Taylor expanded the exponential for $|k_z z| \ll 1$. Since Θ_L interpolates between 0 and 1, Θ'_L can be treated as a probability density function. With this interpretation, the integral in the second line is essentially an average of $1 + 2ik_z z$ over the region $(-L, L)$, weighted by this probability density. The average of $k_z z$ is bounded in magnitude by $k_z L$, giving the final line. Equation (2.36) coincides with the step-function result of eq. (2.26) provided $k_z \simeq \omega \ll L^{-1}$, as expected. Moreover, it highlights how the existence of a kinematic regime where the longitudinal reflection coefficient stays nearly constant is in fact a generic feature of any smooth mass profile, lasting all the way up to $k_z \simeq \omega \sim L^{-1}$.

Given a specific profile, we can use eq. (2.33) to obtain an analytic approximation to the reflection probability. For example, for a wall profile of the familiar kink form

$$\Theta_L(z) = \frac{1}{4} [1 + \tanh(z/L)]^2 , \quad (2.37)$$

one finds

$$R_{l, \text{Born}} \simeq \left(\frac{\Delta m^2}{2m^2} \right)^2 \frac{\pi^2 (k_z L)^2 [1 + (k_z L)^2]}{\sinh^2(\pi k_z L)} . \quad (2.38)$$

Equation (2.38) reproduces the result of eq. (2.26) in the regime $\omega \simeq k_z \ll L^{-1}$, up to

corrections of $\mathcal{O}(k_z^2 L^2)$, whereas in the ultra-relativistic limit it takes the form

$$R_{l, \text{Born}} \simeq \left(\frac{\Delta m^2}{2m^2} \right)^2 \times 4\pi^2 (k_z L)^4 e^{-2\pi k_z L} \quad \text{for} \quad k_z \simeq \omega \gg L^{-1}. \quad (2.39)$$

Indeed, the longitudinal reflection probability dies off exponentially fast – as expected – in the regime of ultra-relativistic motion.

The results of this subsection are best summarized in fig. 2.2, where we show the reflection probability for longitudinal and transverse vectors for the mass profile in eq. (2.37), obtained by numerically solving the corresponding equations of motion. As advertised, the longitudinal reflection coefficient features a plateau for γ -factors in the regime $1 \ll \gamma \ll (Lm)^{-1}$, which is well-approximated by the step-function result of eq. (2.26). The consequences of this behavior for the evolution of bubble walls are the topic of the next section.

2.4 Dark photon friction on bubble walls

We now discuss the implications of the results presented in section 2.3 for the dynamics of expanding vacuum bubbles in the early universe. In section 2.4.1, we compute the pressure on an expanding wall due to the presence of dark photons, and argue that the requirement for run-away walls can be much stronger than previously believed. Section 2.4.2 focuses on the dynamics of bubble walls that reach an equilibrium regime as a result of longitudinal friction, including a self-consistent determination of the equilibrium γ -factor. Once in equilibrium, most of the energy released in the transition goes into making a fraction of the dark photons relativistic. The fate of this dark radiation depends sensitively on the size of self-interactions among the dark photons, as well as between the dark photons and particles in the thermal plasma, as we discuss in sec-




Figure 2.2: Reflection probability for longitudinal (blue) and transverse (orange) massive vectors scattering off a planar wall in the limit of normal incidence, as a function of the Lorentz γ -factor. The change in mass across the interface is 1% and 0.1% for the left and right panels respectively. Solid lines show numerical results for a smooth mass profile of the form $m_V^2(z) = m^2 + \Delta m^2 \Theta_L(z)$, with $\Theta_L(z)$ as in eq. (2.37). Dashed lines correspond to the step function results of eqs. (2.23) and (2.26). As anticipated, the step-function analysis provides an excellent approximation up to $\gamma \sim (mL)^{-1}$. Details on how we obtain our numerical results are summarized in section 2.B.4.

tion 2.4.3. If interaction rates are negligible – an assumption more likely to hold if the sector undergoing vacuum decay is cold – we argue that the reflected dark photons accumulate in a thin “shell” of dark radiation surrounding the expanding bubbles. If they remain relativistic until late times, their contribution to ΔN_{eff} could be observable for phase transitions with strength $\alpha \sim 10^{-2} - 10^{-1}$. Alternatively, if the scalar sector is in equilibrium with the thermal plasma, interactions between the reflected dark photons and ϕ particles can easily be efficient, in which case the energy density in dark radiation gets transferred instead into the thermal fluid.

2.4.1 Maximum Dynamic Pressure

It is helpful to rearrange eq. (2.10) using the relationship $T_\lambda = 1 - R_\lambda$, as follows:

$$\mathcal{P} = \gamma |\vec{v}| n_V \left\{ \frac{1}{3} R_l (k_z + \tilde{k}_z) + (k_z - \tilde{k}_z) + \frac{2}{3} R_\perp (k_z + \tilde{k}_z) \right\}. \quad (2.40)$$

In the relativistic limit, $k_z + \tilde{k}_z \simeq 2\gamma m$ and $k_z - \tilde{k}_z \simeq \Delta m^2 / (2\gamma m)$, and eq. (2.40) reads

$$\mathcal{P} \simeq \frac{2}{3} \gamma^2 \rho_V R_l + \rho_V \frac{\Delta m^2}{2m^2} + \frac{4}{3} \gamma^2 \rho_V R_\perp \quad \text{for} \quad \gamma \gg 1. \quad (2.41)$$

The second term above is just the asymptotic contribution of eq. (2.11). The last term corresponds to reflections of transverse modes. As discussed around eq. (2.23), $R_\perp \propto \gamma^{-4}$ already in the inter-relativistic regime, and therefore this third term falls off as γ^{-2} for large γ . In contrast, the peculiar behavior of the longitudinal reflection probability, R_l , which stays approximately constant in the region of inter-relativistic motion, leads to a contribution to the overall pressure on the expanding interface that grows $\propto \gamma^2$. In total, neglecting the last term above and substituting the expression for R_l appropriate in the inter-relativistic regime (remember eq. (2.26)), we have:

$$\mathcal{P} \simeq \underbrace{\frac{2}{3} \gamma^2 \rho_V \left(\frac{\Delta m^2}{2m^2} \right)^2}_{\text{longitudinal reflections}} + \underbrace{\rho_V \frac{\Delta m^2}{2m^2}}_{\mathcal{P}_\infty} \quad \text{for} \quad 1 \ll \gamma \ll (mL)^{-1}. \quad (2.42)$$

We emphasize that this γ^2 -growing pressure is a transient phenomenon that is only present for γ -factors within the range indicated in eq. (2.42). Once $\gamma \gg (mL)^{-1}$, the longitudinal reflection probability dies off exponentially fast, leaving the second term above as the sole significant contribution to the overall pressure.

A consequence of our previous discussion is that the pressure exerted on the expanding interface will reach a maximum near the end of the inter-relativistic regime, before

dropping down to $\mathcal{P}_\infty \simeq \rho_V \times \Delta m^2/2m^2$ in the ultra-relativistic limit. As advertised in the Introduction, we will refer to this pressure peak as Maximum Dynamic Pressure. Parametrically, it is given by

$$\mathcal{P}_{\text{mdp}} \simeq \frac{2}{3}\gamma_{\text{max}}^2\rho_V R_l + \rho_V \frac{\Delta m^2}{2m^2} \sim \frac{\rho_V}{(Lm)^2} \left(\frac{\Delta m^2}{2m^2} \right)^2 + \rho_V \frac{\Delta m^2}{2m^2} , \quad (2.43)$$

where $\gamma_{\text{max}} \sim (mL)^{-1}$ is the value of γ at Maximum Dynamic Pressure. As a result, for bubble walls to become run-away, the following condition must be satisfied:

$$\text{Run-away criterion:} \quad \Delta V > \mathcal{P}_{\text{mdp}} \quad (\text{this work}) . \quad (2.44)$$

Equation (2.44) replaces eq. (2.3) as a diagnostic of run-away bubble walls in the presence of phase-dependent massive dark photons in transitions with access to the regime of inter-relativistic motion, and is a primary new result of this work.

If the first term in eq. (2.43) is a small correction on top of the asymptotic pressure, \mathcal{P}_∞ , the effect of longitudinal mode reflections will be largely irrelevant to describe the dynamics of the expanding bubbles. However, if the first term dominates then eq. (2.44) can be much stronger than the requirement $\Delta V > \mathcal{P}_\infty$ quoted in eq. (2.3). Parametrically, the effect of longitudinal reflections will dominate the overall pressure on the expanding interface provided

$$mL \ll \sqrt{\frac{\Delta m^2}{m^2}} . \quad (2.45)$$

When $\Delta m^2/m^2 \ll 1$, this is a stronger requirement than the condition in eq. (2.5) for the inter-relativistic regime to be accessible, although the main feature in both cases is that there must be a significant hierarchy between the mass of the dark photon and the energy scale L^{-1} characterizing the thickness of the bubble wall.

Our discussion thus far has only been concerned with the pressure due to a population

of cold dark photons. If the sector undergoing the transition is cold, such that vacuum decay proceeds via quantum tunneling and there is no thermal plasma that interacts with the bubble walls, this will be a good approximation to the overall pressure. Alternatively, if the scalar sector is in equilibrium with the thermal fluid, an additional source of pressure will be present due to interactions between the plasma and the wall. As mentioned in the Introduction, particles in the plasma that gain mass across the wall contribute as $\mathcal{P}_\infty \sim \Delta m^2 T^2$ in the relativistic limit [48]. In perturbative theories this source of friction can easily fall below ΔV , and therefore won't be large enough to obstruct the acceleration of the bubble walls. Potentially more significant are those cases where the phase transition sector features gauge bosons that acquire a mass as they cross the wall. As anticipated in section 2.1, this can lead to an additional source of pressure from the transition radiation emitted as charged particles cross the interface, of the form $\mathcal{P}_\infty \sim \gamma g^2 \Delta m_v T^3$ [50, 51, 56]. For this source of friction to be subdominant to that from longitudinal dark photons, one would need

$$\gamma_{\text{eq}} g^2 \Delta m_v T_*^3 \lesssim \Delta V . \quad (2.46)$$

This condition can be interpreted as an upper bound on T_* relative to the typical energy scale of the phase transition. Parametrically, taking for simplicity $\Delta V \sim v^4$ and $\Delta m_v \sim gv$, we find

$$\frac{T_*}{v} \lesssim \frac{1}{g \gamma_{\text{eq}}^{1/3}} \quad (2.47)$$

$$\sim 10^{-3} \left(\frac{1}{g}\right) \left(\frac{\Delta m^2/m^2}{10^{-4}}\right)^{1/3} \left(\frac{10^{-2}}{\alpha}\right)^{1/6} \left(\frac{100 \text{ GeV}}{T_*}\right)^{1/6} \left(\frac{\rho_{\text{dm}}}{\rho_V}\right)^{1/6} . \quad (2.48)$$

If the relevant gauge couplings are $g = \mathcal{O}(1)$ – as in the Standard Model – then neglecting transition radiation would require $T_* \ll v$, i.e. the transition needs to be significantly super-cooled. More generally, in hidden sectors where the relevant gauge couplings are

$g \ll 1$, the above condition could be satisfied even within ‘standard’ thermal transitions where $T_* \sim v$. A more comprehensive analysis of the class of thermal transitions for which this assumption holds is an interesting direction for future investigation.

Figure 2.3 shows the pressure on an expanding interface due to massive dark photons, relative to its asymptotic value in the limit $\gamma \rightarrow \infty$, in cases where the inter-relativistic kinematic regime identified in eq. (2.5) is accessible during the wall’s expansion. The evolution of bubble walls that fail the run-away condition of eq. (2.44) as a result of longitudinal friction is the topic to which we now turn.




Figure 2.3: The pressure exerted by massive dark photons on an expanding bubble wall can feature a non-monotonic dependence on the wall γ -factor due to reflections of longitudinal modes, reaching a peak at $\gamma \sim (mL)^{-1}$ that is potentially much larger than its asymptotic value in the ultra-relativistic limit. The criterion for run-away bubble walls can therefore be far stronger than simply requiring $\Delta V > \mathcal{P}_\infty$, as discussed around eq. (2.44). **Left:** increasing the relative change in the dark photon mass raises the Maximum Dynamic Pressure without significantly altering its position. **Right:** decreasing the wall thickness increases *both* the height and position of the MDP. This behavior is in qualitative agreement with the discussion in and around eq. (2.43). Both plots have been obtained by evaluating eq. (2.40) after numerically obtaining reflection coefficients corresponding to the smooth wall profile of eq. (2.37).

2.4.2 Equilibrium γ -factor and energy budget

If the Maximum Dynamic Pressure is dominated by the reflection of longitudinal modes, and the run-away criterion of eq. (2.44) is not satisfied, then bubble walls will reach an equilibrium regime once $\mathcal{P}(\gamma_{\text{eq}}) \simeq \Delta V$. Parametrically, the equilibrium γ -factor is given by

$$\gamma_{\text{eq}} \simeq \left(\frac{3\Delta V}{2\rho_V R_l} \right)^{1/2} \quad (2.49)$$

$$\sim 10^9 \left(\frac{10^{-4}}{\Delta m^2/m^2} \right) \left(\frac{\alpha}{10^{-2}} \right)^{1/2} \left(\frac{T_*}{100 \text{ GeV}} \right)^{1/2} \left(\frac{\rho_{\text{dm}}}{\rho_V} \right)^{1/2}. \quad (2.50)$$

How easy is it for bubble walls to reach an equilibrium regime as a result of longitudinal reflections? An obvious self-consistency condition on our determination of γ_{eq} in eq. (2.49) is that it lies below the γ -factor at Maximum Dynamic Pressure, i.e.

$$\gamma_{\text{eq}} \lesssim \frac{1}{mL}. \quad (2.51)$$

Additionally, a “kinematic” condition is that the expanding walls reach equilibrium *before* the bubble walls collide and the phase transition ends. The bubble radius at collision is set by the Hubble scale at the epoch of the phase transition. Demanding that the size of the expanding bubbles at the onset of the equilibrium regime is a fraction $x < 1$ of the Hubble radius, we find

$$R_{\text{eq}} \simeq \gamma_{\text{eq}} R_0 \lesssim x H(T_*)^{-1} \quad \Rightarrow \quad \gamma_{\text{eq}} \lesssim \frac{x}{R_0 H(T_*)}, \quad (2.52)$$

where R_0 is the critical bubble radius.

Equations (2.51) and (2.52) can be interpreted in various ways, but perhaps the most relevant for us is to regard them as *lower bounds* on the fractional change of the

dark photon mass for the effect of longitudinal pressure to be large enough to stop the acceleration of the bubble walls, as follows:

$$\frac{\Delta m^2}{m^2} \gtrsim \left(\frac{\Delta m^2}{m^2} \right)_{\min} \equiv 2 \left(\frac{3\Delta V}{2\rho_V} \right)^{1/2} \times \text{Max} \{ m L, x^{-1} R_0 H(T_*) \} . \quad (2.53)$$

This is illustrated in fig. 2.4, where we show contours of $\sqrt{(\Delta m^2/m^2)_{\min}}$ for various choices of the underlying parameters, as described in the caption. As can be seen in fig. 2.4, even extremely small changes in the mass of the dark photon across the interface can cause enough friction to halt the acceleration of the bubble walls.

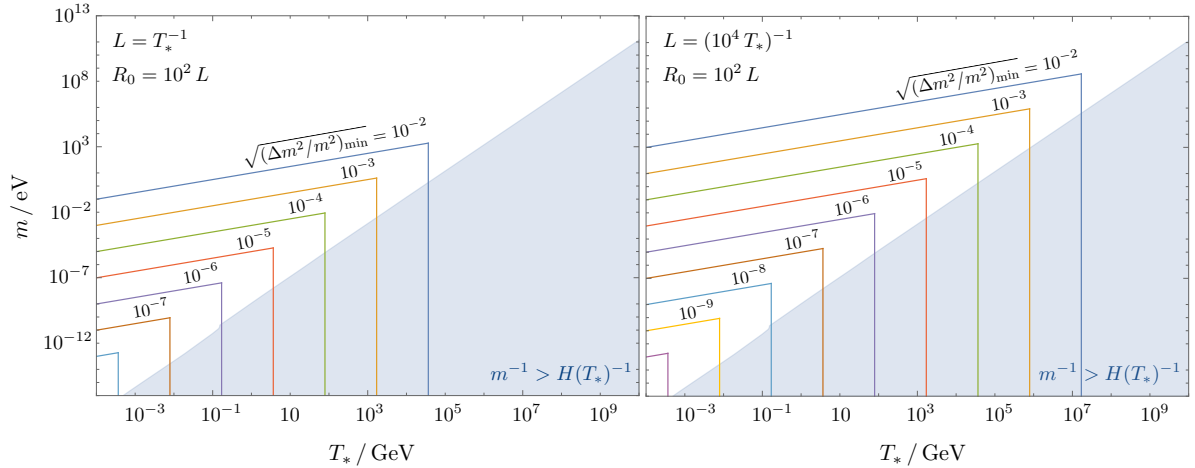


Figure 2.4: Contour lines indicating the lower bound on the fractional change of the dark photon mass for longitudinal reflections to create enough friction to halt the acceleration of the bubble walls. These contours saturate the inequality in eq. (2.53), with the choice $x = 10^{-1}$ for illustration. The energy density in cold dark photons is taken to be that of the dark matter, and the difference in energy densities, ΔV , is taken to be 1% of that in the Standard Model bath (i.e. $\alpha = 10^{-2}$ in eq. (2.6)). The blue shaded area corresponds to dark photon masses such that the corresponding Compton wavelength exceeds the Hubble radius, $m^{-1} > H(T_*)^{-1}$ – a regime that lies outside the validity of our discussion. **Left:** The wall thickness is related to the temperature of the thermal plasma at the phase transition epoch as $L = T_*^{-1}$, as expected in a ‘standard’ thermal transition. **Right:** $L = (10^4 T_*)^{-1} \ll T_*^{-1}$, as could be the case if the phase transition took place within a cold hidden sector with a characteristic energy scale $\gg T_*$, or for a thermal transition featuring significant super-cooling (recall discussion below eq. (2.7)).

After the bubble walls reach an equilibrium speed, they carry a decreasing fraction of the total energy available in the transition:

$$\frac{E_{\text{wall}}}{E_{\text{total}}} \simeq \frac{4\pi R(t)^2 \gamma_{\text{eq}} \sigma}{\frac{4\pi}{3} R(t)^3 \Delta V} \sim \frac{\gamma_{\text{eq}} \sigma}{\Delta V R(t)}, \quad (2.54)$$

where $R(t)$ is the bubble radius at time t . As $R(t)$ grows, $E_{\text{wall}}/E_{\text{total}}$ quickly becomes tiny. By comparison, the above ratio is identically 1 for transitions in vacuum, where all the available energy goes into accelerating the bubble walls [27]. In thermal transitions where the walls reach an equilibrium speed due to friction from the thermal plasma, most of the available energy goes instead into producing motion in the form of sound waves or hydrodynamic turbulence, as mentioned in the Introduction and summarized in [5, 6]. In the case at hand, as we will now show, the energy released in the transition goes instead into accelerating a fraction of the dark photons, turning them into dark radiation.

In the rest frame of the bubble wall, reflected longitudinal modes propagate away from the interface with Lorentz-factor γ_{eq} , while in the rest frame of the dark photons far away from the wall – loosely, the rest frame of the dark matter – the γ -factor of the reflected dark photons is given by

$$\gamma_{\text{dr}} = \frac{1 + |\vec{v}_{\text{eq}}|^2}{1 - |\vec{v}_{\text{eq}}|^2} \simeq 2\gamma_{\text{eq}}^2 \gg \gamma_{\text{eq}}. \quad (2.55)$$

Within a Hubble volume, the average number density of dark photons that become relativistic as the bubble walls sweep across the dark matter from the onset of the equilibrium regime until the bubbles collide can be approximated by

$$\langle n_{\text{dr}} \rangle \simeq \frac{1}{3} R_i n_V \left(1 - \frac{R_{\text{eq}}^3}{R_{\text{coll}}^3} \right). \quad (2.56)$$

The factor in parenthesis takes care of the fact that we are only interested in keeping track

of the fraction of dark matter that gets converted into dark radiation during equilibrium (before, most of the energy released as the bubbles grow goes into accelerating the bubble walls). Obviously, if the bubble walls only reach equilibrium right before they collide, $R_{\text{eq}} \approx R_{\text{coll}}$, then $\langle n_{\text{dr}} \rangle$ will be correspondingly tiny. However, provided the onset of equilibrium takes place well before collision, then $\langle n_{\text{dr}} \rangle \simeq \frac{1}{3} R_l n_V$. For example, taking $R_{\text{eq}} \simeq \gamma_{\text{eq}} R_0$ and $R_{\text{eq}} \simeq x H(T_*)^{-1}$, we find $R_{\text{eq}}^3 / R_{\text{coll}}^3 \sim 10^{-15} \ll 1$, where we have evaluated all parameters as in eq. (2.50), and for illustration we have taken $x = 10^{-2}$ as well as $R_0^{-1} = T_* = 100$ GeV. More generally, notice that having $R_{\text{eq}} \lesssim R_{\text{coll}}$ is a requirement for dark photon reflections to affect significantly the evolution of the bubble walls (as per our discussion around eq. (2.50)), and so we proceed under this assumption.

Right after bubble walls collide, the volume-averaged energy density in dark radiation is therefore

$$\rho_{\text{dr}} = \omega_{\text{dr}} \langle n_{\text{dr}} \rangle \simeq \frac{2}{3} \gamma_{\text{eq}}^2 R_l \rho_V \quad \text{at} \quad T \sim T_*, \quad (2.57)$$

where $\omega_{\text{dr}} = \gamma_{\text{dr}} m$ and γ_{dr} is given in eq. (2.55). Notice the right-hand-side above is equal to ΔV by eq. (2.49) – consistent with energy conservation in the rest frame of the dark matter. Thus, in phase transitions where bubble walls reach an equilibrium regime as a result of friction from the dark photons, most of the difference in vacuum energy densities goes into turning a fraction of the (cold) dark photons into dark radiation.

2.4.3 Fate of the dark radiation

The fate of the reflected dark photons depends sensitively on the values of the various underlying parameters. As a result, general statements about the later evolution of the dark radiation are not possible, and a comprehensive study spanning all available parameter space is well beyond the scope of this work. Instead, we will focus on highlighting the possible outcomes given certain assumptions.

Although sweeping statements are not possible, there are two qualitatively different cases that merit separate consideration, depending on whether the sector undergoing vacuum decay is part of the thermal plasma or belongs in a cold hidden sector such that the transition proceeds via quantum tunneling. We will focus on the cold scenario first. In this case, the absence of a thermal population of particles that interact with the wall background leaves the dark photons as the sole source of pressure on the expanding interface, and the evolution of the dark radiation as the bubbles grow depends primarily on the strength of the self-interactions among the dark photons. As discussed in section 2.A.1, quartic dark photon self-interactions are generally part of the low energy theory, and in the context of an Abelian Higgs UV-completion the relevant quartic coupling is as given in eq. (2.68). Interactions among the dark radiation and the cold dark photons occur at a center-of-mass energy $\sqrt{s}/2 \simeq \gamma_{\text{eq}} m$, which always fall below the cutoff scale of the effective theory. In this kinematic regime, the relevant cross section is as given in eq. (2.72), and one finds that the corresponding interaction rate can be tiny compared to $H(T_*)$ across significant portions of parameter space.

When interactions between the dark matter and the reflected dark photons are negligible, the following picture emerges. Going back to eq. (2.55), notice that (in the dark matter frame) the speed of the dark radiation is ever-so-slightly larger than the speed of the bubble wall:

$$\Delta v_{\text{dr}} \equiv |\vec{v}_{\text{dr}}| - |\vec{v}_{\text{eq}}| = |\vec{v}_{\text{eq}}| \frac{1 - |\vec{v}_{\text{eq}}|^2}{1 + |\vec{v}_{\text{eq}}|^2} \simeq \frac{|\vec{v}_{\text{eq}}|}{2\gamma_{\text{eq}}^2} \ll 1 . \quad (2.58)$$

As a result, from the beginning of the equilibrium regime until the moment of bubble wall collisions, reflected dark photons become uniformly distributed on a shell of thickness ΔL

in front of the interface, with

$$\Delta L \sim \Delta v_{\text{dr}} \times \Delta t \simeq \frac{|\vec{v}_{\text{eq}}| \Delta t}{2\gamma_{\text{eq}}^2} \lesssim \frac{R_{\text{coll}}}{2\gamma_{\text{eq}}^2} \ll R_{\text{coll}} , \quad (2.59)$$

where Δt refers to the time interval between when equilibrium is reached and collision of the bubble walls, and we have used $|\vec{v}_{\text{eq}}| \Delta t \lesssim R_{\text{coll}}$. Thus, as the bubbles grow bigger, a shell of dark radiation forms, moving ever-so-slightly in front of the bubble walls. Although the walls move at constant speed, most of the energy density remains localized in a thin layer close to the surface of the expanding bubbles.

The fate of these shells of radiation depends on the interactions between relativistic dark photons as the bubble walls meet. Now, the relevant center-of-mass-energy is $\sqrt{s}/2 \sim \omega_{\text{dr}} \simeq 2m\gamma_{\text{eq}}^2$, and interactions are often well-described within the high energy Goldstone regime of eq. (2.74). To estimate the interaction rate, notice that although the average number density of dark radiation is well-approximated by $\langle n_{\text{dr}} \rangle \simeq \frac{1}{3} R_l n_V$, the distribution of relativistic dark photons is highly inhomogeneous, with $n_{\text{dr}} \sim \langle n_{\text{dr}} \rangle H^{-1} / \Delta L \gg \langle n_{\text{dr}} \rangle$ in a thin shell surrounding the bubble walls and zero elsewhere. This cross section defines a mean-free-path $\lambda_{\text{mfp}} = (\sigma n_{\text{dr}})^{-1}$. Comparing this length scale to the typical thickness of one of these shells we find, e.g.

$$\frac{\lambda_{\text{mfp}}}{\Delta L} \sim \frac{10^8}{\lambda_{\Phi}^2} \left(\frac{m}{1 \text{ eV}} \right)^3 \left(\frac{10^{-4}}{\Delta m^2 / m^2} \right)^6 \left(\frac{T_*}{100 \text{ GeV}} \right) \left(\frac{\alpha}{10^{-2}} \right)^2 \left(\frac{\rho_{\text{dm}}}{\rho_V} \right)^3 . \quad (2.60)$$

The above expression has a strong dependence on a number of parameters, especially the fractional change in the mass of the dark photon. But interestingly it can remain $\gg 1$ in a large region of the relevant parameter space where longitudinal reflections are relevant to the evolution of the bubble walls. In this case, the shells of dark radiation will pass each other without significant dissipation, a process that could last for much longer than

the usual duration of a cosmological phase transitions. The long-lasting motion of these shells could greatly enhance the strength of the gravitational wave signal, a possibility that motivates more careful exploration of this potential new source of gravitational waves. We will return to this topic in future work.

Moreover, it is interesting to consider the limiting possibility that the reflected dark photons remain relativistic long after the phase transition is over, leading to a dark radiation signal. Making the optimistic assumption that significant losses to gravitational radiation and other “inelastic” processes can be ignored, the requirement that the reflected dark photons remain relativistic at temperatures $T \leq T_*$ can be written as:

$$\gamma_{\text{dr}}(T \leq T_*) \simeq 2\gamma_{\text{eq}}^2 \frac{a(T_*)}{a(T)} \approx 2 \times 10^{13} \left(\frac{10^{-4}}{\Delta m^2/m^2} \right)^2 \left(\frac{\alpha}{10^{-2}} \right) \left(\frac{\rho_{\text{dm}}}{\rho_V} \right) \left(\frac{T}{1 \text{ MeV}} \right) \gtrsim 1. \quad (2.61)$$

Of particular interest are the temperatures of Big Bang Nucleosynthesis ($T_{\text{BBN}} \sim 1 \text{ MeV}$) and recombination ($T_{\text{rec}} \sim 1 \text{ eV}$), when ΔN_{eff} bounds exist. As can be inferred from eq. (2.61), this is easily the case provided α is not too tiny. When this is the case, the corresponding contribution to ΔN_{eff} can be written as

$$\Delta N_{\text{eff}} = \frac{8}{7} \left(\frac{11}{4} \right)^{4/3} \left(\frac{\rho_{\text{dr}}(T_*)}{\rho_{\gamma,0}} \right) \left(\frac{a(T_*)}{a(T_0)} \right)^4 \quad (2.62)$$

$$\simeq 1.4 \left(\frac{\alpha}{0.1} \right) \frac{g_*(T_*)}{g_{*,s}(T_*)^{4/3}} \quad (2.63)$$

$$\simeq 0.3 \left(\frac{\alpha}{0.1} \right) \left(\frac{g_*(100 \text{ GeV})}{g_*(T_*)} \right)^{1/3}, \quad (2.64)$$

where $\rho_{\gamma,0}$ refers to the current energy density in (Standard Model) photons. In the first step above we have used $\rho_{\text{dr}}(T_*) \simeq \Delta V$, as discussed around eqs. (2.56) and (2.57), and in the final step we have ignored the difference between $g_{*,s}(T_*)$ and $g_*(T_*)$. Interestingly,

phase transitions with $\alpha \sim 10^{-1}$ would already be probed by current ΔN_{eff} constraints – though we emphasize that this is under the assumption that the highly relativistic dark photons suffer no significant energy loss since the time of the phase transition other than redshift due to Hubble expansion. Under this assumption, fig. 2.5 shows the prediction for ΔN_{eff} as a function of the characteristic temperature of the phase transition, for various values of α . Strong phase transitions with $\alpha > 0.1$ would be probed by current measurements, while weaker transitions with $\alpha \gtrsim 10^{-2}$ could be probed by CMB-Stage 4 observations [71]. Because bubble acceleration is stalled by the reflection of particles, this




Figure 2.5: In phase transitions where bubble walls reach an equilibrium regime as a result of dark photon reflections, most of the difference in vacuum energy densities goes into turning a fraction of the (cold) dark photons into dark radiation. If the dark radiation remains relativistic until late times, an observable contribution to ΔN_{eff} is possible. Plotted are the maximum possible contributions to ΔN_{eff} , as discussed in the main text, as a function of the thermal plasma at the epoch of the transition, T_* . Solid gray lines correspond to current bounds [72], and dashed gray lines show the projected sensitivities of upcoming observatories [71, 73].

effect is a particularly efficient manner of converting the latent heat of the false vacuum to relativistic matter, thus leading to a significant contribution to ΔN_{eff} if the reflected particles remain relativistic until late times. When instead bubbles of the new phase are run-away or reach a terminal velocity by dispersing energy into the thermal plasma,

only the resulting stochastic background of gravitational waves contributes towards dark radiation, and ΔN_{eff} is a less important observable.

The second scenario highlighted early on in this subsection corresponds to the case where the phase transition sector is in equilibrium with the thermal plasma. In this case, the reflected dark photons will also interact with the population of ϕ 's in the thermal fluid via eq. (2.8) (or its UV-complete version, as described in section 2.A.1). Interactions take place at center-of-mass energies $\sqrt{s}/2 \sim \sqrt{2T_*\omega_{\text{dr}}}/2 \simeq \gamma_{\text{eq}}\sqrt{mT_*}$, and the corresponding cross section is well-approximated by eq. (2.75) or eq. (2.76) depending on whether \sqrt{s} falls above or below the cutoff scale of the effective theory. As discussed in section 2.A.2, the corresponding interaction rate can be well above $H(T_*)$, although once more this depends sensitively on the values of the underlying parameters. At any rate, interactions between the reflected dark photons and the thermal plasma provide an additional obstacle for the dark radiation to propagate undisturbed. It is therefore more likely that the energy density in the dark radiation gets distributed within the primordial plasma, in which case the main source of gravitational radiation would come from motion within the thermal fluid, similar to the case of phase transitions that reach an equilibrium regime as a result of thermal pressure.

Determining more precisely under what circumstances the reflected dark photons remain relativistic until late times, and therefore contribute to ΔN_{eff} , is an important question worthy of further attention.

2.5 Conclusions

In this work, we have discussed a new physical effect that can affect the evolution of cosmological vacuum bubbles expanding against a population of phase-dependent massive dark photons, with a special focus on the case where the dark photons furnish the dark

matter. Namely, the existence of a transient relativistic regime, for sufficiently thin walls, characterized by a constant reflection probability of longitudinal dark photons. The reflection of longitudinal modes creates a pressure on the expanding interface that features a characteristic non-monotonic dependence on the γ -factor of the bubble wall, reaching a peak at intermediate values of γ that we have dubbed Maximum Dynamic Pressure. The existence of a MDP that exceeds the asymptotic value of the pressure in the ultra-relativistic limit can make it much harder for bubble walls to become run-away, even in the absence of a thermal plasma that interacts with the wall background.

Our work opens a number of avenues for future exploration. In phase transitions where bubble walls reach an equilibrium regime as a result of this effect, the later evolution of the reflected dark photons could modify the features of the resulting gravitational wave signal and, in some cases, lead to an observable contribution to ΔN_{eff} if the reflected dark photons remain relativistic until late times. As discussed in section 2.4 (especially section 2.4.3), the extent to which this happens depends on a variety of considerations, most importantly on whether the sector undergoing vacuum decay is cold or, instead, is in thermal equilibrium with the primordial plasma surrounding the expanding bubbles. Understanding more generally the possible implications of the dark radiation for these two broad classes of models is clearly an important topic that deserves further attention.

Moreover, although we have focused exclusively on the case of phase-dependent massive dark photons, it is possible that the phenomenon of MDP on expanding bubble walls could be realized in scenarios beyond this example. Given the impact this can have on the dynamics of bubble walls, which in turn largely determine the features of the resulting gravitational wave signal, this is an important question worthy of further investigation.

Appendices

2.A Beyond the effective theory

2.A.1 Abelian Higgs UV-completion

In this appendix, we briefly discuss how the effective Lagrangian of eqs. (2.7) and (2.8) may arise from a more complete framework. Our goal is not to be exhaustive, but rather to identify the necessary features of underlying models giving rise to our effective theory. The simplest UV-completion of eqs. (2.7) and (2.8) can be written in terms of an Abelian Higgs model:

$$\mathcal{L}_{\text{UV}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + |D_\mu\Phi|^2 - V_\Phi(|\Phi|) + \frac{1}{2}(\partial_\mu\phi)^2 - V_\phi(\phi) + \frac{\eta}{2}\phi^2|\Phi|^2, \quad (2.65)$$

where $D_\mu\Phi \equiv (\partial_\mu - ig'V_\mu)\Phi$. Assuming a Higgs potential of the usual form: $V_\Phi(|\Phi|) = \lambda_\Phi/2(|\Phi|^2 - v'^2/2)^2$, the corresponding radial mode gets a mass $m_\rho = \sqrt{\lambda_\Phi}v'$, whereas $m = g'v'$ is the mass of the dark photon. Integrating out the radial mode, one obtains a low energy effective theory featuring additional interactions among the light degrees of freedom. In particular, a term such as eq. (2.8) – repeated here for convenience – is

generated at tree-level:

$$\mathcal{L} \supset \frac{\kappa}{2} \phi^2 V^\mu V_\mu \quad \text{with} \quad \kappa = \frac{\eta m^2}{m_\rho^2} = \frac{\eta g'^2}{\lambda_\Phi}. \quad (2.66)$$

Notice in particular that the ratio $\Delta m^2/m^2$ is independent of g' , and it is therefore controlled by a different set of parameters of the underlying model compared to the overall mass of the dark photon. In this context, the scale of UV-completion is $\Lambda = m_\rho$, which can be written as

$$m_\rho = \frac{\sqrt{\eta} v}{\sqrt{\Delta m^2/m^2}} \lesssim \frac{4\pi v}{\sqrt{\Delta m^2/m^2}}. \quad (2.67)$$

This reproduces our EFT expectation for the upper bound on the cutoff scale in eq. (2.9) subject to the perturbativity requirement $\eta \lesssim 16\pi^2$.

