Precision Light Flavor Physics from Lattice QCD

David Murphy

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY
2017

Abstract

Precision Light Flavor Physics from Lattice QCD

David Murphy

In this thesis we present three distinct contributions to the study of light flavor physics using the techniques of lattice QCD. These results are arranged into four self-contained papers. The first two papers concern global fits of the quark mass, lattice spacing, and finite volume dependence of the pseudoscalar meson masses and decay constants, computed in a series of lattice QCD simulations, to partially quenched $SU(2)$ and $SU(3)$ chiral perturbation theory (χPT). These fits determine a subset of the low energy constants of chiral perturbation theory — in some cases with increased precision, and in other cases for the first time — which, once determined, can be used to compute other observables and amplitudes in χPT. We also use our formalism to self-consistently probe the behavior of the (asymptotic) chiral expansion as a function of the quark masses by repeating the fits with different subsets of the data.

The third paper concerns the first lattice QCD calculation of the semileptonic $K^0 \to \pi^- \ell^+ \overline{\nu}_\ell$ ($K_{\ell 3}$) form factor at vanishing momentum transfer, $f_+^{K\pi}(0)$, with physical mass domain wall quarks. The value of this form factor can be combined with a Standard Model analysis of the experimentally measured $K^0 \to \pi^- \ell^+ \overline{\nu}_\ell$ decay rate to extract a precise value of the Cabibbo-Kobayashi-Maskawa (CKM) matrix element V_{us}, and to test unitarity of the CKM matrix. We also discuss lattice calculations of the pion and kaon decay constants, which can be used to extract V_{ud} through an analogous Standard Model analysis of experimental constraints on leptonic pion and kaon decays.

The final paper explores the recently proposed exact one flavor algorithm (EOFA). This algorithm has been shown to drastically reduce the memory footprint required to simulate single quark flavors on the lattice relative to the widely used rational hybrid Monte Carlo (RHMC) algorithm, while also offering modest $\mathcal{O}(20\%)$ speed-ups. We independently derive the exact one flavor action, explore its equivalence to the RHMC action, and demonstrate that additional preconditioning techniques can be used to significantly accelerate EOFA simulations. We apply EOFA to the ongoing

RBC/UKQCD calculation of the $\Delta I = 1/2$ $K \to \pi\pi$ decay amplitude, and demonstrate that, in this context, gauge field configurations can be generated a factor of 4.2 times faster using an EOFA-based simulation rather than the previous RHMC-based simulations. We expect that EOFA will help to significantly reduce the statistical error in the first-principles determination of the Standard Model CP-violation parameters ϵ and ϵ' offered by the $K \to \pi\pi$ calculation.

Contents

List of Tables

List of Figures

ACKNOWLEDGMENTS

I am deeply indebted to my advisor, Bob Mawhinney, for many hours of insightful discussions and suggestions leading up to the completion of this thesis. I am grateful to have had other excellent mentors to share this responsibility along the way, including, but not limited to, Szabi Márka, Chris Kelly, Norman Christ, Peter Boyle, Chulwoo Jung, and Andreas Jüttner. Much of this work was collaborative, and I am also indebted to my fellow graduate students and more senior colleagues in the RBC/UKQCD collaboration for their many inputs.

Lastly, I would like to thank my wonderful wife, Prapti, and my family and friends, for many years of unwavering support as I ignored them to study physics.

Chapter 1

Introduction

In this chapter we briefly review the QCD Lagrangian and its properties, as well as the additional couplings between quarks and the electroweak gauge bosons in the full Standard Model of particle physics. We then discuss the lattice field theory framework used throughout this thesis, and outline the basics of a prototypical lattice QCD calculation.

1.1 Quantum Chromodynamics

Quantum chromodynamics (QCD) is widely believed to be the correct theoretical description of the strong nuclear interaction, which is responsible for binding quarks into hadrons and hadrons into nuclei. QCD is a nonabelian gauge theory with gauge group $SU(3)$. The basic degrees of freedom are the quark fields $\psi_f(x)$ — which are spin-1/2 Dirac fermions transforming in the 3-dimensional fundamental representation of $SU(3)$ — and the gluon fields $A^a_\mu(x)$ — which are spin-1 vector bosons transforming in the 8-dimensional adjoint representation of $SU(3)$. The distinct quark species are known as "flavors", and are labeled by the flavor index f; the properties of the six experimentally observed quark flavors are summarized in Table 1.1. QCD is defined by the Lagrangian

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a + \sum_f \overline{\psi}_f \left(i\gamma^\mu D_\mu - m_f \right) \psi_f, \tag{1.1}$$

where D_μ is the gauge-covariant derivative

$$D_\mu \equiv \partial_\mu + igA^a_\mu T^a, \tag{1.2}$$

$G^a_{\mu\nu}$ is the gluon field strength tensor

$$G^a_{\mu\nu} \equiv \partial_\nu A^a_\nu - \partial_\nu A^a_\mu + g f^{abc} A^b_\mu A^c_\nu, \tag{1.3}$$

and $\{T^a\}$ is a basis for the Lie algebra $\mathfrak{su}(3)$, conventionally chosen to satisfy $\mathrm{tr}\left(T^a T^b\right) = 1/2\,\delta^{ab}$. The $SU(3)$ structure constants f^{abc} are defined by the relation

$$\left[T^a, T^b\right] = i f^{abc} T^c, \tag{1.4}$$

and the coupling constant g parametrizes the strength of the quark-gluon and gluon-gluon interactions. One can check that the Lagrangian of Eqn. (1.1) is indeed invariant under an arbitrary, spacetime-dependent $SU(3)$ transformation $V(x)$ provided the quark and gluon fields transform as

$$\begin{cases} \psi(x) \to V(x)\psi(x) \\ A^a_\mu(x) \to V(x)\left(A^a_\mu(x) T^a + \dfrac{i}{g}\partial_\mu\right) V^\dagger(x) \end{cases}, \tag{1.5}$$

and that the classical equations of motion are the Yang-Mills equations

$$\begin{cases} \left(i\gamma^\mu D_\mu - m\right)\psi_f = 0 \\ \partial^\mu G^a_{\mu\nu} + g f^{abc} A^{b\mu} G^c_{\mu\nu} = -g j^a_\nu \end{cases} \tag{1.6}$$

where

$$j^a_\mu \equiv \sum_f \overline{\psi}_f \gamma_\mu T^a \psi_f. \tag{1.7}$$

The classical QCD Lagrangian with N_f quark flavors has an additional global $U(N_f)_L \otimes U(N_f)_R$ symmetry in the massless limit, known as *chiral symmetry*. Introducing chiral projection operators $P_L = \frac{1}{2}\left(1 - \gamma^5\right)$ and $P_R = \frac{1}{2}\left(1 + \gamma^5\right)$ and arranging the quark fields into an N_f-component vector Ψ, the left-handed $\Psi_L = P_L \Psi$ and right-handed $\Psi_R = P_R \Psi$ components naturally decouple as

$$\mathcal{L}_{\mathrm{QCD}} \supset \overline{\Psi} i\gamma^\mu D_\mu \Psi = \overline{\Psi}_L i\gamma^\mu D_\mu \Psi_L + \overline{\Psi}_R i\gamma^\mu D_\mu \Psi_R. \tag{1.8}$$

Independent unitary rotations $\Psi_L \to U_L \Psi_L$ and $\Psi_R \to U_R \Psi_R$ leave the Lagrangian invariant, leading to conserved vector

$$j^\mu = \overline{\Psi}\gamma^\mu \Psi, \quad j^{\mu a} = \overline{\Psi}\gamma^\mu T^a \Psi \tag{1.9}$$

and axial

$$j_5^\mu = \overline{\Psi}\gamma^\mu\gamma^5\Psi, \quad j_5^{\mu a} = \overline{\Psi}\gamma^\mu\gamma^5 T^a\Psi \tag{1.10}$$

currents. A careful calculation [1] demonstrates that the axial singlet transformation associated with the $U(1)_A$ subgroup fails to be a symmetry of the quantized theory, leading to the *Adler-Bell-Jackiw anomaly*

$$\partial_\mu j_5^\mu = -\frac{g^2 N_f}{32\pi^2}\epsilon^{\mu\nu\rho\sigma}F^a_{\mu\nu}F^a_{\rho\sigma}. \tag{1.11}$$

The $SU(N_f)_L \otimes SU(N_f)_R$ subgroup is known to be further spontaneously broken to $SU(N_f)_V$ — as manifested through nonzero vacuum expectation values of the quark condensates $\langle\overline{\psi}_f\psi_f\rangle \neq 0$ — giving rise to $N_f^2 - 1$ Goldstone bosons (the pseudoscalar mesons).

Many of the basic features of QCD are analogous to features of the simpler $U(1)$ gauge theory of electromagnetism, known as quantum electrodynamics (QED). Just as the electron carries electric charge and interacts with the photon, quarks carry "color charge" and interact with gluons. However, the nonabelian $SU(3)$ gauge group of QCD leads to a richer structure: there are 3 types of color charge, conventionally labeled "red", "blue", and "green", and the color-neutral bound states we observe in nature — analogous to electrically neutral molecules — include mesons with quark structure

$$\overline{\psi}^a\psi^a \tag{1.12}$$

and baryons with quark structure

$$\varepsilon^{abc}\psi^a\psi^b\psi^c. \tag{1.13}$$

In the past 40 years this model has been phenomenally successful at explaining the "particle zoo" of hadrons in terms of the more fundamental quarks. The experimentally observed properties of a representative collection of some light mesons and baryons are listed in Table 1.2.

In other ways, QCD is quite unlike QED. While the photons do not self-interact as they mediate the electromagnetic force between charged particles, the non-vanishing $SU(3)$ structure constants f^{abc} lead to gluon self-interactions with the Feynman rules summarized in Figure 1.1. These additional interactions lead to remarkably different behavior in the renormalization group running of the electric and strong coupling constants. For QED, the one-loop running of the fine structure

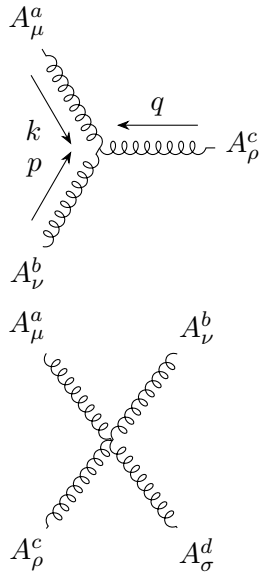

$$g f^{abc} \left[\eta^{\mu\nu} (k-p)^\rho + \eta^{\nu\rho} (p-q)^\mu + \eta^{\rho\mu} (q-k)^\nu\right]$$

$$\begin{aligned} -ig^2 \Big[& f^{abe} f^{cde} (\eta^{\mu\rho}\eta^{\nu\sigma} - \eta^{\mu\sigma}\eta^{\nu\rho}) \\ & + f^{ace} f^{bde} (\eta^{\mu\nu}\eta^{\rho\sigma} - \eta^{\mu\sigma}\eta^{\nu\rho}) \\ & + f^{ade} f^{bce} (\eta^{\mu\nu}\eta^{\rho\sigma} - \eta^{\mu\rho}\eta^{\nu\sigma}) \Big] \end{aligned}$$

Figure 1.1: Feynman rules for the gluon self-interactions.

constant $\alpha \equiv e^2/4\pi$ takes the form [1]

$$\alpha\left(\mu'\right) = \frac{\alpha\left(\mu\right)}{1 - \frac{\alpha(\mu)}{3\pi} \log\left(\frac{\mu'}{\mu}\right)}, \tag{1.14}$$

with, experimentally, $\alpha(\mu = m_e) \approx 1/137$. The theory is weakly coupled — and thus amenable to perturbative calculations — at energy scales relevant to current collider experiments, but slowly grows with μ'. An analogous one-loop QCD calculation demonstrates that the running of the strong coupling constant $\alpha_s \equiv g^2/4\pi$ is instead given by

$$\alpha_s\left(\mu'\right) = \frac{2\pi}{\beta_0 \log\left(\frac{\mu'}{\Lambda_{\mathrm{QCD}}}\right)}, \tag{1.15}$$

with $\beta_0 = 11 - \frac{2}{3}N_f = 7$ [1]. This leads to a remarkably different behavior, known as *asymptotic freedom*: at very high energy scales $\alpha_s \ll 1$, and QCD becomes a perturbative theory of free quarks and gluons, while at low energies $\alpha_s \to \infty$ and the theory is strongly coupled, with quarks permanently bound together into hadrons (*confinement*). As a result, nonperturbative methods are necessary to understand low energy QCD and to calculate hadron properties or strongly interacting matrix elements from first principles. The renormalization scheme-dependent parameter

$\Lambda_{\rm QCD} \sim 200$ MeV characterizes the scale at which the QCD coupling is of order one, and the theory transitions between these two phases.

Name	Symbol	Mass (MeV/c^2)	J^P	B	Q (e)	I_3
Up	u	$2.3\,\binom{+0.7}{-0.5}$	$\frac{1}{2}^+$	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{2}$
Down	d	$4.8\,\binom{+0.5}{-0.3}$	$\frac{1}{2}^+$	$\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{2}$
Strange	s	95(5)	$\frac{1}{2}^+$	$\frac{1}{3}$	$-\frac{1}{3}$	0
Charm	c	1275(25)	$\frac{1}{2}^+$	$\frac{1}{3}$	$\frac{2}{3}$	0
Bottom	b	4180(30)	$\frac{1}{2}^+$	$\frac{1}{3}$	$-\frac{1}{3}$	0
Top	t	173210(510)(710)	$\frac{1}{2}^+$	$\frac{1}{3}$	$\frac{2}{3}$	0

Table 1.1: A summary of some properties of the six known quark flavors. Assigning masses to the quarks is highly non-trivial since all but the top quark are permanently bound into hadrons at low energies — rendering them unobservable in experiments — and thus the quark masses typically depend on a choice of parametrization and renormalization scheme. The masses of the light (up, down, and strange) quarks are most accurately determined using lattice QCD calculations together with experimentally determined hadron masses as inputs. More detail regarding the precise definitions of the quark masses, as well as the values quoted here, can be found in the most recent particle data group (PDG) review [2].

Particle	Quark Composition	J^P	Mass (MeV/c^2)	Mean Lifetime (s)	Primary Decay Mode
$\pi^{\pm}$	+: $u\overline{d}$, −: $d\overline{u}$	0^-	139.57018(35)	$2.6033(5) \times 10^{-8}$	$\mu^+ + \nu_\mu$
π^0	$(u\overline{u} - d\overline{d})/\sqrt{2}$	0^-	134.9766(6)	$8.52(18) \times 10^{-17}$	2γ
$K^{\pm}$	+: $u\overline{s}$, −: $s\overline{u}$	0^-	493.677(16)	$1.2380(21) \times 10^{-8}$	$\mu^+ + \nu_\mu$
K^0	$d\overline{s}$	0^-	497.611(13)	—	—
K^0_S	$(d\overline{s} + s\overline{d})/\sqrt{2}$	0^-	497.611(13)	$8.954(4) \times 10^{-11}$	$\pi^+ + \pi^-$
K^0_L	$(d\overline{s} - s\overline{d})/\sqrt{2}$	0^-	497.611(13)	$5.116(21) \times 10^{-8}$	$\pi^{\pm} + e^{\mp} + \nu_e$
η	$(u\overline{u} + d\overline{d} - 2s\overline{s})/\sqrt{6}$	0^-	547.862(17)	$5.02(19) \times 10^{-19}$	2γ
η'	$(u\overline{u} + d\overline{d} + s\overline{s})/\sqrt{3}$	0^-	957.78(6)	$3.32(15) \times 10^{-21}$	$\pi^+ + \pi^- + \eta$
p^+	uud	$\frac{1}{2}^+$	938.272081(6)	$> 2.1 \times 10^{29}$ years	—
n^0	udd	$\frac{1}{2}^+$	939.565413(6)	880.2(1.0)	$p^+ + e^- + \overline{\nu}_e$
Λ^0	uds	$\frac{1}{2}^+$	1115.683(6)	$2.632(20) \times 10^{-10}$	$p^+ + \pi^-$
Σ^+	uus	$\frac{1}{2}^+$	1189.37(7)	$8.018(26) \times 10^{-11}$	$p^+ + \pi^0$
Σ^0	uds	$\frac{1}{2}^+$	1192.642(24)	$7.4(0.7) \times 10^{-20}$	$\Lambda^0 + \gamma$
Σ^-	dds	$\frac{1}{2}^+$	1197.449(30)	$1.479(11) \times 10^{-10}$	$n^0 + \pi^-$
Ξ^0	uss	$\frac{1}{2}^+$	1314.86(20)	$2.90(9) \times 10^{-10}$	$\Lambda^0 + \pi^0$
Ξ^-	dss	$\frac{1}{2}^+$	1321.71(7)	$1.639(15) \times 10^{-10}$	$\Lambda^0 + \pi^-$
Ω^-	sss	$\frac{3}{2}^+$	1672.45(29)	$8.21(11) \times 10^{-11}$	$\Lambda^0 + K^-$

Table 1.2: Quark content and experimentally observed properties of some representative light unflavored and strange hadrons [2].

1.2 Electroweak Interactions and the CKM Matrix

In the context of the full $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ Standard Model of particle physics, additional interactions between quarks and the electroweak sector are allowed. Yukawa couplings of the quarks to the Higgs field give rise to the quark masses through the Higgs mechanism. In addition, the Standard Model Lagrangian contains interaction terms coupling quarks to the electroweak gauge bosons, which lead to the Feynman rules summarized in Figure 1.2. In particular, flavor changing weak decays coupling quarks to the charged $W^{\pm}_{\mu}$ bosons introduce elements V_{ij} of the Cabibo-

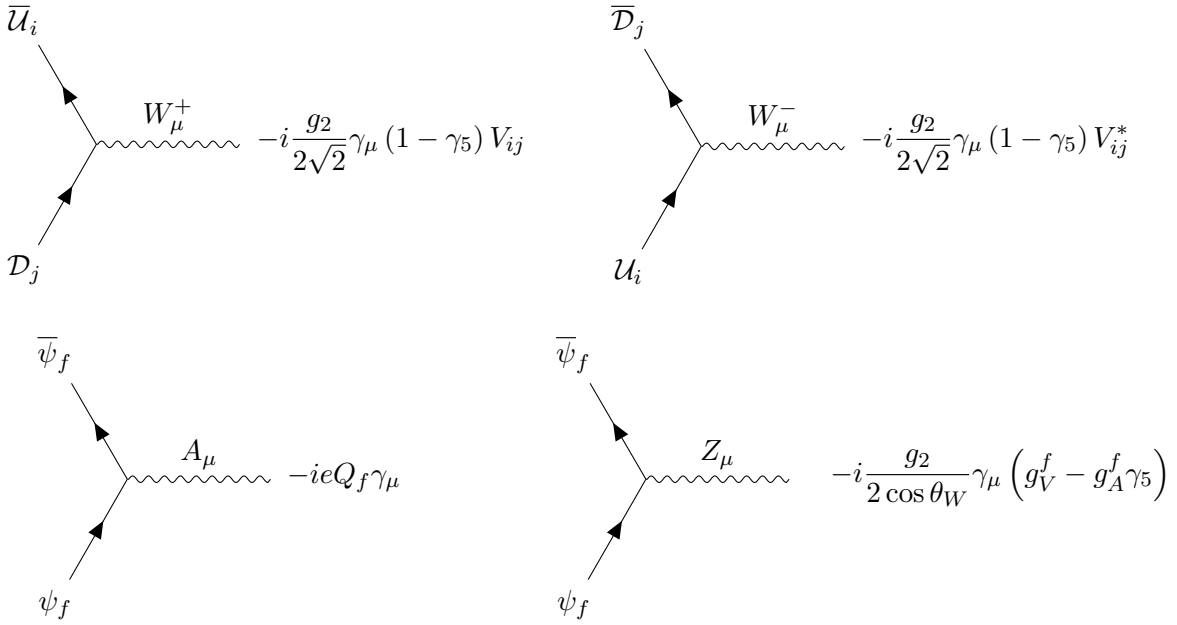

Figure 1.2: Feynman rules for the interactions between quarks and electroweak gauge bosons in the Standard Model. $\mathcal{U}_i \in \{u, c, t\}$ is a charge $2/3\,e$ quark, $\mathcal{D}_i \in \{d, s, b\}$ is a charge $-1/3\,e$ quark, and V_{ij} is the CKM matrix.

Kobayashi-Maskawa (CKM) matrix. The CKM matrix is a 3×3 unitary matrix which encodes the relative probability of a quark with flavor i decaying through the weak interaction to a quark with flavor j; this occurs because the quark eigenstates which couple to the charged electroweak currents are superpositions of the mass eigenstates appearing in the QCD Lagrangian. The CKM matrix can be parametrized by three real angles and a single complex phase, which have been extensively constrained by experiment. Crucially, the complex phase allows for certain types of Standard Model decays which violate time-reversal symmetry, and hence charge-parity (CP) symmetry[1]. A detailed discussion of known CP-violating decay mechanisms — as well as the current state-of-the-art for experimental determinations of the CKM matrix elements — can be found in the most recent PDG review [2].

Precision determinations of the CKM matrix are an important tool in the search for new physics

[1]This follows since conservation of the product CPT can be shown to hold for any reasonable quantum field theory [3].

beyond the Standard Model (BSM), either through tension between different constraints on the same matrix element or through tension with unitarity. Substantial theoretical input is typically required to relate individual CKM matrix elements to experimentally observable processes since free quarks cannot be observed directly in collider experiments. While electroweak contributions can be calculated perturbatively, strongly interacting matrix elements describing the hadronic part of the decay must be calculated with nonperturbative techniques.

1.3 Lattice QCD

The idea of regulating QCD with a finite spacetime lattice dates back to the seminal work of Nobel laureate Kenneth Wilson in the 1970's. Leveraging this theoretical insight with the computational power of modern supercomputers has led to modern lattice QCD, which provides the only known method for performing first-principles QCD calculations in the low energy, nonperturbative regime with fully controlled systematic errors. To formulate lattice QCD from continuum QCD, we first perform a Wick rotation

$$x_0 \to -ix_4 \tag{1.16}$$

from Minkowski spacetime $\mathbb{R}^{1,3}$ to Euclidean spacetime $\mathbb{R}^4$, and then discretize the theory by replacing continuous, infinite volume spacetime with a finite grid. The lattice naturally regulates both infrared (IR) and ultraviolet (UV) divergences, since the spacing between grid points imposes a short distance cutoff, and the total length of the box imposes a long distance cutoff. In a typical lattice calculation the lattice spacing is chosen to be isotropic in all directions; we will denote this distance by a. Likewise, the total number of lattice sites along the three spatial directions L are typically chosen to be equal, while the number of temporal sites T is often chosen to be somewhat larger than L.

The Euclidean QCD action

$$S_E = \int d^4x_E \left[\frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a + \sum_f \overline{\psi}_f \left(\gamma^E_\mu D_\mu + m \right) \psi_f \right] \tag{1.17}$$

has a real and positive definite contribution from the gluonic action, and a positive semi-definite contribution from the fermionic action, allowing n-point correlation functions to be expressed in

terms of a well-defined path integral representation

$$\langle \mathscr{O}_1(x_1)\cdots\mathscr{O}_n(x_n)\rangle = \frac{1}{\mathcal{Z}}\int \mathscr{D}A^a_\mu\Big(\prod_f \mathscr{D}\psi_f\mathscr{D}\overline{\psi}_f\Big)\,(\mathscr{O}_1(x_1)\cdots\mathscr{O}_n(x_n))\,e^{-S_E[A^a_\mu,\overline{\psi}_f,\psi_f]}, \tag{1.18}$$

with the lattice path integral measures defined to be

$$\int \mathscr{D}A^a_\mu \equiv \prod_{x,\mu}\int dA^a_\mu(x), \quad \int \mathscr{D}\psi \equiv \prod_x \int d\psi(x). \tag{1.19}$$

For typical lattice volumes this integral has far too many degrees of freedom to compute directly, but can be evaluated numerically using Monte Carlo techniques. In the following subsections we will elaborate on details of how the Euclidean action is discretized, and how explicit numerical calculations of n-point correlation functions can be performed. We will drop the subscript "E" in the remainder of this work, with the theory understood to be formulated in Euclidean spacetime unless otherwise stated.

1.3.1 Gauge Actions

On the lattice the gauge field is represented by $SU(3)$-valued link variables $U_\mu(x)$ that are understood to reside on the "links" connecting discrete spacetime points. To motivate this, we note the following problem with a naïve discretization of the fermionic contribution to the Euclidean action (Eqn. (1.17)): the derivative will involve a nonlocal bilinear $\overline{\psi}(x)\psi(y)$ which fails to be gauge invariant under the transformations of Eqn. (1.5). In the continuum gauge invariance can be restored by introducing the *Wilson line*

$$U_P(x,y) \equiv \mathcal{P}\exp\left(-ig\int_P dz^\mu A^a_\mu(z)T^a\right), \tag{1.20}$$

where P denotes path ordering along an arbitrary path connecting x and y. The quantity $\overline{\psi}(x)U_P(x,y)\psi(y)$ is invariant under independent gauge transformations $V(x), V(y) \in SU(3)$, as desired, and $U_P(x,y)$ can be thought of as the object responsible for parallel transport of a Dirac spinor $\psi(x)$ to other points on the spacetime manifold. For small separations between adjacent lattice sites we can approximate P with a line segment along the basis vector $\hat{\mu}$ pointing in the direction x^μ, and $U_P(x,y)$ with the gauge link

$$U_\mu(x) \equiv \mathcal{P}e^{-iga A^b_\mu(x)T^b}. \tag{1.21}$$

The *Wilson loop* is likewise defined as a Wilson line around a closed path $U_P(x,x)$ returning to the same spacetime point.

Lattice actions for $SU(3)$ Yang-Mills theory without fermions can be constructed from Wilson loops of varying sizes. The most basic is the 1×1 *plaquette*

$$U_{\mu\nu}(x) \equiv U_\mu(x)U_\nu(x+a\hat{\mu})U_\mu^\dagger(x+a\hat{\nu})U_\nu^\dagger(x), \tag{1.22}$$

from which we can form the Wilson gauge action

$$S_G^W = \beta \,\mathrm{tr} \sum_{\mu<\nu}\sum_x \left[1 - \frac{1}{2}\left(U_{\mu\nu}(x) + U_{\mu\nu}^\dagger(x)\right)\right], \tag{1.23}$$

with $\beta = 6/g^2$. One can check that the plaquette can be identified with a discretized gluon field strength tensor $\mathcal{G}_{\mu\nu}^a$ according to $U_{\mu\nu}(x) = \exp\left(iga^2\mathcal{G}_{\mu\nu}^b T^b\right)$, and that $\mathcal{G}_{\mu\nu}^a \to G_{\mu\nu}^a$ in the continuum limit $a\to 0$, such that Equation (1.23) reduces to the Yang-Mills action (*i.e.* Equation (1.17) with no quark fields) in the same limit [4]. At finite lattice spacing $\mathcal{G}_{\mu\nu}^a$ differs from the continuum $G_{\mu\nu}^a$ by irrelevant operators which are multiplied by powers of a. Improved gauge actions which suppress subsets of these operators can be constructed by adding additional Wilson loops such as the 1×2 *rectangle*

$$R_{\mu\nu} \equiv U_\mu(x)U_\mu(x+a\hat{\mu})U_\nu(x+2a\hat{\mu})U_\mu^\dagger(x+2a\hat{\mu}+a\hat{\nu})U_\mu^\dagger(x+a\hat{\mu}+a\hat{\nu})U_\nu^\dagger(x+a\hat{\nu}). \tag{1.24}$$

All of the QCD simulations presented in this thesis make use of the Iwasaki gauge action

$$S_G^I = \beta \,\mathrm{tr} \sum_x \left[(1-8c_1)\sum_{\mu<\nu}\frac{1}{2}\left(U_{\mu\nu}(x) + U_{\mu\nu}^\dagger(x)\right) + c_1 \sum_{\mu\neq\nu}\frac{1}{2}\left(R_{\mu\nu}(x) + R_{\mu\nu}^\dagger(x)\right)\right], \tag{1.25}$$

where the value $c_1 = -0.331$ was determined by a nonperturbative spin-blocking analysis in Ref. [5]. Some simulations supplement the Iwasaki gauge action with the dislocation suppressing determinant ratio (DSDR), which is an additional term in the gauge action designed to reduce chiral symmetry breaking effects in domain wall fermion calculations with coarse lattice spacings [6].

1.3.2 Fermion Actions

Unlike the case of the gauge action, attempts to straightforwardly discretize the fermion action quickly run into trouble. Considering, for simplicity, a single flavor, the Euclidean Dirac action can

be written as $S = \overline{\psi}_\alpha(x) K_{\alpha\beta}(x,y) \psi_\beta(y)$, with

$$K_{\alpha\beta}(x,y) = \frac{1}{2a} \left(\gamma_\mu\right)_{\alpha\beta} \left(\delta_{y,x+a\hat{\mu}} - \delta_{y,x-a\hat{\mu}}\right) + m\delta_{\alpha\beta}\delta_{xy} \tag{1.26}$$

and α, β denoting spinor indices. By taking the Fourier transform of this operator and inverting in momentum space, it can be shown [4] that the lattice propagator is

$$\left\langle \psi_\alpha(x) \overline{\psi}_\beta(y) \right\rangle = \int_{-\pi/a}^{\pi/a} \frac{d^4p}{(2\pi)^4} \frac{\left[-i \sum_\mu \gamma_\mu \widetilde{p}_\mu + m \right]_{\alpha\beta}}{\sum_\mu \widetilde{p}_\mu^2 + m^2} e^{ip\cdot(x-y)}, \tag{1.27}$$

with

$$\widetilde{p}_\mu = \frac{1}{a} \sin\left(a p_\mu\right). \tag{1.28}$$

In the limit $a \to 0$ we should recover the continuum Dirac propagator, but this is spoiled by the observation that $\widetilde{p}_\mu \approx p_\mu$ not only near the origin, but also when $|p_\mu| \approx \pi/a$. Since this is true of any individual component of the momentum, we see that the naïve fermion action of Equation (1.26) actually describes sixteen degenerate fermion "tastes" in the continuum limit. This is known as the *fermion doubling problem*.

In Refs. [7, 8] Nielsen and Ninomiya provided an elegant characterization of the fermion doubling problem through a famous no-go theorem. They proved that it is not possible to construct a lattice Dirac operator $\mathscr{D}$ for an even dimensional spacetime which is simultaneously:

1. Hermitian
2. Translationally invariant
3. Local, *i.e.* $\mathscr{D}(x,y)$ decays exponentially fast at large distances $|x-y| \gg 1$
4. Consistent with chiral symmetry at vanishing quark mass, *i.e.* respecting $\{\mathscr{D}, \gamma_5\} = 0$
5. Free of doublers

In essence, their proof exploits the Poincaré-Hopf index theorem to demonstrate that conditions 1-4 necessarily lead to doublers for a lattice theory defined on an even-dimensional torus $\mathbb{T}^d$.

A number of fermion actions are in common use in the literature, including *Wilson*, *staggered*, *twisted mass*, *domain wall*, and *overlap* fermions. These actions typically involve trade-offs between

violating particular conditions of the Nielsen-Ninomiya theorem, the size of lattice artifacts at finite lattice spacing, and the relative computational cost of performing a simulation. The best choice of action for a particular calculation is often highly dependent on the details of the target physics and the available computational resources. We will not attempt to provide a general overview, since reviews of each formulation can be found in the literature, but will instead focus on the domain wall fermion action used in this thesis.

Domain Wall Fermions

Domain wall fermions (DWF) avoid the Nielsen-Ninomiya no-go theorem in a particularly clever way: by adding a fictitious fifth spatial direction — conventionally labeled s, with L_s lattice sites along this direction — to sidestep the critical assumption of an even-dimensional spacetime. Shamir and Furman [9, 10], building off of earlier work by Kaplan [11], demonstrated that effective 4D chiral fermions can be recovered at the s-boundaries of a five dimensional theory. While the DWF formalism has the nice property that it can have arbitrarily exact chiral symmetry in the limit $L_s \to \infty$, and is empirically found to maintain excellent chiral symmetry even at modest L_s, this advantage comes at the price of an $\mathcal{O}(L_s)$ increase in the computational cost due to the extra dimension.

The generic domain wall-type fermion action takes the form

$$S_{\text{DWF}}\left[\overline{\psi}, \psi, U\right] = \sum_{xs} \sum_{x's'} \overline{\psi}_{xs} \left(\mathscr{D}_{\text{DWF}}\right)_{xs;x's'} \psi_{x's'}, \tag{1.29}$$

where

$$\left(\mathscr{D}_{\text{DWF}}\right)_{xs;x's'} = b_s \left(D_W\right)_{xx'} \delta_{ss'} + \delta_{xx'}\delta_{ss'} + c_s \left(D_W\right)_{xx'} L_{ss'} - \delta_{xx'} L_{ss'} \tag{1.30}$$

is the DWF Dirac operator,

$$\left(D_W\right)_{xx'} = (4 + M_5)\, \delta_{xx'} - \frac{1}{2} \sum_{\mu} \left[(1 - \gamma_\mu)\, U_\mu(x) \delta_{x+\hat{\mu},x'} + (1 + \gamma_\mu)\, U_\mu^\dagger(x') \delta_{x-\hat{\mu},x'}\right] \tag{1.31}$$

is the four-dimensional Wilson Dirac operator, and

$$L_{ss'} = \left(L_+\right)_{ss'} P_R + \left(L_-\right)_{ss'} P_L \tag{1.32}$$

is the 5D hopping matrix, with

$$(L_+)_{ss'} = (L_-)_{s's} = \begin{cases} -m\delta_{L_s-1,s'}, & s = 0 \\ \delta_{s-1,s'}, & 1 \le s \le L_s - 1 \end{cases} . \tag{1.33}$$

This construction may be regarded as a theory of L_s Wilson fermions of mass $-M_5$ that mix through the "mass" matrix $L_{ss'}$. The gauge field remains a four-dimensional object and is merely replicated for each s-slice. Four dimensional fermion fields q and $\overline{q}$ with mass m and definite chiralities are recovered from the five dimensional quark fields ψ and $\overline{\psi}$ at the boundaries of the fifth dimension

$$\begin{aligned} q_L &= P_L\psi_0 & q_R &= P_R\psi_{L_s-1} \\ \overline{q}_L &= \overline{\psi}_0 P_R & \overline{q}_R &= \overline{\psi}_{L_s-1} P_L \end{aligned} . \tag{1.34}$$

Correlation functions constructed from q and $\overline{q}$ approximate continuum QCD arbitrarily well in the simultaneous continuum and infinite volume limits.

Propagation and mixing of the light left-handed and right-handed modes through the fifth dimension is exponentially suppressed in L_s, but still nonzero when L_s is finite. In addition, the doubler states appear as heavy modes propagating in the five-dimensional bulk. It can be shown that this leads to mild chiral symmetry breaking effects, the largest of which is an additive renormalization of the bare fermion mass $m \to m+m_{\text{res}}$ by the *residual mass* (m_{res}) [12]. Simulating QCD with light pions forces L_s to be taken sufficiently large to keep m_{res} under control. In the limit $L_s \to \infty$, however, the heavy modes propagating in the five-dimensional bulk dominate the spectrum, leading to a divergence. This divergence can be removed by introducing a heavy, Pauli-Villars regulator field: in practice one always computes a determinant ratio

$$\det\left(\frac{\mathscr{D}(m)}{\mathscr{D}(m_{\text{pv}})}\right) \tag{1.35}$$

with $m_{\text{pv}} \gg m$ when simulating QCD with domain wall fermions. This modification can be shown to remove the bulk divergence without affecting the desired low-energy chiral physics [13].

In addition to tuning L_s, the coefficients b_s and c_s can also be chosen to further suppress chiral symmetry breaking, at the expense of making domain wall fermions more expensive to simulate; the ability to achieve the same m_{res} with smaller L_s often justifies the use of these more sophisticated actions. The original *Shamir* DWF construction of Shamir and Furman has $b_s = 1$ and $c_s = 0$ for all s. Other variants commonly used in the literature include:

- *Möbius* DWF [14–16]: $b_s - c_s = 1$ and $b_s + c_s = \alpha$ for all s, where α is a free parameter known as the *Möbius scale*.

- *Optimal* DWF [17]: b_s and c_s are real parameters constructed to minimize chiral symmetry breaking at fixed L_s.

- *zMöbius* DWF [18, 19]: b_s and c_s are complex parameters constructed to minimize chiral symmetry breaking at fixed L_s.

The simulations presented in this thesis make use of either the Shamir or Möbius DWF action.

1.3.3 Boundary Conditions

Completely specifying a lattice simulation requires a choice of boundary conditions for the gauge and fermion fields in addition to a choice of action. Typical simulations apply *periodic boundary conditions* to the gauge field $U_\mu(x + L_\mu\hat{\mu}) = U_\mu(x)$ and to the fermion fields along spatial directions, and *antiperiodic boundary conditions* to the fermion fields $\psi(x + T\hat{t}) = -\psi(x)$ along the temporal direction[2]. More sophisticated boundary conditions are sometimes used in special contexts.

In the $K_{\ell 3}$ calculation we make use of *twisted boundary conditions* for the light quark fields

$$\psi(x + L_i\theta_i) = e^{i\theta_i}\psi(x), \tag{1.36}$$

which modify the allowed quark momentum states

$$p_i = \frac{2\pi n_i}{L} + \frac{\theta_i}{L}, \quad n_i \in \mathbb{Z} \tag{1.37}$$

as can easily be seen by Fourier transforming to momentum space. Twisted boundary conditions allow the momentum of the ground state to be tuned to an arbitrary kinematical point by an appropriate choice of θ_i, which is, in general, both less expensive and more accurate than repeating a simulation for multiple choices of n_i and interpolating to the desired momentum. One important but subtle point is that typical lattice calculations actually apply *partially twisted boundary conditions*,

[2] Lüscher has demonstrated that the quark fields must have antiperiodic temporal boundary conditions to construct a proper transfer matrix representation of the lattice theory, which is typically defined in terms of the path integral representation [20].

where Equation (1.36) is applied only to the valence quark fields. Sachrajda and Villadoro [21] have studied this issue in chiral perturbation theory and demonstrated that partial twisting introduces exponentionally small finite volume errors for single particle states, but introduces large, power law finite volume errors for interacting, multi-particle final states. This makes twisted boundary conditions suitable for $K \to \pi$ decays, but other techniques must be used *e.g.* for the $K \to \pi\pi$ decay.

In the $\Delta I = 1/2$ $K \to \pi\pi$ calculation we instead apply *G-parity boundary conditions* to the light quark doublet, where the G-parity operation is the product of charge conjugation and a 180° isospin rotation about the y-axis $\hat{G} = \hat{C}\exp(i\pi\hat{I}_y)$. It can be shown that the pions are odd eigenstates of the G-parity operation

$$\hat{G}\begin{pmatrix}\pi^+\\ \pi^0\\ \pi^-\end{pmatrix} = (-1)\begin{pmatrix}\pi^+\\ \pi^0\\ \pi^-\end{pmatrix}, \tag{1.38}$$

implying that the allowed pion momenta are odd-integer multiples of $\pm\pi/L$. Together with careful tuning of the ensemble parameters, G-parity boundary conditions allow for simulations of the $K \to \pi\pi$ decay with physical kinematics and the final pions in the ground state. G-parity was introduced as a quantum number long ago by Lee and Yang [22], but has only recently been successfully applied to lattice QCD simulations; Christ and Kim suggested the application of G-parity as a boundary condition for the $\Delta I = 1/2$ $K \to \pi\pi$ decay in Ref. [23].

1.3.4 The Hybrid Monte Carlo Algorithm

In theory, after specifying the gauge action, fermion action, and boundary conditions, arbitrary n-point correlation functions can be computed on a spacetime lattice by evaluating the path integral of Equation (1.18). In practice, a typical lattice has far too many degrees of freedom to perform this calculation directly. Instead, one applies Monte Carlo techniques to ergodically sample a representative Markov chain of gauge field configurations $\{U_\mu^{(i)}\}_{i=1}^N$, for which

$$\langle \mathscr{O}_1(x_1)\cdots\mathscr{O}_n(x_n)\rangle \approx \frac{1}{N}\sum_{i=1}^{N}\mathscr{O}_1(x_1)\cdots\mathscr{O}_n(x_n) \tag{1.39}$$

up to $\mathcal{O}(1/\sqrt{N})$ *statistical errors*. In addition, to avoid having to represent anticommuting Grassman variables in a computer, the fermions are integrated out and reintroduced in terms of bosonic "pseudofermion" fields ϕ as

$$\int \mathscr{D}\psi\mathscr{D}\overline{\psi}e^{-\overline{\psi}M\psi} = \det(M) = \frac{1}{\det(M^{-1})} = \int \mathscr{D}\phi\mathscr{D}\phi^{\dagger}e^{-\phi^{\dagger}M^{-1}\phi}. \tag{1.40}$$

This comes at the cost of applications of M^{-1} — typically through iterative algorithms like conjugate gradient (CG) [24] — rather than M.

The simplest Monte Carlo scheme used in lattice simulations is known as the *Metropolis algorithm*. At each step a single gauge link is randomly chosen and modified. The change in the action ΔS is then computed, and the new gauge field configuration is accepted with probability $P = \min(1, e^{-\Delta S})$. While it can be shown that this simple algorithm is sufficient to properly sample the gauge field [4], it also suffers from a number of drawbacks. In particular, modifying a single gauge link at a time moves through the space of gauge field configurations far too slowly to be of practical use unless the lattice volume is very small, especially since the determinant of the Dirac operator must be recomputed each time the gauge field is modified. This problem is not easily addressed since randomly modifying many links at once leads to large changes in the action, and poor acceptance. Practical lattice simulations require a method to globally update the entire gauge field without unacceptably large changes in the action.

The *Hybrid Monte Carlo* (HMC) algorithm is a global update Monte Carlo technique that is widely used in modern lattice calculations. We will only briefly describe HMC here, since the algorithm is discussed in detail in Section 6.4. After introducing an $SU(3)$-valued conjugate momentum $\pi_\mu(x)$ for the gauge field $U_\mu(x)$, the Hamiltonian equations corresponding to

$$\mathscr{H} = \frac{1}{2}\pi^2 + S[U] \tag{1.41}$$

are formed and integrated along a surface of constant energy in fictitious "molecular dynamics" (MD) time, labeled τ. Numerical errors in the finite precision integration are corrected stochastically with a Metropolis accept/reject step: after a fixed time interval $\Delta\tau$ (an *MD trajectory*) the total change in the Hamiltonian is computed, and the current gauge field is accepted as the next step in the Markov chain with probability $P = \min(1, e^{-\Delta\mathscr{H}})$. This ensures that the algorithm

remains exact even if inexact numerical integration techniques are used to evolve the Hamiltonian system. Ergodicity is achieved by picking a new direction in the phase space $\{(\pi, U)\}$ at the start of each trajectory (*heatbath refreshment*). HMC is in some ways an unusual application of numerical integration techniques to a dynamical system, since keeping $\Delta\mathscr{H}$ arbitrarily small is an inefficient strategy. Instead, we aim for $\Delta\mathscr{H} \sim \mathcal{O}(1)$, ensuring reasonable acceptance while still minimizing the time required to generate a new gauge field configuration.

HMC is observed to work well in practice, and is now a standard lattice technique. Improving the efficiency of HMC for simulations involving single quark flavors is the subject of Chapter 6.

1.3.5 Measuring Correlation Functions

After generating a Markov chain $\{U_\mu^{(i)}\}_{i=1}^N$ of gauge field configurations n-point correlation functions can then be computed on each configuration and used to extract physical observables of interest. In this section we will consider computing the mass of the π^+ meson from the $\langle \pi^- \pi^+ \rangle$ correlation function as a representative example of the general procedure. This produces a sequence of estimates $\{m_{\pi^+}^{(i)}\}_{i=1}^N$ determining m_{π^+} up to a statistical uncertainty which can be estimated using resampling techniques such as the jackknife or bootstrap [25]. In addition, binning — *i.e.* averaging over blocks of measurements in MD time — can be used to remove the effects of autocorrelations. The statistical uncertainty decays with the number of independent configurations sampled as $\sigma_{\langle m \rangle} \sim 1/\sqrt{N}$. The appropriate bin size is typically estimated by computing the *integrated autocorrelation time* [25].

The π^+ two-point function is computed using an interpolating operator which creates states with the same quantum numbers as the $|\pi^+\rangle$ state. A local operator

$$\mathscr{O}_{\pi^+}(x) = \overline{u}(x)\gamma_5 d(x) \tag{1.42}$$

is the simplest and most obvious choice, but one can consider more general, non-local operators

$$\mathscr{O}_{\pi^+}(\vec{p}, \vec{q}, t) = \sum_{\vec{x},\vec{y}} e^{i(\vec{p}\cdot\vec{x} + \vec{q}\cdot\vec{y})} f(\vec{x}, \vec{y}) \left(\overline{u}(\vec{x}, t) g[U] \gamma_5 d(\vec{y}, t)\right) \tag{1.43}$$

and tune to better couple to a particular, desired state. Here $f(\vec{x}, \vec{y})$ is an arbitrary dimensionless weighting function, and $g[U]$ represents an appropriate product of gauge links to make the

interpolating operator gauge invariant. Standard Wick contractions [1] can be used to relate the $\langle \overline{\mathscr{O}}_{\pi^+}(x)\mathscr{O}_{\pi^+}(y)\rangle$ two-point function to a spin-color trace involving a product of quark propagators and gamma matrices. In position space, for the local operator:

$$\left\langle \overbrace{d(x)\gamma_5\overline{u}(x)}\overbrace{u(y)\gamma_5\overline{d}(y)}\right\rangle = \mathrm{tr}\left[\mathscr{D}_d^{-1}(x-y)\gamma_5\mathscr{D}_u^{-1}(y-x)\gamma_5\right]. \tag{1.44}$$

On a typical lattice the full quark propagator $\mathscr{D}_q^{-1}(x-y)$ is too expensive to compute directly, so one instead computes *lattice propagators* $\psi(x)$ by numerically inverting

$$\sum_y \mathscr{D}(x,y)\psi(y) = \eta(x) \tag{1.45}$$

for a number of different sources $\eta(x)$, and uses these to construct an estimate to the right-hand side of Equation (1.44). The source itself can also be tuned to increase the signal for particular states. Common choices in the literature include *point*, *wall*, and *box* sources, which set

$$\eta(x) = \begin{cases} \delta_{\alpha\beta}\,\delta_{ab}, & x \in V \\ 0, & \text{otherwise} \end{cases} \tag{1.46}$$

for different choices of the (hypercubic) subvolume V, where $\delta_{\alpha\beta}$ and δ_{ab} are spin and color delta functions, respectively. Computing a single lattice propagator actually requires 12 inversions of the Dirac operator, one for each of the 3×4 color and spin components.

After forming the $\langle \pi^+\pi^+\rangle$ correlation function, the π^+ mass can be extracted by fitting to its Euclidean time dependence. To derive this analytically, we begin with the two-point correlation function and insert a normalized sum over eigenstates of the Hamiltonian, $1 = \sum_{n=0}^{\infty} |n\rangle\langle n| / 2E_nV$:

$$\begin{aligned} \mathcal{C}_{\pi^+}(\Delta t) &= \left\langle \overline{\mathcal{O}}_{\pi^+}(t)\mathcal{O}_{\pi^+}(\tau)\right\rangle \\ &= \sum_{n=0}^{\infty} \frac{1}{2E_nV}\left\langle 0\left|e^{-\mathscr{H}t}\overline{\mathscr{O}}_{\pi^+}(0)\right|n\right\rangle\left\langle n\left|e^{\mathscr{H}\tau}\mathscr{O}_{\pi^+}(0)\right|0\right\rangle, \\ &= \frac{|Z_0|^2}{2m_{\pi^+}V}e^{-m_{\pi^+}\Delta t} + \sum_{n=1}^{\infty}\frac{|Z_n|^2}{2E_nV}e^{-E_n\Delta t} \end{aligned} \tag{1.47}$$

where $Z_n \equiv \langle n|\mathcal{O}_{\pi^+}|0\rangle$ and $\Delta t \equiv |t-\tau|$. At large time separations $\Delta t \gg 1$ this sum is dominated by the contribution from the lowest energy state. To extract m_{π^+} we identify a suitable window in Δt where this approximation is valid, and fit an exponential function. This procedure is aided by

plots of the *effective mass*

$$m_\pi^{\text{eff}}(\Delta t) \equiv \log\left(\frac{\mathcal{C}_{\pi^+}(\Delta t)}{\mathcal{C}_{\pi^+}(\Delta t+1)}\right), \tag{1.48}$$

which exhibit a *plateau* in the region dominated by the ground state. On a lattice with a finite temporal direction one observes *around-the-world contamination*: states may propagate through one end of the lattice and return through the other end, as prescribed by the boundary conditions. This is easily taken into account by adding a backward-propagating contribution to the exponential ansatz

$$\mathcal{C}_{\pi^+}(\Delta t) \overset{\Delta t \gg 1}{\simeq} \frac{|Z_0|^2}{2m_{\pi^+}V}\left(e^{-m_{\pi^+}\Delta t} + e^{-m_{\pi^+}(T-\Delta t)}\right), \tag{1.49}$$

and likewise modifying the effective mass

$$m_\pi^{\text{eff}}(\Delta t) \equiv \cosh^{-1}\left(\frac{\mathcal{C}_{\pi^+}(\Delta t+1)+\mathcal{C}_{\pi^+}(\Delta t-1)}{2\,\mathcal{C}_{\pi^+}(\Delta t)}\right). \tag{1.50}$$

While we have concentrated on the π^+ mass as a representative example of lattice methods, the same techniques can be extended to extract a variety of masses, form factors, couplings, and matrix elements by fitting an appropriate analytic form to a Euclidean n-point function computed on the lattice. One can even extract information about hadron scattering through the Lüscher formalism, which relates the Euclidean space, finite volume spectrum to infinite volume, Minkowski spacetime scattering parameters [26, 27].

1.3.6 Partial Quenching

Since fermion actions are typically bilinear, it is generally possible to integrate out the quark fields and write the lattice QCD path integral in the form

$$\mathcal{Z} = \int \mathscr{D}U \det\left(\mathcal{M}[U]\right) e^{-S_{\text{eff}}[U]}. \tag{1.51}$$

In early lattice simulations the effects of dynamical fermion loops were neglected by explicitly setting $\det(\mathcal{M}) = 1$, enormously decreasing the cost of Monte Carlo simulations. This was known as the *quenched approximation*, and while it made early lattice studies more tractable, it also introduced uncontrolled systematic errors, and is rarely used in modern calculations.

Partial quenching is a loosely related technique in which the valence quark masses — *i.e.* the masses used when the Dirac operator is inverted to compute lattice propagators and form correlation functions (Section 1.3.5) — are taken to be different from the sea quark masses — *i.e.* the masses of virtual quarks appearing in closed loops — entering into the fermion determinant and thus the generation of gauge field configurations via the HMC algorithm (Section 1.3.4). Since ensemble generation is, in general, more expensive than computing lattice propagators, partially quenched calculations allow computational costs to be reduced by using light valence quark masses and heavier sea quark masses. Unlike the quenched approximation, this can be done in a controlled manner, often by interpolation and/or extrapolation of the valence and sea quark mass dependence of physical observables to the physical point. Partially quenched QCD may be regarded as a more general theory in its own right, which reduces to QCD in the limit of equal valence and sea quark masses.

1.4 Summary of Lattice Ensembles

This thesis makes use of a number of domain wall fermion lattice QCD simulations performed by the RBC/UKQCD collaboration, and introduced in Refs. [6, 28–31]. We briefly summarize the properties of these ensembles in Figure 1.3 by plotting the pion mass against the square of the lattice spacing in physical units. More detail can be found by consulting the aforementioned references.

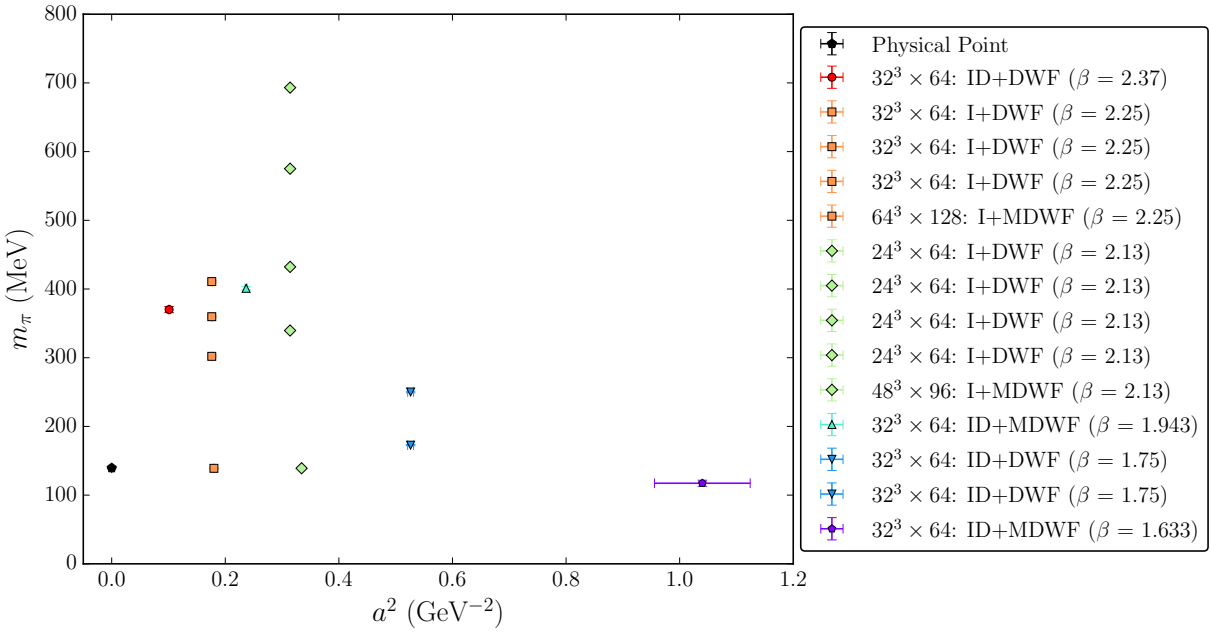

Figure 1.3: Summary of the RBC/UKQCD domain wall fermion ensembles used in this thesis. In the legend the notation I (ID) denotes the Iwasaki (Iwasaki+DSDR) gauge action, and (M)DWF denotes the (Möbius) domain wall fermion action used for the quarks. The values of the pion masses and lattice spacings in physical units are taken from the most recent chiral fits of Ref. [31].

References

[1] M. Peskin and D. Schroeder, *An Introduction To Quantum Field Theory*, Frontiers in Physics (Avalon Publishing, 1995).

[2] C. Patrignani and P. D. Group, "Review of particle physics", Chinese Physics C **40**, 100001 (2016).

[3] R. Streater and A. Wightman, *Pct, spin and statistics, and all that* (Princeton University Press, 1989).

[4] H. J. Rothe, *Lattice Gauge Theories: An Introduction*, World Scientific Lecture Notes in Physics (World Scientific, Singapore, 2012).

[5] Y. Iwasaki and T. Yoshie, "Renormalization Group Improved Action for $SU(3)$ Lattice Gauge Theory and the String Tension", Phys.Lett. **B143**, 449 (1984).

[6] R. Arthur et al., "Domain Wall QCD with Near-Physical Pions", Phys.Rev. **D87**, 094514 (2013).

[7] H. B. Nielsen and M. Ninomiya, "Absence of Neutrinos on a Lattice: (I). Proof by Homotopy Theory", Nucl. Phys. **B185**, [,533(1980)], 20 (1981).

[8] H. B. Nielsen and M. Ninomiya, "A No-Go Theorem for Regularizing Chiral Fermions", Phys. Lett. **B105**, 219–223 (1981).

[9] Y. Shamir, "Chiral Fermions from Lattice Boundaries", Nucl. Phys. **B406**, 90–106 (1993).

[10] V. Furman and Y. Shamir, "Axial Symmetries in Lattice QCD with Kaplan Fermions", Nucl. Phys. **B439**, 54–78 (1995).

[11] D. B. Kaplan, "A Method for Simulating Chiral Fermions on the Lattice", Phys. Lett. **B288**, 342–347 (1992).

[12] T. Blum et al., "Quenched Lattice QCD with Domain Wall Fermions and the Chiral Limit", Phys. Rev. **D69**, 074502 (2004).

[13] G. Fleming, "Finite Temperature QCD with Domain Wall Fermions", PhD thesis (Columbia University, 2001).

[14] R. C. Brower, H. Neff, and K. Orginos, "Mobius Fermions: Improved Domain Wall Chiral fermions", Nucl. Phys. Proc. Suppl. **140**, [,686(2004)], 686–688 (2005).

[15] R. Brower, H. Neff, and K. Orginos, "Möbius Fermions", Nucl.Phys.Proc.Suppl. **153**, 191–198 (2006).

[16] R. C. Brower, H. Neff, and K. Orginos, "The Möbius Domain Wall Fermion Algorithm", (2012).

[17] T.-W. Chiu, "Optimal Lattice Domain-Wall Fermions", Phys. Rev. Lett. **90**, 071601 (2003).

[18] G. McGlynn, "Algorithmic Improvements for Weak Coupling Simulations of Domain Wall Fermions", PoS **LAT2015**, 019 (2015).

[19] G. McGlynn, "Advances in Lattice Quantum Chromodynamics", PhD thesis (Columbia University, 2016).

[20] M. Lüscher, "Construction of a Selfadjoint, Strictly Positive Transfer Matrix for Euclidean Lattice Gauge Theories", Commun. Math. Phys. **54**, 283 (1977).

[21] C. T. Sachrajda and G. Villadoro, "Twisted Boundary Conditions in Lattice Simulations", Phys. Lett. **B609**, 73–85 (2005).

[22] T. D. Lee and C.-N. Yang, "Charge Conjugation, a New Quantum Number G, and Selection Rules Concerning a Nucleon Anti-Nucleon System", Nuovo Cim. **10**, [,233(1956)], 749–753 (1956).

[23] C.-h. Kim and N. H. Christ, "$K \to \pi\pi$ Decay Amplitudes from the Lattice", Nucl. Phys. Proc. Suppl. **119**, [,365(2002)], 365–367 (2003).

[24] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003).

[25] L. Lellouch, R. Sommer, B. Svetitsky, A. Vladikas, and L. Cugliandolo, eds., *Modern Perspectives in Lattice QCD: Quantum Field Theory and High Performance Computing. Proceedings, International School, 93rd Session, Les Houches, France, August 3-28, 2009* (2011).

[26] M. Lüscher, "Signatures of Unstable Particles in Finite Volume", Nuclear Physics B **364**, 237–251 (1991).

[27] M. Lüscher, "Two Particle States on a Torus and their Relation to the Scattering Matrix", Nucl. Phys. **B354**, 531–578 (1991).

[28] C. Allton et al., "Physical Results from 2+1 Flavor Domain Wall QCD and $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D78**, 114509 (2008).

[29] Y. Aoki et al., "Continuum Limit Physics from $2 + 1$ Flavor Domain Wall QCD", Phys. Rev. **D83**, 074508 (2011).

[30] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[31] P. A. Boyle et al., "Low Energy Constants of $SU(2)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD", Phys. Rev. **D93**, 054502 (2016).

Chapter 2

Physics Goals

We discuss the physics goals of this thesis, briefly describing the context and content of the publications following in subsequent chapters. The common thread unifying each of these major results is the development or application of lattice QCD methodology to precision light flavor physics. We focus on first-principles determinations of the low energy constants of chiral perturbation theory (Section 2.1), and of hadronic matrix elements describing weak decays of a kaon to a single-pion (Section 2.2) or two-pion (Section 2.3) final state.

My work, in many cases, has been performed in the context of large, collaborative projects of the RIKEN-Brookhaven-Columbia (RBC) and UKQCD collaborations. For clarity, I briefly summarize my original contributions to each of the projects discussed in this thesis:

1. **Next-to-Next-to Leading Order Chiral Perturbation Theory:**

   - Implemented continuum and finite volume next-to leading order $SU(3)$ PQχPT expressions for the pseudoscalar masses and decay constants in the RBC/UKQCD code base for performing chiral/continuum fits
   - Extended RBC/UKQCD code base to call Johan Bijnens' Fortran libraries of next-to-next-to leading order continuum $SU(2)$ and $SU(3)$ PQχPT formulas, including checks for numerical stability
   - Extracted low energy spectrum on $32^3 \times 64$ $\beta = 1.633$ and $\beta = 1.943$ Iwasaki+DSDR

Möbius DWF ensembles [1]

- Performed and analyzed fits of the pseudoscalar masses and decay constants from the full RBC/UKQCD data set (Figure 1.3) to NLO and NNLO $SU(2)$ and $SU(3)$ PQχPT [1–4]

2. **Leptonic and Semileptonic Kaon Decays:**

   - Extracted the low energy spectrum [5] and $K_{\ell 3}$ form factors [6] from fits to Euclidean two- and three-point correlation functions computed on the physical quark mass $48^3 \times 96$ and $64^3 \times 128$ Möbius domain wall fermion ensembles
   - Performed RBC's chiral and continuum extrapolation of the $K_{\ell 3}$ form factors; independently checked by members of UKQCD
   - Extensively studied fit systematics and parametrizations for the chiral / continuum extrapolation

3. ***K* $\rightarrow \pi\pi$ Decays and the Exact One Flavor Algorithm:**

   - Developed and coded the implementation of EOFA in the RBC/UKQCD code bases (BFM, CPS, Grid) [7]
   - Introduced novel preconditioning technique which significantly accelerates the EOFA algorithm
   - Tuned EOFA for the current $\Delta I = 1/2$ $K \rightarrow \pi\pi$ production ensemble generation calculation, achieving a factor of 4.2 speed-up in the time required to generate an independent gauge field configuration

2.1 Chiral Perturbation Theory

2.1.1 Effective Field Theories and QCD

Lattice QCD, as we have argued in Chapter 1, provides the only known first-principles method for performing fully non-perturbative calculations of hadron properties in terms of the interactions

between constituent quarks. The *effective field theory* (EFT) formalism provides a powerful, orthogonal approach: since it is too difficult at low energies to work analytically in terms of quarks and gluons, we can instead use symmetry principles to write down a theory — matched to QCD — whose fundamental degrees of freedom are the hadrons themselves. The pre-Standard Model Fermi theory of β decay

$$n \to p + e^- + \overline{\nu}_e \tag{2.1}$$

can be considered a prototypical example of an EFT, obtained by integrating out the $W^\pm$ bosons. At energy scales $E \ll m_W$ Fermi's description of β decay in terms of a point-like interaction between hadrons and leptons accurately describes experimental results. It is not until one reaches the threshold for creating $W^\pm$ bosons, $E \sim m_W \sim 80$ GeV, that the Fermi theory breaks down, and a more fundamental description of β decay in terms of the Standard Model is necessary. One might naturally wonder if similar ideas can be successfully applied to low energy QCD.

More generally, the EFT formalism follows from Weinberg's famous "folk theorem" [8]:

> *This remark is based on a "theorem", which as far as I know has never been proven, but which I cannot imagine could be wrong. The "theorem" says that although individual quantum field theories have of course a good deal of content, quantum field theory itself has no content beyond analyticity, unitarity, cluster decomposition, and symmetry. This can be put more precisely in the context of perturbation theory: if one writes down the most general possible Lagrangian, including all terms consistent with assumed symmetry principles, and then calculates matrix elements with this Lagrangian to any given order of perturbation theory, the result will simply be the most general possible S-matrix consistent with analyticity, perturbative unitarity, cluster decomposition and the assumed symmetry principles.*

For the effective theory to have predictive power one must also have a separation of scales $E \ll \Lambda$, and a power counting scheme for arranging contributions to a given matrix element from individual terms in the Lagrangian into a perturbative expansion in powers of the small ratio E/Λ.

We noted, in Section 1.1, that the QCD Lagrangian with N_f quark flavors has a global $SU(N_f)_L \otimes SU(N_f)_R$ symmetry in the massless limit. We also noted that, in nature, this symmetry is spontaneously broken down to the $SU(N_f)_V$ subgroup, as evidenced by the nonzero vacuum

expectation values of the quark condensates $\langle\overline{\psi}_f\psi_f\rangle \neq 0$. The nonzero quark masses further explicitly break this symmetry, giving rise to $N_f^2 - 1$ pseudo Nambu-Goldstone bosons (pNGBs), which should be light if $SU(N_f)_L \otimes SU(N_f)_R$ is indeed an approximate symmetry of QCD. For $N_f = 2$ this is unambiguously the case, as one can observe in Table 1.2: the charged and neutral pion masses differ by only a few percent, and are nearly an order of magnitude lighter than the lightest baryons. Similarly, the approximate $N_f = 3$ symmetry corresponds to the pseudoscalar octet (π, K, η), but in this case one observes that the breaking of the degeneracy between, for example, the pions and the kaons is considerably larger than between the charged and neutral pions, as is the separation of scales between the masses of the kaons or η meson and the mass of the proton, with $m_{K^0}/m_{p^+} \sim 0.5$. For $N_f \geq 4$ the "approximate" symmetry is too poor to be of any phenomenological use[1].

Armed with an approximate symmetry of low-energy QCD — $SU(N_f)_L \otimes SU(N_f)_R$ — and evidence of a separation of scales — $m_q/\Lambda_{\text{QCD}} \ll 1$ for the up, down, and possibly strange quarks — we can set about following Weinberg's prescription for constructing an effective field theory of the pseudoscalar mesons. This theory is known as *chiral perturbation theory* (ChPT or χPT).

2.1.2 The Chiral Perturbation Theory Lagrangian

Following Weinberg, we aim to write down the most general Lagrangian for the pseudoscalar mesons invariant under $SU(N_f)_L \otimes SU(N_f)_R$ transformations. We separately consider the cases $N_f = 2$ and $N_f = 3$, describing the pions and the pseudoscalar octet (π, K, η), respectively. We will only outline the construction, since chiral perturbation theory is a well-developed subject and the details can be found in a number of textbooks [9] and review articles [10–12].

The first step in the construction of the χPT Lagrangian is to write down a parametrization of the pNGB fields. The symmetry breaking pattern $SU(N_f)_L \otimes SU(N_f)_R \to SU(N_f)_V$ results in a vacuum state invariant under $SU(N_f)_V$ transformations, and $N_f^2 - 1$ pNGBs associated with elements of the quotient group $SU(N_f)_L \otimes SU(N_f)_R / SU(N_f)_V$. One typically chooses the *exponential*

[1] One can also frame this heuristic argument in terms of the QCD scale $\Lambda_{\text{QCD}} \sim 200$ MeV: $\frac{1}{2}(m_u + m_d)/\Lambda_{\text{QCD}} \sim 0.02$ points to a large separation of scales, and $m_s/\Lambda_{\text{QCD}} \sim 0.5$ might still be considered a separation of scales, but $m_q/\Lambda_{\text{QCD}} > 1$ for the charm, bottom, and top quarks.

representation[2]

$$U(x) \equiv \exp\left(\frac{i\phi(x)}{f}\right), \tag{2.2}$$

with

$$\phi(x) = \sum_{i=1}^{3} \tau_i \phi_i(x) = \begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 & \pi^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 \end{pmatrix} \tag{2.3}$$

for the $N_f = 2$ theory and

$$\phi(x) = \sum_{a=1}^{8} \lambda_a \phi_a(x) = \begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta & \pi^+ & K^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta & K^0 \\ K^- & \overline{K}^0 & -\frac{2}{\sqrt{6}}\eta \end{pmatrix} \tag{2.4}$$

for the $N_f = 3$ theory. The utility of the exponential representation is that the U field transforms under $SU(N_f)_L \otimes SU(N_f)_R$ in a simple way

$$U(x) \to RU(x)L^\dagger \tag{2.5}$$

where $L \in SU(N_f)_L$ and $R \in SU(N_f)_R$ are independent local transformations. The ground state corresponds to the origin $U_0 = \mathbb{1}$ in this parametrization, and is indeed invariant under vector transformations — $L = R = V \in SU(N_f)$ — but not axial transformations — $L = A^\dagger$, $R = A$, $A \in SU(N_f)$ — consistent with the desired symmetry breaking pattern.

The explicit breaking of chiral symmetry by non-vanishing quark masses is included by introducing an additional (constant) operator $\chi = 2BM$, where $M = \mathrm{diag}(m_u, m_d)$ is the $N_f = 2$ mass matrix and $M = \mathrm{diag}(m_u, m_d, m_s)$ is the $N_f = 3$ mass matrix. This new operator transforms under $SU(N_f)_L \otimes SU(N_f)_R$ in the same manner as U:

$$\chi \to R\chi L^\dagger. \tag{2.6}$$

More generally, the vector and axial currents, as well as the scalar and pseudoscalar densities, can be coupled to the χPT Lagrangian as external sources in a manner which elegantly reproduces all of the Ward-Takahashi identities of QCD [13, 14].

In Weinberg's power counting scheme [8] a *chiral order* is assigned to each term in the Lagrangian by counting the number of derivatives of U which enter: $\partial^n U \sim \mathcal{O}(p^n)$, where p corresponds to

[2]Note: a slightly different normalization in terms of $F = f/\sqrt{2}$ is also frequently found in the literature.

the scale of external momenta carried by the pNGBs. A mass term m^2 is counted as $\mathcal{O}(p^2)$. The Lagrangian can be systematically constructed as an expansion in powers of p

$$\mathcal{L}_{\chi\mathrm{PT}} = \mathcal{L}^{(2)}_{\chi\mathrm{PT}} + \mathcal{L}^{(4)}_{\chi\mathrm{PT}} + \mathcal{L}^{(6)}_{\chi\mathrm{PT}} + \cdots \tag{2.7}$$

by writing down all operators $\mathscr{O}_i^{(n)}$ constructed from U, χ, and derivatives, and invariant under $SU(N_f)_L \otimes SU(N_f)_R$, with chiral order $\mathcal{O}(p^n)$:

$$\mathcal{L}^{(n)}_{\chi\mathrm{PT}} = \sum_i L_i \mathscr{O}_i^{(n)}. \tag{2.8}$$

In terms of this expansion, Weinberg's power counting scheme assigns a dimension D to a diagram with N_L loops and N_n insertions of a vertex originating from $\mathcal{L}^{(n)}_{\chi\mathrm{PT}}$ according to the formula [12]

$$D = 2 + 2N_L + \sum_{n=1}^{\infty} (n-2)\, N_n. \tag{2.9}$$

The $L_i \in \mathbb{R}$ are *a priori* unknown coefficients called *low energy constants* (LECs), which encode the matching of the chiral effective theory to QCD. Conventionally one denotes the LECs of the $SU(2)$ theory as l_i and the LECs of the $SU(3)$ theory as L_i to avoid confusion. Since one must write down all possible operators at a given chiral order, the number of LECs quickly explodes as one moves to successively higher orders in the chiral expansion (Table 3.1).

One typically works with χPT by truncating the Lagrangian to NLO or NNLO. This truncated Lagrangian depends on a finite number of LECs, which must be determined by matching χPT calculations to experimental data or to lattice simulations. Crucially, the truncated Lagrangian can also be renormalized with a finite number of counter terms, and loop divergences absorbed into the LECs, such that they depend on the choice of a renormalization scale μ. One typically chooses dimensional regularization since it respects chiral symmetry. The complete renormalization of the truncated χPT Lagrangian to NLO [13, 14] and NNLO [15] has been performed using a background field method and heat kernel techniques. One can also find explicit RG equations for the running of the LECs in the same references.

While we argued heuristically in Section 2.1.1 that there is evidence of a separation of scales between the masses of the lightest pseudoscalar mesons and a "typical" QCD mass — such as the mass of the proton — identifying the expansion parameter corresponding to Weinberg's power

counting scheme is somewhat subtle. A simple method is to consider the coefficients multiplying a generic loop integral: expanding U in terms of the pNGB fields introduces powers of $1/f^2$, and the loop integral itself introduces an overall numerical factor[3] of $1/(4\pi)^2$. We may therefore estimate $\Lambda_\chi \sim 4\pi f \sim 1.5$ GeV as the *chiral scale*, and regard χPT as a dual, asymptotic expansion in powers of p/Λ_χ and $m_{\text{pNGB}}/\Lambda_\chi$.

2.1.3 Example: The Pion Mass at Leading and Next-to Leading Order

As an explicit example, we consider computing the pion mass up to NLO in $SU(2)$ χPT with degenerate up and down quark masses $m_u = m_d \equiv m_l$. The $\mathcal{O}(p^2)$ Lagrangian is

$$\mathcal{L}^{(2)}_{\chi\text{PT}} = c_1 \operatorname{tr}\left(\partial_\mu U \partial^\mu U^\dagger\right) + c_2 \operatorname{tr}\left(\chi U^\dagger + U\chi^\dagger\right), \tag{2.10}$$

with the values of c_1 and c_2 fixed in terms of B and f by the requirement that the kinetic and mass terms of the pion fields are canonically normalized. Expanding

$$U(x) = \sum_{n=0}^{\infty} \frac{1}{n!} \left(\frac{i}{f}\right)^n \phi^n \tag{2.11}$$

and inserting this sum into $\mathcal{L}^{(2)}_{\chi\text{PT}}$, keeping terms up to $\mathcal{O}(\phi^2)$, we obtain

$$\mathcal{L}^{(2)}_{\chi\text{PT}} \supset \frac{1}{2}\left(\partial_\mu \pi^0\right)\left(\partial^\mu \pi^0\right) + \frac{1}{2}\left(2Bm_l\right)\left(\pi^0\right)^2 \tag{2.12}$$

with $c_1 = c_2 = f^2/4$. Thus, at leading order,

$$m_\pi^2 = 2Bm_l + \mathcal{O}(p^4). \tag{2.13}$$

Lattice QCD calculations demonstrate that the leading order χPT prediction $m^2_{\text{pNGB}} \propto m_q$, known as the Gell-Mann-Oakes-Renner relation, is surprisingly accurate over a wide range of quark masses.

At NLO the calculation is more involved; we will simply describe the method and refer the reader to Ref. [9] for the details. From the power counting formula of Equation (2.9) we see that there are two ways to construct a diagram which contributes at $\mathcal{O}(p^4)$: a one loop diagram constructed from $\mathcal{L}^{(2)}_{\chi\text{PT}}$, or a tree diagram constructed from $\mathcal{L}^{(4)}_{\chi\text{PT}}$ (Figure 2.1). The latter contribution can be

[3] See, for example, Appendix A.4 of Ref. [16].

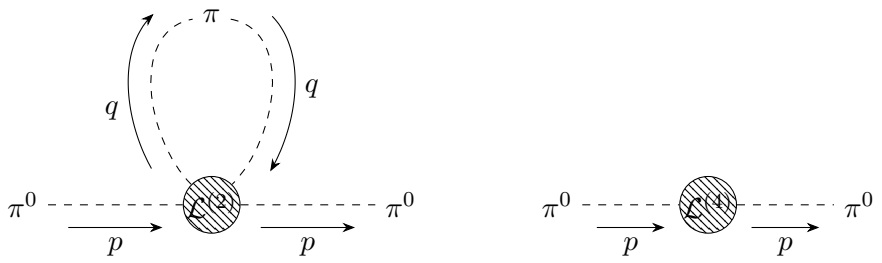

(a) One loop correction from $\mathcal{L}^{(2)}_{\chi\mathrm{PT}}$ (b) Tree level correction from $\mathcal{L}^{(4)}_{\chi\mathrm{PT}}$

Figure 2.1: Corrections to m_π^2 at NLO in $SU(2)$ chiral perturbation theory. The particle in the loop can be any member of the pion triplet $\pi \in \{\pi^-, \pi^0, \pi^+\}$.

straightforwardly calculated from the $\mathcal{O}(p^4)$ Lagrangian [14]

$$\begin{aligned}\mathcal{L}^{(4)}_{\chi\mathrm{PT}} = l_1 \left[\mathrm{tr}\left(\partial_\mu U \partial^\mu U^\dagger\right)\right]^2 + l_2 \,\mathrm{tr}\left(\partial_\mu U \partial_\nu U^\dagger\right)\mathrm{tr}\left(\partial^\mu U \partial^\nu U^\dagger\right) \\ + l_3 \left[\mathrm{tr}\left(\chi U^\dagger + U\chi^\dagger\right)\right]^2 + h_1 \,\mathrm{tr}\left(\chi^\dagger \chi\right),\end{aligned} \tag{2.14}$$

where we have intentionally excluded terms involving external sources or which vanish in the isospin-symmetric limit, by expanding U in powers of the pion fields up to $\mathcal{O}(\phi^2)$. In particular, the operators proportional to l_3 generate terms with the structure $m^4\phi^2$. Calculating the loop correction is more involved: expanding $\mathcal{L}^{(2)}_{\chi\mathrm{PT}}$ to $\mathcal{O}(\phi^4)$ generates a four-point vertex from operators of the form $\phi^2(\partial\phi)^2$ and $m^2\phi^4$. After working out the Feynman rule for this vertex, one can then compute the loop diagram *e.g.* in dimensional regularization, isolating the finite part and absorbing the divergence into renormalized LECs defined at a scale μ. The final result for the pion mass is [9]

$$m_\pi^2 = \chi_l \left\{1 + \frac{4\chi_l}{f^2} l_3 + \frac{\chi_l}{16\pi^2 f^2} \log\left(\frac{\chi_l}{\mu^2}\right)\right\} + \mathcal{O}(p^6), \tag{2.15}$$

where $\chi_l \equiv 2Bm_l$ is the leading order expression.

At next-to leading order and beyond the corrections include analytic terms which are products of LECs and quark masses, as well as non-analytic *chiral logarithms* — such as the $m_q^2 \log(m_q)$ term entering into m_π^2 at NLO — that arise from loop corrections and do not introduce new LECs. While the analytic terms can be reproduced from a simple Taylor expansion in the quark masses, the chiral logarithms are a unique prediction of the full machinery of χPT. We also note that since both the renormalized LECs and the logarithms depend on an arbitrary choice of the renormalization scale

μ, one may vary the relative sizes of these terms by varying μ. The complete expression at a given chiral order, including both types of contributions, however, is independent of μ.

2.1.4 Partially Quenched Chiral Perturbation Theory and Other Extensions

In Section 1.3.6 we introduced the idea of a partially quenched lattice calculation, in which the valence quarks entering into fermion propagators and the sea quarks entering into virtual quark loops are allowed to have different masses. Partial quenching can, in some cases, be used to reduce the costs of lattice calculations with controlled systematics that vanish in the *unitary limit* of equal valence and sea quark masses. These systematics can also be described in a natural way within chiral perturbation theory, leading to an extension known as *partially quenched chiral perturbation theory* (PQChPT or PQχPT).

In the framework of PQχPT the effects of partial quenching are included analytically by generalizing to a supersymmetric theory with N_{sea} and N_{val} sea and valence quarks, respectively. In addition, the theory contains N_{val} unphysical bosonic ghost quarks, which are introduced to exactly cancel the contributions of the valence quarks to closed fermion loops. The $SU(N_f)_L \otimes SU(N_f)_R$ chiral symmetry of massless QCD is promoted to a graded $SU(N_{\text{val}}+N_{\text{sea}}|N_{\text{val}})_L \otimes SU(N_{\text{val}}+N_{\text{sea}}|N_{\text{val}})_R$ supersymmetry, and the most general effective Lagrangian consistent with this enhanced symmetry is constructed order-by-order, in analogy to χPT. In this thesis we make use of the full next-to-next-to leading order expressions for the partially quenched psuedoscalar meson masses and decay constants computed by Bijnens et al. [17–20].

As an example, we may again consider the pseudoscalar meson mass at next-to leading order in PQχPT with $N_{\text{val}} = N_{\text{sea}} = 2$. Assuming degenerate sea quark masses m_l, but non-degenerate valence quark masses m_x and m_y, the resulting pseudoscalar meson mass is [21]

$$m_{xy}^2 = \frac{1}{2}(\chi_x + \chi_y)\left\{1 + \frac{32}{f^2}\chi_l\left(2\hat{L}_6^{(2)} - \hat{L}_4^{(2)}\right) + \frac{8}{f^2}(\chi_x + \chi_y)\left(2\hat{L}_8^{(2)} - \hat{L}_5^{(2)}\right) + \frac{1}{16\pi^2 f^2}\left[\frac{\chi_x - \chi_l}{\chi_x - \chi_y}\chi_x \log\left(\frac{\chi_x}{\mu^2}\right) + \frac{\chi_y - \chi_l}{\chi_y - \chi_x}\chi_y \log\left(\frac{\chi_y}{\mu^2}\right)\right]\right\} + \mathcal{O}(p^6), \tag{2.16}$$

where $\chi_q \equiv 2Bm_q$. We use the notation $\hat{L}_i^{(2)}$ to emphasize that the LECs of $N_{\text{val}} = N_{\text{sea}} = 2$ PQχPT

are not, in general, the same as the LECs l_i of $SU(2)$ χPT. However, by taking the unitary limit and matching to χPT one can derive expressions for the l_i in terms of linear combinations of the $\hat{L}_i^{(2)}$. While Equation (2.16) can be shown to reduce to Equation (2.15) in the limit $m_x = m_y = m_l$, this limit is somewhat subtle, both due to the matching of the partially quenched and unquenched LECs and due to the indeterminate forms of the non-analytic chiral logarithms.

Many other variants of chiral perturbation theory exist in the literature. Of particular interest are variants which describe lattice systematics — such as the use of a finite spacetime volume or discrete lattice spacing — within an effective field theory framework. In Chapters 3 and 4 we will also make use of next-to leading order *finite volume chiral perturbation theory* (FVChPT or FVχPT) [21] to parametrize the dependence of our results on the spatial lattice extent L.

2.1.5 Chiral Perturbation Theory and Lattice QCD

Chiral perturbation theory has been used to parametrize the quark mass dependence of lattice QCD calculations since the earliest days of the field. Until very recently calculations with physical pion masses were simply too expensive to perform even with the most powerful existing supercomputers, forcing the use of unphysical, heavy quark masses. One could still make physical predictions by using χPT as an ansatz to extrapolate simulations with heavier-than-physical quarks down to the physical point. The reliability of these extrapolations was difficult to address, however, and evidence that $SU(3)$ χPT in particular poorly described existing lattice data was presented in Ref. [21].

More recently, it has become possible to simulate QCD directly with physical quark masses. While in some cases χPT has still been used to make modest percent-level corrections for slight mistunings in the input quark masses or finite volume effects — *e.g.* in Ref. [5] — the field has moved away from relying heavily on the machinery of χPT to make physical predictions. In this work we invert the procedure, and instead use the wide range of RBC/UKQCD domain wall QCD simulations (Figure 1.3) to systematically probe χPT at next-to-next-to leading order. These studies have two goals: first, to determine as many of the low energy constants from first-principles as possible using the available lattice data, and second, to study the convergence properties of the (asymptotic) χPT expansion over a wide range of quark masses. The values of the low energy

constants are of great interest in the phenomenological community, since, once they have been determined at a given order, any other processes of interest involving the pseudoscalar mesons can be computed to the same order in χPT to make genuine predictions[4]. More generally, this is an interesting test of the effective field theory formalism: we systematically compare predictions computed non-perturbatively in the full, UV complete theory (QCD) and in its low energy effective description (χPT).

In Chapters 3 and 4 we describe in detail fits of RBC/UKQCD lattice data for pseudoscalar meson masses and decay constants to the more general PQχPT at next-to and next-to-next-to leading order. The generic ansatz we use to describe the lattice data is written schematically in Equation (3.9), and includes the continuum PQχPT expressions up to NNLO, the finite volume corrections computed in NLO FVχPT, and terms $\propto a^2$ describing the leading discretization errors. A complete summary of the lattice data, including partially quenched measurements with non-unitary valence quark masses, can be found in Appendix 3.B. These fits determine 9 of the NLO and 8 linearly independent combinations of the NNLO low energy constants of PQχPT, some of which were previously unknown. The unquenched LECs of χPT are then recovered using the expressions summarized in Appendices 3.A and 4.A. The values we obtain — as well as values from other recent lattice and phenomenological determinations — are summarized in Figures 3.6, 3.7, 4.8, and 4.9 and in Tables 3.10 and 4.9.

By repeating the fits with different subsets of the data we are also able to self-consistently study the behavior of the chiral expansion as a function of the quark masses. We find that the $SU(2)$ expansion is quite robust, satisfying the expected hierarchy LO $\gg$ NLO $\gg$ NNLO at the physical point, and continuing to describe the lattice data up to a heavy scale $m_\pi \sim 450$ MeV before showing obvious distress. We also find that NNLO $SU(3)$ PQχPT can be reliably fit to our data with percent scale accuracy up to the physical point, in contrast to earlier, unsuccessful

[4]This is especially true for calculations which remain intractable on the lattice. Examples of experimental and phenomenological interest include scattering phenomena beyond the 2 particle initial and final state scattering described by the Lüscher formalism — although progress is being made in this direction [22] — and decays with complicated multi-particle final states such as $K_{\ell 4}$.

results[5] from the lattice community. Our conclusions regarding the reliability of the $SU(2)$ and $SU(3)$ expansions at next-to-next-to leading order are consistent with other lattice studies [23–27]. Finally, in Tables 3.9 and 4.8 we have used the large set of LECs determined in the full NNLO fits to make additional predictions from χPT at NLO for meson-meson scattering parameters and QCD isospin breaking effects.

2.1.6 Ongoing Work

One limitation of our current data set is that, while our ensembles span a broad range of light quark masses, all of the strange quark masses are near the physical m_s. We are currently generating a $32^3 \times 64 \times 24$ ensemble with a physical pion mass and a 300 MeV kaon mass to better constrain the strange quark dependence of our $SU(3)$ fits. In addition, in the time since the current fits were performed, Johan Bijnens has developed and released an improved library of one-loop and two-loop calculations in chiral perturbation theory [28]. This library includes some new results — including the full next-to-next-to leading order $SU(3)$ finite volume corrections — which were previously unavailable. We expect the additional ensemble and the two-loop finite volume corrections to improve the quality of our fits to $SU(3)$ PQχPT.

Determining additional low energy constants requires new observables to be included in the fits. We have recently completed a calculation of the $I = 2$ $\pi\pi$ scattering length a_0^2 on the full RBC/UKQCD DWF ensemble set summarized in Figure 1.3. Preliminary results including subsets of this data were presented at Lattice 2015 [2] and at ICHEP 2016 [4], and were found to dramatically improve the accuracy of the $SU(2)$ LECs $\bar{\ell}_1$ and $\bar{\ell}_2$, which contribute to a_0^2 at NLO. In addition, data for the pion vector and scalar form factors and for current-current correlation

[5] In Ref. [21] it was demonstrated that the value of the LEC f_0 obtained from fits to the available lattice data was unreasonably low and inconsistent with other published values, causing the authors to question the applicability of $SU(3)$ χPT altogether. In our current fits we find a more reasonable value of f_0 consistent with the literature. In Section 4.4.6 we demonstrate that by successively removing light data from the fits we can continuously interpolate between our current value of f_0 and the old, inconsistent value, suggesting that these earlier studies were simply applying NLO $SU(3)$ PQχPT at heavy quark masses outside its range of validity, and rightfully concluding that the resulting fits were unreliable.

functions are available on some ensembles, which constrain the remaining next-to leading order LECs $\overline{\ell}_5$ and $\overline{\ell}_6$. We plan to revisit our NNLO $SU(2)$ and $SU(3)$ fits using this additional data and the $SU(3)$-specific improvements discussed in the previous paragraph in the near future.

2.2 Leptonic and Semileptonic Kaon Decays

Leptonic ($\pi \to \ell\overline{\nu}_\ell$ or $K \to \ell\overline{\nu}_\ell$) and semileptonic ($K \to \pi\ell\overline{\nu}_\ell$) pion and kaon decays currently provide the most accurate determinations of the CKM matrix elements V_{ud} and V_{us}, and contribute to stringent tests of the Standard Model through CKM unitarity. The leading order Standard Model processes contributing to the leptonic ($K_{\ell 2}$) and semileptonic ($K_{\ell 3}$) decay modes are summarized by the Feynman diagrams of Figure 2.2. At leading order the decay rate for the leptonic decay of

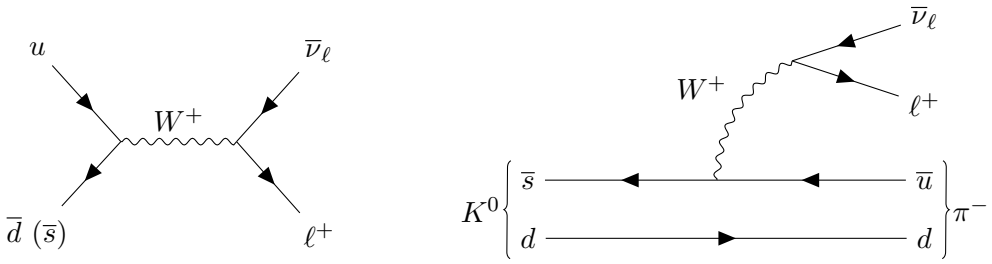

(a) $K_{\ell 2}$: $\pi^+ (K^+) \to \ell^+\overline{\nu}_\ell$ (b) $K_{\ell 3}$: $K^0 \to \pi^-\ell^+\overline{\nu}_\ell$

Figure 2.2: Leading order Standard Model diagrams contributing to the leptonic π^+ and K^+ decays (left) and semileptonic K^0 decay (right).

a pseudoscalar meson (P) composed of quark flavors q and q' is

$$\Gamma(P \to \ell\overline{\nu}_\ell) = \frac{G_F^2}{8\pi} f_P^2 m_\ell^2 m_P \left(1 - \frac{m_\ell^2}{m_P^2}\right)^2 \left|V_{qq'}\right|^2, \tag{2.17}$$

where G_F is the Fermi decay constant, m_ℓ is the lepton mass, m_P is the meson mass, $V_{qq'}$ is a CKM matrix element, and f_P is the pseudoscalar meson decay constant, which encapsulates the strongly-coupled, hadronic part of the decay process. Explicitly, for the pion and kaon:

$$\langle 0|\,\overline{d}\gamma_\mu\gamma_5 u\,\left|\pi^+(\vec{p})\right\rangle = ip_\mu f_{\pi^+}, \quad \langle 0|\,\overline{s}\gamma_\mu\gamma_5 u\,\left|K^+(\vec{p})\right\rangle = ip_\mu f_{K^+}. \tag{2.18}$$

Taking a ratio of the experimentally measured $K^\pm$ and $\pi^\pm$ decay rates leads to a relation between the ratio of decay constants $f_{K^\pm}/f_{\pi^\pm}$ and the ratio of CKM matrix elements $|V_{us}|/|V_{ud}|$, since the Fermi decay constant and masses are accurately known. The current best experimental constraint on this ratio, including an estimate of the $\mathcal{O}(1\%)$ radiative corrections from loop diagrams, is [29]

$$\frac{|V_{us}|}{|V_{ud}|}\frac{f_{K^\pm}}{f_{\pi^\pm}} = 0.27599(29)(24). \tag{2.19}$$

The accurate determination of $f_{K^\pm}/f_{\pi^+}$ from first principles — allowing a clean extraction of the CKM ratio — is a natural target of lattice calculations.

The Standard Model analysis of the $K_{\ell 3}$ decay rate proceeds similarly, but is more complicated due to the three-particle final state. In nature, both the $K^+ \to \pi^0 \ell^+ \overline{\nu}_\ell$ and $K^0 \to \pi^- \ell^+ \overline{\nu}_\ell$ decays are allowed. In anticipation of the lattice calculation, however, we work in the isospin symmetric limit of equal light quark masses $m_u = m_d \equiv m_l$ and neglect electromagnetism, with the small corrections from these effects computed in chiral perturbation theory and reintroduced in a form made explicit below. In this limit the charged and neutral pions are degenerate, as are the charged and neutral kaons. The hadronic $K \to \pi$ part of the decay is parametrized by the $K_{\ell 3}$ form factors $f_\pm^{K\pi}(q^2)$, which are defined in terms of the *vector matrix element*

$$\langle \pi(p_\pi)| \,\overline{s}\gamma^\mu u\, |K(p_K)\rangle \equiv f_+^{K\pi}(q^2)\left(p_K^\mu + p_\pi^\mu\right) + f_-^{K\pi}(q^2)\left(p_K^\mu - p_\pi^\mu\right), \tag{2.20}$$

where $q^\mu \equiv p_K^\mu - p_\pi^\mu$ is the momentum transfer between the kaon and pion. The full $K_{\ell 3}$ decay rate, after reintroducing electromagnetic and isospin breaking corrections, is:

$$\Gamma\left(K \to \pi \ell \overline{\nu}_\ell\right) = \frac{G_F^2 m_K^5}{192\pi^3} C_K^2 S_{\mathrm{EW}} \left|V_{us}\right|^2 \left(f_+^{K\pi}(0)\right)^2 I_{K\ell} \left(1 + \delta_{\mathrm{EM}}^{K\ell} + \delta_{SU(2)}^{K\pi}\right)^2. \tag{2.21}$$

The terms appearing in this formula are:

- C_K: Clebsch-Gordan coefficient, with $C_{K^0} = 1$ and $C_{K^\pm} = 1/\sqrt{2}$
- S_{EW}: short-distance electroweak corrections, computed perturbatively in the Standard Model
- $I_{K\ell}$: phase-space integral, determined experimentally
- $\delta_{\mathrm{EM}}^{K\ell}$: long-distance electromagnetic corrections, computed in χPT and dependent on the final lepton state ℓ

- $\delta^{K\pi}_{SU(2)}$: isospin breaking corrections, computed in χPT

We refer the interested reader to Ref. [30] for additional detail regarding the numerical values and methodology for determining $S_{\rm EW}$, $I_{K\ell}$, $\delta^{K\ell}_{\rm EM}$, and $\delta^{K\pi}_{SU(2)}$.

In addition to the vector matrix element of Equation (2.20), one can also consider the *scalar matrix element*

$$\langle \pi(p_\pi)| \, \bar{s}u \, |K(p_K)\rangle \equiv \frac{m_K^2 - m_\pi^2}{m_s - m_u} f_0^{K\pi}(q^2). \tag{2.22}$$

Acting on Equation (2.20) with q^μ, and applying the Ward-Takahashi identity

$$q^\mu \, \langle V_\mu \rangle = (m_s - m_u) \, \langle \bar{s}u \rangle \tag{2.23}$$

results in a relationship between the three form factors

$$f_0^{K\pi}(q^2) = f_+^{K\pi}(q^2) + \frac{q^2}{m_K^2 - m_\pi^2} f_-^{K\pi}(q^2). \tag{2.24}$$

In particular, at $q^2 = 0$ one has $f_0^{K\pi}(0) = f_+^{K\pi}(0)$, and so $f_+^{K\pi}(0)$ can also be determined directly from the scalar matrix element.

Excluding $f_+^{K\pi}(0)$, the other factors appearing in the analytic expression for the $K_{\ell 3}$ decay rate are either known from experiment or else can be computed using perturbation theory, leading to the experimental constraint [31]

$$|V_{us}| \, f_+^{K\pi}(0) = 0.21654(41). \tag{2.25}$$

Extracting $|V_{us}|$ itself requires knowledge of at least one of the strongly coupled, hadronic matrix element of Equations (2.20) and (2.22), which are most accurately computed using lattice QCD. However, the uncertainty in lattice determinations of $f_+^{K\pi}(0)$ remains the dominant source of error in both $|V_{us}|$ and in the Standard Model unitarity test[6] $\delta_u \equiv 1 - |V_{ud}|^2 - |V_{us}|^2 - |V_{ub}|^2$, making this a topic of continued interest in the field.

[6] If the CKM matrix is indeed unitary, as the Standard Model predicts, then $\delta_u = 0$ up to statistical and systematic uncertainties. Tension with unitarity could potentially signal new physics beyond the Standard Model.

2.2.1 Lattice Calculations of f_π, f_K, and $f_+^{K\pi}(0)$

The extraction of the isospin-symmetric decay constants f_π and f_K at zero momentum from the temporal components of Equation (2.18) is a straightforward calculation on the lattice, since this requires only two-point correlation functions. One minor complication in the domain wall fermion formalism, however, is the need to compute renormalization coefficients for the axial (Z_A) and vector (Z_V) currents. In Ref. [32] it was demonstrated that $Z_A \approx Z_V$ up to small $\mathcal{O}(a^2 m_{\text{res}}^2)$ corrections, so we choose to renormalize both currents using Z_V, with a method discussed below. In addition, the lattice two-point functions contain implicit errors from the finite lattice spacing, finite volume, and typically unphysical quark masses used in the simulations, which must be removed with an extrapolation/interpolation of the lattice data. The most recent RBC/UKQCD result in the continuum, infinite volume, physical quark mass limit

$$\frac{f_K}{f_\pi} = 1.1945(45) \tag{2.26}$$

was presented in Ref. [5]. The 2016 Flavor Lattice Averaging Group (FLAG) review [33] estimates the sub-percent correction from isospin breaking effects as $\delta_{SU(2)} = -0.004(1)$ using chiral perturbation theory. Applying the correction to the RBC/UKQCD result gives

$$\frac{f_{K^\pm}}{f_{\pi^\pm}} = \frac{f_K}{f_\pi}\sqrt{1+\delta_{SU(2)}} = 1.1921(46), \tag{2.27}$$

allowing the ratio of CKM matrix elements $|V_{us}|/|V_{ud}|$ to be extracted from Equation (2.19).

The extraction of the $K_{\ell 3}$ form factor $f_+^{K\pi}(0)$ from the Euclidean three-point correlation functions of the vector current (Eqn. (2.20)) or scalar density (Eqn. (2.22)) is more involved. Since a detailed discussion of the most recent RBC/UKQCD $K_{\ell 3}$ calculation [6] performed with the physical quark mass $48^3 \times 96$ and $64^3 \times 128$ Möbius domain wall fermion ensembles is the subject of Chapter 5, we will simply paraphrase here. This was the first RBC/UKQCD $K_{\ell 3}$ calculation to use the domain wall fermion action with physical quark masses, as well as the first time the calculation was performed directly at zero momentum transfer. This was achieved using a stationary kaon and moving pion, with the pion momentum tuned to satisfy $q^2 = 0$ using twisted boundary conditions (Section 1.3.3):

$$|\vec{p}_\pi| = \frac{m_K^2 - m_\pi^2}{2m_K}. \tag{2.28}$$

We computed three-point correlation functions for both the vector and scalar matrix elements by evaluating the quark line diagram depicted in Figure 5.1. In the diagram, the scalar matrix element corresponds to an insertion of $\Gamma = \mathbb{1}$ and the vector matrix element corresponds to an insertion of $\Gamma = \gamma_\mu$. The upper quark line has momentum $p^i = |\vec{p}_\pi|/\sqrt{3}$ in each of the three spatial directions, allowing improved statistics by averaging the x, y, and z components of the vector matrix element. The vector current renormalization coefficient was computed from the analogous matrix elements $\langle \pi | \overline{u} \gamma_4 u | \pi \rangle$ (Z_V^π) and $\langle K | \overline{s} \gamma_4 s | K \rangle$ (Z_V^K) with stationary initial and final states: charge conservation implies that $f_+^{\pi\pi}(0) = f_+^{KK}(0) = 1$ and $f_-^{\pi\pi}(0) = f_-^{KK}(0) = 0$, allowing Z_V to be determined.

After computing the vector and scalar matrix elements on the $48^3 \times 96$ and $64^3 \times 128$ ensembles, as well as the pion and kaon matrix elements used to determine Z_V, the $K_{\ell 3}$ form factors $f_+^{K\pi}(0)$ and $f_-^{K\pi}(0)$ were determined by fitting to the Euclidean time dependence of Equations (2.20) and (2.22) over an appropriate range of pion-kaon separations and operator insertion times. We then combined this data with older $K_{\ell 3}$ calculations performed on ensembles with heavy pions to make small, $\mathcal{O}(1\%)$ interpolations to correct for slight mistuning in the input quark masses. We observe that next-to leading order $SU(3)$ chiral perturbation theory poorly describes our data, and, lacking sufficient data to constrain the full next-to-next-to leading order expression, use a polynomial in the $SU(3)$-breaking quantity $\Delta m^2 \equiv m_K^2 - m_\pi^2$ instead. This procedure was performed independently for $f_+^{K\pi}(0)$ as determined by the vector matrix element renormalized with Z_V^π and Z_V^K, as well as for $f_+^{K\pi}(0)$ as determined by the scalar matrix element. We then performed a simultaneous extrapolation to a common continuum limit. Our final result was

$$f_+^{K\pi}(0) = 0.9685(34)(14), \tag{2.29}$$

where the first uncertainty is statistical and the second is an estimate of the remaining finite volume error. The details can be found in Chapter 5.

In Figure 2.3 we reprint the most recent FLAG summary plots of theoretical determinations of $f_{K^\pm}/f_{\pi^\pm}$ and $f_+^{K\pi}(0)$. The results of this thesis — framed in yellow — represent the current state-of-the-art for $N_f = 2 + 1$ calculations.

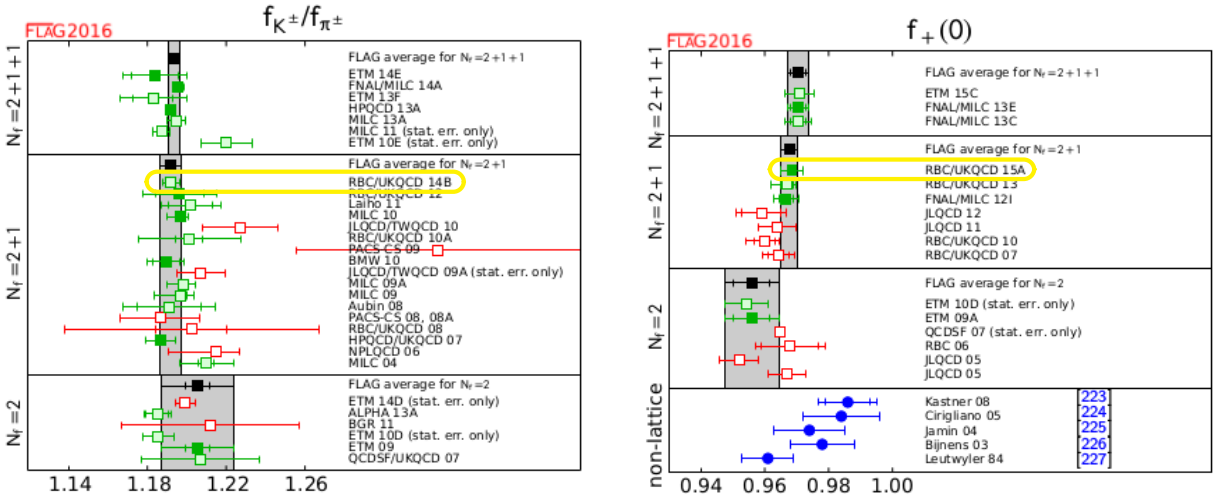

Figure 2.3: Summary plots reprinted from the FLAG 2016 review [33]. The results discussed in this thesis — RBC/UKQCD 14B [5] and RBC/UKQCD 15A [6] — are framed in yellow.

2.2.2 Standard Model Constraints

Combining the final result for $f_+^{K\pi}(0)$ from Equation (2.29) with the experimental constraint of Equation (2.25), we can extract

$$|V_{us}| = 0.22358(41)_{\text{expt}}(85)_{\text{latt}}. \tag{2.30}$$

If we instead use the FLAG global average of $N_f = 2+1$ lattice calculations of $f_+^{K\pi}(0)$ [33] — which includes our result — the value is consistent, with a somewhat reduced lattice error: $|V_{us}| = 0.22378(41)_{\text{expt}}(62)_{\text{latt}}$. While the lattice error is beginning to approach the experimental error in the determination of $|V_{us}|$, it is clear that further refinement is necessary, especially in light of the ongoing KLOE-2 experiment [34] which promises to further tighten the experimental constraints. Further progress will require the inclusion of isospin breaking and electromagnetic effects in the lattice calculations, since these are expected to contribute at a level comparable to the current sub-percent total uncertainty in $f_+^{K\pi}(0)$. The exact one flavor algorithm may prove useful in enabling this next generation of $K_{\ell 3}$ calculations by reducing the cost of performing a complete calculation with physical up and down quark masses.

Having determined $|V_{us}|$, we can then extract $|V_{ud}|$ by combining this result with Equations

(2.19) and (2.27). Using the RBC/UKQCD results for $f_+^{K\pi}(0)$ and $f_{K^\pm}/f_{\pi^\pm}$ as the inputs we obtain

$$|V_{ud}| = 0.9657(22)_{\text{expt}}(52)_{\text{latt}}, \tag{2.31}$$

or $|V_{ud}|= 0.9665(22)_{\text{expt}}(46)_{\text{latt}}$ using the $N_f = 2+1$ FLAG lattice averages. $|V_{ud}|$ can also be determined experimentally from a number of super-allowed nuclear β decays, leading to the significantly more precise result $|V_{ud}|= 0.97417(21)$ after averaging [35]. While the lattice determination is an interesting cross-check, it seems unlikely that the errors could be reduced to this level of precision in the near future.

The last CKM matrix element involving the up quark is $|V_{ub}|\sim 4\times 10^{-3}$ [35]. Given the current uncertainties in $|V_{ud}|$ and $|V_{us}|$, we may neglect $|V_{ub}|$ altogether when testing first-row unitarity through the quantity $\delta_u \equiv 1-|V_{ud}|^2-|V_{us}|^2-|V_{ub}|^2$. Using the RBC/UKQCD results for $|V_{ud}|$ and $|V_{us}|$ we have

$$\delta_u = 0.017(3)_{\text{expt}}(7)_{\text{latt}}, \tag{2.32}$$

or $\delta_u = 0.016(3)_{\text{expt}}(6)_{\text{latt}}$ using the CKM matrix elements determined from the FLAG lattice averages. If we instead use the more precise value of $|V_{ud}|$ from super-allowed β decays we find

$$\delta_u = 0.0010(4)_{\text{expt}}(6)_{\text{latt}} \tag{2.33}$$

using the RBC/UKQCD inputs and $\delta_u = 0.0009(4)_{\text{expt}}(4)_{\text{latt}}$ using the FLAG inputs. Regardless of the choice of inputs we observe a mild (1-2)σ tension with CKM unitarity.

2.3 $K \to \pi\pi$ Decays and the Exact One Flavor Algorithm

The final result of this thesis is the exploration and refinement of a recently proposed algorithm for accelerating the hybrid Monte Carlo simulations used in lattice QCD, known as the exact one flavor algorithm (EOFA) [36–38]. We begin with a discussion of CP-violating neutral kaon decays in the Standard Model, largely following Ref. [9], and briefly review the RBC/UKQCD collaboration's recent first-principles calculations of the $K \to \pi\pi$ decay amplitudes and direct CP-violation parameter ϵ'. We then paraphrase the results of our studies of EOFA: in particular, we

have achieved a factor of 4.2 reduction in the cost of generating a statistically independent gauge field configuration for the ongoing $\Delta I = 1/2$ $K \to \pi\pi$ calculation. We also briefly discuss other potential physics projects which might similarly benefit from EOFA. The details of EOFA and our refinements and benchmarks are the subject of Chapter 6.

2.3.1 The Phenomenology of CP-Violating Kaon Decays in the Standard Model

In the Standard Model, the neutral kaon states $|K^0\rangle$ and $|\overline{K}^0\rangle$ mix through the one-loop "box" diagrams of Figure 2.4. As a result, a beam of neutral kaons will oscillate between $|K^0\rangle$ and $|\overline{K}^0\rangle$

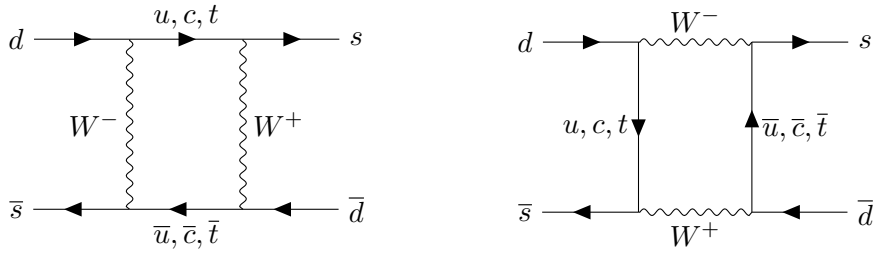

Figure 2.4: Standard Model diagrams contributing to $K^0 - \overline{K}^0$ mixing at lowest order.

as it propagates. In addition, the kaons will tend to decay through the weak interaction. We can model this phenomenologically with a time-dependent, mixed state vector

$$|\psi(t)\rangle = a(t)\left|K^0\right\rangle + b(t)\left|\overline{K}^0\right\rangle \tag{2.34}$$

whose evolution is governed by the dynamical system[7]

$$i\frac{d}{dt}\begin{pmatrix} a(t) \\ b(t) \end{pmatrix} = \begin{pmatrix} M_{11} - \frac{i}{2}\Gamma_{11} & M_{12} - \frac{i}{2}\Gamma_{12} \\ M_{12}^* - \frac{i}{2}\Gamma_{12}^* & M_{11} - \frac{i}{2}\Gamma_{11} \end{pmatrix}\begin{pmatrix} a(t) \\ b(t) \end{pmatrix}. \tag{2.35}$$

The normalized eigenstates found by diagonalizing this system are

$$|K_S\rangle = \frac{1}{\sqrt{1+|\overline{\epsilon}|^2}}\left(\frac{1+\overline{\epsilon}}{\sqrt{2}}\left|K^0\right\rangle + \frac{1-\overline{\epsilon}}{\sqrt{2}}\left|\overline{K}^0\right\rangle\right), \quad |K_L\rangle = \frac{1}{\sqrt{1+|\overline{\epsilon}|^2}}\left(\frac{1+\overline{\epsilon}}{\sqrt{2}}\left|K^0\right\rangle - \frac{1-\overline{\epsilon}}{\sqrt{2}}\left|\overline{K}^0\right\rangle\right), \tag{2.36}$$

[7]The form of the mixing matrix $M - \frac{i}{2}\Gamma$ is partially determined by CPT-invariance and the requirement that it is Hermitian. It can be shown that the former constrains the diagonal elements to be equal, and that the latter forces $M_{11} = M_{11}^*$, $M_{12} = M_{21}^*$, and $M_{22} = M_{22}^*$, and likewise for the matrix elements of Γ.

where we have introduced a new parameter $\bar{\epsilon}$ defined implicitly by the relation

$$\frac{1+\bar{\epsilon}}{1-\bar{\epsilon}} = \left[\frac{M_{12} - \frac{i}{2}\Gamma_{12}}{M_{12}^* - \frac{i}{2}\Gamma_{12}^*}\right]^{1/2}. \tag{2.37}$$

It is more revealing, however, to write these in terms of the CP-eigenstates[8] $K_\pm^0\rangle \equiv \frac{1}{\sqrt{2}}\left(|K^0\rangle \mp |\overline{K}^0\rangle\right)$:

$$|K_S\rangle = \frac{1}{\sqrt{1+|\bar{\epsilon}|^2}}\left(\left|K_+^0\right\rangle + \bar{\epsilon}\left|K_-^0\right\rangle\right), \quad |K_L\rangle = \frac{1}{\sqrt{1+|\bar{\epsilon}|^2}}\left(\bar{\epsilon}\left|K_+^0\right\rangle + \left|K_-^0\right\rangle\right). \tag{2.38}$$

The parameter $\bar{\epsilon}$ is experimentally known to be small but nonzero, and measures the slight discrepancy between the eigenstates of the CP operation and the eigenstates of the weak interaction due to CP-violating decay modes. Assuming exact CP-symmetry, we would have $\bar{\epsilon} = 0$, and conservation of CP would require $|K_+^0\rangle = |K_S\rangle$ to decay exclusively to the CP-even $|\pi\pi\rangle$ final states, and $|K_-^0\rangle = |K_L\rangle$ to decay exclusively to the CP-odd $|\pi\pi\pi\rangle$ final states. In practice, this is approximately true since $\bar{\epsilon}$ is small, and explains why the lifetime of the $|K_S\rangle$ is dramatically shorter than that of the $|K_L\rangle$ — the allowed phase space of the $K_L \to \pi\pi\pi$ decay is significantly more restricted — hence the names "L" (long) and "S" (short). However, one also observes the $K_L \to \pi\pi$ decay to occur. *Indirect CP-violation*, *i.e.* the decay $K_+^0 \to \pi\pi$ through the small $\bar{\epsilon}|K_+^0\rangle$ component of the $|K_L\rangle$ state, was first demonstrated experimentally in the 1960's by Cronin and Fitch, for which they were awarded the 1980 Nobel Prize in Physics. Later experimental work also confirmed that *direct CP-violation*, *i.e.* the explicitly CP-violating decay $K_-^0 \to \pi\pi$, also occurs with nonzero probability.