Of course, the low energy effective theory contains interactions among the light degrees of freedom beyond eq. (2.8). Of particular interest are self-interactions among the dark photons. Indeed, a quartic interaction is generated at tree-level, of the form

$$\mathcal{L} \supset \lambda_V (V_\mu V^\mu)^2 \quad \text{with} \quad \lambda_V \sim \frac{g'^4}{\lambda_\Phi} = \lambda_\Phi \frac{m^4}{m_\rho^4} \lll 1. \quad (2.68)$$

Interactions among the dark photons, as well as between the dark photons and a potential population of ϕ particles in the thermal plasma, are relevant (*i*) to assess the extent to which our assumptions that the dark matter is cold and non-interacting are self-consistent, and (*ii*) to determine the evolution of the reflected dark photons, which is the topic of section 2.4.3. For convenience, we summarize the relevant interaction rates in the remainder of this appendix.

2.A.2 Interaction rates

At low momenta, eq. (2.68) leads to a self-interaction among the dark photons of the form

$$\sigma_{VV} \sim \frac{\lambda_V^2}{8\pi m^2} = \frac{\lambda_\Phi^2}{8\pi m_\rho^2} \left(\frac{m}{m_\rho}\right)^6 \quad \text{for} \quad |\vec{k}| \ll m. \quad (2.69)$$

The corresponding dark matter self-interaction rate is always tiny compared to the Hubble rate in the early universe in the region of parameter space relevant to this work. Potentially more relevant are interactions between the dark photons and ϕ particles in the thermal fluid. In particular, processes of the form $V\phi \rightarrow V\phi$ could be efficient at transferring energy from the plasma to the dark matter, unless the corresponding interaction rates are tiny. The relevant cross section is of the form, parametrically:

$$\sigma_{V\phi} \sim \frac{\kappa^2}{8\pi T_*^2} \sim \frac{1}{8\pi T_*^2} \left(\frac{\Delta m^2}{m^2}\right)^2 \frac{m^4}{T_*^4}, \quad (2.70)$$

where $\kappa \simeq \Delta m^2/v^2$, as discussed below eq. (2.8), and we've taken $m_\phi \sim v \sim T_*$, which should be a good approximation at the epoch of the phase transition (barring a tiny quartic coupling for ϕ). Assuming a thermal population of ϕ particles, with number density $n_\phi \sim T_*^3$, we find

$$\frac{\Gamma_{V\phi}}{H(T_*)} \sim 10^{-39} \left(\frac{\Delta m^2/m^2}{10^{-4}}\right)^2 \left(\frac{m}{1 \text{ eV}}\right)^4 \left(\frac{100 \text{ GeV}}{T_*}\right)^5. \quad (2.71)$$

Clearly, this ratio can be very small in much of the parameter space of interest (see left panel in fig. 2.4).

Let us now discuss interactions involving the (highly relativistic) reflected dark photons. These will be relevant in section 2.4.3 when we discuss the fate of the dark radiation, and for convenience we summarize the relevant results here. At large center-of-mass energies, i.e. $\sqrt{s} \gg m$, the two-to-two scattering cross section among the dark photons is

dominated by the scattering of longitudinally-polarized vectors, and so it will be sufficient for our purposes to just consider the process $V_l V_l \rightarrow V_l V_l$. Two limiting kinematic regimes are of interest. In the region $m \ll \sqrt{s} \ll m_\rho$, the relevant cross section is of the form

$$\sigma_{ll} \sim \frac{\lambda_V^2}{8\pi s} \left(\frac{\sqrt{s}}{m}\right)^8 \sim \frac{\lambda_\Phi^2}{8\pi m_\rho^2} \left(\frac{\sqrt{s}}{m_\rho}\right)^6 \quad \text{for} \quad m \ll \sqrt{s} \ll m_\rho. \quad (2.72)$$

This will be relevant for scattering between the dark radiation and the cold dark photons, where the relevant center-mass-energy is given by $\sqrt{s}/2 \simeq \gamma_{\text{eq}} m$. Comparing the relevant interaction rate to $H(T_*)$, we find

$$\frac{\Gamma_{ll}}{H(T_*)} \sim 10^{-12} \frac{\lambda_\Phi^2}{\eta^4} \left(\frac{m}{1 \text{ eV}}\right)^5 \left(\frac{10^{-4}}{\Delta m^2/m^2}\right)^2 \left(\frac{T_*}{v}\right)^8 \left(\frac{100 \text{ GeV}}{T_*}\right)^4 \left(\frac{\alpha}{10^{-2}}\right)^3 \left(\frac{\rho_{\text{dm}}}{\rho_V}\right)^2, \quad (2.73)$$

where we have evaluated the cross section at $\sqrt{s}/2 = \gamma_{\text{eq}} m$, and substituted m_ρ as given in eq. (2.67). Notice that the ratio T_*/v could be $\lll 1$ if the phase transition sector is cold and decoupled from the thermal plasma, as discussed below eq. (2.7), which would further suppress the above ratio. If instead $\sqrt{s} \gg m_\rho$ then the cross section must be computed in the complete theory of eq. (2.65), and we find

$$\sigma_{ll} \sim \frac{\lambda_\Phi^2}{8\pi s} \quad \text{for} \quad \sqrt{s} \gg m_\rho, \quad (2.74)$$

as expected by virtue of Goldstone Equivalence. This is the form of the cross section used in deriving eq. (2.60).

Another important class of interactions are those between the dark radiation and the ϕ particles in the plasma when the phase transition sector is thermal. In particular, if the process $V_l \phi \rightarrow V \phi$ happens efficiently, the energy density stored in the dark radiation would quickly become distributed among the thermal fluid, as discussed in section 2.4.3.

The typical s -parameter is of the form $s \sim 2T_*\omega_{\text{dr}}$. In the regime $\sqrt{s} \ll m_\rho$,

$$\sigma_{l\phi} \sim \frac{\kappa^2}{8\pi s} \left(\frac{\sqrt{s}}{m}\right)^4 \quad \text{for} \quad m \ll \sqrt{s} \ll m_\rho, \quad (2.75)$$

whereas in the high-energy limit $\sqrt{s} \gg m_\rho$,

$$\sigma_{l\phi} \sim \frac{\eta^2}{8\pi s} \quad \text{for} \quad \sqrt{s} \gg m_\rho. \quad (2.76)$$

As an example, if there's a full thermal distribution of ϕ particles with $n_\phi \sim T_*^3$, then the relevant interaction rate, relative to $H(T_*)$ is given by

$$\frac{\Gamma_{l\phi}}{H(T_*)} \sim 10^8 \eta^2 \left(\frac{\Delta m^2}{m^2}\right)^2 \left(\frac{1 \text{ eV}}{m}\right)^2 \left(\frac{100 \text{ GeV}}{T_*}\right)^2 \left(\frac{10^{-2}}{\alpha}\right)^2 \left(\frac{\rho_V}{\rho_{\text{dm}}}\right)^2, \quad (2.77)$$

where we have used eq. (2.76) in obtaining $\Gamma_{l\phi}$, as appropriate given our choice of parameters on the right-hand-side. Clearly, interactions between the dark radiation and the thermal plasma can be very efficient in parts of the relevant parameter space but – again – this conclusion depends strongly on the choices of the various scales and couplings.

2.B Reflection probabilities – supplemental material

This appendix contains additional details supplementing the discussion of section 2.3.

2.B.1 Longitudinal reflection from Goldstone Equivalence

In the main text we described the theory of a massive vector in eqs. (2.7) and (2.13) in terms of the Proca Lagrangian, featuring three degrees of freedom for the three physical polarizations of a massive spin-1 particle. However, this is often understood as a particular choice of gauge, after gauge redundancy is restored via $V_\mu \longrightarrow V_\mu - \partial_\mu\theta$. Although

$\theta = 0$ (unitary gauge) is naively simplest, often it is convenient to keep the ‘Goldstone’ θ . The Goldstone Equivalence Theorem states that in the relativistic limit $\omega \gg m$ correct amplitudes can be obtained by identifying the physical longitudinal vector with the scalar θ , with an error of order $\mathcal{O}(m^2/\omega^2)$. Using the usual gauge-fixing parameter ξ , the theory is described by

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}m_V^2(z)(V_\mu - \partial_\mu\theta)^2 - \frac{1}{2\xi}(\partial_\mu V^\mu + m^2\xi\theta)^2 . \quad (2.78)$$

Proceeding with the equation of motion for θ :

$$\partial_\mu (m_V^2(z)\partial^\mu\theta) + \xi m^4\theta = 0 , \quad (2.79)$$

where we have neglected mixing with V^μ in order to test the spirit of the Goldstone Equivalent in our context. Working in the step function limit, the scattering solution is simple for $z \neq 0$:

$$\theta = e^{-i\omega t} \begin{cases} e^{ik_z z} + r_G e^{-ik_z z}, & z < 0 \\ t_G e^{i\tilde{k}_z z}, & z > 0 \end{cases} , \quad (2.80)$$

$$k_z = \sqrt{\omega^2 - \xi m^2} , \quad \tilde{k}_z = \sqrt{\omega^2 - \xi m^4/\tilde{m}^2} , \quad (2.81)$$

and need only be supplemented by the matching conditions:

$$m_V^2(z)\partial_z\theta \quad \text{and} \quad \theta \quad \text{are continuous at } z = 0 , \quad (2.82)$$

derived by integrating eq. (2.79) once and twice respectively. A little algebra gives the Goldstone reflection probability as

$$R_G = |r_G|^2 = \left(\frac{\tilde{k}_z \tilde{m}^2 - k_z m^2}{\tilde{k}_z \tilde{m}^2 + k_z m^2} \right)^2 \xrightarrow{\omega \gg m, \tilde{m}} \left(\frac{\tilde{m}^2 - m^2}{\tilde{m}^2 + m^2} \right)^2 + \mathcal{O}(\tilde{m}^2/\omega^2) \quad (2.83)$$

which is ξ (gauge) independent and matches the leading order result derived in the main text, eq. (2.26), consistent with the spirit of the GET. This result suggests that the enhanced reflection of the inter-relativistic limit described in this work might fundamentally be a property of Nambu-Goldstone bosons more generally, a topic that we will return to in future work.

2.B.2 Scattering on a δ' potential

As anticipated in section 2.3.3, the effective scattering potential for the longitudinal component in the limit of vanishing wall thickness takes the form

$$U_l(z) \xrightarrow{L \rightarrow 0} \kappa \delta'(z) \quad \text{with} \quad \kappa \equiv -\frac{\Delta m^2}{2m^2}, \quad (2.84)$$

and eq. (2.30) reads

$$(\partial_z^2 + k_z^2)\lambda(z) = \kappa \delta'(z)\lambda(z). \quad (2.85)$$

We can obtain matching conditions for λ and λ' by integrating (twice) over this equation.

First, integrating eq. (2.85) from $z_0 < 0$ to z , we find

$$\int_{z_0}^z d\hat{z} [(\partial_{\hat{z}}^2 + k_z^2)\lambda(\hat{z})] = \int_{z_0}^z d\hat{z} \kappa \delta'(\hat{z})\lambda(\hat{z}) \quad (2.86)$$

$$\Rightarrow \lambda'(z) - \lambda'(z_0) + k_z^2 \int_{z_0}^z d\hat{z} \lambda(\hat{z}) = \kappa \lambda(z)\delta(z) - \frac{\kappa}{2} [\lambda'(0^+) + \lambda'(0^-)] \Theta(z). \quad (2.87)$$

Integrating eq. (2.87) from $z = -\epsilon$ to $z = +\epsilon$ and taking the limit $\epsilon \rightarrow 0$, we find:

$$\lambda(0^+) - \lambda(0^-) = -\frac{\kappa}{2} [\lambda(0^+) + \lambda(0^-)]. \quad (2.88)$$

Similarly, taking $z_0 = -\epsilon$ and $z = +\epsilon$ in eq. (2.87), with $\epsilon \rightarrow 0$, one finds:

$$\lambda'(0^+) - \lambda'(0^-) \simeq \frac{\kappa}{2} [\lambda'(0^+) + \lambda'(0^-)] , \quad (2.89)$$

where in the right-hand-side we have only kept terms of $\mathcal{O}(\kappa)$. From the last two equations, one finds $r_l \simeq \kappa$, and therefore

$$R_l \simeq \kappa^2 = \left(\frac{\Delta m^2}{2m^2} \right)^2 , \quad (2.90)$$

in agreement with eq. (2.26) at leading order in $\kappa \ll 1$.

2.B.3 The Born approximation

In section 2.3.3, we found it convenient to recast the equations of motion for both transverse and longitudinal components as a one-dimensional Schrödinger equation, of the form

$$(\partial_z^2 + k_z^2)\psi(z) = U(z)\psi(z) , \quad (2.91)$$

with $U(z) \rightarrow U_l(z)$ and $\psi \rightarrow \xi$ for the longitudinal component (c.f. eq. (2.30)), and $U(z) \rightarrow U_\perp(z)$ and $\psi \rightarrow v_\perp^\mu$ for the transverse modes (c.f. eq. (2.29)). It is an important feature that the effective scattering potential vanishes in the limit $z \rightarrow -\infty$, as will become clear in due time.

As is well-known, it is possible to recast eq. (2.91) as an integral equation, as follows

$$\psi(z) = \psi_0(z) + \int_{-\infty}^{\infty} dz' G(z-z')U(z')\psi(z') , \quad (2.92)$$

where ψ_0 is any function satisfying the free particle equation, $(\partial_z^2 + k_z^2)\psi_0 = 0$, and $G(z)$

is a Green's function for the differential operator on the right-hand-side of eq. (2.91), i.e.

$$(\partial_z^2 + k_z^2)G(z) = \delta(z) . \quad (2.93)$$

It is straightforward to check that eq. (2.92) is equivalent to eq. (2.91) by applying $(\partial_z^2 + k_z^2)$ to both sides and using eq. (2.93). In order to solve for $G(z)$, note that for $z \neq 0$ eq. (2.93) is just the free particle equation. The matching conditions at $z = 0$ are that G must be continuous and G' must have unit jump. Putting this together, the solutions are

$$G(z) = \pm \frac{1}{2ik_z} e^{\pm ik_z |z|} , \quad (2.94)$$

up to addition of functions that satisfy the free particle equation. It turns out that we only need one of these two solutions: the one with the + sign.

To represent an incoming plane wave, we take $\psi_0(z) = e^{ik_z z}$. Then our integral equation is

$$\psi(z) = e^{ik_z z} + \frac{1}{2ik_z} \int_{-\infty}^{\infty} dz' e^{ik_z |z-z'|} U(z') \psi(z') . \quad (2.95)$$

So far, we have made no approximations. However, if the correction to the incoming wavefunction ψ_0 is small (meaning the reflection coefficient is tiny), then we can plug this equation for ψ into itself in the integral and truncate higher order terms, leaving

$$\psi(z) \simeq e^{ik_z z} + \frac{1}{2ik_z} \int_{-\infty}^{\infty} dz' e^{ik_z |z-z'|} U(z') e^{ik_z z'} \quad (2.96)$$

$$= e^{ik_z z} + \frac{e^{ik_z z}}{2ik_z} \int_{-\infty}^z dz' U(z') + \frac{e^{-ik_z z}}{2ik_z} \int_z^{\infty} dz' e^{2ik_z z'} U(z') . \quad (2.97)$$

We wish to extract a reflection probability from this 'first Born approximation'. To this end, consider the limit $z \rightarrow -\infty$. The second term above clearly vanishes, given that $U(z') \rightarrow 0$ as $z' \rightarrow -\infty$ as emphasized below eq. (2.91). Furthermore, the third term can

be identified as the reflected plane wave piece, with associated probability

$$R_{l, \text{Born}} = \frac{1}{4k_z^2} \left| \int_{-\infty}^{\infty} dz e^{2ik_z z} U(z) \right|^2. \quad (2.98)$$

Keep in mind that $R_l \simeq R_{l, \text{Born}}$ only when $R_l \ll 1$.

2.B.4 Numerical methods

Finding the reflection and transmission coefficients for the longitudinal mode reduces to solving eq. (2.30) with boundary conditions corresponding to plane-wave behavior far away from the bubble wall. i.e.

$$\lambda(z) = \begin{cases} e^{ik_z z} + r_l e^{-ik_z z} & \text{as } z \rightarrow -\infty \\ t_l e^{i\tilde{k}_z z} & \text{as } z \rightarrow +\infty \end{cases} \quad (2.99)$$

It will be convenient to rescale λ by an overall complex constant, and define $f(z) \equiv t_l^{-1} \lambda(z)$. The equation of motion for $f(z)$ remains as in eq. (2.30), and the boundary conditions can now be written as $f(z_0) = 1$ and $\partial_z f(z_0) = i\tilde{k}_z$ for some $z_0 > 0$ far away to the right of the wall. Our numerical solution will lie in the window $z \in [-z_0, z_0]$, where z_0 is set to be much larger than the incident particle wavelength (constraining for low energies) and the thickness of the bubble wall (constraining for high energies).

Lastly, we need a prescription for extracting the reflection probability from the numerical solution. Consider evaluating the asymptotic expression for $f(z)$ at the following

points, where $n \in \mathbb{N}$, $n \rightarrow \infty$:

$$f\left(-\frac{n\pi}{k_z}\right) \rightarrow (t_l^{-1} + t_l^{-1} r_l) (-1)^n \quad (2.100)$$

$$f\left(-\frac{(n - \frac{1}{2})\pi}{k_z}\right) \rightarrow (t_l^{-1} - t_l^{-1} r_l) (-1)^n (i) \quad (2.101)$$

These equations can then be easily inverted to obtain the scattering coefficients r_l and t_l , and ultimately the reflection and transmission probabilities.

Chapter 3

P Not PQ

3.1 Introduction

The current upper bound on the size of the neutron electric dipole moment (EDM) is $|d_n| < 1.8 \cdot 10^{-26} e \cdot \text{cm}$ [74].¹ In turn, this severely constrains the size of the QCD vacuum angle, which is required to be

$$\bar{\theta} = \theta_s + \theta_q \lesssim 10^{-10} , \quad (3.1)$$

where θ_q is the argument of the determinant of the quark mass matrix, and θ_s the coefficient of the $G\tilde{G}$ operator,

$$\mathcal{L} \supset \frac{\theta_s \alpha_s}{4\pi} \text{tr} \left(G^a \tilde{G}^a \right) . \quad (3.2)$$

In the Standard Model (SM), $\theta_q = \arg \det(y_u y_d)$, with $y_{u,d}$ the Yukawa matrices in the up- and down-quark sectors. $\bar{\theta}$ provides a physical, basis-independent measurement of

¹Here we quote the direct limit; the inferred bound $|d_n| < 1.6 \cdot 10^{-26} e \cdot \text{cm}$ from the ^{199}Hg EDM limit [75] is comparable assuming no additional contributions to the atomic EDM.

CP -violation in the strong sector of the SM.

That $\bar{\theta}$ is constrained to be so tiny is one of the most puzzling features of the SM, and it is known as the strong CP problem. It stands alongside the electroweak hierarchy problem and the cosmological constant problem as one of the three great naturalness puzzles that remain unsolved. Although numerically the strong CP problem is orders of magnitude less severe than either of its siblings, it is considerably more robust against anthropic arguments.² As such, it has drawn renewed attention during an era in which LHC null results are challenging naturalness-based approaches to the electroweak hierarchy.

Although one could argue that $\theta_s = 0$ on the basis that QCD interactions otherwise preserve CP , a similar argument cannot be made for a vanishing θ_q . For example, if CP were a good symmetry of the Yukawa sector then the Yukawa matrices would need to be real. However, real Yukawas would lead to a vanishing phase in the CKM matrix, in direct conflict with the $\mathcal{O}(1)$ CP -violation observed in the electroweak sector of the SM. Besides CP , a non-zero $\bar{\theta}$ also violates P . Again, the fact that the strong sector preserves parity may allow us to impose $\theta_s = 0$. P invariance in the Yukawa sector would require the Yukawa matrices to be hermitian, in which case $\theta_q = 0$ too, while still allowing for a non-zero CKM phase. However, the fact that P is maximally violated by the electroweak interactions severely weakens this line of reasoning as an attempt to argue for a small $\bar{\theta}$.

So although $\bar{\theta}$ is a measurement of both P and CP violation by strong dynamics, the above discussion highlights how the origin of the strong CP problem in the SM ultimately lies in the features of the *electroweak* sector. It is the fact that electroweak interactions maximally violate both P and CP that precludes an understanding of the bound in eq. (3.1) based on the underlying symmetries of the SM.

With this in mind, it is natural to attempt an understanding of the smallness of $\bar{\theta}$ in the context of theories with an extended electroweak sector. If either P or CP are

²See, however, [76] for arguments to the contrary.

good symmetries of the extended theory, then $\bar{\theta}$ will be forced to vanish. Of course, to account for the P and CP violation we observe in nature, they must eventually be broken, and a non-zero $\bar{\theta}$ will be radiatively generated. If the induced $\bar{\theta}$ is small enough, this class of theories offer a symmetry-based solution to the strong CP problem. Concrete implementations of this idea based on spontaneously broken P and CP were first proposed in [77, 78, 79] (building on [80]) and [81, 82] respectively. It is the former that will be the focus of this work.³

There is another good reason to consider solutions to strong CP based on the restoration of spacetime symmetries, namely that these may be realized as gauge symmetries in the context of string theory [86, 87]. As such, they can only be broken spontaneously and not explicitly. Depending on the scale of spontaneous symmetry breaking, the apparent lack of P and CP violation in the strong sector could therefore be fully, or partially, explained in this context. Clearly, a resolution to the strong CP problem along these lines would be especially attractive: it would allow us to understand the smallness of $\bar{\theta}$ as an accident resulting from the underlying structure of the UV-completion, as opposed to being the result of a model-building effort specially designed to address eq. (3.1).

From the bottom-up, there are a number of ways the SM can be extended to accommodate spontaneously broken P . However, in order to address the strong CP problem, a necessary feature of all of them is the presence of an $SU(2)_R$ gauge factor, as well as an extended matter content that mirrors that of the SM. Crucially, the $SU(3)$ quantum numbers of SM fermions and their mirror counterparts must be the same, making the presence of additional colored particles an irreducible feature of these models. With this extended field content, parity enforces the Yukawa couplings in the two sectors to be identical. To be phenomenologically acceptable, parity must be broken at some scale v' above the weak scale, with the additional gauge bosons and mirror quarks being suffi-

³For recent exploration along these lines, see also [83, 84, 85].

ciently heavy to evade experimental constraints. Naïvely, bounds on the mass of colored particles would seemingly require $y_u v' \gtrsim 1$ TeV [88, 89], in turn setting a lower bound $v' \gtrsim 10^8$ GeV. But a parametric separation of scales between v and v' entails an irreducible amount of fine-tuning $\Delta^{-1} \simeq 2v^2/v'^2$, which would become $\Delta^{-1} \lesssim 10^{-12}$ for such a stringent bound on v' . Considering that the goal is to naturally explain a number of $\mathcal{O}(10^{-10})$, parity would hardly seem to remain an attractive solution to strong CP.

In this chapter, we show that the conclusion of the previous paragraph is premature, and that a parity-breaking scale as low as 18 TeV is consistent with all experimental constraints. This significantly improves the level of fine-tuning, and leaves an open window for symmetry-based solutions to strong CP that are based on spontaneously broken parity. The leading constraint on the low-tuning version of these models comes not from bounds on colored particles, but from direct searches for Z' and W' resonances at the LHC [90, 91]. Future searches for heavy gauge bosons at current and future colliders are the most promising probes of this class of theories, with a 100 TeV proton collider guaranteed to make a discovery if the level of fine-tuning is better than $\Delta^{-1} \sim 10^{-5}$ [92, 93]. Overall, the viability of these parity-based models makes collider experiments a central testing ground for solutions to strong CP.

Another attractive feature of this class of solutions to the strong CP problem is that they are robust against the effects of symmetry-breaking higher dimensional operators (HDOs) that may arise from short-distance physics associated with a gravitational UV completion. If parity is a gauge symmetry of the underlying theory, we are led to consider only those HDOs proportional to the source of spontaneous symmetry breaking. On the other hand, if parity were global, the expectation that quantum gravity violates all global symmetries [94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106] suggests we should include all HDOs that explicitly violate P . Although the nature of the operators is different in the gauge and global implementations, the conclusion will be the same: in

both cases, the leading HDOs with $\mathcal{O}(1)$ coefficients may be present without destabilizing the solution to strong CP.

This stands in stark contrast with the reality of what has traditionally been the most popular solution to the strong CP problem: the QCD axion [107, 108, 109, 110, 111, 112, 113]. In this case, the parameter $\bar{\theta}$ is promoted to the status of dynamical field, the axion, which is a pseudo-Nambu-Goldstone boson of a spontaneously broken $U(1)_{PQ}$ *global* symmetry. A potential for the axion is induced non-perturbatively by QCD dynamics, and its vacuum expectation value (vev) adjusts such that $\bar{\theta} = 0$, thereby solving strong CP. To work, the QCD axion potential must dominate to 1 part in 10^{10} , overwhelming any other contributions that may arise from additional degrees of freedom. New dynamics responsible for, say, dark matter, baryogenesis, or addressing the hierarchy problem, cannot significantly contribute to the axion potential. Similarly, Planck-suppressed HDOs that break $U(1)_{PQ}$ must be exceptionally suppressed [114, 115, 116, 117]. The mechanism is not robust. The need for $U(1)_{PQ}$ to be a high quality global symmetry has become known as the “axion quality problem”. Attempts to turn the QCD axion into a high quality axion are valuable [118, 119, 120, 121, 122, 123, 124, 125], but hardly helpful in making a small $\bar{\theta}$ appear natural.

The goal of this work is to identify the most natural parity-based solution to the strong CP problem, and highlight its experimental consequences. We do so by following a strategy that combines the traditional notion of naturalness with the expectation that gravity violates all global symmetries. The former singles out a specific implementation of the spectrum of parity-symmetric models, and underscores the central role of collider experiments in exploring solutions to strong CP. The latter opens up an entirely new avenue of exploration for parity solutions to the strong CP problem, ranging from EDM experiments to gravitational wave observatories, depending on the degree to which the symmetry remains approximate.

To this end, this chapter is organized as follows. In section 3.2 we review the main features of parity-based solutions to strong CP, and discuss how a low symmetry breaking scale can be realized while complying with experimental constraints. We focus on the main phenomenological signatures of these models that are relevant for collider and flavor experiments in section 3.3. In section 3.4, we discuss the size of radiative corrections to both $\bar{\theta}$ and the EDM of elementary fermions, including charged leptons, depending on the details of the parity-breaking sector. We explore the effect of Planck-suppressed HDOs on this class of models in section 3.5, paying special attention to a potential gravitational wave signal from the spontaneous breaking of parity. Section 3.6 contains our conclusions. Finally, a series of appendices contain results that have been crucial in our analysis, but may be skipped on a first reading of the chapter.

3.2 P to solve strong CP

In this section, we introduce the main features of symmetry-based solutions to the strong CP problem based on parity. In section 3.2.1 we review the basic idea, as first introduced in [77, 78, 79]. We focus on the scalar potential in section 3.2.2, with an emphasis on the implications for fine-tuning of the weak scale that arise as a result of the breaking of parity. In section 3.2.3 we discuss how the scale of additional colored particles can be decoupled from the parity-breaking scale, in turn minimizing the level of fine-tuning.

3.2.1 Parity as a solution to the strong CP problem

A symmetry-based solution to the strong CP problem based on parity requires extending the SM both in terms of matter content and gauge interactions. The minimal

implementation of this idea is based on the gauge group

$$SU(3) \times SU(2)_L \times SU(2)_R \times U(1)_{\hat{Y}}, \quad (3.3)$$

as well as a doubling of the matter content of the SM into a ‘mirror’ sector with identical quantum numbers, except that $SU(2)_L$ doublets are now doublets of $SU(2)_R$. Table 1 summarizes the gauge charges in the quark and Higgs sectors of the theory. (Analogous charge assignments apply in the lepton sector, which we don’t make explicit.) Crucially, the Higgs sector of the theory does not introduce additional sources of CP -violation — indeed, the freedom to perform both $SU(2)_L$ and $SU(2)_R$ gauge transformations allows us to expand around a vacuum where both vevs are real.⁴

	$Q = \begin{pmatrix} u \\ d \end{pmatrix}$	U^\dagger	D^\dagger	H	$Q'^\dagger = \begin{pmatrix} u'^\dagger \\ d'^\dagger \end{pmatrix}$	U'	D'	H'^*
$SU(3)$	3	3	3	\cdot	3	3	3	\cdot
$SU(2)_L$	2	\cdot	\cdot	2	\cdot	\cdot	\cdot	\cdot
$SU(2)_R$	\cdot	\cdot	\cdot	\cdot	2	\cdot	\cdot	2
$U(1)_{\hat{Y}}$	$\frac{1}{6}$	$\frac{2}{3}$	$-\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{2}{3}$	$-\frac{1}{3}$	$\frac{1}{2}$

Table 1: Quantum numbers in the quark and Higgs sectors. Mirror sector fields are distinguished with a prime. We use notation such that all of Q , U , and D (as well as their mirror counterparts Q' , U' , and D') are left-handed, two-component Weyl fermions, whereas daggered fields are always right-handed.

With this additional field content, the theory admits an alternative definition of parity that combines the action of the ‘ordinary’ parity transformation with an internal symmetry that exchanges the fields of the SM and mirror sectors. Explicitly, in the

⁴The model we have just introduced is *not* the minimal parity-symmetric extension of the SM, but rather the minimal extension that solves the strong CP problem. The most minimal extension of the SM that admits a generalized parity symmetry was introduced in [126], and does not require a doubling of the fermion sector. The scalar sector, however, requires the addition of an additional Higgs field transforming as a bifundamental of the two $SU(2)$ factors, in order to write Yukawa couplings in a gauge invariant fashion. As discussed in [79], the vev of this field will in general be complex, precluding a vanishing θ_q .

gauge, quark, and Higgs sectors:

$$\mathbf{W}_L^\mu \leftrightarrow \mathbf{W}_{R\mu}, \quad (3.4)$$

$$Q, U, D \leftrightarrow Q'^\dagger, U'^\dagger, D'^\dagger, \quad (3.5)$$

$$H \leftrightarrow H'^*, \quad (3.6)$$

and similarly for leptons. Since $SU(3)$ and $U(1)_{\hat{Y}}$ interactions are not mirrored, the corresponding gauge fields transform as usual under parity. Unlike ordinary parity in the SM, this ‘generalized’ parity transformation is now a good symmetry of the gauge sector of the theory, thanks to the extended electroweak sector and matter content.

In this context, the strong CP problem is solved as follows. On the one hand, parity requires that $\theta_s = 0$, just as one may argue in the SM based on the properties of the strong sector alone. On the other hand, the presence of additional colored particles results in an extended quark mass matrix. In particular, Yukawa terms can be written for both the SM and mirror sectors, of the form

$$\mathcal{L} \supset - \left\{ (y_u)_{ij} Q_i H U_j + (y'_u)_{ij} Q_i'^\dagger H'^* U_j'^\dagger \right\} + \text{h.c.}, \quad (3.7)$$

and similarly for down-type quarks and leptons. As a result, the tree-level value of θ_q in these models is given by

$$\theta_q = \arg \det(y_u y_d) + \arg \det(y'_u{}^* y'_d{}^*). \quad (3.8)$$

Crucially, demanding that Yukawa interactions preserve parity, which is now a good symmetry of the extended electroweak sector, enforces the Yukawa couplings in the two

sectors to be identical, i.e.

$$y'_f = y_f. \quad (3.9)$$

In turn, this implies $\theta_q = 0$, as per eq. (3.8), forcing $\bar{\theta}$ to vanish at tree-level in parity-symmetric models.

With the field content outlined in table 1, the theory admits an additional fermion mass term involving only the $SU(2)$ -singlets, of the form

$$\mathcal{L} \supset -(\mathcal{M}_u)_{ij} U_i U'_j + \text{h.c.} \quad (3.10)$$

(with analogous terms for down-type quarks and leptons), where invariance under generalized parity requires that the vector-like mass matrix be hermitian, i.e. $\mathcal{M}_f^\dagger = \mathcal{M}_f$.⁵

In general, the expression for θ_q can be conveniently written as

$$\theta_q = \arg \det(\mathbb{M}_u \mathbb{M}_d), \quad (3.11)$$

where \mathbb{M}_u and \mathbb{M}_d are 6×6 matrices, of the form

$$\mathbb{M}_f = \begin{pmatrix} 0 & \frac{v'}{\sqrt{2}} y'_f{}^* \\ \frac{v}{\sqrt{2}} y_f^T & \mathcal{M}_f \end{pmatrix}, \quad \text{for } f = u, d. \quad (3.12)$$

Due to the zero in the upper-left block of the overall 6×6 mass matrix, the expression for θ_q remains as in eq. (3.8). As we will discuss in section 3.2.3, the presence of vector-like masses is crucial in implementing a version of the model with low fine-tuning.⁶

⁵Note that a non-hermitian mass matrix is compatible with softly broken parity; we will explore the consequences of such soft breaking in section 3.4.1.

⁶A variation on the model we have so far discussed entails extending the gauge group in eq. (3.3) with an additional $U(1)$, as first discussed in [79]. In this case, SM and mirror fields are charged under different $U(1)$ factors, which transform into each other under parity. Although this seems like a minimal modification of the model presented here, this two- $U(1)$ version does not allow for the vector-like mass terms of eq. (3.10), in turn precluding the implementation of a low parity-breaking scale.

To obtain a phenomenologically viable model, parity must be broken, with different vevs in the mirror and SM Higgs sectors. This will induce a non-vanishing $\bar{\theta}$ beyond tree-level, which must be small enough if the theory is to remain a bona-fide solution to strong CP. The size of radiative corrections depends on the details of how parity is spontaneously broken. If P is broken without breaking CP , then the radiatively induced $\bar{\theta}$ will be no larger than in the SM [79], where $\bar{\theta} < 10^{-19}$ [127]. On the other hand, if CP is also spontaneously broken (e.g. through the vev of a pseudo-scalar) then a larger $\bar{\theta}$, as well as a neutron EDM *independent* of $\bar{\theta}$, may be radiatively generated. Even in this latter case, we will see that radiative corrections can be small enough to remain compatible with experimental constraints. Given that the final size of $\bar{\theta}$ is a somewhat model-dependent feature of this class of models, we defer a more detailed discussion of this issue to section 3.4.

More generally, discussing the leading effect of broken parity on the fine-tuning of the electroweak sector does not require committing to a specific implementation of spontaneous symmetry breaking. For this purpose, it will be enough to focus on the features of the Higgs sector, to which we now turn.