In the context of the Standard Model, indirect and direct CP-violation are parametrized in terms of the quantities ϵ and ϵ', respectively. These parameters are defined in terms of the experimentally accessible ratios of decay amplitudes $\mathscr{A}(K^i \to \pi^j\pi^k) \equiv \langle\pi^j\pi^k|\mathscr{H}_W|K^i\rangle$, where $\mathscr{H}_W$ is the weak Hamiltonian:

$$\frac{\mathscr{A}(K_L \to \pi^+\pi^-)}{\mathscr{A}(K_S \to \pi^+\pi^-)} \equiv \epsilon + \epsilon', \quad \frac{\mathscr{A}(K_L \to \pi^0\pi^0)}{\mathscr{A}(K_S \to \pi^0\pi^0)} \equiv \epsilon - 2\epsilon'. \tag{2.39}$$

The current best experimental constraints on the magnitudes are $|\epsilon| = 2.228(11)\times10^{-3}$ and $\mathrm{Re}(\epsilon'/\epsilon) = 1.66(23)\times10^{-3}$ [35], indicating that CP-violation is indeed a small effect. For our purposes, we

[8] We adopt the standard phase conventions $CP|K^0\rangle = -|\overline{K}^0\rangle$ and $CP|\overline{K}^0\rangle = -|K^0\rangle$.

would like to invert this, and instead write ϵ and ϵ' in terms of quantities which are accessible on the lattice. Since we work in the isospin symmetric limit[9] $m_u = m_d \equiv m_l$, it is useful to classify the two pion final states $|\pi\pi^I\rangle$ according to isospin: $I = 0$, $I = 1$, or $I = 2$. The $I = 1$ state can be excluded, since it is antisymmetric under interchanging the two pions and thus would violate bosonic symmetry. We parametrize the two allowed decay channels as

$$\left\langle \pi\pi^I \right| \mathscr{H}_W \left| K^0 \right\rangle = \sqrt{2}\mathscr{A}_I e^{i\delta_0^I}, \quad \left\langle \pi\pi^I \right| \mathscr{H}_W \left| \overline{K}^0 \right\rangle = -\sqrt{2}\mathscr{A}_I^* e^{i\delta_0^I}. \tag{2.40}$$

Expressing $|\pi^+\pi^-\rangle$ and $|\pi^0\pi^0\rangle$ in terms of the $|\pi\pi^I\rangle$ states, $|K_L\rangle$ and $|K_S\rangle$ in terms of $|K^0\rangle$ and $|\overline{K}^0\rangle$, and inverting Equation (2.39), working to leading order in $\bar{\epsilon}$ and the imaginary parts of the decay amplitudes, one can show [39]

$$\begin{cases} \epsilon = \bar{\epsilon} + i\dfrac{\mathrm{Im}(\mathscr{A}_0)}{\mathrm{Re}(\mathscr{A}_0)} \\ \epsilon' = \dfrac{i}{\sqrt{2}}\dfrac{\mathrm{Re}(\mathscr{A}_2)}{\mathrm{Re}(\mathscr{A}_0)}\left(\dfrac{\mathrm{Im}(\mathscr{A}_2)}{\mathrm{Re}(\mathscr{A}_2)} - \dfrac{\mathrm{Im}(\mathscr{A}_0)}{\mathrm{Re}(\mathscr{A}_0)}\right) e^{i(\delta_0^2 - \delta_0^0)} \end{cases}, \tag{2.41}$$

where δ_0^I is the isospin I S-wave $\pi - \pi$ scattering phase shift. Since the kaon has isospin 1/2, $\mathscr{A}_2$ is also referred to as the "$\Delta I = 3/2$ amplitude" and $\mathscr{A}_0$ as the "$\Delta I = 1/2$ amplitude".

2.3.2 Computing the $K \to \pi\pi$ Decay Amplitudes on the Lattice

The strategy of the lattice calculation is to compute the $K^0 \to (\pi\pi)_I$ matrix elements of the weak Hamiltonian and the S-wave scattering phase shifts, from which ε and ε' can then be computed using Equation (2.41) and compared to the experimental results. This is a complex and technical calculation, the details of which are somewhat tangential to the novel work performed in this thesis, so we will only paraphrase here. The interested reader can consult Refs. [39–41] for additional detail of the $\Delta I = 3/2$ calculation, and Refs. [42–44] for additional detail of the $\Delta I = 1/2$ calculation.

Typical lattice cutoffs $a^{-1} \sim \mathcal{O}(1-3 \text{ GeV})$ preclude the 80 GeV W boson as one of the simulated degrees of freedom. Instead, one works in an $N_f = 2 + 1$ low-energy effective theory without the W boson and heavy quark flavors, and expands the weak Hamiltonian as

$$\mathscr{H}_W = \frac{G_F}{2} V_{us}^* V_{ud} \sum_{i=1}^{10} C_i(\mu) Q_i(\mu). \tag{2.42}$$

[9]Note: in this limit the charged and neutral pions are degenerate, as are the charged and neutral kaons, so we may unambiguously refer to "the pion" (π) and "the kaon" (K).

In this notation the V_{ij} are CKM matrix elements, the $C_i(\mu)$ are $\overline{\text{MS}}$-renormalized Wilson coefficients computed perturbatively at next-to-leading (NLO) order in the full Standard Model [45], and $\{Q_i(\mu)\}_{i=1}^{10}$ is a basis of ten $\Delta S = 1$ four-quark operators listed explicitly *e.g.* in Ref. [43]. The parameter μ is a renormalization scale: while $\mathscr{H}_W$ is independent of μ, the Wilson coefficients and four-quark operators individually depend on μ. The $K \to (\pi\pi)_I$ decay amplitudes can be computed in terms of three-point lattice correlation functions $\langle \pi\pi^I(x_{\pi\pi})|Q_i(x_Q)|K(x_K)\rangle$. In addition, one also computes the energies of the $I = 0$ and $I = 2$ two-pion states and the masses of the pion and kaon from the two-point correlation functions $\langle \pi\pi^I(t_1)|\pi\pi^I(t_2)\rangle$, $\langle \pi(t_1)|\pi(t_2)\rangle$, and $\langle K(t_1)|K(t_2)\rangle$, respectively.

In addition to computing bare Euclidean space matrix elements on the lattice, one must also match these to the physical matrix elements that define the Standard Model CP-violation parameters ϵ and ϵ'. The Lüscher formalism can be used to directly relate the finite volume energy shift $\delta E_\pi^I \equiv E_{\pi\pi}^I - 2m_\pi$ computed in Euclidean space to the infinite volume, Minkowski space $\pi - \pi$ scattering phase shift $\delta_0^I(p)$ at a particular kinematic point p determined by the ensemble parameters. Relating the bare, Euclidean, finite volume three-point matrix elements $M_i^{\text{lat}} \equiv \langle \pi\pi^I|Q_i|K\rangle$ to their $\overline{\text{MS}}$-renormalized, Minkowski, infinite volume analogues determining the decay amplitudes $\mathscr{A}_I$ is more involved, and proceeds in steps. First, the multiplicative Lellouch-Lüscher factor [46] is used to relate the bare finite volume, Euclidean matrix elements to the bare infinite volume, Minkwoski space matrix elements. One then performs a matching calculation: first the bare matrix elements are renormalized in variants of the non-perturbative regularization independent (RI) scheme, which is then perturbatively matched to $\overline{\text{MS}}$ at a high scale μ where QCD perturbation theory is known to be reliable. Systematic errors are estimated by using multiple variants of the intermediate RI scheme and analyzing the influence on the final $\overline{\text{MS}}$ results. One can then compute the physical decay amplitudes $\mathscr{A}_I$, and ultimately ϵ and ϵ'.

To date the $\Delta I = 3/2$ calculation has been performed on two, independent lattice ensembles with different cutoffs, allowing an additional continuum extrapolation to be performed [41]. The $\Delta I = 1/2$ calculation of Ref. [44] has been performed on a single ensemble and the resulting finite lattice spacing systematic has been estimated as part of the error budget. Both calculations have

physical kinematics and quark masses, and, together, predict $\mathrm{Re}(\epsilon'/\epsilon) = 1.38(5.15)(4.59) \times 10^{-4}$ [44], where the first error is statistical and the second systematic, suggesting a tantalizing 2.2σ discrepancy with the experimental value $\mathrm{Re}(\epsilon'/\epsilon) = 1.66(23) \times 10^{-3}$ [35]. Reducing the error and enabling a more precise comparison is a major goal of the RBC/UKQCD collaboration in the next few years, and substantial computational effort is currently underway both to increase the number of measurements on the existing $\Delta I = 1/2$ ensemble — driving down the statistical error — and to generate a second ensemble with a different lattice spacing, allowing the continuum limit of $\mathscr{A}_0$ to be taken.

2.3.3 The Calculation of $\mathscr{A}_0$ and G-Parity Boundary Conditions

While the $\Delta I = 1/2$ and $\Delta I = 3/2$ calculations outlined above are closely related, the $\Delta I = 1/2$ calculation is substantially more expensive to perform with physical kinematics and controlled statistical errors, for two major reasons. The first is the appearance of *disconnected diagrams* among the possible Wick contractions of the $\langle \pi\pi^0 | Q_i | K \rangle$ three-point functions, which do not appear in the corresponding $I = 2$ contractions[10]. Figure 2.5 shows an example of such a diagram. Disconnected

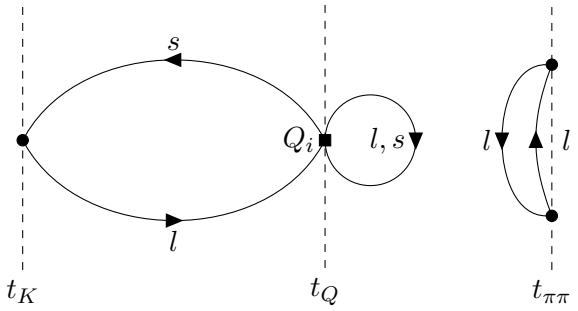

Figure 2.5: An example of a disconnected diagram contributing to $\mathscr{A}_0$. Lines denote light (l) or strange (s) quark propagators, circles mark the locations of the pions and kaon, and the box denotes an insertion of one of the Q_i operators.

diagrams are notoriously difficult to compute accurately on the lattice. Intuitively, they factorize as

[10]Disconnected diagrams appear in the $I = 0$ case because the $|\pi\pi^0\rangle$ state has the same quantum numbers as the vacuum, allowing contractions where the two pions annihilate into the vacuum and then reappear.

a product of two pieces with independent statistical fluctuations — $\langle \pi\pi^I | Q_i | K \rangle \sim \mathrm{tr}(\cdots) \times \mathrm{tr}(\cdots)$ — and thus require specialized methods and/or very long Monte Carlo simulations to suppress this enhanced noise.

The second issue making the $\Delta I = 1/2$ calculation more difficult than the $\Delta I = 3/2$ calculation is related to the respective techniques used to achieve physical kinematics. Since $2m_\pi < m_K$, the physical $K \to \pi\pi$ decay must involve final state pions with nonzero momenta. As we discussed in Section 1.3.3, Sachrajda and Villadoro [47] have demonstrated that partially twisted boundary conditions cannot be applied to the $|\pi\pi\rangle$ final state. In the $\Delta I = 3/2$ calculation a clever trick has been used to circumvent this problem: the Wigner-Eckhart theorem can be exploited to relate the desired $\langle \pi\pi^2 | Q_i | K \rangle$ matrix elements to unphysical $\langle \pi^+\pi^+ | Q_i | K^+ \rangle$ matrix elements [41]. If, in addition, antiperiodic boundary conditions are applied to the d quark in the spatial directions, the resulting $|\pi^+\pi^+\rangle$ ground state has zero total momentum, while the individual pions have momenta $|\vec{p}_\pi| = \sqrt{3}\pi/L$, and, crucially, $E^{I=2}_{\pi\pi} \approx m_K$ for the ensembles used in the most recent calculation of Ref. [41]. This setup allows for simulations with physical kinematics and well controlled statistical errors.

The same trick based on the Wigner-Eckhart theorem cannot be extended to the $\Delta I = 1/2$ calculation. The strategy of the most recent $\Delta I = 1/2$ calculation [44] has been to instead adopt the G-parity boundary conditions (GPBCs) introduced in Section 1.3.3. Since the pion is G-parity odd, its ground state has momentum $\pm\pi/L$ in spatial directions with GPBCs. As a result, after introducing GPBCs the ensemble parameters can be carefully tuned to achieve $E^{I=0}_{\pi\pi} \approx m_K$ with the pions in the ground state, ensuring once again that we have physical kinematics without substantially enhancing the statistical error.

GPBCs also mix quark flavors at the lattice boundary — the light quark doublet transforms as $(u, d) \mapsto (\overline{d}, -\overline{u})$ under the G-parity operation — implying that the G-parity Dirac operator inherently describes two quark flavors rather than one. In the HMC algorithm (Section 1.3.4) one typically chooses $M = \mathscr{D}^\dagger \mathscr{D}$ as the fermion matrix appearing in the pseudofermion path integral of Equation (1.40), since this M is Hermitian and positive-definite[11] whereas $\mathscr{D}$ is not. Since

[11] These properties are essential for the stability of numerical algorithms such as conjugate gradient.

the Dirac operator without GPBCs describes a single quark flavor, M describes two degenerate quark flavors, and can be used directly to simulate the isospin symmetric light quarks of most lattice calculations. In the G-parity case, however, $\mathscr{D}^\dagger\mathscr{D}$ describes four degenerate quark flavors, so one actually computes $[\det(\mathscr{D}^\dagger\mathscr{D})]^{1/2}$ for the light quark pair and $[\det(\mathscr{D}^\dagger\mathscr{D})]^{1/4}$ for the strange quark. While there is a standard and widely used algorithm for computing arbitrary roots of the fermion determinant — known as the rational HMC (RHMC) algorithm — it is substantially more expensive than a standard HMC simulation of a degenerate quark pair with $M = \mathscr{D}^\dagger\mathscr{D}$. One major result of this thesis is to explore an alternative algorithm for computing square roots of the fermion determinant, and demonstrate that it can substantially reduce the cost of G-parity simulations.

2.3.4 Rational Hybrid Monte Carlo and the Exact One Flavor Algorithm

On the lattice, dynamical fermions are included in simulations by expressing the fermionic determinant as a path integral over bosonic pseudofermion fields. For the isospin-symmetric light quark pair, one typically uses the *quotient action*

$$\mathrm{Quo}(m_1, m_2) \equiv \det\left(\frac{\mathscr{D}^\dagger\mathscr{D}(m_1)}{\mathscr{D}^\dagger\mathscr{D}(m_2)}\right), \tag{2.43}$$

and for single quark flavors the *rational quotient action*

$$\mathrm{RatQuo}_{1/n}(m_1, m_2) \equiv \left[\det\left(\frac{\mathscr{D}^\dagger\mathscr{D}(m_1)}{\mathscr{D}^\dagger\mathscr{D}(m_2)}\right)\right]^{1/n} \tag{2.44}$$

with $n = 2$. For G-parity simulations, one instead uses $\mathrm{RatQuo}_{1/2}$ for the light quarks and $\mathrm{RatQuo}_{1/4}$ for the strange quark. In the standard RHMC algorithm the root is approximated by a rational function

$$\left(\mathscr{D}^\dagger\mathscr{D}\right)^{1/n}\psi \approx \left(\alpha_0 + \sum_{i=1}^{N}\frac{\alpha_i}{\mathscr{D}^\dagger\mathscr{D} + \beta_i}\right)\psi \tag{2.45}$$

with coefficients $\alpha_i, \beta_i \in \mathbb{R}$ constructed via the Remez algorithm [48]. The matrix inverse implied by Equation (2.45) is well-defined, since $\mathscr{D}^\dagger\mathscr{D}$ is Hermitian and positive-definite, and since $\mathscr{D}$ itself has a bounded spectrum due to the lattice regularization. A naïvely formulated implementation of RHMC is prohibitively expensive, since each evaluation of the rational quotient action involves N independent CG inversions to compute $(\mathscr{D}^\dagger\mathscr{D} + \beta_i)^{-1}\psi$. This cost can be somewhat ameliorated

with a *multishift conjugate gradient* solver[12] [49], but RHMC simulations remain substantially more expensive than HMC simulations of the quotient action for the same quark mass. This is especially true for the G-parity light quarks.

In Chapter 6 we explore an alternative to RHMC proposed recently by the TWQCD collaboration [37] for simulating square-rooted fermion determinants with Wilson or domain wall quarks; this describes degenerate quark pairs with GPBCs or single quark flavors without GPBCs. TWQCD's construction applies the *Schur decomposition* (Eqn. (6.A.1)) to the spin structure of the Dirac operator, ultimately arriving at a factorization

$$\det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m_2)}\right) = \frac{1}{\det\left(H_1\right)} \cdot \frac{1}{\det\left(H_2\right)}, \tag{2.46}$$

with H_1 and H_2 manifestly Hermitian and positive-definite. Their exact one flavor algorithm is equivalent to RHMC in the sense that it computes the same determinant ratio, but has the advantage that it avoids the need for computing an overall square root of the fermion determinant, and thus the additional costs associated with rational approximations and multishift CG when evaluating the pseudofermion action. In a subsequent study TWQCD demonstrated that EOFA can provide modest $\mathcal{O}(20\%)$ performance improvements over RHMC, as well as a reduced memory footprint, after retuning the integrator used to evolve the HMC equations of motion [38].

In Appendix 6.A we provide a complete derivation of the EOFA action, following the outline of the derivation in Ref. [37], but filling in additional detail. We then elaborate on HMC with the exact one flavor algorithm, and perform statistical tests of the equivalence of RHMC and EOFA (Section 6.4). We further check explicitly, using a series of inexpensive $16^3 \times 32 \times 8$ and $16^3 \times 32 \times 16$ ensembles with heavy pion masses, that simulations performed with either the RHMC action or the

[12]Multishift CG exploits the observation that Krylov spaces are shift-invariant — $K_n(\mathscr{D}^\dagger\mathscr{D}, \phi) = K_n(\mathscr{D}^\dagger\mathscr{D} + \beta_i, \phi)$, where $K_n(M, \phi) \equiv \text{span}\{\phi, M\phi, \ldots, M^{n-1}\phi\}$ — to simultaneously invert a family of linear systems $(\mathscr{D}^\dagger\mathscr{D} + \beta_i)\psi = \phi$ for all β_i, with a convergence rate controlled by the least well-conditioned system. While this is substantially less expensive than performing independent standard CG inversions, it has a couple of important drawbacks: each iteration is more expensive since additional linear algebra is required to form the current search and solution vectors for each subsystem, and the memory footprint is much larger since all of these vectors need to be stored simultaneously. In addition, one is forced to use zero for each initial guess, rendering acceleration techniques such as forecasting and implicitly restarted algorithms inapplicable.

EOFA action give rise to consistent values for low energy observables, such as the average plaquette and pseudoscalar meson masses. After introducing and refining a number of techniques for tuning and accelerating EOFA simulations (Section 6.6), we benchmark EOFA against RHMC using two state-of-the-art RBC/UKQCD ensemble generation calculations with physical quark masses: the first is a coarse $24^3 \times 64 \times 24$ ensemble with non-GPBCs, and the second is the $32^3 \times 64 \times 12$ G-parity ensemble used to compute $\mathscr{A}_0$ in Ref. [44] (Section 6.7). In both cases we observe a significant speed-up — by a factor of 3.5 (5.0) per MD trajectory for the strange (light) quark determinant on the 24^3 (32^3) ensemble — after replacing RHMC with our highly-optimized implementation of EOFA. The key to these performance improvements is a novel preconditioning technique which substantially reduces the cost associated with inverting the Dirac operator in the context of EOFA (Appendix 6.C).

2.3.5 Ongoing Work

EOFA is currently being used in the production ensemble generation runs associated with the $\Delta I = 1/2$ $K \to \pi\pi$ calculation. The factor of 5.0 reduction in the cost of computing the light quark determinant per MD trajectory translates to a factor of 4.2 reduction in the total job time per MD trajectory, implying that we will be able to generate over four times as many gauge field configurations for the same computational cost after switching to EOFA. In addition, substantial effort has been made by others to increase the performance of the $\Delta I = 1/2$ measurement code, further reducing the cost of this calculation. We plan, within the next year, to increase the 216 existing measurements on the $a^{-1} = 1.4$ GeV, $32^3 \times 64 \times 12$ ensemble to $\mathcal{O}(1000)$ measurements, as well as to generate $\mathcal{O}(1000)$ measurements on a new, $24^3 \times 64 \times 24$ ensemble with a second lattice spacing of $a^{-1} \approx 1$ GeV. These new calculations will both drastically reduce the statistical error on $\mathscr{A}_0$ and allow a continuum limit to be taken, sharpening the test of ϵ' and Standard Model CP-violation first reported in Ref. [44].

While accelerating the $\Delta I = 1/2$ $K \to \pi\pi$ calculation was the short term goal of exploring EOFA, the algorithm is more general, and, as we have shown, can be used to accelerate any RHMC calculation involving a square-rooted fermion determinant. We expect EOFA may also prove useful

in at least two other contexts: $N_f = 2+1+1$ simulations which include a dynamical charm quark, and $N_f = 1+1+1$ (or $N_f = 1+1+1+1$) simulations with physical, non-isospin symmetric up and down quarks. Substantial progress toward practical $N_f = 2+1+1$ simulations using a very large $80^2 \times 96 \times 192 \times 32$ ensemble was made in Greg McGlynn's Ph.D. thesis [50], but this calculation has ultimately proven to be too expensive for the current generation of supercomputing resources[13]. We intend to explore the potential performance improvements associated with using EOFA for the strange and charm quarks on this ensemble. Isospin broken $N_f = 1+1+1$ calculations have been rare in lattice QCD to date, but there are a number of observables — such as the masses and decay constants of the low energy hadron spectrum, and the K_{l3} form factor $f_+^{K\pi}(q^2)$ — which can be computed to sufficiently high precision in current lattice QCD calculations that the small corrections arising from electromagnetic and isospin breaking effects must now be addressed to make further progress. We likewise intend to explore using EOFA to generate a coarse, $a^{-1} \approx 1$ GeV $24^3 \times 64 \times 24$ Möbius DWF ensemble with physical up and down quarks, allowing for exploratory calculations including isospin breaking effects.

References

[1] P. A. Boyle et al., "Low Energy Constants of $SU(2)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2+1$ Domain Wall QCD", Phys. Rev. **D93**, 054502 (2016).

[2] R. Mawhinney and D. Murphy, "NLO and NNLO Low Energy Constants for $SU(2)$ Chiral Perturbation Theory", PoS **LAT2015**, 061 (2015).

[3] R. Mawhinney and D. Murphy, "NLO and NNLO Low Energy Constants for $SU(3)$ Chiral Perturbation Theory", PoS **LAT2015**, 062 (2015).

[4] D. Murphy, "Light flavor physics from domain wall lattice qcd", in Ichep 2016 (Aug. 2016).

[5] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[13] The substantially heavier mass of the charm quark — $m_c \approx 1.3$ GeV — forces the use of very fine lattice spacings $a^{-1} \gtrsim 3$ GeV to simulate charm physics with reasonable discretization errors.

[6] P. A. Boyle et al., "The Kaon Semileptonic Form Factor in $N_f = 2 + 1$ Domain Wall Lattice QCD with Physical Light Quark Masses", JHEP **06**, 164 (2015).

[7] D. J. Murphy, "Domain Wall Fermion Simulations with the Exact One-Flavor Algorithm", PoS **LAT2016**, 272 (2016).

[8] S. Weinberg, "Phenomenological Lagrangians", Physica **A96**, 327–340 (1979).

[9] J. F. Donoghue, E. Golowich, and B. R. Holstein, *Dynamics of the standard model*, Cambridge Books Online (Cambridge University Press, 1992).

[10] S. Scherer, "Introduction to Chiral Perturbation Theory", Adv. Nucl. Phys. **27**, 277 (2003).

[11] S. Scherer and M. R. Schindler, "A Chiral Perturbation Theory Primer", (2005).

[12] B. Borasoy, "Introduction to Chiral Perturbation Theory", Springer Proc. Phys. **118**, 1–26 (2008).

[13] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory to One Loop", Annals Phys. **158**, 142 (1984).

[14] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory: Expansions in the Mass of the Strange Quark", Nucl.Phys. **B250**, 465 (1985).

[15] J. Bijnens, G. Colangelo, and G. Ecker, "Renormalization of Chiral Perturbation Theory to Order p^6", Annals Phys. **280**, 100–139 (2000).

[16] M. Peskin and D. Schroeder, *An Introduction To Quantum Field Theory*, Frontiers in Physics (Avalon Publishing, 1995).

[17] Bijnens, Johan and Lähde, Timo A., "Masses and Decay Constants of Pseudoscalar Mesons to Two Loops in Two-Flavor Partially Quenched Chiral Perturbation Theory", Phys.Rev. **D72**, 074502 (2005).

[18] Bijnens, Johan and Danielsson, Niclas and Lähde, Timo A., "The Pseudoscalar Meson Mass to Two Loops in Three-Flavor Partially Quenched χPT", Phys.Rev. **D70**, 111503 (2004).

[19] Bijnens, Johan and Lähde, Timo A., "Decay Constants of Pseudoscalar Mesons to Two Loops in Three-Flavor Partially Quenched χPT", Phys.Rev. **D71**, 094502 (2005).

[20] Bijnens, Johan and Danielsson, Niclas and Lähde, Timo A., "Three-Flavor Partially Quenched Chiral Perturbation Theory at NNLO for Meson Masses and Decay Constants", Phys.Rev. **D73**, 074509 (2006).

[21] C. Allton et al., "Physical Results from 2+1 Flavor Domain Wall QCD and $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D78**, 114509 (2008).

[22] M. Hansen, "Extracting Three-Body Observables from Finite-Volume Quantities", PoS **LATTICE2015**, 008 (2016).

[23] S. Borsanyi, S. Durr, Z. Fodor, S. Krieg, A. Schafer, et al., "$SU(2)$ Chiral Perturbation Theory Low-Energy Constants from 2+1 Flavor Staggered Lattice Simulations", Phys.Rev. **D88**, 014513 (2013).

[24] Dürr, Stephan and others, "Lattice QCD at the Physical Point Meets $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D90**, 114504 (2014).

[25] Brandt, Bastian B. and Jüttner, Andreas and Wittig, Hartmut, "The Pion Vector Form Factor from Lattice QCD and NNLO Chiral Perturbation Theory", JHEP **11**, 034 (2013).

[26] A. Bazavov et al., "MILC Results for Light Pseudoscalars", PoS **CD09**, 007 (2009).

[27] A. Bazavov et al., "Results for Light Pseudoscalar Mesons", PoS **LATTICE2010**, 074 (2010).

[28] J. Bijnens, "CHIRON: A Package for ChPT Numerical Results at Two Loops", Eur. Phys. J. **C75**, 27 (2015).

[29] J. L. Rosner, S. Stone, and R. S. Van de Water, "Leptonic Decays of Charged Pseudoscalar Mesons - 2015", Submitted to: Particle Data Book (2015).

[30] M. Antonelli et al., "An Evaluation of $|V_{us}|$ and Precise Tests of the Standard Model from World Data on Leptonic and Semileptonic Kaon Decays", Eur. Phys. J. **C69**, 399–424 (2010).

[31] M. Moulson, "Experimental Determination of V_{us} from Kaon Decays", in 9th International Workshop on the CKM Unitarity Triangle (CKM2016) Mumbai, India, November 28-December 3, 2016 (2017).

[32] Y. Aoki et al., "Continuum Limit Physics from $2+1$ Flavor Domain Wall QCD", Phys. Rev. **D83**, 074508 (2011).

[33] S. Aoki et al., "Review of Lattice Results Concerning Low-Energy Particle Physics", Eur. Phys. J. **C77**, 112 (2017).

[34] G. Amelino-Camelia et al., "Physics with the KLOE-2 Experiment at the Upgraded DAϕNE", Eur. Phys. J. **C68**, 619–681 (2010).

[35] C. Patrignani and P. D. Group, "Review of particle physics", Chinese Physics C **40**, 100001 (2016).

[36] K. Ogawa, T.-W. Chiu, and T.-H. Hsieh, "One-Flavor Algorithm for Wilson and Domain-Wall Fermions", PoS **LAT2009**, 033 (2009).

[37] Y.-C. Chen and T.-W. Chiu, "Exact Pseudofermion Action for Monte Carlo Simulation of Domain-Wall Fermion", Phys. Lett. **B738**, 55–60 (2014).

[38] Y.-C. Chen and T.-W. Chiu, "One-Flavor Algorithms for Simulation of Lattice QCD with Domain-Wall Fermion: EOFA versus RHMC", PoS **IWCSE2013**, 059 (2014).

[39] M. Lightman, "Delta I Equals Three Halfs Kaon To Two Pion Decays Using Lattice Quantum Chromodynamics with Domain Wall Fermions", PhD thesis (Columbia University, 2011).

[40] T. Blum et al., "Lattice Determination of the $K \to (\pi\pi)_{I=2}$ Decay Amplitude A_2", Phys. Rev. **D86**, 074513 (2012).

[41] T. Blum et al., "$K \to \pi\pi$ $\Delta I = 3/2$ Decay Amplitude in the Continuum Limit", Phys. Rev. **D91**, 074502 (2015).

[42] Q. Liu, "Kaon to Two Pions Decays from Lattice QCD: $\Delta I = 1/2$ Rule and CP Violation", PhD thesis (Columbia University, 2012).

[43] D. Zhang, "Kaon to Two Pion Decay from Lattice QCD and CP Violation", PhD thesis (Columbia University, 2015).

[44] Z. Bai et al., "Standard Model Prediction for Direct CP Violation in $K \to \pi\pi$ Decay", Phys. Rev. Lett. **115**, 212001 (2015).

[45] G. Buchalla, A. J. Buras, and M. E. Lautenbacher, "Weak Decays Beyond Leading Logarithms", Rev. Mod. Phys. **68**, 1125–1144 (1996).

[46] L. Lellouch and M. Lüscher, "Weak Transition Matrix Elements from Finite-Volume Correlation Functions", Communications in Mathematical Physics **219**, 31–44 (2001).

[47] C. T. Sachrajda and G. Villadoro, "Twisted Boundary Conditions in Lattice Simulations", Phys. Lett. **B609**, 73–85 (2005).

[48] M. Clark, "The Rational Hybrid Monte Carlo Algorithm", PhD thesis (The University of Edinburgh, 2005).

[49] B. Jegerlehner, "Krylov Space Solvers for Shifted Linear Systems", (1996).

[50] G. McGlynn, "Advances in Lattice Quantum Chromodynamics", PhD thesis (Columbia University, 2016).

Chapter 3

Next-to-Next-to Leading Order $SU(2)$ Chiral Perturbation Theory

The Low Energy Constants of $SU(2)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD

P.A. Boyle[1], N.H. Christ[2], N. Garron[3], C. Jung[4], A. Jüttner[5], C. Kelly[6], R.D. Mawhinney[2], G. McGlynn[2], D.J. Murphy[2], S. Ohta[7,8,6], A. Portelli[5], and C.T. Sachrajda[5]

[1] *SUPA, School of Physics, The University of Edinburgh, Edinburgh EH9 3JZ, UK*

[2] *Department of Physics, Columbia University, New York, NY 10027, USA*

[3] *Centre for Mathematical Sciences, Plymouth University, Plymouth, PL4 8AA, UK*

[4] *Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA*

[5] *School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK*

[6] *RIKEN-BNL Research Center, Brookhaven National Laboratory, Upton, NY 11973, USA*

[7] *Theory Center, KEK, Tsukuba, Ibaraki 305-0801, Japan*

[8]*Department of Particle and Nuclear Physics, SOKENDAI, Hayama, Kanagawa 240-0193, Japan*

Abstract

We have performed fits of the pseudoscalar masses and decay constants, from a variety of RBC-UKQCD domain wall fermion ensembles, to $SU(2)$ partially quenched chiral perturbation theory at next-to leading order (NLO) and next-to-next-to leading order (NNLO). We report values for 9 NLO and 8 linearly independent combinations of NNLO partially quenched low energy constants, which we compare to other lattice and phenomenological determinations. We discuss the size of successive terms in the chiral expansion and use our large set of low energy constants to make predictions for mass splittings due to QCD isospin breaking effects and the S-wave $\pi\pi$ scattering lengths. We conclude that, for the range of pseudoscalar masses explored in this work, 115 MeV $\lesssim m_{\rm PS} \lesssim$ 430 MeV, the NNLO $SU(2)$ expansion is quite robust and can fit lattice data with percent-scale accuracy.

3.1 Introduction

Effective field theories (EFT) formalize the intuitive idea that to understand physics at a particular energy scale E, the full details of physics at much higher energy scales $\Lambda \gg E$ are not needed. After identifying the relevant degrees of freedom associated with scale E, one can write down a low-energy approximation, which differs from the full theory up to corrections which are powers in E/Λ. If the separation of scales is large, the approximation is arbitrarily good, and the precise form of the E/Λ corrections need not be specified. In practice, high energy degrees of freedom do not need to be integrated out of the theory explicitly: it suffices to write down the most general low-energy effective Lagrangian containing all terms consistent with the symmetries of the full theory [1]. An early, successful example is the Fermi theory of β decay, which can be regarded as a low-energy approximation to the standard model obtained by integrating out the W boson [2]. Effective field theories are widely employed in modern physics, and the standard model itself is an EFT likely modified by some yet-unknown new physics at sufficiently high energies. Renormalization plays an important role in defining effective field theories, both in understanding how heavy particle masses

can enter at low scales via the Appelquist and Carazzone decoupling theorem [3], and in handling higher loop calculations in the low energy effective Lagrangian. Correctly matching EFTs across particle mass thresholds is a crucial detail of precision calculations in the standard model [4].

In this paper we discuss the physics of light pseudoscalar mesons, which played an important role in the development of the theory of the strong interactions — Quantum Chromodynamics (QCD) — and in the development of effective field theory techniques in general. The EFT of the light pseudoscalar mesons — Chiral Perturbation Theory (ChPT) — is both a prototypical example of an EFT and a theory whose corrections in powers of E/Λ can be determined, since lattice techniques enable direct QCD calculations. These correction terms contain "low energy constants" (LECs) which must be determined by matching to QCD. In this paper we fit lattice QCD data for the light pseudoscalar mesons to the corresponding ChPT formulas to determine the LECs and to gain information about the accuracy of ChPT as an approximation to QCD at low energies. While the primary focus is the physics of QCD, it is also of general interest to explore a system where the reliability of calculating in an EFT truncated to some order — we consider next-to leading order (NLO) and next-to-next-to leading order (NNLO) — can be directly tested against calculations in the full theory.

QCD is highly nonlinear in the low energy regime, and lattice QCD provides the only known technique for calculating hadronic properties from first principles[1]. The QCD vacuum dynamically breaks the $SU(N_f)_L \times SU(N_f)_R$ chiral symmetry of QCD with N_f massless quarks, at least when $N_f \leq 6$, giving rise to $N_f^2 - 1$ pseudo-Goldstone bosons, which are the pseudoscalar mesons. The scale of this meson physics is lighter than the scales of other phenomenon in QCD provided the quark masses are not too large, suggesting an effective field theory description (ChPT). For quarks of nonzero mass, one is naturally led to consider an effective field theory expansion in powers of the masses and momenta. One obtains $SU(2)$ ChPT [5] or $SU(3)$ ChPT [6] depending on whether or not the strange quark is included. The $SU(2)$ theory allows for explicit calculations of pion physics, while the $SU(3)$ theory describes the pseudoscalar meson octet (π, K, η). The matching of ChPT to QCD is encapsulated in the a priori unknown LECs, which parametrize the contributions from

[1] A perturbative expansion in powers of the strong coupling constant, g_s, is only useful at very high energies.

the various operators appearing in the ChPT Lagrangian.

Historically, ChPT has been an important tool for lattice QCD practitioners, as the limitations of available computational resources required the use of unphysically heavy quarks to make calculations practical. Until recently, a typical lattice calculation was performed at several, heavy values of the input quark masses, and then extrapolated with ChPT to the quark masses found in nature to make physical predictions. This is the approach taken in all but the most recent of the RBC-UKQCD collaboration's domain wall QCD simulations [7–9]. The reliability of ChPT as an approximation to QCD at the heavy, simulated points was largely left as an open question by these studies.

Recent advances in algorithms and computers have enabled computations directly at physical quark masses, minimizing the need for sophisticated chiral extrapolations. In the RBC-UKQCD collaboration's recent analysis of two physical mass Möbius domain wall fermion ensembles [10] $SU(2)$ ChPT was only used to correct for small mistunings in the simulation parameters, resulting in modest $\mathcal{O}(1\%)$ corrections to the simulated pseudoscalar masses and decay constants. While ChPT-based extrapolations may no longer be necessary in lattice QCD, the availability of lattice data ranging from physical to much heavier than physical quark mass allows for a complementary study of the applicability of ChPT as a low energy approximation to QCD. In this paper we seek to:

1. Determine as many of the low energy constants of $SU(2)$ ChPT as possible from our data, and

2. Systematically study the behavior and range of applicability of the $SU(2)$ ChPT expansion up to next-to-next-to leading order (NNLO).

Exploratory fits of an earlier RBC-UKQCD domain wall QCD data set to NNLO $SU(2)$ ChPT were first performed in Ref. [11], but suffered from numerical instabilities in the form of large NNLO corrections. More recently, the BMW collaboration has studied the pion mass and decay constant in $SU(2)$ ChPT up to NNLO using staggered [12] and Wilson [13] fermions. Fits of the pion mass, decay constant, and vector form factor computed using $\mathcal{O}(a)$-improved Wilson quarks to

NNLO $SU(2)$ ChPT were performed by Brandt, Jüttner, and Wittig [14]. Our domain wall fermion analyses complement these studies, providing an additional fermion discretization with excellent chiral symmetry properties. In addition, we perform our fits using the more general formalism of partially quenched chiral perturbation theory (PQChPT), from which we can also readily extract the low energy constants of ordinary ChPT. An analogous study of fits of RBC/UKQCD domain wall fermion data to $SU(3)$ partially quenched ChPT at NLO and NNLO will be the topic of a subsequent paper [15].

We briefly discuss some of the issues that arise in fitting our data to ChPT, which we will elaborate on in later sections. First, given that perturbative expansions of four-dimensional field theories generally produce asymptotic series rather than convergent series, the ChPT expansion is expected not to be convergent, with new counterterms arising at each loop order due to the non-renormalizability of the theory. One can hope that the series has the correct hierarchy to give accurate results when truncated to the first few orders — *i.e.* that each subsequent term is of smaller magnitude than the one that precedes it — for the range of quark masses probed in a typical lattice simulation, but this is not guaranteed. Second, if a large data set with quark masses less than some bound is fit to a given order of ChPT, statistical tests of the goodness of fit will become arbitrarily poor as the statistical resolution of the data is improved. This occurs because truncations of the ChPT expansion are only an approximation to QCD — eventually the data will be more accurate than the ChPT expansion can describe at a given order unless additional, higher order terms are added. This means that statistical goodness of fit criterion may initially show a reasonable fit to a small data set — when the statistical errors exceed the systematic errors from truncating the expansion — and then produce arbitrarily poor fits as more measurements are added and the statistical errors become smaller than the truncation errors. Finally, our fit procedure only gives us a self-consistent view of the properties of the expansion: we have data corresponding to a particular range of quark masses, which we fit to ChPT, and then ask whether the resulting expansion is sensible. While we have some freedom to vary the range of quark masses included in our fits, lattice QCD can, in principle, provide arbitrarily accurate data at arbitrarily small quark masses. For the time being we remain far from that situation.

3.2 Partially Quenched Chiral Perturbation Theory at Next-to-Next-to-Leading Order

The basic degrees of freedom in QCD are the quark fields, q_f, which transform in the fundamental representation of (color) $SU(3)$ and carry a flavor index f, and the gluon fields, A^a_μ, which transform in the adjoint representation of (color) $SU(3)$ and mediate the strong nuclear force. In the limit of vanishing quark masses, the QCD Lagrangian with N_f flavors of quarks

$$\mathcal{L}_{QCD} = -\frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a + \sum_f \overline{q}_f i\gamma^\mu D_\mu q_f \tag{3.1}$$

has an exact $SU(N_f)_L \times SU(N_f)_R$ symmetry[2]. This symmetry is spontaneously broken down to a single $SU(N_f)_V$ subgroup by the QCD vacuum, giving rise to $N_f^2 - 1$ Goldstone bosons: these are the pions (π^+, π^0, π^-) for $N_f = 2$, and the pseudoscalar octet[3] $(\pi^+, \pi^0, \pi^-, K^+, K^0, \overline{K}^0, K^-, \eta_8)$ for $N_f = 3$. The full $SU(N_f)_L \times SU(N_f)_R$ symmetry of the massless Lagrangian is also explicitly broken by the nonzero masses of the quarks in nature, generating masses for the (pseudo-)Goldstone bosons.

ChPT is the low-energy effective theory whose degrees of freedom are precisely the Goldstone bosons of QCD. The Goldstone fields can be parametrized in the exponential representation

$$U(x) = \exp\left(\frac{i}{f}\phi(x)\right), \quad \phi(x) \in su(N_f) \tag{3.2}$$

with

$$\phi(x) = \begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 & \pi^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 \end{pmatrix} \tag{3.3}$$

[2]Naively, the classical Lagrangian (3.1) has an even larger $U(N_f)_L \times U(N_f)_R$ symmetry, but the $U(1)_A$ component is broken by the chiral anomaly and fails to be a symmetry of the quantum theory.

[3]We use the notation η_8 to emphasize that this is the pseudo-Goldstone boson associated with the eighth generator of $SU(3)$, not the physical η meson detected in particle experiments. In reality flavor $SU(3)$ is not an exact symmetry of nature, and the states $\eta_1 = (u\overline{u} + d\overline{d} + s\overline{s})/\sqrt{3}$ and $\eta_8 = (u\overline{u} + d\overline{d} - 2s\overline{s})/\sqrt{6}$ mix to form the physical η and η'.

for the $SU(2)$ theory, and

$$\phi(x) = \begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta_8 & \pi^+ & K^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta_8 & K^0 \\ K^- & \overline{K}^0 & -\frac{2}{\sqrt{6}}\eta_8 \end{pmatrix} \tag{3.4}$$

for the $SU(3)$ theory. A *chiral order* is assigned to each term by counting the number of derivatives of U which enter: $\partial^n U \sim p^n$, where p corresponds to external momenta carried by the Goldstone bosons. One can then systematically construct the ChPT Lagrangian order-by-order in this power counting scheme

$$\mathcal{L}_{\text{ChPT}} = \underbrace{\mathcal{L}^{(2)}_{\text{ChPT}}}_{\text{LO}} + \underbrace{\mathcal{L}^{(4)}_{\text{ChPT}}}_{\text{NLO}} + \underbrace{\mathcal{L}^{(6)}_{\text{ChPT}}}_{\text{NNLO}} + \cdots \tag{3.5}$$

by writing down all possible terms of $\mathcal{O}(p^n)$,

$$\mathcal{L}^{(n)}_{\text{ChPT}} = \sum_i \alpha_i \mathscr{O}^{(n)}_i, \tag{3.6}$$

where $\alpha_i \in \mathbb{R}$ are the low energy constants, and $\mathscr{O}^{(n)}_i \sim p^n$ is constructed from U and its derivatives, and is invariant under the $SU(N_f)_L \times SU(N_f)_R$ symmetry. Gasser and Leutwyler further showed that by coupling the quark mass matrix, vector and axial currents, and scalar and pseudoscalar densities to the ChPT Lagrangian as external sources one can elegantly reproduce the Ward identities of QCD by taking appropriate functional derivatives [5, 6]. While this construction produces the most general effective Lagrangian consistent with the underlying symmetries of QCD, the numerical values of the low energy constants (LECs) are a priori unknown, and must be determined phenomenologically or by fits to lattice simulations.

The first detailed, next-to-leading order ChPT calculations were performed by Gasser and Leutwyler in Ref. [5] for the $SU(2)$ case, and Ref. [6] for the $SU(3)$ case. They compute a number of two-point and four-point correlation functions which allow them to determine the pseudoscalar masses and decay constants, scattering lengths, and other low-energy observables of interest. These calculations were then extended to NNLO in [16], where the $\mathcal{O}(p^6)$ Lagrangian was first explicitly constructed, and in Ref. [17] ($SU(2)$) and Ref. [18] ($SU(3)$). We will make use of two further generalizations of chiral perturbation theory: finite volume ChPT and partially quenched ChPT.

In finite volume (FV) ChPT the spatial $\mathbb{R}^3$ of Minkowski spacetime is replaced with a cubic box of volume L^3. This discretizes the allowed momentum states, requiring continuous integrals over momenta to be replaced with sums. Corrections to infinite volume ChPT results can be computed as functions of L, and must vanish in the $L \to \infty$ limit. Since, in a typical lattice QCD simulation, the pion correlation length is comparable to L, finite volume effects are often one of the dominant systematic errors when trying to make physical predictions, and FV ChPT is important to remove or bound these errors. In our fits we parametrize the chiral ansätze for the pseudoscalar masses and decay constants as

$$\begin{aligned} m_{xy}^2 &= (m_{xy}^2)^\infty + \Delta_{m_{xy}^2}^L \\ f_{xy} &= (f_{xy})^\infty + \Delta_{f_{xy}}^L \end{aligned} \tag{3.7}$$

where $(X)^\infty$ denotes the infinite volume result, and $\Delta_X^L \equiv (X)^L - (X)^\infty$ is the finite volume correction for a box of size L. Explicit formulae for Δ_X^L are known to NNLO [19–21], but we will only make use of the NLO results summarized in the appendices of Ref. [7] for our fits.

Partial quenching is a technique used in lattice simulations to lower the simulated pion mass without substantially increasing computational cost. On the lattice one is free to independently vary the sea and valence quark masses: the former enter the fermion determinant used to generate gauge field configurations, and the latter appear in fermion propagators when computing correlation functions. In practice $m_{\text{val}} < m_{\text{sea}}$ is often used since reducing the sea quark masses is more expensive than reducing the valence quark masses. One can regard partially quenched QCD as a theory in its own right, which reduces to ordinary QCD in the unitary limit $m_{\text{val}} = m_{\text{sea}}$.

In the framework of ChPT partial quenching is included analytically by generalizing to a supersymmetric theory with N_{sea} and N_{val} sea and valence quarks, respectively. The theory also contains N_{val} unphysical bosonic ghost quarks which exactly cancel the contributions from the fermionic valence quarks to closed fermion loops. The $SU(N_f)_L \times SU(N_f)_R$ symmetry of ordinary massless QCD is promoted to a graded $SU(N_{\text{val}} + N_{\text{sea}}|N_{\text{val}})_L \times SU(N_{\text{val}} + N_{\text{sea}}|N_{\text{val}})_R$ symmetry, and the most general effective Lagrangian consistent with this symmetry is constructed order-by-order, in analogy to ordinary ChPT. The original construction of the PQChPT Lagrangian is discussed in Ref. [22], and in Ref. [23] NLO expressions for the pion mass and decay constant are calculated. For

our NLO PQChPT fits we use the explicit $SU(2)$ formulae collected in Ref. [7]. Bijnens, Danielsson, and Lähde further generalized the PQChPT expressions for the partially quenched pseudoscalar masses and decay constants to NNLO: these calculations are presented in Ref. [24] for the $SU(2)$ case and Ref. [25–27] for the $SU(3)$ case. We make use of Fortran codes provided by Bijnens to compute these expressions in our NNLO fits. By explicitly taking the unitary limit $m_{\rm val} = m_{\rm sea}$ in the PQChPT Lagrangian and matching to the ChPT Lagrangian one can write down explicit relations between the PQChPT and ChPT LECs. We collect these results in Appendix 3.A.1.

In Table 3.1 we summarize the counting of LECs up to NNLO in $SU(2)$ and $SU(3)$ ChPT and PQChPT, and introduce our notation.

	ChPT	ChPT	PQChPT	PQChPT
N_f	2	3	2	3
LO	B, f	B_0, f_0	B, f	B_0, f_0
NLO	l_i	L_i	$\hat{L}_i^{(2)}$	$\hat{L}_i^{(3)}$
	7	10	11	11
NNLO	c_i	C_i	$\hat{K}_i^{(2)}$	$\hat{K}_i^{(3)}$
	53	90	112	112

Table 3.1: Counting of the LECs in ChPT and PQChPT up to NNLO, from [28]. The notations $\{l_i, c_i\}$ for the $SU(2)$ ChPT LECs and $\{L_i, C_i\}$ for the $SU(3)$ ChPT LECs are conventional in the literature. Similarly, we use the notation $\{\hat{L}_i^{(N_f)}, \hat{K}_i^{(N_f)}\}$ to distinguish the more general partially quenched LECs.

3.3 Lattice Setup

In this analysis we make use of a number of RBC/UKQCD domain wall fermion ensembles with a wide range of unitary pion masses, 117 MeV $\leq m_\pi \leq$ 432 MeV, physical volumes, (2.005(11) fm)$^3 \leq L^3 \leq$ (6.43(26) fm)^3, and inverse lattice spacings, 0.98(4) GeV $\leq a^{-1} \leq$ 3.14(2) GeV. In all cases we work in the isospin symmetric limit of QCD, with two, degenerate dynamical light quark flavors

of bare mass m_l, and a single dynamical heavy flavor of bare mass m_h ($N_f = 2 + 1$). Many of these ensembles have been analyzed in earlier publications which describe the ensemble generation, fits to extract the spectrum, and earlier chiral extrapolations based on NLO chiral perturbation theory [7–10]. We also include two new Möbius domain wall fermion ensembles; details of the ensemble generation and fits to extract the spectrum are discussed in Appendix 3.C.

In Table 3.2 we list the 12 ensembles included in this analysis and summarize the actions and input parameters. In all cases we use the Iwasaki gauge action (I) [29], and on some ensembles supplement this with the dislocation suppressing determinant ratio (I+DSDR) [30, 31]. The DSDR term suppresses dislocations ("tears") in the gauge field, representing tunneling between different topological sectors, that give rise to enhanced chiral symmetry breaking in domain wall fermion calculations, and occur more frequently at strong coupling. We simulate QCD with $N_f = 2 + 1$ quark flavors using the domain wall fermion formalism, with either the Shamir (DWF) [32, 33] or Möbius (MDWF) [34–36] kernel. The details of how the low-energy QCD spectrum has been extracted from fits to various Green's functions can be found in Ref. [7] for the 24I ensembles, Ref. [8] for the 32I ensembles, Ref. [9] for the 32ID ensembles, Ref. [10] for the 48I, 64I, and 32I-fine ensembles, and in Appendix 3.C for the 32ID-M1 and 32ID-M2 ensembles. In addition, detailed discussions of the Möbius kernel and the properties of MDWF simulations of QCD can be found in Ref. [10].

In Appendix 3.B we list fit values at the simulated quark masses in lattice units for the pseudoscalar masses and decay constants, Ω baryon mass, residual mass, and Wilson flow scales on each ensemble. On the older 24I, 32I, and 32ID ensembles these measurements were performed for a number of different partially quenched valence quark mass combinations which are listed explicitly in the appendix. In addition, reweighting in the dynamical heavy quark mass was used to determine the m_h dependence and allow for a small, linear interpolation from the simulated m_h to the physical value. On the newer ensembles — 32I-fine, 48I, 64I, 32ID-M1, and 32ID-M2 — we perform a single set of unitary measurements of the same observables, and do not reweight in m_h.

Ensemble	Action	β	$L^3 \times T \times L_s$	am_l	am_h	$m_\pi L$	m_π (MeV)
24I	DWF+I	2.13	$24^3 \times 64 \times 16$	0.005	0.04	4.568(13)	339.6(1.2)
	DWF+I	2.13	$24^3 \times 64 \times 16$	0.01	0.04	5.814(12)	432.2(1.4)
32I	DWF+I	2.25	$32^3 \times 64 \times 16$	0.004	0.03	4.062(11)	302.0(1.1)
	DWF+I	2.25	$32^3 \times 64 \times 16$	0.006	0.03	4.8377(82)	359.7(1.2)
	DWF+I	2.25	$32^3 \times 64 \times 16$	0.008	0.03	5.526(12)	410.8(1.5)
32ID	DWF+I+DSDR	1.75	$32^3 \times 64 \times 32$	0.001	0.046	3.9992(69)	172.7(9)
	DWF+I+DSDR	1.75	$32^3 \times 64 \times 32$	0.0042	0.046	5.7918(79)	250.1(1.2)
32I-fine	DWF+I	2.37	$32^3 \times 64 \times 12$	0.0047	0.0186	3.773(42)	370.1(4.4)
48I	MDWF+I	2.13	$48^3 \times 96 \times 24$	0.00078	0.0362	3.8633(63)	139.1(4)
64I	MDWF+I	2.25	$64^3 \times 128 \times 12$	0.000678	0.02661	3.7778(84)	139.0(5)
32ID-M1	MDWF+I+DSDR	1.633	$32^3 \times 64 \times 24$	0.00022	0.0596	3.780(15)	117.3(4.4)
32ID-M2	MDWF+I+DSDR	1.943	$32^3 \times 64 \times 12$	0.00478	0.03297	6.236(21)	401.0(2.3)

Table 3.2: Summary of ensembles included in this analysis and input parameters. Here β is the gauge coupling, $L^3 \times T \times L_s$ is the lattice volume decomposed into the length of the spatial (L), temporal (T), and fifth (L_s) dimensions, and am_l and am_h are the bare, input light and heavy quark masses. The value of m_π quoted is the unitary pion mass in physical units, where we have used the lattice spacings listed in Table 3.3.

3.4 The Global Fit Procedure

In Ref. [8–10] we have developed a “global fit" procedure for performing a combined chiral fit and continuum extrapolation of lattice data, the details of which we will summarize here. The global fit also allows us to convert predictions from our simulations, which are performed in dimensionless lattice units, into physical units by determining the lattice spacing a on each ensemble. While we have historically focused on using this construction to make physical predictions from our simulations, viewing chiral perturbation theory as a tool to parametrize the quark mass dependence of low-energy QCD observables, here we will adopt a slightly different view and regard the fit to ChPT itself as our primary interest.

Our canonical global fit, which we have most recently used in Ref. [10], includes the pion and kaon masses[4] m_π and m_K, the pion and kaon decay constants f_π and f_K, the omega baryon mass m_Ω, and the Wilson flow scales [37] $t_0^{1/2}$ and w_0. Partially quenched next-to-leading order $SU(2)$ chiral perturbation theory with finite volume corrections is used to perform the chiral fit to the valence quark (m_x, m_y) and light dynamical quark (m_l) mass dependence of m_π and f_π. The input dynamical heavy quark mass is carefully tuned during the ensemble generation to closely correspond to the physical strange quark mass, however, any slight mistuning introduces small errors in our simulated values of m_π and f_π, which are not described by $SU(2)$ PQChPT. We account for this by reweighting (see Section II.D of Ref. [8]) in the heavy quark determinant to generate a series of values of each observable for several m_h near the simulated mass, and then supplement the chiral $SU(2)$ ansatz with a term linear in m_h, allowing us to interpolate the reweighted data to the physical strange quark mass. NLO $SU(2)$ heavy meson PQChPT with finite volume corrections [7, 38] is used for m_K and f_K. The chiral fits to m_Ω and the Wilson flow scales are performed using a simple analytic ansatz which is linear in the quark masses. Discretization effects are included by adding a term linear in a^2 to each fit form, allowing us to ultimately take the continuum limit $a \to 0$. The raw simulation data is in dimensionless lattice units which are different for each ensemble, reflecting the different (physical) lattice spacings. We account for this by performing the chiral fits in the bare, dimensionless lattice units of a single reference ensemble, which we choose to be our $32^3 \times 64$ Iwasaki (32I) lattice (Table 3.2). The choice of reference ensemble is arbitrary, and for well-behaved fits should have no influence on predictions for physical observables or for the values of the low energy constants. We introduce additional fit parameters

$$R_a^e \equiv \frac{a^r}{a^e}, \quad Z_l^e \equiv \frac{1}{R_a^e}\frac{(a\tilde{m}_l)^r}{(a\tilde{m}_l)^e}, \quad Z_h^e \equiv \frac{1}{R_a^e}\frac{(a\tilde{m}_h)^r}{(a\tilde{m}_h)^e} \tag{3.8}$$

to convert between bare lattice units on the reference ensemble r and other ensembles e, where a is the lattice spacing and $\tilde{m}_q = m_q + m_{\text{res}}$ is the total quark mass[5].

[4]Note: we work in the isospin symmetric limit of QCD, where $m_u = m_d \equiv m_l$, and neglect electromagnetic corrections. In this limit the charged and neutral pions are degenerate, as are the charged and neutral kaons, so we can speak unambiguously of "the pion" and "the kaon".

[5]In the domain wall fermion formalism a finite fifth dimension introduces a small chiral symmetry breaking, leading

The chiral ansätze discussed above reflect a simultaneous expansion in the quark masses, lattice volume (L), and lattice spacing (a), about the infinite volume, continuum, chiral limit. Our power-counting scheme counts the dominant discretization term — which is proportional to a^2 for domain wall fermions — as the same order as the NLO continuum PQChPT corrections. While we include continuum PQChPT terms up to $\mathcal{O}(p^6)$ in our NNLO fits, cross terms proportional to $X^{\rm NLO}\times\Delta_X^{\rm NLO}$ and $X^{\rm NLO}\times a^2$ are neglected since they are higher-order in our power-counting, and are empirically observed to be small. The full chiral ansatz for $X\in\{m_\pi^2, f_\pi\}$, for example, including the finite volume and a^2 terms, has the generic form

$$X(\tilde{m}_q, L, a^2) \simeq X_0 \Big(1 + \underbrace{X^{\rm NLO}(\tilde{m}_q) + X^{\rm NNLO}(\tilde{m}_q)}_{\text{NNLO Continuum PQChPT}} + \underbrace{\Delta_X^{\rm NLO}(\tilde{m}_q, L)}_{\text{NLO FV corrections}} + \underbrace{c_X a^2}_{\text{Lattice spacing}} \Big) \tag{3.9}$$

where X_0 is the leading order value of X in the continuum and infinite-volume limits, and "$\simeq$" denotes equality up to truncation of higher order terms. Since the Iwasaki and I+DSDR actions have, in general, different discretization errors for a given value of the lattice spacing, we fit independent a^2 coefficients for each observable X, denoted c_X^I and c_X^{ID}, respectively. The NLO $SU(2)$ ansätze are written in complete detail in Appendix H of Ref. [10]; the generalization to NNLO is straightforward. Appendix B of the same reference also discusses how to write a given chiral ansatz in our dimensionless formalism.

The procedure for performing a global fit is as follows:

1. The valence quark mass dependence of $m_{\rm res}$ is fit to a linear ansatz on each ensemble. We then extrapolate $m_{\rm res}$ to the chiral limit $m_q \to 0$, and use this value in the remainder of the analysis.

2. A simultaneous chiral/continuum fit of m_π^2, m_K^2, f_π, f_K, m_Ω, $t_0^{1/2}$ and w_0 is performed on all ensembles using the ansätze described in the preceding paragraph. The quark mass

to an additive renormalization of the input quark masses by $m_{\rm res}$ (the residual mass). In Appendix 3.C we briefly discuss how $m_{\rm res}$ is extracted.

dependence is parametrized in terms of $\tilde{m}_q = m_q + m_{\rm res}$. This step also determines the ratios of lattice scales R_a^e and $Z^e_{\{l,h\}}$ and the dependence on a^2.

3. Three of the quantities from 2 are defined to have no a^2 corrections and establish our continuum scaling trajectory by matching onto their known, physical values[6]. In the analysis of [10] we have used m_π, m_K, and m_Ω, and implemented this condition by numerically inverting the chiral fit to determine input bare valence quark masses $m_l^{\rm phys}$ and $m_h^{\rm phys}$ such that the ratios m_π/m_Ω and m_K/m_Ω take their physical values.

4. From 3 we obtain m_Ω at $m_l^{\rm phys}$ and $m_h^{\rm phys}$ on the reference ensemble; we then use the ratio $m_\Omega^r/m_\Omega^{\rm phys}$ to determine the lattice spacing a^r in physical units. Together with the ratios of lattice scales from 2 we can determine the lattice spacings on the other ensembles, as well as extrapolate observables to the physical quark mass, continuum limit in physical units.

The fits described in steps 1 and 2 are performed using uncorrelated nonlinear χ^2 minimization with the Levenberg-Marquardt algorithm [40, 41]. Due to the large number of data points in our fits we have a very nearly singular correlation matrix that we cannot reliably invert, as would be required to perform fits with a fully correlated χ^2; we show an example of one of our correlation matrices in Appendix 3.D. As a result, the χ^2/dof that we present cannot be interpreted as the goodness-of-fit, and instead we will present histograms showing the distribution of the data around our fit. These histograms provide a simple summary of the fit quality, and, in particular, highlight any data that is far from the fit function. The numerical inversion in step 3 is performed by minimizing

$$\chi^2 = \left[\left(\frac{m_\pi}{m_\Omega}\right)(\tilde{m}_l, \tilde{m}_h) - \left(\frac{m_\pi}{m_\Omega}\right)^{\rm PDG}\right]^2 + \left[\left(\frac{m_K}{m_\Omega}\right)(\tilde{m}_l, \tilde{m}_h) - \left(\frac{m_K}{m_\Omega}\right)^{\rm PDG}\right]^2, \tag{3.10}$$

where PDG denotes the experimental value from [39]. Statistical errors on the fit parameters are computed using the superjackknife resampling technique [42]. The choices of which quantities are used to determine the physical quark masses in step 3 and the lattice spacing in step 4 are arbitrary,

[6] For reference, our values for the "physical", isospin symmetric masses and decay constants, excluding QED effects, are: $m_\pi^{\rm phys} = 135.0\,{\rm MeV}$ (PDG π^0 mass), $m_K^{\rm phys} = 495.7\,{\rm MeV}$ (average of the PDG K^0 and $K^\pm$ masses), $m_\Omega^{\rm phys} = 1672.45\,{\rm MeV}$ (PDG Ω^- mass), $f_\pi^{\rm phys} = 130.4\,{\rm MeV}$ (PDG π^- decay constant), and $f_K^{\rm phys} = 156.1\,{\rm MeV}$ (PDG K^- decay constant) [39].

and all results should agree in the continuum limit regardless of this choice.

The matching to our chosen scaling trajectory results in values of the physical quark masses, m_l^{phys} and m_h^{phys}, as well as corresponding values of the leading-order chiral parameter B, that are normalized in the native units of our 32I ensemble. In order to be useful to others, these quantities must be renormalized into a more convenient scheme such as $\overline{\text{MS}}$. As described in Refs. [8–10] we achieve this by first renormalizing in variants of the non-perturbative Rome-Southampton regularization-invariant momentum scheme with symmetric kinematics (RI/SMOM) [43–47]. The matching factors between these schemes and $\overline{\text{MS}}$ can be computed using standard continuum perturbation theory with dimensional regularization applied at a high energy scale, typically $\mu \sim 3$ GeV, at which perturbation theory is known to be reliable. We use the RI/SMOM intermediate scheme for our central values. The only significant systematic error on the result is due to the truncation of the perturbative series to two-loop order in the computation of the RI/SMOM$\to \overline{\text{MS}}$ matching factors. In order to estimate the size of this effect we compare the resulting $\overline{\text{MS}}$ values to those computed using the RI/SMOM$_{\gamma^\mu}$ intermediate scheme, taking the full difference as a conservative estimate[7].

Renormalized quark masses are obtained by taking the product

$$m_f^{\overline{\text{MS}}} = Z_m^{\overline{\text{MS}},\ 32\text{I}} m_f^{\text{phys}} + \mathcal{O}(a^2)\,, \tag{3.11}$$

where $f \in \{l, h\}$ and $Z_m^{\overline{\text{MS}},\ 32\text{I}}$ is the quark mass renormalization coefficient computed on the 32I ensemble. This determination of $m_f^{\overline{\text{MS}}}$ contains $\text{O}(a^2)$ errors because the renormalization factors have only been computed at a single lattice spacing. Using the quantities Z_l and Z_h defined in Eqn. (3.8), we can also compute the renormalized physical quark mass using renormalization factors calculated on the 24I ensemble as follows:

$$m_f^{\overline{\text{MS}}} = \frac{Z_m^{\overline{\text{MS}},\ 24\text{I}}}{Z_f^{24\text{I}}} m_f^{\text{phys}} + \mathcal{O}(a^2)\,, \tag{3.12}$$

[7]For more detail regarding the SMOM and SMOM$_{\gamma^\mu}$ schemes we refer the reader to Refs. [8, 46].

Combining these two equations, we can compute a value for the quark mass that is free from $\mathcal{O}(a^2)$ errors:

$$m_f^{\overline{\rm MS}} = Z_{mf}^{\overline{\rm MS}} m_f^{\rm phys} + \mathcal{O}(a^4)\,, \tag{3.13}$$

where

$$Z_{mf}^{\overline{\rm MS}} = \lim_{a\to 0} \left\{ Z_m^{\overline{\rm MS}}(a)/Z_f(a) \right\}\,, \tag{3.14}$$

and the $a \to 0$ limit is taken by performing a linear extrapolation using the two available lattice spacings. Similarly, the renormalized value of B can be obtained as

$$B^{\overline{\rm MS}} = B^{\rm fit}/Z_{ml}^{\overline{\rm MS}}\,. \tag{3.15}$$

Note that the fact that domain wall fermions are non-perturbatively $\mathcal{O}(a)$ improved and have good chiral symmetry eliminates dependence on odd-powers of the lattice spacing.

For this analysis we use the values of Z_{ml} and Z_{mh} computed in Ref. [10], and for more details we refer the reader to Section V.C and Appendix F of that work. Note that the calculation of these quantities necessarily involves the computed values of the lattice spacing, which differ between the various fits we perform. For the analyses presented in this document we do not recompute Z_{mf} for each fit; however our lattice spacings are all in excellent agreement with those in the aforementioned work, hence we choose to neglect the small systematic error associated with this mismatch.

While the fits discussed in this work are in many ways an extension of the analysis presented in Ref. [10], there are a few important differences we would like to emphasize. First, in Ref. [10] chiral perturbation theory was used only to make modest, $\mathcal{O}(1\%)$ corrections to the spectrum computed on the physical quark mass $48^3 \times 96$ (48I) and $64^3 \times 128$ (64I) lattices. This was achieved using an overweighting procedure, in which the contributions

$$\chi_e^2 = \alpha_e \sum_i \left(\frac{y_e^i - f_e^i}{\sigma_e^i} \right)^2 \tag{3.16}$$

to $\chi^2 = \sum_e \chi_e^2$ from each ensemble were multiplied by tunable, independent parameters α_e. By choosing $\alpha_{48\rm I}, \alpha_{64\rm I} \gg 1$ and $\alpha_e = 1$ otherwise, the fit was effectively forced to pass through the

48I and 64I data, using information from the other ensembles only to make a small correction to the physical point. In this work we are interested more generally in the applicability of chiral perturbation theory to describe the quark mass dependence of the QCD spectrum, and thus we do not employ overweighting. Second, in Ref. [10] the Wilson flow scales $t_0^{1/2}$ and w_0 were introduced into the global fit procedure, which we do not include in any of the fits presented in Section 3.5. While the inclusion of the Wilson flow scales leads to a marked improvement in the determination of the lattice spacings, they do not constrain the ChPT LECs, and are unnecessary for our computationally demanding NNLO fits.

Since the 32ID-M1 and 32ID-M2 lattices have not appeared in our earlier global fit analyses, we have updated our canonical global fit from Ref. [10] to include these ensembles and determine their properties. We note that even though the 32ID-M2 ensemble has a relatively heavy unitary pion mass ($m_\pi = 401.0(2.3)\,\mathrm{MeV}$) that lies outside the 370 MeV cut used in this fit, the overweighting procedure results in a fit that is insensitive to heavy ensembles, and we can safely assume that this discrepancy will not lead to any significant systematics. This provides an explicit check that our fits in this work, including the new ensembles, are consistent with our earlier work, and we indeed see that the lattice spacings and other parameters are consistent with Ref. [10]. This fit also establishes a baseline relative to the global fit performed in Ref. [10], by which we can judge the consistency of the new fits discussed in Section 3.5. The values we obtain for the physical box sizes, lattice spacings, and residual mass in the chiral limit are summarized in Table 3.3.

Ensemble	L (fm)	a^{-1} (GeV)	am_l^{phys}	am_h^{phys}	am_{res}
24I	2.6496(73)	1.7844(49)	-0.001770(79)	0.03225(18)	0.003038(78)
32I	2.6466(93)	2.3820(84)	0.000261(13)	0.02480(18)	0.000662(11)
32ID	4.573(22)	1.3784(68)	-0.000106(16)	0.04625(48)	0.0018478(73)
32I-fine	2.005(11)	3.144(17)	0.000057(16)	0.01846(32)	0.0006300(59)
48I	5.468(12)	1.7293(36)	0.0006982(80)	0.03580(16)	0.0006102(40)
64I	5.349(16)	2.3572(69)	0.0006213(77)	0.02542(17)	0.0003116(23)
32ID-M1	6.43(26)	0.981(39)	0.00107(26)	0.0850(68)	0.002170(16)
32ID-M2	3.067(16)	2.055(11)	-0.003429(16)	0.02358(33)	0.0044660(46)

Table 3.3: Physical box sizes, inverse lattice spacings, bare, unrenormalized quark masses, and residual mass in the chiral limit for the ensembles included in this work. These numbers are obtained by repeating the global fit analysis published in Ref. [10], including the new 32ID-M1 and 32ID-M2 ensembles.

3.5 Fits to $SU(2)$ PQχPT

In this section we discuss global fits based on $SU(2)$ partially quenched chiral perturbation theory. These fits include:

1. The pion mass and decay constant, fit to NLO or NNLO PQChPT, with NLO finite volume corrections in both cases.

2. The kaon mass and decay constant, fit to NLO heavy-meson PQChPT with NLO finite volume corrections.

3. The Ω baryon mass, fit to a linear, analytic ansatz.

m_π, m_K, and m_Ω are used as the three inputs to determine the physical quark masses and lattice spacings; this leaves f_π and f_K as predictions. We consider two different cuts on the heaviest unitary pion mass included in the fit: 370 MeV and 450 MeV. Any ensemble with a unitary pion mass greater than the cut is excluded from the fit completely. Likewise, all partially quenched

"pion" measurements with $m_{xy} > m_\pi^{\rm cut}$ are excluded even if the unitary pion mass is within the cut. The data we use for the fits with a 370 MeV cut is the same as the data used in the fits with a 370 MeV cut in Ref. [10], with the addition of the new 32ID-M1 ensemble. We do not include any additional kaon or Ω baryon data when we raise the mass cut, since these quantities are described by NLO (kaon) or linear (Ω) ansatzäe in all of the fits that we have performed — the heavier 450 MeV cut is intended to test the full partially quenched NNLO expressions for m_π and f_π by using all of our available data.

In Sections 3.5.1-3.5.3 we present the fit results, including our values for the partially quenched NLO and NNLO LECs. In Section 3.5.4 we examine the range of applicability of NNLO $SU(2)$ ChPT and the relative sizes of the terms in the chiral expansion. Finally, in Section 3.5.5 we compute the unquenched $SU(2)$ ChPT LECs from these results, and also discuss other predictions we can make from $SU(2)$ ChPT. All fits discussed in this section were performed by minimizing the uncorrelated χ^2; in Appendix 3.D we repeat the fits using a weighted χ^2 to explore systematic effects associated with correlations in the data. These weighted fits are also defined by Eqn. (3.16), but rather than choosing $\alpha_e \gg 1$ to overweight the physical point ensembles as we did in Ref. [10], here we underweight the 24I, 32I, and 32ID ensembles by a factor $\alpha_e = 1/N_e$, where N_e is the number of nondegenerate (partially quenched) pseudoscalar mass measurements on ensemble e. This has the effect of capturing some of the most important correlations — those between partially quenched measurements with different combinations of valence quarks on a given ensemble, and between reweightings in m_h of the same observable — as we argue in Appendix 3.D, while avoiding the numerical instabilities that plague fully correlated fits.

3.5.1 Fit Parameters

Tables 3.4 - 3.7 summarize the fit parameters, including a statistical error computed with the superjackknife resampling technique [10]. These include the χ^2/dof, physical quark masses, and inverse lattice spacings in physical units (Table 3.4), the ratios of quark masses and lattice spacings between the reference 32I ensemble and the other ensembles (Table 3.5), the PQChPT LECs (Table 3.6), and additional fit parameters describing the continuum and chiral scaling of the kaon

and Ω baryon data (Table 3.7). We generally observe excellent consistency comparing ensemble properties across the fits we have performed — the physical quark masses and lattice spacings from Table 3.4, and the ratios of lattice scales from Table 3.5, for example — with the notable exception of the NLO fit with a 450 MeV cut, for which we observe systematic shifts outside our statistical errors. This is not surprising, however, since we do not expect NLO ChPT to accurately describe the lattice data up to such a heavy scale, and indeed we see a large increase in the χ^2/dof for this particular fit.

While NLO fits constrain the four LECS $\{\hat{L}_4^{(2)}, \hat{L}_5^{(2)}, \hat{L}_6^{(2)}, \hat{L}_8^{(2)}\}$, NNLO fits constrain nine NLO LECs — $\{\hat{L}_i^{(2)}\}_{i=0}^{8}$ — as well as eight linear combinations of twelve NNLO LECs, which are listed explicitly in Table 3.6. We have set $\hat{K}_{22}^{(2)} = \hat{K}_{27}^{(2)} = \hat{K}_{39}^{(2)} = \hat{K}_{40}^{(2)} = 0$ when we perform the fits for simplicity, so that each linear combination reduces to a single, independent LEC. We also impose the constraint[8] $\hat{L}_{11}^{(2)} = -l_4/4$, which is required for the PQChPT Lagrangian to reduce to the unquenched ChPT Lagrangian in the unitary limit [48]. We perform independent fits at the two chiral scales $\Lambda_\chi = 770\,\text{MeV}$ and $\Lambda_\chi = 1\,\text{GeV}$, and report the PQChPT LECs at both scales. Since $\hat{L}_7^{(2)}$ and $\hat{L}_8^{(2)}$ are scale-independent, comparing the results for the fit with $\Lambda_\chi = 770\,\text{MeV}$ and the fit with $\Lambda_\chi = 1\,\text{GeV}$ provides a further consistency check.

We note that the 32ID-M1 ensemble has previously appeared in Ref. [49], where a simple estimate of the lattice spacing — $a = m_\Omega/m_\Omega^{\text{PDG}}$, with m_Ω at the simulated heavy quark mass — was used to convert the spectrum from lattice units to physical units. We find a 10% discrepancy between this lattice spacing and the lattice spacings obtained from our global fits and reported in Table 3.4. This arises from the 33% difference between the input bare heavy quark mass $am_h = 0.0596$ and the physical bare heavy quark masses determined from the global fits (also reported in Table 3.4): there is an $\mathcal{O}(10\%)$ shift in the ratio $m_\Omega/m_\Omega^{\text{PDG}}$ when m_Ω is adjusted from the simulated point to the physical point.

We note that $Z_l = Z_h = R_a = 1$ by definition on the 32I ensemble. We have constrained $Z_l^{64I} = Z_h^{64I} = 1$ since the Möbius parameters and gauge coupling on the 64I ensemble have been

[8] We have experimented with fits where $\hat{L}_{11}^{(2)}$ is left as a free parameter, but we find that $\hat{L}_{11}^{(2)} \neq -l_4/4$ well outside of statistics.

chosen such that the 64I action is identical to the 32I action up to small chiral symmetry breaking effects. As we argue in Ref. [10], these chiral symmetry breaking effects lead to a small shift in the lattice spacings, so we do not constrain $R_a^{64I} = 1$. Likewise, we constrain $Z_l^{24I} = Z_l^{48I}$ and $Z_h^{24I} = Z_h^{48I}$ for the same reason, but do not set $R_a^{24I} = R_a^{48I}$.

The observation that $Z_l, Z_h \sim 0.7$ for the 32ID-M1 ensemble in Table 3.5 suggests that this lattice is at sufficiently strong coupling that the five-dimensional domain wall fermion fields are no longer tightly bound to the domain walls, and instead leak into the fifth (s) dimension. As a result, somewhat larger input masses are required to achieve the same effective mass for the physical four-dimensional quark fields defined on the domain walls. We choose to include this ensemble in our fits since we do not observe any significant systematics if it is removed, and it is our only ensemble with lighter-than-physical pions, which probes the regime where chiral curvature is most pronounced.

		NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
	χ^2/dof	0.36(10)	1.14(27)	0.21(9)	0.29(10)
	$N_{\text{parameters}}$	42	45	55	58
	N_{data}	668	889	668	889
	am_l^{phys}	-0.001774(82)	-0.001764(77)	-0.001772(81)	-0.001767(80)
24I	am_h^{phys}	0.03209(40)	0.03239(32)	0.03210(38)	0.03219(35)
	a^{-1}	1.784(14) GeV	1.781(12) GeV	1.784(13) GeV	1.782(13) GeV
	am_l^{phys}	0.000272(15)	0.000244(18)	0.000282(14)	0.000282(14)
32I	am_h^{phys}	0.02512(29)	0.02424(43)	0.02537(27)	0.02550(27)
	a^{-1}	2.360(17) GeV	2.405(22) GeV	2.349(16) GeV	2.344(16) GeV
	am_l^{phys}	-0.000098(20)	-0.000105(21)	-0.000098(20)	-0.000097(18)
32ID	am_h^{phys}	0.04652(58)	0.04633(61)	0.04637(53)	0.04624(50)
	a^{-1}	1.374(8) GeV	1.377(9) GeV	1.376(8) GeV	1.377(7) GeV
	am_l^{phys}	0.000091(32)	0.000059(32)	0.000098(32)	0.000095(32)
32I-fine	am_h^{phys}	0.01936(67)	0.01784(66)	0.01977(68)	0.01993(70)
	a^{-1}	3.079(44) GeV	3.176(48) GeV	3.059(44) GeV	3.051(43) GeV
	am_l^{phys}	0.000685(14)	0.000706(12)	0.000688(13)	0.000695(13)
48I	am_h^{phys}	0.03547(33)	0.03595(24)	0.03550(31)	0.03562(27)
	a^{-1}	1.737(8) GeV	1.726(6) GeV	1.736(7) GeV	1.733(6) GeV
	am_l^{phys}	0.000625(10)	0.000604(15)	0.0006352(92)	0.000635(10)
64I	am_h^{phys}	0.02556(23)	0.02486(40)	0.02579(21)	0.02590(21)
	a^{-1}	2.352(9) GeV	2.379(17) GeV	2.343(8) GeV	2.339(8) GeV
	am_l^{phys}	0.00094(12)	0.00110(12)	0.00087(11)	0.00086(11)
32ID-M1	am_h^{phys}	0.0823(35)	0.0860(32)	0.0800(30)	0.0797(30)
	a^{-1}	1.002(20) GeV	0.978(17) GeV	1.015(17) GeV	1.017(18) GeV
	am_l^{phys}	—	-0.003404(35)	—	-0.003367(37)
32ID-M2	am_h^{phys}	—	0.02486(97)	—	0.0255(11)
	a^{-1}	—	2.025(34) GeV	—	1.990(35) GeV

Table 3.4: The (uncorrelated) χ^2/dof, unrenormalized physical quark masses in bare lattice units (without m_{res} included), and the values of the inverse lattice spacing a^{-1} in physical units, obtained from fits to $SU(2)$ PQChPT with the stated pion mass cuts.

		NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
24I	Z_l	0.980(11)	0.959(11)	0.9842(97)	0.979(10)
	Z_h	0.9711(82)	0.950(10)	0.9756(78)	0.9770(73)
	R_a	0.7561(61)	0.7402(73)	0.7596(58)	0.7604(56)
32I	Z_l	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	Z_h	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	R_a	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
32ID	Z_l	0.9162(79)	0.908(11)	0.9212(76)	0.9186(84)
	Z_h	0.9157(66)	0.9028(99)	0.9218(61)	0.9258(59)
	R_a	0.5822(45)	0.5725(64)	0.5858(41)	0.5877(40)
32I-fine	Z_l	0.994(30)	0.995(31)	0.995(30)	1.001(30)
	Z_h	0.989(21)	1.021(20)	0.980(21)	0.978(21)
	R_a	1.305(16)	1.320(16)	1.302(16)	1.302(16)
48I	Z_l	0.980(11)	0.959(11)	0.9842(97)	0.979(10)
	Z_h	0.9711(82)	0.950(10)	0.9756(78)	0.9770(73)
	R_a	0.7360(69)	0.7174(76)	0.7391(65)	0.7393(62)
64I	Z_l	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	Z_h	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	R_a	0.9968(57)	0.9892(52)	0.9973(57)	0.9981(57)
32ID-M1	Z_l	0.708(15)	0.682(15)	0.720(14)	0.719(14)
	Z_h	0.719(15)	0.694(15)	0.733(13)	0.737(13)
	R_a	0.4246(83)	0.4067(77)	0.4321(74)	0.4338(74)
32ID-M2	Z_l	—	1.013(13)	—	1.013(16)
	Z_h	—	1.009(14)	—	1.028(18)
	R_a	—	0.8419(97)	—	0.849(12)

Table 3.5: Ratios of lattice spacings (R_a) and light and heavy quark masses (Z_l, Z_h) between each ensemble and the reference 32I ensemble.

LEC	Λ_χ	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
B	—	4.229(35) GeV	4.270(41) GeV	4.189(43) GeV	4.203(44) GeV
f		0.1213(15) GeV	0.1236(20) GeV	0.1207(17) GeV	0.1215(16) GeV
$10^3\hat{L}_0^{(2)}$	1 GeV	—	—	-3.8(2.5)	1.0(1.1)
$10^3\hat{L}_1^{(2)}$		—	—	0.52(71)	-0.62(52)
$10^3\hat{L}_2^{(2)}$		—	—	-4.1(1.7)	0.06(74)
$10^3\hat{L}_3^{(2)}$		—	—	1.1(1.4)	-1.56(87)
$10^3\hat{L}_4^{(2)}$		-0.211(79)	-0.038(51)	-0.31(25)	-0.56(22)
$10^3\hat{L}_5^{(2)}$		0.438(72)	0.501(43)	0.37(34)	0.60(28)
$10^3\hat{L}_6^{(2)}$		-0.175(48)	-0.054(31)	-0.19(13)	-0.38(10)
$10^3\hat{L}_7^{(2)}$		—	—	-1.30(48)	-0.75(27)
$10^3\hat{L}_8^{(2)}$		0.594(36)	0.581(22)	0.52(16)	0.69(13)
$10^3\hat{L}_0^{(2)}$	770 MeV	—	—	-3.7(2.8)	1.1(1.1)
$10^3\hat{L}_1^{(2)}$		—	—	0.63(90)	-0.52(53)
$10^3\hat{L}_2^{(2)}$		—	—	-3.9(2.0)	0.27(78)
$10^3\hat{L}_3^{(2)}$		—	—	1.3(1.3)	-1.42(85)
$10^3\hat{L}_4^{(2)}$		-0.004(79)	0.169(51)	-0.10(27)	-0.35(22)
$10^3\hat{L}_5^{(2)}$		0.852(72)	0.915(43)	0.78(35)	1.02(28)
$10^3\hat{L}_6^{(2)}$		-0.019(48)	0.101(31)	-0.04(14)	-0.23(10)
$10^3\hat{L}_7^{(2)}$		—	—	-1.30(52)	-0.75(26)
$10^3\hat{L}_8^{(2)}$		0.594(36)	0.581(22)	0.52(17)	0.69(13)
$10^6\left(\hat{K}_{17}^{(2)}-\hat{K}_{39}^{(2)}\right)$	1 GeV	—	—	-7.3(2.0)	-7.6(1.1)
$10^6\left(\hat{K}_{18}^{(2)}+6\hat{K}_{27}^{(2)}-\hat{K}_{40}^{(2)}\right)$		—	—	14.5(7.9)	19.2(4.7)
$10^6\hat{K}_{19}^{(2)}$		—	—	11(16)	-0.9(4.2)
$10^6\hat{K}_{20}^{(2)}$		—	—	-12(10)	-3.2(2.8)
$10^6\left(\hat{K}_{21}^{(2)}+2K_{22}^{(2)}\right)$		—	—	-7.6(6.9)	4.9(4.1)
$10^6\hat{K}_{23}^{(2)}$		—	—	-12.4(3.5)	-2.8(1.4)
$10^6\hat{K}_{25}^{(2)}$		—	—	6.7(4.5)	1.3(1.7)
$10^6\left(\hat{K}_{26}^{(2)}+6\hat{K}_{27}^{(2)}\right)$		—	—	3.2(6.8)	11.2(3.6)
$10^6\left(\hat{K}_{17}^{(2)}-\hat{K}_{39}^{(2)}\right)$	770 MeV	—	—	-6.2(1.5)	-5.33(77)
$10^6\left(\hat{K}_{18}^{(2)}+6\hat{K}_{27}^{(2)}-\hat{K}_{40}^{(2)}\right)$		—	—	8.3(6.6)	14.5(3.9)
$10^6\hat{K}_{19}^{(2)}$		—	—	3(12)	-3.9(2.3)
$10^6\hat{K}_{20}^{(2)}$		—	—	-5.0(7.5)	0.0(1.8)
$10^6\left(\hat{K}_{21}^{(2)}+2K_{22}^{(2)}\right)$		—	—	-6.8(7.0)	6.2(3.2)
$10^6\hat{K}_{23}^{(2)}$		—	—	-7.2(3.2)	-0.2(1.2)
$10^6\hat{K}_{25}^{(2)}$		—	—	2.4(3.4)	-1.0(1.1)
$10^6\left(\hat{K}_{26}^{(2)}+6\hat{K}_{27}^{(2)}\right)$		—	—	1.9(6.4)	10.1(3.1)

Table 3.6: $SU(2)$ PQChPT LECs fit at two different chiral scales — $\Lambda_\chi = 1\,\text{GeV}$ and $\Lambda_\chi = 770\,\text{MeV}$ — in units of the canonical size at a given order in the chiral expansion. The LECs $\hat{L}_7^{(2)}$ and $\hat{L}_8^{(2)}$ have no scale dependence. The value of B quoted here is unrenormalized.

Parameter	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
$m^{(K)}$	0.4863(27) GeV	0.4861(43) GeV	0.4862(24) GeV	0.4862(25) GeV
$f^{(K)}$	0.1501(17) GeV	0.1535(22) GeV	0.1490(17) GeV	0.1488(16) GeV
$10^3\lambda_1$	3.2(1.0)	3.64(98)	3.2(1.0)	3.3(1.0)
$10^3\lambda_2$	28.17(65)	28.45(65)	28.27(78)	28.76(74)
$10^3\lambda_3$	-3.9(1.1)	-3.22(98)	-3.8(1.1)	-3.9(1.0)
$10^3\lambda_4$	5.69(31)	5.82(32)	5.70(31)	5.83(33)
c_f^I	0.059(47) GeV^2	-0.028(51) GeV^2	0.081(48) GeV^2	0.065(45) GeV^2
c_f^{ID}	-0.013(17) GeV^2	-0.058(19) GeV^2	0.013(15) GeV^2	0.012(16) GeV^2
$c_{f^{(K)}}^I$	0.049(39) GeV^2	-0.035(38) GeV^2	0.070(41) GeV^2	0.069(36) GeV^2
$c_{f^{(K)}}^{ID}$	-0.005(15) GeV^2	-0.044(14) GeV^2	0.011(15) GeV^2	0.019(15) GeV^2
c_{m_h,m_π^2}	1.6(2.7)	0.1(2.2)	1.4(2.7)	0.9(2.1)
c_{m_h,f_π}	0.14(11)	0.061(89)	0.221(97)	0.257(80)
c_{m_y,m_K^2}	3.915(22) GeV	3.981(34) GeV	3.895(20) GeV	3.884(20) GeV
c_{m_h,m_K^2}	0.008(52) GeV	0.046(58) GeV	0.022(51) GeV	0.026(56) GeV
c_{m_y,f_K}	0.2926(62)	0.2983(59)	0.2906(64)	0.2987(56)
c_{m_h,f_K}	0.067(50)	0.073(52)	0.062(51)	0.096(48)
$m^{(\Omega)}$	1.6646(47) GeV	1.6643(91) GeV	1.6643(37) GeV	1.6644(36) GeV
c_{m_l,m_Ω}	3.54(74)	3.73(67)	3.68(74)	3.66(76)
c_{m_y,m_Ω}	5.650(59)	5.794(67)	5.585(55)	5.550(55)
c_{m_h,m_Ω}	2.31(62)	3.19(55)	1.83(61)	1.64(63)

Table 3.7: Additional fit parameters in physical units and adjusted to the physical strange quark mass. Here $\{m^{(K)}, f^{(K)}\}$ and $\{\lambda_i\}$ are the LO and NLO LECs of heavy-meson $SU(2)$ PQChPT evaluated at the chiral scale $\Lambda_\chi = 1\,\text{GeV}$. c_f^I and c_f^{ID} are the a^2 coefficients of f_π for the Iwasaki and Iwasaki+DSDR gauge actions, respectively, and likewise for $c_{f^{(K)}}^I$ and $c_{f^{(K)}}^{ID}$. The notation $c_{m_q,X}$ denotes the coefficient of a term linear in m_q for quantity X, and $m^{(\Omega)}$ is the constant term in the (linear) m_Ω ansatz. We emphasize that the distinction between "NLO" and "NNLO" fits, as well as the mass cut, applies only to m_π and f_π: the kaon and Ω baryon data and fit forms are the same in all of the fits.

3.5.2 Histograms

In Figure 3.1 we plot stacked histograms of the deviation of each data point Y_i from the fit prediction Y_i^{fit} in units of the standard deviation of the data σ_{Y_i}:

$$X_i \equiv \frac{Y_i - Y_i^{\text{fit}}}{\sigma_{Y_i}}. \tag{3.17}$$

This can be thought of as the signed square root of the contribution to χ^2 from each data point, where the sign indicates whether the fit is overestimating (-) or underestimating (+) the data. The distributions of m_π^2 and f_π, in particular, give an overall impression of how well partially quenched $SU(2)$ chiral perturbation theory truncated to a given order is able to describe all of our (in general partially quenched) lattice data. We observe excellent agreement between the data and the NLO fit when we use a pion mass cut of 370 MeV, however, when we raise the mass cut to 450 MeV, the NLO fit clearly starts to break down, as evidenced by the larger χ^2/dof and broader histogram with many 3σ and 4σ outliers. The NNLO ansatz appears to have no difficulty describing our full data set.

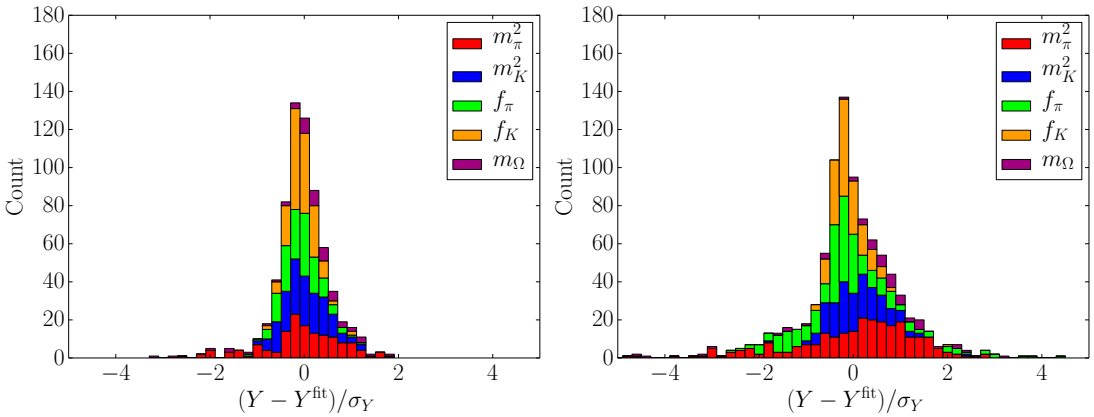

(a) NLO, $m_\pi^{\text{cut}} = 370\,\text{MeV}$

(b) NLO, $m_\pi^{\text{cut}} = 450\,\text{MeV}$

(c) NNLO, $m_\pi^{\text{cut}} = 370\,\text{MeV}$

(d) NNLO, $m_\pi^{\text{cut}} = 450\,\text{MeV}$

Figure 3.1: Stacked histograms of the signed deviation of the data from the fit in units of the standard deviation.

3.5.3 Unitary Chiral Extrapolation

In Figures 3.2 and 3.3 we overlay the unitary measurements of m_π^2/m_l and f_π on each ensemble with the ChPT prediction obtained using the LECs from each fit. The fit results have also been used to correct each lattice measurement from the simulated point to the continuum, infinite volume, and physical strange quark mass limit. The light quark mass has been renormalized in the $\overline{\text{MS}}$ scheme at 3 GeV using the renormalization coefficient computed in Ref. [10].

The influence of the NNLO terms is most clear in the chiral fits to f_π (Figure 3.3), which, in general, exhibit a more pronounced nonlinearity in the light quark mass than the chiral fits to m_π^2. While we observe that both m_π^2 and f_π are consistent between the NLO and NNLO fits with a mass cut of 370 MeV, when the mass cut is raised to 450 MeV the NLO and NNLO ansätze accommodate the additional heavy data differently. For the NLO case the entire m_π^2 and f_π curves are systematically shifted upward to higher energy — as one can see by comparing this fit to the adjacent NLO fit with $m_\pi^{\text{cut}} = 370\,\text{MeV}$ in Figures 3.2 and 3.3 — providing further evidence that this heavy data has extended into a regime where NLO PQChPT is no longer reliable. A similar comparison between the NNLO fits suggests that the heavy data influences these fits by smoothing out the curvature of f_π in the heavy mass regime $m_l^{\overline{\text{MS}}} \gtrsim 0.025\,\text{MeV}$.

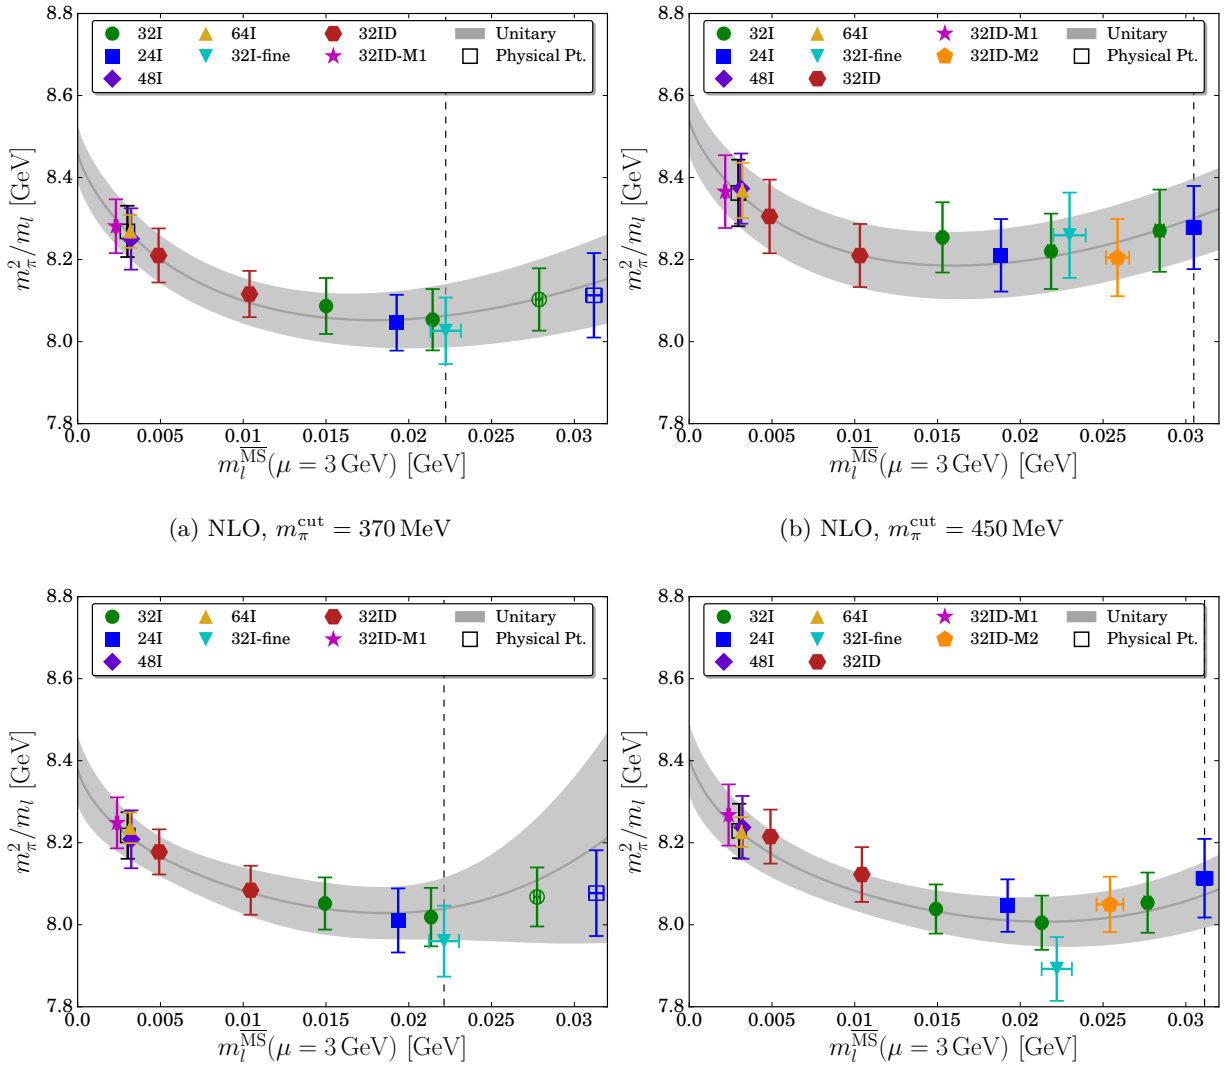

(a) NLO, $m_\pi^{\rm cut} = 370\,{\rm MeV}$

(b) NLO, $m_\pi^{\rm cut} = 450\,{\rm MeV}$

(c) NNLO, $m_\pi^{\rm cut} = 370\,{\rm MeV}$

(d) NNLO, $m_\pi^{\rm cut} = 450\,{\rm MeV}$

Figure 3.2: Chiral extrapolation of unitary m_π^2 data. The fit has been used to correct each data point from the simulated strange quark mass to the physical strange quark mass, as well as to take the infinite volume limit. Filled symbols correspond to sub-ensembles which were included in the fit, and open symbols correspond to sub-ensembles which were excluded from the fit based on the pion mass cut. The dashed vertical line corresponds to the heaviest unitary point included in the fit. "Physical point" is the prediction for the physical pion mass obtained by interpolating the fit to $m_l^{\rm phys}$.

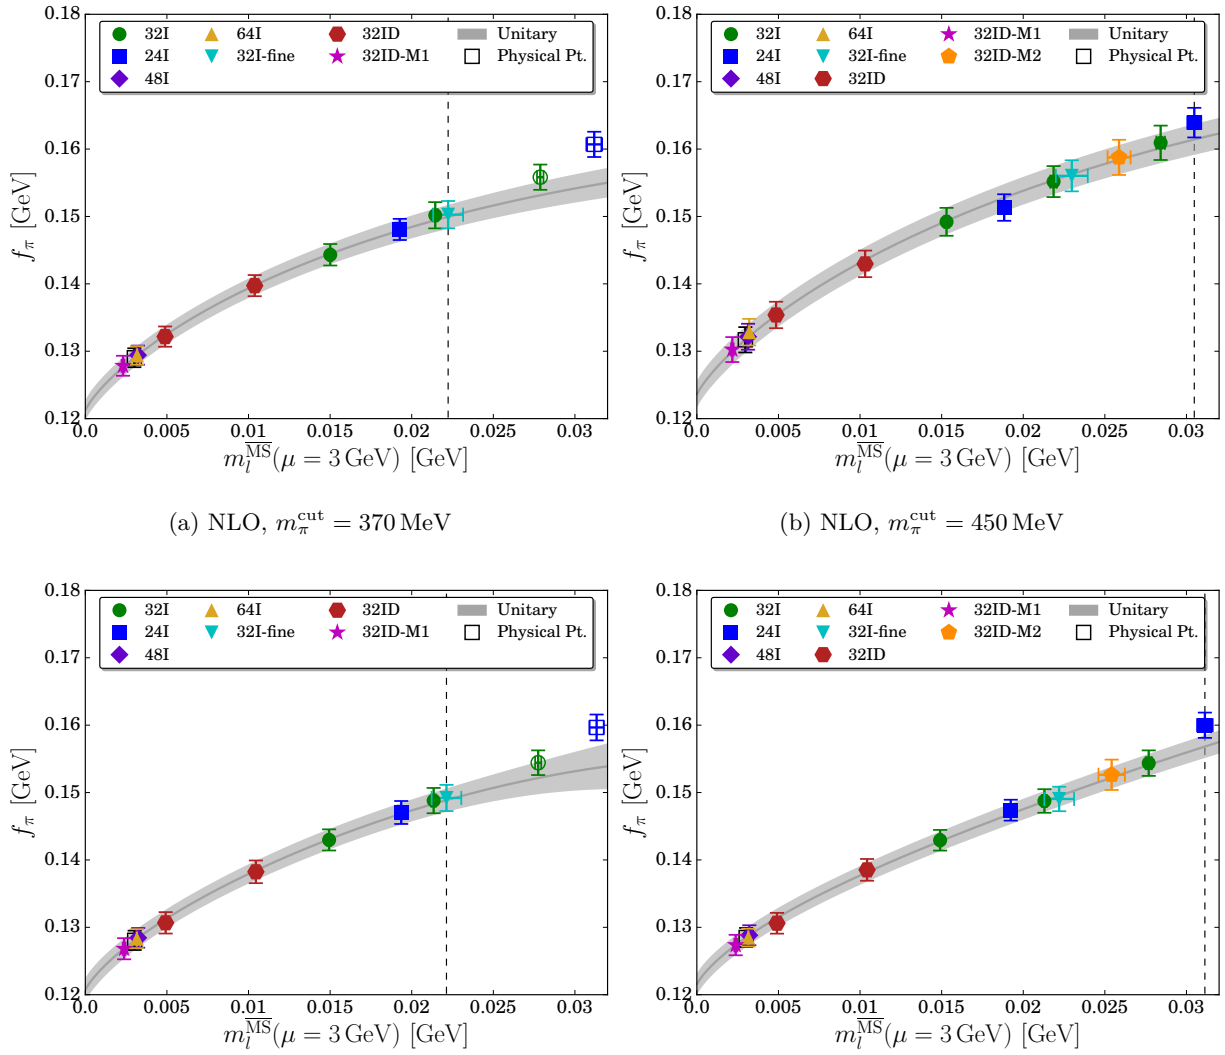

(a) NLO, $m_\pi^{\rm cut} = 370\,{\rm MeV}$

(b) NLO, $m_\pi^{\rm cut} = 450\,{\rm MeV}$

(c) NNLO, $m_\pi^{\rm cut} = 370\,{\rm MeV}$

(d) NNLO, $m_\pi^{\rm cut} = 450\,{\rm MeV}$

Figure 3.3: Chiral extrapolation of unitary f_π data. The fit has been used to correct each data point from the simulated strange quark mass to the physical strange quark mass, as well as to take the infinite volume and continuum limits. Filled symbols correspond to sub-ensembles which were included in the fit, and open symbols correspond to sub-ensembles which were excluded from the fit based on the pion mass cut. The dashed vertical line corresponds to the heaviest unitary point included in the fit. "Physical point" is the prediction for the physical pion decay constant obtained by interpolating the fit to $m_l^{\rm phys}$.

3.5.4 Chiral Expansion

Chiral perturbation theory is an effective field theory with an asymptotic series expansion. For ChPT to have any practical use it must be applied in a regime where the expansion is *well-ordered* in the sense that $|\text{LO}| > |\text{NLO}| > |\text{NNLO}| > \cdots$, since calculations beyond one or two loops are generally intractable, and higher order terms must be neglected. This in turn restricts the range of quark masses for which ChPT is applicable. While the very light masses of the up and down quarks suggest that the $SU(2)$ expansion ought to be well-ordered at the physical point, one expects that there is an upper limit, beyond which the N^{3}LO and higher order terms can no longer be discarded if one expects ChPT to describe low-energy QCD with high precision. In this section we use our NNLO fits to probe this scale.

In Figure 3.4 we plot the relative sizes of the LO, NLO, and NNLO terms for the pion mass and decay constant as a function of the liqht quark mass, using the LECs from Table 3.6. The heaviest unitary ensemble included in the fit is indicated with a dashed vertical line. We observe that the NLO and NNLO terms contribute to m_π^2 with opposite sign, but to f_π with the same sign: this behavior is expected from the lattice data, which suggests that the tree-level prediction $m_\pi^2 \propto m_l$ works reasonably well even for heavier-than-physical m_l, but not for the markedly nonlinear f_π. We also observe that the NNLO terms are generally statistically consistent with zero for the fit with the lighter mass cut, indicating that the ensembles with $m_\pi \gtrsim 350\,\text{MeV}$ are important for constraining the NNLO terms in our fits. This should be viewed as an artifact of our data set rather than a statement about $SU(2)$ chiral perturbation theory; one ought to be able to constrain the LECs to any order with data arbitrarily close to the chiral limit provided one has enough high-precision measurements[9]. Both mass cuts give consistent results for $m_l/m_l^{\text{phys}} \lesssim 8.0$, where the fits are directly constrained by lattice data. At the physical point we find

$$\begin{aligned} \frac{m_\pi^2}{\chi_l} &= 1.0000 - 0.0245(41) + 0.0034(10) \\ \frac{f_\pi}{f} &= 1.0000 + 0.0586(35) - 0.0011(7) \end{aligned} \tag{3.18}$$

[9] In fact, one could argue that the mass cut should be taken so that only the lightest quark masses are used since systematic deviations between the predictions of ChPT and full QCD vanish in the chiral limit.

for the decomposition into LO+NLO+NNLO, normalized by LO. The errors on the more restrictive fit quickly grow when we extrapolate to heavier m_l, so we focus on the $m_\pi^{\rm cut} = 450\,\text{MeV}$ result to test the breakdown of the expansion at heavy quark masses.

While both the NLO and NNLO terms remain small relative to LO — at most $\mathcal{O}(20\%)$ — even up to very heavy $m_\pi \sim 500\,\text{MeV}$, the NLO and NNLO terms start to become comparable in size for $m_\pi \gtrsim 450\,\text{MeV}$. In figure 3.5 we plot the ratios NLO/LO and NNLO/NLO as a function of the light quark mass. If we conservatively define "distress" in the chiral expansion as $|\text{NNLO}| \simeq 0.5|\text{NLO}|$ within statistical error, we find that this corresponds to $m_l/m_l^{\rm phys} \approx 10.9$ ($m_\pi \approx 445\,\text{MeV}$) for f_π. A more relaxed definition of $|\text{NNLO}| \simeq 0.8|\text{NLO}|$ corresponds to $m_l/m_l^{\rm phys} \approx 14.2$ ($m_\pi \approx 520\,\text{MeV}$). The situation for m_π^2 is more subtle: while it is true that we similarly observe an increase in the relative sizes of the NNLO and NLO terms as the light quark mass is increased, they are contributing with opposite sign, and the sum NLO + NNLO remains less than 10% of the LO contribution even at very heavy $m_\pi \gtrsim 500\,\text{MeV}$. We conclude that it is f_π, which exhibits stronger nonlinearity than m_π^2, that sets an upper limit on the applicability of NNLO $SU(2)$ ChPT, of roughly $m_\pi \sim 450 - 500\,\text{MeV}$. We note that the BMW collaboration has performed a similar test by fitting $SU(2)$ ChPT to unitary lattice data computed with $\mathcal{O}(a)$-improved Wilson fermions up to $m_\pi \sim 500\,\text{MeV}$, and finds results consistent with our own [13].