3.2.2 Scalar sector and fine-tuning

For the time being, we will parametrize the necessary breaking of parity through an explicit soft term in the scalar potential. Of course, such soft breaking should ultimately be the result of some spontaneous symmetry breaking dynamics, as we will make more explicit in section 3.4. In this spirit, the most general scalar potential involving the $SU(2)_L$ and $SU(2)_R$ Higgs doublets takes the form

$$V(H, H') = -m_H^2(|H|^2 + |H'|^2) + \lambda(|H|^2 + |H'|^2)^2 + \kappa(|H|^4 + |H'|^4) + \mu^2|H|^2. \quad (3.13)$$

At this level, eq. (3.13) is identical to the scalar potential of theories of Neutral Naturalness, such as Twin Higgs [128, 129]. The first two terms respect both parity and a larger accidental $SU(4)$ (really, $O(8)$) symmetry, while κ respects the former but not the latter. The parameter μ^2 softly breaks parity. In the interest of a non-trivial vacuum structure, we take $m_H^2 > 0$. Depending on the relative signs and sizes of the quartic couplings, the tree-level vacua for $\mu^2 = 0$ either preserve parity (with $v' = v$) or spontaneously break parity (with $v \neq 0, v' = 0$ or $v = 0, v' \neq 0$). A vacuum with $v' \gg v \neq 0$ may be obtained by deforming the theory away from the parity-symmetric vacuum with nonzero μ^2 and $\lambda, \kappa > 0$. At tree-level, the vevs in the SM and mirror Higgs sectors are then given by

$$v^2 = \frac{m_H^2 - \mu^2(1 + \lambda/\kappa)}{2\lambda + \kappa}, \quad \text{and} \quad v'^2 = \frac{m_H^2 + \mu^2\lambda/\kappa}{2\lambda + \kappa}, \quad (3.14)$$

where $v^2 \ll v'^2$ is necessary in order to obtain a phenomenologically viable model. After spontaneous symmetry breaking, the spectrum of the theory contains two scalar fields, h and h' , with masses $m_h \simeq 2\sqrt{\kappa}v$ and $m_{h'} \simeq \sqrt{2\lambda}v'$ respectively, as well as six Goldstones that become the longitudinal components of the gauge bosons of our extended electroweak sector. The physical gauge boson spectrum contains Z' and W' resonances, which are heavier than their SM counterparts by a factor of v'/v . We defer further details of the scalar and gauge sectors to section 3.A.1.

It is clear from eq. (3.14) that to obtain a hierarchy of scales between v and v' we need to introduce a tree-level tuning between the parity-preserving and parity-breaking mass-squared terms. Heuristically, the necessary fine-tuning is given by

$$\Delta^{-1} \equiv \frac{m_H^2 - \mu^2(1 + \lambda/\kappa)}{m_H^2} \simeq \frac{(2\lambda + \kappa)v^2}{(\lambda + \kappa)v'^2} \simeq \frac{2v^2}{v'^2}. \quad (3.15)$$

This is an irreducible contribution to the fine-tuning in this class of models. Insofar

as it involves the sensitivity of the weak scale v to underlying parameters, it may be classified as a tuning associated with the electroweak hierarchy problem, although it is not necessarily the only such contribution. For example, a hierarchy of scales $v'^2 \ll M_{Pl}^2$ would constitute an additional source of fine-tuning in the absence of a stabilizing mechanism. Similarly, additional hierarchies of scales or couplings in the sector responsible for spontaneously breaking P might necessitate similar accurate cancellations. However, these are issues that could, at least *in principle*, be addressed at some higher scale above v' , provided the necessary dynamics do not spoil the smallness of $\bar{\theta}$ [130]. In contrast, eq. (3.15) is forced on us independently of the UV-completion. Although it is tempting to attach the tuning in eq. (3.15) to the electroweak hierarchy problem and attribute it to anthropic selection (the perspective advocated in e.g. [84, 85]), this necessarily entails some favorable assumptions about the properties of an anthropic landscape. Here we prefer to render unto strong CP the things that are strong CP's, and take the irreducible tuning in eq. (3.15) at face value as a measure of the degree to which a parity model naturally explains the small value of $\bar{\theta}$ without reintroducing tuning elsewhere.

With this in mind, in this chapter we focus on implementations of parity solutions to strong CP where the level of fine-tuning, as parametrized in eq. (3.15), is as mild as possible. This will concentrate our attention on a specific mechanism to generate fermion masses that in turn endows these models with characteristic phenomenology, as we discuss next.

3.2.3 Fermion masses and a low parity-breaking scale

Fermion mass terms arise from the Yukawa couplings of eq. (3.7), as well as from the vector-like mass involving the $SU(2)$ -singlets. In total:

$$\mathcal{L} \supset - \left\{ \frac{v}{\sqrt{2}} (y_u)_{ij} u_i U_j + \frac{v'}{\sqrt{2}} (y_u)_{ij}^* u'_i U'_j + (\mathcal{M}_u)_{ij} U_i U'_j \right\} + \text{h.c.}, \quad (3.16)$$

where we have already set $y'_u = y_u$, as mandated by generalized parity, and analogous mass terms are present both for down-type quarks and leptons.

The structure of eq. (3.16) allows for two limiting realizations of the fermion spectrum. If the overall scale of the vector-like mass matrix is $M \ll v, v'$, then fermion masses are generated mainly through the Yukawa terms, as in the SM. In this case, mirror fermions would be an exact copy of the SM, just heavier by a factor of v'/v . Demanding that the lightest mirror quark is heavy enough to comply with current experimental constraints requires $m_u \times v'/v \gtrsim 1$ TeV [88, 89], in turn setting a lower bound $v' \gtrsim 10^8$ GeV. As advertised in the Introduction, this sets the level of fine-tuning in the electroweak sector to $\Delta^{-1} \lesssim 10^{-12}$. The phenomenology of parity solutions to strong CP in this regime was discussed recently in [84].

On the other hand, the limit $M \gg v, v'$, allows for a see-saw realization of the fermion spectrum, consisting of three light (SM-like) fermions, and three heavy fermions with mass of order M . A sufficiently high scale for the mass of additional colored particles can now be achieved by increasing M , not v' . This allows for a much lower parity-breaking scale, and therefore a much better level of fine-tuning. See-saw implementations of fermion masses, for both quarks and leptons, are discussed in [131, 132, 133, 134], and it was in fact in this context that a parity-based solution to the strong CP problem was first proposed [77, 78]. It is this second realization of the fermion spectrum that we

concentrate on in this work.⁷

The up-quark sector requires special consideration, since the see-saw mechanism cannot be applied to the top quark while maintaining perturbative Yukawas. So let us discuss the down-quark and lepton sectors first. (We will use notation appropriate to the down-quark sector, but emphasize that the same results apply for leptons.) To leading order in both v/M and v'/M , the masses of the light and heavy fermions are obtained by diagonalizing the 3×3 hermitian matrices

$$\frac{vv'}{2}y_d^*\mathcal{M}_d^{-1}y_d^T, \quad \text{and} \quad \mathcal{M}_d, \quad (3.17)$$

respectively. We make the simplifying assumption that there are no significant hierarchies in the eigenvalues of \mathcal{M}_d , and therefore the heavy quarks appear at a common scale $\sim M$. Parametrically, light quark masses are then of the form $m_{d_i} \sim |y|^2 vv'/M$. The see-saw mechanism generates fermion masses $m_{d_i} \ll v$ while allowing for much larger Yukawa couplings than in the SM, which is obviously one of the main attractions of this class of models. Generating the b quark mass through the see-saw mechanism while maintaining perturbativity sets an upper bound on the ratio M/v' , parametrically:

$$m_b \sim |y|^2 \frac{vv'}{M} \lesssim \frac{vv'}{M} \quad \Rightarrow \quad \frac{M}{v'} \lesssim \frac{v}{m_b} \sim 10^2. \quad (3.18)$$

Rotating from the flavor to the mass eigenbasis in the fermion sector can be conveniently performed step by step at each order in perturbation theory, and we present a detailed discussion of this procedure in section 3.A.2. At zeroth order in $v^{(i)}/M$, it is necessary to perform unitary transformations acting separately on the $SU(2)$ -singlet and

⁷The vector-like masses of eq. (3.16) provide a soft breaking of generalized parity, precluding the existence of degrees of freedom that would be stabilized by the internal part of this symmetry. As a result, these models do not feature a natural dark matter candidate whose presence is linked to the resolution of the strong CP problem, unlike some of the viable parameter space of the QCD axion.

doublet fields, of the form:

$$d \rightarrow \mathcal{O}_d^\dagger d, \quad d' \rightarrow \mathcal{O}_d^T d', \quad \text{and} \quad D' \rightarrow \mathcal{O}_{D'}^\dagger D', \quad D \rightarrow \mathcal{O}_{D'}^T D. \quad (3.19)$$

\mathcal{O}_d and $\mathcal{O}_{D'}$ are 3×3 unitary matrices acting on flavor space that diagonalize the first and second matrices of eq. (3.17), respectively. At first order in $v^{(l)}/M$, a further rotation is required that mixes the $SU(2)$ -singlet and doublet fields as follows:

$$\begin{pmatrix} d \\ D' \end{pmatrix} \rightarrow \begin{pmatrix} \mathbb{1}_3 & \epsilon_d^\dagger \\ -\epsilon_d & \mathbb{1}_3 \end{pmatrix} \begin{pmatrix} d \\ D' \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} d' \\ D \end{pmatrix} \rightarrow \begin{pmatrix} \mathbb{1}_3 & \epsilon_d'^\dagger \\ -\epsilon_d' & \mathbb{1}_3 \end{pmatrix} \begin{pmatrix} d' \\ D \end{pmatrix}, \quad (3.20)$$

where ϵ_d and ϵ_d' are 3×3 matrices with entries of $\mathcal{O}(v/M)$ and $\mathcal{O}(v'/M)$ respectively, and whose explicit expressions are given in eq. (3.95).

Using Dirac notation, the left- and right-handed components of the light and heavy mass eigenstates are then given by

$$d_{iL} = \begin{pmatrix} d_i \\ 0 \end{pmatrix}, \quad d_{iR} = \begin{pmatrix} 0 \\ d_i^\dagger \end{pmatrix}, \quad \text{and} \quad D_{iL} = \begin{pmatrix} D'_i \\ 0 \end{pmatrix}, \quad D_{iR} = \begin{pmatrix} 0 \\ D_i^\dagger \end{pmatrix}. \quad (3.21)$$

In particular, notice that the right-handed components of the light (SM-like) fermions consist of the corresponding component of the $SU(2)_R$ -doublets, up to corrections of $\mathcal{O}(v'/M)$. This feature plays a crucial role in the phenomenology of these models. In particular, it leads to unsuppressed couplings between $SU(2)_R$ gauge bosons, and the right-handed currents of the SM-like fermions. As we will discuss in section 3.3.1, this leads to the most stringent bound on the parity-breaking scale.

As far as the up-quark sector is concerned, the see-saw mechanism can be implemented for the u and c quarks, with the corresponding heavy partners appearing at the scale $\sim M$.

The mass eigenstates for the first two generations are as in eq. (3.21). The top sector, on the other hand, cannot be significantly “see-sawed”. Instead, it consists of light and heavy top partners with tree-level masses $m_t \simeq y_t v / \sqrt{2}$ and $m_{t'} \simeq m_t \times v' / v$, respectively. In Dirac notation, and at zeroth order in $v^{(i)}/M$, the mass eigenstates are now purely made of SM and mirror sector fields, i.e.

$$t_L = \begin{pmatrix} u_3 \\ 0 \end{pmatrix}, \quad t_R = \begin{pmatrix} 0 \\ U_3^\dagger \end{pmatrix}, \quad \text{and} \quad t'_L = \begin{pmatrix} U'_3 \\ 0 \end{pmatrix}, \quad t'_R = \begin{pmatrix} 0 \\ u'_3{}^\dagger \end{pmatrix}. \quad (3.22)$$

As usual, rotation matrices in the quark sector are constrained by the requirement that the CKM matrix is reproduced appropriately, which in this case implies $V = \mathcal{O}_u \mathcal{O}_d^\dagger$, up to corrections of $\mathcal{O}(v^2/M^2)$. Further details concerning the mass diagonalization procedure in the fermion sector can be found in section 3.A.2.

This finalizes our discussion of the main characteristics of parity solutions to strong CP that feature low fine-tuning in the electroweak sector. Before moving on, we include in fig. 1 a schematic representation of the typical spectrum of these models. Amusingly, the combination of parity and the see-saw mechanism leads to a spectrum of partner particles strikingly reminiscent of a “natural” left-right Twin Higgs model [129, 135] with light top and W/Z partners.

3.3 Dial P for Phenomenology

We now turn to the phenomenology of natural parity solutions to strong CP, beginning with direct bounds from the LHC in section 3.3.1 before turning to indirect flavor constraints in section 3.3.2. The collider and flavor phenomenology of similar left-right models has been the topic of previous work [77, 78, 134, 136, 135, 137, 138], and our focus here will be on those “irreducible” signatures that are mandated by the structure




Figure 1: Schematic illustration of the particle spectrum of parity solutions to the strong CP problem in their least tuned version. The lightest exotic particles are W' and Z' resonances, followed by a top partner appearing at a scale of order v' . A mirror Higgs is also part of the low-lying spectrum, with $m_{h'} \simeq \sqrt{2\lambda}v'$. (For illustration, we have chosen $\lambda = \mathcal{O}(1)$, but note that h' could be much lighter if $\lambda \ll 1$.) Additional mirror quarks have masses of order the see-saw scale, $M \gg v'$. The lepton sector must also be “see-sawed”, with mirror leptons similarly appearing well above v' (although we emphasize that the see-saw scales in the quark and lepton sectors need not coincide).

of the theory in its capacity as a solution to strong CP. A more in depth study of collider and flavor signatures in light of forthcoming data can illuminate the additional structure of these models, and it is a worthwhile direction for continued study.

3.3.1 Collider bounds

The doubling of the electroweak sector gives rise to a plethora of experimental signatures at colliders, ranging from additional vector bosons (the W' and Z' of spontaneously broken $SU(2)_R$) to vector-like quarks (the $SU(2)$ -singlet fermions) to additional Higgs bosons. Ultimately, given that the Z' and W' gauge bosons acquire masses exclusively from $SU(2)_R$ breaking, and inherit couplings to the SM-like quarks and leptons, collider searches for these additional vectors place the most solid and strongest direct bounds on the models under consideration.

Neutral currents

The Z' resonance inherits couplings to both the left- and right-handed currents of SM-like fermions. In the down-type quark and lepton sectors, these are both flavor diagonal and generation universal. After rotating to the mass eigenbasis in both the gauge and fermion sectors, as outlined in section 3.A, we find:

$$\mathcal{L} \supset gZ'_\mu \sum_{i=1}^3 (z'_{d_R} \bar{d}_{iR} \gamma^\mu d_{iR} + z'_{e_R} \bar{e}_{iR} \gamma^\mu e_{iR} + z'_{\nu_R} \bar{\nu}_{iR} \gamma^\mu \nu_{iR}) + \{R \rightarrow L\}. \quad (3.23)$$

As discussed in section 3.2.3, the see-saw implementation of fermion masses leads to unsuppressed couplings between the $SU(2)_R$ gauge bosons and right-handed fermions. Up to corrections of $\mathcal{O}(\sin^2 \theta_w)$, these are identical to the couplings between the SM Z and left-handed currents. Specifically:

$$z'_{d_R} \simeq -\frac{g}{2} (1 + \mathcal{O}(s_w^2)), \quad z'_{e_R} \simeq -\frac{g}{2} (1 + \mathcal{O}(s_w^2)), \quad \text{and} \quad z'_{\nu_R} \simeq \frac{g}{2} (1 + \mathcal{O}(s_w^2)), \quad (3.24)$$

where $s_w \equiv \sin \theta_w$, and we have ignored corrections of $\mathcal{O}(v^2/v'^2)$. On the other hand, couplings of the Z' to left-handed currents are now suppressed:

$$z'_{d_L} \simeq -\frac{g}{6} \frac{s_w t_w}{\sqrt{\cos 2\theta_w}} = \mathcal{O}(s_w^2), \quad \text{and} \quad z'_{e_L} = z'_{\nu_L} = -3z'_{d_L}. \quad (3.25)$$

The situation in the up-quark sector is somewhat different. Now, couplings between the Z' and the right-handed fermion currents are no longer universal. Instead, we have:

$$\mathcal{L} \supset gZ'_\mu \left(\sum_{i=1}^3 z'_{u_L} \bar{u}_{iL} \gamma^\mu u_{iL} + \sum_{i=1}^2 z'_{u_R} \bar{u}_{iR} \gamma^\mu u_{iR} + z'_{t_R} \bar{t}_R \gamma^\mu t_R \right). \quad (3.26)$$

As before, Z' couplings to first and second generation right-handed currents are unsuppressed, and are given by

$$z'_{u_R} \simeq \frac{g}{2} (1 + \mathcal{O}(s_w^2)) , \quad (3.27)$$

whereas those to left-handed fermions, as well as to the right-handed top, now read

$$z'_{u_L} = \frac{z'_{t_R}}{4} \simeq -\frac{g}{6} \frac{s_w t_w}{\sqrt{\cos 2\theta_w}} = \mathcal{O}(s_w^2) . \quad (3.28)$$

Bounds on the Z' mass from its production at the LHC will therefore be similar to those found in the so-called Sequential Standard Model, which features a Z' resonance that is just a heavy copy of the SM Z . In the present model, couplings of the Z' to SM fermions are similar to those of the Z after the replacement $L \leftrightarrow R$ — a replacement that does not affect either the production cross section or the decay rates into light fermions. The most constraining limits thus come from [90], where a search focused on leptonic final states sets a lower bound $m_{Z'} \gtrsim 5$ TeV. In turn, this translates into a lower limit on the scale of parity breaking of order $v' \gtrsim 13$ TeV.

Charged currents

W' gauge bosons interact with the right-handed SM fermions in a way that mirrors the interactions between their left-handed counterparts and SM W . In the lepton sector:

$$\mathcal{L} \supset \frac{g}{\sqrt{2}} \sum_{i,j=1}^3 (B_{ij} W_\mu^+ \bar{\nu}_{iL} \gamma^\mu e_{jL} + B'_{ij} W_\mu^+ \bar{\nu}_{iR} \gamma^\mu e_{jR}) + \text{h.c.}, \quad (3.29)$$

where $B = B' = \mathcal{O}_\nu \mathcal{O}_e^\dagger$, up to corrections of order $v^{(\prime)2}/M^2$. As far as the quark sector is concerned, the up-type sector again requires special consideration. We find:

$$\mathcal{L} \supset \frac{g}{\sqrt{2}} \sum_{i,j=1}^3 W_\mu^+ V_{ij} \bar{u}_{iL} \gamma^\mu d_{jL} + \text{h.c.}, \quad (3.30)$$

with $V = \mathcal{O}_u \mathcal{O}_d^\dagger + \mathcal{O}(v^2/M^2)$, whereas

$$\mathcal{L} \supset \frac{g}{\sqrt{2}} W_\mu^+ \sum_{j=1}^3 \left(\sum_{i=1}^2 V'_{ij} \bar{u}_{iR} \gamma^\mu d_{jR} + \Delta V'_{3j} \bar{t}_R \gamma^\mu d_{jR} \right) + \text{h.c.} \quad (3.31)$$

Up to corrections of $\mathcal{O}(v^2/M^2)$, we have $V' = V$, and $\Delta V'_{3j} = (\epsilon_u'^* V)_{3j}$. The 3×3 matrix ϵ_u' , whose entries are suppressed by a factor of $v'/M \ll 1$, is given explicitly in eq. (3.101).

As with the Z' , we expect bounds on the W' to be comparable to those in the Sequential Standard Model. Current direct searches set stringent constraints on such W' resonances, of order $m_{W'} \gtrsim 6$ TeV [91]. In turn, this sets the strongest limit on the scale of parity breaking: $v' \gtrsim 18$ TeV. Although direct searches for vector-like quarks and additional Higgs bosons are also germane, they lead to significantly weaker bounds on the scale of parity breaking compared to W' and Z' searches. For example, null results in searches for vector-like top partners [88, 89] lead to $v' \gtrsim 2$ TeV, with comparable bounds coming from searches for SM-singlet scalars.

Looking to the future, a 100 TeV pp collider such as the proposed FCC-hh should be sensitive to W' and Z' bosons as heavy as ~ 40 TeV [92, 93], corresponding to $v' \gtrsim 120$ TeV. This would comprehensively cover the most natural parameter space consistent with current data, and the non-observation of heavy vectors at such a collider would suggest that parity solutions are tuned at the $\Delta^{-1} \sim 10^{-5}$ level. In this respect, future colliders provide a decisive test of parity solutions to the strong CP problem.

3.3.2 Flavor constraints

In the SM, flavor-changing neutral currents (FCNCs) are absent at tree-level, appearing only at one-loop, and being additionally suppressed by the GIM mechanism. As a result, precision measurements of flavor-violating processes often imply stringent constraints on extensions of the SM. In the class of models under consideration, FCNCs arise already at tree-level, mediated by the Z and Z' gauge bosons, as well as the scalars h and h' . However, their size is suppressed by factors of the Yukawa couplings of the relevant fermions, making their effect negligible. At one-loop, FCNCs proceeding via box diagrams involving W' gauge bosons and mirror up-type quarks can lead to deviations in kaon properties, in turn setting the leading constraints on the flavor structure of these models.

Tree-level FCNCs

Rotating from the gauge to the mass eigenbasis in the fermion and gauge boson sectors, as specified in section 3.A, leads to the presence of flavor-changing interactions between the Z and the SM-like fermions. For example, in the down-quark sector there are new interactions of the form

$$\mathcal{L} \supset \frac{g}{2c_w} (\epsilon_d^\dagger \epsilon_d)_{ij} Z_\mu \bar{d}_{iL} \gamma^\mu d_{jL}, \quad (3.32)$$

where ϵ_d is a 3×3 matrix acting on flavor space whose explicit form is given in eq. (3.95). Integrating out the Z , the effective hamiltonian relevant to describe $|\Delta F| = 1$ processes, such as the leptonic decay of B mesons, now contains additional terms, of the form

$$\Delta \mathcal{H}_{\text{eff}} \simeq -\sqrt{2} G_F (\epsilon_d^\dagger \epsilon_d)_{32} \cos(2\theta_w) (\bar{b}_L \gamma^\mu s_L) (\bar{\mu}_L \gamma_\mu \mu_L) + \text{h.c.} \quad (3.33)$$

(An analogous term involving right-handed muons is also present, but suppressed by a factor of s_w^2 , so we neglect it in the subsequent discussion.)

The deviation with respect to the SM prediction for the branching fraction of the process $B_s^0 \rightarrow \mu^+ \mu^-$ as a result of the operator in eq. (3.33) can be written as

$$r_{\mu\mu} \equiv \frac{\text{BR}(B_s^0 \rightarrow \mu^+ \mu^-)_{\text{BSM}}}{\text{BR}(B_s^0 \rightarrow \mu^+ \mu^-)_{\text{SM}}} - 1 \simeq \frac{|C_{10}^{(\text{SM})} + C_{10}^{(\text{BSM})}|^2}{|C_{10}^{(\text{SM})}|^2} - 1, \quad (3.34)$$

where $C_{10}^{(\text{SM})}$ and $C_{10}^{(\text{BSM})}$ are the SM and BSM contributions to the Wilson coefficient of the four-fermion operator $(\bar{b}_L \gamma^\mu s_L)(\bar{\mu} \gamma_\mu \gamma^5 \mu)$. In the SM

$$C_{10}^{(\text{SM})} = \frac{G_F}{2\sqrt{2}} \frac{\alpha}{4\pi} (V_{tb}^* V_{ts}) \tilde{C}_{10}^{(\text{SM})}, \quad (3.35)$$

with $\tilde{C}_{10}^{(\text{SM})} \simeq 4.41$ [139], whereas from eq. (3.33) we have

$$C_{10}^{(\text{BSM})} \simeq \frac{G_F}{\sqrt{2}} \cos(2\theta_w) (\epsilon_d^\dagger \epsilon_d)_{32}. \quad (3.36)$$

A stringent upper bound on the size of the $(\epsilon_d^\dagger \epsilon_d)_{32}$ coefficient arises from the requirement that the masses of the down-type quarks are correctly reproduced in this model. From eq. (3.95), we have

$$(\epsilon_d^\dagger \epsilon_d)_{32} = \frac{v^2}{2} \sum_i \frac{(\tilde{y}_d)_{3i} (\tilde{y}_d)_{2i}^*}{m_{D_i}^2} \sim \frac{v^2}{M^2} \sum_i (\tilde{y}_d)_{3i} (\tilde{y}_d)_{2i}^* \lesssim \frac{v}{M} \frac{\sqrt{m_b m_s}}{v'} \ll 1, \quad (3.37)$$

where in the last step we have made use of the upper bound in eq. (3.94). We then have, parametrically

$$r_{\mu\mu} \sim \frac{2|C_{10}^{(\text{BSM})}|}{|C_{10}^{(\text{SM})}|} \lesssim 10^{-3} \left(\frac{18 \text{ TeV}}{v'} \right)^2 \left(\frac{v'}{M} \right). \quad (3.38)$$

This effect is much smaller than the theoretical and experimental errors on $\text{BR}(B_s^0 \rightarrow$




Figure 2: Additional box diagrams contributing to kaon mixing in the models under consideration include any of the up-type quarks propagating inside the loop, as well as (left) two W bosons, (center) one W and one W' , and (right) two W' s. Quantitatively, it is diagrams with one W and one W' that are the most relevant.

$\mu^+\mu^-$), which are both on the order of 10% [140].⁸

The effects of Z -mediated FCNCs on other processes are even more suppressed. For example, $|\Delta F| = 2$ processes such as kaon mixing require two insertions of the (tiny) flavor-violating coefficient. In the lepton sector, even $|\Delta F| = 1$ decays are virtually unobservable, as the effect is now suppressed by the masses of the relevant leptons. FCNCs mediated by the SM Higgs are similarly negligible, since the corresponding Wilson coefficients feature the same suppression as those from Z exchange, on top the smaller coupling between the Higgs and light fermions. Flavor-changing interactions mediated by the Z' and h' are further suppressed by an additional factor of $m_Z^2/m_{Z'}^2$ and $m_h^2/m_{h'}^2$ respectively, making them irrelevant. Overall, the strong suppression of the tree-level FCNCs that occurs naturally in these models makes their effects negligible.

One-loop FCNCs

Another source of FCNCs beyond those present in the SM arises at one loop. In these models, the familiar box diagram that describes meson mixing in the SM is now accompanied by similar diagrams that include W' s as well as the additional (heavy) up-type quarks running inside the loop, as we show in fig. 2. Given the level of experimental precision in measurements of kaon mixing parameters, even a modification to this process

⁸The effective operator $(\bar{b}_R\gamma^\mu s_R)(\bar{\mu}\gamma_\mu\gamma^5\mu)$ is also generated after integrating out the Z , with a Wilson coefficient $C'_{10}{}^{(\text{BSM})} \sim C_{10}{}^{(\text{BSM})}$ that enters into eq. (3.34) in a similar manner. The presence of this operator does not quantitatively affect our analysis.

at the loop level can be a significant source of constraints.

The relevant interactions are those involving the SM-like down-type quarks and both the W and W' gauge bosons. They are given by

$$\mathcal{L} \supset \frac{g}{\sqrt{2}} W_\mu^+ \sum_{j,i=1}^3 (V_{ij} \bar{u}_{iL} \gamma^\mu d_{jL} + \Delta V_{ij} \bar{U}_{iL} \gamma^\mu d_{jL}) + \text{h.c.}, \quad (3.39)$$

where $U_{3L} = t'_L$ here and $\Delta V = \epsilon_u V$, as well as

$$\begin{aligned} \mathcal{L} \supset \frac{g}{\sqrt{2}} W_\mu'^+ \sum_{j=1}^3 \left(\sum_{i=1}^2 V_{ij} \bar{u}_{iR} \gamma^\mu d_{jR} + V_{3j} \bar{t}'_R \gamma^\mu d_{jR} \right. \\ \left. + \sum_{i=1}^2 \Delta V'_{ij} \bar{U}_{iR} \gamma^\mu d_{jR} + \Delta V'_{3j} \bar{t}'_R \gamma^\mu d_{jR} \right), \end{aligned} \quad (3.40)$$

with $\Delta V' = \epsilon'_u V$. The entries of the 3×3 matrices ϵ_u and ϵ'_u are $\mathcal{O}(v/M)$ and $\mathcal{O}(v'/M)$ respectively, and explicit expressions can be found in eq. (3.101).

The detailed expressions, including loop functions, relevant to estimate the contributions to the kaon mixing parameters Δm_K and $|\epsilon_K|$ can be found in section 3.C. Additional box diagrams including two W s or two W' s always lead to a contribution which is much smaller than that of the SM, and can therefore be neglected. The leading contribution arises from diagrams including one W and one W' . In this case, there is an “irreducible” contribution to both parameters (irreducible in the sense that it can only be “turned-off” by increasing v'), which comes from the u and c quarks, whose couplings to the W' gauge boson are set to be equal to those of the CKM matrix as a result of generalized parity. The size of this correction reads

$$(\Delta m_K)_{u,c} \approx -6 \cdot 10^{-16} \text{ GeV} \left(\frac{6 \text{ TeV}}{m_{W'}} \right)^2, \quad \text{and} \quad |\epsilon_K|_{u,c} \approx 7 \cdot 10^{-5} \left(\frac{6 \text{ TeV}}{m_{W'}} \right)^2, \quad (3.41)$$

which in both cases is an order of magnitude below the theoretical error in the corre-

sponding SM prediction, for values of $m_{W'}$ consistent with the direct bounds discussed in section 3.3.1.

Contributions from box diagrams that involve additional members of the up-quark sector additionally depend on the see-saw scale M , as well as the size of both diagonal and off-diagonal entries in the up-type Yukawa matrices. As far as Δm_K is concerned, the leading contribution comes from diagrams where the u and c quarks propagate inside the loop, and so it is roughly equal to the result in eq. (3.41), even for a see-saw scale M that sits only slightly above v' . In contrast, the contribution to $|\epsilon_K|$ can be large, and it is dominated by diagrams where the t quark propagates inside the loop. Choosing the individual entries in the Yukawa couplings to saturate the upper bound given in eq. (3.100), $|\epsilon_K|$ sets a lower bound on M that can range between 750 TeV and 1000 TeV (depending on whether the leading contribution interferes destructively or constructively with the SM result) for $v' \sim 18$ TeV. This value of M sits comfortably within the upper bound $M \lesssim 10^2 v'$, which follows from the requirement of perturbative Yukawas, as discussed around eq. (3.18). Alternatively, even for $v' = 18$ TeV and $M = 40$ TeV, an additional suppression by a factor of $\mathcal{O}(0.1)$ in the off-diagonal elements of the up-type Yukawas with respect to their upper bound is enough to bring the predicted value of $|\epsilon_K|$ within the allowed range.

Overall, the class of parity solutions to the strong CP problem that we focus on in this work can comfortably satisfy existing constraints from flavor physics. Flavor-changing processes are, nevertheless, an interesting probe of the structure of these models, and a more in-depth investigation is a promising avenue for future work.

3.4 Broken parity and the neutron EDM

As we discussed in section 3.2, parity-symmetric theories predict a vanishing $\bar{\theta}$, therefore offering a potential solution to the strong CP problem. However, the breaking of parity that is necessary for phenomenological reasons implies that, although zero at tree-level, a non-zero $\bar{\theta}$ may be generated radiatively. In this section, we investigate in detail the size of radiative corrections to both $\bar{\theta}$, and the EDM of elementary fermions. We focus on the effect of non-gravitational interactions, and leave gravitational considerations to section 3.5.

The size of radiative corrections to the $\bar{\theta}$ parameter is a somewhat model-dependent question, as it depends on the details of how parity is broken. For instance, we could regard generalized parity to be a global symmetry that is only broken softly by dimensionful parameters, as in eq. (3.13). More realistically, we might expect that the breaking of parity is spontaneous, and not explicit. This must certainly be the case if, for example, parity were a gauge symmetry of the UV theory. Even in this case, there are two qualitatively different options: either parity is broken without breaking CP (e.g. through a symmetry-breaking sector with two scalar fields that obtain asymmetric vevs); or both parity and CP are broken simultaneously (e.g. through the vev of a pseudo-scalar). The former situation is quantitatively similar to the global case. In the latter, however, the symmetry-breaking sector can introduce an additional source of CP -violation beyond that present in the SM, and a non-vanishing $\bar{\theta}$ can arise already at one loop.

In the remainder of this section, we discuss the three qualitatively different possibilities for the breaking of parity, with a focus on the implications for the size of radiative corrections to the neutron EDM.

3.4.1 Softly broken parity

We will first discuss the possibility of parity being broken softly, only as a result of dimensionful parameters. Performing this analysis will give us an understanding of the irreducible effects that will be present in any theory where the breaking of parity happens dynamically.

There are two potential sources of soft breaking. One corresponds to the μ^2 term in the scalar potential of eq. (3.13), which splits the Higgs vevs in the SM and mirror sectors. As anticipated in the introduction, if this was the only source of parity-violation, radiative corrections to $\bar{\theta}$ would be no larger than in the SM [79]. Another potential source of soft breaking are the vector-like mass matrices of eq. (3.10). Relaxing the requirement that these be hermitian introduces a soft breaking of both generalized parity and CP . In this case, a correction to the EDMs of elementary charged fermions (both quarks and leptons) arises already at one loop, whereas $\bar{\theta}$ remains zero both at the tree- and one-loop levels. In turn, this translates into a contribution to the neutron EDM *independent* of $\bar{\theta}$.

Taking the vector-like mass matrices of the $SU(2)$ -singlets to be general complex matrices, WLOG we may write them as

$$\mathcal{M}'_f = \mathcal{M}_f + i\Delta\mathcal{M}_f, \quad (3.42)$$

where both $\mathcal{M}'_f = \mathcal{M}_f$ and $\Delta\mathcal{M}'_f = \Delta\mathcal{M}_f$. If $\Delta\mathcal{M}_f$ is non-vanishing, \mathcal{M}'_f is no longer hermitian, therefore (softly) breaking both generalized parity and CP . At one-loop, a non-zero $\Delta\mathcal{M}_f$ leads to a non-vanishing contribution to the EDM of elementary fermions, with the relevant diagrams depicted in fig. 3. The result is dominated by diagrams where the mirror Higgs, h' , and the heavy mirror fermions propagate inside the loop. We present




Figure 3: One-loop diagrams generating a non-zero EDM for the d quark in the presence of soft breaking of generalized parity through non-hermitian vector-like mass matrices for the $SU(2)$ -singlets. The leading contribution arises from the diagram where h' , and the heavy mirror quarks, D_j , propagate inside the loop. Analogous diagrams are present in both the up-quark and lepton sectors.

a detailed calculation in section 3.B.1. For any of the light SM fermions, we find

$$\frac{d_f}{e} \simeq \frac{n_f Q_f m_f}{32\pi^2 M^2} \times \mathcal{O}\left(\frac{|\Delta\mathcal{M}|}{M}\right). \quad (3.43)$$

where n_f is the number of mirror fermions appearing at the see-saw scale in each fermion sector (i.e. $n_d = n_e = 3$, and $n_u = 2$), and $|\Delta\mathcal{M}|$ refers to the typical size of the entries in the $\Delta\mathcal{M}$ matrix.

Taking the soft-breaking through $\Delta\mathcal{M}$ to be $\mathcal{O}(1)$, we find, parametrically,

$$d_u, d_d \sim 10^{-28} \left(\frac{40 \text{ TeV}}{M}\right)^2 e \cdot \text{cm}. \quad (3.44)$$

In turn, this will translate into an EDM for the neutron of approximately the same size. For illustration, we have normalized the above expression to a value of M that is roughly a factor of two larger than the current lower bound on v' . The corresponding result lies two orders of magnitude below the current experimental bound on d_n , and could fall within reach of future experiments depending on the value of the see-saw scale (see e.g. [141] for a survey of prospective molecule-based searches promising orders-of-magnitude improvement in sensitivity to hadronic CPV).

The see-saw mechanism must also be implemented in the charged lepton sector. If it

were not, the mirror partner of the SM electron would appear at a scale $m_{e'} \simeq m_e \times v'/v$, which would be as low as ~ 40 MeV for the least fine-tuned version of the model where $v' \simeq 18$ TeV. Since mirror fermions carry the same electromagnetic charge as their SM counterparts, this possibility is obviously ruled out. As a result, a non-zero electron EDM is also a generic prediction of this class of theories. Parametrically

$$d_e \sim 10^{-29} \left(\frac{90 \text{ TeV}}{M} \right)^2 e \cdot \text{cm}, \quad (3.45)$$

where we have chosen the see-saw scale in the lepton sector so as to saturate the current upper bound on the electron EDM, which is $|d_e| < 1.1 \cdot 10^{-29} e \cdot \text{cm}$ [142].

Although an EDM is generated at one loop for the various elementary fermions, $\bar{\theta}$ remains zero at this order. At tree-level, it is easy to see that $\bar{\theta} = 0$, even in the presence of non-hermitian vector-like mass matrices. Working in the flavor basis, the full 6×6 mass matrices in both the up- and down-quark sectors are only modified with respect to eq. (3.12) by replacing \mathcal{M}_f with \mathcal{M}'_f in the bottom-right block. We then have

$$\det \mathbb{M}_f = \det \begin{pmatrix} \mathbb{0}_3 & \frac{v'}{\sqrt{2}} y_f'^* \\ \frac{v}{\sqrt{2}} y_f^T & \mathcal{M}'_f \end{pmatrix} \propto \det (y_f'^* y_f^T) \quad \text{for } f = u, d, \quad (3.46)$$

which is real regardless of \mathcal{M}'_f , since $y'_f = y_f$. This is clearly an accidental consequence of the zero appearing in the upper-left corner of the quark mass matrix — the gauge structure of the theory does not allow for relevant operators with the appropriate quantum numbers to fill that block. The vanishing of $\bar{\theta}$ at one-loop is less immediately obvious. The relevant calculation was performed in [78], and it is also apparent as a byproduct of our EDM calculation in section 3.B.1. As already emphasized in [78], a non-zero correction to $\bar{\theta}$ could appear at the two-loop order, and would lead to an additional contribution to the neutron EDM that could be comparable in size to the one discussed here.

3.4.2 Spontaneously broken parity and CP

Perhaps more compellingly — and necessarily, if parity is a gauge symmetry — the breaking of parity can be accomplished through the vacuum expectation value of an additional field. The most minimal realization actually entails the breaking of both parity and CP through the vev of a pseudo-scalar field ϕ . The soft term in eq. (3.13) is generated by pseudo-scalar couplings to the Higgs sector, of the form

$$V \supset \mu_\phi \phi (|H|^2 - |H'|^2) + \lambda_\phi \phi^2 (|H|^2 + |H'|^2). \quad (3.47)$$

The first term above splits the two vevs, and to obtain $v' \gg v$ entails

$$\mu_\phi v_\phi \sim \kappa v'^2. \quad (3.48)$$

A natural possibility is to take $v_\phi \sim \mu_\phi \sim v'$, with $\kappa = \mathcal{O}(1)$. However, $\kappa \ll 1$ is also possible, especially since this coupling breaks the otherwise accidental $SU(4)$ global symmetry of the scalar potential in eq. (3.13). Indeed, the quartic coupling of the SM-like Higgs is $\lambda_{SM} \sim 2\kappa$, suggesting $\kappa \lesssim 0.1$ and thus a pseudo-scalar vev v_ϕ that is numerically somewhat smaller than v' .

Crucially, there is an additional operator consistent with all symmetries that involves ϕ and the $SU(2)$ -singlet fermions, of the form [130]:

$$\mathcal{L} \supset i(\bar{y}_d)_{ij} \phi D_i D'_j + \text{h.c.}, \quad (3.49)$$

and similarly for up-type quarks and leptons. The \bar{y}_f matrices must be hermitian in order to respect generalized parity. When ϕ gets a vev, this term breaks both parity *and* CP . In the notation of section 3.4.1, a non-hermitian contribution to the vector-like

masses in the fermion sector is generated, of the form $\Delta\mathcal{M}_f = \bar{y}_f v_\phi$. More importantly, new interactions involving the pseudo-scalar lead to a non-zero contribution to $\bar{\theta}$ already at one loop, which sets stringent constraints on the size of these couplings. The relevant diagrams are those on the left of fig. 3, minus the external photon line, and allowing for ϕ to propagate inside the loop. In section 3.B.2, we present a detailed calculation, performed in the mass eigenbasis, of the one-loop correction to the quark mass matrix, and the corresponding correction to $\bar{\theta}$, in the context of left-right models with a see-saw fermion structure. In the remainder of this section, we will reproduce the parametric contribution to $\bar{\theta}$ from the down-quark sector using a spurion analysis that the reader might find more instructive.

If we parametrize the one-loop correction to the 6×6 mass matrix in the down-quark sector in terms of 3×3 blocks, as follows

$$\Delta\mathbb{M}_d \equiv \begin{pmatrix} \Delta_{d'd} & \Delta_{d'D'} \\ \Delta_{Dd} & \Delta_{DD'} \end{pmatrix}, \quad (3.50)$$

then the corresponding contribution to $\bar{\theta}$ from the down-quark sector can be written as

$$\begin{aligned} \theta_d &\equiv \arg \det (\mathbb{M}_d + \Delta\mathbb{M}_d) \\ &\simeq \text{Im tr} (\mathbb{M}_d^{-1} \Delta\mathbb{M}_d) \\ &= \text{Im tr} \left\{ - \left(\frac{vv'}{2} y_d'^* \mathcal{M}_d^{-1} y_d^T \right)^{-1} \Delta_{d'd} + \left(\frac{v}{\sqrt{2}} y_d^T \right)^{-1} \Delta_{Dd} + \left(\frac{v'}{\sqrt{2}} y_d'^* \right)^{-1} \Delta_{d'D'} \right\}. \end{aligned} \quad (3.51)$$

Notice that the $\Delta_{DD'}$ block does not contribute to θ_d at this order, which again is a consequence of the zero in the upper-left corner of \mathbb{M}_d .

We will now estimate the size of the Δ matrices appearing in eq. (3.51) through a

spurion analysis, as follows. The Lagrangian will remain invariant under $SU(3)$ flavor transformations on the various quark fields, of the form

$$d \rightarrow \mathcal{R}_Q d, \quad d' \rightarrow \mathcal{R}_{Q'} d', \quad \text{and} \quad D \rightarrow \mathcal{R}_D D, \quad D' \rightarrow \mathcal{R}_{D'} D', \quad (3.52)$$

provided the various Yukawa couplings, as well as the vector-like mass matrix, similarly transform in an appropriate manner. The correct transformation rules for these objects are given by

$$\begin{aligned} y_d^T &\rightarrow \mathcal{R}_D^* y_d^T \mathcal{R}_Q^\dagger, & y_d'^* &\rightarrow \mathcal{R}_{Q'}^* y_d'^* \mathcal{R}_{D'}^\dagger, \\ \mathcal{M}_d &\rightarrow \mathcal{R}_D^* \mathcal{M}_d \mathcal{R}_{D'}^\dagger, & \bar{y}_d &\rightarrow \mathcal{R}_D^* \bar{y}_d \mathcal{R}_{D'}^\dagger. \end{aligned} \quad (3.53)$$

On the other hand, the Δ matrices of eq. (3.51) must similarly transform as follows:

$$\Delta_{d'd} \rightarrow \mathcal{R}_{Q'}^* \Delta_{d'd} \mathcal{R}_Q^\dagger, \quad \Delta_{d'D'} \rightarrow \mathcal{R}_{Q'}^* \Delta_{d'D'} \mathcal{R}_{D'}^\dagger, \quad \text{and} \quad \Delta_{Dd} \rightarrow \mathcal{R}_D^* \Delta_{Dd} \mathcal{R}_Q^\dagger. \quad (3.54)$$

It is now straightforward to identify the leading objects that transform as in eq. (3.54) and contain a single insertion of \bar{y}_d . These are of the form

$$\Delta_{d'd} \sim \frac{v v' v_\phi}{16\pi^2} (y_d'^* \mathcal{M}_d^{-1} \bar{y}_d \mathcal{M}_d^{-1} y_d^T), \quad (3.55)$$

whereas

$$\Delta_{d'D'} \sim \frac{v' v_\phi}{16\pi^2} (y_d'^* \mathcal{M}_d^{-1} \bar{y}_d), \quad \text{and} \quad \Delta_{Dd} \sim \frac{v v_\phi}{16\pi^2} (\bar{y}_d \mathcal{M}_d^{-1} y_d^T). \quad (3.56)$$

We have also included numerical factors to account for the loop suppression, as well as to take into account that the contribution to $\bar{\theta}$ must not diverge in the limits where either v or v' vanish. Plugging this back into eq. (3.51), we find that all three terms give a

contribution of the same size. Parametrically:

$$\theta_d \sim \frac{v_\phi}{16\pi^2} \text{tr} (\mathcal{M}_d^{-1} \bar{y}_d) \sim \frac{|\bar{y}_d| v_\phi}{16\pi^2 M}. \quad (3.57)$$

This result is consistent with the more detailed calculation of the contribution to $\bar{\theta}$ from the quark sector presented in section 3.B.2.

Requiring that $\bar{\theta} \lesssim 10^{-10}$ sets an upper bound on the typical size of the entries of the \bar{y} matrices in the quark sector, of the form

$$\bar{y} \lesssim 10^{-8} \frac{M}{v_\phi} \lesssim 10^{-6}, \quad (3.58)$$

where in the last step we have assumed that $v_\phi \sim v'$, and have taken into account the upper bound on the see-saw scale M as given in eq. (3.18).

This result bring us to the following conclusion: if the spontaneous breaking of parity also implies breaking CP , then any interaction between the quark and symmetry breaking sectors must be extremely weak. Fortunately, if $\bar{y} = 0$ at tree-level, a non-zero value of \bar{y} will not be generated radiatively. Indeed, \bar{y} and the vector-like mass matrices \mathcal{M} are the only two parameters that violate the \mathbb{Z}_2 symmetry acting on the matter fields of the mirror sector. The breaking through \mathcal{M} , however, is soft, and therefore will not translate into a non-zero \bar{y} at the loop order. In this sense, a vanishing \bar{y} is technically natural.

3.4.3 Spontaneously broken parity alone

A less minimal possibility is to spontaneously break parity while preserving CP through the addition of two scalar fields, σ and σ' , whose vevs differ. This can be

achieved if this symmetry breaking sector has a scalar potential of the form

$$V_\sigma = -\frac{m_\sigma^2}{2}(\sigma^2 + \sigma'^2) + \frac{\lambda_1}{4}(\sigma^2 + \sigma'^2)^2 + \frac{\lambda_2}{4}\sigma^2\sigma'^2. \quad (3.59)$$

where for simplicity we have forbidden cubic terms by imposing an additional \mathbb{Z}_2 symmetry. If $\lambda_2 > 0$, the vacua lie at $\langle\sigma\rangle = 0$, $\langle\sigma'\rangle = \pm\sqrt{m_\sigma^2/\lambda_1}$ and viceversa. This option is not viable for the Higgs potential itself, which requires both v and v' to be nonzero, but is perfectly adequate for an additional scalar sector.

Parity breaking can then be translated into the Higgs sector by writing appropriate couplings of the form

$$V \supset \lambda_\sigma(\sigma^2|H|^2 + \sigma'^2|H'|^2) + \lambda'_\sigma(\sigma'^2|H|^2 + \sigma^2|H'|^2). \quad (3.60)$$

These terms are compatible with the generalized parity introduced in section 3.2, acting additionally as $\sigma \leftrightarrow \sigma'$. Provided $\lambda_\sigma \neq \lambda'_\sigma$, this will generate the soft term in eq. (3.13) proportional to $\lambda_\sigma - \lambda'_\sigma$. For example, in the vacuum with $\langle\sigma'\rangle \neq 0$, $v' \gg v$ corresponds to

$$(\lambda'_\sigma - \lambda_\sigma)\langle\sigma'\rangle^2 \sim \kappa v'^2. \quad (3.61)$$

As this scenario breaks P without breaking CP (and the additional \mathbb{Z}_2 symmetry acting on the σ s forbids marginal couplings between σ, σ' and fermion bilinears), there are no significant additional contributions to the neutron EDM. There is, of course, the possibility of collider signatures coming from the Higgs portal coupling in eq. (3.60), most notably mixing between the Higgs and the scalar that acquires a vev, as well as invisible decays of the Higgs if kinematically allowed. The two scalars acquire masses of order $\sqrt{2\lambda_1}\langle\sigma'\rangle \sim \sqrt{\lambda_1}v'$ and $\sqrt{\lambda_2/2}\langle\sigma'\rangle \sim \sqrt{\lambda_2}v'$, respectively, and it is certainly possible for one to be lighter than half the Higgs mass depending on the values of $\lambda_{1,2}$.

3.5 Strong CP and quantum gravity

As we discussed in the Introduction, the strong CP problem arises out of the difficulty of reconciling the smallness of $\bar{\theta}$ with the $\mathcal{O}(1)$ violation of both parity and CP by the electroweak sector. In turn, all attempts to address this puzzle are themselves based on the introduction of an additional symmetry beyond those of the SM. However, there is strong evidence that within a theory of quantum gravity, global symmetries cannot be exact — they must be either broken, or gauged. The origin of this statement goes back a long way [94, 95, 96, 97, 98, 99, 100, 101, 102], and to some extent it has recently been established [103, 104]. Of course, the single most pressing issue for phenomenology is to establish a lower bound on the amount of global symmetry violation that must be present in the IR. Attempts at finding such a “universal” lower bound have been made [105, 106], but a fully satisfactory answer remains elusive. Absent a full understanding of how quantum gravity affects global symmetries at low energies, we can at least attempt to assess the robustness of an EFT against global symmetry violation by considering the impact of HDOs suppressed by the appropriate power of M_{Pl} . This both constrains the viable parameter space of parity solutions to strong CP and illustrates the sense in which P , rather than $U(1)_{PQ}$, provides a solution to the strong CP problem that is robust against the expected intrusion of quantum gravity. Beyond imposing constraints, these higher-dimensional operators also lead to new experimental signatures associated with the spontaneous breaking of parity, which we explore in section 3.5.2.

3.5.1 Constraints from Planck-suppressed operators

The observation that the breaking of global symmetries by quantum gravity can have a profound impact on the validity of the QCD axion solution to strong CP was first made in [114, 115, 116, 117]. Planck-suppressed HDOs that violate the $U(1)_{PQ}$ symmetry

carried by the field Φ , the phase of which is the axion, are of the form

$$\mathcal{L} \supset \frac{\eta}{M_{Pl}^{d-4}} |\Phi|^{d-n} \Phi^n + \text{h.c.} \quad (3.62)$$

Here, d is the operator dimension, n its units of $U(1)_{PQ}$ charge (so $n \geq 1$ in order to break the symmetry), and η a coupling that will in general feature arbitrary real and imaginary parts. HDOs of this form contribute to the axion potential, and, in general, will displace the axion vev away from the value leading to a small $\bar{\theta}$. Following [115], requiring that the shift in the axion vev is small enough so as not to spoil the solution to strong CP translates into the following upper bound

$$|\eta| \left(\frac{f_a}{\sqrt{2} M_{Pl}} \right)^d \lesssim 10^{-81} \bar{\theta} \lesssim 10^{-91}, \quad (3.63)$$

where f_a is the scale of $U(1)_{PQ}$ spontaneous symmetry breaking (alternatively, the axion decay constant), which is experimentally constrained to be between 10^8 and 10^{17} GeV [140]. Focusing on operators of dimension $d = 5$, this translates into an upper bound on the size of η , of the form

$$|\eta| \lesssim 10^{-55} \left(\frac{10^{12} \text{GeV}}{f_a} \right)^5 \left(\frac{\bar{\theta}}{10^{-10}} \right). \quad (3.64)$$

In other words, for all experimentally allowed values of the axion decay constant, the $U(1)_{PQ}$ symmetry must remain an approximate global symmetry to an exceptional degree. This is clearly one of the most significant drawbacks of the axion solution to strong CP.

In the remainder of this section we study the effect of Planck-suppressed HDOs on parity solutions to the strong CP problem. We consider separately the cases where parity is global or gauged. The nature of the HDOs under consideration will be different, but

in both cases we will see that even $\mathcal{O}(1)$ coefficients are compatible with solving strong CP.

Parity as a global symmetry

If we regard parity as a global symmetry, then we must consider the effect of HDOs that explicitly violate P . The relevant dimension-5 HDOs were already identified in [143], and they are of the form

$$\mathcal{L} \supset \frac{1}{M_{Pl}} [(\alpha_u)_{ij}(H'Q'_i)(HQ_j) + (\alpha_d)_{ij}(H'^{\dagger}Q'_i)(H^{\dagger}Q_j)] + \text{h.c.} \quad (3.65)$$

Notice that if $\alpha_f = \alpha_f^{\dagger}$ then the above terms would be parity-symmetric. In general, however, the α_f 's will *not* be hermitian, and it is under this assumption that we proceed.

Setting the Higgs to their vevs, eq. (3.65) leads to a correction to the quark mass matrix that, for arbitrary α_f 's, does not respect generalized parity. The leading contribution to $\bar{\theta}$ will come from the contributions to the up- and down-quark masses, which are of the form

$$\delta m_u \simeq \frac{vv'(\alpha_u)_{11}}{2M_{Pl}}, \quad \text{and} \quad \delta m_d \simeq \frac{vv'(\alpha_d)_{11}}{2M_{Pl}}. \quad (3.66)$$

In turn,

$$\theta_q \simeq \frac{\text{Im}(\delta m_u)}{m_u} + \frac{\text{Im}(\delta m_d)}{m_d} \sim 10^5 \frac{|\alpha|v'}{2M_{Pl}}, \quad (3.67)$$

where in the last step we have used $m_u/v \sim m_d/v \sim 10^{-5}$. Requiring that the above contribution is smaller than the current bound on $\bar{\theta}$ translates into an upper bound on the parity breaking scale:

$$v' \lesssim \frac{20 \text{ TeV}}{|\alpha|} \left(\frac{\bar{\theta}}{10^{-10}} \right). \quad (3.68)$$

Notice this upper bound is (just) compatible with the lower bound $v' \gtrsim 18 \text{ TeV}$ from direct searches of W' gauge bosons, as discussed in section 3.3.1. As a result, if *global*

generalized parity is responsible for solving strong CP, an $\mathcal{O}(1)$ violation of the symmetry due to gravitational effects would imply a contribution to $\bar{\theta}$ accessible in near-future experiments.

Parity as a gauge symmetry

If parity is instead a gauge symmetry of the underlying theory, HDOs that explicitly violate P are therefore not allowed. Planck-suppressed operators such as those in eq. (3.65) might still be generated, but only with $\alpha_f = \alpha_f^\dagger$, and therefore will not contribute to $\bar{\theta}$. Instead, the operators of interest must be proportional to the source of spontaneous symmetry breaking. If the latter takes place via the vev of a pseudo-scalar, as discussed in section 3.4.2, then there are two dimension-5 HDOs that satisfy this requirement, namely:

$$\mathcal{L} \supset \eta_s \frac{\phi \alpha_s}{4\pi M_{Pl}} \text{tr} \left(G^a \tilde{G}^a \right), \quad (3.69)$$

and

$$\mathcal{L} \supset \frac{i\phi}{M_{Pl}} \left\{ (\zeta_u)_{ij} Q_i H U_j + (\zeta'_u)_{ij} Q'_i H' U'_j + (\zeta_d)_{ij} H^\dagger Q_i D_j + (\zeta'_d)_{ij} H'^\dagger Q'_i U'_j \right\} + \text{h.c.}, \quad (3.70)$$

with $\eta_s \in \mathbb{R}$, and $\zeta'_f = \zeta_f^*$ so as to satisfy generalized parity.

The operator of eq. (3.69) will generate a contribution to θ_s after spontaneous symmetry breaking, of the form

$$\theta_s \simeq \frac{\eta_s v_\phi}{M_{Pl}}. \quad (3.71)$$

Assuming that $v_\phi \sim v'$, demanding that this contribution is smaller than the current bound on $\bar{\theta}$ leads to an upper bound on the parity breaking scale

$$v_\phi \sim v' \lesssim \frac{10^9 \text{ GeV}}{\eta_s} \left(\frac{\bar{\theta}}{10^{-10}} \right), \quad (3.72)$$

which is clearly well above current bounds on v' .

At the same time, once ϕ gets its vev, the operator of eq. (3.70) leads to an extra contribution to the Yukawa couplings of the up- and down-type quarks in the SM and mirror sectors that are *not* parity-symmetric. In turn, this will lead to an additional contribution to the mass eigenvalues of the light quarks which will in general contain an imaginary component. For example, in the down quark sector

$$\text{Im}(\delta m_{d_i}) \simeq vv' \text{Im} \left\{ \frac{iv_\phi}{M_{Pl}} \sum_j \frac{(y_d)_{ij}^* (\zeta_d)_{ij}}{m_{D_j}} \right\} \sim \frac{vv'}{M} \frac{|\zeta_d| v_\phi}{M_{Pl}} |(y_d)_{i\star}|, \quad (3.73)$$

where $|(y_d)_{i\star}|$ refers to the typical size of the entries in the i -th row of the y_d matrix. The leading contribution to θ_q will come from the up and down quarks. In total:

$$\theta_q \simeq \frac{\text{Im}(\delta m_u)}{m_u} + \frac{\text{Im}(\delta m_d)}{m_d} \simeq \frac{v_\phi}{M_{Pl}} \frac{vv'}{M} \left\{ \frac{|\zeta_u| (y_u)_{1\star}}{m_u} + \frac{|\zeta_d| (y_d)_{1\star}}{m_d} \right\}. \quad (3.74)$$

Taking into account the upper bound on the entries of the Yukawa couplings necessary to reproduce the light quark masses (see eq. (3.94)), as well as the requirement that $v' \lesssim M$ in order to implement the see-saw mechanism, the previous equation implies

$$\theta_q \lesssim 10^2 \frac{|\zeta| v_\phi}{M_{Pl}}, \quad (3.75)$$

where we have set $m_u/v \sim m_d/v \sim 10^{-5}$, and have assumed that $|\zeta_u| \sim |\zeta_d|$. In turn, taking $v_\phi \sim v'$, this sets an upper bound on the scale of spontaneous symmetry breaking:

$$v_\phi \sim v' \lesssim \frac{10^7 \text{ GeV}}{|\zeta|} \left(\frac{\bar{\theta}}{10^{-10}} \right). \quad (3.76)$$

As before, this is fully compatible with current experimental bounds on the parity-breaking scale, even for $\mathcal{O}(1)$ coefficients of the corresponding HDOs.

3.5.2 Gravitational waves from the spontaneous breaking of parity

Beyond providing additional constraints on the parameter space of parity solutions to strong CP, the expected effects of gravity also introduce new experimental signatures. Here we highlight one possibility, namely the impact of HDOs when parity is a spontaneously broken global symmetry. The spontaneous breaking of discrete symmetries can lead to the formation of a network of domain walls in the early universe, provided the reheating temperature after inflation is above the scale of spontaneous symmetry breaking [144]⁹. If the spontaneously broken symmetry is global, but otherwise exact, a domain wall configuration interpolates between two distinct vacua that are degenerate, making these defects topologically stable objects. The formation of such networks can be fatal on two grounds. On the one hand, the energy density in domain walls redshifts slower than that of matter or radiation, and would eventually dominate the universe’s energy budget. If this happened before the current epoch, the rapid expansion of the subsequent domain-wall-dominated era would be at odds with observation. On the other hand, even if only a subdominant component of the total energy density was in the form of domain walls today, their effect on large-scale density fluctuations rules out defects with characteristic scales above ~ 1 MeV [147]. These considerations are often referred to as the “domain wall problem” of theories with spontaneously broken discrete symmetries.

These problems are largely solved when we take into consideration that, within a theory of quantum gravity, we expect all symmetries to be either broken or gauged [94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104] — an expectation that includes spacetime symmetries [103, 104]. In this context, the domain wall network is unstable, rendering

⁹This statement relies on the restoration of the spontaneously broken symmetry at high temperatures. Scenarios where symmetry restoration does not take place have been explored in [145, 146]. In these cases, topological defects would not form via the mechanism of [144].

its earlier formation largely unproblematic (see e.g. [148], and also [145, 146]). Moreover, the significant amount of gravitational radiation emitted in the process results in a stochastic gravitational wave background that may be within reach of current and future observatories. We discuss this possibility in the remainder of this section.

We will focus first on the scenario where parity is a global symmetry that is only explicitly broken by gravitational effects. At low energies, the symmetry-breaking dynamics will enter the effective potential for ϕ through HDOs that violate parity. One such operator is of the form

$$V \supset \epsilon \frac{\phi^5}{M_{Pl}} . \quad (3.77)$$

This breaks the degeneracy between the two previously degenerate vacua, corresponding to $\langle \phi \rangle = \pm v_\phi$. Parametrically, the energy difference now reads

$$\delta V \sim \frac{\epsilon v_\phi^5}{M_{Pl}} . \quad (3.78)$$

If the reheating temperature is above the scale of spontaneous symmetry breaking, then we expect that a network of domain walls will be formed once the temperature of the universe drops below $T \sim v_\phi$ [144]. Numerical [149, 150, 151, 152, 153] and analytical [154, 155] studies suggest that, shortly after formation, the network evolves according to a scaling solution, with $\rho_{DW}(t) \simeq \sigma/t$, and typical domain wall size comparable to the Hubble scale $H(t)^{-1}$. σ corresponds to the tension of the domain walls, which in our model is of the form $\sigma \sim \sqrt{\kappa_\phi} v_\phi^3$, where κ_ϕ refers to the quartic coupling in the ϕ potential. Two competing effects determine the network's subsequent evolution. On the one hand, the pressure difference between the two vacua exerts a force per unit area of order $\sim \delta V$. On the other, the tension per unit area acting on a wall with curvature radius R is $\sim \sigma/R$. In the scaling regime, $R(t) \sim H(t)^{-1} \sim t$ (assuming the universe is

radiation dominated), and therefore the effect of tension decreases with time. Eventually, the pressure difference between the two vacua dominates, causing the network to collapse at a time of order

$$t_{\text{coll.}} \sim \frac{\sigma}{\delta V} \sim \frac{\sqrt{\kappa_\phi} M_{Pl}}{\epsilon v_\phi^2}. \quad (3.79)$$

Clearly, the domain wall network could be very long-lived if $\epsilon \lll 1$. The requirement that collapse takes place either before the universe becomes domain wall dominated, or before the start of BBN (so as to avoid energy injection into the SM plasma that would disrupt light element formation), sets a lower bound on ϵ as a function of the spontaneous symmetry breaking scale. This is depicted in fig. 4, where the BBN and domain-wall-domination restrictions dominate for values of v_ϕ below and above ~ 73 TeV respectively. As can be appreciated in the figure, in the region of parameter space where the fine-tuning is better than 10^{-10} (that is, $v_\phi \sim v' \lesssim 3.5 \cdot 10^4$ TeV), ϵ may be as small as $\mathcal{O}(10^{-13})$.

The collapse of a domain wall network leads to the production of gravitational waves [156, 157]. On dimensional grounds, one would expect the energy density in gravitational radiation to be of the form $\rho_{\text{gw}} \sim G_N \sigma^2$ (the mandatory power of G_N times the necessary factors of σ to make up dimensions), an expectation that is largely upheld by numerical analysis [158, 159, 160, 161]. The resulting gravitational wave spectrum has an extended shape, peaking at a frequency corresponding to the Hubble size at the time of collapse (corresponding to the typical size of the domain walls), and falling off as $1/f$ for larger frequencies. At the present epoch, the peak frequency of the gravitational wave signal is given by

$$\begin{aligned} f_* &\simeq 10^{-9} \text{ Hz} \left(\frac{T_{\text{coll.}}}{10^{-2} \text{ GeV}} \right) \left(\frac{g_*(T_{\text{coll.}})}{10} \right)^{1/6} \\ &\sim 10^{-9} \text{ Hz} \left(\frac{v_\phi}{18 \text{ TeV}} \right) \left(\frac{\epsilon}{10^{-12}} \right)^{1/2} \left(\frac{1}{\kappa_\phi} \right)^{1/4}, \end{aligned} \quad (3.80)$$




Figure 4: Constraints on the size of the coefficient of the Planck-suppressed HDO of eq. (3.77), as a function of the scale of spontaneous symmetry breaking, v_ϕ . The region $v_\phi \lesssim 18$ TeV (green) is in conflict with direct bounds on the mass of W' and Z' resonances, under the assumption that $v_\phi \sim v'$, as discussed in section 3.3.1. Values of ϵ that are too small (blue) do not destabilize the domain wall network early enough to either avoid a domain-wall-dominated era, or to ensure collapse before the onset of BBN, and are therefore ruled out. The region of parameter space in pink is experimentally allowed, but the level of fine-tuning in the electroweak sector worsens as v_ϕ is increased (corresponding to a darker shade). The dashed line at $v_\phi \simeq 3.5 \cdot 10^4$ TeV corresponds to a fine-tuning of $\mathcal{O}(10^{-10})$ in the electroweak sector. (For illustration, we have set the quartic coupling in the pseudo-scalar potential to be $\kappa_\phi = 1$ in this plot.)

and the energy density in gravitational radiation at frequency peak reads [162]

$$\begin{aligned}\Omega_{\text{gw}}h^2(f_*) &\simeq 2 \cdot 10^{-10} \left(\frac{\sigma}{(20 \text{ TeV})^3} \right)^2 \left(\frac{10^{-2} \text{ GeV}}{T_{\text{coll.}}} \right)^4 \left(\frac{10}{g_*(T_{\text{coll.}})} \right)^{4/3} \\ &\sim 10^{-10} \left(\frac{v_\phi}{18 \text{ TeV}} \right)^2 \left(\frac{10^{-12}}{\epsilon} \right)^2 \left(\frac{\kappa_\phi}{1} \right)^2,\end{aligned}\tag{3.81}$$

where $T_{\text{coll.}}$ refers to the temperature of the SM plasma at a time $t_{\text{coll.}}$, and we have assumed that network collapse takes place during radiation domination.¹⁰

Figure 5 shows the region that can be spanned by the peak of the stochastic gravitational wave background in the f_* vs. $\Omega_{\text{gw}}h^2(f_*)$ plane, together with the sensitivity curves of a number of gravitational wave experiments. The lower bound on ϵ depicted in fig. 4 translates into a lower bound on f_* for each value of the symmetry breaking scale (e.g. $f_* \gtrsim 10^{-9}$ Hz for $v_\phi \simeq 18$ TeV). As can be seen in fig. 5, a region of parameter space with low v' falls within reach of gravitational wave observatories probing the low frequency regime such as SKA [163], NANOGrav [164], and the EPTA [41]. As ϵ is increased, the collapse of the domain wall network occurs earlier, further suppressing the current value of the energy density in gravitational radiation by the corresponding redshift factor.

Our discussion so far applies in the context of global discrete symmetries provided that they either do not descend from a continuous symmetry, or that, if they do, the symmetry breaking scale of the continuous factor is above the reheating temperature, so that a network of cosmic strings is not formed in the early universe. On the other hand, if the reheating temperature is larger than the scale set by the tension of the strings, μ , then a string network will be formed first, with the strings later joined by domain walls. The entire string-wall network now evolves together, and the problem features an

¹⁰In the second steps of eqs. (3.80) and (3.81), we have substituted $T_{\text{coll.}}$ by the corresponding expression in terms of the model's fundamental parameters, while ignoring a weak dependence on g_* .




Figure 5: In blue, region of parameter space where the frequency peak of the stochastic gravitational wave signal, f_* , as well as the corresponding energy density, $\Omega_{\text{gw}} h^2(f_*)$, can fall in light of the experimental constraints on the various model parameters summarized in fig. 4. Dotted lines correspond to constant ϵ . The region to the right of the dashed line corresponding to $v_\phi = 3.5 \cdot 10^4$ TeV features a level of fine-tuning worse than 1 part in 10^{10} , and it is therefore less attractive. Sensitivity curves for a variety of gravitational wave experiments are shown, including the pulsar timing arrays EPTA [41], NANOGrav [164], SKA [163] (observation time of 5, 10 and 20 years as indicated), as well as the space-based interferometers LISA [165], and DECIGO [166]. (For illustration, we have set the quartic coupling in the pseudo-scalar potential to be $\kappa_\phi = 1$ in this plot.)

additional time scale, given by

$$t_* \sim \frac{\mu}{\sigma}. \quad (3.82)$$

At $t \sim t_*$, the force per unit length on a string of radius $R(t) \sim t$, given by $\sim \mu/t$, becomes comparable to the wall tension. The system then becomes dominated by the tension of the domain walls, causing the network to shrink, and break down into pieces that will further decay into gravitational waves (or, potentially, also massive particles, depending on their size and the relevant particle spectrum) [156, 157]. If this timescale is shorter than $t_{\text{coll.}}$, the earlier destruction of the network of defects could move any potential gravitational wave signal into an unobservable regime.

The discussion of the previous paragraph is especially relevant if parity is instead realized as a gauge symmetry, for which explicit breaking is no longer allowed. Naïvely, one would hope that the gauge case would be cosmologically more benign: the gauge equivalence of the two vacua makes them no longer distinct, eliminating the topological stability of the domain walls. Indeed, domain walls can be destroyed by a process in which a string loop is nucleated on the wall, further growing to destroy the entire defect. However, the corresponding nucleation probability is proportional to $e^{-\mu^3/\sigma^2}$ [167], which will be exceedingly small for any reasonable separation of scales between the string and wall tensions, therefore rendering gauge domain walls effectively stable. It is therefore crucial that the reheating temperature is above the string tension scale, so that a string network is formed that can later result in the entire collapse of the subsequent string-wall network. The cosmological implications, as well as potential gravitational wave signatures, of a discrete parity symmetry that is gauged will be further explored in future work.

3.6 Conclusions

The strong CP problem remains one of the great naturalness problems of the Standard Model, and is perhaps the most compelling in light of its resistance to straightforward anthropic explanations. Fully satisfying solutions to the problem remain elusive given the expected violation of global symmetries in a theory of quantum gravity, which demands extensive effort to protect the Peccei-Quinn symmetry underlying axion-based approaches. In this work we have pursued a possibility that is more transparently robust against the effects of quantum gravity, revisiting parity-based solutions to the strong CP problem. Our approach highlights the experimental signatures associated with the most natural regions of parameter space in these models, as well as ancillary signatures that are dependent upon the detailed mechanism of parity breaking.

The notion of naturalness within this parameter space is governed by the tuning associated with the separation of scales of $SU(2)_L$ and $SU(2)_R$ breaking, which are related by generalized parity. Given this tuning, “see-saw” vector-like masses for the $SU(2)$ -singlet fermions play a key role in allowing the scale of $SU(2)_R$ breaking to be lowered toward its most natural value consistent with experimental constraints. Within this framework, the LHC provides the strongest test of natural parity-based solutions to the strong CP problem, probing the scale of $SU(2)_R$ breaking through searches for W' and Z' vector bosons as well as vector-like quarks and additional Higgs bosons. This leaves parity solutions tuned at the $\sim 10^{-3}$ level, which while not fully natural remains a significant improvement in explaining the observed $\bar{\theta} \lesssim 10^{-10}$. The extended reach for heavy resonances at future colliders such as FCC-hh will decisively test these parity solutions at the level of $\sim 10^{-5}$ tuning. Constraints on new sources of flavor violation play a complementary role, with additional sensitivity to the scale of vector-like fermions and the underlying model of flavor.