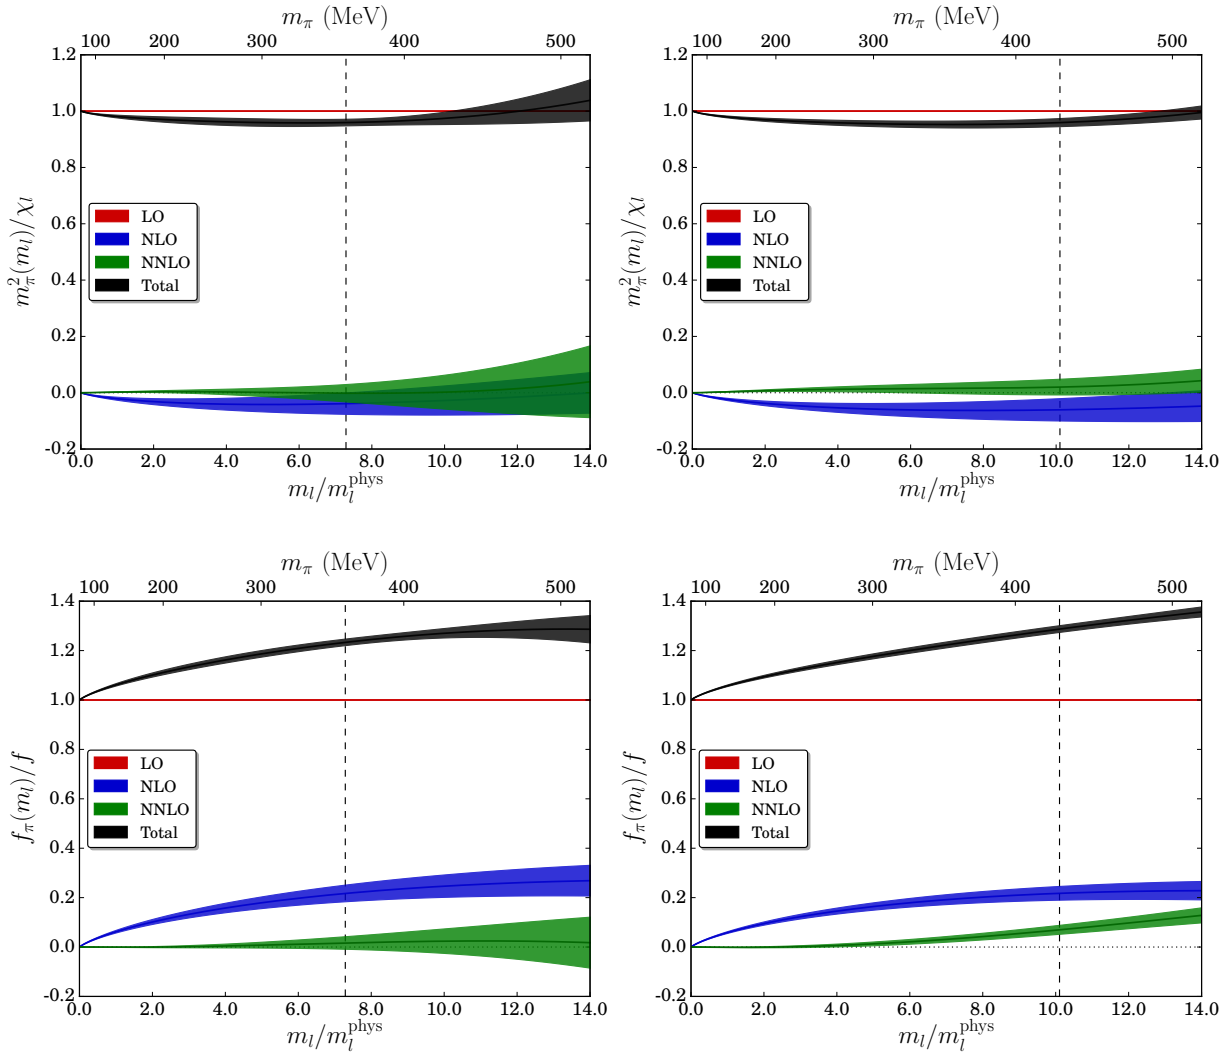

Figure 3.4: Decomposition of the $SU(2)$ chiral expansion into LO, NLO, and NNLO terms, normalized by LO. The pion mass (top) and pion decay constant (bottom) are plotted as a function of the light quark mass, using the LECs obtained from a fit with a pion mass cut of 370 MeV (left) and 450 MeV (right). The vertical dashed line corresponds to the heaviest unitary point included in the fit, and the horizontal dotted line marks zero.

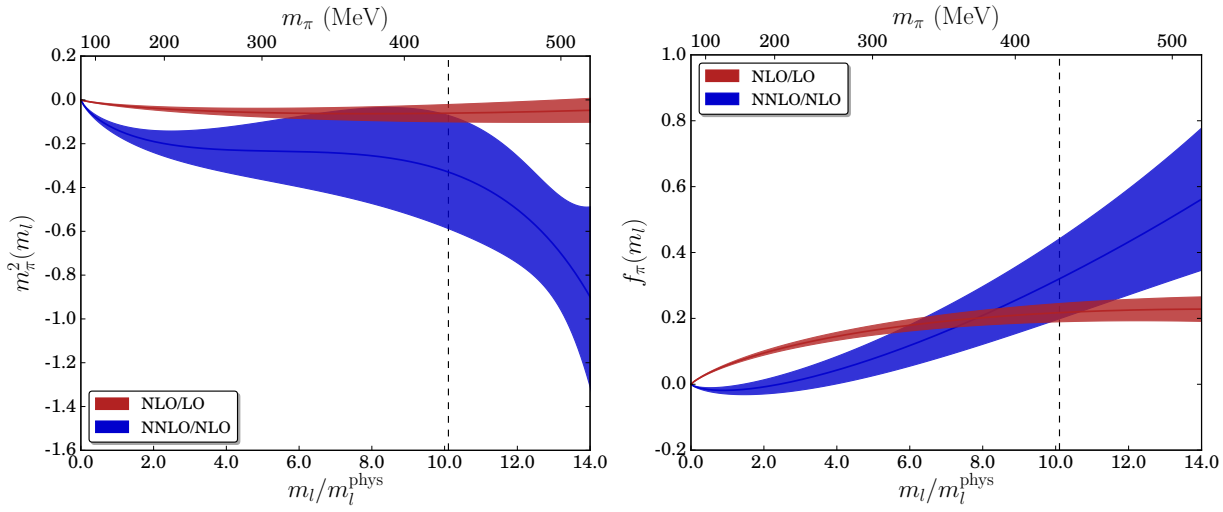

Figure 3.5: Relative sizes of the LO, NLO, and NNLO terms in the $SU(2)$ chiral expansion for m_π^2 (left) and f_π (right) using the LECs obtained from a fit with a pion mass cut of 450 MeV. The vertical dashed line corresponds to the heaviest unitary point included in the fit.

3.5.5 Predictions

Unquenched LECs

In Table 3.8 we use the relations listed in Appendix 3.A.1 to compute the unquenched $SU(2)$ LECs $\{l_i\}_{i=1}^7$ which can be determined from the partially quenched LECs in Table 3.6. Traditionally, values for the scale independent LECs $\{\overline{\ell}_i\}_{i=1}^6$ are quoted rather than $\{l_i\}_{i=1}^6$; we also compute these using relations listed explicitly in Appendix 3.A.2. There is no analogous $\overline{\ell}_7$ since l_7 is already scale independent. We also compute the renormalized leading order LEC B in the $\overline{\text{MS}}$ scheme at $\mu = 2.0\,\text{GeV}$, and the $\overline{\text{MS}}$ renormalized quark condensate

$$\Sigma = -\left\langle \overline{\psi}_l \psi_l \right\rangle\big|_{m_l \to 0} = \frac{B f^2}{2}. \tag{3.19}$$

We use the renormalization coefficients Z_{ml} computed in Ref. [10] to first renormalize B and Σ in the SMOM and SMOM$_{\gamma^\mu}$ schemes, which are then matched perturbatively to $\overline{\text{MS}}$. The difference in central value between the two intermediate schemes is used to assign a systematic error associated with the renormalization procedure.

In Figures 3.6 and 3.7 we compare our preferred determinations of the leading order and next-to leading order unquenched $SU(2)$ LECs (blue circles) to the 2013 $N_f = 2 + 1$ FLAG lattice averages [8, 9, 12, 50–53] (black squares) and two phenomenological fits (green diamonds): the first is Gasser and Leutwyler's original determination of the $SU(2)$ LECs in Ref. [5], and the second is Colangelo et al.'s updated fit of experimental pion scattering and scalar charge radius data to NNLO $SU(2)$ ChPT and the Roy equations [54]. We also include our final prediction for each LEC, including the full statistical and systematic error budget discussed in Section 3.6 summed in quadrature ("prediction"). For consistency with FLAG we quote our values for the dimensionless ratio f_π/f rather than f.

We generally observe excellent consistency between our fits, and find that our results for the LO LECs, $\overline{\ell}_3$, and $\overline{\ell}_4$ — which by now are standard lattice calculations — compare favorably with the FLAG averages and phenomenological fits. We find that $\overline{\ell}_3$ and $\overline{\ell}_4$ are determined more precisely by the NLO fits than the NNLO fits, which is not surprising: at two-loop order the NLO LECs can enter into the expressions for the pion mass and decay constant quadratically or as terms which are

LEC	Λ_χ	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
$B^{\overline{\rm MS}}(\mu = 2\,{\rm GeV})$		2.804(34)(30) GeV	2.831(37)(30) GeV	2.778(40)(30) GeV	2.787(40)(30) GeV
f	—	121.3(1.5) MeV	123.6(2.0) MeV	120.7(1.7) MeV	121.5(1.6) MeV
$\Sigma^{1/3,\,\overline{\rm MS}}(\mu = 2\,{\rm GeV})$		274.2(2.8)(1.0) MeV	278.6(3.8)(1.0) MeV	272.5(3.0)(1.0) MeV	274.0(2.8)(1.0) MeV
$10^3 l_1$		—	—	11.9(9.6)	-7.6(3.9)
$10^3 l_2$		—	—	-32(17)	4.3(6.8)
$10^3 l_3$	1 GeV	1.89(30)	2.08(21)	2.1(1.0)	1.46(78)
$10^3 l_4$		0.06(51)	1.70(34)	-1.0(1.6)	-2.07(94)
$10^3 l_7$		—	—	16.6(7.3)	6.5(3.8)
$10^3 l_1$		—	—	13(11)	-7.1(4.0)
$10^3 l_2$		—	—	-31(19)	5.4(6.9)
$10^3 l_3$	770 MeV	1.07(30)	1.25(21)	1.3(1.0)	0.63(78)
$10^3 l_4$		3.38(51)	5.01(34)	2.3(1.6)	1.24(95)
$10^3 l_7$		—	—	16.6(7.9)	6.5(3.7)
$\bar{\ell}_1$		—	—	15.3(9.1)	-3.2(3.7)
$\bar{\ell}_2$		—	—	-11.0(7.9)	6.0(3.2)
$\bar{\ell}_3$	—	2.81(19)	2.69(13)	2.66(64)	3.08(49)
$\bar{\ell}_4$		4.015(81)	4.274(54)	3.84(25)	3.68(15)

Table 3.8: Unquenched $SU(2)$ LECs computed from partially quenched $SU(2)$ fits. Missing entries are not constrained by the fits at a given order. For B and Σ the first error is statistical and the second is a systematic uncertainty in the perturbative matching to $\overline{\rm MS}$.

a product of an LEC and a chiral logarithm, whereas at one-loop order they enter only as simple linear, analytic terms.

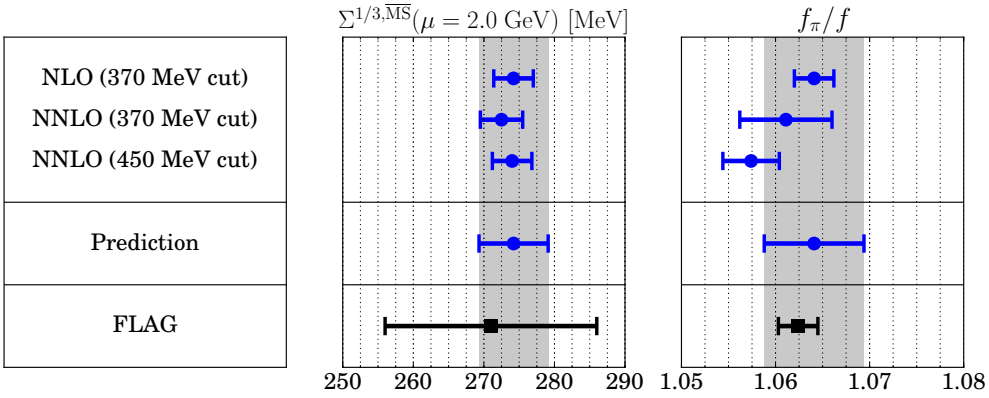

Figure 3.6: Leading order $SU(2)$ ChPT LECs compared to the 2013 FLAG lattice averages.

From our NNLO fits we are also able to constrain $\bar{\ell}_1$, $\bar{\ell}_2$, and the scale-independent NLO LEC l_7. This is, to the authors' knowledge, the first direct prediction for l_7: Gasser and Leutwyler provide the order of magnitude estimate $l_7 \sim 5 \times 10^{-3}$ [5], which is consistent with our predictions (*e.g.* $l_7 = 6.5(3.7) \times 10^{-3}$ from the fit with a 450 MeV cut). While our results for $\bar{\ell}_1$ and $\bar{\ell}_2$ are consistent with the phenomenological results, these LECs are determined much more precisely by the $\pi\pi$ scattering-based phenomenological fits. In this sense the lattice and phenomenological results are nicely complementary. We have begun to sharpen our predictions for $\bar{\ell}_1$ and $\bar{\ell}_2$ by including additional observables — *e.g.* $\pi\pi$ scattering lengths and pion form factors — which can be computed on the lattice and provide stronger constraints on these LECs [55].

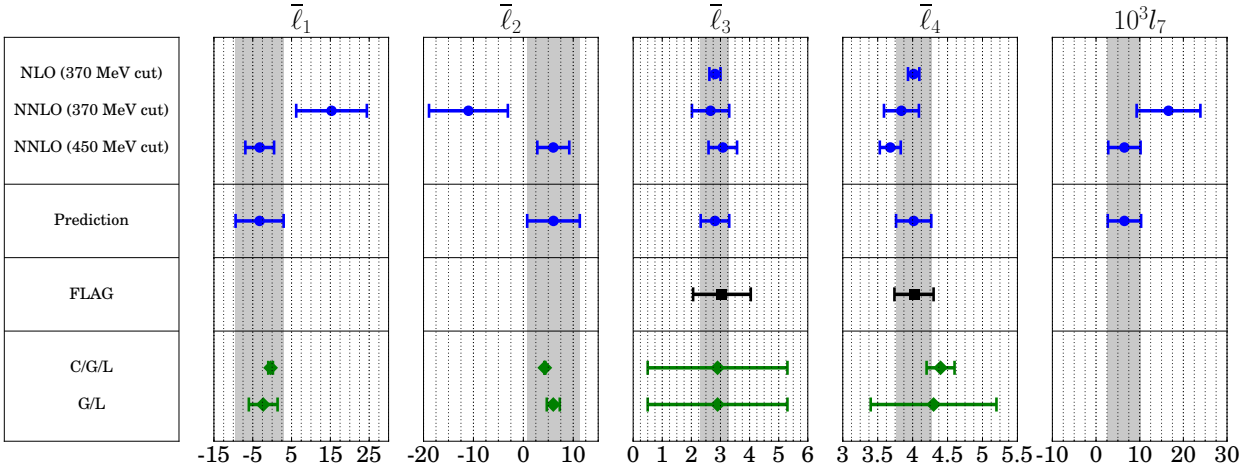

Figure 3.7: Next-to leading order $SU(2)$ ChPT LECs compared to the 2013 FLAG lattice averages and two phenomenological determinations.

Other Physical Predictions

Table 3.9 summarizes a number of predictions based on our results for the $SU(2)$ LECs from the previous section: f_π, f_K, and the ratios f_K/f_π and f_π/f are obtained directly from the global fit by interpolating our lattice results to the physical point. The final three quantities — the $I = 0$ (a_0^0) and $I = 2$ (a_0^2) $\pi\pi$ scattering lengths, and the pion mass splitting due to QCD isospin breaking effects — are one-loop ChPT predictions computed using Appendix 3.A.3 and the values of the LECs $\{\bar{\ell}_i\}_{i=1}^4$ and l_7 from Table 3.8.

	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
f_π	0.1290(14) GeV	0.1317(19) GeV	0.1281(14) GeV	0.1285(14) GeV
f_K	0.1540(15) GeV	0.1575(18) GeV	0.1530(15) GeV	0.1527(14) GeV
f_K/f_π	1.1937(54)	1.1962(74)	1.1944(78)	1.1884(67)
f_π/f	1.0641(21)	1.0658(21)	1.0611(49)	1.0574(30)
$m_\pi a_0^0$	—	—	0.170(20)	0.1987(86)
$m_\pi a_0^2$	—	—	-0.0577(90)	-0.0404(33)
$[m_{\pi^\pm}^2 - m_{\pi^0}^2]_{\rm QCD}/\Delta m_{du}^2$	—	—	80(35)	31(17)

Table 3.9: Predictions from NLO and NNLO fits and $SU(2)$ ChPT. $\Delta m_{du} \equiv m_d - m_u$. We emphasize that the distinction between "NLO" and "NNLO" fits, as well as the mass cut, applies only to m_π and f_π: the kaon and Ω baryon data and fit forms are the same in all of these fits.

The RBC-UKQCD collaboration has historically observed that, if f_π and f_K are determined from fits to heavy lattice data which is extrapolated down to the physical point, the predictions for f_π and f_K are systematically low compared to the physical values $f_\pi^{\rm phys} = 130.7\,{\rm MeV}$ and $f_K^{\rm phys} = 156.1\,{\rm MeV}$, which we observe in Table 3.9 as well. We have also found, however, that either overweighting the contributions to χ^2 from the physical pion mass 48I and 64I ensembles[10] or normalizing the contributions to χ^2 from each ensemble by the number of partially quenched measurements performed on that ensemble — effectively underweighting the heavy pion mass 24I and 32I ensembles — as we explore in Appendix 3.D, removes this discrepancy, and results in predictions for f_π and f_K consistent with their physical values. We conclude that two effects are responsible: 1) the large number of partially quenched measurements on the 24I, 32I, and 32ID ensembles causes the heavier data to dominate an unweighted, uncorrelated fit, and 2) chiral fits which are dominated by heavy data can exhibit excessive curvature near the physical point, leading to predictions which are systematically low.

The last three predictions in Table 3.9 allow for an interesting test of chiral perturbation theory. Since our NNLO fits determine the LECs $\bar{\ell}_1$, $\bar{\ell}_2$, and l_7 without containing any direct information

[10]This procedure was introduced in Ref. [10] to make small corrections for quark mass mistunings on the physical point ensembles.

about $\pi\pi$ scattering or isospin breaking, we can compute these quantities to NLO as predictions from our fits. The $\pi\pi$ scattering lengths computed from our preferred NNLO fit with a 450 MeV cut can be compared to recent experimental results based on measurements of K_{e4} and $K^{\pm} \to \pi^{\pm}\pi^0\pi^0$ decays: $m_\pi a_0^0 = 0.221(5)$ and $m_\pi a_0^2 = -0.043(5)$ [56]. Our prediction for the $\pi^{\pm} - \pi^0$ mass splitting is less straightforward to interpret directly since the largest contribution to the physical splitting arises from electromagnetic effects which we do not take into account. If we take a reasonable estimate of the up/down mass difference $\Delta m_{du} \equiv m_d - m_u \sim 2.5\,\mathrm{MeV}$, we can compare our prediction — $[m^2_{\pi^\pm} - m^2_{\pi^0}]_{\mathrm{QCD}} = 195(112)\,\mathrm{MeV}^2$ from the fit with the heavier mass cut — to the physical mass difference $m^2_{\pi^\pm} - m^2_{\pi^0} = 1261\,\mathrm{MeV}^2$ [39], which suggests that $\sim 15(9)\%$ of the total mass splitting arises from QCD isospin breaking effects. When combined with the leading-order prediction for the electromagnetic corrections computed by Bijnens and Danielsson in partially quenched ChPT [57], $[m^2_{\pi^\pm} - m^2_{\pi^0}]_{\mathrm{EM}} = 1000\,\mathrm{MeV}^2$, we find excellent agreement with the physical mass splitting.

3.6 Error Budget and Final Results for the Unquenched $SU(2)$ LECs

In this section we discuss the error budget for our determination of the leading and next-to leading order unquenched $SU(2)$ low energy constants, and report our final values including all systematics. In particular, we assign the following error to each LEC in table 3.10:

- *Influence of heavy data as determined by underweighting correlated data in the fits*: While our global fits are uncorrelated, we know that the partially quenched measurements on a given ensemble are highly correlated since they are computed with the same set of field configurations. If we were fitting to a function which exactly represented our data, as opposed to an expansion with some limited precision, our uncorrelated fits would not introduce any systematic bias into our answers. Since this is not the case, changing the weighting of the heavy mass ensembles, which contain highly correlated partially quenched measurements, gives us an estimate of the systematic effects on our results due to the worsening systematic disagree-

ment betweeen PQChPT and QCD at heavier quark masses. We estimate the impact on our fits by taking the difference in central value between the LECs of an unweighted, uncorrelated fit (Section 3.5) and the LECs of a fit where the contributions to χ^2 from ensembles with multiple partially quenched measurements have been systematically underweighted to capture the dominant effects of correlations (Appendix 3.D).

We also assign additional errors to the LECs which are determined by both NLO and NNLO fits (B, f, Σ, $\overline{\ell}_3$, and $\overline{\ell}_4$):

- *Influence of heavy data as determined by varying the mass cut*: We also estimate the dependence of the LECs on the choice of mass cut by taking the difference in central value between an NNLO fit with a unitary pion mass cut of 370 MeV and an NNLO fit with a cut of 450 MeV where applicable. For the LECs where we can estimate the influence of the heavy data using both methods we take the larger estimate as the systematic included in our error budget.

- *Truncation of the (continuum) chiral expansion*: We estimate the influence of truncating N^{3}LO and higher terms by taking the difference in central value between an NLO fit and an NNLO fit, both with a unitary pion mass cut of 370 MeV.

- *Finite volume effects*: As a conservative bound on the influence of NNLO and higher order FV corrections, as well as neglected cross terms — *e.g.* (NLO continuum ChPT) $\times$ (NLO FV correction) — we compute the difference in central value between an NLO PQChPT fit with NLO FV corrections and an NLO PQChPT fit with no FV corrections, both with a unitary pion mass cut of 370 MeV.

We do not attempt to quantify the latter set of systematics for the LECs which only enter into the $SU(2)$ ChPT expressions for the pion mass and decay constant at two loop order — $\overline{\ell}_1$, $\overline{\ell}_2$, and l_7 — since these LECs typically have $\mathcal{O}(50\%)$ or larger statistical errors, and are perhaps more accurately regarded as bounds than high-precision determinations. Likewise, we do not attempt to quantify systematic errors for the partially quenched LECs (Section 3.5.1) or for our predictions of the $\pi\pi$ scattering lengths and isospin breaking effects (Section 3.5.5), but one could, in principle, assign an analogous error budget.

$\boldsymbol{B^{\overline{\rm MS}}(\mu = 2\,\mathrm{GeV})}$	$\mathbf{2.804(34)(40)\,GeV}$
$\boldsymbol{f}$	$\mathbf{121.3(1.5)(2.1)\,MeV}$
$\boldsymbol{\Sigma^{1/3,\,\overline{\rm MS}}(\mu = 2\,\mathrm{GeV})}$	$\mathbf{274.2(2.8)(4.0)\,MeV}$
$\boldsymbol{f_\pi/f}$	$\mathbf{1.0641(21)(49)}$
$\bar{\ell}_1$	$-3.2(3.7)(5.0)$
$\bar{\ell}_2$	$6.0(3.2)(4.2)$
$\boldsymbol{\bar{\ell}_3}$	$\mathbf{2.81(19)(45)}$
$\boldsymbol{\bar{\ell}_4}$	$\mathbf{4.02(8)(24)}$
$10^3 l_7$	$6.5(3.8)(0.2)$

Table 3.10: Final predictions for the unquenched $SU(2)$ LECs including all statistical and systematic errors. The reported errors are the statistical (left) and the total systematic (right) obtained by summing the contributions we discuss in the text in quadrature. Bold entries correspond to LECs which enter into both NLO and NNLO fits, for which we assign the full error budget; for the other entries the mass cut, chiral truncation, and finite volume systematics are assumed to be negligible compared to the statistical error and are not quantified. The central values and statistical errors of B, f, $\Sigma^{1/3}$, $\bar{\ell}_3$, and $\bar{\ell}_4$ are from an NLO fit with a 370 MeV cut, while the central values and statistical errors of $\bar{\ell}_1$, $\bar{\ell}_2$, and l_7 are from an NNLO fit with a 450 MeV cut. We also include our prediction for the ratio f_π/f.

3.7 Conclusions

In this work we have performed fits of pseudoscalar masses and decay constants from a series of RBC-UKQCD domain wall fermion ensembles to the corresponding formulae in next-to-next-to leading order $SU(2)$ partially quenched chiral perturbation theory. We reported values for a large set of partially quenched low-energy constants, and used these values to compute the unquenched leading and next-to leading order LECs. We also examined the range of quark masses for which NLO and NNLO ChPT accurately describe our lattice data, and used the newly determined LECs from NNLO fits to make one-loop predictions for isospin breaking effects and $\pi\pi$ scattering lengths,

which we compare to other lattice and experimental results. We have observed that $SU(2)$ PQChPT generally describes the included range of partially quenched data with percent-scale accuracy: to emphasize this point we plot in Figure 3.8 histograms of the percent deviation between the data and fit

$$\Delta \equiv \frac{(Y - Y^{\text{fit}})}{(Y + Y^{\text{fit}})/2} \times 100 \tag{3.20}$$

for our preferred fits, NLO PQChPT with a unitary pion mass cut of 370 MeV and NNLO PQChPT with a 450 MeV cut.

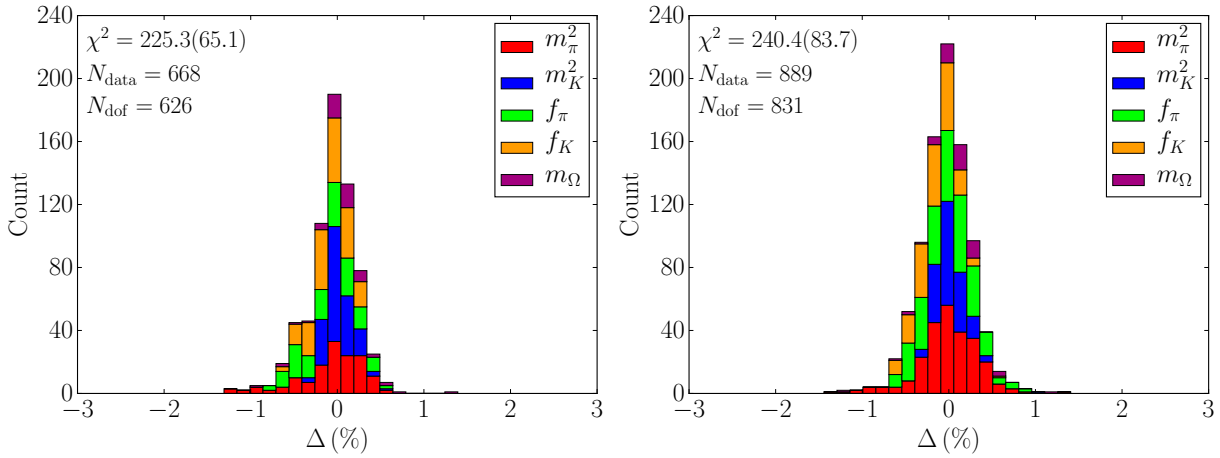

(a) NLO, $m_\pi^{\text{cut}} = 370\,\text{MeV}$ (b) NNLO, $m_\pi^{\text{cut}} = 450\,\text{MeV}$

Figure 3.8: Percent deviation between fits and data. We plot stacked histograms of the quantity $\Delta \equiv 200 \times (Y - Y^{\text{fit}})/(Y + Y^{\text{fit}})$.

We have observed that NNLO $SU(2)$ PQChPT can be reliably fit to our data for the pion mass and decay constant without the need for additional terms or constraints to stabilize the fits, and we determine values for 8 linear combinations of NNLO low energy constants. The values we obtained for the unquenched $SU(2)$ LECs were consistent between our NLO and NNLO fits and with other lattice and phenomenological determinations reported in the literature. At the physical light quark mass we found that the chiral expansions for the pion mass and decay constant behave like rapidly convergent series. After probing the breakdown of the chiral expansion at heavy light

quark mass we concluded that NLO $SU(2)$ PQChPT is sufficient to describe our lattice data up to $m_\pi \sim \mathcal{O}(350\,\mathrm{MeV})$, beyond which we observe some deviation between the NLO prediction for the pion decay constant and our lattice data. Likewise, we concluded that NNLO $SU(2)$ PQChPT remains consistent with our data up to $m_\pi \sim \mathcal{O}(450\,\mathrm{MeV})$. By 500 MeV, the NNLO corrections to the pion decay constant have grown to the point that they are comparable in size to the NLO corrections, indicating that the chiral expansion truncated to NNLO is unreliable at this scale. Of course, all statements regarding the values of LECs and the behavior of the $SU(2)$ chiral expansion made in this work are subject to the statistical precision, finite volume errors, and cutoff effects inherent in our lattice data. These points will need to be revisited and reassessed in the future as more and increasingly precise data becomes available.

We also note that our fits in this work only make use of the pseudoscalar masses and decay constants. Future work will incorporate a calculation of the $I = 2$ $\pi\pi$ scattering length and the pion vector form factor on many of the domain wall fermion ensembles considered here. Including these results in our chiral fits will give first-principles determinations of the scattering length a_0^2, the pion charge radius $\langle r^2 \rangle_V^\pi$, and the $SU(2)$ LEC $\overline{\ell}_6$, as well as sharpen the predictions for $\overline{\ell}_1$ and $\overline{\ell}_2$, which are currently determined most precisely by phenomenological fits to experimental data. A forthcoming paper will also explore analogous fits of the pseudoscalar masses and decay constants to $SU(3)$ partially quenched chiral perturbation theory at next-to-next-to leading order.

Acknowledgements

We thank members of the RBC-UKQCD collaboration for helpful discussions and support of this work.

The 32ID-M1 and 32ID-M2 ensembles were generated using the IBM Blue Gene/Q (BG/Q) "Mira" system at the Argonne Leadership Class Facility (ALCF), and the "Vulcan" BG/Q system at Lawrence Livermore National Laboratory (LLNL). Measurements, on these ensembles as well as older ensembles used in this work, were performed using Vulcan, the "DiRAC" BG/Q system in the Advanced Computing Facility at the University of Edinburgh, and the BG/Q computers of the RIKEN-BNL Research Center and Brookhaven National Lab.

The software used includes the CPS QCD code (`http://qcdoc.phys.columbia.edu/cps.html`) [58], supported in part by the USDOE SciDAC program, and the BAGEL (`http://www2.ph.ed.ac.uk/~paboyle/bagel/Bagel.html`) assembler kernel generator for high-performance optimized kernels and fermion solvers [59]. Gauge fixing was performed using the "GLU" (Gauge Link Utility) codebase (`https://github.com/RJhudspith/GLU`). Next-to-next-to leading order partially-quenched $SU(2)$ and $SU(3)$ chiral perturbation theory expressions for the pseudoscalar masses and decay constants were computed using Fortran routines provided by J. Bijnens. Gauge fixing and fitting were performed using the Columbia University CUTH cluster.

N.H.C., R.D.M., G.M., and D.J.M. are supported in part by U.S. DOE grant #DE-SC0011941. C.K. is supported by a RIKEN foreign postdoctoral research (FPR) grant. A.P. and C.T.S. were supported in part by UK STFC Grant ST/L000296/1. P.B. has been in part supported by STFC Grants ST/M006530/1, ST/L000458/1, ST/K005790/1, and ST/K005804/1. C.J. is supported by DOE contract #AC-02-98CH10886(BNL). N.G. is supported by the Leverhulme Research grant RPG-2014-118. A.J. acknowledges funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement No. 279757.

References

[1] Steven Weinberg, "Phenomenological Lagrangians", Physica A: Statistical Mechanics and its Applications **96**, 327–340 (1979).

[2] J. F. Donoghue, E. Golowich, and B. R. Holstein, *Dynamics of the standard model*, Cambridge Books Online (Cambridge University Press, 1992).

[3] T. Appelquist and J. Carazzone, "Infrared Singularities and Massive Fields", Phys. Rev. D **11**, 2856–2861 (1975).

[4] H. Georgi, "Effective Field Theory", Ann.Rev.Nucl.Part.Sci. **43**, 209–252 (1993).

[5] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory to One Loop", Annals Phys. **158**, 142 (1984).

[6] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory: Expansions in the Mass of the Strange Quark", Nucl.Phys. **B250**, 465 (1985).

[7] C. Allton et al., "Physical Results from 2+1 Flavor Domain Wall QCD and $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D78**, 114509 (2008).

[8] Y. Aoki et al., "Continuum Limit Physics from $2 + 1$ Flavor Domain Wall QCD", Phys. Rev. **D83**, 074508 (2011).

[9] R. Arthur et al., "Domain Wall QCD with Near-Physical Pions", Phys.Rev. **D87**, 094514 (2013).

[10] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[11] R. Mawhinney, "NLO and NNLO Chiral Fits for 2+1 Flavor DWF Ensembles", PoS **LAT2009**, 081 (2009).

[12] S. Borsanyi, S. Durr, Z. Fodor, S. Krieg, A. Schafer, et al., "$SU(2)$ Chiral Perturbation Theory Low-Energy Constants from 2+1 Flavor Staggered Lattice Simulations", Phys.Rev. **D88**, 014513 (2013).

[13] Dürr, Stephan and others, "Lattice QCD at the Physical Point Meets $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D90**, 114504 (2014).

[14] Brandt, Bastian B. and Jüttner, Andreas and Wittig, Hartmut, "The Pion Vector Form Factor from Lattice QCD and NNLO Chiral Perturbation Theory", JHEP **11**, 034 (2013).

[15] C. Kelly, R. Mawhinney, and D. Murphy, "The Low Energy Constants of $SU(3)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD", in preparation.

[16] H. Fearing and S. Scherer, "Extension of the Chiral Perturbation Theory Meson Lagrangian to Order p^6", Phys.Rev. **D53**, 315–348 (1996).

[17] J. Bijnens, G. Colangelo, G. Ecker, J. Gasser, and M. Sainio, "Pion Pion Scattering at Low-Energy", Nucl.Phys. **B508**, 263–310 (1997).

[18] G. Amoros, J. Bijnens, and P. Talavera, "Two Point Functions At Two Loops in Three Flavor Chiral Perturbation Theory", Nucl.Phys. **B568**, 319–363 (2000).

[19] G. Colangelo, S. Durr, and C. Haefeli, "Finite Volume Effects for Meson Masses and Decay Constants", Nucl. Phys. **B721**, 136–174 (2005).

[20] G. Colangelo and C. Haefeli, "Finite Volume Effects for the Pion Mass at Two Loops", Nucl. Phys. **B744**, 14–33 (2006).

[21] J. Bijnens and T. Rössler, "Finite Volume at Two-Loops in Chiral Perturbation Theory", JHEP **1501**, 034 (2015).

[22] C. W. Bernard and M. F. L. Golterman, "Partially quenched gauge theories and an application to staggered fermions", Phys. Rev. D **49**, 486–494 (1994).

[23] S. R. Sharpe, "Enhanced Chiral Logarithms in Partially Quenched QCD", Phys.Rev. **D56**, 7052–7058 (1997).

[24] Bijnens, Johan and Lähde, Timo A., "Masses and Decay Constants of Pseudoscalar Mesons to Two Loops in Two-Flavor Partially Quenched Chiral Perturbation Theory", Phys.Rev. **D72**, 074502 (2005).

[25] Bijnens, Johan and Danielsson, Niclas and Lähde, Timo A., "The Pseudoscalar Meson Mass to Two Loops in Three-Flavor Partially Quenched χPT", Phys.Rev. **D70**, 111503 (2004).

[26] Bijnens, Johan and Lähde, Timo A., "Decay Constants of Pseudoscalar Mesons to Two Loops in Three-Flavor Partially Quenched χPT", Phys.Rev. **D71**, 094502 (2005).

[27] Bijnens, Johan and Danielsson, Niclas and Lähde, Timo A., "Three-Flavor Partially Quenched Chiral Perturbation Theory at NNLO for Meson Masses and Decay Constants", Phys.Rev. **D73**, 074509 (2006).

[28] J. Bijnens, "Chiral Perturbation Theory Beyond One Loop", Prog.Part.Nucl.Phys. **58**, 521–586 (2007).

[29] Y. Iwasaki and T. Yoshie, "Renormalization Group Improved Action for $SU(3)$ Lattice Gauge Theory and the String Tension", Phys.Lett. **B143**, 449 (1984).

[30] P. M. Vranas, "Gap Domain Wall Fermions", Phys.Rev. **D74**, 034512 (2006).

[31] D. Renfrew, T. Blum, N. Christ, R. Mawhinney, and P. Vranas, "Controlling Residual Chiral Symmetry Breaking in Domain Wall Fermion Simulations", PoS **LATTICE2008**, 048 (2008).

[32] D. B. Kaplan, "A Method for Simulating Chiral Fermions on the Lattice", Phys. Lett. **B288**, 342–347 (1992).

[33] Y. Shamir, "Chiral Fermions from Lattice Boundaries", Nucl. Phys. **B406**, 90–106 (1993).

[34] R. C. Brower, H. Neff, and K. Orginos, "Mobius Fermions: Improved Domain Wall Chiral fermions", Nucl. Phys. Proc. Suppl. **140**, [,686(2004)], 686–688 (2005).

[35] R. Brower, H. Neff, and K. Orginos, "Möbius Fermions", Nucl.Phys.Proc.Suppl. **153**, 191–198 (2006).

[36] R. C. Brower, H. Neff, and K. Orginos, "The Möbius Domain Wall Fermion Algorithm", (2012).

[37] M. Lüscher, "Properties and Uses of the Wilson Flow in Lattice QCD", JHEP **08**, 071 (2010).

[38] A. Roessl, "Pion Kaon Scattering Near the Threshold in Chiral $SU(2)$ Perturbation Theory", Nucl.Phys. **B555**, 507–539 (1999).

[39] K.A. Olive and Particle Data Group, "Review of Particle Physics", Chinese Physics C **38**, 090001 (2014).

[40] K. Levenberg, "A Method for the Solution of Certain Non-Linear Problems in Least Squares", Quarterly Journal of Applied Mathmatics **II**, 164–168 (1944).

[41] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters", Journal of the Society for Industrial and Applied Mathematics **11**, 431–441 (1963).

[42] J. Bratt et al., "Nucleon Structure from Mixed Action Calculations Using 2+1 Flavors of Asqtad Sea and Domain Wall Valence Fermions", Phys.Rev. **D82**, 094502 (2010).

[43] G. Martinelli, C. Pittori, C. T. Sachrajda, M. Testa, and A. Vladikas, "A General Method for Nonperturbative Renormalization of Lattice Operators", Nucl.Phys. **B445**, 81–108 (1995).

[44] Y. Aoki, P. Boyle, N. Christ, C. Dawson, M. Donnellan, et al., "Non-Perturbative Renormalization of Quark Bilinear Operators and B_K Using Domain Wall Fermions", Phys.Rev. **D78**, 054510 (2008).

[45] R. Arthur and P. Boyle, "Step Scaling with Off-Shell Renormalisation", Phys.Rev. **D83**, 114511 (2011).

[46] C. Sturm, Y. Aoki, N. Christ, T. Izubuchi, C. Sachrajda, et al., "Renormalization of Quark Bilinear Operators in a Momentum-Subtraction Scheme with a Nonexceptional Subtraction Point", Phys.Rev. **D80**, 014501 (2009).

[47] R. Arthur, P. Boyle, N. Garron, C. Kelly, and A. Lytle, "Opening the Rome-Southampton Window for Operator Mixing Matrices", Phys.Rev. **D85**, 014501 (2012).

[48] J. Bijnens, G. Colangelo, and G. Ecker, "Renormalization of Chiral Perturbation Theory to Order p^6", Annals Phys. **280**, 100–139 (2000).

[49] T. Bhattacharya et al., "QCD Phase Transition with Chiral Quarks and Physical Quark Masses", Phys.Rev.Lett. **113**, 082001 (2014).

[50] S. Aoki, Y. Aoki, C. Bernard, T. Blum, G. Colangelo, et al., "Review of Lattice Results Concerning Low-Energy Particle Physics", Eur.Phys.J. **C74**, 2890 (2014).

[51] A. Bazavov et al., "MILC Results for Light Pseudoscalars", PoS **CD09**, 007 (2009).

[52] A. Bazavov et al., "Staggered Chiral Perturbation Theory in the Two-Flavor Case and $SU(2)$ Analysis of the MILC Data", PoS **LATTICE2010**, 083 (2010).

[53] S. R. Beane, W. Detmold, P. M. Junnarkar, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, A. Torok, and A. Walker-Loud, "$SU(2)$ Low-Energy Constants from Mixed-Action Lattice QCD", Phys. Rev. **D86**, 094509 (2012).

[54] G. Colangelo, J. Gasser, and H. Leutwyler, "$\pi\pi$ Scattering", Nucl.Phys. **B603**, 125–179 (2001).

[55] R. Mawhinney and D. Murphy, "NLO and NNLO Low Energy Constants for $SU(2)$ Chiral Perturbation Theory", PoS **LAT2015**, 061 (2015).

[56] B. Bloch-Devaux, "Precision Measurement of $\pi\pi$ Scattering Lengths in K_{e4} Decays by NA48/2", PoS **KAON09**, 033 (2009).

[57] J. Bijnens and N. Danielsson, "Electromagnetic Corrections in Partially Quenched Chiral Perturbation Theory", Phys. Rev. **D75**, 014505 (2007).

[58] C. Jung (RBC and UKQCD Collaborations), "Overview of Columbia Physics System", PoS **LAT2013**, 417 (2014).

[59] P. Boyle, "The BAGEL Assembler Generation Library", Comput.Phys.Commun. **180**, 2739–2748 (2009).

[60] S. Schaefer, R. Sommer, and F. Virotta, "Critical Slowing Down and Error Analysis in Lattice QCD Simulations", Nuclear Physics B **845**, 93–119 (2011).

[61] P. de Forcrand, M. Garcia Perez, and I.-O. Stamatescu, "Topology of the $SU(2)$ Vacuum: A Lattice Study Using Improved Cooling", Nucl.Phys. **B499**, 409–449 (1997).

[62] T. Blum, T. Izubuchi, and E. Shintani, "New Class of Variance-Reduction Techniques using Lattice Symmetries", Phys. Rev. **D88**, 094503 (2013).

[63] A. Stathopoulos and K. Orginos, "Computing and Deflating Eigenvalues While Solving Multiple Right Hand Side Linear Systems in Quantum Chromodynamics", SIAM J.Sci.Comput. **32**, 439–462 (2010).

3.A ChPT Relations

In this appendix we collect various relations used in the analysis in the body of the paper. We do not explicitly reprint the expressions for the pseudoscalar masses and decay constants and the corresponding finite volume corrections used in the chiral fits: instead we refer the reader to the appendices of Ref. [7]. The NNLO pseudoscalar masses and decay constants were computed using Fortran routines provided by J. Bijnens.

3.A.1 Relations Between PQChPT and ChPT LECs at NLO

The $SU(N_f)$ ChPT Lagrangian can be recovered from the more general $SU(N_f)$ PQChPT Lagrangian in the limit of equal sea and valence quark masses. Here we have collected the explicit expressions relating the NLO LECs in this limit from Ref. [48]. The analogous expressions for the NNLO LECs can be found in the same reference, but we do not use them here. For $N_f = 2$, the NLO ChPT LECs $\{l_i\}_{i=1}^{7}$ are related to the NLO PQChPT LECs $\{\hat{L}_i^{(2)}\}_{i=0}^{12}$ by

$$\begin{aligned}
l_1 &= -2\hat{L}_0^{(2)} + 4\hat{L}_1^{(2)} + 2\hat{L}_3^{(2)} & l_4 &= 4\left(2\hat{L}_4^{(2)} + \hat{L}_5^{(2)}\right) & & \\
l_2 &= 4\left(\hat{L}_0^{(2)} + \hat{L}_2^{(2)}\right) & l_5 &= \hat{L}_{10}^{(2)} & l_7 &= -8\left(2\hat{L}_7^{(2)} + \hat{L}_8^{(2)}\right) \\
l_3 &= 4\left(-2\hat{L}_4^{(2)} - \hat{L}_5^{(2)} + 4\hat{L}_6^{(2)} + 2\hat{L}_8^{(2)}\right) & l_6 &= -2\hat{L}_9^{(2)} & &
\end{aligned} \tag{C.21}$$

and the additional constraints $\hat{L}_{11}^{(2)} = -l_4/4$ and $\hat{L}_{12}^{(2)} = 0$.

3.A.2 Scale Independent $SU(2)$ LECs

Conventionally, one quotes values of the scale independent $SU(2)$ LECs $\{\bar{\ell}_i\}_{i=1}^{6}$ rather than $\{l_i\}_{i=1}^{6}$. These are obtained by running the $\{l_i\}_{i=1}^{6}$ from the energy scale at which they are defined, μ, to the physical pion mass using

$$\bar{\ell}_i = \gamma_i l_i - \log\left(\frac{m_\pi^2}{\mu^2}\right), \tag{C.22}$$

where the coefficients

$$\gamma_1 = 96\pi^2, \quad \gamma_2 = 48\pi^2, \quad \gamma_3 = -64\pi^2, \quad \gamma_4 = 16\pi^2, \quad \gamma_5 = -192\pi^2, \quad \gamma_6 = -96\pi^2, \tag{C.23}$$

were computed in Ref. [5]. The remaining LEC l_7 has no scale dependence.

3.A.3 One-Loop $SU(2)$ Predictions

While NLO fits to the pion mass and decay constant constrain the unquenched $SU(2)$ LECs l_3 and l_4, NNLO fits also constrain l_1, l_2, and l_7, allowing us to make additional one-loop predictions [5]. At NLO l_1 and l_2 determine quantities related to $\pi\pi$ scattering. The s-wave scattering lengths a_0^I

in the isospin channels $I = 0$ and $I = 2$, for example, are given by[11]

$$\begin{aligned} m_\pi a_0^0 &= \frac{7\chi_l}{16\pi f^2}\left[1 + \frac{16\chi_l}{7f^2}\left(5l_1 + 5l_2 + 3l_3\right) + \frac{\chi_l}{16\pi^2 f^2}\left(5 - 4\log\left(\frac{\chi_l}{\Lambda_\chi^2}\right)\right)\right] \\ m_\pi a_0^2 &= -\frac{\chi_l}{8\pi f^2}\left[1 - \frac{16\chi_l}{f^2}\left(l_1 + l_2\right) - \frac{\chi_l}{16\pi^2 f^2}\left(1 - 8\log\left(\frac{\chi_l}{\Lambda_\chi^2}\right)\right)\right] \end{aligned} . \tag{C.24}$$

The LEC l_7 controls the size of the pion mass splitting due to the difference between the up and down quark masses,

$$\left[m_{\pi^\pm}^2 - m_{\pi^0}^2\right]_{\text{QCD}} = \left(m_d - m_u\right)^2 \frac{4B^2}{f^2} l_7. \tag{C.25}$$

We use the subscript "QCD" to emphasize that this is only the contribution to the mass splitting from QCD isospin breaking. The dominant contribution is due to electromagnetic effects, and enters at $\mathcal{O}(m_d - m_u)$.

3.B Summary of Lattice Data Included in Chiral Fits

In this appendix we collect the results for fits of the pseudoscalar masses and decay constants, the Ω baryon mass, the ratio $R(t)$ (Eqn. (C.28)) which determines m_{res} in the chiral limit, and the Wilson flow scales on each ensemble in lattice units. Earlier results for the 24I ensemble can be found in Ref. [7], but differ from the current work in that the number of configurations has been approximately doubled and the spectrum re-analyzed in later works. For the other ensembles, these fits are identical to results we have published in earlier analyses: these can be found in Ref. [8] for the 32I ensembles, Ref. [9] for the 32ID ensembles, and Ref. [10] for the 48I, 64I, and 32I-fine ensembles. The 32ID-M1 and 32ID-M2 ensembles have not appeared in any of our earlier global fits.

3.B.1 Pseudoscalar Masses, Decay Constants, and Ω Baryon Mass

[11] Note: for consistency with the chiral interpolations in our global fits we choose to parametrize the expansions for the scattering lengths in terms of the light quark mass m_l rather than the more commonly used ratio m_π / f_π.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.005	0.04	0.001	0.001	0.13914(63)	0.08140(46)	—
0.005	0.04	0.001	0.005	0.16693(60)	0.08316(41)	—
0.005	0.04	0.001	0.01	0.19602(59)	0.08526(40)	—
0.005	0.04	0.001	0.02	0.24402(61)	0.08897(41)	—
0.005	0.04	0.001	0.03	0.28430(64)	0.09222(45)	—
0.005	0.04	0.001	0.04	0.31990(69)	0.09511(49)	—
0.005	0.04	0.005	0.005	0.19035(56)	0.08468(38)	—
0.005	0.04	0.005	0.01	0.21609(54)	0.08666(37)	—
0.005	0.04	0.005	0.02	0.26026(53)	0.09027(37)	—
0.005	0.04	0.005	0.03	0.29833(54)	0.09347(39)	—
0.005	0.04	0.005	0.04	0.33245(55)	0.09632(41)	—
0.005	0.04	0.01	0.01	0.23894(51)	0.08858(35)	—
0.005	0.04	0.01	0.02	0.27945(49)	0.09215(36)	—
0.005	0.04	0.01	0.03	0.31524(49)	0.09533(37)	—
0.005	0.04	0.01	0.04	0.34777(50)	0.09816(39)	—
0.005	0.04	0.02	0.02	0.31487(47)	0.09572(36)	—
0.005	0.04	0.02	0.03	0.34722(46)	0.09890(38)	—
0.005	0.04	0.02	0.04	0.37722(46)	0.10175(40)	—
0.005	0.04	0.03	0.03	0.37705(45)	0.10213(40)	0.9629(37)
0.005	0.04	0.03	0.04	0.40512(44)	0.10502(42)	—
0.005	0.04	0.04	0.04	0.43165(42)	0.10796(43)	1.0134(31)

Table 3.11: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 24I $am_l = 0.005$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.01	0.04	0.001	0.001	0.14342(68)	0.08531(45)	—
0.01	0.04	0.001	0.005	0.17087(63)	0.08712(41)	—
0.01	0.04	0.001	0.01	0.19972(60)	0.08921(42)	—
0.01	0.04	0.001	0.02	0.24751(60)	0.09288(46)	—
0.01	0.04	0.001	0.03	0.28773(63)	0.09609(53)	—
0.01	0.04	0.001	0.04	0.32333(70)	0.09898(60)	—
0.01	0.04	0.005	0.005	0.19399(57)	0.08841(39)	—
0.01	0.04	0.005	0.01	0.21954(53)	0.09024(39)	—
0.01	0.04	0.005	0.02	0.26358(50)	0.09370(41)	—
0.01	0.04	0.005	0.03	0.30164(50)	0.09684(44)	—
0.01	0.04	0.005	0.04	0.33577(53)	0.09969(48)	—
0.01	0.04	0.01	0.01	0.24223(49)	0.09193(38)	—
0.01	0.04	0.01	0.02	0.28264(45)	0.09529(39)	—
0.01	0.04	0.01	0.03	0.31839(45)	0.09838(41)	—
0.01	0.04	0.01	0.04	0.35091(46)	0.10118(43)	—
0.01	0.04	0.02	0.02	0.31795(41)	0.09859(39)	—
0.01	0.04	0.02	0.03	0.35023(40)	0.10165(39)	—
0.01	0.04	0.02	0.04	0.38018(40)	0.10443(40)	—
0.01	0.04	0.03	0.03	0.37997(39)	0.10471(39)	0.9785(44)
0.01	0.04	0.03	0.04	0.40797(38)	0.10751(40)	—
0.01	0.04	0.04	0.04	0.43443(38)	0.11035(40)	1.0276(36)

Table 3.12: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 24I $am_l = 0.01$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.004	0.03	0.002	0.002	0.09757(38)	0.05983(30)	—
0.004	0.03	0.002	0.004	0.11330(37)	0.06090(29)	—
0.004	0.03	0.002	0.006	0.12707(37)	0.06192(29)	—
0.004	0.03	0.002	0.008	0.13945(37)	0.06286(30)	—
0.004	0.03	0.002	0.025	0.21797(44)	0.06905(34)	—
0.004	0.03	0.002	0.03	0.23631(47)	0.07048(35)	—
0.004	0.03	0.004	0.004	0.12694(35)	0.06181(29)	—
0.004	0.03	0.004	0.006	0.13926(34)	0.06274(29)	—
0.004	0.03	0.004	0.008	0.15058(34)	0.06363(29)	—
0.004	0.03	0.004	0.025	0.22518(37)	0.06969(32)	—
0.004	0.03	0.004	0.03	0.24301(39)	0.07112(33)	—
0.004	0.03	0.006	0.006	0.15051(33)	0.06363(29)	—
0.004	0.03	0.006	0.008	0.16100(33)	0.06449(30)	—
0.004	0.03	0.006	0.025	0.23227(33)	0.07050(32)	—
0.004	0.03	0.006	0.03	0.24963(35)	0.07193(33)	—
0.004	0.03	0.008	0.008	0.17081(32)	0.06534(30)	—
0.004	0.03	0.008	0.025	0.23920(32)	0.07132(31)	—
0.004	0.03	0.008	0.03	0.25614(32)	0.07276(32)	—
0.004	0.03	0.025	0.025	0.29296(27)	0.07750(32)	0.7332(23)
0.004	0.03	0.025	0.03	0.30733(27)	0.07902(32)	—
0.004	0.03	0.03	0.03	0.32118(27)	0.08058(32)	0.7597(21)

Table 3.13: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 32I $am_l = 0.004$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.006	0.03	0.002	0.002	0.09888(38)	0.06070(33)	—
0.006	0.03	0.002	0.004	0.11439(32)	0.06179(32)	—
0.006	0.03	0.002	0.006	0.12802(30)	0.06282(32)	—
0.006	0.03	0.002	0.008	0.14031(29)	0.06377(32)	—
0.006	0.03	0.002	0.025	0.21843(31)	0.06987(35)	—
0.006	0.03	0.002	0.03	0.23673(34)	0.07129(36)	—
0.006	0.03	0.004	0.004	0.12782(28)	0.06263(31)	—
0.006	0.03	0.004	0.006	0.14003(27)	0.06354(31)	—
0.006	0.03	0.004	0.008	0.15127(26)	0.06442(31)	—
0.006	0.03	0.004	0.025	0.22559(27)	0.07038(32)	—
0.006	0.03	0.004	0.03	0.24338(28)	0.07178(33)	—
0.006	0.03	0.006	0.006	0.15118(26)	0.06439(30)	—
0.006	0.03	0.006	0.008	0.16160(25)	0.06523(30)	—
0.006	0.03	0.006	0.025	0.23266(25)	0.07113(31)	—
0.006	0.03	0.006	0.03	0.24999(26)	0.07254(32)	—
0.006	0.03	0.008	0.008	0.17136(25)	0.06605(30)	—
0.006	0.03	0.008	0.025	0.23961(25)	0.07192(31)	—
0.006	0.03	0.008	0.03	0.25652(25)	0.07334(31)	—
0.006	0.03	0.025	0.025	0.29338(23)	0.07793(30)	0.7392(22)
0.006	0.03	0.025	0.03	0.30775(23)	0.07941(31)	—
0.006	0.03	0.03	0.03	0.32161(22)	0.08092(31)	0.7655(20)

Table 3.14: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 32I $am_l = 0.006$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.008	0.03	0.002	0.002	0.10008(46)	0.06211(40)	—
0.008	0.03	0.002	0.004	0.11564(44)	0.06310(38)	—
0.008	0.03	0.002	0.006	0.12933(43)	0.06408(36)	—
0.008	0.03	0.002	0.008	0.14167(44)	0.06501(36)	—
0.008	0.03	0.002	0.025	0.22029(54)	0.07127(37)	—
0.008	0.03	0.002	0.03	0.23875(58)	0.07276(39)	—
0.008	0.03	0.004	0.004	0.12910(41)	0.06382(35)	—
0.008	0.03	0.004	0.006	0.14134(40)	0.06467(34)	—
0.008	0.03	0.004	0.008	0.15261(40)	0.06551(33)	—
0.008	0.03	0.004	0.025	0.22728(45)	0.07151(33)	—
0.008	0.03	0.004	0.03	0.24519(48)	0.07296(34)	—
0.008	0.03	0.006	0.006	0.15250(39)	0.06545(33)	—
0.008	0.03	0.006	0.008	0.16293(38)	0.06625(32)	—
0.008	0.03	0.006	0.025	0.23419(41)	0.07212(32)	—
0.008	0.03	0.006	0.03	0.25160(42)	0.07354(33)	—
0.008	0.03	0.008	0.008	0.17268(37)	0.06702(31)	—
0.008	0.03	0.008	0.025	0.24099(38)	0.07280(31)	—
0.008	0.03	0.008	0.03	0.25795(39)	0.07422(32)	—
0.008	0.03	0.025	0.025	0.29429(32)	0.07847(31)	0.7399(30)
0.008	0.03	0.025	0.03	0.30862(32)	0.07993(31)	—
0.008	0.03	0.03	0.03	0.32243(31)	0.08140(31)	0.7664(27)

Table 3.15: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 32I $am_l = 0.008$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.001	0.046	0.0001	0.0001	0.10423(23)	0.0938(12)	—
0.001	0.046	0.0001	0.001	0.11512(22)	0.0944(12)	—
0.001	0.046	0.0001	0.0042	0.14718(22)	0.0964(12)	—
0.001	0.046	0.0001	0.008	0.17755(24)	0.0984(12)	—
0.001	0.046	0.0001	0.035	0.31783(45)	0.1090(13)	—
0.001	0.046	0.0001	0.045	0.35642(56)	0.1121(14)	—
0.001	0.046	0.0001	0.055	0.39150(67)	0.1149(14)	—
0.001	0.046	0.001	0.001	0.12497(22)	0.0950(12)	—
0.001	0.046	0.001	0.0042	0.15485(21)	0.0969(12)	—
0.001	0.046	0.001	0.008	0.18385(22)	0.0988(12)	—
0.001	0.046	0.001	0.035	0.32120(39)	0.1092(13)	—
0.001	0.046	0.001	0.045	0.35939(47)	0.1123(14)	—
0.001	0.046	0.001	0.055	0.39418(56)	0.1151(14)	—
0.001	0.046	0.0042	0.0042	0.17949(21)	0.0986(12)	—
0.001	0.046	0.0042	0.008	0.20483(21)	0.1005(12)	—
0.001	0.046	0.0042	0.035	0.33342(30)	0.1107(13)	—
0.001	0.046	0.0042	0.045	0.37030(34)	0.1137(14)	—
0.001	0.046	0.0042	0.055	0.40411(38)	0.1164(14)	—
0.001	0.046	0.008	0.008	0.22725(21)	0.1024(12)	—
0.001	0.046	0.008	0.035	0.34760(26)	0.1126(14)	—
0.001	0.046	0.008	0.045	0.38315(28)	0.1156(14)	—
0.001	0.046	0.008	0.055	0.41594(30)	0.1183(14)	—
0.001	0.046	0.035	0.035	0.43684(21)	0.1231(15)	1.1608(42)
0.001	0.046	0.035	0.045	0.46618(22)	0.1262(15)	—
0.001	0.046	0.035	0.055	0.49409(22)	0.1291(16)	—
0.001	0.046	0.045	0.045	0.49404(21)	0.1294(16)	1.2130(37)
0.001	0.046	0.045	0.055	0.52070(21)	0.1324(16)	—
0.001	0.046	0.055	0.055	0.54632(21)	0.1354(16)	1.2641(34)

Table 3.16: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 32ID $am_l = 0.001$ ensemble.

am_l	am_h	am_x	am_y	am_{xy}	af_{xy}	am_{xxx}
0.0042	0.046	0.0001	0.0001	0.10581(27)	0.0973(12)	—
0.0042	0.046	0.0001	0.001	0.11668(25)	0.0977(12)	—
0.0042	0.046	0.0001	0.0042	0.14870(26)	0.0994(12)	—
0.0042	0.046	0.0001	0.008	0.17913(27)	0.1013(13)	—
0.0042	0.046	0.0001	0.035	0.31972(52)	0.1118(14)	—
0.0042	0.046	0.0001	0.045	0.35808(62)	0.1147(15)	—
0.0042	0.046	0.0001	0.055	0.39279(71)	0.1173(15)	—
0.0042	0.046	0.001	0.001	0.12654(24)	0.0981(12)	—
0.0042	0.046	0.001	0.0042	0.15638(24)	0.0997(12)	—
0.0042	0.046	0.001	0.008	0.18544(26)	0.1015(12)	—
0.0042	0.046	0.001	0.035	0.32302(44)	0.1118(14)	—
0.0042	0.046	0.001	0.045	0.36102(52)	0.1148(14)	—
0.0042	0.046	0.001	0.055	0.39549(59)	0.1173(15)	—
0.0042	0.046	0.0042	0.0042	0.18099(25)	0.1011(12)	—
0.0042	0.046	0.0042	0.008	0.20634(26)	0.1028(13)	—
0.0042	0.046	0.0042	0.035	0.33502(32)	0.1129(14)	—
0.0042	0.046	0.0042	0.045	0.37182(35)	0.1158(14)	—
0.0042	0.046	0.0042	0.055	0.40549(40)	0.1184(15)	—
0.0042	0.046	0.008	0.008	0.22872(26)	0.1044(13)	—
0.0042	0.046	0.008	0.035	0.34906(27)	0.1144(14)	—
0.0042	0.046	0.008	0.045	0.38459(28)	0.1174(14)	—
0.0042	0.046	0.008	0.055	0.41736(31)	0.1200(15)	—
0.0042	0.046	0.035	0.035	0.43813(22)	0.1243(15)	1.1695(48)
0.0042	0.046	0.035	0.045	0.46748(21)	0.1274(15)	—
0.0042	0.046	0.035	0.055	0.49540(21)	0.1302(16)	—
0.0042	0.046	0.045	0.045	0.49534(21)	0.1305(16)	1.2220(41)
0.0042	0.046	0.045	0.055	0.52200(20)	0.1334(16)	—
0.0042	0.046	0.055	0.055	0.54759(19)	0.1363(16)	1.2735(36)

Table 3.17: Partially quenched pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements on the 32ID $am_l = 0.0042$ ensemble.