The detailed mechanism of parity breaking gives rise to additional signatures within reach of near-future tabletop experiments and gravitational wave observatories. Soft parity and CP -violating terms give rise to EDMs for elementary fermions at one-loop, both quarks and charged leptons, which provide a pathway to discovery in precision searches for CP -violation in molecular systems. Spontaneous violation of parity and CP through the vev of a pseudo-scalar gives rise to additional one-loop contributions to $\bar{\theta}$, which provides an additional pathway to discovery and already requires the source of parity violation to be sequestered from the quark sector (albeit in a technically natural way). The expected violation of global symmetries in a theory of quantum gravity further shapes the viable parameter space and potential experimental signatures through the impact of various Planck-suppressed operators whose form depends on the underlying parity-breaking mechanism. If parity is a global symmetry that is broken both spontaneously (by a pseudo-scalar vev) and explicitly (by gravitational effects), collapse of the domain wall network associated with the spontaneous breaking of parity can generate a gravity wave signal accessible at low-frequency gravitational wave observatories. In this respect, the violation of global symmetries by gravitational effects is a feature of parity-based solutions to the strong CP problem, rather than a bug. Taken together, these experimental opportunities warrant further exploration of generalized parity as a solution to strong CP.

Appendices

3.A Mass eigenstates

3.A.1 Gauge and Higgs sectors

With the gauge group of eq. (3.3), and the Higgs sector specified in table 1, spontaneous symmetry breaking takes place in two steps, as follows

$$SU(2)_L \times SU(2)_R \times U(1)_{\hat{Y}} \xrightarrow{v' \neq 0} SU(2)_L \times U(1)_Y \xrightarrow{v \neq 0} U(1)_{EM}. \quad (3.83)$$

The physical spectrum contains SM-like Z , W^\pm , and γ gauge bosons, as well as exotic Z' and W'^\pm excitations. At tree-level, no mixing occurs in the charged gauge boson sector, and the mass eigenstates are given in terms of the gauge eigenbasis by the usual expression:

$$W^\pm = \frac{1}{\sqrt{2}}(W^1 \mp iW^2), \quad (3.84)$$

and similarly in the W' sector. Tree-level masses are of the form $m_W = gv/2$ and $m_{W'} = gv'/2$, where we have assumed that $g' = g$, as mandated by generalized parity. By contrast, in the neutral gauge boson sector mixing between SM and mirror fields takes place already at tree-level. At zeroth order in a v/v' expansion, the gauge eigenstates

can be written in the mass eigenbasis as follows

$$\begin{pmatrix} W_\mu^3 \\ W_\mu^3 \\ \hat{B}_\mu \end{pmatrix} = \begin{pmatrix} \frac{\sqrt{\cos 2\theta_w}}{\cos \theta_w} & -\sin \theta_w \tan \theta_w & \sin \theta_w \\ 0 & \cos \theta_w & \sin \theta_w \\ -\tan \theta_w & -\tan \theta_w \sqrt{\cos 2\theta_w} & \sqrt{\cos 2\theta_w} \end{pmatrix} \begin{pmatrix} Z'_\mu \\ Z_\mu \\ A_\mu \end{pmatrix}, \quad (3.85)$$

where $\sin^2 \theta_w \simeq 0.231$ as usual. Corrections to the above expression arise at $\mathcal{O}(v^2/v'^2)$. Masses for the SM-like Z and mirror Z' are given by

$$m_Z = \frac{gv}{2\cos \theta_w} + \mathcal{O}\left(\frac{v^2}{v'^2}\right), \quad \text{and} \quad m_{Z'} = \frac{gv' \cos \theta_w}{2\sqrt{\cos 2\theta_w}} + \mathcal{O}\left(\frac{v^2}{v'^2}\right). \quad (3.86)$$

After electroweak symmetry breaking, the Higgs sector consists of two real scalar fields, h and h' , with masses given by $m_h \simeq 2\sqrt{\kappa}v$ and $m_{h'} \simeq \sqrt{2\lambda}v'$. Rotating from the gauge to the mass eigenbasis can be performed as follows

$$\begin{pmatrix} h \\ h' \end{pmatrix} \rightarrow \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} h \\ h' \end{pmatrix}, \quad (3.87)$$

with mixing angle $\alpha \sim v/v'$.

3.A.2 Fermion sector

Rotating from the flavor to the mass eigenbasis in the fermion sector requires solving the eigenvalue problem for the 6×6 matrices $\mathbb{M}_f^\dagger \mathbb{M}_f$, and $\mathbb{M}_f \mathbb{M}_f^\dagger$, with \mathbb{M}_f as given in eq. (3.12). This can be conveniently done as a perturbation expansion in $v/M, v'/M \ll 1$. In this section, we summarize the relevant results of this procedure. We focus first on the down-quark and lepton sectors (although we will use notation appropriate to the down-quark sector, we emphasize that identical results apply for leptons). The singularities

of the up sector as related to the top quark merit a separate discussion that we present later.

Down-type quarks and leptons

The mass eigenvalues in the down-quark sector can be found by diagonalizing the two 3×3 matrices

$$\frac{v'v}{2}y_d'^*\mathcal{M}_d^{-1}y_d^T, \quad \text{and} \quad \mathcal{M}_d. \quad (3.88)$$

In full generality, i.e. without yet imposing generalized parity, the above matrices are not necessarily hermitian, and two unitary matrices are needed in order to bring them into real diagonal form. This corresponds to the unitary transformations

$$d \rightarrow \mathcal{O}_d^\dagger d, \quad d' \rightarrow \mathcal{O}_{d'}^\dagger d', \quad \text{and} \quad D \rightarrow \mathcal{O}_D^\dagger D, \quad D' \rightarrow \mathcal{O}_{D'}^\dagger D'. \quad (3.89)$$

By definition, the rotation matrices are such that

$$\mathfrak{m}_d \equiv \mathcal{O}_{d'}^* \left(\frac{v'v}{2}y_d'^*\mathcal{M}_d^{-1}y_d^T \right) \mathcal{O}_d^\dagger = \text{diag}(m_{d_i}), \quad (3.90)$$

and

$$\mathfrak{m}_D \equiv \mathcal{O}_D^* \mathcal{M}_d \mathcal{O}_{D'}^\dagger = \text{diag}(m_{D_i}), \quad (3.91)$$

where m_{d_i} and m_{D_i} are the masses of the SM and exotic heavy quarks respectively. As advertised in section 3.2.3, we will make the simplifying assumption that all three mirror quarks appear at a common scale $m_{D_i} \sim M \gg v, v'$. Imposing generalized parity makes both matrices in eq. (3.88) hermitian. In this case, a single unitary matrix suffices to make them diagonal, and we have $\mathcal{O}_{d'} = \mathcal{O}_d^*$ and $\mathcal{O}_D = \mathcal{O}_{D'}^*$.

It is convenient to define two new matrices corresponding to the Yukawa couplings in

this new basis

$$\tilde{y}_d \equiv \mathcal{O}_d^* y_d \mathcal{O}_D^\dagger, \quad \text{and} \quad \tilde{y}'_d \equiv \mathcal{O}_{d'} y'_d \mathcal{O}_{D'}^T. \quad (3.92)$$

With this definition, the tree-level masses of the SM-like fermions read

$$m_{d_i} = \frac{vv'}{2} \sum_j \frac{(\tilde{y}'_d)^*_{ij} (\tilde{y}_d)_{ij}}{m_{D_j}} = \frac{vv'}{2} \sum_j \frac{|(\tilde{y}_d)_{ij}|^2}{m_{D_j}}, \quad (3.93)$$

where the last step holds provided we impose generalized parity. From this expression, we can find an upper bound on the individual entries in the Yukawa matrix, of the form

$$|(\tilde{y}_d)_{ij}| \lesssim \left(\frac{2m_{d_i} m_{D_j}}{vv'} \right)^{1/2} \sim \left(\frac{m_{d_i} M}{vv'} \right)^{1/2}. \quad (3.94)$$

As advertised in section 3.2.3, bringing the full 6×6 matrix of eq. (3.12) into diagonal form requires a further transformation that mixes the $SU(2)$ -doublet and singlet fields, as specified in eq. (3.20). In terms of the \tilde{y}_d and \tilde{y}'_d couplings defined earlier, the 3×3 blocks appearing in eq. (3.20) can be written as

$$\epsilon_d = \frac{v}{\sqrt{2}} m_D^{-1} \tilde{y}_d^T, \quad \text{and} \quad \epsilon'_d = \frac{v'}{\sqrt{2}} m_D^{-1} \tilde{y}'_d{}^\dagger, \quad (3.95)$$

whose entries are of $\mathcal{O}(v/M)$ and $\mathcal{O}(v'/M)$ respectively. When generalized parity is only broken by the different vev's in the SM and mirror sectors, we have $\epsilon'_d = (v'/v)\epsilon_d^*$.

Up-type quarks

The diagonalization procedure in the up-quark sector is analogous to that for down-type quarks and leptons, although this time accommodating for the singularities of the third generation for which the see-saw mechanism cannot be implemented.

As before, at zeroth order in $v^{(l)}/M$, we perform transformations of the form

$$u \rightarrow \mathcal{O}_u^\dagger u, \quad u' \rightarrow \mathcal{O}_{u'}^\dagger u', \quad \text{and} \quad U \rightarrow \mathcal{O}_U^\dagger U, \quad U' \rightarrow \mathcal{O}_{U'}^\dagger U'. \quad (3.96)$$

On the one hand, the matrices \mathcal{O}_U and $\mathcal{O}_{U'}$ must be chosen such that the vector-like mass matrix \mathcal{M}_u is brought into diagonal form. In this case, we make the assumption that two of the eigenvalues of \mathcal{M}_u are $m_{U_1}, m_{U_2} \sim M$, whereas the third one is much smaller, and for simplicity we will take it to vanish in what follows. On the other hand, the matrices \mathcal{O}_u and $\mathcal{O}_{u'}$ must now be such that

$$(\tilde{y}_u'^* \hat{m}_U^{-1} \tilde{y}_u^T)_{ij} = \delta_{ij} (\tilde{y}_u'^* \hat{m}_U^{-1} \tilde{y}_u^T)_{ii}, \quad \text{and} \quad (\tilde{y}_u)_{i3} = (\tilde{y}_u')_{i3} = 0 \quad (3.97)$$

for $i, j = 1, 2$, and where $\hat{m}_U^{-1} \equiv \text{diag}(m_{U_1}^{-1}, m_{U_2}^{-1}, 0)$, and the \tilde{y}_u and \tilde{y}_u' matrices are defined as in eq. (3.92). Moreover, we define $y_t \equiv \tilde{y}_{33}$, and $y_{t'} \equiv \tilde{y}_{33}'$, which we may choose to be real and positive.

With this preliminaries, the tree-level mass eigenvalues in the top sector read

$$m_t = \frac{y_t}{\sqrt{2}} v, \quad \text{and} \quad m_{t'} = \frac{y_{t'}}{\sqrt{2}} v', \quad (3.98)$$

with $y_{t'} = y_t$ if we impose generalized parity. For the first and second generation, we have instead

$$m_{u_i} = \frac{vv'}{2} \sum_{j=1}^2 \frac{(\tilde{y}_u'^*)_{ij} (\tilde{y}_u)_{ij}}{m_{U_j}} = \frac{vv'}{2} \sum_{j=1}^2 \frac{|(\tilde{y}_u)_{ij}|^2}{m_{U_j}}, \quad (3.99)$$

where the last step holds provided we impose generalized parity. As before, we can now obtain an upper bound on the individual Yukawa entries, of the form

$$|(\tilde{y}_u)_{ij}| \lesssim \left(\frac{2m_{u_i} m_{U_j}}{vv'} \right)^{1/2} \sim \left(\frac{m_{u_i} M}{vv'} \right)^{1/2} \quad \text{for} \quad i, j = 1, 2. \quad (3.100)$$

A further transformation mixing the $SU(2)$ -singlet and doublet components is again necessary in order to diagonalize the full 6×6 mass matrix, which can be written as in eq. (3.20). The corresponding ϵ_u and ϵ'_u blocks are now given by

$$\begin{aligned} (\epsilon_u)_{ij} &= \frac{v}{\sqrt{2}} (\hat{m}_U^{-1} \tilde{y}^T)_{ij} , & (\epsilon_u)_{3j} &= -\frac{vv'}{2} \frac{m_t (\tilde{y}'^* \hat{m}_U^{-1} \tilde{y}^T)_{3j}}{m_t^2 - \delta_{j3} m_t^2} , \\ (\epsilon'_u)_{ij} &= \frac{v'}{\sqrt{2}} (\hat{m}_U^{-1} \tilde{y}'^T)_{ij}^* , & (\epsilon'_u)_{3j} &= -\frac{vv'}{2} \frac{m_t (\tilde{y}^* \hat{m}_U^{-1} \tilde{y}'^T)_{3j}^*}{m_t^2 - \delta_{j3} m_t^2} . \end{aligned} \quad (3.101)$$

for $i = 1, 2$ and $j = 1, 2, 3$. Just as in the down-quark sector, if generalized parity is only broken by the difference between v and v' , we have $\epsilon'_u = (v'/v) \epsilon_u^*$.

3.B Radiatively induced EDM

3.B.1 One-loop EDM

We will now present a calculation of the one-loop correction to the EDM of elementary charged fermions that arises under the assumption that parity is only broken softly, both in the scalar potential and through the presence of non-hermitian vector-like masses for the $SU(2)$ -singlets. The relevant diagrams are those featured in fig. 3. We will concentrate first on diagrams where either h or h' propagate inside the loop.

In full generality, a Dirac fermion f that interacts with another fermion ψ , and a neutral scalar ϕ through Yukawa couplings of the form

$$\mathcal{L} \supset L(\bar{f}_R \psi_L) \phi + R(\bar{f}_L \psi_R) \phi + \text{h.c.}, \quad (3.102)$$

will receive a one-loop EDM given by

$$\frac{d_f}{e} = \frac{Q_\psi}{16\pi^2} \frac{m_\psi}{m_\phi^2} A(r) \text{Im}(LR^*), \quad (3.103)$$

where $r = m_\psi^2/m_\phi^2$, and the loop function A is given by

$$A(r) = \frac{1}{2(1-r)^2} \left(3 - r + \frac{2 \log r}{1-r} \right). \quad (3.104)$$

In the model we are considering, the Yukawa interactions involving both light and heavy fermions can be written as

$$\mathcal{L} \supset - \sum_{s=h,h'} s \left(\bar{d}_R \hat{\eta}^s d_L + \bar{d}_R \hat{\beta}^s D_L + \bar{D}_R \hat{\gamma}^s d_L + \bar{D}_R \hat{\delta}^s D_L \right) + \text{h.c.}, \quad (3.105)$$

where we are using notation specific to the down-quark sector, but analogous expressions apply for up-quarks and leptons (although the specific form of the Yukawa matrices will differ). It will be convenient to write the above matrices as $\hat{\omega}^s = R_{1s} \omega^h + R_{2s} \omega^{h'}$ (for $\omega = \eta, \beta, \gamma, \delta$), where $R_{11} = R_{22} = \cos \alpha$ and $R_{12} = -R_{21} = \sin \alpha$, with $\alpha \sim v/v'$ the mixing angle the Higgs sector. For the down-quark sector, the ω matrices can be conveniently written as follows

$$\eta^h = -\frac{m_d}{v}, \quad \beta^h = -\frac{1}{\sqrt{2}} (m_d \tilde{y}^* m_D^{-1}), \quad \gamma^h = \frac{\tilde{y}^T}{\sqrt{2}}, \quad \delta^h = \frac{v}{2} (\tilde{y}^T \tilde{y}^* m_D^{-1}), \quad (3.106)$$

and

$$\eta^{h'} = -\frac{m_d}{v'}, \quad \beta^{h'} = \frac{\tilde{y}'^*}{\sqrt{2}}, \quad \gamma^{h'} = -\frac{1}{\sqrt{2}} m_D^{-1} \tilde{y}'^T m_d, \quad \delta^{h'} = \frac{v'}{2} (m_D^{-1} \tilde{y}'^T \tilde{y}'^*). \quad (3.107)$$

The one-loop correction to the EDM of one of the SM-like quarks, d_i , is dominated by diagrams where the heavy mirror quarks propagate inside the loop. Since $m_{D_j} \sim M \gg m_h, m_{h'}$, we can expand the loop function $A(r)$ in the limit $r \gg 1$. Keeping the first two

terms, we find

$$\frac{d_{d_i}}{e} \simeq - \sum_{j,s} \frac{Q_d}{32\pi^2 m_{D_j}} \left(1 - \frac{m_s^2}{m_{D_j}^2}\right) \text{Im} \left(\hat{\beta}_{ij}^s \hat{\gamma}_{ji}^s \right). \quad (3.108)$$

The last factor in the previous expression can be written as

$$\text{Im}(\hat{\beta}_{ij}^s \hat{\gamma}_{ji}^s) = R_{1s} R_{2s} \text{Im}(\beta_{ij}^h \gamma_{ji}^{h'} + \beta_{ij}^{h'} \gamma_{ji}^h), \quad (3.109)$$

where we have taken into account that $\beta_{ij}^h \gamma_{ji}^h, \beta_{ij}^{h'} \gamma_{ji}^{h'} \in \mathbb{R}$, so those combinations don't appear on the right-hand-side. When summing over s in eq. (3.108), the contribution from the leading term in the $m_s^2/m_{D_j}^2 \ll 1$ expansion vanishes since $\sum_s R_{1s} R_{2s} = 0$. The leading contribution to d_{d_i} then reads

$$\begin{aligned} \frac{d_{d_i}}{e} &\simeq \sum_{j,s} \frac{Q_d}{32\pi^2} \frac{m_s^2}{m_{D_j}^3} R_{1s} R_{2s} \text{Im} \left(\beta_{ij}^h \gamma_{ji}^{h'} + \beta_{ij}^{h'} \gamma_{ji}^h \right) \\ &\simeq \sum_j \frac{Q_d}{32\pi^2} \frac{m_{h'}^2}{m_{D_j}^3} \sin \alpha \text{Im} \left(\beta_{ij}^{h'} \gamma_{ji}^h \right), \end{aligned} \quad (3.110)$$

where in the last step we have neglected the first term in parenthesis since it is suppressed by a factor of $\mathcal{O}((vv'/M^2)^2)$ with respect to the second, and we have only kept the contribution from h' , since the contribution from h is suppressed by an additional factor of $m_h^2/m_{h'}^2$.

From eqs. (3.106) and (3.107), we find

$$\beta_{ij}^{h'} \gamma_{ji}^h = \frac{1}{2} \tilde{y}_{ij}'^* \tilde{y}_{ij}. \quad (3.111)$$

When generalized parity is a good symmetry, we have $\tilde{y}' = \tilde{y}$, and therefore the above term is real, in turn leading to a vanishing EDM. In the presence of soft breaking through non-hermitian vector-like masses, the relationship $\tilde{y}' = \tilde{y}$ no longer holds, even if the Yukawa couplings in the flavor basis remain identical since the breaking is soft. To see

that the equality of the Yukawa couplings in the SM and mirror sectors no longer holds in the mass basis, it is useful to remind ourselves of the fermion mass diagonalization procedure discussed in section 3.A.2. When generalized parity remains unbroken, the unitary matrices of eq. (3.89) are such that $\mathcal{O}_F = \mathcal{O}_{F'}^*$ and $\mathcal{O}_{f'} = \mathcal{O}_f^*$. However, the non-hermiticity of the vector-like mass matrix means the unitary matrices needed to bring the 3×3 matrices of eq. (3.88) into real diagonal form will no longer satisfy this simple relation. Instead, writing the new vector-like mass matrices as $\mathcal{M}_f + i\Delta\mathcal{M}_f$, with both \mathcal{M}_f and $\Delta\mathcal{M}_f$ hermitian, the new unitary matrices are modified as follows

$$\begin{aligned} \mathcal{O}_{F'} &\rightarrow \tilde{\mathcal{O}}_{F'} = \mathcal{O}_{F'} + \Delta_{F'}, & \mathcal{O}_F &\rightarrow \tilde{\mathcal{O}}_F = \mathcal{O}_{F'}^* - \Delta_{F'}^*, \\ \mathcal{O}_f &\rightarrow \tilde{\mathcal{O}}_f = \mathcal{O}_f + \Delta_f, & \mathcal{O}_{f'} &\rightarrow \tilde{\mathcal{O}}_{f'} = \mathcal{O}_f^* - \Delta_f^*. \end{aligned} \quad (3.112)$$

The Δ matrices arise at $\mathcal{O}(|\Delta\mathcal{M}|/M)$, and are given by

$$(\Delta_{F'})_{ij} = i \sum_{k \neq i} \frac{[\Delta\tilde{\mathcal{M}}, \mathfrak{m}_F]_{ik}}{m_{F_i}^2 - m_{F_k}^2} (\mathcal{O}_{F'})_{kj}, \quad \text{and} \quad (\Delta_f)_{ij} = i \sum_{k \neq i} \frac{[\Delta\tilde{m}, \mathfrak{m}_f]_{ik}}{m_{f_i}^2 - m_{f_k}^2} (\mathcal{O}_f)_{kj}, \quad (3.113)$$

where

$$\Delta\tilde{\mathcal{M}} \equiv \mathcal{O}_{F'} \Delta\mathcal{M} \mathcal{O}_{F'}^\dagger \quad \text{and} \quad \Delta\tilde{m} \equiv \frac{vv'}{2} \tilde{y}^* \mathfrak{m}_F^{-1} \Delta\tilde{\mathcal{M}} \mathfrak{m}_F^{-1} \tilde{y}^T. \quad (3.114)$$

Using the above expressions in the definition of \tilde{y} , it is possible to write

$$\tilde{y}'^* = \tilde{y}^* (\mathbb{1} + \xi), \quad (3.115)$$

where ξ is a matrix with entries of $\mathcal{O}(|\Delta\mathcal{M}|/M)$. Explicitly, after some massaging,

$$\tilde{y}'_{ij}^* = \sum_l \tilde{y}_{il}^* \left(\delta_{lj} + i \frac{\Delta\tilde{\mathcal{M}}_{lj}}{m_{F_l} + m_{F_j}} (1 - \delta_{lj}) + i \frac{vv'}{2} \sum_{n, k \neq i} \frac{\tilde{y}_{kj}^* \tilde{y}_{kn}}{m_{f_i} + m_{f_k}} \frac{\Delta\tilde{\mathcal{M}}_{ln}}{m_{F_l} m_{F_n}} \right). \quad (3.116)$$

In total:

$$\text{Im}(\tilde{y}'_{ij*}\tilde{y}_{ij}) = |\tilde{y}_{ij}|^2 \times \mathcal{O}\left(\frac{\Delta\mathcal{M}}{M}\right). \quad (3.117)$$

Plugging this back into eq. (3.110), we have

$$\frac{d_{d_i}}{e} \simeq \sum_j \frac{Q_d}{32\pi^2} \frac{m_{h'}^2}{m_{D_j}^3} \sin\alpha \frac{1}{2} \text{Im}(\tilde{y}'_{ij*}\tilde{y}_{ij}) \simeq \frac{n_d Q_d}{32\pi^2} \frac{m_{d_i}}{M^2} \times \mathcal{O}\left(\frac{\Delta\mathcal{M}}{M}\right), \quad (3.118)$$

with $n_d = 3$ the number of mirror fermions appearing at the see-saw scale in the down-quark sector. The above expression also applies to the lepton sector, after making the obvious substitutions. In the up-quark sector, the expressions for the Yukawa couplings are somewhat different to those in eqs. (3.106) and (3.107), but can be similarly found by following the flavor-to-mass-basis rotation procedure outlined in section 3.A.2. In the end, diagrams where h' and the mirror partners of the u and c quarks propagate inside the loop give the leading contribution to the one-loop EDM. Thus, the above expression also applies for the up-quark sector, this time with $n_u = 2$ instead.

Additional contributions arise from diagrams where Z and Z' propagate inside the loop (see fig. 3). In this case, the leading contribution arises from diagrams involving Z' as well as heavy mirror fermions. In total, the final result is parametrically the same as that in eq. (3.118), except for an additional suppression by a factor of $g^2 \sin^2 \theta_w$.

Although the potential one-loop correction to $\bar{\theta}$ that could arise as a result of the soft breaking through non-hermitian vector-like masses was already shown to vanish in [78], this can also be seen from the calculation we have just performed. The relevant diagrams contributing to the quark mass matrix, and therefore to $\bar{\theta}$, are those of fig. 3, minus the external photon line. So although the appropriate loop function will be different, the overall correction will be similarly proportional to eq. (3.111). Using eq. (3.90) to rewrite

\tilde{y}' in terms of \tilde{y} , and the diagonal mass matrices, we find

$$\text{Im}(\delta m_{d_i}) \propto \text{Im}(\tilde{y}'_{ij}{}^* \tilde{y}_{ij}) \propto m_{d_i} \text{Im}((\tilde{y}^{-1})_{ji} \tilde{y}_{ij}). \quad (3.119)$$

As a result, the corresponding contribution to $\bar{\theta}$ from the down-quark sector reads

$$\sum_i \frac{\text{Im}(\delta m_{d_i})}{m_{d_i}} \propto \text{Im} \left(\sum_i (\tilde{y}^{-1})_{ji} \tilde{y}_{ij} \right) = 0. \quad (3.120)$$

Notice the sum over quark flavors is crucial in the above cancellation.

3.B.2 One-loop $\bar{\theta}$

The calculation of the one-loop correction to the quark mass matrix, and, in turn, to $\bar{\theta}$, proceeds along similar lines to the EDM calculation we have just discussed. The leading contribution to $\bar{\theta}$ comes from corrections to the light quark masses, and it is due to diagrams where either h' or ϕ propagate inside the loop.

Rotating from the gauge to the mass basis in the scalar sector requires performing a transformation $s_i \rightarrow R_{ij} s_j$, with $s_i = \{h, h', \phi\}$, and R is a 3×3 orthogonal matrix that we parametrize in terms of the various mixing angles as

$$R = \begin{pmatrix} c_\alpha c_\beta & c_\alpha s_\beta s_\gamma + s_\alpha c_\gamma & -c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma \\ -s_\alpha c_\beta & -s_\alpha s_\beta s_\gamma + c_\alpha c_\gamma & s_\alpha s_\beta c_\gamma + c_\alpha s_\gamma \\ s_\beta & -c_\beta s_\gamma & c_\beta c_\gamma \end{pmatrix} \quad (3.121)$$

where $c_\alpha = \cos \alpha$, $s_\alpha = \sin \alpha$, etc. Parametrically, we expect $c_\alpha \sim c_\beta \sim c_\gamma = \mathcal{O}(1)$, whereas $s_\alpha \sim v/v'$, $s_\beta \sim v/v_\phi$, and $s_\gamma \sim v_\phi/v'$. In the down-quark sector, the Yukawa interactions of eq. (3.105) need to be extended to include ϕ in the sum, and the $\hat{\omega}^s$

matrices are now given by

$$\hat{\omega}^s = R_{1s}\omega^h + R_{2s}\omega^{h'} + R_{3s}\omega^\phi \quad \text{for} \quad \omega = \eta, \beta, \gamma, \delta. \quad (3.122)$$

The expressions for ω^h and $\omega^{h'}$ are as in eqs. (3.106) and (3.107), whereas for ϕ we have

$$\begin{aligned} \eta^\phi &= -i \frac{vv'}{2} (\tilde{y}'^* m_D^{-1} \tilde{y} m_D^{-1} \tilde{y}^T), & \beta^\phi &= \frac{iv'}{\sqrt{2}} (\tilde{y}'^* m_D^{-1} \tilde{y}), \\ \gamma^\phi &= \frac{iv}{\sqrt{2}} (\tilde{y} m_D^{-1} \tilde{y}^T), & \delta_\phi &= -i \tilde{y}, \end{aligned} \quad (3.123)$$

where $\tilde{y} \equiv \mathcal{O}_{F'} \bar{y} \mathcal{O}_{F'}^\dagger$, as usual.

In the notation of eq. (3.102), the one-loop correction to $\text{Im}(\delta m_f)$ is given by

$$\text{Im}(\delta m_f) = \frac{m_\psi}{16\pi^2} \mathcal{F}(m_\psi, m_\phi) \text{Im}(LR^*), \quad (3.124)$$

where the loop function \mathcal{F} now reads

$$\mathcal{F}(m_\psi, m_\phi) = \frac{1}{m_\psi^2 - m_\phi^2} \left[m_\psi^2 \left(\log \frac{m_\psi^2}{\mu^2} - 1 \right) - m_\phi^2 \left(\log \frac{m_\phi^2}{\mu^2} - 1 \right) \right]. \quad (3.125)$$

For the case at hand, the leading one-loop correction to the mass of the SM-like fermions involves diagrams where the heavy mirror partners appearing at scale M propagate inside the loop. Specifically, in the down-quark sector, we have

$$\text{Im}(\delta m_{d_i}) = \sum_{s,j} \frac{m_{D_j}}{16\pi^2} \mathcal{F}(m_{D_j}, m_s) \text{Im}(\hat{\beta}_{ij}^s \hat{\gamma}_{ji}^s). \quad (3.126)$$

In analogy to the discussion in the previous section, the leading term in \mathcal{F} in the limit $m_{D_j} \gg m_s$ is independent of m_s , and its contribution to $\text{Im}(\delta m_f)$ vanishes as a result of the orthogonality of the mixing matrix in the scalar sector. The leading correction to

$\text{Im}(\delta m_{d_i})$ then reads

$$\begin{aligned}\text{Im}(\delta m_{d_i}) &\simeq \sum_{j,s} \frac{m_s^2}{16\pi^2 m_{D_j}} \log \frac{m_{D_j}^2}{m_s^2} \text{Im}(\hat{\beta}_{ij}^s \hat{\gamma}_{ji}^s) \\ &\simeq \sum_{j,s=h',\phi} \frac{m_s^2}{16\pi^2 m_{D_j}} \log \frac{m_{D_j}^2}{m_s^2} \text{Im}(s_\gamma \beta_{ij}^{h'} \gamma_{ji}^\phi + s_\gamma s_\alpha \beta_{ij}^\phi \gamma_{ji}^h),\end{aligned}\tag{3.127}$$

where in the last step we have neglected the contribution from h , which is suppressed by a factor of m_h^2/m_s^2 compared to that from h' and ϕ . The two terms inside the parenthesis are given by

$$s_\gamma \text{Im}(\beta_{ij}^{h'} \gamma_{ji}^\phi) = s_\gamma \frac{v}{2} \frac{\tilde{y}'_{ij} \tilde{y}_{jk} \tilde{y}_{ik}}{m_{D_k}} \sim s_\gamma \frac{m_{d_i} \bar{y}}{v'} \sim \frac{v_\phi m_{d_i} \bar{y}}{v'^2},\tag{3.128}$$

where in the last step we have substituted $s_\gamma \sim v_\phi/v'$, as we expect when $v_\phi \lesssim v'$, and

$$s_\gamma s_\alpha \text{Im}(\beta_{ij}^\phi \gamma_{ji}^h) = s_\gamma s_\alpha \frac{v'}{2} \frac{\tilde{y}'_{ik} \tilde{y}_{kj} \tilde{y}_{ij}}{m_{D_k}} \sim s_\gamma s_\alpha \frac{m_{d_i} \bar{y}}{v} \sim \frac{v_\phi m_{d_i} \bar{y}}{v'^2}.\tag{3.129}$$

Both terms are therefore of the same order. When $v_\phi \lesssim v'$, the contribution from ϕ to $\text{Im}(\delta m_{d_i})$ is subleading to that from h' . Setting $m_{h'} \simeq \sqrt{2\lambda} v' \sim v'$, we then have

$$\text{Im}(\delta m_{d_i}) \sim \frac{m_{d_i}}{16\pi^2} \frac{\bar{y} v_\phi}{M} \log \frac{M^2}{m_{h'}^2},\tag{3.130}$$

and the contribution to $\bar{\theta}$ from the down-quark sector reads

$$\bar{\theta} \simeq \sum_i \frac{\text{Im}(\delta m_{d_i})}{m_{d_i}} \sim \frac{1}{16\pi^2} \frac{\bar{y} v_\phi}{M} \log \frac{M^2}{m_{h'}^2}.\tag{3.131}$$

The above expression agrees with the parametric estimate presented in section 3.4.2, except for the log factor that is not captured in our spurion analysis.

3.C Kaon mixing

The Δm_K and $|\epsilon_K|$ parameters characterizing the kaon sector can be written as

$$\Delta m_K = 2\text{Re}(m_{12}^K), \quad \text{and} \quad |\epsilon_K| = \frac{\kappa_\epsilon |\text{Im}(m_{12}^K)|}{\sqrt{2}\Delta m_K}, \quad (3.132)$$

where $m_{12}^K \equiv \frac{1}{2m_K} \langle K^0 | \mathcal{H}_{\text{eff}} | \bar{K}^0 \rangle$, and \mathcal{H}_{eff} refers to the effective hamiltonian appropriate to describe kaon mixing.

In the SM, \mathcal{H}_{eff} is generated at one-loop through box diagrams involving two W gauge bosons. The corresponding contribution reads

$$\mathcal{H}_{\text{eff}} \supset -\frac{G_F^2 m_W^2}{4\pi^2} (\bar{d}_L \gamma_\mu s_L) (\bar{d}_L \gamma^\mu s_L) \sum_{\alpha, \beta} \lambda_\alpha \lambda_\beta F(x_\alpha, x_\beta) + \text{h.c.}, \quad (3.133)$$

where the loop function F is given by [168]

$$F(x_\alpha, x_\beta) = \frac{x_\alpha^2 \log x_\alpha}{(x_\beta - x_\alpha)(1 - x_\alpha)^2} \left(1 - 2x_\beta + \frac{x_\alpha x_\beta}{4} \right) + \{x_\alpha \leftrightarrow x_\beta\} \\ + \frac{1}{(1 - x_\alpha)(1 - x_\beta)} \left(\frac{7x_\alpha x_\beta}{4} - 1 \right), \quad (3.134)$$

and $\lambda_\alpha = V_{\alpha d}^* V_{\alpha s}$ for $\alpha = u, c, t$. In the present model, the sum over α and β in eq. (3.133) must be extended to include the additional members of the up-quark sector. The corresponding couplings can be read off from eq. (3.39), and are given by

$$\lambda_\alpha = \Delta V_{\alpha d}^* \Delta V_{\alpha s} \quad \text{for} \quad \alpha = U, C, T. \quad (3.135)$$

An additional contribution to \mathcal{H}_{eff} arises from diagrams involving one W and one W' .

In this case:

$$\mathcal{H}_{\text{eff}} \supset -\frac{G_F^2 m_W^2}{4\pi^2} \beta (\bar{d}_R s_L)(\bar{d}_L s_R) \sum_{\alpha, \beta} \lambda_\alpha^{LR} \lambda_\beta^{RL} \tilde{F}(\beta, x_\alpha, x_\beta) + \text{h.c.}, \quad (3.136)$$

where $\beta \equiv m_W^2/m_{W'}^2 = v^2/v'^2$, and the loop function now reads [169]

$$\tilde{F}(\beta, x_\alpha, x_\beta) = \sqrt{x_\alpha x_\beta} \{ (1 + \beta) I_2(x_\alpha, x_\beta, \beta) - (4 + \beta x_\alpha x_\beta) I_1(x_\alpha, x_\beta, \beta) \}, \quad (3.137)$$

with

$$\begin{aligned} I_1(\beta, x_\alpha, x_\beta) &= \frac{x_\alpha \log x_\alpha}{(1 - x_\alpha)(1 - \beta x_\alpha)(x_\alpha - x_\beta)} + \{x_\alpha \leftrightarrow x_\beta\} - \frac{\beta \log \beta}{(1 - \beta)(1 - \beta x_\alpha)(1 - \beta x_\beta)}, \\ I_2(\beta, x_\alpha, x_\beta) &= \frac{x_\alpha^2 \log x_\alpha}{(1 - x_\alpha)(1 - \beta x_\alpha)(x_\alpha - x_\beta)} + \{x_\alpha \leftrightarrow x_\beta\} - \frac{\log \beta}{(1 - \beta)(1 - \beta x_\alpha)(1 - \beta x_\beta)}. \end{aligned} \quad (3.138)$$

The relevant couplings follow from the interactions in eqs. (3.39) and (3.40). For the u and c quarks, we have $\lambda_\alpha^{LR} = \lambda_\alpha^{RL} = \lambda_\alpha$, whereas for their heavy partners

$$\lambda_\alpha^{LR} = \Delta V_{\alpha d}^* \Delta V'_{\alpha s}, \quad \text{and} \quad \lambda_\alpha^{RL} = \Delta V_{\alpha d}' \Delta V_{\alpha s}, \quad \text{for} \quad \alpha = U, C. \quad (3.139)$$

In the top sector, on the other hand, we have

$$\lambda_t^{LR} = V_{td}^* \Delta V'_{3s}, \quad \lambda_t^{RL} = \Delta V'_{3d} V_{ts}, \quad \text{and} \quad \lambda_T^{LR} = \Delta V_{3d}^* V_{ts}, \quad \lambda_T^{RL} = V_{td}^* \Delta V_{3s}. \quad (3.140)$$

Chapter 4

The Power of Covariantized Phase Space

4.1 Introduction

Novel representations of collider data play an increasingly important role in the modern machine learning (ML) era. In particular, a physically-motivated notion of the distance between collider events is a powerful input to ML-based methods for event classification and anomaly detection. Recently, optimal transport distances between energy flow distributions [170, 171, 172] or spectral functions [173] have given rise to a new geometry of collider events, with widespread applications. However, there is another natural distance between collider events that has yet to be fully explored: the one furnished by the phase space manifold on which scattering amplitudes are defined.