Ensemble	am_l	am_h	am_{ll}	am_{lh}	af_{ll}	af_{lh}	am_{hhh}
32I-fine	0.0047	0.0186	0.1179(13)	0.1772(12)	0.04846(32)	0.05358(22)	0.5522(29)
48I	0.00078	0.0362	0.08049(13)	0.28853(14)	0.075799(84)	0.090396(86)	0.97018(96)
64I	0.000678	0.02661	0.05903(13)	0.21531(17)	0.055505(95)	0.066534(99)	0.71811(73)
32ID-M1	0.00022	0.0596	0.11812(46)	0.42313(49)	0.12489(23)	0.14673(33)	1.5290(31)
32ID-M2	0.00478	0.03297	0.19487(64)	0.30792(64)	0.07771(22)	0.08716(21)	0.9148(34)

Table 3.18: Unitary pseudoscalar mass, pseudoscalar decay constant, and Ω baryon mass measurements.

3.B.2 R

Ensemble	am_l	am_h	aR
24I	0.005	0.04	0.003154(15)
	0.01	0.04	0.003187(24)
32I	0.004	0.03	0.0006697(34)
	0.006	0.03	0.0006589(30)
	0.008	0.03	0.0006676(34)
32ID	0.001	0.046	0.0018510(43)
	0.0042	0.046	0.0018735(48)
32I-fine	0.0047	0.0186	0.0006300(59)
48I	0.00078	0.0362	0.0006102(40)
64I	0.000678	0.02661	0.0003116(23)
32ID-M1	0.00022	0.0596	0.002170(16)
32ID-M2	0.00478	0.03297	0.0044660(46)

Table 3.19: Summary of measurements of R (Equation (C.28)) at the simulated quark masses on each ensemble. This quantity is equal to m_{res} in the chiral limit.

3.B.3 Wilson Flow Scales

Ensemble	am_l	am_h	${t_0}^{1/2}$	w_0
24I	0.005	0.04	1.31625(57)	1.4911(15)
	0.01	0.04	1.30501(65)	1.4653(14)
32I	0.004	0.03	1.7422(11)	2.0124(26)
	0.006	0.03	1.73622(86)	1.9963(19)
	0.008	0.03	1.7286(11)	1.9793(24)
32ID	0.001	0.046	1.02682(25)	1.21778(72)
	0.0042	0.046	1.02245(27)	1.20420(73)
32I-fine	0.0047	0.0186	2.2860(63)	2.664(16)
48I	0.00078	0.0362	1.29659(39)	1.5013(10)
64I	0.000678	0.02661	1.74448(98)	2.0502(26)
32ID-M1	0.00022	0.0596	0.78719(16)	0.88865(78)
32ID-M2	0.00478	0.03297	1.4841(16)	1.7151(33)

Table 3.20: Summary of Wilson flow measurements.

3.C Analysis of the 32ID-M1 and 32ID-M2 Ensembles

Here we present details of an analysis of the 32ID-M1 and 32ID-M2 ensembles. These lattices were originally generated for scale setting in the context of QCD thermodynamics calculations, and have not appeared in any of our previous chiral fits.

3.C.1 Evolution

The Möbius domain wall action [36] introduces two new scaling parameters, b and c, into the kernel of the domain wall action. If $b - c = 1$, the kernel is identical to the Shamir kernel of conventional domain wall fermions up to a scaling coefficient $\alpha = b + c$. In Ref. [10] we show that a Möbius DWF simulation with $b - c = 1$, a fifth-dimensional extent of L_s, and a scaling coefficient α is directly equivalent to a simulation with Shamir DWF and fifth-dimensional extent αL_s up to small terms that vanish in the $L_s \to \infty$ limit. For the same cost we can therefore use Möbius DWF

to simulate with substantially reduced explicit chiral symmetry breaking simply by increasing α, without deviating from the scaling trajectory of our conventional Shamir ensembles.

In Table 3.21 we summarize the Möbius scale $\alpha = b + c$, the average plaquette and quark condensates, and evolution parameters for the 32ID-M1 and 32ID-M2 ensembles. Both ensembles were generated using an exact hybrid Monte Carlo algorithm with five intermediate Hasenbusch masses — (0.008, 0.04, 0.12, 0.30, 0.60) — for the two, degenerate flavors of light quarks, and a rational approximation for the strange quark determinant. Integration of the gauge and fermion fields was performed using a three-level nested force gradient integrator (FGI QPQPQ): the top level corresponds to updates of the fermion force, the middle level corresponds to DSDR updates, and the bottom level corresponds to gauge field updates, with equal numbers of updates of each level per HMC trajectory. Details regarding the implementation of the DSDR term can be found in Ref. [9].

	32ID-M1	32ID-M2
α	4.0	4.0
Steps per HMC traj.	18	10
$\Delta\tau$	0.056	0.1
Metropolis acceptance	89%	68%
$\langle$Plaquette$\rangle$	0.4681561(65)	0.5671088(24)
$\langle\overline{\psi}_l\psi_l\rangle$	0.0019387(73)	0.0010403(9)
$\langle\overline{\psi}_l\gamma_5\psi_l\rangle$	-0.000008(13)	-0.000007(2)

Table 3.21: The Möbius scale ($\alpha = b + c$), integration parameters, and the measured ensemble averages of the plaquette and quark condensates on the 32ID-M1 and 32ID-M2 ensembles. Here $\Delta\tau$ is the MD time step.

In Figures 3.9 and 3.10 we plot the evolution of the average plaquette, light quark chiral condensate $\langle\overline{\psi}_l\psi_l\rangle$, light quark pseudoscalar condensate $\langle\overline{\psi}_l\gamma_5\psi_l\rangle$, pion propagator evaluated at the fixed time slice $t/a = 20$, square of the topological charge Q^2, and the clover discretized Yang-Mills

action density $E = \mathrm{tr}(F_{\mu\nu}F_{\mu\nu})$ evaluated at the Wilson flow times $t = t_0$ and $t = w_0^2$, as a function of the molecular dynamics simulation time (MD time). Following [60] and our most recent analysis [10] we consider the square of the topological charge rather than the topological charge itself, since this is a parity even observable and our HMC algorithm is parity invariant. We measured the topological charge by cooling the gauge fields with 60 rounds of APE smearing using a smearing coefficient of 0.45, and then measured the topological charge density using the five-loop-improved discretization introduced in Ref. [61].

In Figure 3.11 we plot the integrated autocorrelation times obtained from each of these observables. The integrated autocorrelation time for an observable $Y(t)$ with mean $\overline{Y}$ and variance σ_Y^2 is defined to be

$$\tau_{\mathrm{int}}(\Delta_{\mathrm{cut}}) = \frac{1}{2} + \sum_{\Delta=1}^{\Delta_{\mathrm{cut}}} C(\Delta), \tag{C.26}$$

where

$$C(\Delta) = \left\langle \frac{\left(Y(t) - \overline{Y}\right)\left(Y(t+\Delta) - \overline{Y}\right)}{\sigma_Y^2} \right\rangle_t \tag{C.27}$$

is the autocorrelation at lag Δ, and Δ_{cut} is a cutoff on the maximum lag. The quantity $2\tau_{\mathrm{int}}$ estimates the number of MD time units separating statistically uncorrelated measurements of Y. The error on the integrated autocorrelation time is estimated by bootstrap resampling the set of measurements of $(Y(t) - \overline{Y})(Y(t+\Delta) - \overline{Y})$ with fixed Δ, binned over 20 (40) MD time units on the 32ID-M1 (32ID-M2) ensembles. This choice of binning corresponds to the separation between measurements of the spectrum, and was chosen based on increasing the bin size until the error bars in Figure 3.11 were observed to stabilize and stop growing. More detail regarding this procedure can be found in Ref. [9].

We conclude from the autocorrelation analysis that our separation of 20 (40) MD time units between measurements of the spectrum on the 32ID-M1 (32ID-M2) ensemble is sufficient to ensure that the measurements are uncorrelated, and so we do not perform any further binning. While one should worry about the long autocorrelation time associated with the topological charge on the 32ID-M2 ensemble, we note that our ChPT fits depend only on the measured values of masses

and decay constants, and the long range observables in Figure 3.11 — the pion propagator and quark condensates, for example — suggest an autocorrelation time well within our measurement separation. One should additionally worry that this significant autocorrelation time associated with Q^2 and the poor sampling of topological sectors evidenced by Figure 3.10 suggests statistical errors on the 32ID-M2 ensemble may be underestimated. We choose to still include this ensemble in some of our fits[12] for a number of reasons: in particular, it allows us to overconstrain the linear a^2-scaling terms associated with the DSDR gauge action since it provides an additional DSDR ensemble with a third, independent lattice spacing. In addition, we observe that our results for the LECs of $SU(2)$ PQChPT are completely consistent when we consider the same fit performed with and without the 32ID-M2 ensemble, suggesting that the influence of any undesirable effects of undersampling on our conclusions regarding ChPT are negligible.

3.C.2 Spectrum

We measure and fit the spectrum with the same analysis package previously used to analyze the 48I, 64I, and 32I-fine ensembles in Ref. [10]. This analysis package uses the all-mode averaging (AMA) technique introduced by Blum, Izubuchi, and Shintani [62]. Five *exact* light quark propagators were computed per trajectory using a deflated mixed-precision conjugate gradient solver [63] with 1000 low-mode deflation vectors and a tight stopping precision $r = 10^{-8}$, while *sloppy* light quark propagators with a reduced stopping precision $r = 10^{-4}$ were computed for all time slices. The cheaper strange quark propagators were computed to the tight residual $r = 10^{-8}$ on all time slices using the ordinary conjugate gradient algorithm with no deflation. AMA correlation functions were then computed by time-translational averaging of the sloppy propagators, using the available exact propagators to correct for bias. In all cases we use Coulomb gauge-fixed wall sources (W), and either local (L) or wall sinks.

We have computed the low-energy QCD spectrum for 21 configurations separated by 20 MD

[12]Because of the heavy pion mass $m_\pi \sim 400$ MeV this ensemble is excluded completely from the fits with a 370 MeV mass cut.

time units each on the 32ID-M1 ensemble, and 24 configurations separated by 40 MD time units each on the 32ID-M2 ensemble. These measurements include the residual mass (m_{res}), light-light and heavy-light pseudoscalar masses (m_{ll}, m_{lh}) and decay constants (f_{ll}, f_{lh}), the axial and vector current renormalization coefficients (Z_A, Z_V), the Ω baryon mass (m_{hhh}), and the Wilson flow scales ($t_0^{1/2}$, w_0). Since the analysis package has been discussed in detail in our previous work we paraphrase the fits which were preformed below, and refer the reader to [10] for additional detail. In the following we use the notation "$\simeq$" to denote equality up to excited state contamination for a suitably chosen plateau range. These fits are performed by minimizing an uncorrelated χ^2 (Eqn. C.41) where the correlation functions and fit forms are listed explicitly below.

1. The ratio
$$R(t) = \frac{\langle 0|\Sigma_{\vec{x}} j^a_{5q}(\vec{x},t)|\pi\rangle}{\langle 0|\Sigma_{\vec{x}} j^a_{5}(\vec{x},t)|\pi\rangle}, \tag{C.28}$$
where j^a_{5q} is the pseudoscalar density evaluated at the midpoint of the fifth dimension, and j^a_5 is the physical pseudoscalar density constructed from the surface fields. The residual mass is obtained by averaging over a range of values of t and extrapolating R to the chiral limit.

2. The light-light and heavy-light pseudoscalar masses from
$$\langle 0|\mathcal{O}_1^{s_1}(t)\mathcal{O}_2^{s_2}(0)|0\rangle \simeq \frac{\langle 0|\mathcal{O}_1^{s_1}|X\rangle\langle X|\mathcal{O}_2^{s_2}|0\rangle}{2m_X V}\left(e^{-m_X t} \pm e^{-m_X(T-t)}\right). \tag{C.29}$$
Here $\mathcal{O}_i^{s_i}$ denotes the interpolating operator and smearing, and X denotes the state to which the interpolating operator couples. We perform simultaneous fits to the $\langle PP^{LW}\rangle$, $\langle PP^{WW}\rangle$, and $\langle AP^{LW}\rangle$ correlators for both the light-light and heavy-light pseudoscalar states. The sign is +(-) for the PP(AP) correlator.

3. The ratio $Z_A/Z_{\mathscr{A}}$ — where Z_A ($Z_{\mathscr{A}}$) is a renormalization coefficient relating the local four-dimensional (non-local five-dimensional) axial current to the Symanzik-improved axial current — from
$$\frac{1}{2}\left[\frac{C_{\mathscr{A}}(t-1)+C_{\mathscr{A}}(t)}{2\,C_A(t-\frac{1}{2})} + \frac{2\,C_{\mathscr{A}}(t)}{C_A(t+\frac{1}{2})+C_A(t-\frac{1}{2})}\right] \simeq \frac{Z_A}{Z_{\mathscr{A}}}, \tag{C.30}$$
where $C_{\mathscr{A}}(t) \equiv \langle 0|\sum_{\vec{x}}\partial_\mu \mathscr{A}^a_\mu(\vec{x},t)|\pi\rangle$ and $C_A(t-\frac{1}{2}) \equiv \langle 0|\sum_{\vec{x}}\partial_\mu A^a_\mu(\vec{x},t)|\pi\rangle$. This is the procedure we introduced in [10] to extract Z_A on our Möbius domain wall fermion ensembles; in

our earlier analyses with plain domain wall fermions we extracted Z_A directly from matrix elements of the four-dimensional and five-dimensional axial currents.

4. The renormalization coefficient Z_V relating the local four-dimensional vector current to the Symanzik-improved vector current from

$$\frac{\langle \pi(\Delta t) | \pi(0) \rangle}{\langle \pi(\Delta t) | V_0(t) | \pi(0) \rangle} \simeq Z_V. \tag{C.31}$$

Here V_0 is the temporal component of the light quark electromagnetic current $V_\mu = \overline{q}_l \gamma_\mu q_l$. While Eqn. (C.31) is technically equal to the ratio $Z_V / Z_\mathscr{V}$, where $Z_\mathscr{V}$ relates the non-local five-dimensional vector current to the Symanzik current, the five-dimensional current is exactly conserved on the lattice, implying $Z_\mathscr{V} = 1$.

5. The renormalized light-light and heavy-light pseudoscalar decay constants

$$f_X = Z_V \sqrt{\frac{2}{m_X V} \frac{\left(\mathscr{N}_{AP}^{LW}\right)^2}{\mathscr{N}_{PP}^{WW}}}, \tag{C.32}$$

where we have defined

$$\mathscr{N}_{\mathscr{O}_1 \mathscr{O}_2}^{s_1 s_2} \equiv \frac{\langle 0 | \mathscr{O}_1^{s_1} | X \rangle \langle X | \mathscr{O}_2^{s_2} | 0 \rangle}{2 m_X V}. \tag{C.33}$$

We choose to renormalize the decay constants by Z_V rather than Z_A, which differ by small terms of $\mathcal{O}(m_{\text{res}}^2)$ since the five-dimensional axial current differs from unity by terms of $\mathcal{O}(m_{\text{res}})$, introducing $\mathcal{O}(m_{\text{res}})$ errors into the determination of Z_A via Eqn. (C.30). This point is discussed in further detail in Ref. [8].

6. The Ω baryon mass from the two-point correlation function

$$\mathscr{C}_{\Omega\Omega}^{s_1 s_2}(t) = \sum_{i=1}^{3} \sum_{\vec{x}} \langle 0 | \mathscr{O}_\Omega^{s_1}(\vec{x}, t)_i \overline{\mathscr{O}}_\Omega^{s_2}(0)_i | 0 \rangle \tag{C.34}$$

with the interpolating operator $\mathscr{O}_\Omega(x)_i = \epsilon_{abc} \left(s_a^\top(x) C \gamma_i s_b(x) \right) s_c(x)$. This correlator was computed for both a Coulomb gauge-fixed wall source and a Z_3 box source ($Z_3 B$), and, in both cases, a local sink. The correlators were then projected onto the positive parity component

$$\mathscr{P}_+ \mathscr{C}_{\Omega\Omega}^{s_1 s_2} = \frac{1}{4} \operatorname{tr} \left[\frac{1}{2} \left(1 + \gamma_4 \right) \mathscr{C}_{\Omega\Omega}^{s_1 s_2} \right] \tag{C.35}$$

and simultaneously fit to a double exponential ansatz with common mass terms

$$\begin{cases} \mathscr{C}_{\Omega\Omega}^{LW}(t) = \mathscr{N}_{\Omega\Omega}^{LW} e^{-m_{hhh}t} + \mathscr{N}_{\Omega\Omega}^{LW\prime} e^{-m'_{hhh}t} \\ \mathscr{C}_{\Omega\Omega}^{LZ_3B}(t) = \mathscr{N}_{\Omega\Omega}^{LZ_3B} e^{-m_{hhh}t} + \mathscr{N}_{\Omega\Omega}^{LZ_3B\prime} e^{-m'_{hhh}t} \end{cases}, \tag{C.36}$$

where m_{hhh} is the Ω baryon mass and m'_{hhh} is the mass of the first excited state in the positive parity channel.

7. The Wilson flow scales, $t_0^{1/2}$ and w_0, defined by

$$t^2\langle E(t)\rangle\big|_{t=t_0} = 0.3 \tag{C.37}$$

and

$$t\frac{d}{dt}\left(t^2\langle E(t)\rangle\right)\Big|_{t=w_0^2} = 0.3 \tag{C.38}$$

respectively, where $E = \frac{1}{2}\,\mathrm{tr}(F_{\mu\nu}F_{\mu\nu})$ is the clover discretized Yang-Mills action density.

The fit results are summarized in Table 3.22. The corresponding effective mass ./su2_chpt/plots are shown in Figures 3.12-3.20.

	32ID-M1	32ID-M2
am_{ll}	0.11812(46)	0.19487(64)
am_{lh}	0.42313(49)	0.30792(64)
af_{ll}	0.12489(23)	0.07771(22)
af_{lh}	0.14673(33)	0.087164(21)
Z_A	0.73195(39)	0.70087(14)
Z_V	0.72482(52)	0.70593(92)
am_{hhh}	1.5290(31)	0.9148(34)
am'_{hhh}	1.917(39)	1.215(36)
aR	0.002170(16)	0.0044660(46)
$t_0^{1/2}/a$	0.78719(16)	1.4841(16)
w_0/a	0.88865(78)	1.7151(33)
m_{ll}/m_{hhh}	0.07725(34)	0.21303(90)
m_{lh}/m_{hhh}	0.27673(65)	0.3366(12)
f_{ll}/m_{hhh}	0.08248(14)	0.08496(41)
f_{lh}/m_{hhh}	0.09690(26)	0.09529(40)

Table 3.22: Summary of fit results in lattice units. Here R is defined by Equation (C.28), which becomes m_{res} when extrapolated to the chiral limit.

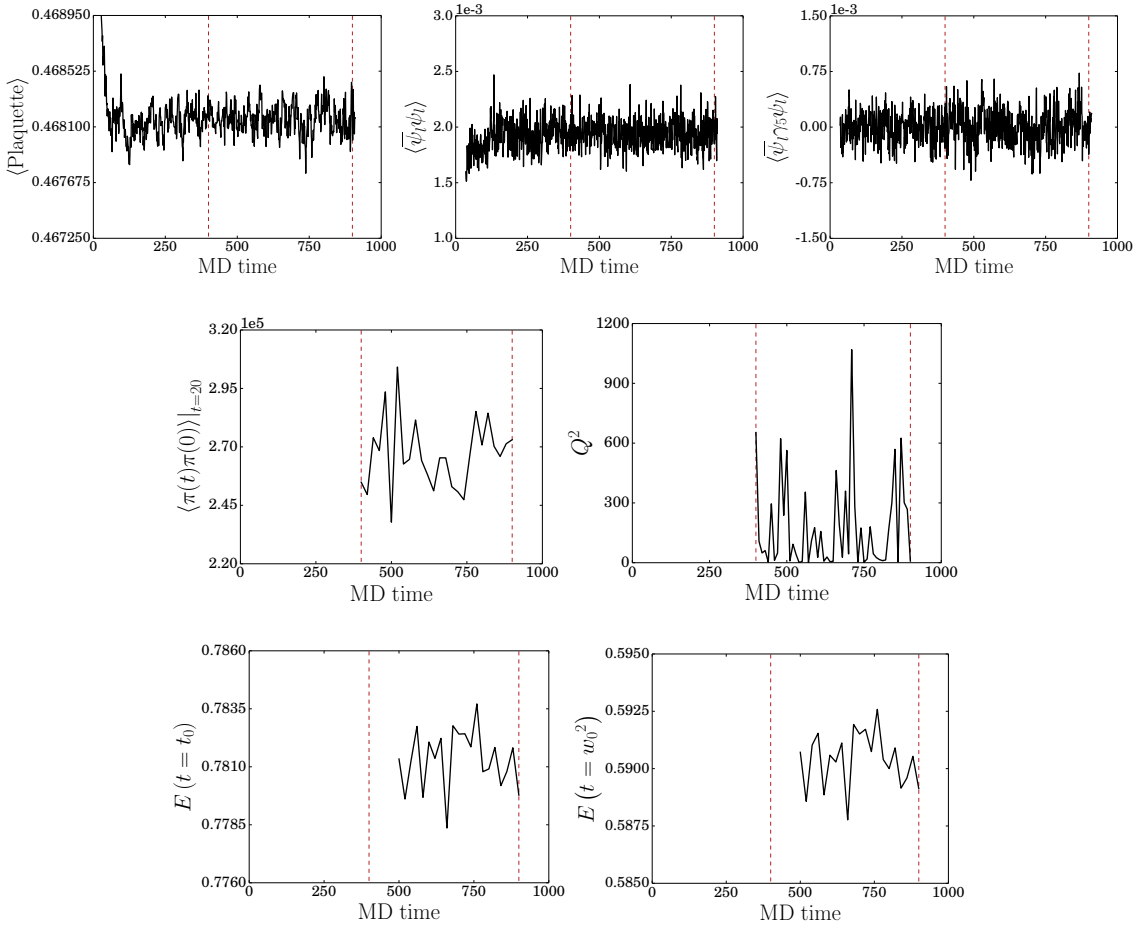

Figure 3.9: Molecular dynamics evolution of the plaquette, chiral and pseudoscalar condensates, pion propagator at $t/a = 20$, square of the topological charge, and clover discretized action density computed at the Wilson flow times t_0 and w_0^2 as a function of MD time on the 32ID-M1 ensemble. The first three quantities were computed every MD time step as part of the evolution. The topological charge and Wilson flow scales were computed every 10 and 20 MD time steps, respectively, after the ensemble was thermalized. The dashed vertical lines mark the range of MD times used to perform calculations of the spectrum.

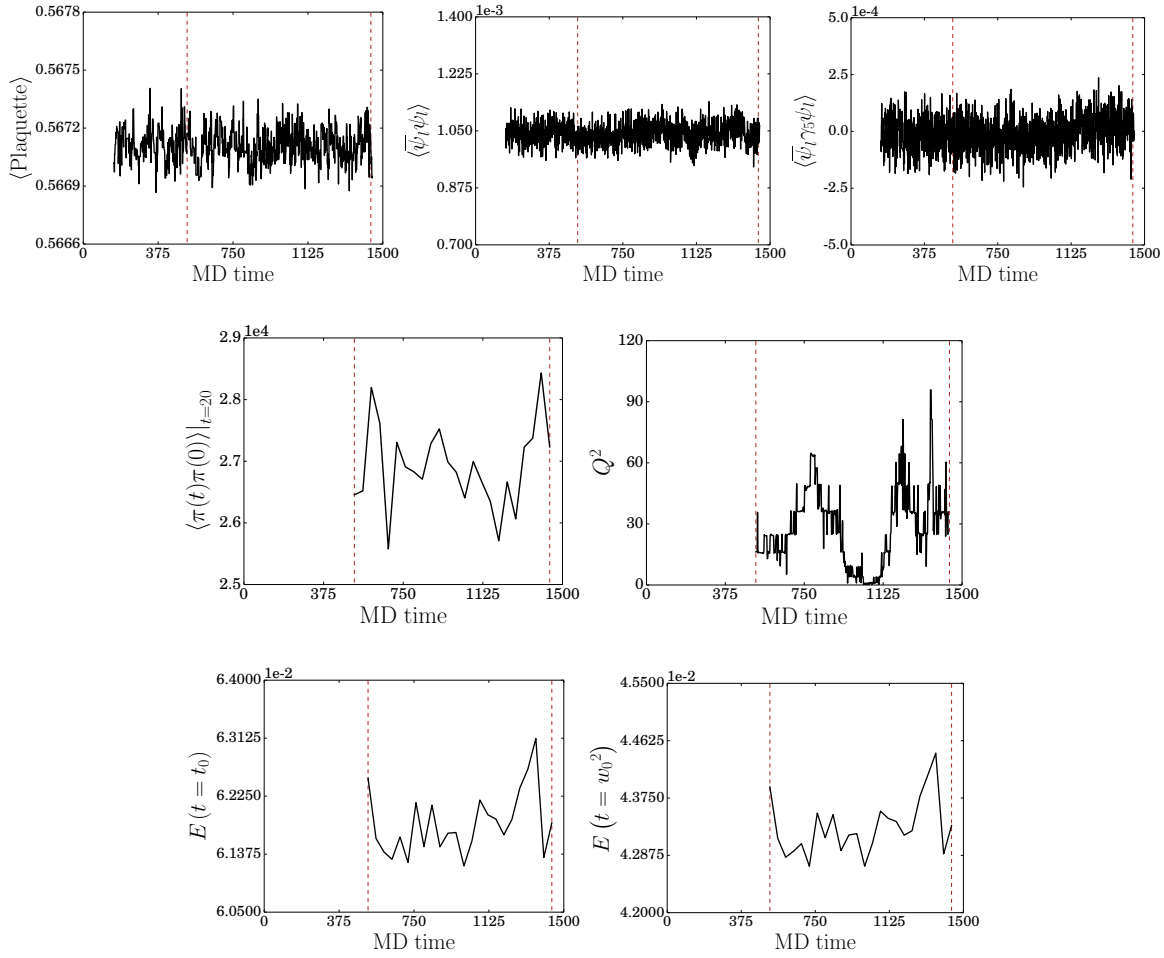

Figure 3.10: Molecular dynamics evolution of the plaquette, chiral and pseudoscalar condensates, pion propagator at $t/a = 20$, square of the topological charge, and clover discretized action density computed at the Wilson flow times t_0 and w_0^2 as a function of MD time on the 32ID-M2 ensemble. The first three quantities were computed every MD time step as part of the evolution. The topological charge and Wilson flow scales were computed every 2 and 40 MD time steps, respectively, after the ensemble was thermalized. The dashed vertical lines mark the range of MD times used to perform calculations of the spectrum.

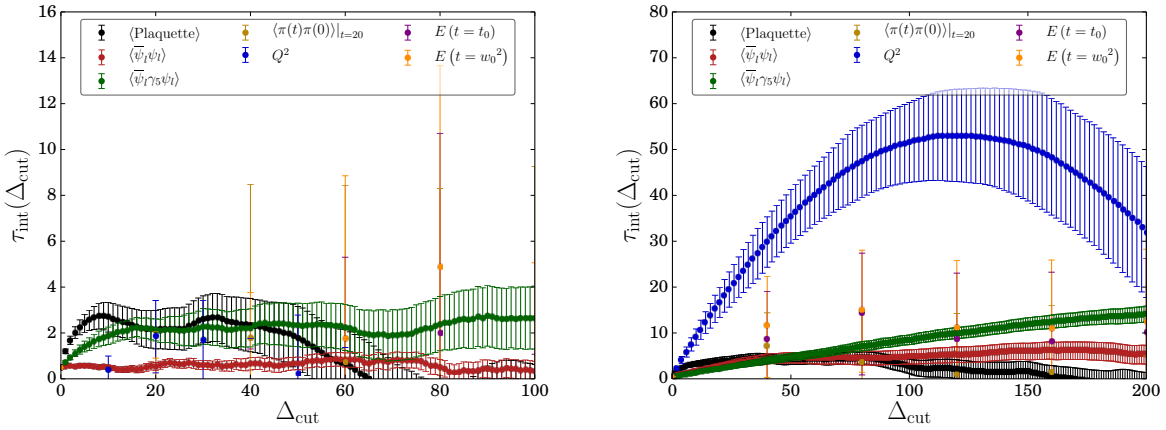

(a) 32ID-M1

(b) 32ID-M2

Figure 3.11: Integrated autocorrelation times for the observables plotted in Figures 3.9 and 3.10.

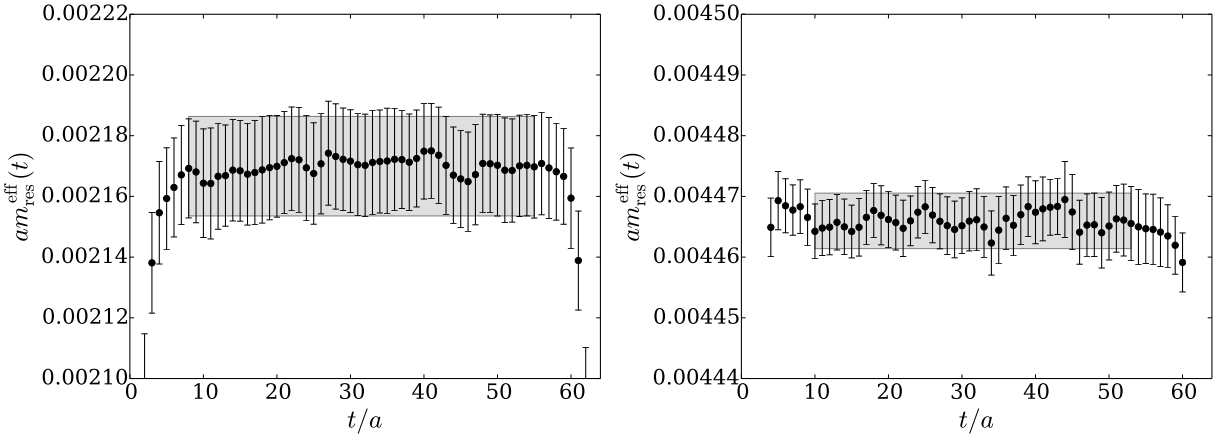

Figure 3.12: The residual mass, from Eqn. (C.28), on the 32ID-M1 (left) and 32ID-M2 (right) ensembles.

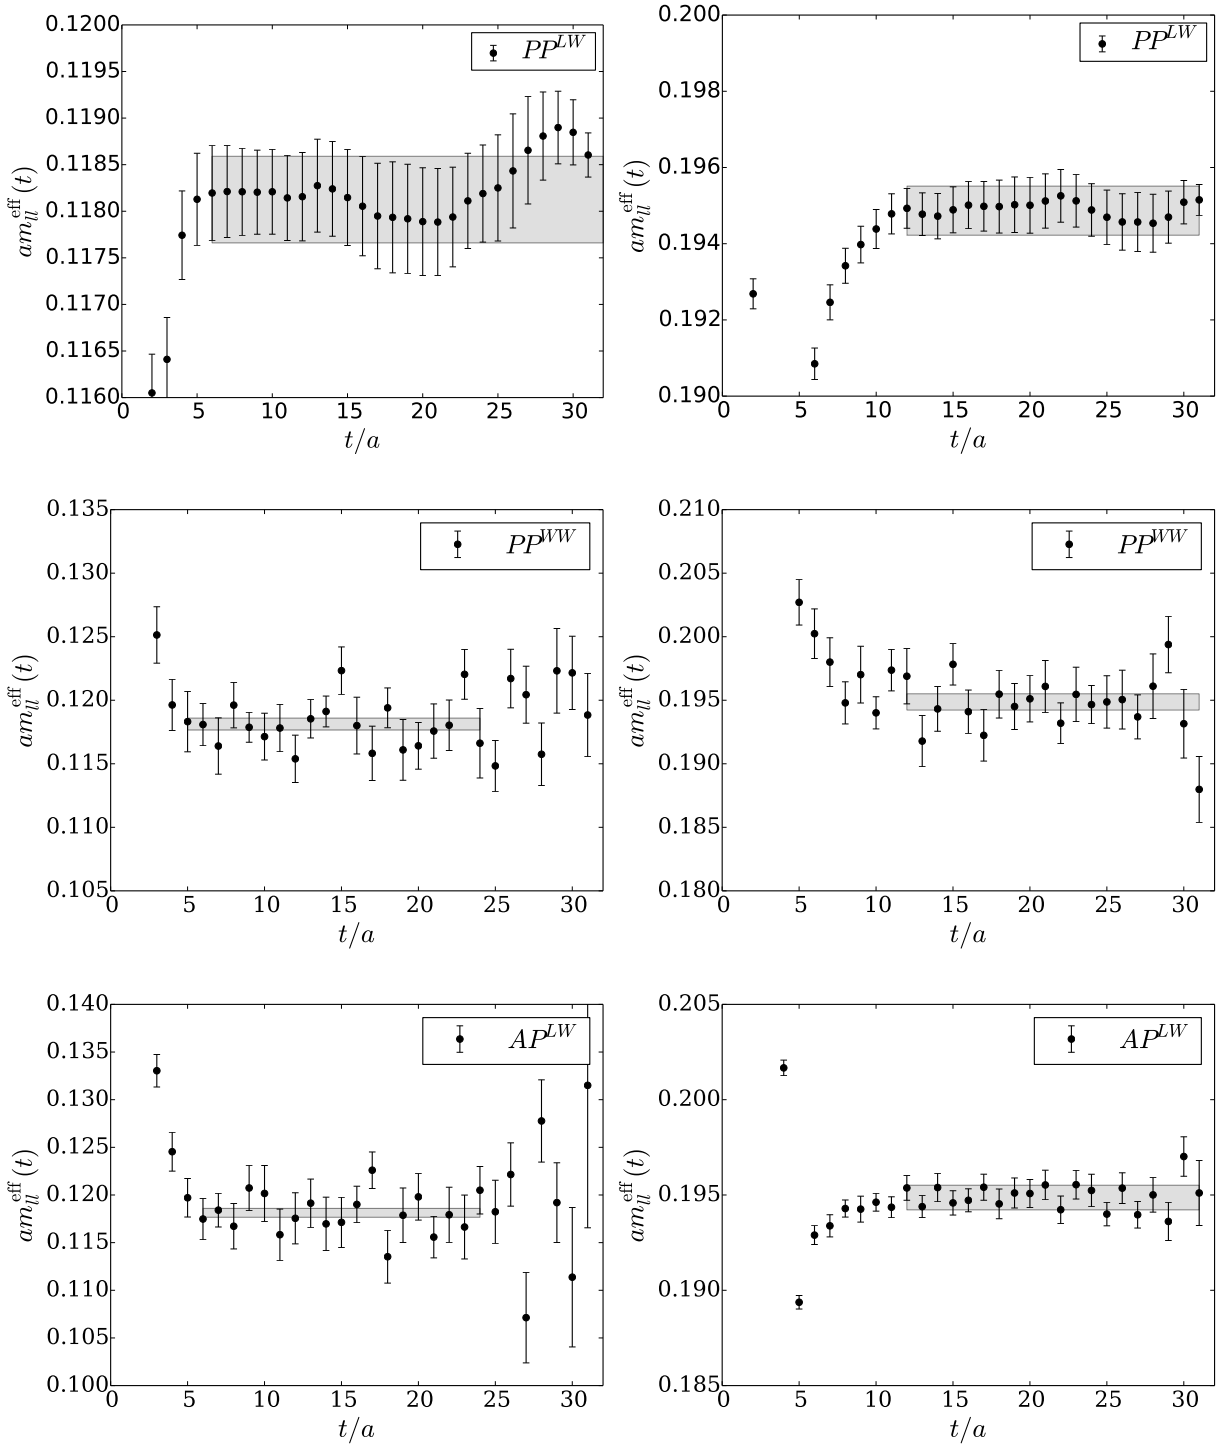

Figure 3.13: Light-light pseudoscalar mass on the 32ID-M1 (left) and 32ID-M2 (right) ensembles. We simultaneously fit a common mass m_{ll} to the three correlators $\langle PP^{LW}\rangle$, $\langle PP^{WW}\rangle$, and $\langle AP^{LW}\rangle$ on each ensemble.

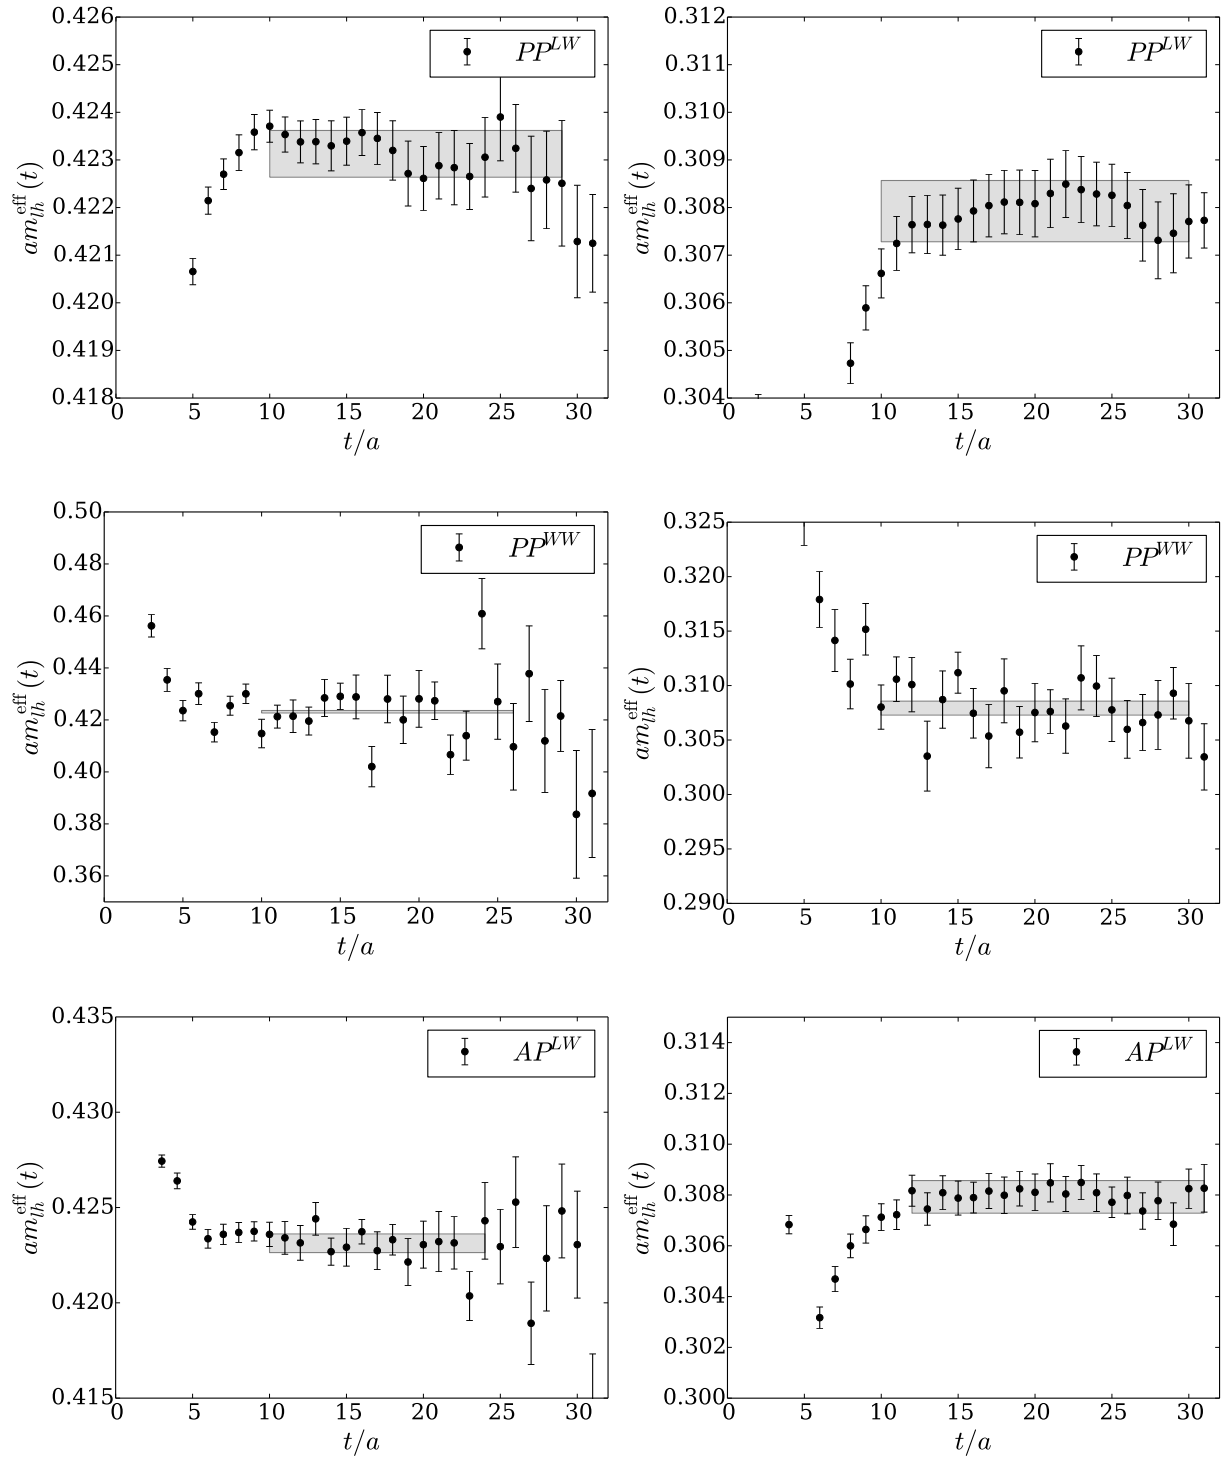

Figure 3.14: Heavy-light pseudoscalar mass on the 32ID-M1 (left) and 32ID-M2 (right) ensembles. We simultaneously fit a common mass m_{lh} to the three correlators $\langle PP^{LW}\rangle$, $\langle PP^{WW}\rangle$, and $\langle AP^{LW}\rangle$ on each ensemble.

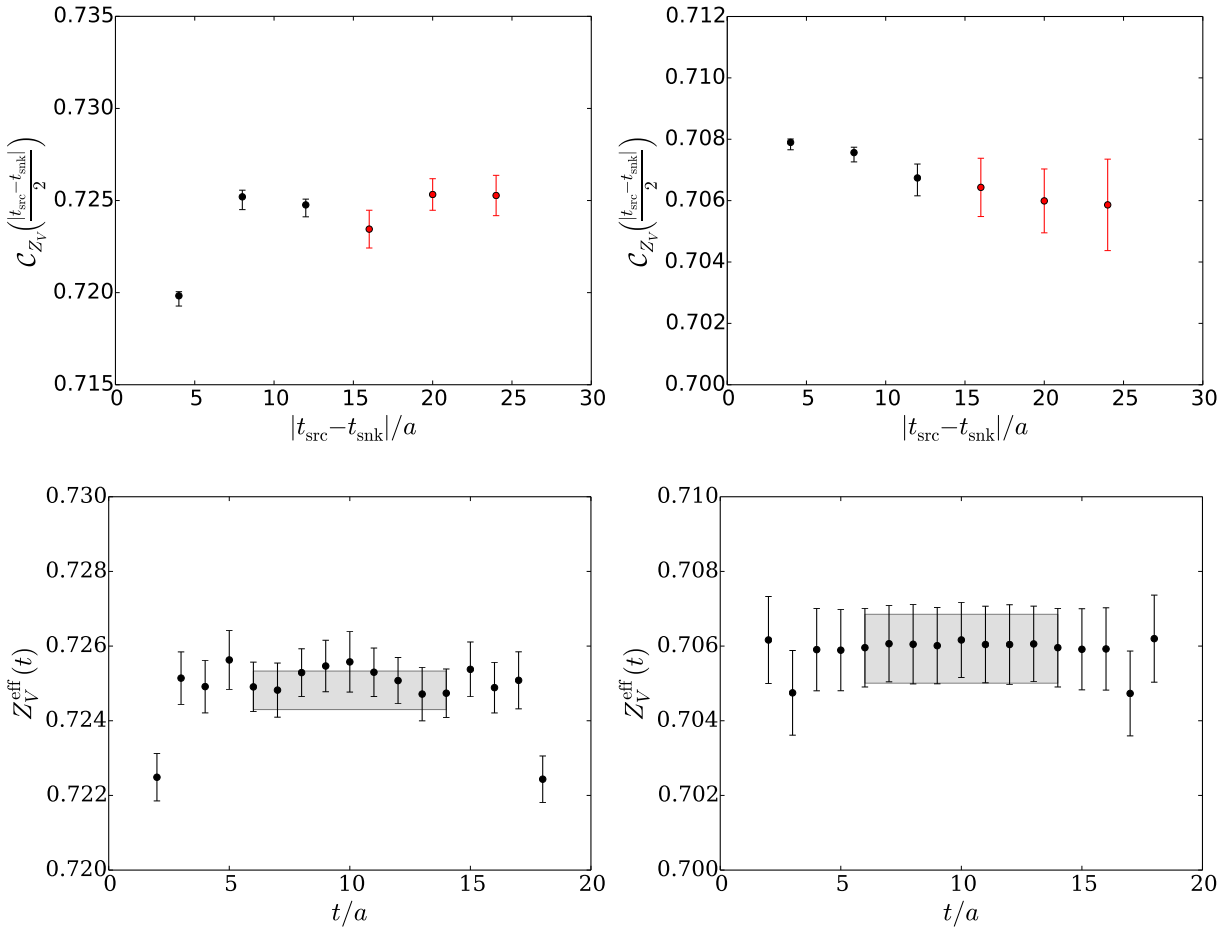

Figure 3.15: The vector current renormalization coefficient on the 32ID-M1 (left) and 32ID-M2 (right) ensembles. In the upper plot we show the dependence of the ratio (C.31) on the source-sink separation: the point plotted for each separation is evaluated at the midpoint $t = |t_{\rm src} - t_{\rm snk}|/2a$. Points which were included in the fit are marked in red. In the lower plot we show an example of the fit to Z_V overlaying the ratio (C.31) for one of the source-sink separations included in the fit.

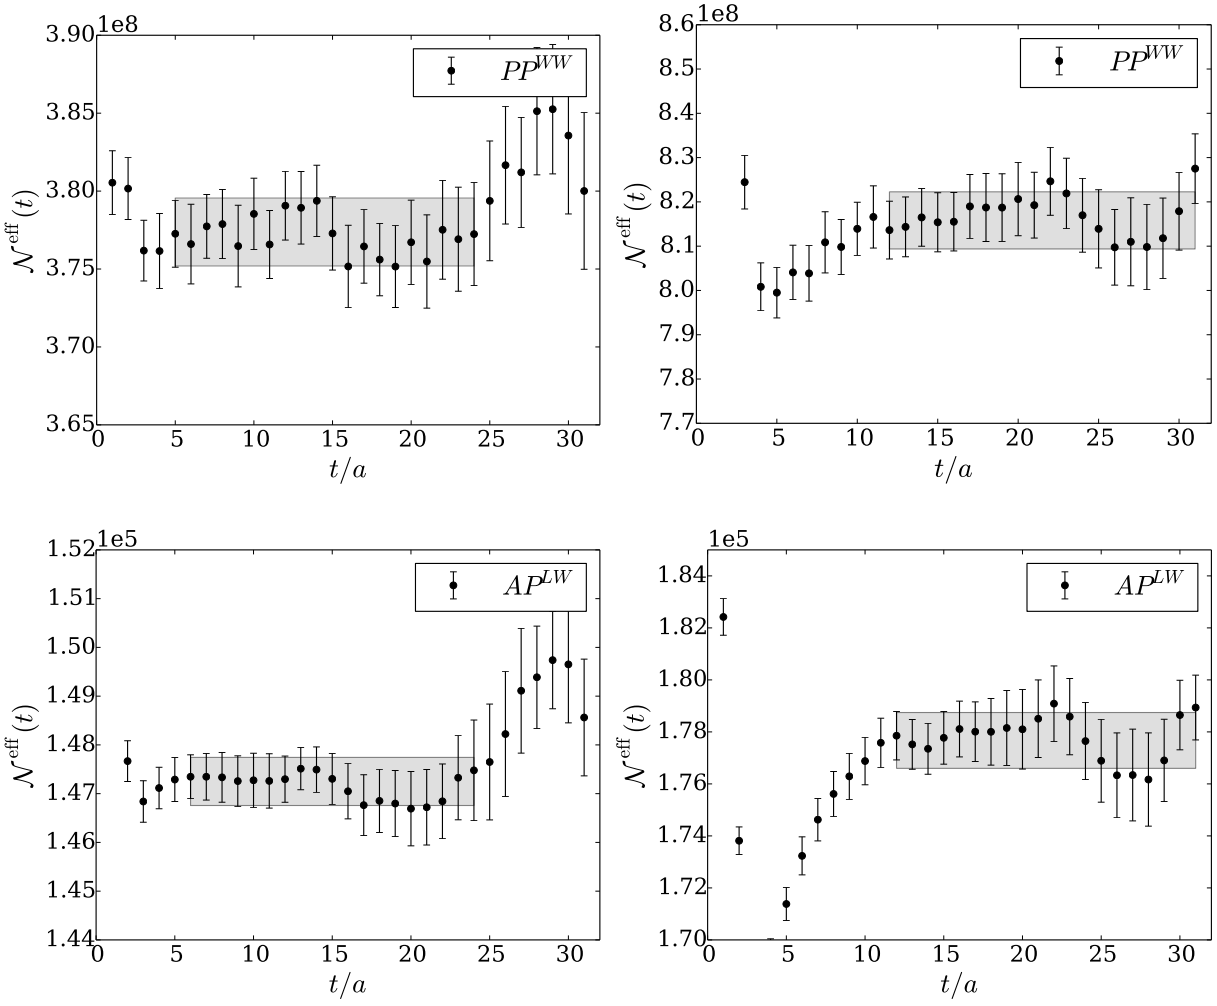

Figure 3.16: Light-light effective amplitudes $\mathscr{N}^{\rm eff}_{\mathscr{O}_1\mathscr{O}_2}(t) \equiv \langle \mathscr{O}_1(t)\mathscr{O}_2(0)\rangle/(e^{-m_{\rm eff}t} \pm e^{-m_{\rm eff}(T-t)})$ on the 32ID-M1 (left) and 32ID-M2 (right) ensembles. The sign is +(-) for the PP(AP) correlator. These are related to the light-light pseudoscalar decay constant according to Eqn. (C.32).