Traditionally, little attention has been paid to the intrinsic geometry of the phase space manifold itself, except within the context of the structure of S-matrix elements (especially at the boundaries of phase space). The deeper potential of the phase space manifold was recently highlighted in Ref. [7], which presented a covariant description of

the massless N -body phase space manifold Π_N and demonstrated that it is isomorphic to the product space of a $N - 1$ -dimensional simplex Δ_{N-1} and a $2N - 3$ -dimensional hypersphere S^{2N-3} , $\Pi_N \cong \Delta_{N-1} \times S^{2N-3}$. Explicit global coordinates can then be constructed on the phase space, promoting it to a metric space. This offers a new way to organize collider data that is firmly grounded in the underlying physics and furnishes a novel representation of collider data for any downstream machine learning applications.

In this paper we develop and demonstrate a practical metric for collider events using the global phase space coordinates presented in Ref. [7], tailored in part to the structure of events at hadron colliders. Constructing a true metric from these phase space coordinates poses a number of challenges, which we address in turn. Given that the dimension of the phase space manifold depends on the particle multiplicity, exclusive clustering of events into a fixed number of N “particles” enables the calculation of distances between different events on a phase space manifold of uniform dimension $3N - 4$. Ensuring that the phase space metric satisfies the identity of indiscernibles requires minimizing distances over $N!$ permutations of particles for each pair of events, an exponentially computationally intensive task that we efficiently approximate by exactly minimizing distances on the simplex part of the phase space manifold and azimuthally rotating events to minimize distances on the hypersphere part. Together, these steps ensure that the metric on phase space satisfies the identity of indiscernibles and define a physically principled prescription for event pre-processing. We further tailor the metric to global phase space coordinates relevant to hadron colliders, including both the center-of-mass frame for total events and a collinear frame for individual boosted jets.

Different types of collider events exhibit distinctive structures in phase space, making the phase space distance between events a potentially powerful tool for event classification and anomaly detection. We demonstrate this potential with a number of examples in which the phase space distances among ensembles of signal and background events are

inputs to simple ML algorithms for event classification. Distances on the phase space metric lead to competitive classification both at the level of entire collider events and within the substructure of boosted jets. Taken together, our results suggest that distances on a suitably-defined phase space manifold provide a practical and effective representation of collider data with numerous applications.

This paper is organized as follows: In section 4.2 we review the prescription for constructing the phase space manifold and explicit global coordinates in the center-of-mass frame presented in in Ref. [7] and develop an analogous prescription in the collinear limit for boosted objects. We then address the practical aspects of constructing a phase space metric for collider events in section 4.3, including the representation of different events in terms of a fixed number of particles and the efficient minimization of distances over particle permutations. We demonstrate the classification power of the phase space metric coupled to simple ML algorithms in section 4.4 with an event-level Standard Model $t\bar{t}$ vs. QCD dijet study, an event-level beyond-the-Standard Model isotropic Hidden Valley vs. QCD dijet study, and a boosted W boson vs. QCD jet study. We summarize our conclusions and sketch future directions in section 4.5, reserving technical details about the minimization over particle permutations for an appendix.

4.2 Intrinsic geometry of N -body phase space

In this section, we review the formalism of Ref. [7] for establishing the phase space manifold in the center-of-mass (CM) frame. We restrict our study to assume that experimentally detected particles are all massless, which will be a good approximation in the high energy collider environment of the LHC. For our purposes in this paper, we will need more explicit results for application on realistic collider event data. We will thus provide all necessary mappings from momentum space coordinates to the natural

coordinates on simplex-hypersphere formulation of phase space and the corresponding Riemannian metric. Additionally, we will construct the manifold on the phase space of a boosted jet, effectively in the infinite-momentum frame. The coordinate of the jet's axis is special, breaking the event's $O(3)$ invariance in the CM frame, and also the mass of the jet is not necessarily measured and constrained, which introduces a new coordinate.

4.2.1 Phase space in the CM frame: a product manifold

We start with the volume form of N -body phase space in terms of the four-momenta p_i of the N massless particles constituting an event in Cartesian coordinates, where

$$d\Pi_N = (2\pi)^{4-3N} \left[\prod_{i=1}^N d^4 p_i \delta^+(p_i^2) \right] \delta^{(4)} \left(Q - \sum_{i=1}^N p_i \right). \quad (4.1)$$

Here $Q = (Q, 0, 0, 0)$ is the total energy of the event in the CM frame and $\delta^+(p_i^2) = \delta(p_i^2) \Theta(p_{0i})$ enforces that all particles are on-shell with positive energy. The on-shell constraints can be satisfied by re-expressing the four-momenta in terms of the two spinors for each particle, where

$$\lambda_i^\alpha = \frac{1}{\sqrt{p_{0i} + p_{3i}}} \begin{pmatrix} p_{0i} + p_{3i} \\ p_{1i} + ip_{2i} \end{pmatrix}, \quad \tilde{\lambda}_i^{\dot{\alpha}} = \frac{1}{\sqrt{p_{0i} + p_{3i}}} \begin{pmatrix} p_{0i} + p_{3i} \\ p_{1i} - ip_{2i} \end{pmatrix}, \quad (4.2)$$

with $i = 1, \dots, N$ labeling the particle and $\alpha, \dot{\alpha} = 1, 2$ labeling the spinor components. Note that we have here $\tilde{\lambda}_i^{\dot{\alpha}} = \lambda_i^{\alpha*}$ because the momentum matrix $p_i = \lambda_i \tilde{\lambda}_i$ is Hermitian.

We now define two N -component vectors

$$\vec{u} \equiv (\lambda_1^1 \lambda_2^1 \dots \lambda_N^1), \quad \text{with components } u_i = \sqrt{p_{0i} + p_{3i}}, \quad (4.3)$$

$$\vec{v} \equiv (\lambda_1^2 \lambda_2^2 \dots \lambda_N^2), \quad \text{with components } v_i = \frac{p_{1i} + ip_{2i}}{\sqrt{p_{0i} + p_{3i}}}. \quad (4.4)$$

Notice that \vec{u} is a real vector while \vec{v} is complex. Written in terms of these \vec{u}, \vec{v} vectors, eq. (4.1) now becomes

$$\begin{aligned} d\Pi_N &= (2\pi)^{4-3N} Q^{2N-4} \left[\prod_{i=1}^N \frac{d^2\lambda_i^1 d^2\lambda_i^2}{\text{U}(1)} \right] \\ &= (2\pi)^{4-3N} Q^{2N-4} \frac{d^N u d^N v}{\text{U}(1)^N} \delta(1 - |\vec{u}|^2) \delta(1 - |\vec{v}|^2) \delta^{(2)}(\vec{u}^\dagger \vec{v}), \end{aligned} \quad (4.5)$$

where the division by $\text{U}(1)$ represents implicit restriction to one element of the little group action on the spinors. By introducing the new coordinate

$$\rho_i \equiv u_i^2, \quad (4.6)$$

the \vec{u} part of the volume form reduces to

$$\frac{d^N u}{\text{U}(1)^N} \delta(1 - |\vec{u}|^2) = \frac{1}{2^N} \prod_{i=1}^N [d\rho_i] \delta\left(1 - \sum_{i=1}^N \rho_i\right) = \frac{1}{2^N} d\Delta_{N-1}, \quad (4.7)$$

where $d\Delta_{N-1}$ is the flat measure on a $N - 1$ -dimensional simplex. We therefore obtain the simplex part of the phase space manifold.

For the terms in eq. (4.5) involving \vec{v} , we first enforce the orthogonality condition $\delta^{(2)}(\vec{u}^\dagger \vec{v})$ to reduce the dimension of \vec{v} by 1 via the removal of the last component of \vec{v} , i.e., v_N . Thus, we have

$$d^N v \delta(1 - |\vec{v}|^2) \delta^{(2)}(\vec{u}^\dagger \vec{v}) = \frac{d^{N-1} v}{\rho_N} \delta(1 - |\vec{v}|^2 - |v_N|^2). \quad (4.8)$$

Note that the \vec{v} on the right hand side of eq. (4.8) now has only $N - 1$ components, and v_N can be expressed in terms of the components of \vec{u} and \vec{v} . We then perform a change

of variables from \vec{v} to \vec{v}' , where \vec{v}' is an $N - 1$ -component complex unit vector,

$$\frac{d^{N-1}v}{\rho_N} \delta(1 - |\vec{v}|^2 - |v_N|^2) = d^{N-1}v' \delta(1 - |\vec{v}'|^2) = dS^{2N-3}, \quad (4.9)$$

which is manifestly the measure for the sphere S^{2N-3} .

Putting it all together, we have shown that the phase space is a product manifold of the simplex Δ_{N-1} and the sphere S^{2N-3} :

$$d\Pi_N = \frac{(2\pi)^4}{Q^4} \left(\frac{Q^2}{16\pi^3} \right)^N d\Delta_{N-1} dS^{2N-3}. \quad (4.10)$$

One can easily check that dimensionality works since the sum of the simplex and sphere dimensions gives us $(N - 1) + (2N - 3) = 3N - 4$, which is the correct phase space dimension for a system of N massless particles. Also, because we will always compare events with the same total energy Q , we will often rescale energies so that $Q = 1$, and so the overall dimensionality of the phase space metric is eliminated.

4.2.2 Phase space in the CM frame: metric and explicit coordination

From this factorization of the phase space manifold into a simplex and a sphere, we can then introduce explicit, global coordinates. First, let's explicitly work out the full transformation from \vec{v} to \vec{v}' . This can be expressed by

$$\begin{pmatrix} v'_1 \\ v'_2 \\ \vdots \\ v'_{N-1} \end{pmatrix} = \frac{1}{u_N(1 - u_N^2)} \hat{R} \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_{N-1} \end{pmatrix} \quad (4.11)$$

where the matrix \hat{R} is

$$\hat{R} = \begin{pmatrix} u_1^2 + u_N(1 - u_1^2 - u_N^2) & u_1 u_2(1 - u_N) & \cdots & u_1 u_{N-1}(1 - u_N) \\ u_1 u_2(1 - u_N) & u_2^2 + u_N(1 - u_2^2 - u_N^2) & \cdots & u_2 u_{N-1}(1 - u_N) \\ \vdots & \vdots & \ddots & \vdots \\ u_1 u_{N-1}(1 - u_N) & u_2 u_{N-1}(1 - u_N) & \cdots & u_{N-1}^2 + u_N(1 - u_{N-1}^2 - u_N^2) \end{pmatrix}. \quad (4.12)$$

Note that the above matrix \hat{R} is a real, symmetric and positive-definite, and the reality of \hat{R} ensures that the real and imaginary components of \vec{v} are not mixed with each other. Furthermore, the determinant of the matrix is u_N^{-1} , and inverting \hat{R} gives a Jacobian $J = u_N^2$, which ensures that $d^{N-1}v = u_N^2 d^{N-1}v' = \rho_N d^{N-1}v'$.

The vector \vec{v}' can now be conveniently expressed by a generalized spherical coordinate system,

$$\begin{aligned} v'_1 &= e^{-i\xi_1} \cos \eta_1 \\ v'_2 &= e^{-i\xi_2} \sin \eta_1 \cos \eta_2 \\ &\vdots \\ v'_{N-2} &= e^{-i\xi_{N-2}} \sin \eta_1 \cdots \sin \eta_{N-3} \cos \eta_{N-2} \\ v'_{N-1} &= e^{-i\xi_{N-1}} \sin \eta_1 \cdots \sin \eta_{N-3} \sin \eta_{N-2}, \end{aligned} \quad (4.13)$$

where $\xi_i \in [0, 2\pi]$ and $\eta_i \in [0, \pi/2]$. These coordinates explicitly enforce the normalization $|\vec{v}'|^2 = 1$. In the case of $N = 2$, the sphere part of the manifold is $S^{2 \cdot 2 - 3} = S^1$, and so no η coordinate remains, and we're left with only one parameter for the sphere part, i.e., ξ_1 .

To summarize, we have shown that the $(3N - 4)$ degrees of freedom of the phase space manifold are composed of two groups of coordinates, with $(N - 1)$ ρ coordinates on the simplex Δ^{N-1} , and $(N - 1)$ ξ together with $(N - 2)$ η coordinates on the hypersphere

S^{2N-3} .

Armed with a system of explicit coordinates, we can now easily define the corresponding metric on the phase space manifold. The line element of the simplex is

$$ds_{\Delta_{N-1}}^2 = \sum_{i=1}^N d\rho_i^2. \quad (4.14)$$

The summation above is over N coordinates despite the fact that the simplex has dimension $N - 1$. Indeed, these coordinates are not all independent but are instead constrained via the δ -function in eq. (4.7). With a suitable change of variables, the simplex line element could be written using only $N - 1$ coordinates, but this is not necessary for our purposes. The line element of the hypersphere follows a recursive description

$$ds_{S^{2N-3}}^2 = d\eta_{N-2}^2 + \cos^2 \eta_{N-2} d\xi_{N-1}^2 + \sin^2 \eta_{N-2} ds_{S^{2N-5}}^2 \quad (4.15)$$

starting from $ds_{S^1}^2 = d\xi_1^2$.

In this paper, we are more interested in the distances between pairs of events assigned according to their positions on the phase space manifold, so these line elements can be promoted to honest Riemannian metrics that satisfy symmetry, identity of indiscernibles, and the triangle inequality. For two events that we label A and B , their metric distance on the simplex is

$$d_{\Delta}(\vec{\rho}_A, \vec{\rho}_B) \equiv \sqrt{(\vec{\rho}_A - \vec{\rho}_B)^2} = \sqrt{\sum_{i=1}^N (\rho_{iA} - \rho_{iB})^2}. \quad (4.16)$$

Distances on the hypersphere are determined by the angle between the two events, and

therefore the metric on the hypersphere is

$$d_S(\vec{v}'_A, \vec{v}'_B) \equiv \cos^{-1}(\Re \vec{v}'_A \dagger \vec{v}'_B) = \cos^{-1} \left(\Re \sum_{i=1}^{N-1} v'_{iA}{}^* v'_{iB} \right), \quad (4.17)$$

where $\Re \vec{v}'_A \dagger \vec{v}'_B$ is the real part of the inner product of the vectors \vec{v}'_A and \vec{v}'_B . Any convex linear combination of these two metrics is itself a metric, and could therefore be a metric on the phase space manifold. However, the volume element of phase space is fixed, which constrains the determinant of the metric.

By demanding that the metric on phase space reproduces the dependence on multiplicity N in the volume element of eq. (4.10), we can express the metric between two events A and B on phase space as dependent on a single parameter $c > 0$, where

$$d_{\Pi}(\vec{\rho}_A, \vec{v}'_A; \vec{\rho}_B, \vec{v}'_B) = \sqrt{\frac{1}{16\pi^2} \left(\frac{c}{4}\right)^{\frac{3-2N}{3N-4}} d_{\Delta}^2(\vec{\rho}_A, \vec{\rho}_B) + \frac{1}{4\pi^2} \left(\frac{c}{4}\right)^{\frac{N-1}{3N-4}} d_S^2(\vec{v}'_A, \vec{v}'_B)}. \quad (4.18)$$

Varying the value of c then reweights the simplex and sphere contributions to the distance between events. With only the volume of phase space as a constraint, there is no natural value for c , and so in our studies we will consider a few values to identify where on phase space are events most different.

Actually, to ensure that d_{Π} is a metric on phase space, we must further enforce permutation invariance of the particles in the events. Specifically, if the events were identical, $A = B$, then for two distinct orderings of the particles, the metric would in general be non-zero. This would immediately fail the identity of indiscernibles requirement, and so we must additionally enforce that the metric is minimized over permutations of particles in the events. In practice, global minimization over $N!$ permutations is exponentially computationally intensive and unfeasible even for reasonably small N . Nevertheless, in section 4.3 we will discuss simple pre-processing steps that can be implemented to

approximate the global permutation minimum, and, at the very least, ensure that the distance from an event to itself is 0.

4.2.3 Phase space in the collinear limit

The phase space for an isolated jet, as a collimated, high-energy collection of particles, is distinct from the phase space of a high-energy event. In particular, our starting approximation for a jet will be in the infinite momentum limit, in which momentum transverse to its axis is parametrically smaller than the momentum along the axis. In this limit, there is no Lorentz transformation that can be performed to boost to the rest frame of the jet, and so this requires a new analysis of the corresponding manifold on which particle momenta live.

Without loss of generality, we assume that the net momentum of the jet lies along the $+\hat{z}$ axis and introduce the lightcone momentum components

$$p^- \equiv p_0 + p_3, \quad p^+ \equiv p_0 - p_3. \quad (4.19)$$

Then, for our purposes here, we define a jet as a collection of particles for which their net $-$ component is parametrically larger than the $+$ component: $p^- \gg p^+$. To just find a jet, we require some net p^- but are inclusive to the value of p^+ , without further restrictions on the structure of the jet. Additionally, as the axis of the jet is defined to be the direction of net momentum, the net momentum transverse to the axis is 0. We denote these components of momentum transverse to the jet axis as \vec{p}_\perp .

With this set-up, the differential phase space volume of a jet that consists of N

massless particles is then:

$$d\Pi_N^{(J)} = (2\pi)^{4-3N} \left[\prod_{i=1}^N d^4 p_i \delta^+(p_i^2) \right] dp^+ \delta \left(p^- - \sum_{i=1}^N p_i^- \right) \delta \left(p^+ - \sum_{i=1}^N p_i^+ \right) \delta^{(2)} \left(\sum_{i=1}^N \vec{p}_{\perp i} \right). \quad (4.20)$$

As with phase space in the CM frame, we introduce vectors \vec{u} and \vec{v} to automatically incorporate the on-shell constraints, where now the individual components are:

$$u_i = \sqrt{\frac{p_i^-}{p^-}}, \quad v_i = \frac{p_{1i} + i p_{2i}}{\sqrt{p^+ p_i^-}}. \quad (4.21)$$

Note that unlike their definition for phase space in the CM frame, these vectors are dimensionless, which is preferred because to good approximation jets are scale invariant.

Phase space volume in these coordinates then becomes

$$d\Pi_N^{(J)} = (2\pi)^{4-3N} (p^-)^{2N-4} \left(\frac{p^+}{p^-} \right)^{N-2} \frac{d^N u d^N v}{U(1)^N} dp^+ \delta(1 - |\vec{u}|^2) \delta(1 - |\vec{v}|^2) \delta^{(2)}(\vec{u}^\dagger \vec{v}). \quad (4.22)$$

Here we can go a bit further and simplify the phase space volume by incorporating some additional physics. In the soft and/or collinear limit that dominates the description of a jet, perturbative QCD is approximately scale-invariant, which imposes structure on the squared matrix element $|\mathcal{M}|^2$ as a function of the present scales. By dimensional analysis, the squared matrix element for N final state particles must have overall dependence on p^- as

$$|\mathcal{M}|^2 \propto (p^-)^{4-2N}, \quad (4.23)$$

in order to ensure that probabilities are dimensionless. Further, assuming scale invari-

ance, probabilities are unchanged if the ratio

$$\chi \equiv \frac{p^+}{p^-} \quad (4.24)$$

is rescaled as $\chi \rightarrow \lambda\chi$, for any $\lambda > 0$, and so the matrix element must have dependence on χ as

$$|\mathcal{M}|^2 \sim \chi^{1-N} (p^-)^{4-2N}. \quad (4.25)$$

Including this knowledge of the matrix element, the phase space now becomes

$$d\Pi_N^{(J)} |\mathcal{M}|^2 \sim (2\pi)^{4-3N} d\ln \chi \frac{d^N u d^N v}{\text{U}(1)^N} \delta(1 - |\vec{u}|^2) \delta(1 - |\vec{v}|^2) \delta^{(2)}(\vec{u}^\dagger \vec{v}). \quad (4.26)$$

If one is inclusive over particle number N , then this exact scale invariance in QCD is modified by an anomalous dimension (and the running coupling), proportional to α_s . In general, the matrix element will exhibit singularities on the phase space when pairs of particles become collinear, but those details depend precisely on the specific dynamics of the jet ensemble of interest.

One can now proceed to construct explicit global coordinates on the jet's phase space manifold. As in the case of events in the CM frame, we will have coordinates ρ, v', ξ, η expressed with the same formulae as in eqs. (4.6), (4.11) and (4.13). Therefore, we again end up with $(N - 1)$ ρ coordinates for the simplex and $(2N - 3)$ ξ and η coordinates for the hypersphere. For a jet, we will also have an additional χ coordinate whose form is dictated by the specific structure of the squared matrix element. As a result, the phase space for collinear jets is a product of three parts: the simplex, the hypersphere, and a scale factor.

For two jets A and B , the metric of the scale factor χ is simply the Euclidean metric

on the real line:

$$d_\chi(\chi_A, \chi_B) = |\log(\chi_A) - \log(\chi_B)|. \quad (4.27)$$

Now, with three components to the jet phase space manifold and one constraint on its volume form eq. (4.26), we can express the metric as a convex sum in terms of two parameters $c_1, c_2 > 0$, where

$$\begin{aligned} d_{\Pi(J)}^2(\vec{\rho}_A, \vec{v}'_A, \chi_A; \vec{\rho}_B, \vec{v}'_B, \chi_B) & \quad (4.28) \\ &= \frac{1}{16\pi^2} \left(\frac{c_1}{4}\right)^{\frac{3-2N}{3N-3}} \left(\frac{c_2}{16\pi^2}\right)^{-\frac{1}{3N-3}} d_\Delta^2(\vec{\rho}_A, \vec{\rho}_B) + \frac{1}{4\pi^2} \left(\frac{c_1}{4}\right)^{\frac{N}{3N-3}} \left(\frac{c_2}{16\pi^2}\right)^{-\frac{1}{3N-3}} d_S^2(\vec{v}'_A, \vec{v}'_B) \\ &+ \left(\frac{c_1}{4}\right)^{\frac{3-2N}{3N-3}} \left(\frac{c_2}{16\pi^2}\right)^{\frac{3N-4}{3N-3}} d_\chi^2(\chi_A, \chi_B), \end{aligned}$$

where d_Δ is again given by eq. (4.16) and d_S by eq. (4.17). In some jet studies, there is an expected natural mass scale, for example, when searching for resonances that decay hadronically, and in those cases, one would impose a narrow cut on the jet mass. As such, the value of χ would vary very little on a jet-by-jet basis in the relevant ensemble, and so the contribution to the metric from d_χ can typically be ignored. Additionally, as discussed for events in the CM frame, permutation invariance of the particles must be imposed externally to ensure that the distance between a jet and itself is 0. Pre-processing steps that ensure the identity of indiscernibles will be discussed in the next section.

4.3 Implementation of the phase space metric for classification

In this section, we describe the practical implementation of the metric on phase space that we explore in a specific case study later. This will involve significant pre-processing steps that can get a long way toward permutation invariance and exact minimization required for the metric on phase space reviewed and established in the previous section to indeed satisfy the requirements of a metric. We also describe the general techniques employed for post-processing the phase space representation of the events for the task of event classification, which will necessarily introduce machine learning to establish boundaries between signal- and background-rich regions.

Pre- and post-processing of particle physics collider event data for any analysis, with machine learning or not, has long been employed to simplify the representation of the data. Typically, pre-processing is done in an ad hoc manner, modding out by symmetries that are expected to hold to good approximation, but are rarely ever quantitatively justified. On the phase space manifold with explicit global coordinates and a local metric, these pre-processing steps are motivated by the necessary minimization over particle ordering and orientation to enforce the properties of the metric. The steps discussed here can therefore inform a more principled pre-processing prescription in other contexts.

Throughout this section, we will refer to the complete collection of particles produced in a collision as an “event”. We will later consider both analysis of complete collision events as well as individual jets excised from events, but refer to either case as “events”, with the understanding that the appropriate description of phase space is used depending on the particular application.

4.3.1 Pre-processing of collider events

We first discuss the pre-processing that we perform on events. These pre-processing steps ensures that the distance between an event and itself vanishes according to the metric on phase space.

Event representation

This formulation of a metric on phase space explicitly fixes the particle multiplicity N . So, the first step of pre-processing is to map any given event onto N -body phase space. We accomplish this by clustering the event into N jets with the k_T clustering algorithm. Further, for each identified jet, we also set its four-momentum to be massless (fixing hadronic coordinates).

For small N , we might expect that this masslessness restriction is a poor approximation, as each jet may consist of relatively hard, relative wide angle particles that generate a relatively large mass for the jets. However, as multiplicity N increases, it is expected that this masslessness restriction weakens as small clusters of hadrons are resolved that have high energy, but are collinear with one another. At any rate, in principle mass can be incorporated in the phase space manifold and calculations in perturbative QCD assume partons are massless. We leave a detailed analysis of the limitations of this massless approximation to future work.

Particle ordering

Events on which only energies and momenta are measured have no natural ordering of their constituent particles and so any function of them must be permutation invariant. For the phase space metric, this permutation invariance is implemented as a minimization of the distance between two events over all possible permutations of particles. This

minimization ensures that the metric then satisfies both the triangle inequality and the identity of indiscernibles. However, for every pair of events with N particles on which their distance is calculated, we would have to scan over all possible $N!$ distinct orderings to identify the minimum. Such an exhaustive search is clearly computationally impossible even for moderate N , so we will propose a different approach here. While minimization over an exponentially-large number of permutations is ripe for machine learning, we present a much more pedestrian proposal here that is provably correct for the simplex submanifold, and will motivate prescriptions on all of phase space.

Recall that the metric on the simplex is the Euclidean metric in N dimensions:

$$d_{\Delta}^2(\vec{\rho}_A, \vec{\rho}_B) = \min_{\sigma \in S_N} \sum_{i=1}^N (\rho_{iA} - \rho_{\sigma(i)B})^2. \quad (4.29)$$

Of course, the simplex is $N - 1$ dimensional by energy conservation, but we will work with this manifestly permutation invariant form. Also, here we have explicitly included the minimization over the permutations of particles in event B , where σ is an element of the symmetric group S_N . Without loss of generality, we can assume that the particles in event A are, say, ordered monotonically decreasing in ρ value. Then, the permutation of the particles of event B that minimizes this distance is also when the particles in event B are ordered in decreasing ρ value. That is,

$$d_{\Delta}^2(\vec{\rho}_A, \vec{\rho}_B) = \sum_{\substack{i=1 \\ \rho_{iA} > \rho_{(i+1)A} \\ \rho_{iB} > \rho_{(i+1)B}}}^N (\rho_{iA} - \rho_{iB})^2. \quad (4.30)$$

We provide an elementary proof that this ordering does indeed render the distance a metric on the simplex in section 4.A.

d_{Π}^2 is of course a convex linear combination of the squared metrics on the simplex and

the hypersphere, and so ordering particles according to their ρ value does not guarantee minimization over permutations of the complete metric on phase space. Nevertheless, this pre-processing step is very simple and computationally cheap, and at the very least approximates a local minimum over permutations. To practically implement this ordering, recall that the definition of ρ is

$$\rho = p_0 + p_3, \quad (4.31)$$

the sum of the energy and z -component of momentum of the particle. For events in the CM frame, there is possibly an $O(3)$ rotation symmetry of particle momenta. For events at a hadron collider, this $O(3)$ symmetry is broken to $O(2)$ about the beams and reflection of the beams, but still, the direction of the $+z$ -axis is not unambiguously defined. In our analyses, then, we will consider the following three different particle orderings to probe the sensitivity to the choice of axis:

- Descending order in $p_0 + p_3$,
- Descending order in $p_0 - p_3$,
- Descending order in $p_T \equiv \sqrt{p_1^2 + p_2^2}$.

Azimuthal rotation

Even with the particle ordering that minimizes distances on the simplex, there are still minimization ambiguities on the hypersphere that can be addressed with pre-processing. Recall that the metric on the hypersphere is

$$d_S(\vec{v}'_A, \vec{v}'_B) = \min_{\arg(\vec{v}'_B)} \cos^{-1} \left(\Re \left[\vec{v}'_A{}^\dagger \vec{v}'_B \right] \right). \quad (4.32)$$

Here, the minimization is a bit subtle. Event-by-event, there is an $O(2)$ azimuthal symmetry about the beam, for which the momentum transverse to the beam can be rotated by a common angle or reflected. In the metric, then, $SO(2)$ rotation invariance is manifest as minimization over the argument of the vector \vec{v}'_B , for example, and we fix the vector \vec{v}'_A because only their relative phase is important. We will address reflection invariance shortly.

Then, because \cos^{-1} monotonically decreases as its argument increases from 0, the minimization over the phase of \vec{v}'_B is implemented by maximizing

$$\max_{\arg(\vec{v}'_B)} \Re \left[\vec{v}'_A \dagger \vec{v}'_B \right] \leq \left| \vec{v}'_A \dagger \vec{v}'_B \right|. \quad (4.33)$$

Thus, because the real part of this vector dot product is bounded from above by its absolute value, even when maximizing over the argument of \vec{v}'_B , we should just use the absolute value in calculating the metric on the hypersphere. This prescription will minimize the distance on the hypersphere for any given particle ordering, but does not ensure that that ordering is a global minimum. Nevertheless, with the ordering prescription on the simplex, this will produce at least a local minimum of the metric on the entire phase space manifold.

To account for reflections in the transverse space, note that a reflection of $p_y \rightarrow -p_y$ corresponds to complex conjugation in the vector \vec{v} . Then, to minimize the hypersphere metric over reflections with a given particle ordering, we simply choose the larger of $\left| \vec{v}'_A \dagger \vec{v}'_B \right|$ and $\left| \vec{v}'_A \top \vec{v}'_B \right|$. Again, we only reflect/complex conjugate vector \vec{v}'_A in the product because the metric is only sensitive to the relative orientation between events. Finally, for a given particle ordering, the form of the metric we take on the hypersphere is

$$d_S(\vec{v}'_A, \vec{v}'_B) = \cos^{-1} \left(\max \left[\left| \vec{v}'_A \dagger \vec{v}'_B \right|, \left| \vec{v}'_A \top \vec{v}'_B \right| \right] \right). \quad (4.34)$$

Crucially, the particle ordering on the simplex ensures that the distance between an event and itself vanishes, and this prescription for minimization over azimuthal symmetries ensures that the distance on the hypersphere also vanishes. Therefore, these pre-processing prescriptions explicitly ensure that the metric on phase space satisfies the quality of identity of indiscernibles.

4.3.2 Processing the phase space manifold

With the pre-processing steps discussed above, we now turn to discussion of practical evaluation of the coordinates on the phase space manifold and the corresponding distances between realistic events.

Definition of the simplex and sphere coordinates

If events enjoy invariance to $O(3)$ transformations of the celestial sphere and all particles are detected, then there is a continuous infinity of possible choices for the \vec{u} and \vec{v} coordinates, all of which are equally valid. This is because, in the center-of-mass frame, the net momentum is 0 and so the net momentum transverse to any fixed axis is also 0. However, in this paper, we either consider events produced at a hadron collider, for which numerous particles are lost down the beam region and undetected, or of individual jets for which a center-of-mass frame does not exist. These systems then break the infinite degeneracy of coordinate systems and reduce to just a couple of different choices that respect the symmetries of the events in the lab frame.

The possible choices are exclusively set by the coordinates on the simplex, \vec{u} , or equivalently, ρ . Given fixed coordinates for which the \hat{z} -axis lies along the beams and

one beam is fixed to be the $+\hat{z}$ direction, there are two possibilities:

$$\rho = p_0 + p_3, \quad \text{or} \quad \rho = p_0 - p_3. \quad (4.35)$$

Note that either choice produces a valid metric, because there is no absolute frame for the coordinates on phase space. Instead, the choice of coordinates should be determined by the particular task at hand; here, we would choose the coordinates based on optimization of discrimination power. However, lacking infinite computing resources, we will simply study the response of our classification task with different choices of coordinates. Note that this choice of coordinates ensures that the momentum transverse to the z axis, that defines the coordinates \vec{v} , is (very nearly) 0, up to the small residual transverse momentum lost down the beams. However, the visible particles in general have a large net momentum along the z axis, which will need to be boosted away to render the events in the center-of-mass frame. Such a boost does not affect the direction of the z axis, and further the ρ coordinate transforms homogeneously under a z -axis boost, where

$$\rho \rightarrow e^\eta \rho, \quad (4.36)$$

where the (pseudo)rapidity of the visible particles in the lab frame is η .

Given a possible natural choice for the \vec{u} coordinates, the \vec{v} coordinates are unique. Conservation of transverse momentum requires that \vec{v} is transverse to \vec{u} , $\vec{u}^\dagger \vec{v} = 0$, and so \vec{v} must be formed from a complex linear combination of p_1, p_2 . We can in principle rotate the x and y axes about the z axis however we desire, naively producing another continuous infinity of axis choices. However, focused on ultimately evaluating the metric distance between events, we have already optimized this azimuthal rotation in the implementation of the metric on the hypersphere. Global event rotations about the z axis, as would be

used to define new (x, y) axes, cannot affect the metric distance, and so we can happily take whatever axes are provided to us in the representation of the data.

For studies on jets, we simply note that all of these considerations apply there, too.

Compute distances

At this point, we can actually compute the simplex and sphere distances separately to get 2 distance matrices (or 3 for jet level because of the χ coordinate). These distances are calculated according to eqs. (4.16), (4.17) and (4.27)

4.3.3 Post-processing for classification

Relative weightings

We must now choose the relative coefficient between the simplex and sphere parts in the metric to determine how to combine them into a total distance. In principle, one must also vary the coefficient in front of the χ component for the jet-level metric, though we do not consider this variation in the present work.

There is again a continuous infinity of coefficient choices here. We do not need to choose the weighting to minimize distance. Rather, any finite choice represents a valid metric. In a perfect world with infinite computational power, we would want to continuously vary the weighting(s) to maximize final AUC. Here, we will simply choose three different benchmark weightings.

Machine learning algorithm

Once the phase space coordinate of an event or a jet is obtained, one needs a classifier to determine whether it belongs to the signal or the background. In principle, any classifier can be employed, such as deep neural networks which are known to deliver

amazing tagging performance. Given that our purpose here is to better understand the power of phase space, it is more desirable to choose a basic classifier which can elucidate the underlying metrics and coordinates themselves. Therefore, we only consider the classifiers k -Nearest Neighbors (kNN) and support vector machine (SVM), leaving the experimentation of more sophisticated models to future study.