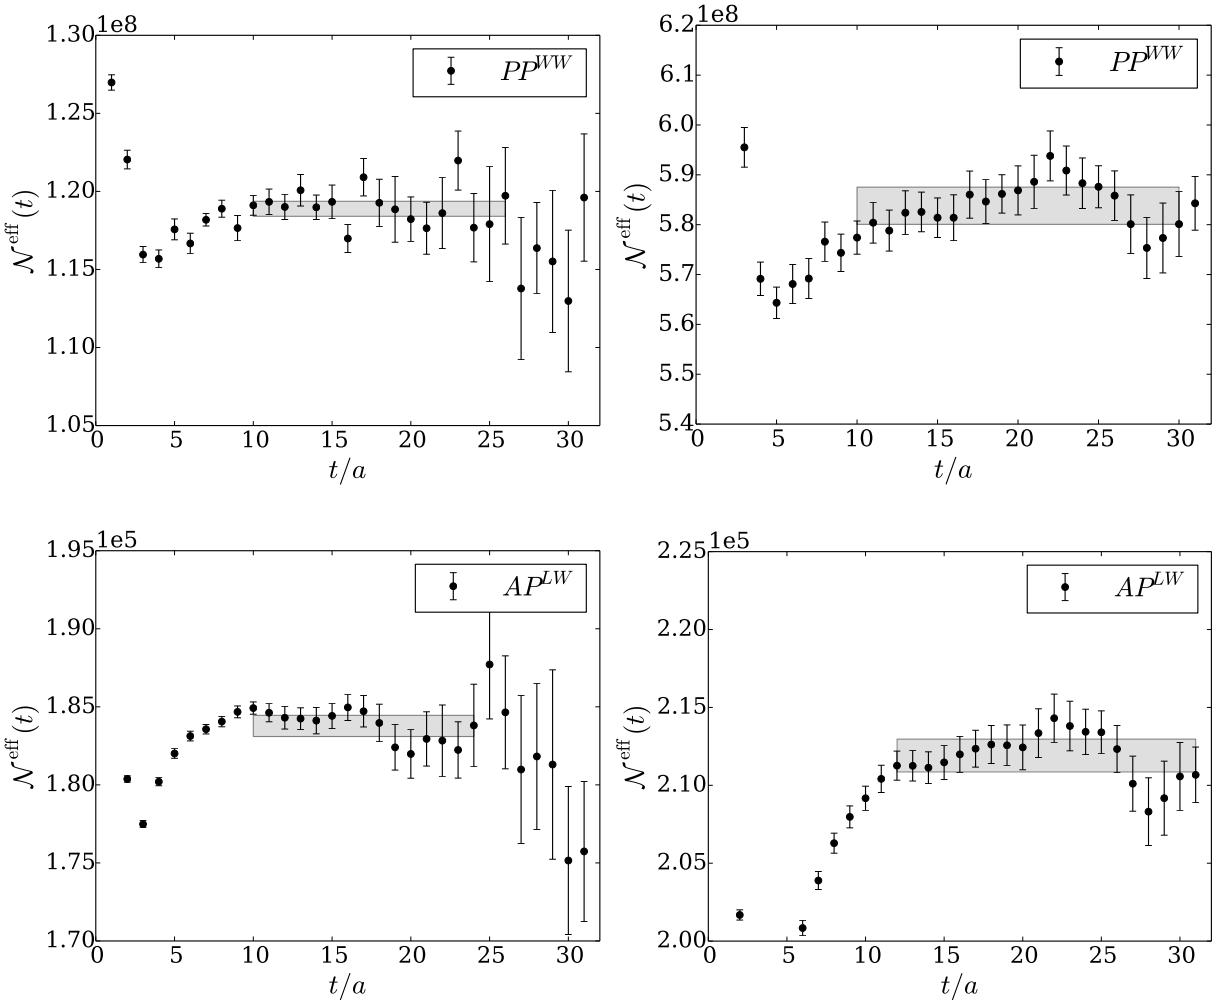

Figure 3.17: Heavy-light effective amplitudes on the 32ID-M1 (left) and 32ID-M2 (right) ensembles.

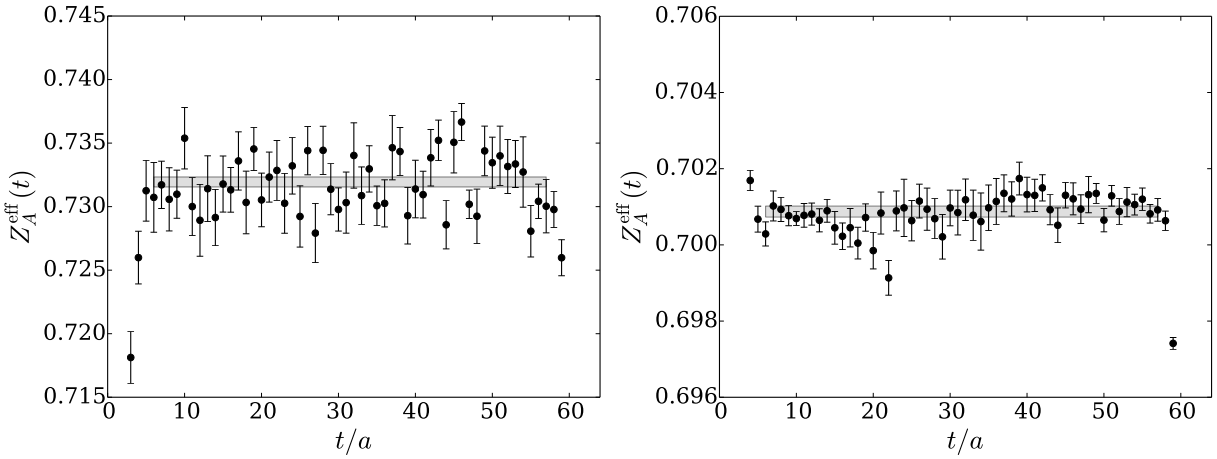

Figure 3.18: The axial current renormalization coefficient, from Eqn. (C.30), on the 32ID-M1 (left) and 32ID-M2 (right) ensembles.

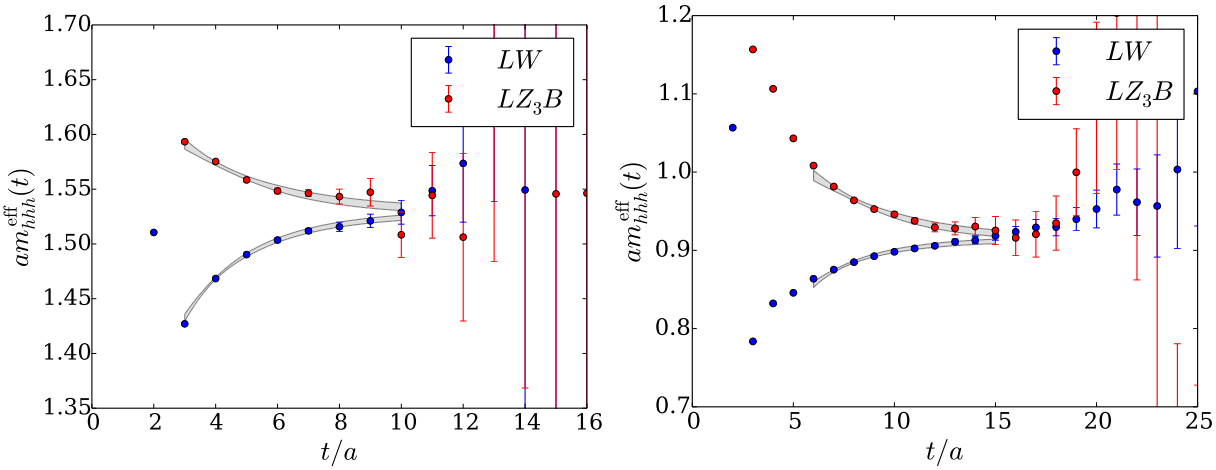

Figure 3.19: The Ω baryon mass on the 32ID-M1 (left) and 32ID-M2 (right) ensembles. The wall source and Z_3 box source correlators are simultaneously fit to double exponential ansätze with common mass terms (Eqn. (C.36)). Here we overlay the data with the effective mass curves obtained from the fit.

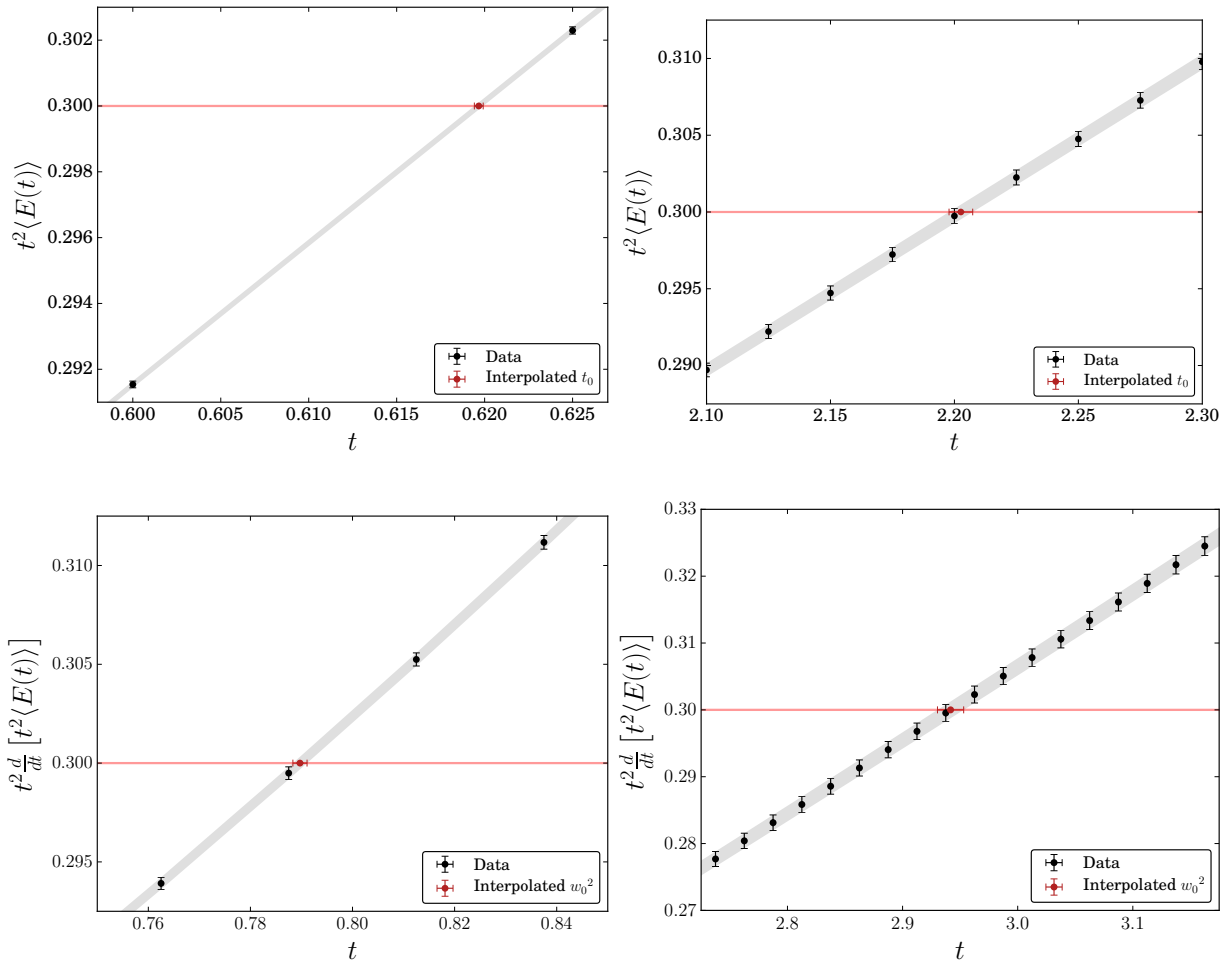

Figure 3.20: The Wilson flow scales $t_0^{1/2}$ (top) and w_0 (bottom) on the 32ID-M1 (left) and 32ID-M2 (right) ensembles.

3.D Fits with Weighted χ^2

Correlations among data in a non-linear least squares fit are included by minimizing the *correlated* χ^2

$$\chi^2 = \sum_{ij} \left(\frac{y^i - f(\vec{\beta})^i}{\sigma^i} \right) \left(\rho^{-1}\right)^{ij} \left(\frac{y^j - f(\vec{\beta})^j}{\sigma^j} \right) \tag{C.39}$$

over the the space of model parameters $\vec{\beta}$, where

$$\rho^{ij} = \frac{\langle \left(y^i - \mu^i\right) \left(y^j - \mu^j\right) \rangle}{\sigma^i \sigma^j} \tag{C.40}$$

is the correlation matrix. In the limit that the data is completely uncorrelated $\rho^{ij} = \delta^{ij}$, and we recover the familiar *uncorrelated* χ^2

$$\chi^2 = \sum_{i} \left(\frac{y^i - f(\vec{\beta})^i}{\sigma^i} \right)^2 . \tag{C.41}$$

In practice, correlations between data points computed on the same ensemble are often so strong that $\rho^{ij} \approx 1 \, \forall i, j$ is nearly singular, and minimization of the correlated χ^2 defined by Eqn. (C.39) is numerically unstable. This pathology can be tamed by ignoring the correlations and instead minimizing the uncorrelated χ^2, or by removing modes with small eigenvalues from ρ^{ij} until the minimization algorithm becomes stable. In either case one loses a rigorous interpretation of χ^2 as a statistical measure of the goodness-of-fit.

In Figure 3.21 we plot the correlation matrix ρ^{ij} and its eigenvalue spectrum computed from the data included in fits with a 370 MeV cut. We find, as expected, that the correlation matrix is extremely singular due to strong correlations associated with partial quenching and reweighting: the eigenvalues span 15 orders of magnitude, and the condition number is $\text{cond}(\rho^{ij}) = 1.85 \times 10^{17}$. In Figure 3.22 we further plot the sub-blocks of ρ^{ij} corresponding to the 32I ensembles as an example of the cross-correlations present in our data, for example, between the light-light and heavy-light pseudoscalar masses. We conclude that we are unable to accurately invert the correlation matrix, much less attempt fully correlated fits as defined by Equation C.39.

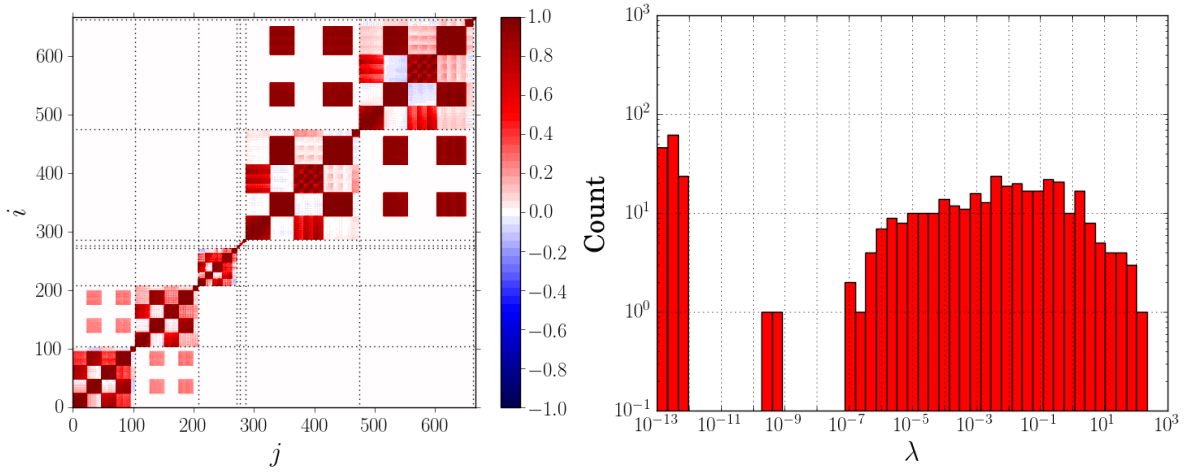

Figure 3.21: Left: the correlation matrix ρ^{ij} corresponding to fits with a 370 MeV cut. The dashed lines mark the division into sub-blocks by ensemble. From left to right these are: 32I ($m_l = 0.004$), 32I ($m_l = 0.006$), 24I ($m_l = 0.005$), 48I, 64I, 32I-fine, 32ID ($m_l = 0.001$), 32ID ($m_l = 0.0042$), and 32ID-M1. Right: the eigenvalue spectrum of ρ^{ij}.

The fits discussed in Section 3.5 were performed by minimizing the uncorrelated χ^2 (Eqn. (C.41)). We expect, however, that our data is highly correlated, in particular between measurements of partially quenched observables on the same ensemble but with different choices of the valence quark masses, and between different reweightings in m_h of the same observable. These particular classes of correlations are especially troublesome since our partially quenched measurements and m_h reweightings were performed on the relatively heavy pion mass ensembles (24I, 32I, and 32ID) — a naive uncorrelated fit will tend to give too much weight to this data, which is far from the chiral limit where ChPT is exact. In this appendix we repeat these fits, normalizing the contributions to χ^2 by the number of nondegenerate pseudoscalar mass measurements (N_e) associated with a given ensemble (e):

$$\chi^2_e = \frac{1}{N_e} \sum_i \left(\frac{y^i_e - f^i_e}{\sigma^i_e} \right)^2 , \tag{C.42}$$

where $\chi^2 = \sum_e \chi^2_e$ and the N_e are listed in Table 3.23. This can be loosely regarded as the limit of extreme correlation, in which all of the partially quenched measurements on a given ensemble are effectively weighted as a single point by inflating their statistical errors $\sigma^i_e \to \sqrt{N_e}\sigma^i_e$. We use the

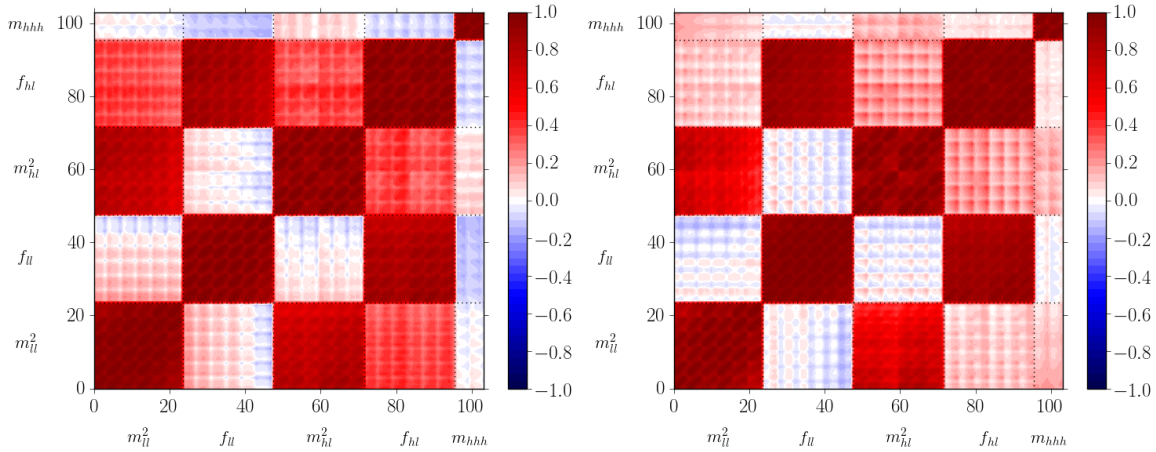

(a) 32I, $m_l = 0.004$

(b) 32I, $m_l = 0.006$

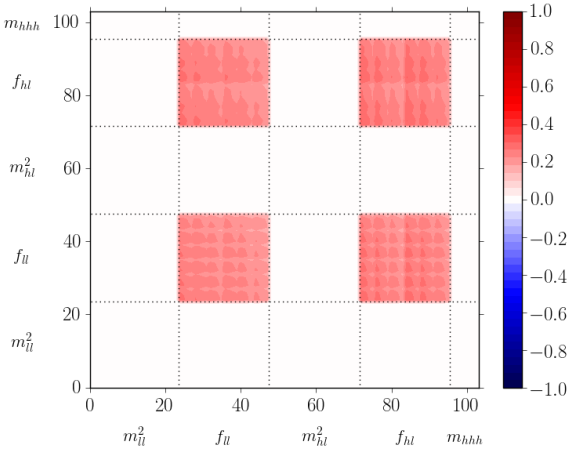

(c) 32I, $m_l = 0.004 \times m_l = 0.006$

Figure 3.22: Sub-blocks of the correlation matrix corresponding to the 32I ensembles. Panel (c) shows the cross-correlations between the $m_l = 0.004$ and $m_l = 0.006$ ensembles induced by the use of Z_V extrapolated to the chiral limit to normalize the decay constants.

difference in central values between these two schemes to assign a systematic error associated with our inability to fully resolve the true correlation matrix to our fits.

This scheme for weighting χ^2 can be further understood by analyzing the correlation matrix in the limit that the off-diagonal terms are completely dominated by the correlations between

Mass Cut	24I	32I	32ID	32I-fine	48I	64I	32ID-M1	32ID-M2
370 MeV	12	48	80	1	1	1	1	—
450 MeV	48	120	80	1	1	1	1	1

Table 3.23: The value of N_e — the number of non-degenerate quark mass combinations (m_x, m_y, m_l, m_h) used for pseudoscalar measurements entering into the fits — for each ensemble and mass cut. There are four values of m_h for each fixed (m_x, m_y, m_l) obtained by reweighting in the heavy sea quark determinant.

partially quenched measurements on the same ensemble. To clarify this discussion, we write the full correlation matrix (Eqn. (C.40)) as $\rho_{(e,i,a);(e',i',a')}$, where e indexes the ensemble, i indexes the valence quark mass combination, and a indexes the observable. If the data is both highly correlated and dominated by the correlations between the partially quenched data the correlation matrix will have a block structure

$$\rho_{(e,i,a),(e',i',a')} = (\rho_{PQ})^{i'j'}_{(e,a)} \delta_{ee'} \delta_{aa'} \tag{C.43}$$

where $(\rho_{PQ})^{i'j'}_{(e,a)}$ is the $N_{(e,a)} \times N_{(e,a)}$ sub-matrix describing the correlations between partially quenched measurements of observable a on ensemble e; in the left panel of Figure 3.21, for example, these are the extremely correlated blocks lying along the main diagonal. We expect these correlations to be sufficiently strong that the blocks $(\rho_{PQ})^{i'j'}_{(e,a)}$ will be nearly singular, which we can write in general as

$$(\rho_{\mathrm{PQ}})^{i'j'}_{(e,a)} = \begin{pmatrix} 1 & 1-\epsilon^{12}_{(e,a)} & \cdots & 1-\epsilon^{1N_{(e,a)}}_{(e,a)} \\ 1-\epsilon^{21}_{(e,a)} & 1 & \cdots & 1-\epsilon^{2N_{(e,a)}}_{(e,a)} \\ \vdots & \vdots & \ddots & \vdots \\ 1-\epsilon^{N_{(e,a)}1}_{(e,a)} & 1-\epsilon^{N_{(e,a)}2}_{(e,a)} & \cdots & 1 \end{pmatrix} \tag{C.44}$$

where the $\epsilon^{i'j'}_{(e,a)} \ll 1$ measure the small deviations from unity of the off-diagonal entries. To simplify the analysis we set $\epsilon^{i'j'}_{(e,a)} = \epsilon$ everywhere and work to leading order in ϵ. In this limit each of the $(\rho_{\mathrm{PQ}})^{i'j'}_{(e,a)}$ has a single eigenvector $(1, 1, \ldots, 1)$ with eigenvalue $N_{(e,a)} - (N_{(e,a)} - 1)\epsilon$, representing the mode where all $N_{(e,a)}$ data points are completely correlated. The remaining $N_{(e,a)} - 1$ eigenvectors

are degenerate with eigenvalue ϵ and span the subspace of correlations in the data orthogonal to the completely correlated mode; their poor statistical resolution can be understood as a source of the numerical instabilities observed in fully correlated fits. Since $(\rho_{\rm PQ}^{ij})_{(e,a)}$ is a real, symmetric matrix it can be diagonalized by an orthogonal transformation:

$$(\rho_{\rm PQ}^{ij})_{ee} = \left(\begin{array}{c|c} 1 & \\ 1 & \\ \vdots & Q \\ 1 & \end{array}\right) \left(\begin{array}{cccc} N_{(e,a)} - (N_{(e,a)}-1)\epsilon & 0 & \cdots & 0 \\ 0 & \epsilon & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \epsilon \end{array}\right) \left(\begin{array}{c|ccc} 1 & 1 & \cdots & 1 \\ \hline & & & \\ & Q^\top & & \\ & & & \end{array}\right). \tag{C.45}$$

Here Q is an orthogonal matrix whose columns correspond to an appropriate choice of the $N_{(e,a)}-1$ degenerate eigenvectors with eigenvalue ϵ. Eqn. (C.42) follows from the fully correlated Eqn. (C.39) if we define new blocks $(\tilde{\rho}_{\rm PQ})_{(e,a)}^{i'j'}$ by making the replacements $N_{(e,a)} - (N_{(e,a)}-1)\epsilon \approx N_{(e,a)}$ for the largest eigenvalue, and $\epsilon \to N_{(e,a)}$ for the remaining $N_{(e,a)}-1$ eigenvalues:

$$(\tilde{\rho}_{\rm PQ})_{(e,a)}^{i'j'} = \left(\begin{array}{c|c} 1 & \\ 1 & \\ \vdots & Q \\ 1 & \end{array}\right) \left(\begin{array}{cccc} N_{(e,a)} & 0 & \cdots & 0 \\ 0 & N_{(e,a)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & N_{(e,a)} \end{array}\right) \left(\begin{array}{c|ccc} 1 & 1 & \cdots & 1 \\ \hline & & & \\ & Q^\top & & \\ & & & \end{array}\right) = N_{(e,a)}\delta^{i'j'}, \tag{C.46}$$

and substitute $(\rho_{\rm PQ})_{(e,a)}^{i'j'} \to (\tilde{\rho}_{\rm PQ})_{(e,a)}^{i'j'}$. Effectively, in Eqn. (C.42) we are treating the modes associated with the largest eigenvalue of each of the $(\rho_{\rm PQ})_{(e,a)}^{i'j'}$ exactly up to terms of $\mathcal{O}(\epsilon)$, and underweighting the subdominant modes by a factor $\sim \epsilon/N_{(e,a)}$. We find in practice that this stabilizes the fits while still capturing some of the important effects of correlations in the data. More generally, one expects that the $\epsilon_{(e,a)}^{i'j'}$ in Eqn. (C.44) are not all equal — breaking the degeneracy between the $N_{(e,a)}-1$ smallest eigenvalues of $(\rho_{\rm PQ})_{(e,a)}^{i'j'}$ — and that some of the off-diagonal entries in the full correlation matrix, representing other kinds of correlations, are non-zero; these effects are $\mathcal{O}(\epsilon)$ and do not change the argument presented here.

In the remainder of the appendix we summarize the results of fits performed by minimizing the normalized χ^2 (Eqn. (C.42)).

3.D.1 Fit Parameters

		NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
	χ^2/dof	0.011(5)	0.049(13)	0.008(4)	0.007(4)
24I	am_l^{phys}	-0.001767(79)	-0.001774(81)	-0.001765(79)	-0.001765(79)
	am_h^{phys}	0.03236(32)	0.03206(29)	0.03237(32)	0.03238(30)
	a^{-1}	1.777(13) GeV	1.797(12) GeV	1.777(13) GeV	1.777(12) GeV
32I	am_l^{phys}	0.000263(14)	0.000254(13)	0.000265(14)	0.000266(13)
	am_h^{phys}	0.02485(24)	0.02469(18)	0.02491(23)	0.02496(21)
	a^{-1}	2.371(16) GeV	2.398(14) GeV	2.369(16) GeV	2.365(15) GeV
32ID	am_l^{phys}	-0.000131(27)	-0.000156(25)	-0.000121(25)	-0.000120(26)
	am_h^{phys}	0.04547(86)	0.04496(75)	0.04557(80)	0.04544(82)
	a^{-1}	1.389(13) GeV	1.400(12) GeV	1.387(12) GeV	1.389(12) GeV
32I-fine	am_l^{phys}	0.000077(30)	0.000060(30)	0.000073(31)	0.000082(33)
	am_h^{phys}	0.01884(60)	0.01830(58)	0.01881(59)	0.01907(65)
	a^{-1}	3.110(43) GeV	3.172(42) GeV	3.114(43) GeV	3.094(44) GeV
48I	am_l^{phys}	0.0006959(86)	0.0007012(75)	0.0006983(84)	0.0007001(81)
	am_h^{phys}	0.03574(18)	0.03588(14)	0.03575(17)	0.03580(16)
	a^{-1}	1.731(4) GeV	1.728(3) GeV	1.730(4) GeV	1.729(4) GeV
64I	am_l^{phys}	0.0006175(78)	0.0006219(64)	0.0006192(74)	0.0006198(67)
	am_h^{phys}	0.02530(17)	0.02552(13)	0.02535(17)	0.02539(14)
	a^{-1}	2.362(7) GeV	2.354(5) GeV	2.360(7) GeV	2.358(6) GeV
32ID-M1	am_l^{phys}	0.000825(68)	0.000731(47)	0.000808(65)	0.000797(51)
	am_h^{phys}	0.0791(16)	0.0753(10)	0.0784(16)	0.0778(12)
	a^{-1}	1.020(10) GeV	1.039(7) GeV	1.024(10) GeV	1.029(7) GeV
32ID-M2	am_l^{phys}	—	-0.003417(20)	—	-0.003413(23)
	am_h^{phys}	—	0.02435(48)	—	0.02422(55)
	a^{-1}	—	2.048(19) GeV	—	2.030(22) GeV

Table 3.24: The (uncorrelated) χ^2/dof, unrenormalized physical quark masses in bare lattice units (without m_{res} included), and the values of the inverse lattice spacing a^{-1} in physical units, obtained from fits to $SU(2)$ PQChPT with the stated pion mass cuts.

		NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
24I	Z_l	0.9710(53)	0.9698(46)	0.9702(51)	0.9691(50)
	Z_h	0.9618(39)	0.9642(32)	0.9626(38)	0.9626(37)
	R_a	0.7495(39)	0.7493(36)	0.7501(38)	0.7515(39)
32I	Z_l	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	Z_h	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	R_a	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
32ID	Z_l	0.9225(90)	0.9310(87)	0.9189(83)	0.9170(86)
	Z_h	0.9209(85)	0.9279(75)	0.9210(82)	0.9228(84)
	R_a	0.5857(60)	0.5838(55)	0.5855(57)	0.5872(60)
32I-fine	Z_l	0.998(30)	1.003(31)	1.003(31)	0.997(33)
	Z_h	0.999(19)	1.012(20)	1.001(19)	0.994(21)
	R_a	1.311(16)	1.323(16)	1.315(16)	1.308(17)
48I	Z_l	0.9710(53)	0.9698(46)	0.9702(51)	0.9691(50)
	Z_h	0.9618(39)	0.9642(32)	0.9626(38)	0.9626(37)
	R_a	0.7299(51)	0.7205(43)	0.7304(50)	0.7311(48)
64I	Z_l	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	Z_h	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$	$\equiv 1.0$
	R_a	0.9963(60)	0.9816(52)	0.9963(58)	0.9968(57)
32ID-M1	Z_l	0.719(12)	0.7291(86)	0.720(11)	0.7192(84)
	Z_h	0.7303(100)	0.7552(71)	0.7345(98)	0.7368(78)
	R_a	0.4301(57)	0.4332(41)	0.4323(57)	0.4351(44)
32ID-M2	Z_l	—	1.023(11)	—	1.027(12)
	Z_h	—	1.0300(84)	—	1.0405(93)
	R_a	—	0.8541(59)	—	0.8585(64)

Table 3.25: Ratios of lattice spacings (R_a) and light and heavy quark masses (Z_l, Z_h) between each ensemble and the reference 32I ensemble.

LEC	Λ_χ	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
B	—	4.246(22) GeV	4.234(18) GeV	4.235(26) GeV	4.238(22) GeV
f		0.12298(93) GeV	0.12153(77) GeV	0.1226(13) GeV	0.1229(11) GeV
$10^3\hat{L}_0^{(2)}$	1 GeV	—	—	-4.5(4.8)	-0.2(2.0)
$10^3\hat{L}_1^{(2)}$		—	—	0.7(1.2)	-0.30(57)
$10^3\hat{L}_2^{(2)}$		—	—	-4.4(3.3)	-0.9(1.3)
$10^3\hat{L}_3^{(2)}$		—	—	1.4(2.5)	-0.8(1.2)
$10^3\hat{L}_4^{(2)}$		-0.193(77)	0.024(55)	-0.36(36)	-0.48(20)
$10^3\hat{L}_5^{(2)}$		0.479(82)	0.448(48)	0.94(49)	0.69(29)
$10^3\hat{L}_6^{(2)}$		-0.165(48)	-0.004(35)	-0.25(17)	-0.345(99)
$10^3\hat{L}_7^{(2)}$		—	—	-1.60(80)	-0.78(36)
$10^3\hat{L}_8^{(2)}$		0.604(41)	0.532(24)	0.81(22)	0.73(14)
$10^3\hat{L}_0^{(2)}$	770 MeV	—	—	-4.5(5.1)	-0.1(2.0)
$10^3\hat{L}_1^{(2)}$		—	—	0.8(1.2)	-0.20(58)
$10^3\hat{L}_2^{(2)}$		—	—	-4.2(3.4)	-0.7(1.4)
$10^3\hat{L}_3^{(2)}$		—	—	1.6(2.7)	-0.6(1.3)
$10^3\hat{L}_4^{(2)}$		0.014(77)	0.231(55)	-0.15(36)	-0.27(20)
$10^3\hat{L}_5^{(2)}$		0.893(82)	0.862(48)	1.35(48)	1.11(29)
$10^3\hat{L}_6^{(2)}$		-0.010(48)	0.151(35)	-0.09(17)	-0.189(99)
$10^3\hat{L}_7^{(2)}$		—	—	-1.61(84)	-0.78(36)
$10^3\hat{L}_8^{(2)}$		0.604(41)	0.532(24)	0.81(22)	0.73(14)
$10^6\left(\hat{K}_{17}^{(2)}-\hat{K}_{39}^{(2)}\right)$	1 GeV	—	—	-10.1(2.7)	-8.2(1.3)
$10^6\left(\hat{K}_{18}^{(2)}+6\hat{K}_{27}^{(2)}-\hat{K}_{40}^{(2)}\right)$		—	—	20(13)	18.5(5.2)
$10^6\hat{K}_{19}^{(2)}$		—	—	6(25)	-2.9(8.0)
$10^6\hat{K}_{20}^{(2)}$		—	—	-15(16)	-3.9(4.4)
$10^6\left(\hat{K}_{21}^{(2)}+2K_{22}^{(2)}\right)$		—	—	-5.3(7.7)	3.4(3.9)
$10^6\hat{K}_{23}^{(2)}$		—	—	-10.3(5.7)	-2.6(2.3)
$10^6\hat{K}_{25}^{(2)}$		—	—	3.7(7.1)	-0.0(2.8)
$10^6\left(\hat{K}_{26}^{(2)}+6\hat{K}_{27}^{(2)}\right)$		—	—	6.3(7.7)	10.7(3.3)
$10^6\left(\hat{K}_{17}^{(2)}-\hat{K}_{39}^{(2)}\right)$	770 MeV	—	—	-8.3(2.1)	-6.1(1.0)
$10^6\left(\hat{K}_{18}^{(2)}+6\hat{K}_{27}^{(2)}-\hat{K}_{40}^{(2)}\right)$		—	—	12(10)	13.7(3.9)
$10^6\hat{K}_{19}^{(2)}$		—	—	-5(17)	-6.4(4.9)
$10^6\hat{K}_{20}^{(2)}$		—	—	-7(11)	-0.1(2.7)
$10^6\left(\hat{K}_{21}^{(2)}+2K_{22}^{(2)}\right)$		—	—	-5.6(8.0)	4.2(3.5)
$10^6\hat{K}_{23}^{(2)}$		—	—	-5.6(4.8)	-0.0(19)
$10^6\hat{K}_{25}^{(2)}$		—	—	-1.0(5.2)	-2.3(1.9)
$10^6\left(\hat{K}_{26}^{(2)}+6\hat{K}_{27}^{(2)}\right)$		—	—	4.2(7.0)	9.5(2.9)

Table 3.26: $SU(2)$ PQChPT LECs fit at two different chiral scales — $\Lambda_\chi = 1\,\text{GeV}$ and $\Lambda_\chi = 770\,\text{MeV}$ — in units of the canonical size at a given order in the chiral expansion. The LECs $\hat{L}_7^{(2)}$ and $\hat{L}_8^{(2)}$ have no scale dependence. The value of B quoted here is unrenormalized.

Parameter	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
$m^{(K)}$	0.4863(21) GeV	0.4857(17) GeV	0.4863(21) GeV	0.4863(18) GeV
$f^{(K)}$	0.15201(94) GeV	0.15108(81) GeV	0.15187(92) GeV	0.15121(86) GeV
$10^3\lambda_1$	3.1(1.0)	4.56(80)	3.06(99)	3.2(1.0)
$10^3\lambda_2$	28.62(45)	28.36(42)	28.57(65)	28.87(57)
$10^3\lambda_3$	-4.01(98)	-2.33(77)	-3.91(97)	-4.04(87)
$10^3\lambda_4$	5.74(38)	6.18(48)	5.74(39)	5.93(47)
c_f^I	0.007(22) GeV^2	0.022(19) GeV^2	0.018(25) GeV^2	0.021(24) GeV^2
c_f^{ID}	-0.012(13) GeV^2	0.016(10) GeV^2	-0.000(16) GeV^2	0.008(14) GeV^2
$c_{f^{(K)}}^I$	0.004(17) GeV^2	0.009(16) GeV^2	0.006(17) GeV^2	0.017(16) GeV^2
$c_{f^{(K)}}^{ID}$	-0.003(11) GeV^2	0.0131(82) GeV^2	-0.001(11) GeV^2	0.0175(78) GeV^2
c_{m_h,m_π^2}	3.5(3.9)	0.1(3.4)	3.5(3.3)	0.2(2.9)
c_{m_h,f_π}	0.09(12)	0.116(96)	0.14(12)	0.184(92)
c_{m_y,m_K^2}	3.939(18) GeV	3.953(15) GeV	3.934(18) GeV	3.930(16) GeV
c_{m_h,m_K^2}	0.040(67) GeV	0.167(76) GeV	0.048(66) GeV	0.017(63) GeV
c_{m_y,f_K}	0.2903(88)	0.2944(86)	0.2879(84)	0.3228(93)
c_{m_h,f_K}	0.067(50)	0.042(44)	0.050(57)	0.108(44)
$m^{(\Omega)}$	1.6645(36) GeV	1.6614(25) GeV	1.6643(34) GeV	1.6651(29) GeV
c_{m_l,m_Ω}	3.63(64)	5.05(57)	3.73(65)	3.33(63)
c_{m_y,m_Ω}	5.678(81)	5.39(12)	5.633(78)	5.537(74)
c_{m_h,m_Ω}	1.99(48)	1.23(41)	1.80(52)	1.52(40)

Table 3.27: Additional fit parameters in physical units and adjusted to the physical strange quark mass. Here $\{m^{(K)}, f^{(K)}\}$ and $\{\lambda_i\}$ are the LO and NLO LECs of heavy-meson $SU(2)$ PQChPT evaluated at the chiral scale $\Lambda_\chi = 1\,\text{GeV}$. c_f^I and c_f^{ID} are the a^2 coefficients of f_π for the Iwasaki and Iwasaki+DSDR gauge actions, respectively, and likewise for $c_{f^{(K)}}^I$ and $c_{f^{(K)}}^{ID}$. The notation $c_{m_q,X}$ denotes the coefficient of a term linear in m_q for quantity X, and $m^{(\Omega)}$ is the constant term in the (linear) m_Ω ansatz.

3.D.2 Predictions

LEC	Λ_χ	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
$B^{\overline{\rm MS}}(\mu = 2\,{\rm GeV})$		2.815(33)(30) GeV	2.808(31)(30) GeV	2.808(36)(30) GeV	2.811(35)(30) GeV
f	—	123.0(9) MeV	121.5(8) MeV	122.6(1.3) MeV	122.9(1.1) MeV
$\Sigma^{1/3,\,\overline{\rm MS}}(\mu = 2\,{\rm GeV})$		277.2(1.8)(1.0) MeV	274.7(1.5)(1.0) MeV	276.3(2.1)(1.0) MeV	276.9(1.9)(1.0) MeV
$10^3 l_1$		—	—	15(19)	-2.4(7.6)
$10^3 l_2$		—	—	-35(32)	-5(13)
$10^3 l_3$	1 GeV	1.82(26)	2.22(20)	1.62(79)	1.36(56)
$10^3 l_4$		0.37(52)	1.98(36)	0.8(1.5)	-1.05(99)
$10^3 l_7$		—	—	19(12)	6.7(5.4)
$10^3 l_1$		—	—	16(19)	-1.8(7.7)
$10^3 l_2$		—	—	-35(33)	-3(13)
$10^3 l_3$	770 MeV	0.99(26)	1.39(20)	0.78(79)	0.54(56)
$10^3 l_4$		3.68(52)	5.29(36)	4.2(1.6)	2.26(99)
$10^3 l_7$		—	—	19(13)	6.7(5.4)
$\bar{\ell}_1$		—	—	18(18)	1.8(7.2)
$\bar{\ell}_2$		—	—	-13(15)	1.9(6.2)
$\bar{\ell}_3$	—	2.86(16)	2.61(12)	2.98(50)	3.14(35)
$\bar{\ell}_4$		4.064(82)	4.318(57)	4.14(24)	3.84(16)

Table 3.28: Unquenched $SU(2)$ LECs computed from partially quenched $SU(2)$ fits. Missing entries are not constrained by the fits at a given order. For B and Σ the first error is statistical and the second is a systematic uncertainty in the perturbative matching to $\overline{\rm MS}$.

	NLO (370 MeV cut)	NLO (450 MeV cut)	NNLO (370 MeV cut)	NNLO (450 MeV cut)
f_π	0.13074(84) GeV	0.12986(71) GeV	0.13032(94) GeV	0.13011(89) GeV
f_K	0.15587(79) GeV	0.15542(70) GeV	0.15577(78) GeV	0.15508(71) GeV
f_K/f_π	1.1922(41)	1.1968(39)	1.1953(59)	1.1919(56)
f_π/f	1.0631(18)	1.0686(13)	1.0631(44)	1.0583(29)
$m_\pi a_0^0$	—	—	0.153(33)	0.185(14)
$m_\pi a_0^2$	—	—	-0.057(13)	-0.0431(53)
$[m_{\pi^\pm}^2 - m_{\pi^0}^2]_{\rm QCD}/\Delta m_{du}^2$	—	—	91(57)	32(26)

Table 3.29: Predictions from NLO and NNLO fits and $SU(2)$ ChPT. $\Delta m_{du} \equiv m_d - m_u$.

Chapter 4

Next-to-Next-to Leading Order $SU(3)$ Chiral Perturbation Theory

The Low Energy Constants of $SU(3)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD

C. Kelly[1], R.D. Mawhinney[1], and D.J. Murphy[1]

[1] *Department of Physics, Columbia University, New York, NY 10027, USA*

Abstract

We have performed fits of the pseudoscalar masses and decay constants, from a variety of RBC-UKQCD domain wall fermion ensembles, to $SU(3)$ partially quenched chiral perturbation theory at next-to-next-to leading order, following the approach of our recent paper on the $SU(2)$ theory [1]. These ensembles cover a wide range of unitary pion masses, ranging from a lighter than physical 117 MeV up to 432 MeV. We report values for 9 NLO and 10 linearly independent

combinations of NNLO partially quenched low energy constants, which we compare to other lattice and phenomenological determinations. We discuss the convergence of the expansion and use our large set of low energy constants to make predictions for mass and decay constant splittings due to QCD isospin breaking effects in the kaon sector, and for the S-wave πK scattering lengths. We find that, with the inclusion of new data near the physical point, we are able to successfully fit NLO $SU(3)$ PQChPT to data with two light valence quarks, in contrast to earlier, unsuccessful RBC-UKQCD fits to a heavier subset of our current data. We also find that after including NNLO terms we are able to fit data up to the scale of the physical kaon with percent-scale accuracy. We conclude that, for the range of pseudoscalar masses and decay constants explored in this work, the NNLO $SU(3)$ expansion is accurate but is likely nearing the limits of its applicability at the scale of the physical kaon.

4.1 Introduction

Quantum Chromodynamics (QCD), the quantum field theory describing quarks and their interactions via the strong nuclear force, is widely believed to give the correct theoretical description of the mesons and baryons observed in nature. However, the same highly nonlinear dynamics which successfully predicts that quarks will bind together into hadrons at low energies also makes many analytic calculations intractable. Perturbative expansions of QCD correlation functions in terms of the strong coupling constant fail at low energies, and an alternative approach is needed to analytically compute the properties of hadrons in this regime.

Lattice QCD provides the only known first-principles approach for performing fully non-perturbative QCD calculations: after reformulating QCD on a finite, discrete spacetime lattice powerful supercomputers can be used to solve the resulting equations directly. A typical lattice QCD calculation is performed with several values of the lattice spacing, simulation volume, and input quark masses. Contact with experimental results can then be made by interpolating or extrapolating the simulated points to the infinite volume, continuum, and physical quark mass limit. In this work we make use of a series of lattice ensembles generated by the RBC-UKQCD collaboration and used to compute, in particular, properties of the light pseudoscalar meson spectrum. These ensembles utilize the domain wall fermion (DWF) formalism, which provides excellent chiral symmetry prop-

erties at the expense of introducing a fifth spacetime dimension. With the recent development of simulations performed directly at physical masses RBC-UKQCD has shown that lattice QCD can predict observables such as the pion and kaon decay constants with sub-percent accuracy and in agreement with experiment [2].

An alternative approach to low energy QCD is to apply effective field theory (EFT) techniques. Observing that the QCD Lagrangian with N_f massless quarks has an exact $SU(N_f)_L \times SU(N_f)_R$ symmetry, and that the physical hadron spectrum contains $N_f^2 - 1$ approximately degenerate light pseudoscalar mesons, suggests an EFT description, known as chiral perturbation theory (ChPT). In this formalism the light pseudoscalars are realized as pseudo-Goldstone bosons associated with the spontaneous breaking of chiral symmetry $SU(N_f)_L \times SU(N_f)_R \to SU(N_f)_V$ by the QCD vacuum; these mesons are light rather than massless since chiral symmetry is also explicitly broken by the non-zero quark masses found in nature. One can regard ChPT as a low-energy effective field theory expansion in powers of masses and momenta which becomes exact in the limit of massless quarks. In this work we focus on the $N_f = 3$ variant of ChPT, which allows for systematic calculations of the properties of the pseudoscalar octet (π, K, η) in terms of the up, down, and strange quark masses.

The $SU(3)$ ChPT Lagrangian can be constructed following a general prescription introduced by Weinberg [3]: one picks a power counting scheme and writes down the most general Lagrangian for the light pseudoscalar mesons consistent with an $SU(3)_L \times SU(3)_R$ symmetry order-by-order. The operators which appear in this Lagrangian are parameterized by *a priori* unknown low-energy constants (LECs). These LECs encode the matching of ChPT to QCD and must be determined by experimental or lattice constraints. $SU(3)$ ChPT was first explicitly constructed and explored at next-to leading order (NLO) by Gasser and Leutwyler [4]. Full calculations of the pseudoscalar masses and decay constants at next-to-next-to leading order (NNLO) in the more general framework of partially quenched ChPT (PQChPT) — where the valence and sea quark masses are allowed to differ — were later performed by Bijnens, Danielsson, and Lähde [5]. We make use of Fortran codes provided by Bijnens to compute the NNLO terms in our fits.

In this paper we follow up our recent analysis of NNLO $SU(2)$ PQChPT [1] with an analogous

study of the $SU(3)$ case. In particular, we seek to use RBC-UKQCD's set of domain wall fermion ensembles to:

1. Determine as many of the low energy constants of $SU(3)$ ChPT as possible from our data, and

2. Systematically study the behavior and range of applicability of the $SU(3)$ ChPT expansions up to next-to-next-to leading order.

Earlier NLO $SU(3)$ PQChPT fits to a smaller RBC-UKQCD domain wall QCD data set with relatively heavy pion masses ($250-420$ MeV) were performed in Ref. [6], but were deemed unreliable. These fits predicted a suspiciously low value of the $SU(3)$ chiral decay constant $f_0 \sim 93.5\,\text{MeV}$ as well as large NLO corrections that were $\sim 70\%$ of LO at a scale of 400 MeV. In subsequent RBC-UKQCD works chiral extrapolations based on $SU(3)$ ChPT were abandoned in favor of more reliable $SU(2)$ ChPT-based extrapolations. In addition, the MILC collaboration has studied the pseudoscalar mass and decay constant in $SU(3)$ NNLO PQChPT using staggered fermions [7, 8]. Here we revisit NLO and NNLO partially quenched $SU(3)$ ChPT fits using the current RBC-UKQCD domain wall fermion data set, which now includes physical and even lighter-than-physical pion mass ensembles. Our DWF ensembles, which preserve continuum chiral symmetries even at finite lattice spacing, provide an ideal laboratory for testing ChPT fits and the reliability of the $SU(3)$ expansion.

4.2 Lattice Setup

In Table 4.1 we list the 12 ensembles included in this analysis and summarize the actions and input parameters. These ensembles comprise the same data set used for our recent $SU(2)$ analysis [1], and cover a wide range of unitary pion masses, physical volumes, and inverse lattice spacings. We use the Iwasaki gauge action (I) [9] in all cases, and on some ensembles supplement this with the dislocation suppressing determinant ratio (I+DSDR) [10, 11]. We simulate QCD with two degenerate light quark flavors of bare mass m_l and a single heavy quark flavor of bare mass m_h using the domain wall fermion formalism, with either the Shamir (DWF) [12, 13] or Möbius (MDWF) [14–16] kernel.

Additional details of the ensemble generation and fits to extract the low-energy QCD spectrum can be found in Ref. [17] for the 24I ensembles, Ref. [18] for the 32I ensembles, Ref. [19] for the 32ID ensembles, Ref. [2] for the 48I, 64I, and 32I-fine ensembles, and Ref. [1] for the 32ID-M1 and 32ID-M2 ensembles. Appendix B of Ref. [1] contains an explicit summary of the pseudoscalar mass, pseudoscalar decay constant, and Ω baryon measurements which enter into the chiral fits.

Ensemble	Action	β	$L^3 \times T \times L_s$	am_l	am_h	$m_\pi L$	m_π (MeV)
24I	DWF+I	2.13	$24^3 \times 64 \times 16$	0.005	0.04	4.568(13)	339.6(1.2)
	DWF+I	2.13	$24^3 \times 64 \times 16$	0.01	0.04	5.814(12)	432.2(1.4)
32I	DWF+I	2.25	$32^3 \times 64 \times 16$	0.004	0.03	4.062(11)	302.0(1.1)
	DWF+I	2.25	$32^3 \times 64 \times 16$	0.006	0.03	4.8377(82)	359.7(1.2)
	DWF+I	2.25	$32^3 \times 64 \times 16$	0.008	0.03	5.526(12)	410.8(1.5)
32ID	DWF+I+DSDR	1.75	$32^3 \times 64 \times 32$	0.001	0.046	3.9992(69)	172.7(9)
	DWF+I+DSDR	1.75	$32^3 \times 64 \times 32$	0.0042	0.046	5.7918(79)	250.1(1.2)
32I-fine	DWF+I	2.37	$32^3 \times 64 \times 12$	0.0047	0.0186	3.773(42)	370.1(4.4)
48I	MDWF+I	2.13	$48^3 \times 96 \times 24$	0.00078	0.0362	3.8633(63)	139.1(4)
64I	MDWF+I	2.25	$64^3 \times 128 \times 12$	0.000678	0.02661	3.7778(84)	139.0(5)
32ID-M1	MDWF+I+DSDR	1.633	$32^3 \times 64 \times 24$	0.00022	0.0596	3.780(15)	117.3(4.4)
32ID-M2	MDWF+I+DSDR	1.943	$32^3 \times 64 \times 12$	0.00478	0.03297	6.236(21)	401.0(2.3)

Table 4.1: Summary of ensembles included in this analysis and input parameters. Here β is the gauge coupling, $L^3 \times T \times L_s$ is the lattice volume decomposed into the length of the spatial (L), temporal (T), and fifth (L_s) dimensions, and am_l and am_h are the bare, input light and heavy quark masses. The value of m_π quoted is the unitary pion mass in physical units, where we have used the lattice spacings from our canonical NLO $SU(2)$ global fit [1].

4.3 The $SU(3)$ Global Fit Procedure

In Ref. [2, 18, 19] we have developed a "global fit" procedure for performing a combined chiral fit and continuum extrapolation of lattice data, which was adapted to study partially quenched $SU(2)$ chiral perturbation theory at next-to-next-to leading order in Ref. [1]. In addition to determining low energy constants of ChPT, the global fit also allows us to convert predictions from our simulations, which are performed in dimensionless lattice units, into physical units by determining the lattice spacing a on each ensemble. In this section we briefly review the global fit procedure and highlight some of the differences between the $SU(3)$ case and the $SU(2)$ case.

The $SU(3)$ global fits performed in this work include data for the (in general partially quenched) pseudoscalar mass (m_{xy}) and decay constant (f_{xy}), as well as the omega baryon mass m_Ω. Partially quenched NLO or NNLO $SU(3)$ ChPT with NLO finite volume corrections is used to perform the chiral fit to the valence quark (m_x, m_y) and sea quark (m_l, m_h) mass dependence[1] of m_{xy} and f_{xy}. The chiral fit to m_Ω is performed using a simple analytic ansatz which is linear in the quark masses. Since the raw simulation data is in lattice units which are different for each ensemble we choose our $32^3 \times 64$ Iwasaki (32I) lattice as a reference ensemble, and introduce additional fit parameters

$$R_a^e \equiv \frac{a^r}{a^e}, \quad Z_l^e \equiv \frac{1}{R_a^e}\frac{(a\tilde{m}_l)^r}{(a\tilde{m}_l)^e}, \quad Z_h^e \equiv \frac{1}{R_a^e}\frac{(a\tilde{m}_h)^r}{(a\tilde{m}_h)^e} \tag{4.1}$$

to convert between bare lattice units on the reference ensemble (r) and other ensembles (e). The $SU(3)$ chiral ansatzäe for $X \in \{m_{xy}^2, f_{xy}\}$ have the generic form

$$X(\tilde{m}_q, L) \simeq X_0 \left(1 + \underbrace{X^{\text{NLO}}(\tilde{m}_q) + X^{\text{NNLO}}(\tilde{m}_q)}_{\text{NNLO Continuum PQChPT}} + \underbrace{\Delta_X^{\text{NLO}}(\tilde{m}_q, L)}_{\text{NLO FV corrections}} \right), \tag{4.2}$$

reflecting a simultaneous expansion in the total quark masses and lattice volume (L).