The supervised kNN [174] model works simply by using a majority vote of the closest k neighbors of an unlabeled data point to determine its membership. Here, k is a hyperparameter of the model to be tuned by a validation set, with a smaller value indicating that only local information is needed in classification. The only input kNN relies upon is the notion of a pairwise distance, which in our case is the phase space metric. We also test SVM, which is a more sophisticated algorithm that works by constructing a hyperplane to maximally separate the data points belonging to different classes.

4.4 $t\bar{t}$ vs. QCD dijet event classification

Now we apply the preceding formalism to event classification. For clarity of presentation, we limit ourselves to supervised classification tasks, but we emphasize that metrization and coordination of the phase space manifold have far-reaching potentials beyond tagging. This will be briefly discussed in section 4.5, and a full development of various applications will be the topic for future studies.

The separation of signals from backgrounds is a key task for the search of new physics beyond the Standard Model (BSM) at the Large Hadron Collider (LHC). Oftentimes, the majority of background processes are made up of QCD events, whose dynamics in the perturbative regime is dominated by the emission of soft and collinear quarks and gluons due to the small 't Hooft coupling $\lambda = g_s^2 N_C$ of QCD at high collider scales. Much of the kinematic information from this initial perturbative showering is retained through

subsequent hadronization, and the final state particles of QCD events, mostly populated by light mesons, display jet-like structure that occupy a certain marked region in the phase space manifold.

In contrast, various kinds of signals (including both SM and BSM events) may show distinct phase space structures due to the unique nature of their respective underlying physics. For example, top-quark pair production, a typical SM process, gives rise to an event configuration with six prongs (on average), inducing a more uniform distribution than the QCD background. This will in turn differentiate the phase space region occupied by $t\bar{t}$ events from that of QCD events. In this section we perform the classification task for top pair production vs. QCD dijets as a proof of concept for our phase space formalism.

One can of course go further and consider classifying more exotic signal events. For example, many BSM models can produce high-multiplicity events whose final state particles have a uniform distribution on N -body phase space. Hidden Valley models, which emerge from solutions to string theory constructions and the hierarchy problem, constitute one such class of BSM models [175]. With large 't Hooft couplings in contrast to QCD, Hidden Valleys can yield uniform radiation patterns and are a good signature to compare against SM backgrounds such as QCD events. This event-level classification task is the subject of ongoing work, as is the jet-level task distinguishing two-pronged boosted W boson jets from QCD background jets.

In what follows, we take SM decays of $t\bar{t}$ pairs to be our signal events, with our background coming from QCD dijets. A total of 10k events are simulated, about half of which are signal events.

Proton-proton collision events at $\sqrt{s} = 14$ TeV are generated in MADGRAPH5 2.9.6 [176], with top quarks being pair produced via $pp \rightarrow t\bar{t}$ and QCD dijet events generated through $pp \rightarrow jj$, where j represents a quark (u, d, c, s) or gluon. The top (or anti-top) quarks subsequently decay via $t \rightarrow W^+q$ (or $\bar{t} \rightarrow W^-\bar{q}$) where $q = b, s, d$ and \bar{q}

represents the corresponding anti-quarks. We restrict W boson decays to hadronic ones with $W \rightarrow jj$. Thus, top events in principle have six jets.

Particles are then hadronized and further decayed in PYTHIA 8.303 [177], where default tuning and showering parameters are used. Only final state particles with pseudo-rapidity $|y| < 2$ in the lab frame are kept. We select events whose scalar transverse momentum is $\sum p_T > 400$ GeV and cluster them into jets using FASTJET 3.4.0 [178].

Event pre-processing

To construct the phase space manifold, we first need all events to be represented by a fixed number of N constituents. This is accomplished by clustering each event into a certain number of jets. Here we use the k_T jet clustering algorithm with $R = 0.6$ to exclusively cluster each event into the $N = 3, 4, 5, 10, 15, 20, 25, 30$ hardest jets, in order to examine the effect of the number of jets used to represent one event. For comparison, we also test inclusive clustering with $N = 2, 3, 4, 5$ jets. Only events with the indicated number of jets are kept. Each event is then represented by N jets, which we term the N -jet representation, where a jet is treated as a pseudo-particle with its kinematic information given by the jet axis and its mass manually set to zero in order to comply with the requirements of the phase space formalism. Finally, we boost each event (now represented by N massless jets) to its CM frame.

Of course, in the above process, kinematic information will be distorted and lost to some degree. Ideally, the phase space formalism should be generalized to incorporate cases with massive constituent particles. However, the added complexity of doing so makes the prospect of finding a closed-form solution for the phase space coordinates unlikely. We leave the extension to massive particles as a future project and content ourselves with the massless approximation for the time being.

Figure 1 visualizes a top pair production event and a QCD dijet event randomly

selected from our 10k dataset. The events with all the constituent particles are shown first, followed by the N -jet representation on the $y - \phi$ plane with $N = 2, 3, 4, 5, 10, 20, 30$.




Figure 1: Visualization in the $y - \phi$ plane of a $t\bar{t}$ event (red) and QCD dijet event (green), represented by all the final state particles (leftmost plot) or N exclusive jets ($N = 3, 4, 5, 10, 15, 20, 25, 30$). The size of each dot is proportional to the p_T of the constituents.

Results

Figures 2 to 5 show the AUC scores, indicating the classification performance of phase space coordinates coupled with either kNN or SVM, using N -jet representations under various approximation schemes for particle permutations. We test two different simplex

coordinate definitions, given by $\rho \equiv E \pm p_z$.

From these figures, we wish to point out several key observations. First, our results are quite robust to different simplex coordinate definitions and machine learning algorithms. Using SVM does give better results than kNN (particularly for $N \geq 10$), which is to be expected as it is a more sophisticated algorithm, but the improvement is modest for smaller N values.

Second, we see that the simplex and sphere encode complementary information, as the AUC values for their combination are generally higher than for either manifold component alone. This synergy is also robust to different weightings between the sphere and simplex.

Third, the overall classification performance is quite good for $N = 3, 4, 5$. Clustering into $N = 2$ jets encodes almost no information and the performance deteriorates after the threshold $N_{\text{opt}} \sim 4 - 5$ is passed. The optimal performance AUC is quite remarkable, given that (1) the PS approach is entirely grounded in phase space itself, which does not incorporate any dynamics of the underlying collision process; (2) the particle-level details are completely discarded and an event is only approximated by a few jets which are artificially set to be massless; and (3) the computational cost for obtaining the PS coordinates is relatively cheap.

As a comparison, [179] introduced event isotropy as a new event shape observable. Roughly speaking, it quantifies how similar a given radiation pattern of an event is to a uniform N -body phase space configuration. Our phase space approach gives moderately better results than event isotropy on the $t\bar{t}$ vs. QCD dijet event tagging task, with the caution that the underlying datasets of the two studies are similar but not the same, and that our purpose and focus are different from theirs.

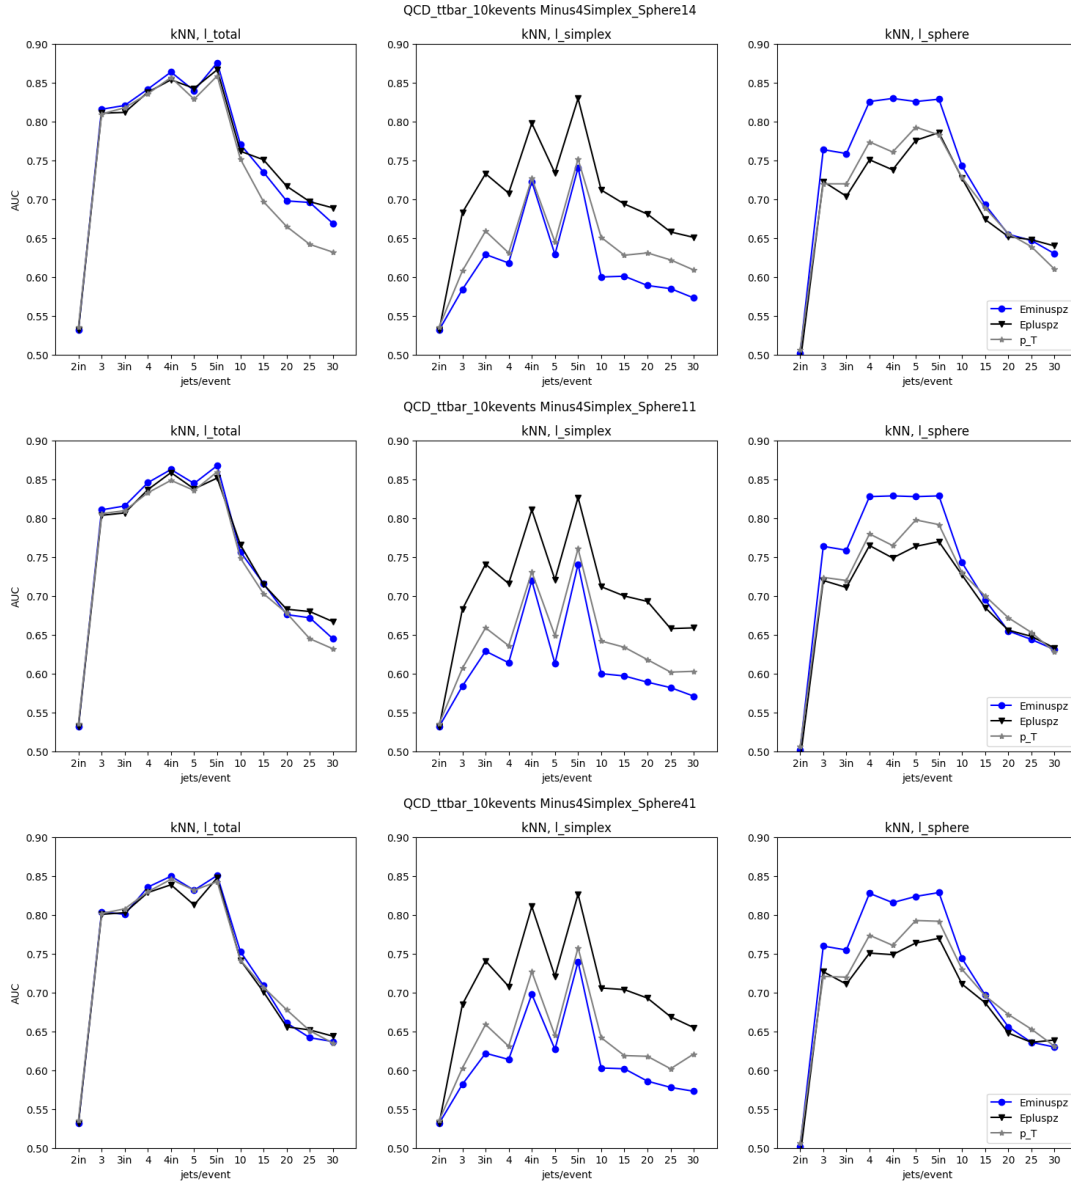


Figure 2: AUC values for kNN classification of QCD dijet events vs. $t\bar{t}$ events. The simplex coordinate is chosen to be $\rho \equiv E - p_z$. On the horizontal axis of each plot, we vary the number of jets N used to represent each event. Exclusive clustering is used unless otherwise noted (e.g. “2in”). The blue, black, and gray curves represent results with particles in a jet ordered by their $E - p_z$, $E + p_z$, and p_T values, respectively. The top, middle, and bottom panels differ in the relative sphere-to-simplex weighting in the metric (see eq. (4.18)), with $c = 0.25, 1, 4$, respectively. The left, middle, and right columns show results for the machine learning classifier applied to the total, simplex, and sphere distance matrices, respectively.

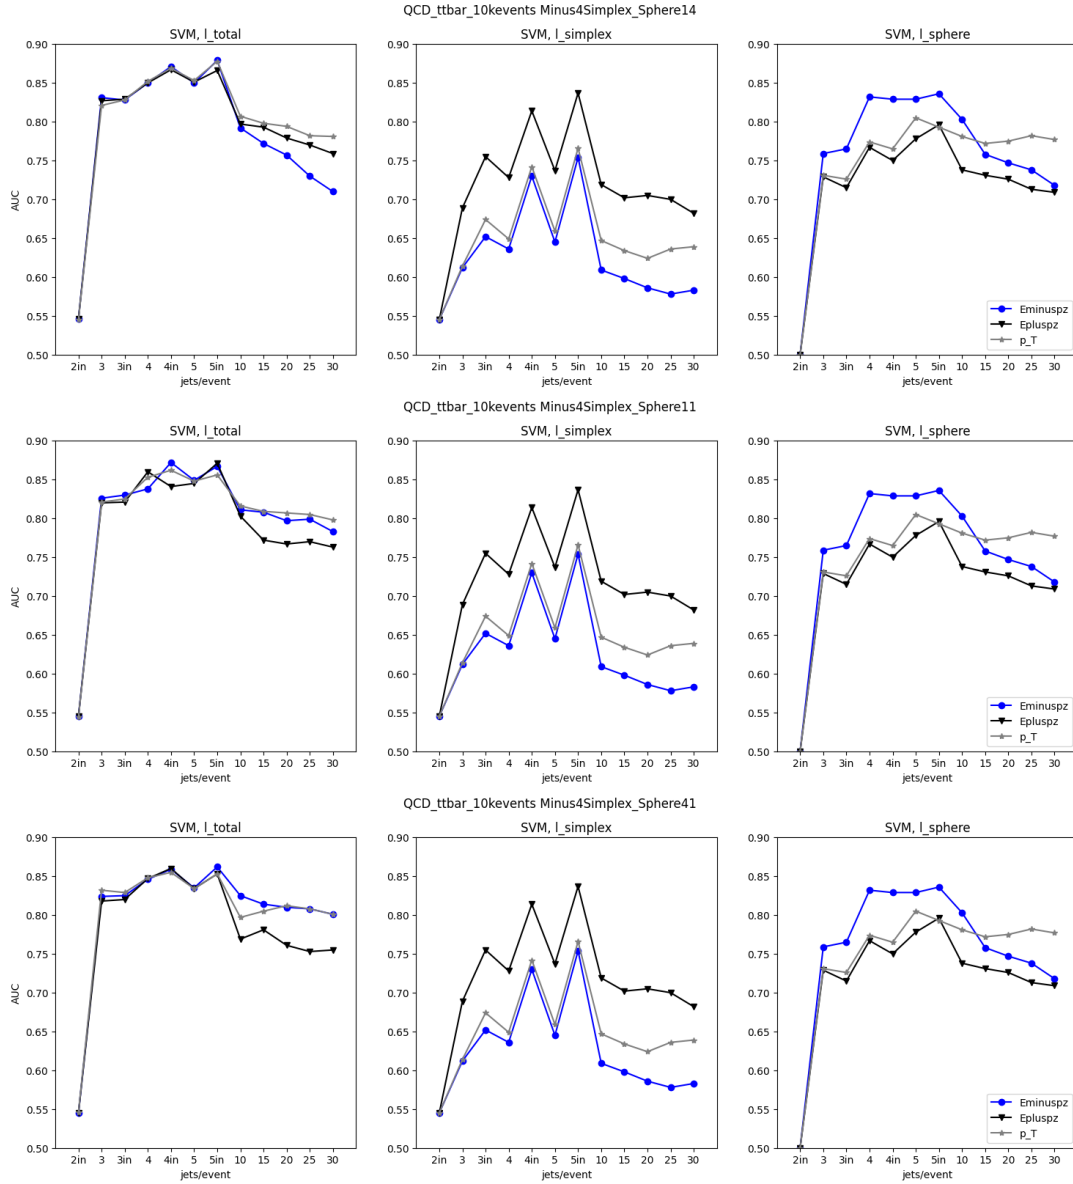


Figure 3: AUC values for SVM classification of QCD dijet events vs. $t\bar{t}$ events. The simplex coordinate is chosen to be $\rho \equiv E - p_z$. Remaining details are as in fig. 2.

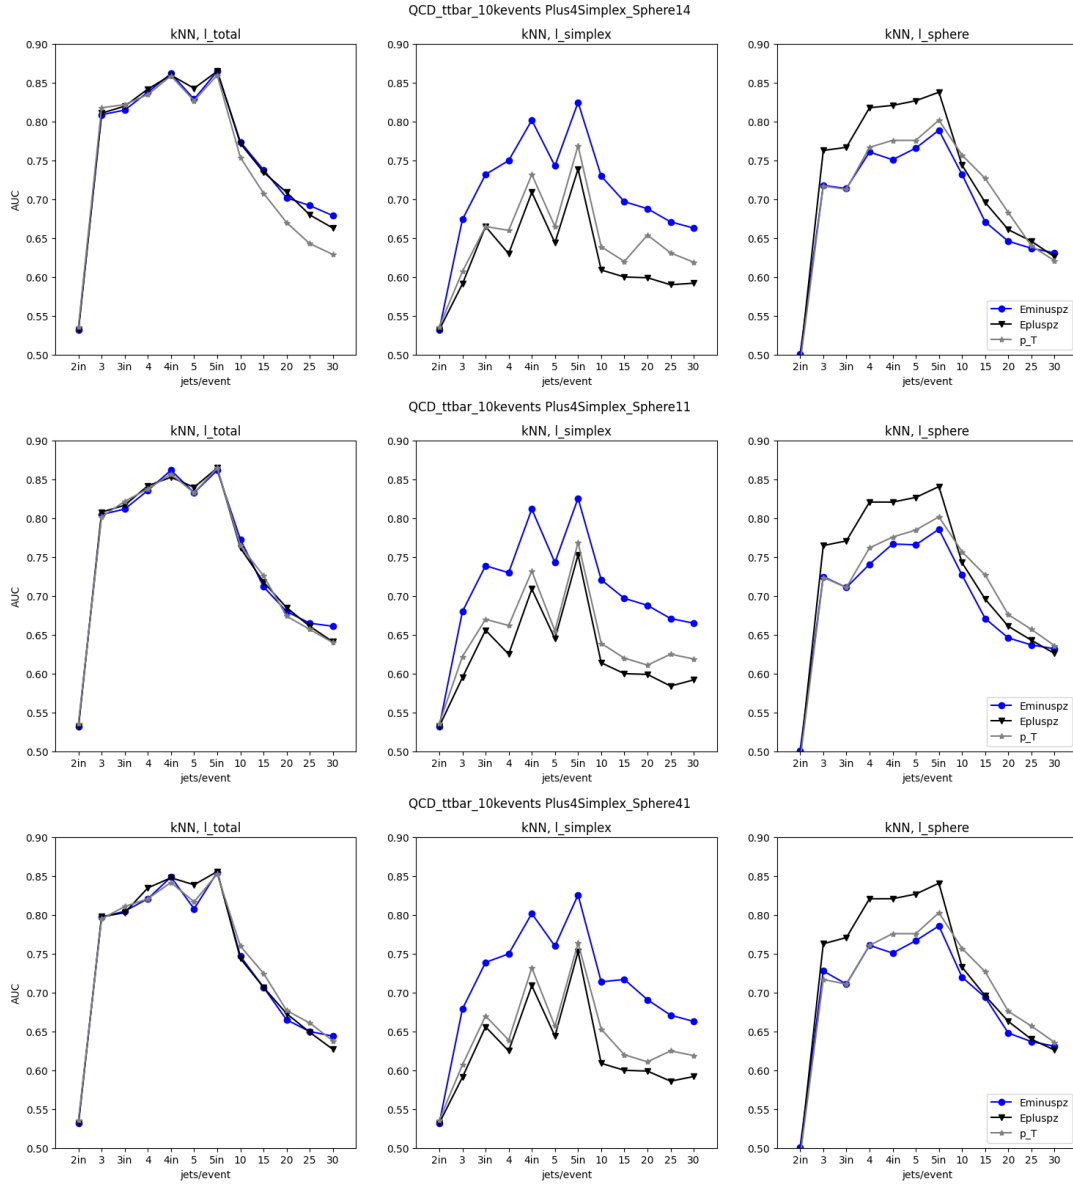


Figure 4: AUC values for kNN classification of QCD dijet events vs. $t\bar{t}$ events. The simplex coordinate is chosen to be $\rho \equiv E + p_z$. Remaining details are as in fig. 2.

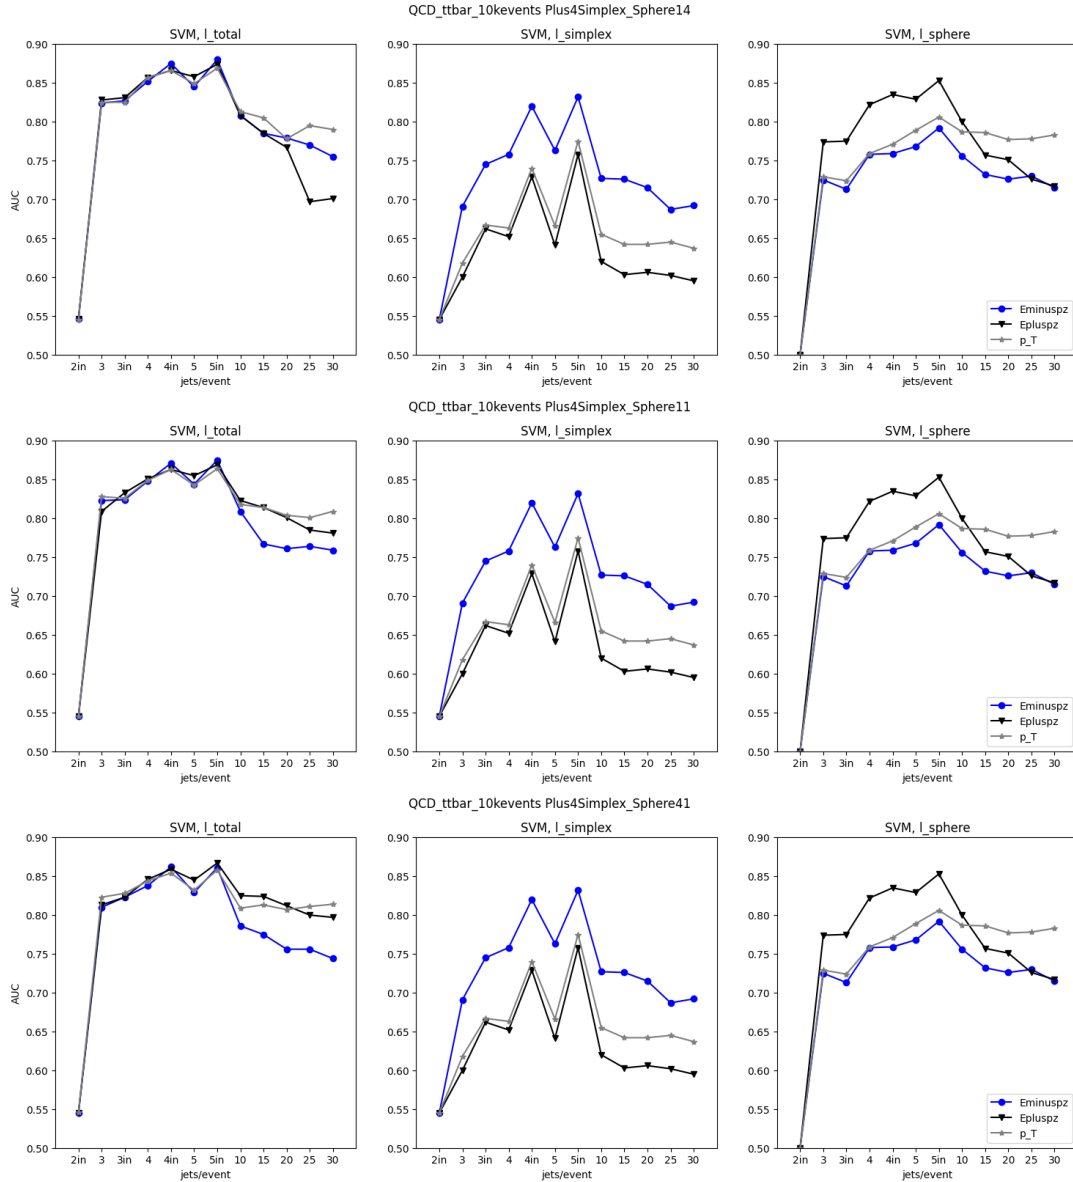


Figure 5: AUC values for SVM classification of QCD dijet events vs. $t\bar{t}$ events. The simplex coordinate is chosen to be $\rho \equiv E + p_z$. Remaining details are as in fig. 2.

4.5 Conclusions

The phase space manifold furnishes a natural and powerful notion of the distance between collider events. In this paper, we have developed a practical prescription for

defining and computing phase space distances at hadron colliders, both for boosted objects and entire events. These distances illuminate the distinctive phase space structures of different scattering amplitudes and lead to competitive event classification when coupled to simple machine learning algorithms.

There are a number of promising directions for future study. In converting explicit global coordinates on the phase space manifold to a practical metric, we have made a number of choices that leave room for exploration. For example, we have fixed the dimensionality of the phase space manifold for a given ensemble of events by clustering into N jets, mapping events onto an N -jet representation that influences the effectiveness of ML-based event classification. Alternative approaches to treating events of different intrinsic dimensionality may prove fruitful. Furthermore, the phase space distances computed between these N -jet representations assume massless phase space, which is not guaranteed to be a good approximation in all regimes; this could be circumvented by generalizing the phase space metric to massive particles. Finally, we forego an exact (and computationally prohibitive) minimization of the distance over particle permutations in favor of an approximate procedure that is only guaranteed to find local minima. Alternative approaches to the minimization problem (potentially using machine learning) may ultimately prove effective.

Further opportunities also abound within our prescription for computing phase space distances. Analytic and numerical study of the phase space structure of different scattering amplitudes is likely to be productive. And while we have demonstrated the value of phase space distances in classification tasks, they are equally likely to be useful in other collider physics applications such as unsupervised anomaly detection.

Long a backdrop to the calculation of scattering amplitudes, the geometry of phase space is beyond ready to step into the limelight.

Appendices

4.A Proof of minimal ordering for Euclidean metric

In this appendix, we present a simple proof regarding the permutation-invariant Euclidean metric.

Lemma 1. *The permutation-invariant Euclidean metric distance between two sets of points $\{x_i\}_{i=1}^N, \{y_i\}_{i=1}^N$ can be expressed as*

$$d^2(\{x_i\}_{i=1}^N, \{y_i\}_{i=1}^N) = \min_{\sigma \in S_N} \sum_{i=1}^N (x_i - y_{\sigma(i)})^2 = \sum_{\substack{i=1 \\ x_i > x_{i+1} \\ y_i > y_{i+1}}}^N (x_i - y_i)^2. \quad (4.37)$$

That is, the ordering of points that minimizes their distance is when they are in descending order.

Proof. We will assume that all points are distinct, i.e., $x_i = x_j$ implies $i = j$. With probability 1, all points are distinct in our application of this theorem to the study of phase space because the number of points is finite and drawn from a continuous probability distribution. The case with degeneracies can be considered as well, but the main result does not change.

Without loss of generality, we can assume that the x_i points are ordered, $x_1 > x_2 > \dots > x_N$, while the y_i points are in random order. Expanding out the square, the metric

can be expressed as

$$\min_{\sigma \in S_N} \sum_{i=1}^N (x_i - y_{\sigma(i)})^2 = \sum_{i=1}^N (x_i^2 + y_i^2) - 2 \max_{\sigma \in S_N} \sum_{i=1}^N x_i y_{\sigma(i)}. \quad (4.38)$$

We only need to consider the final term on the right for the minimization of the metric. Let's consider two terms in that sum and determine if permuting them increases its value. That is, we will consider $x_i > x_j$ for $j > i$ and focus on the terms

$$\sum_{i=1}^N x_i y_{\sigma(i)} \supset x_i y_{\sigma(i)} + x_j y_{\sigma(j)}. \quad (4.39)$$

Now, there are obviously two cases: either $y_{\sigma(i)} > y_{\sigma(j)}$ or $y_{\sigma(j)} > y_{\sigma(i)}$. Further, we can assume, without loss of generality, that $\sigma(j) > \sigma(i)$. First, if $y_{\sigma(i)} > y_{\sigma(j)}$, then note that

$$(x_i - x_j)(y_{\sigma(i)} - y_{\sigma(j)}) > 0, \quad (4.40)$$

or that

$$x_i y_{\sigma(i)} + x_j y_{\sigma(j)} > x_i y_{\sigma(j)} + x_j y_{\sigma(i)}. \quad (4.41)$$

That is, the points are already ordered to minimize the metric, with $i < j$ and $\sigma(i) < \sigma(j)$.

In the other case, where $y_{\sigma(j)} > y_{\sigma(i)}$, we have

$$(x_i - x_j)(y_{\sigma(j)} - y_{\sigma(i)}) > 0, \quad (4.42)$$

or that

$$x_i y_{\sigma(j)} + x_j y_{\sigma(i)} > x_i y_{\sigma(i)} + x_j y_{\sigma(j)}. \quad (4.43)$$

In this case, we can swap the indices on the y points, $\sigma(i) \leftrightarrow \sigma(j)$, and after this swap, the metric is minimized with the index ordering $i < j$ and $\sigma(i) < \sigma(j)$.

We can perform this pairwise comparison of points and their contribution to the metric, swapping indices on the y points as necessary, until no swap increases the value of the sum in eq. (4.39). After this procedure, we have therefore reordered all y points such that $y_i > y_j$ with $i < j$, for all $1 \leq i < j \leq N$. That is, all y points have been sorted in descending order, just like the x points. Further, note that this sorting is the global minimum because any other ordering necessarily increases the metric distance. This then proves the equality of eq. (4.37). \square

Bibliography

- [1] I. Garcia Garcia, G. Koszegi, and R. Petrossian-Byrne, *Reflections on Bubble Walls*, arXiv:2212.1057.
- [2] H. Al Ali *et. al.*, *The muon Smasher's guide*, *Rept. Prog. Phys.* **85** (2022), no. 8 084201, [arXiv:2103.1404].
- [3] N. Craig, I. Garcia Garcia, G. Koszegi, and A. McCune, *P not PQ*, *JHEP* **09** (2021) 130, [arXiv:2012.1341].
- [4] B. J. Harder, B. Good, M. Schmitt, B. Kowalski, G. Koszegi, M. T. Johnson, and K. T. Faber, *Deposition of electrically conductive zirconium monoxide via plasma spray-physical vapor deposition*, *Journal of the American Ceramic Society* **105** (2022), no. 5 3568–3580, [<https://ceramics.onlinelibrary.wiley.com/doi/pdf/10.1111/jace.18309>].
- [5] C. Caprini *et. al.*, *Science with the space-based interferometer eLISA. II: Gravitational waves from cosmological phase transitions*, *JCAP* **04** (2016) 001, [arXiv:1512.0623].
- [6] C. Caprini *et. al.*, *Detecting gravitational waves from cosmological phase transitions with LISA: an update*, *JCAP* **03** (2020) 024, [arXiv:1910.1312].
- [7] A. J. Larkoski and T. Melia, *Covariantizing phase space*, *Phys. Rev. D* **102** (2020), no. 9 094014, [arXiv:2008.0650].
- [8] E. Witten, *Cosmic Separation of Phases*, *Phys. Rev. D* **30** (1984) 272–285.
- [9] C. J. Hogan, *Gravitational radiation from cosmological phase transitions*, *Mon. Not. Roy. Astron. Soc.* **218** (1986) 629–636.
- [10] A. Kosowsky, M. S. Turner, and R. Watkins, *Gravitational radiation from colliding vacuum bubbles*, *Phys. Rev. D* **45** (1992) 4514–4535.
- [11] A. Kosowsky, M. S. Turner, and R. Watkins, *Gravitational waves from first order cosmological phase transitions*, *Phys. Rev. Lett.* **69** (1992) 2026–2029.

- [12] M. Kamionkowski, A. Kosowsky, and M. S. Turner, *Gravitational radiation from first order phase transitions*, *Phys. Rev. D* **49** (1994) 2837–2851, [astro-ph/9310044].
- [13] P. Schwaller, *Gravitational Waves from a Dark Phase Transition*, *Phys. Rev. Lett.* **115** (2015), no. 18 181101, [arXiv:1504.0726].
- [14] M. Breitbach, J. Kopp, E. Madge, T. Opferkuch, and P. Schwaller, *Dark, Cold, and Noisy: Constraining Secluded Hidden Sectors with Gravitational Waves*, *JCAP* **07** (2019) 007, [arXiv:1811.1117].
- [15] A. J. Helmboldt, J. Kubo, and S. van der Woude, *Observational prospects for gravitational waves from hidden or dark chiral phase transitions*, *Phys. Rev. D* **100** (2019), no. 5 055025, [arXiv:1904.0789].
- [16] R. Jinno, H. Seong, M. Takimoto, and C. M. Um, *Gravitational waves from first-order phase transitions: Ultra-supercooled transitions and the fate of relativistic shocks*, *JCAP* **10** (2019) 033, [arXiv:1905.0089].
- [17] A. Azatov, D. Barducci, and F. Sgarlata, *Gravitational traces of broken gauge symmetries*, *JCAP* **07** (2020) 027, [arXiv:1910.0112].
- [18] W.-C. Huang, M. Reichert, F. Sannino, and Z.-W. Wang, *Testing the dark $SU(N)$ Yang-Mills theory confined landscape: From the lattice to gravitational waves*, *Phys. Rev. D* **104** (2021), no. 3 035005, [arXiv:2012.1161].
- [19] Y. Nakai, M. Suzuki, F. Takahashi, and M. Yamada, *Gravitational Waves and Dark Radiation from Dark Phase Transition: Connecting NANOGraV Pulsar Timing Data and Hubble Tension*, *Phys. Lett. B* **816** (2021) 136238, [arXiv:2009.0975].
- [20] J. Halverson, C. Long, A. Maiti, B. Nelson, and G. Salinas, *Gravitational waves from dark Yang-Mills sectors*, *JHEP* **05** (2021) 154, [arXiv:2012.0407].
- [21] M. Reichert, F. Sannino, Z.-W. Wang, and C. Zhang, *Dark confinement and chiral phase transitions: gravitational waves vs matter representations*, *JHEP* **01** (2022) 003, [arXiv:2109.1155].
- [22] F. Huang, V. Sanz, J. Shu, and X. Xue, *LIGO as a probe of dark sectors*, *Phys. Rev. D* **104** (2021), no. 9 095001, [arXiv:2102.0315].
- [23] F. Ertas, F. Kahlhoefer, and C. Tasillo, *Turn up the volume: listening to phase transitions in hot dark sectors*, *JCAP* **02** (2022), no. 02 014, [arXiv:2109.0620].
- [24] J. B. Dent, B. Dutta, S. Ghosh, J. Kumar, and J. Runburg, *Sensitivity to dark sector scales from gravitational wave signatures*, *JHEP* **08** (2022) 300, [arXiv:2203.1173].

- [25] M. Fairbairn, E. Hardy, and A. Wickens, *Hearing without seeing: gravitational waves from hot and cold hidden sectors*, *JHEP* **07** (2019) 044, [arXiv:1901.1103].
- [26] R. Jinno, B. Shakya, and J. van de Vis, *Gravitational Waves from Feebly Interacting Particles in a First Order Phase Transition*, arXiv:2211.0640.
- [27] S. R. Coleman, *The Fate of the False Vacuum. 1. Semiclassical Theory*, *Phys. Rev. D* **15** (1977) 2929–2936. [Erratum: *Phys.Rev.D* 16, 1248 (1977)].
- [28] S. J. Huber and T. Konstandin, *Gravitational Wave Production by Collisions: More Bubbles*, *JCAP* **09** (2008) 022, [arXiv:0806.1828].
- [29] J. R. Espinosa, T. Konstandin, J. M. No, and G. Servant, *Energy Budget of Cosmological First-order Phase Transitions*, *JCAP* **06** (2010) 028, [arXiv:1004.4187].
- [30] M. Hindmarsh, S. J. Huber, K. Rummukainen, and D. J. Weir, *Numerical simulations of acoustically generated gravitational waves at a first order phase transition*, *Phys. Rev. D* **92** (2015), no. 12 123009, [arXiv:1504.0329].
- [31] M. Hindmarsh, *Sound shell model for acoustic gravitational wave production at a first-order phase transition in the early Universe*, *Phys. Rev. Lett.* **120** (2018), no. 7 071301, [arXiv:1608.0473].
- [32] M. Hindmarsh, S. J. Huber, K. Rummukainen, and D. J. Weir, *Shape of the acoustic gravitational wave power spectrum from a first order phase transition*, *Phys. Rev. D* **96** (2017), no. 10 103520, [arXiv:1704.0587]. [Erratum: *Phys.Rev.D* 101, 089902 (2020)].
- [33] M. Hindmarsh and M. Hijazi, *Gravitational waves from first order cosmological phase transitions in the Sound Shell Model*, *JCAP* **12** (2019) 062, [arXiv:1909.1004].
- [34] C. Caprini, R. Durrer, T. Konstandin, and G. Servant, *General Properties of the Gravitational Wave Spectrum from Phase Transitions*, *Phys. Rev. D* **79** (2009) 083519, [arXiv:0901.1661].
- [35] G. Barenboim and W.-I. Park, *Gravitational waves from first order phase transitions as a probe of an early matter domination era and its inverse problem*, *Phys. Lett. B* **759** (2016) 430–438, [arXiv:1605.0378].
- [36] A. Hook, G. Marques-Tavares, and D. Racco, *Causal gravitational waves as a probe of free streaming particles and the expansion of the Universe*, *JHEP* **02** (2021) 117, [arXiv:2010.0356].
- [37] Y. Gouttenoire, G. Servant, and P. Simakachorn, *Kination cosmology from scalar fields and gravitational-wave signatures*, arXiv:2111.0115.

- [38] S. Weinberg, *Damping of tensor modes in cosmology*, *Phys. Rev. D* **69** (2004) 023503, [astro-ph/0306304].
- [39] D. E. Morrissey and M. J. Ramsey-Musolf, *Electroweak baryogenesis*, *New J. Phys.* **14** (2012) 125003, [arXiv:1206.2942].
- [40] **NANOGrav** Collaboration, Z. Arzoumanian *et. al.*, *The NANOGrav 11-year Data Set: Pulsar-timing Constraints On The Stochastic Gravitational-wave Background*, *Astrophys. J.* **859** (2018), no. 1 47, [arXiv:1801.0261].
- [41] L. Lentati *et. al.*, *European Pulsar Timing Array Limits On An Isotropic Stochastic Gravitational-Wave Background*, *Mon. Not. Roy. Astron. Soc.* **453** (2015), no. 3 2576–2598, [arXiv:1504.0369].
- [42] P. B. Arnold, *One loop fluctuation - dissipation formula for bubble wall velocity*, *Phys. Rev. D* **48** (1993) 1539–1545, [hep-ph/9302258].
- [43] G. D. Moore and T. Prokopec, *How fast can the wall move? A Study of the electroweak phase transition dynamics*, *Phys. Rev. D* **52** (1995) 7182–7204, [hep-ph/9506475].
- [44] N. Turok, *Electroweak bubbles: Nucleation and growth*, *Phys. Rev. Lett.* **68** (1992) 1803–1806.
- [45] M. Dine, R. G. Leigh, P. Y. Huet, A. D. Linde, and D. A. Linde, *Towards the theory of the electroweak phase transition*, *Phys. Rev. D* **46** (1992) 550–571, [hep-ph/9203203].
- [46] B.-H. Liu, L. D. McLerran, and N. Turok, *Bubble nucleation and growth at a baryon number producing electroweak phase transition*, *Phys. Rev. D* **46** (1992) 2668–2688.
- [47] G. D. Moore and T. Prokopec, *Bubble wall velocity in a first order electroweak phase transition*, *Phys. Rev. Lett.* **75** (1995) 777–780, [hep-ph/9503296].
- [48] D. Bodeker and G. D. Moore, *Can electroweak bubble walls run away?*, *JCAP* **05** (2009) 009, [arXiv:0903.4099].
- [49] L. Leitaó and A. Megevand, *Hydrodynamics of ultra-relativistic bubble walls*, *Nucl. Phys. B* **905** (2016) 45–72, [arXiv:1510.0774].
- [50] D. Bodeker and G. D. Moore, *Electroweak Bubble Wall Speed Limit*, *JCAP* **05** (2017) 025, [arXiv:1703.0821].
- [51] A. Azatov and M. Vanvlasselaer, *Bubble wall velocity: heavy physics effects*, *JCAP* **01** (2021) 058, [arXiv:2010.0259].

- [52] W.-Y. Ai, B. Garbrecht, and C. Tamarit, *Bubble wall velocities in local equilibrium*, *JCAP* **03** (2022), no. 03 015, [arXiv:2109.1371].
- [53] M. B. Mancha, T. Prokopec, and B. Świeżewska, *Field-theoretic derivation of bubble-wall force*, *Journal of High Energy Physics* **2021** (jan, 2021).
- [54] F. Bigazzi, A. Caddeo, T. Canneti, and A. L. Cotrone, *Bubble wall velocity at strong coupling*, *JHEP* **08** (2021) 090, [arXiv:2104.1281].
- [55] Y. Bea, J. Casalderrey-Solana, T. Giannakopoulos, D. Mateos, M. Sanchez-Garitaonandia, and M. Zilhão, *Bubble wall velocity from holography*, *Phys. Rev. D* **104** (2021), no. 12 L121903, [arXiv:2104.0570].
- [56] Y. Gouttenoire, R. Jinno, and F. Sala, *Friction pressure on relativistic bubble walls*, *JHEP* **05** (2022) 004, [arXiv:2112.0768].
- [57] A. Azatov, G. Barni, S. Chakraborty, M. Vanvlasselaer, and W. Yin, *Ultra-relativistic bubbles from the simplest Higgs portal and their cosmological consequences*, *JHEP* **10** (2022) 017, [arXiv:2207.0223].
- [58] B. Laurent and J. M. Cline, *First principles determination of bubble wall velocity*, *Phys. Rev. D* **106** (2022), no. 2 023501, [arXiv:2204.1312].
- [59] S. De Curtis, L. D. Rose, A. Guiggiani, A. G. Muyor, and G. Panico, *Bubble wall dynamics at the electroweak phase transition*, *JHEP* **03** (2022) 163, [arXiv:2201.0822].
- [60] S.-J. Wang and Z.-Y. Yuwen, *Hydrodynamic backreaction force of cosmological bubble expansion*, *Phys. Rev. D* **107** (2023), no. 2 023501, [arXiv:2205.0249].
- [61] W. Hu, R. Barkana, and A. Gruzinov, *Cold and fuzzy dark matter*, *Phys. Rev. Lett.* **85** (2000) 1158–1161, [astro-ph/0003365].
- [62] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, *Ultralight scalars as cosmological dark matter*, *Phys. Rev. D* **95** (2017), no. 4 043541, [arXiv:1610.0829].
- [63] A. E. Nelson and J. Scholtz, *Dark Light, Dark Matter and the Misalignment Mechanism*, *Phys. Rev. D* **84** (2011) 103501, [arXiv:1105.2812].
- [64] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, *WISPy Cold Dark Matter*, *JCAP* **06** (2012) 013, [arXiv:1201.5902].
- [65] P. W. Graham, J. Mardon, and S. Rajendran, *Vector Dark Matter from Inflationary Fluctuations*, *Phys. Rev. D* **93** (2016), no. 10 103520, [arXiv:1504.0210].

- [66] P. Agrawal, N. Kitajima, M. Reece, T. Sekiguchi, and F. Takahashi, *Relic Abundance of Dark Photon Dark Matter*, *Phys. Lett. B* **801** (2020) 135136, [arXiv:1810.0718].
- [67] I. Baldes, Y. Gouttenoire, and F. Sala, *String Fragmentation in Supercooled Confinement and Implications for Dark Matter*, *JHEP* **04** (2021) 278, [arXiv:2007.0844].
- [68] I. Baldes, Y. Gouttenoire, and F. Sala, *Hot and heavy dark matter from a weak scale phase transition*, *SciPost Phys.* **14** (2023) 033, [arXiv:2207.0509].
- [69] O. Lebedev, H. M. Lee, and Y. Mambrini, *Vector Higgs-portal dark matter and the invisible Higgs*, *Phys. Lett. B* **707** (2012) 570–576, [arXiv:1111.4482].
- [70] G. R. Farrar and J. W. McIntosh, Jr., *Scattering from a domain wall in a spontaneously broken gauge theory*, *Phys. Rev. D* **51** (1995) 5889–5904, [hep-ph/9412270].
- [71] K. N. Abazajian, *et al.*, *CMB-S4 Science Book, First Edition*, 2016.
- [72] **Planck** Collaboration, N. Aghanim *et. al.*, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6, [arXiv:1807.0620].
- [73] **Simons Observatory** Collaboration, P. Ade *et. al.*, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056, [arXiv:1808.0744].
- [74] **nEDM** Collaboration, C. Abel *et. al.*, *Measurement of the permanent electric dipole moment of the neutron*, *Phys. Rev. Lett.* **124** (2020), no. 8 081803, [arXiv:2001.1196].
- [75] B. Graner, Y. Chen, E. Lindahl, and B. Heckel, *Reduced Limit on the Permanent Electric Dipole Moment of Hg199*, *Phys. Rev. Lett.* **116** (2016), no. 16 161601, [arXiv:1601.0433]. [Erratum: *Phys.Rev.Lett.* 119, 119901 (2017)].
- [76] N. Kaloper and J. Terning, *Landscaping the Strong CP Problem*, *JHEP* **03** (2019) 032, [arXiv:1710.0174].
- [77] K. Babu and R. N. Mohapatra, *CP Violation in Seesaw Models of Quark Masses*, *Phys. Rev. Lett.* **62** (1989) 1079.
- [78] K. Babu and R. N. Mohapatra, *A Solution to the Strong CP Problem Without an Axion*, *Phys. Rev. D* **41** (1990) 1286.
- [79] S. M. Barr, D. Chang, and G. Senjanovic, *Strong CP problem and parity*, *Phys. Rev. Lett.* **67** (1991) 2765–2768.

- [80] R. N. Mohapatra and G. Senjanovic, *Natural Suppression of Strong p and t Noninvariance*, *Phys. Lett. B* **79** (1978) 283–286.
- [81] A. E. Nelson, *Naturally Weak CP Violation*, *Phys. Lett. B* **136** (1984) 387–391.
- [82] S. M. Barr, *Solving the Strong CP Problem Without the Peccei-Quinn Symmetry*, *Phys. Rev. Lett.* **53** (1984) 329.
- [83] S. Chakdar, K. Ghosh, S. Nandi, and S. K. Rai, *Collider signatures of mirror fermions in the framework of a left-right mirror model*, *Phys. Rev. D* **88** (2013), no. 9 095005, [arXiv:1305.2641].
- [84] R. T. D’Agnolo and A. Hook, *Finding the Strong CP problem at the LHC*, *Phys. Lett. B* **762** (2016) 421–425, [arXiv:1507.0033].
- [85] L. J. Hall and K. Harigaya, *Implications of Higgs Discovery for the Strong CP Problem and Unification*, *JHEP* **10** (2018) 130, [arXiv:1803.0811].
- [86] M. Dine, R. G. Leigh, and D. A. MacIntire, *Of CP and other gauge symmetries in string theory*, *Phys. Rev. Lett.* **69** (1992) 2030–2032, [hep-th/9205011].
- [87] K.-w. Choi, D. B. Kaplan, and A. E. Nelson, *Is CP a gauge symmetry?*, *Nucl. Phys. B* **391** (1993) 515–530, [hep-ph/9205202].
- [88] **ATLAS** Collaboration, M. Aaboud *et. al.*, *Combination of the searches for pair-produced vector-like partners of the third-generation quarks at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. Lett.* **121** (2018), no. 21 211801, [arXiv:1808.0234].
- [89] **CMS** Collaboration, A. M. Sirunyan *et. al.*, *Search for vector-like T and B quark pairs in final states with leptons at $\sqrt{s} = 13$ TeV*, *JHEP* **08** (2018) 177, [arXiv:1805.0475].
- [90] **ATLAS** Collaboration, G. Aad *et. al.*, *Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Lett. B* **796** (2019) 68–87, [arXiv:1903.0624].
- [91] **ATLAS** Collaboration, G. Aad *et. al.*, *Search for a heavy charged boson in events with a charged lepton and missing transverse momentum from pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **100** (2019), no. 5 052013, [arXiv:1906.0560].
- [92] C. Helsens and M. Selvaggi, *Search for high-mass resonances at FCC-hh*, Tech. Rep. CERN-ACC-2019-0028, CERN, Geneva, Oct, 2018.

- [93] FCC Collaboration, A. Abada *et. al.*, *FCC Physics Opportunities: Future Circular Collider Conceptual Design Report Volume 1*, *Eur. Phys. J. C* **79** (2019), no. 6 474.
- [94] Y. Zeldovich, *A New Type of Radioactive Decay: Gravitational Annihilation of Baryons*, *Phys. Lett. A* **59** (1976) 254.
- [95] Y. Zeldovich, *A Novel Type of Radioactive Decay: Gravitational Baryon Annihilation*, *Zh. Eksp. Teor. Fiz.* **72** (1977) 18–21.
- [96] T. Banks and L. J. Dixon, *Constraints on String Vacua with Space-Time Supersymmetry*, *Nucl. Phys. B* **307** (1988) 93–108.
- [97] S. B. Giddings and A. Strominger, *Axion Induced Topology Change in Quantum Gravity and String Theory*, *Nucl. Phys. B* **306** (1988) 890–907.
- [98] K.-M. Lee, *Wormholes and Goldstone Bosons*, *Phys. Rev. Lett.* **61** (1988) 263–266.
- [99] L. Abbott and M. B. Wise, *Wormholes and Global Symmetries*, *Nucl. Phys. B* **325** (1989) 687–704.
- [100] S. R. Coleman and K.-M. Lee, *WORMHOLES MADE WITHOUT MASSLESS MATTER FIELDS*, *Nucl. Phys. B* **329** (1990) 387–409.
- [101] R. Kallosh, A. D. Linde, D. A. Linde, and L. Susskind, *Gravity and global symmetries*, *Phys. Rev. D* **52** (1995) 912–935, [hep-th/9502069].
- [102] T. Banks and N. Seiberg, *Symmetries and Strings in Field Theory and Gravity*, *Phys. Rev. D* **83** (2011) 084019, [arXiv:1011.5120].
- [103] D. Harlow and H. Ooguri, *Symmetries in quantum field theory and quantum gravity*, arXiv:1810.0533.
- [104] D. Harlow and H. Ooguri, *Constraints on Symmetries from Holography*, *Phys. Rev. Lett.* **122** (2019), no. 19 191601, [arXiv:1810.0533].
- [105] S. Fichtel and P. Saraswat, *Approximate Symmetries and Gravity*, *JHEP* **01** (2020) 088, [arXiv:1909.0200].
- [106] T. Daus, A. Hebecker, S. Leonhardt, and J. March-Russell, *Towards a Swampland Global Symmetry Conjecture using weak gravity*, *Nucl. Phys. B* **960** (2020) 115167, [arXiv:2002.0245].
- [107] R. Peccei and H. R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.

- [108] R. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Instantons*, *Phys. Rev. D* **16** (1977) 1791–1797.
- [109] F. Wilczek, *Problem of Strong P and T Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279–282.
- [110] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223–226.
- [111] J. E. Kim, *Weak Interaction Singlet and Strong CP Invariance*, *Phys. Rev. Lett.* **43** (1979) 103.
- [112] M. A. Shifman, A. Vainshtein, and V. I. Zakharov, *Can Confinement Ensure Natural CP Invariance of Strong Interactions?*, *Nucl. Phys. B* **166** (1980) 493–506.
- [113] M. Dine, W. Fischler, and M. Srednicki, *A Simple Solution to the Strong CP Problem with a Harmless Axion*, *Phys. Lett. B* **104** (1981) 199–202.
- [114] S. M. Barr and D. Seckel, *Planck scale corrections to axion models*, *Phys. Rev. D* **46** (1992) 539–549.
- [115] M. Kamionkowski and J. March-Russell, *Planck scale physics and the Peccei-Quinn mechanism*, *Phys. Lett. B* **282** (1992) 137–141, [hep-th/9202003].
- [116] R. Holman, S. D. Hsu, T. W. Kephart, E. W. Kolb, R. Watkins, and L. M. Widrow, *Solutions to the strong CP problem in a world with gravity*, *Phys. Lett. B* **282** (1992) 132–136, [hep-ph/9203206].
- [117] S. Ghigna, M. Lusignoli, and M. Roncadelli, *Instability of the invisible axion*, *Phys. Lett. B* **283** (1992) 278–281.
- [118] E. Chun and A. Lukas, *Discrete gauge symmetries in axionic extensions of the SSM*, *Phys. Lett. B* **297** (1992) 298–304, [hep-ph/9209208].
- [119] L. Randall, *Composite axion models and Planck scale physics*, *Phys. Lett. B* **284** (1992) 77–80.
- [120] H.-C. Cheng and D. E. Kaplan, *Axions and a gauged Peccei-Quinn symmetry*, hep-ph/0103346.
- [121] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *String Axiverse*, *Phys. Rev. D* **81** (2010) 123530, [arXiv:0905.4720].
- [122] H. Fukuda, M. Ibe, M. Suzuki, and T. T. Yanagida, *A "gauged" U(1) Peccei-Quinn symmetry*, *Phys. Lett. B* **771** (2017) 327–331, [arXiv:1703.0111].
- [123] L. Di Luzio, E. Nardi, and L. Ubaldi, *Accidental Peccei-Quinn symmetry protected to arbitrary order*, *Phys. Rev. Lett.* **119** (2017), no. 1 011801, [arXiv:1704.0112].

- [124] B. Lillard and T. M. P. Tait, *A Composite Axion from a Supersymmetric Product Group*, *JHEP* **11** (2017) 005, [arXiv:1707.0426].
- [125] B. Lillard and T. M. Tait, *A High Quality Composite Axion*, *JHEP* **11** (2018) 199, [arXiv:1811.0308].
- [126] G. Senjanovic and R. N. Mohapatra, *Exact Left-Right Symmetry and Spontaneous Violation of Parity*, *Phys. Rev. D* **12** (1975) 1502.
- [127] J. R. Ellis and M. K. Gaillard, *Strong and Weak CP Violation*, *Nucl. Phys. B* **150** (1979) 141–162.
- [128] Z. Chacko, H.-S. Goh, and R. Harnik, *The Twin Higgs: Natural electroweak breaking from mirror symmetry*, *Phys. Rev. Lett.* **96** (2006) 231802, [hep-ph/0506256].
- [129] Z. Chacko, H.-S. Goh, and R. Harnik, *A Twin Higgs model from left-right symmetry*, *JHEP* **01** (2006) 108, [hep-ph/0512088].
- [130] A. Albaid, M. Dine, and P. Draper, *Strong CP and SUZ₂*, *JHEP* **12** (2015) 046, [arXiv:1510.0339].
- [131] A. Davidson and K. C. Wali, *Universal Seesaw Mechanism?*, *Phys. Rev. Lett.* **59** (1987) 393.
- [132] A. Davidson and K. C. Wali, *Family Mass Hierarchy From Universal Seesaw Mechanism*, *Phys. Rev. Lett.* **60** (1988) 1813.
- [133] A. Davidson, S. Ranfone, and K. C. Wali, *Quark Masses and Mixing Angles From Universal Seesaw Mechanism*, *Phys. Rev. D* **41** (1990) 208.
- [134] S. Ranfone, *The Three generation seesaw model for quarks. 2.*, *Phys. Rev. D* **42** (1990) 3819–3828.
- [135] H.-S. Goh and S. Su, *Phenomenology of The Left-Right Twin Higgs Model*, *Phys. Rev. D* **75** (2007) 075010, [hep-ph/0611015].
- [136] Y. Kiyo, T. Morozumi, P. Parada, M. Rebelo, and M. Tanimoto, *Quark mass hierarchy, FCNC and CP violation in a seesaw model*, *Prog. Theor. Phys.* **101** (1999) 671–706, [hep-ph/9809333].
- [137] A. Maiezza and M. Nemevšek, *Strong P invariance, neutron electric dipole moment, and minimal left-right parity at LHC*, *Phys. Rev. D* **90** (2014), no. 9 095002, [arXiv:1407.3678].
- [138] G. Senjanovic and V. Tello, *Strong CP violation: problem or blessing?*, arXiv:2004.0403.

- [139] A. Khodjamirian, T. Mannel, A. Pivovarov, and Y.-M. Wang, *Charm-loop effect in $B \rightarrow K^{(*)}\ell^+\ell^-$ and $B \rightarrow K^*\gamma$* , *JHEP* **09** (2010) 089, [arXiv:1006.4945].
- [140] **Particle Data Group** Collaboration, P. Zyla *et. al.*, *Review of Particle Physics*, *PTEP* **2020** (2020), no. 8 083C01.
- [141] N. Hutzler *et. al.*, *Searches for new sources of CP violation using molecules as quantum sensors*, arXiv:2010.0870.
- [142] **ACME** Collaboration, V. Andreev *et. al.*, *Improved limit on the electric dipole moment of the electron*, *Nature* **562** (2018), no. 7727 355–360.
- [143] Z. G. Berezhiani, R. N. Mohapatra, and G. Senjanovic, *Planck scale physics and solutions to the strong CP problem without axion*, *Phys. Rev. D* **47** (1993) 5565–5570, [hep-ph/9212318].
- [144] T. Kibble, *Topology of Cosmic Domains and Strings*, *J. Phys. A* **9** (1976) 1387–1398.
- [145] G. R. Dvali and G. Senjanovic, *Is there a domain wall problem?*, *Phys. Rev. Lett.* **74** (1995) 5178–5181, [hep-ph/9501387].
- [146] G. R. Dvali, A. Melfo, and G. Senjanovic, *Nonrestoration of spontaneously broken P and CP at high temperature*, *Phys. Rev. D* **54** (1996) 7857–7866, [hep-ph/9601376].
- [147] Y. Zeldovich, I. Kobzarev, and L. Okun, *Cosmological Consequences of the Spontaneous Breakdown of Discrete Symmetry*, *Zh. Eksp. Teor. Fiz.* **67** (1974) 3–11.
- [148] B. Rai and G. Senjanovic, *Gravity and domain wall problem*, *Phys. Rev. D* **49** (1994) 2729–2733, [hep-ph/9301240].
- [149] W. H. Press, B. S. Ryden, and D. N. Spergel, *Dynamical Evolution of Domain Walls in an Expanding Universe*, *Astrophys. J.* **347** (1989) 590–604.
- [150] T. Garagounis and M. Hindmarsh, *Scaling in numerical simulations of domain walls*, *Phys. Rev. D* **68** (2003) 103506, [hep-ph/0212359].
- [151] J. Oliveira, C. Martins, and P. Avelino, *The Cosmological evolution of domain wall networks*, *Phys. Rev. D* **71** (2005) 083509, [hep-ph/0410356].
- [152] P. Avelino, C. Martins, and J. Oliveira, *One-scale model for domain wall network evolution*, *Phys. Rev. D* **72** (2005) 083506, [hep-ph/0507272].
- [153] A. Leite and C. Martins, *Scaling Properties of Domain Wall Networks*, *Phys. Rev. D* **84** (2011) 103523, [arXiv:1110.3486].

- [154] M. Hindmarsh, *Analytic scaling solutions for cosmic domain walls*, *Phys. Rev. Lett.* **77** (1996) 4495–4498, [hep-ph/9605332].
- [155] M. Hindmarsh, *Level set method for the evolution of defect and brane networks*, *Phys. Rev. D* **68** (2003) 043510, [hep-ph/0207267].
- [156] A. Vilenkin, *Gravitational Field of Vacuum Domain Walls and Strings*, *Phys. Rev. D* **23** (1981) 852–857.
- [157] T. Vachaspati, A. E. Everett, and A. Vilenkin, *Radiation From Vacuum Strings and Domain Walls*, *Phys. Rev. D* **30** (1984) 2046.
- [158] M. Gleiser and R. Roberts, *Gravitational waves from collapsing vacuum domains*, *Phys. Rev. Lett.* **81** (1998) 5497–5500, [astro-ph/9807260].
- [159] J. F. Dufaux, A. Bergman, G. N. Felder, L. Kofman, and J.-P. Uzan, *Theory and Numerics of Gravitational Waves from Preheating after Inflation*, *Phys. Rev. D* **76** (2007) 123517, [arXiv:0707.0875].
- [160] T. Hiramatsu, M. Kawasaki, and K. Saikawa, *Gravitational Waves from Collapsing Domain Walls*, *JCAP* **05** (2010) 032, [arXiv:1002.1555].
- [161] M. Kawasaki and K. Saikawa, *Study of gravitational radiation from cosmic domain walls*, *JCAP* **09** (2011) 008, [arXiv:1102.5628].
- [162] K. Saikawa, *A review of gravitational waves from cosmic domain walls*, *Universe* **3** (2017), no. 2 40, [arXiv:1703.0257].
- [163] G. Janssen *et. al.*, *Gravitational wave astronomy with the SKA*, *PoS AASKA14* (2015) 037, [arXiv:1501.0012].
- [164] **NANOGrav** Collaboration, Z. Arzoumanian *et. al.*, *The NANOGrav 11-year Data Set: Pulsar-timing Constraints On The Stochastic Gravitational-wave Background*, *Astrophys. J.* **859** (2018), no. 1 47, [arXiv:1801.0261].
- [165] **LISA** Collaboration, P. Amaro-Seoane *et. al.*, *Laser Interferometer Space Antenna*, arXiv:1702.0078.
- [166] N. Seto, S. Kawamura, and T. Nakamura, *Possibility of direct measurement of the acceleration of the universe using 0.1-Hz band laser interferometer gravitational wave antenna in space*, *Phys. Rev. Lett.* **87** (2001) 221103, [astro-ph/0108011].
- [167] J. Preskill and A. Vilenkin, *Decay of metastable topological defects*, *Phys. Rev. D* **47** (1993) 2324–2342, [hep-ph/9209210].
- [168] T. Inami and C. Lim, *Effects of Superheavy Quarks and Leptons in Low-Energy Weak Processes $K_L \rightarrow \bar{\mu}\mu$, $K^+ \rightarrow \pi^+\nu\bar{\nu}$ and $K^0 \leftrightarrow \bar{K}^0$* , *Prog. Theor. Phys.* **65** (1981) 297. [Erratum: *Prog.Theor.Phys.* 65, 1772 (1981)].

- [169] G. Ecker and W. Grimus, *CP Violation and Left-Right Symmetry*, *Nucl. Phys. B* **258** (1985) 328–360.
- [170] P. T. Komiske, E. M. Metodiev, and J. Thaler, *Metric Space of Collider Events*, *Phys. Rev. Lett.* **123** (2019), no. 4 041801, [arXiv:1902.0234].
- [171] P. T. Komiske, E. M. Metodiev, and J. Thaler, *The hidden geometry of particle collisions*, *Journal of High Energy Physics* **2020** (jul, 2020).
- [172] T. Cai, J. Cheng, K. Craig, and N. Craig, *Which metric on the space of collider events?*, *Phys. Rev. D* **105** (Apr, 2022) 076003.
- [173] A. J. Larkoski and J. Thaler, *A Spectral Metric for Collider Geometry*, arXiv:2305.0375.
- [174] T. Cover and P. Hart, *Nearest neighbor pattern classification*, *IEEE Transactions on Information Theory* **13** (1967), no. 1 21–27.
- [175] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Physics Letters B* **651** (aug, 2007) 374–379.
- [176] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *Journal of High Energy Physics* **2014** (Jul, 2014).
- [177] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *An introduction to pythia 8.2*, *Computer Physics Communications* **191** (Jun, 2015) 159–177.
- [178] M. Cacciari, G. P. Salam, and G. Soyez, *Fastjet user manual*, *The European Physical Journal C* **72** (Mar, 2012).
- [179] C. Cesarotti and J. Thaler, *A Robust Measure of Event Isotropy at Colliders*, *JHEP* **08** (2020) 084, [arXiv:2004.0612].
- [180] J. Ellis, M. Lewicki, J. M. No, and V. Vaskonen, *Gravitational wave energy budget in strongly supercooled phase transitions*, *JCAP* **06** (2019) 024, [arXiv:1903.0964].
- [181] A. Wulzer, *An equivalent gauge and the equivalence theorem*, *Nuclear Physics B* **885** (aug, 2014) 97–126.
- [182] D. J. Griffiths and D. F. Schroeter, *Introduction to Quantum Mechanics*. Cambridge University Press, 3 ed., 2018.
- [183] J. Ellis, M. Lewicki, J. M. No, and V. Vaskonen, *Gravitational wave energy budget in strongly supercooled phase transitions*, *Journal of Cosmology and Astroparticle Physics* **2019** (jun, 2019) 024–024.

- [184] S. Höche, J. Kozaczuk, A. J. Long, J. Turner, and Y. Wang, *Towards an all-orders calculation of the electroweak bubble wall velocity*, *JCAP* **03** (2021) 009, [arXiv:2007.1034].
- [185] V. A. Kuzmin and M. E. Shaposhnikov, *Photon Condensation in the Hot Universe and the Longitudinal Background Radiation*, in *International Cosmic Ray Conference*, vol. 2 of *International Cosmic Ray Conference*, p. 226, Jan., 1979.
- [186] L. Bass and E. Schrödinger, *Must the photon mass be zero?*, *j-PROC-R-SOC-LOND-SER-A-MATH-PHYS-SCI* **232** (Oct., 1955) 1–6.
- [187] J. Ellis, M. Lewicki, and J. M. No, *Gravitational waves from first-order cosmological phase transitions: lifetime of the sound wave source*, *Journal of Cosmology and Astroparticle Physics* **2020** (jul, 2020) 050–050.
- [188] M. Redi and A. Tesi, *Dark Photon Dark Matter without Stueckelberg Mass*, arXiv:2204.1427.
- [189] P. W. Graham, J. Mardon, and S. Rajendran, *Vector dark matter from inflationary fluctuations*, *Physical Review D* **93** (may, 2016).
- [190] A. E. Nelson and J. Scholtz, *Dark light, dark matter, and the misalignment mechanism*, *Physical Review D* **84** (nov, 2011).
- [191] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, *WISPy cold dark matter*, *Journal of Cosmology and Astroparticle Physics* **2012** (jun, 2012) 013–013.
- [192] G. Arcadi, A. Djouadi, and M. Kado, *The Higgs-portal for vector dark matter and the effective field theory approach: A reappraisal*, *Phys. Lett. B* **805** (2020) 135427, [arXiv:2001.1075].
- [193] G. D. Kribs, G. Lee, and A. Martin, *Effective Field Theory of Stückelberg Vector Bosons*, arXiv:2204.0175.
- [194] G. R. Farrar and J. W. McIntosh, *Scattering from a domain wall in a spontaneously broken gauge theory*, *Physical Review D* **51** (may, 1995) 5889–5904.
- [195] C. G. Callan, Jr. and S. R. Coleman, *The Fate of the False Vacuum. 2. First Quantum Corrections*, *Phys. Rev. D* **16** (1977) 1762–1768.
- [196] S. R. Coleman and F. De Luccia, *Gravitational Effects on and of Vacuum Decay*, *Phys. Rev. D* **21** (1980) 3305.
- [197] E. Witten, *Instability of the Kaluza-Klein Vacuum*, *Nucl. Phys. B* **195** (1982) 481–492.

- [198] M. Dine, P. J. Fox, and E. Gorbatov, *Catastrophic decays of compactified space-times*, *JHEP* **09** (2004) 037, [hep-th/0405190].
- [199] M. Dine, J. A. P. Law-Smith, S. Sun, D. Wood, and Y. Yu, *Obstacles to Constructing de Sitter Space in String Theory*, *JHEP* **02** (2021) 050, [arXiv:2008.1239].
- [200] A. H. Guth and E. J. Weinberg, *Could the Universe Have Recovered from a Slow First Order Phase Transition?*, *Nucl. Phys. B* **212** (1983) 321–364.
- [201] **Supernova Cosmology Project** Collaboration, S. Perlmutter *et. al.*, *Measurements of Ω and Λ from 42 high redshift supernovae*, *Astrophys. J.* **517** (1999) 565–586, [astro-ph/9812133].
- [202] **SDSS** Collaboration, W. J. Percival *et. al.*, *Baryon Acoustic Oscillations in the Sloan Digital Sky Survey Data Release 7 Galaxy Sample*, *Mon. Not. Roy. Astron. Soc.* **401** (2010) 2148–2168, [arXiv:0907.1660].
- [203] C. Baker *et. al.*, *An Improved experimental limit on the electric dipole moment of the neutron*, *Phys. Rev. Lett.* **97** (2006) 131801, [hep-ex/0602020].
- [204] **ATLAS** Collaboration, M. Aaboud *et. al.*, *Combination of searches for heavy resonances decaying into bosonic and leptonic final states using 36 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *Phys. Rev. D* **98** (2018), no. 5 052008, [arXiv:1808.0238].
- [205] A. Vilenkin and A. Everett, *Cosmic Strings and Domain Walls in Models with Goldstone and PseudoGoldstone Bosons*, *Phys. Rev. Lett.* **48** (1982) 1867–1870.
- [206] A. E. Everett and A. Vilenkin, *Left-right Symmetric Theories and Vacuum Domain Walls and Strings*, *Nucl. Phys. B* **207** (1982) 43–53.
- [207] **ATLAS** Collaboration, M. Aaboud *et. al.*, *Search for single production of vector-like quarks decaying into Wb in pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *JHEP* **05** (2019) 164, [arXiv:1812.0734].
- [208] N. Arkani-Hamed, T. Han, M. Mangano, and L.-T. Wang, *Physics opportunities of a 100 TeV proton-proton collider*, *Phys. Rept.* **652** (2016) 1–49, [arXiv:1511.0649].
- [209] T. Golling *et. al.*, *Physics at a 100 TeV pp collider: beyond the Standard Model phenomena*, *CERN Yellow Rep.* (2017), no. 3 441–634, [arXiv:1606.0094].
- [210] J. Crowder and N. J. Cornish, *Beyond LISA: Exploring future gravitational wave missions*, *Phys. Rev. D* **72** (2005) 083005, [gr-qc/0506015].

- [211] R. Kleiss, W. Stirling, and S. Ellis, *A new monte carlo treatment of multiparticle phase space at high energies*, *Computer Physics Communications* **40** (1986), no. 2 359–373.
- [212] P. T. Komiske, E. M. Metodiev, and J. Thaler, *Energy flow polynomials: a complete linear basis for jet substructure*, *Journal of High Energy Physics* **2018** (Apr, 2018).
- [213] P. T. Komiske, E. M. Metodiev, and J. Thaler, *Energy flow networks: deep sets for particle jets*, *Journal of High Energy Physics* **2019** (Jan, 2019).
- [214] P. T. Komiske, E. M. Metodiev, and J. Thaler, *Cutting Multiparticle Correlators Down to Size*, *Phys. Rev. D* **101** (2020), no. 3 036019, [arXiv:1911.0449].
- [215] P. T. Komiske, R. Mastandrea, E. M. Metodiev, P. Naik, and J. Thaler, *Exploring the Space of Jets with CMS Open Data*, *Phys. Rev. D* **101** (2020), no. 3 034009, [arXiv:1908.0854].
- [216] T. Cai, J. Cheng, N. Craig, and K. Craig, *Linearized optimal transport for collider events*, *Phys. Rev. D* **102** (Dec, 2020) 116019.
- [217] T. Cai, J. Cheng, B. Schmitzer, and M. Thorpe, *The linearized hellinger–kantorovich distance*, 2021.
- [218] C. Cortes and V. Vapnik, *Support-vector networks.*, *Machine Learning* **20** (1995), no. 3 273–297.