The procedure for performing an $SU(3)$ global fit is as follows:

1. The valence quark mass dependence of m_{res} is fit to a linear ansatz on each ensemble. We then extrapolate m_{res} to the chiral limit $m_q \to 0$, and use this value in the remainder of the analysis.

[1] Note that this differs from the $SU(2)$ case, where a series of reweightings in m_h was computed for each combination of (m_x, m_y, m_l), and used to perform a small linear interpolation in m_h to the physical strange quark mass. For $SU(3)$ fits the m_h dependence of m_{xy} and f_{xy} is described directly by ChPT, rendering this procedure unnecessary.

2. A simultaneous chiral/continuum fit of m_{xy}^2, f_{xy}, and m_Ω is performed on all ensembles using the ansätze described in the preceding paragraph. The quark mass dependence is parametrized in terms of $\tilde{m}_q = m_q + m_{\text{res}}$. This step also determines the ratios of lattice scales R_a^e and the scaling coefficients $Z_{\{l,h\}}^e$.

3. We match onto a continuum scaling trajectory by numerically inverting the chiral fit to determine input bare valence quark masses m_l^{phys} and m_s^{phys} such that the ratios m_π/m_Ω and f_π/m_Ω take their PDG experimental values [20].

4. From 3 we obtain m_Ω at m_l^{phys} and m_s^{phys} on the reference ensemble; we then use the ratio $m_\Omega^r/m_\Omega^{\text{phys}}$ to determine the lattice spacing a^r in physical units. Together with the ratios of lattice scales from 2 we can determine the lattice spacings on the other ensembles, as well as extrapolate observables to the physical point in physical units.

Renormalization-scheme dependent quantities — in particular, the physical quark masses m_l^{phys} and m_s^{phys}, and the leading-order $SU(3)$ LEC B_0 — are converted to $\overline{\text{MS}}$ using the same procedure and renormalization coefficients we have used in our earlier $SU(2)$ fits: we first renormalize in variants of the non-perturbative Rome-Southampton regularization-invariant momentum scheme with symmetric kinematics (RI-SMOM), and then perturbatively match to $\overline{\text{MS}}$. We direct the interested reader to Section 4 of Ref. [1] for additional detail.

While this procedure is largely the same as that of the $SU(2)$ global fits described in our earlier work, there are a few important differences. First, for $SU(3)$ fits we match onto a continuum scaling trajectory by forcing m_π and f_π to take their physical values rather than m_π and m_K, as we have conventionally used for $SU(2)$ fits. This is motivated by our expectation that $SU(3)$ ChPT ought to be reliable at the pion scale, but may not be as reliable at the kaon scale. Second, the $SU(2)$ analogues of the chiral ansatzäe defined schematically by Eqn. (4.2) also included discretization terms proportional to a^2 for f_π and f_K, which were not used to set the scale or match onto a scaling trajectory. Our $SU(3)$ fits contain exactly three independent observables — m_{xy}, f_{xy}, and m_Ω — all of which are used to determine the lattice spacings and continuum scaling trajectory, and thus we did not similarly include a^2 terms.

The fits described in steps 1 and 2 are performed using uncorrelated nonlinear χ^2 minimization with the Levenberg-Marquardt algorithm [21, 22]. In Appendix D of Ref. [1] we demonstrated that the large number of data points in our fits leads to an ill-conditioned correlation matrix, forcing us to perform fits by minimizing the uncorrelated χ^2:

$$\chi^2 = \sum_i \left(\frac{y^i - f^i}{\sigma^i} \right)^2 . \tag{4.3}$$

As a result, the χ^2/dof that we report cannot be interpreted as a rigorous statistical measure of the goodness-of-fit, and instead we present histograms which show the distribution of the data around our fits. In the same appendix we also argue that there is a potential systematic bias introduced by our use of uncorrelated fits: some subsets of our data consist of partially quenched measurements computed on a single ensemble with different valence quark mass combinations, and this data is strongly correlated. Furthermore, since these partially quenched measurements were computed on our older, relatively heavy pion mass ensembles — which are far from the chiral limit where ChPT is exact — a naive uncorrelated fit will tend to give too much weight to this data and potentially bias the results. We argued that we could estimate the associated systematic error by introducing a second set of weighted fits

$$\chi^2_e = \frac{1}{N_e} \sum_{j \in e} \left(\frac{y^j_e - f^j_e}{\sigma^j_e} \right)^2 , \ \chi^2 = \sum_e \chi^2_e \tag{4.4}$$

where the contribution to χ^2 from each ensemble e was normalized by the number of nondegenerate valence mass combinations computed on that ensemble (N_e), and we used the difference in central value between fit parameters from a weighted and an unweighted fit as an estimate of this systematic. We adopt the same procedure here: the results of unweighted fits are reported in Section 4.4, the results of weighted fits in Appendix 4.B, and the difference between the two fits is factored into our final error budget. We direct the reader to Appendix D of our previous work for more detail regarding this procedure.

4.4 Fits to $SU(3)$ PQχPT

In this section we discuss global fits of the pseudoscalar mass, psuedoscalar decay constant, and Ω baryon to $SU(3)$ partially quenched chiral perturbation theory. In particular, pseudoscalar quantities are fit to NLO or NNLO $SU(3)$ PQChPT with NLO finite volume corrections, while the Ω baryon mass is fit to a linear, analytic ansatz. We determine the physical quark masses and the lattice spacings by constraining m_π, f_π, and m_Ω to take their experimental, PDG values; we can then predict m_K and f_K from $SU(3)$ ChPT with the LECs and physical quark masses obtained from our fits. We remind the reader that, in contrast to the $SU(2)$ fits, we do not determine a^2 scaling coefficients for f_π and f_K in our $SU(3)$ fits. This is because there are exactly three observables in the fits and three constraints are required to match onto a continuum scaling trajectory in our formalism. We also do not include any data which has been reweighted in m_h, since the m_h dependence of observables is directly parametrized by $SU(3)$ PQChPT.

We consider two different mass cuts: a 370 MeV cut which uses exactly the same (pion) data as the corresponding $SU(2)$ fit from Ref. [1], and a 510 MeV cut which uses all partially quenched measurements with $m_{xy} \leq 510\,\text{MeV}$, including the unitary kaon on our lightest ensembles. These mass cuts were motivated by RBC-UKQCD's earlier experience with fitting $SU(3)$ PQChPT to the 24I lattice data, where it was observed that good fits to the partially quenched pion were possible, but that these fits broke down as heavier measurements up to the kaon scale were included [17]. We emphasize, however, that the ensembles included in this earlier analysis were all at very heavy quark mass — the lightest unitary kaon, for example, had a mass of $m_K \sim 600\,\text{MeV}$, well outside the cuts in this work.

In contrast to the $SU(2)$ case, where the RBC-UKQCD collaboration has extensive experience with successful fits of ChPT to lattice data, much less is known about the applicability of the $SU(3)$ theory. The significantly heavier mass of the strange quark implies that the $SU(3)$ expansion in m_h may converge much more slowly than the $SU(2)$ expansion in m_l, and indeed, the RBC-UKQCD collaboration's earlier attempts to fit NLO $SU(3)$ PQChPT to data from the 24I ensemble were unsuccessful [17]. In light of this we relax our expectations somewhat for the $SU(3)$ fits: we regard the results in this section as an exploratory study of whether or not we can reliably fit NLO and/or

NNLO $SU(3)$ PQChPT to our data set — which now includes physical pion mass ensembles — and whether or not $SU(3)$ ChPT can make predictions at the physical point with percent-scale accuracy once values for the low-energy constants have been determined.

In Section 4.4.1 we discuss implementation details which are specific to the NNLO $SU(3)$ fits. In Sections 4.4.2-4.4.4 we present the fit results, including our values for the partially quenched NLO and NNLO LECs. In Section 4.4.5 we examine the range of applicability of NNLO $SU(3)$ ChPT and the relative sizes of the terms in the chiral expansion. In Section 4.4.6 we discuss our new results in relation to our previous conclusions regarding $SU(3)$ ChPT from Ref. [17]. Finally, in Section 4.4.7 we compute the unquenched $SU(2)$ and $SU(3)$ ChPT LECs from these results, and also discuss other predictions we can make. All fits are performed using the Marquardt-Levenberg algorithm to minimize the uncorrelated χ^2; in Appendix 4.B we repeat the fits using a weighted χ^2 to explore systematic effects associated with correlations in the data.

4.4.1 Implementation Details for NNLO $SU(3)$ Fits

There is an additional numerical obstacle to implementing the global fit procedure for NNLO $SU(3)$ fits, which we do not observe to be an issue in the other cases: while the PQChPT expressions for the pseudoscalar mass m_{xy}^2 and decay constant f_{xy} are smooth functions of the quark masses (m_x, m_y, m_l, m_h), there are particular limits which must be handled with care in the numerical implementation. As a concrete example, consider the mass of a meson containing a light sea quark and a valence quark $m_x \neq m_l$ at NLO:

$$\begin{aligned} m_{lx}^2 = \frac{1}{2}(\chi_l + \chi_x)\Bigg\{1 + \frac{48}{f_0^2}\left(2\hat{L}_6^{(3)} - \hat{L}_4^{(3)}\right)\overline{\chi} + \frac{8}{f_0^2}\left(2\hat{L}_8^{(3)} - \hat{L}_5^{(3)}\right)(\chi_l + \chi_x) \\ + \frac{1}{24\pi^2 f_0^2}\left[\frac{\chi_x - \chi_h}{\chi_x - \chi_\eta}\chi_x \log\left(\frac{\chi_x}{\Lambda_\chi^2}\right) + \frac{\chi_\eta - \chi_h}{\chi_\eta - \chi_x}\chi_\eta \log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right)\right]\Bigg\} \end{aligned} \tag{4.5}$$

where $\chi_q \equiv 2B_0 m_q$, $\overline{\chi} \equiv 2B_0(2m_l + m_h)/3$, and $\chi_\eta \equiv 2B_0(m_l + 2m_h)/3$. While the limit $\chi_x \to \chi_\eta$ has a completely well-defined analytic expression

$$\begin{aligned} m_{lx}^2 = \frac{1}{2}(\chi_l + \chi_\eta)\Bigg\{1 + \frac{48}{f_0^2}\left(2\hat{L}_6^{(3)} - \hat{L}_4^{(3)}\right)\overline{\chi} + \frac{8}{f_0^2}\left(2\hat{L}_8^{(3)} - \hat{L}_5^{(3)}\right)(\chi_l + \chi_\eta) \\ + \frac{1}{24\pi^2 f_0^2}\left[\chi_\eta - \chi_h + (2\chi_\eta - \chi_h)\log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right)\right]\Bigg\}, \end{aligned} \tag{4.6}$$

numerical instabilities are encountered if one instead naively tries to compute m_{lx}^2 from (4.5) with $m_x \approx m_\eta$. At NLO one can explicitly work out all such cases and take care to use an appropriate analytic form of the expression for the pseudoscalar mass or decay constant for each set of input quark masses, but at NNLO this exercise is considerably more complicated, and has not been worked out in all cases in the Fortran routines provided by J. Bijnens.

We use a simple scheme to catch and interpolate around quark mass combinations where the routines become singular, which we illustrate here using the pseudoscalar mass. For a given combination of quark masses we compute $m_{xy}^2(m_x, m_y, m_l, m_h)$, as well as the four nearest-neighbors in the valence quark plane, $m_{xy}^2(m_x \pm \Delta m_x, m_y \pm \Delta m_y, m_l, m_h)$, where $\Delta m_x/m_x \ll 1$ and $\Delta m_y/m_y \ll 1$ are small perturbations. We then compute the two-point

$$\begin{cases} m_{\overline{x}y}^2 = \frac{1}{2}\left[m_{xy}^2(m_x + \Delta m_x, m_y, m_l, m_h) + m_{xy}^2(m_x - \Delta m_x, m_y, m_l, m_h)\right] \\ m_{x\overline{y}}^2 = \frac{1}{2}\left[m_{xy}^2(m_x, m_y + \Delta m_y, m_l, m_h) + m_{xy}^2(m_x, m_y - \Delta m_y, m_l, m_h)\right] \end{cases} \tag{4.7}$$

and four-point

$$\begin{aligned} m_{\overline{xy}}^2 = \frac{1}{4}\big[&m_{xy}^2(m_x + \Delta m_x, m_y, m_l, m_h) + m_{xy}^2(m_x, m_y + \Delta m_y, m_l, m_h) \\ &+ m_{xy}^2(m_x - \Delta m_x, m_y, m_l, m_h) + m_{xy}^2(m_x, m_y - \Delta m_y, m_l, m_h)\big] \end{aligned} \tag{4.8}$$

interpolations, as well as the slopes $\partial m_{xy}^2/\partial m_x$ and $\partial m_{xy}^2/\partial m_y$. While one should generally have $m_{xy}^2 \cong m_{\overline{x}y}^2 \cong m_{x\overline{y}}^2 \cong m_{\overline{xy}}^2$ and $\partial m_{xy}^2/\partial m_x = \partial m_{xy}^2/\partial m_y$ to arbitrary precision if Δm_x and Δm_y are sufficiently small, this will fail if either the direct calculation or one of the nearest-neighbors happens to be evaluated at a point where the NNLO routines are poorly behaved. By checking these constraints we can decide which of $\{m_{xy}^2, m_{\overline{xy}}^2, m_{\overline{x}y}^2, m_{x\overline{y}}^2\}$ to take as the value of the pseudoscalar mass, depending on which comparisons fail. In practice this interpolation scheme is sufficient to catch and accurately approximate cases where the direct calculation encounters numerical instabilities.

4.4.2 Fit Parameters

Tables 4.2-4.4 summarize the parameters of the fits we have performed, where the errors are purely statistical. These include the χ^2/dof, physical quark masses, and inverse lattice spacings in physical

units (Table 4.2), the ratios of quark masses and lattice spacings between the reference 32I ensemble and the other ensembles (Table 4.3), and the $SU(3)$ PQChPT LECs (Table 4.4). The counting of the low-energy constants is slightly different between the $SU(2)$ and $SU(3)$ cases, since the additional degree of freedom associated with the heavy dynamical quark in $SU(3)$ breaks some of the degeneracy of the NNLO LECs. For $SU(3)$ there are ten non-degenerate linear combinations at NNLO, listed explicitly in Table 4.4. We set $\hat{K}_{39}^{(3)} = \hat{K}_{40}^{(3)} = 0$ in the fits to simplify the 10 linear combinations to 10 independent fit parameters.

For the NNLO fits, we consider two different schemes. In the first scheme, which we refer to as "free" fits, all fit parameters are allowed to vary freely with no constraints. In the second scheme, which we refer to as fits with "frozen LO LECs", the leading-order low-energy constants B_0 and f_0 are frozen superjackknife-block-by-superjackknife-block to the values of B_0 and f_0 obtained from the NLO fit with the same mass cut. This was motivated by the observation that the chiral decay constant f_0 tends to run off to $f_0 \sim 130\,\text{MeV}$ in the free NNLO fits, indicating either an intrinsic lack of reliability in the series when NNLO terms are added — since LO and NLO effects are mixed into NNLO effects — or insufficient numerical data to fully stabilize our fits. For these reasons, we regard the frozen fit as our preferred NNLO $SU(3)$ fit for examining the ordering of the chiral expansion and quoting values of the low energy constants. We note that the MILC collaboration has adopted a similar strategy to stabilize their own NNLO $SU(3)$ fits: they first perform "low-mass" fits to determine lower order LECs, then freeze these LECs and perform a second set of "high-mass" fits to determine the remaining, higher order LECs [7, 8].

We observe excellent χ^2/dof for both the NLO and NNLO fits with $m_{xy}^{\text{cut}} = 370\,\text{MeV}$, indicating that the $SU(3)$ theory has no trouble describing the partially quenched pseudoscalar mass and decay constant up to this scale with two light valence quarks and the strange dynamical quark mass near its physical value. When we extend the mass cut up to the kaon scale, however, the NLO fit becomes obviously strained, with $\chi^2/\text{dof} = 6.5$. Including NNLO terms leads to a more reasonable $\chi^2/\text{dof} = 1.2$ for the free fit and $\chi^2/\text{dof} = 2.3$ for the frozen fit. We also observe the expected hierarchy of terms in the chiral LECs, with LO $\sim \mathcal{O}(1)$, NLO $\sim \mathcal{O}(10^{-3})$, and NNLO $\sim \mathcal{O}(10^{-6})$. Since the region of applicability of $SU(3)$ PQChPT remains poorly understood, it is a priori unclear

whether this tension between chiral perturbation theory and the lattice data arises from unitary quark masses which are too far from the $SU(3)$ chiral limit, or from valence quark masses that differ too much from the dynamical quark masses. We have explored restricting the degree of partial quenching — we can place a second cut, for example, on the ratios $m_q^{\text{val}}/m_q^{\text{sea}}$ — but find that pruning partially quenched points from our analysis does little to improve the χ^2/dof, at the expense of poorly resolved NNLO fits.

				Free		**Frozen LO LECs**
		NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
	χ^2/dof	0.46	6.48	0.34	1.16	2.35
	$N_{\text{parameters}}$	26	29	41	44	42
	N_{data}	94	201	94	201	201
	am_l^{phys}	-0.001848(46)	-0.001818(44)	-0.001827(80)	-0.001883(62)	-0.001881(65)
24I	am_h^{phys}	0.03342(46)	0.03441(42)	0.0358(43)	0.03178(75)	0.0312(12)
	a^{-1}	1.735(19) GeV	1.718(12) GeV	1.710(42) GeV	1.789(29) GeV	1.828(50) GeV
	am_l^{phys}	0.000285(14)	0.000286(10)	0.000303(18)	0.000287(20)	0.000247(18)
32I	am_h^{phys}	0.02515(45)	0.02523(31)	0.02588(44)	0.02586(40)	0.02483(45)
	a^{-1}	2.312(24) GeV	2.318(16) GeV	2.285(39) GeV	2.336(21) GeV	2.393(28) GeV
	am_l^{phys}	-0.000179(40)	-0.000146(36)	-0.00013(17)	-0.000221(56)	-0.000098(66)
32ID	am_h^{phys}	0.0415(25)	0.0438(25)	0.047(11)	0.0401(25)	0.0414(21)
	a^{-1}	1.402(15) GeV	1.387(15) GeV	1.374(38) GeV	1.432(18) GeV	1.424(31) GeV
	am_l^{phys}	0.000082(28)	0.000067(27)	0.000065(46)	0.000078(32)	0.000026(54)
32I-fine	am_h^{phys}	0.0199(17)	0.01858(81)	0.0192(20)	0.0186(12)	0.0179(13)
	a^{-1}	3.089(32) GeV	3.097(30) GeV	3.096(53) GeV	3.087(35) GeV	3.140(96) GeV
	am_l^{phys}	0.0007018(68)	0.0007225(58)	0.000715(47)	0.000697(33)	0.000727(33)
48I	am_h^{phys}	0.03617(25)	0.03690(27)	0.0383(36)	0.03537(58)	0.0355(13)
	a^{-1}	1.725(2) GeV	1.721(2) GeV	1.710(14) GeV	1.737(9) GeV	1.737(20) GeV
	am_l^{phys}	0.0006240(75)	0.0006279(54)	0.0006333(82)	0.000641(17)	0.000615(14)
64I	am_h^{phys}	0.02507(47)	0.02522(31)	0.02558(58)	0.02618(28)	0.02555(51)
	a^{-1}	2.351(6) GeV	2.350(5) GeV	2.343(8) GeV	2.339(5) GeV	2.359(12) GeV
	am_l^{phys}	0.000730(36)	0.000612(23)	0.00071(14)	0.00048(13)	0.00086(20)
32ID-M1	am_h^{phys}	0.083(11)	0.0783(50)	0.091(33)	0.0628(38)	0.0686(74)
	a^{-1}	1.045(7) GeV	1.043(3) GeV	1.038(23) GeV	1.077(15) GeV	1.043(29) GeV
	am_l^{phys}	—	-0.003349(16)	—	-0.003378(40)	-0.003420(44)
32ID-M2	am_h^{phys}	—	0.02618(68)	—	0.02542(78)	0.0238(11)
	a^{-1}	—	1.953(16) GeV	—	2.011(31) GeV	2.087(46) GeV

Table 4.2: The (uncorrelated) χ^2/dof, unrenormalized physical quark masses in bare lattice units (without m_{res} included), and the values of the inverse lattice spacing a^{-1} in physical units, obtained from fits to $SU(3)$ PQChPT with the stated pion mass cuts.

				Free		Frozen LO LECs
		NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
24I	Z_l	0.9717(71)	0.9627(36)	0.977(26)	0.981(12)	0.941(16)
	Z_h	0.940(18)	0.930(15)	0.911(99)	0.991(18)	0.972(28)
	R_a	0.7506(31)	0.7415(27)	0.7485(71)	0.7658(79)	0.764(16)
32I	Z_l	1.0	1.0	1.0	1.0	1.0
	Z_h	1.0	1.0	1.0	1.0	1.0
	R_a	1.0	1.0	1.0	1.0	1.0
32ID	Z_l	0.943(14)	0.939(10)	0.940(72)	0.960(23)	0.880(26)
	Z_h	0.981(48)	0.948(40)	0.90(20)	1.030(73)	0.990(38)
	R_a	0.6067(74)	0.5985(82)	0.6015(98)	0.6130(69)	0.595(13)
32I-fine	Z_l	1.000(30)	1.023(30)	1.028(49)	1.019(33)	1.061(60)
	Z_h	0.943(91)	1.009(43)	0.99(12)	1.042(58)	1.046(53)
	R_a	1.336(14)	1.336(13)	1.355(38)	1.321(17)	1.312(36)
48I	Z_l	0.9717(71)	0.9627(36)	0.977(30)	0.981(12)	0.941(16)
	Z_h	0.940(18)	0.930(15)	0.911(98)	0.991(18)	0.972(28)
	R_a	0.7463(78)	0.7425(53)	0.7485(96)	0.7436(57)	0.726(12)
64I	Z_l	1.0	1.0	1.0	1.0	1.0
	Z_h	1.0	1.0	1.0	1.0	1.0
	R_a	1.0170(88)	1.0141(59)	1.025(14)	1.0012(77)	0.986(12)
32ID-M1	Z_l	0.7259(74)	0.7610(34)	0.740(31)	0.782(26)	0.691(36)
	Z_h	0.67(10)	0.715(52)	0.63(25)	0.886(36)	0.826(74)
	R_a	0.4520(74)	0.4500(43)	0.454(16)	0.4610(60)	0.436(12)
32ID-M2	Z_l	—	1.0125(84)	—	1.018(17)	1.001(22)
	Z_h	—	1.003(20)	—	1.031(35)	1.033(28)
	R_a	—	0.8427(35)	—	0.8609(78)	0.872(14)

Table 4.3: Ratios of lattice spacings (R_a) and light and heavy quark masses (Z_l, Z_h) between each ensemble and the reference 32I ensemble.

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
B_0	4.220(93) GeV	3.914(46) GeV	3.885(80) GeV	4.018(38) GeV	4.041(58) GeV
f_0	0.1144(28) GeV	0.1120(16) GeV	0.1282(83) GeV	0.1275(24) GeV	0.1156(18) GeV
$10^3\hat{L}_0^{(3)}$	—	—	0.27(76)	-1.40(18)	-0.26(47)
$10^3\hat{L}_1^{(3)}$	—	—	-0.14(18)	-0.76(11)	-0.308(80)
$10^3\hat{L}_2^{(3)}$	—	—	-1.21(33)	-0.647(52)	-0.68(26)
$10^3\hat{L}_3^{(3)}$	—	—	-0.48(29)	0.88(15)	-0.04(28)
$10^3\hat{L}_4^{(3)}$	-0.102(59)	-0.044(34)	-0.26(12)	-0.54(11)	-0.190(55)
$10^3\hat{L}_5^{(3)}$	0.934(73)	0.913(32)	0.67(92)	1.01(12)	0.86(10)
$10^3\hat{L}_6^{(3)}$	-0.070(40)	0.018(24)	-0.032(52)	-0.239(46)	-0.117(41)
$10^3\hat{L}_7^{(3)}$	—	—	-0.23(27)	-0.202(81)	-0.13(10)
$10^3\hat{L}_8^{(3)}$	0.639(31)	0.466(11)	0.05(44)	0.469(48)	0.364(62)
$10^6\left(\hat{K}_{17}^{(3)}-\hat{K}_{39}^{(3)}\right)$	—	—	-4(10)	0.25(76)	-0.6(1.3)
$10^6\left(\hat{K}_{18}^{(3)}-\hat{K}_{40}^{(3)}\right)$	—	—	-4.4(5.4)	-2.06(36)	1.7(1.2)
$10^6\hat{K}_{19}^{(3)}$	—	—	1(26)	-2.6(1.3)	-2.5(1.1)
$10^6\hat{K}_{20}^{(3)}$	—	—	-3(11)	-1.40(63)	-1.40(71)
$10^6\hat{K}_{21}^{(3)}$	—	—	-3.4(3.0)	-2.74(75)	0.0(1.4)
$10^6\hat{K}_{22}^{(3)}$	—	—	1.53(65)	2.39(40)	1.11(24)
$10^6\hat{K}_{23}^{(3)}$	—	—	-2(12)	-1.3(1.5)	-1.5(2.5)
$10^6\hat{K}_{25}^{(3)}$	—	—	0.8(9.8)	-0.67(48)	-0.17(68)
$10^6\hat{K}_{26}^{(3)}$	—	—	-4.8(3.0)	-2.59(34)	-0.15(89)
$10^6\hat{K}_{27}^{(3)}$	—	—	0.51(20)	0.81(10)	0.548(79)
$m^{(\Omega)}$	1.6648(35) GeV	1.6668(35) GeV	1.6657(35) GeV	1.667(11) GeV	1.6668(99) GeV
c_{m_l,m_Ω}	3.5(1.1)	2.55(60)	3.0(2.1)	2.3(2.3)	2.61(97)
c_{m_y,m_Ω}	5.28(19)	5.38(17)	5.61(90)	5.13(84)	5.27(27)
c_{m_h,m_Ω}	-0.8(1.9)	0.1(1.2)	-1.5(5.3)	3.1(3.4)	4.5(2.1)

Table 4.4: $SU(3)$ PQChPT LECs fit at the chiral scale $\Lambda_\chi = 770\,\text{MeV}$ in units of the canonical size at a given order in the chiral expansion. The parameters $m^{(\Omega)}$ and c_{m_q,m_Ω} are the constant term and m_q slopes for the (linear) m_Ω ansatz, respectively. The value of B_0 quoted here is unrenormalized.

4.4.3 Histograms

In Figure 4.1 we plot stacked histograms of the deviation of each data point Y_i from the fit prediction Y_i^{fit} in units of the standard deviation of the data σ_{Y_i}:

$$X_i \equiv \frac{Y_i - Y_i^{\text{fit}}}{\sigma_{Y_i}}. \tag{4.9}$$

This can be thought of as the signed square root of the contribution to χ^2 from each data point, where the sign indicates whether the fit is overestimating $(-)$ or underestimating $(+)$ the data. Here we see quite clearly that the NNLO terms are necessary for a reasonable fit when we extend the mass cut up to the kaon scale, while, for the lighter mass cut, we obtain a good fit using either the NLO or NNLO ansatz. We also observe a noticeable difference in spread between the histograms for the NNLO fits with frozen and unfrozen LO LECs.

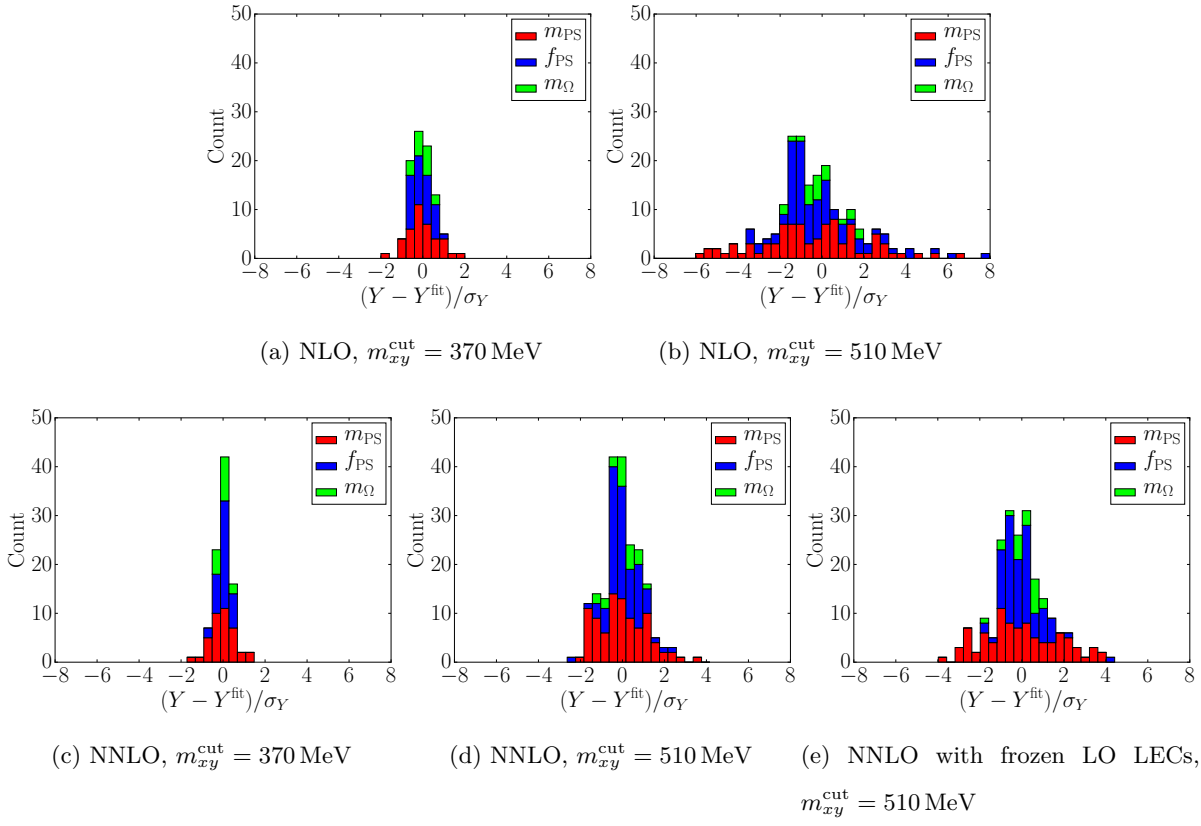

(a) NLO, $m_{xy}^{\text{cut}} = 370\,\text{MeV}$

(b) NLO, $m_{xy}^{\text{cut}} = 510\,\text{MeV}$

(c) NNLO, $m_{xy}^{\text{cut}} = 370\,\text{MeV}$

(d) NNLO, $m_{xy}^{\text{cut}} = 510\,\text{MeV}$

(e) NNLO with frozen LO LECs, $m_{xy}^{\text{cut}} = 510\,\text{MeV}$

Figure 4.1: Stacked histograms of the signed deviation of the data from the fit in units of the standard deviation.

In Figure 4.13 we also plot histograms of the percent deviation between the data and the fit for the fits shown in panels (a) and (e) here. In particular, we observe that the $\mathcal{O}(4\sigma)$ discrepancies observed in panel (e) correspond to $\mathcal{O}(2\%)$ discrepancies when parametrized in terms of the percent difference between data and fit, indicating that even for our worst outliers the NNLO expansion is still accurate at the level of a few percent for the more inclusive fit.

4.4.4 Unitary Chiral Extrapolation

In Figures 4.2 and 4.3 we plot the unitary pseudoscalar mass and decay constant measurements on each ensemble together with the ChPT predictions obtained using the LECs from each fit. We separately plot curves for the light quark mass dependence of the pion and the strange quark mass

dependence of the kaon, where the fit has been used to correct the strange (light) quark mass to its physical value for the pion (kaon) data. In both cases the fit is also used to correct the data to the infinite volume limit. No explicit continuum correction is made. We also plot a dotted horizontal line which corresponds to the PDG value of the kaon mass or decay constant, which we compare to the prediction from the fit (marked "physical point"). The quark masses have been renormalized in the $\overline{\text{MS}}$ scheme at 3 GeV using the renormalization coefficients from Ref. [2]. We find that the $SU(3)$ low energy constants are too poorly determined by the NNLO fit with $m_{xy}^{\text{cut}} = 370\,\text{MeV}$ to extrapolate beyond the range of quark masses directly constrained by lattice data, and so we do not include plots for this fit.

In Figure 4.3 we observe a clear tension between the (unitary) pion decay constant measured on the heaviest 24I and 32I ensembles ($m_\pi \gtrsim 400\,\text{MeV}$) and the NLO $SU(3)$ ansatz. Even in the fit with a more aggressive mass cut, where the values of the LECs have been directly constrained by these heavy points, the curvature of the ChPT formula is simply too large to match the lattice data. The situation is improved at NNLO, but suggests a large, rapidly growing NNLO correction. We conclude that NLO fits are unreliable at this scale.

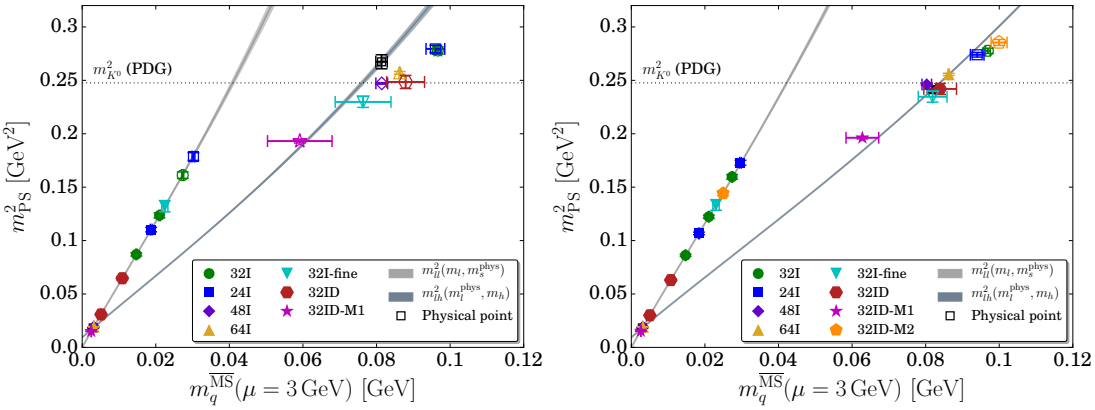

(a) NLO, $m_{xy}^{\rm cut} = 370\,{\rm MeV}$

(b) NLO, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

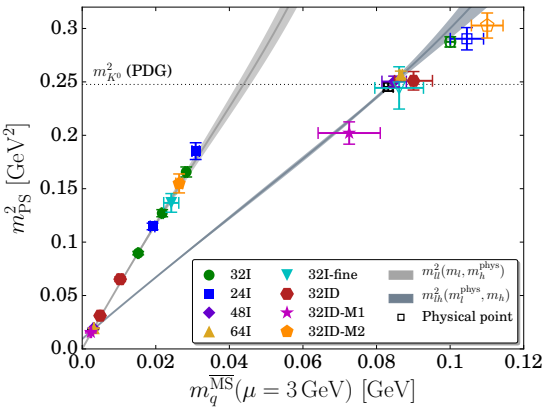

(c) NNLO with frozen LO LECs, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

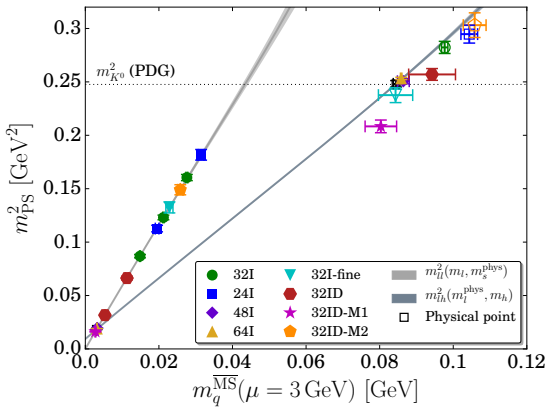

(d) NNLO, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

Figure 4.2: Unitary chiral extrapolation of pseudoscalar meson mass data. The left curve (light gray) shows the light quark mass dependence of m_π^2 with $m_h = m_s^{\rm phys}$ fixed, and the right curve (dark gray) shows the heavy quark mass dependence of m_K^2 with $m_l = m_l^{\rm phys}$ fixed. The fit has been used to correct each data point from the simulated heavy (light) quark mass to the physical heavy (light) quark mass for the pion (kaon), as well as to take the infinite volume limit. Filled symbols correspond to sub-ensembles that were included in the fit, and open symbols correspond to sub-ensembles that were excluded from the fit based on the pseudoscalar mass cut. "Physical point" is the prediction for the physical pion and kaon masses obtained by interpolating the fit to $m_l^{\rm phys}$ and $m_s^{\rm phys}$.

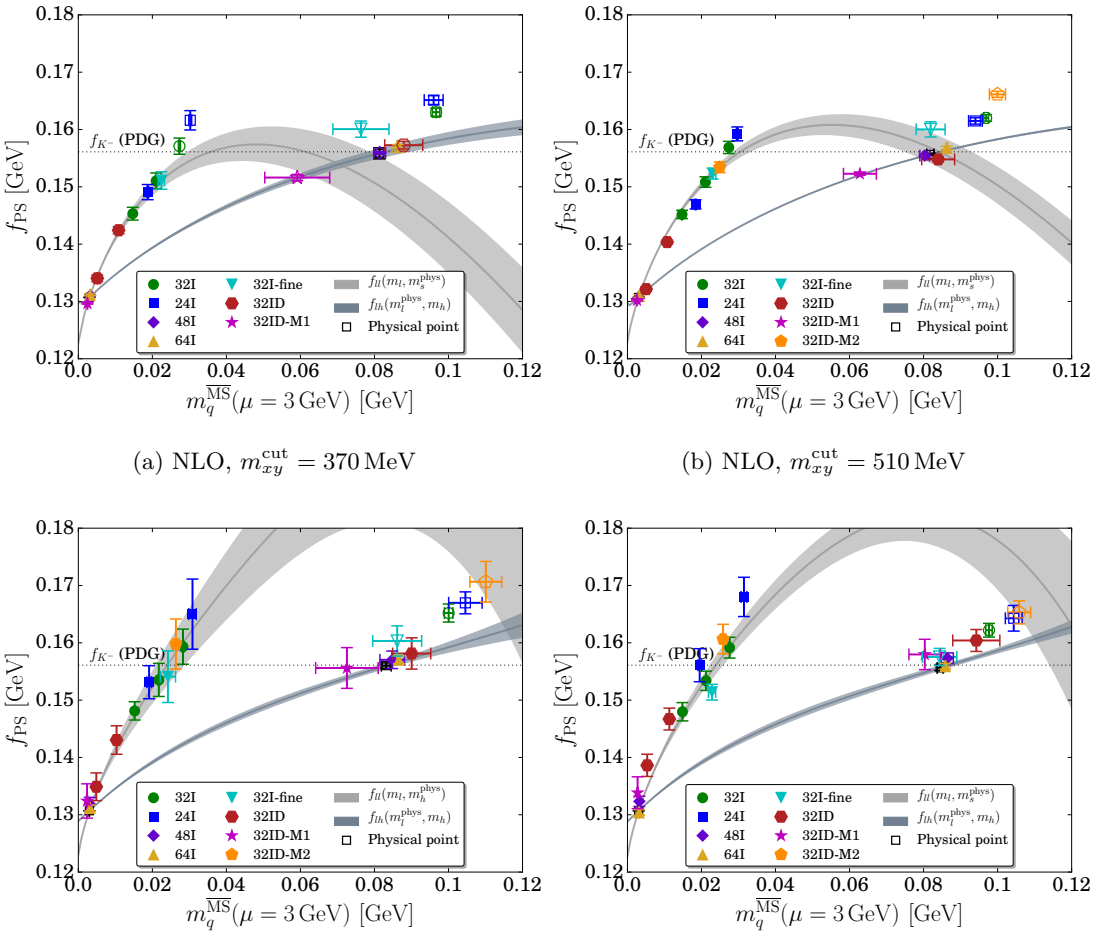

(a) NLO, $m_{xy}^{\rm cut} = 370\,{\rm MeV}$

(b) NLO, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

(c) NNLO with Frozen LO LECs, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

(d) NNLO, $m_{xy}^{\rm cut} = 510\,{\rm MeV}$

Figure 4.3: Unitary chiral extrapolation of pseudoscalar decay constant data. The left curve (light gray) shows the light quark mass dependence of f_π with $m_h = m_s^{\rm phys}$ fixed, and the right curve (dark gray) shows the heavy quark mass dependence of f_K with $m_l = m_l^{\rm phys}$ fixed. The fit has been used to correct each data point from the simulated heavy (light) quark mass to the physical heavy (light) quark mass for the pion (kaon), as well as to take the infinite volume limit. Filled symbols correspond to sub-ensembles that were included in the fit, and open symbols correspond to sub-ensembles that were excluded from the fit based on the pseudoscalar mass cut. "Physical point" is the prediction for the physical pion and kaon decay constants obtained by interpolating the fit to $m_l^{\rm phys}$ and $m_s^{\rm phys}$.

4.4.5 Chiral Expansion

In this section we probe the hierarchy of terms in the $SU(3)$ chiral expansion as a function of the quark masses. This is somewhat more complicated in the $SU(3)$ theory than in the $SU(2)$ theory since, for the $SU(3)$ case, the unitary pion and kaon masses and decay constants depend on both m_l and m_h. To simplify the presentation of our results we consider one-dimensional parametrizations of the chiral expansion rather than two-dimensional plots where m_l and m_h are independent degrees of freedom. For pion quantities we fix $m_h = m_s^{\text{phys}}$ and only plot the light quark mass dependence, and for kaon quantities we parametrize the quark mass dependence in terms of ξ, where $(m_l, m_h) = (\xi m_l^{\text{phys}}, \xi m_s^{\text{phys}})$ ($\xi = 1$ corresponds to the physical kaon). Other limits are also of interest: we can, for example, consider the heavy quark chiral limit $m_h \to 0$, or the $SU(3)$-symmetric limit $m_l = m_h \equiv m$. We note that our data set constrains the light quark direction more strongly than the heavy quark direction: our ensembles cover a large range of dynamical light quark masses — $0.7 \lesssim m_l/m_l^{\text{phys}} \lesssim 10.3$ — but the simulated dynamical heavy quark masses are all near the physical strange quark, with $0.7 \lesssim m_h/m_s^{\text{phys}} \lesssim 1.3$.

In Figure 4.4 we plot the chiral expansion for the unitary pion as a function of the light quark mass, with the strange quark mass fixed at its physical value. We include results from the NLO fit with a 370 MeV cut, and the NNLO fits with a 510 MeV cut, with and without frozen leading-order LECs. The dashed vertial line indicates the heaviest quark mass which corresponds to an ensemble included in the fit. Comparing panels (d) and (f) we see that, while the free NNLO fit (e)-(f) is able to achieve a somewhat better χ^2/dof than the constrained NNLO fit (c)-(d), it accomplishes this by rearranging terms in such a way that the expected hierarchy |LO|> |NLO|> |NNLO|> $\cdots$ for the light-light decay constant is lost. We interpret this as further evidence of unreliability in the free NNLO fit, and focus on the constrained NNLO fit in the remainder of the section.

While the qualitative behavior we observe for the $SU(3)$ expansion is similar to the behavior we observed for the $SU(2)$ theory in Section 5.4 of Ref. [1] — in particular, the NLO and NNLO contributions to m_π^2 enter with similar magnitudes but opposite signs, leaving the total approximately linear in the light quark mass, while both the NLO and NNLO contributions to the pion decay constant are positive and add — the individual terms are larger in the $SU(3)$ case. At m_l^{phys}

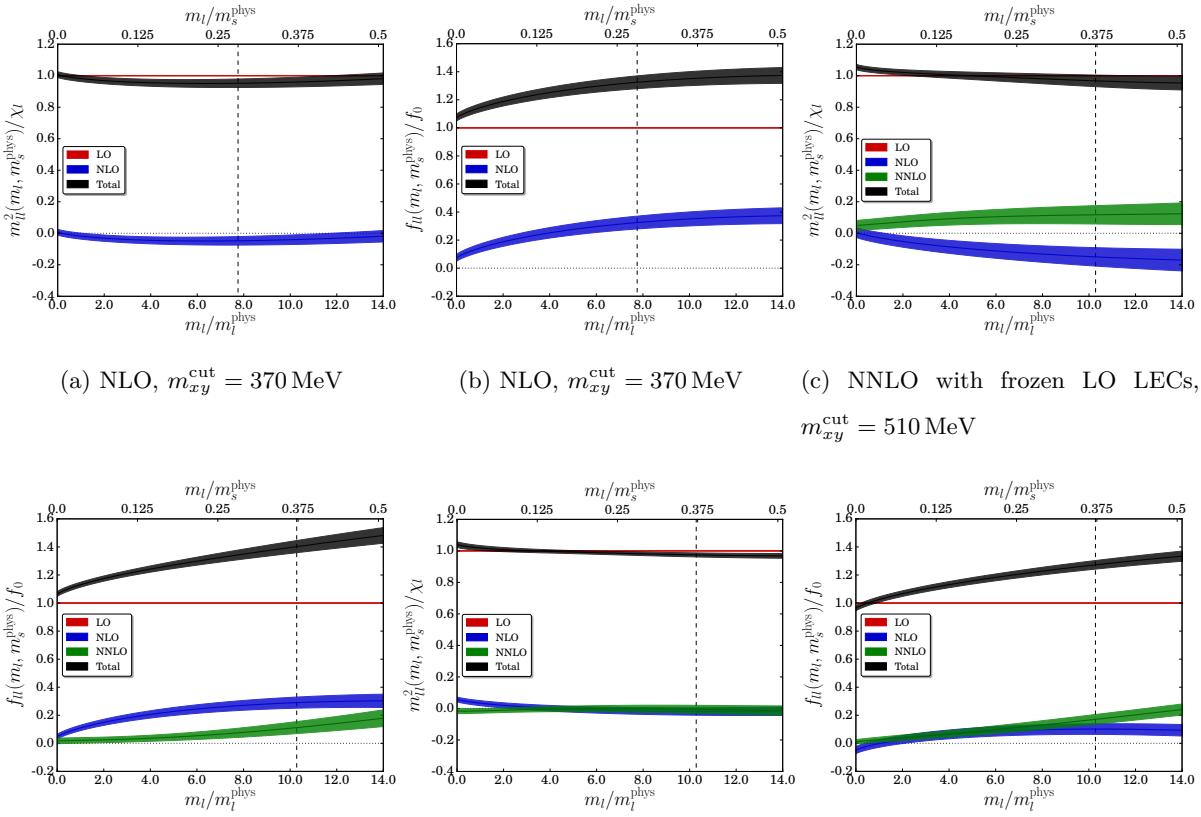

(a) NLO, $m_{xy}^{\rm cut} = 370\,\text{MeV}$

(b) NLO, $m_{xy}^{\rm cut} = 370\,\text{MeV}$

(c) NNLO with frozen LO LECs, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

(d) NNLO with frozen LO LECs, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

(e) NNLO, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

(f) NNLO, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

Figure 4.4: Decomposition of the terms in the $SU(3)$ chiral expansion into LO, NLO, and NNLO terms, normalized by LO, with the heavy (dynamical) quark fixed at the physical strange quark mass. The light-light pseudoscalar mass (left) and decay constant (right) are plotted as a function of the light quark mass, using the LECs obtained from an NLO fit with a pseudoscalar mass cut of 370 MeV (top) and from NNLO fits with a pseudoscalar mass cut of 510 MeV with (middle) and without (bottom) frozen LO LECs. The vertical dashed line corresponds to the heaviest unitary point included in the fit, and the horizontal dotted line marks zero.

we find

$$\begin{aligned} \frac{m_\pi^2}{\chi l} &= 1.000 - 0.029(34) + 0.061(34) \\ \frac{f_\pi}{f_0} &= 1.000 + 0.110(19) + 0.021(20) \end{aligned} \tag{4.10}$$

for the decomposition into LO+NLO+NNLO, normalized by the LO term. As with the $SU(2)$ expansion, the most obvious signal of a breakdown in the series occurs for f_π, where the NLO and NNLO terms become comparable in size for sufficiently heavy m_l with $m_h = m_s^{\text{phys}}$ fixed. We find that NLO $\simeq 0.5|$NNLO$|$ corresponds to $m_l/m_l^{\text{phys}} \approx 9.2$ ($m_\pi \approx 400\,\text{MeV}$) and NLO $\simeq 0.8|$NNLO$|$ corresponds to $m_l/m_l^{\text{phys}} \approx 13.2$ ($m_\pi \approx 500\,\text{MeV}$), indicating that this happens at somewhat lighter pion mass for the $SU(3)$ expansion than for the $SU(2)$ expansion.

Our results for the chiral expansion of the kaon mass and decay constant are plotted in Figure 4.5. At the physical point $\xi = 1$ we find

$$\begin{aligned} \frac{m_K^2}{(\chi_l + \chi_h)/2} &= 1.000 - 0.130(43) + 0.090(41) \\ \frac{f_K}{f_0} &= 1.000 + 0.315(33) + 0.035(30) \end{aligned} \tag{4.11}$$

We note that for m_K^2 the long-observed linearity in the quark masses is realized in our fits as a near-cancellation between a negative NLO term and a positive NNLO term, just as we observe for m_π^2. For f_K we see that the NLO term itself is large, even for $\xi < 1$. We can further decompose the NLO contribution into terms which are analytic in the quark masses and terms which contain chiral logs:

$$\begin{aligned} \frac{f_K}{f_0} = 1 + \underbrace{\frac{24}{f_0^2}\hat{L}_4^{(3)}\overline{\chi} + \frac{4}{f_0^2}\hat{L}_5^{(3)}(\chi_l + \chi_h)}_{\text{analytic}} \\ \underbrace{- \frac{3}{64\pi^2 f_0^2}\left[(\chi_l + \chi_h)\log\left(\frac{\chi_l + \chi_h}{2\Lambda_\chi^2}\right) + \chi_l \log\left(\frac{\chi_l}{\Lambda_\chi^2}\right) + \chi_\eta \log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right)\right]}_{\text{chiral logs}} + \cdots \end{aligned} \tag{4.12}$$

Numerically, at the physical point, we find

$$\frac{f_K}{f_0} = 1.000 + \underbrace{0.071}_{\text{analytic}} + \underbrace{0.244}_{\text{chiral logs}} + \cdots \tag{4.13}$$

For our choice of Λ_χ this large NLO term mostly arises from the chiral logs, but we note that by choosing a different scale Λ'_χ one can change the relative contribution of the analytic and logarithmic terms. The total NLO contribution, however, remains fixed.

In Figure 4.6 we plot three additional limits of the $SU(3)$ expansion: the heavy quark chiral limit ($m_h \to 0$), the degenerate quark mass limit ($m_l = m_h$), and the heavy sea quark mass dependence

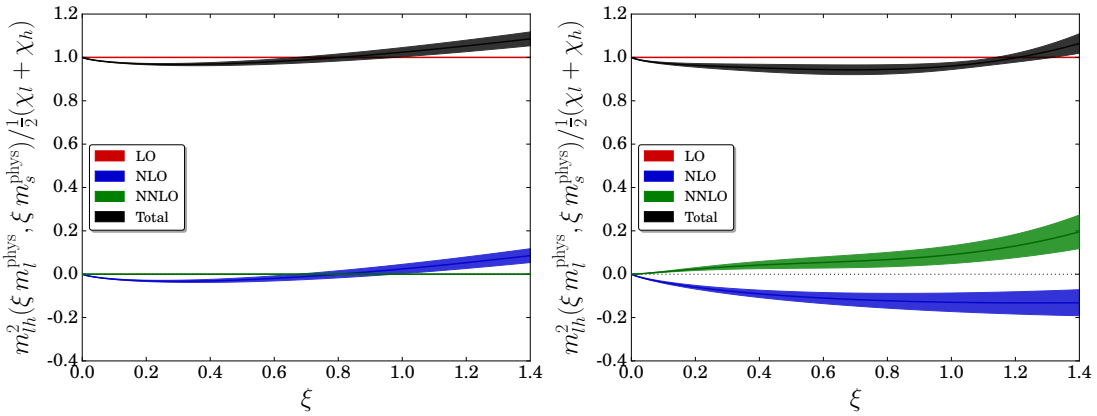

(a) NLO, $m_{xy}^{\rm cut} = 370\,\text{MeV}$ (b) NNLO with frozen LO LECs, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

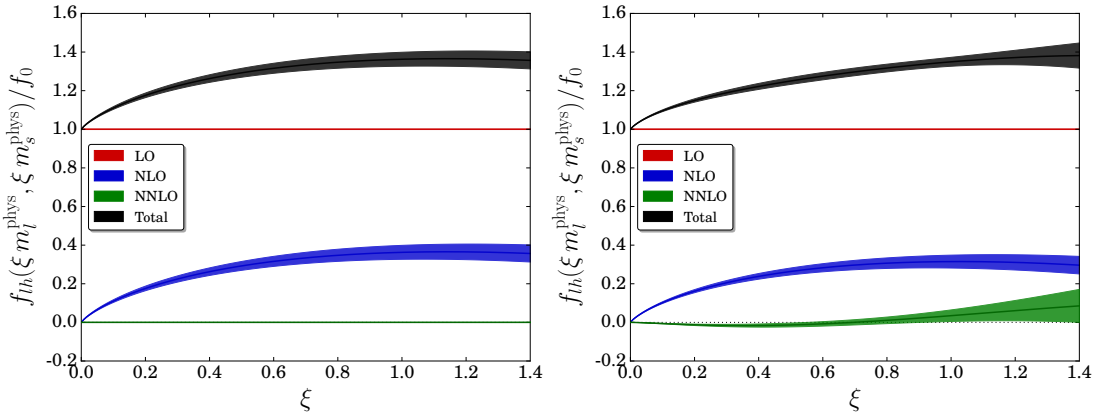

(c) NLO, $m_{xy}^{\rm cut} = 370\,\text{MeV}$ (d) NNLO with frozen LO LECs, $m_{xy}^{\rm cut} = 510\,\text{MeV}$

Figure 4.5: Decomposition of the terms in the $SU(3)$ chiral expansion for the unitary heavy-light mass and decay constant into LO, NLO, and NNLO terms, normalized by LO. The quark mass dependence is parametrized in terms of ξ, where $\xi = 0$ corresponds to the chiral limit and $\xi = 1$ corresponds to the physical kaon, using the LECs obtained from an NLO fit with a pseudoscalar mass cut of 370 MeV (left) and from an NNLO fit with a pseudoscalar mass cut of 510 and frozen LO LECs (right). The horizontal dotted line marks zero.

of the pion mass and decay constant with $m_l \equiv m_l^{\rm phys}$ fixed. These plots were generated using the values of the low energy constants from the NNLO $SU(3)$ fit with a pseudoscalar mass cut of 510 MeV and frozen leading order LECs. The degenerate limit is particularly interesting since the

chiral logarithms are exaggerated. We observe that the expansion has clearly broken down well before the quark mass reaches the scale of the physical strange quark, corresponding to the mass and decay constant of the $s\overline{s}$ state. We also observe that the dependence of the pion mass and decay constant on the heavy quark mass is small. This is to be expected, since pion observables only depend on m_h through K and η loops.

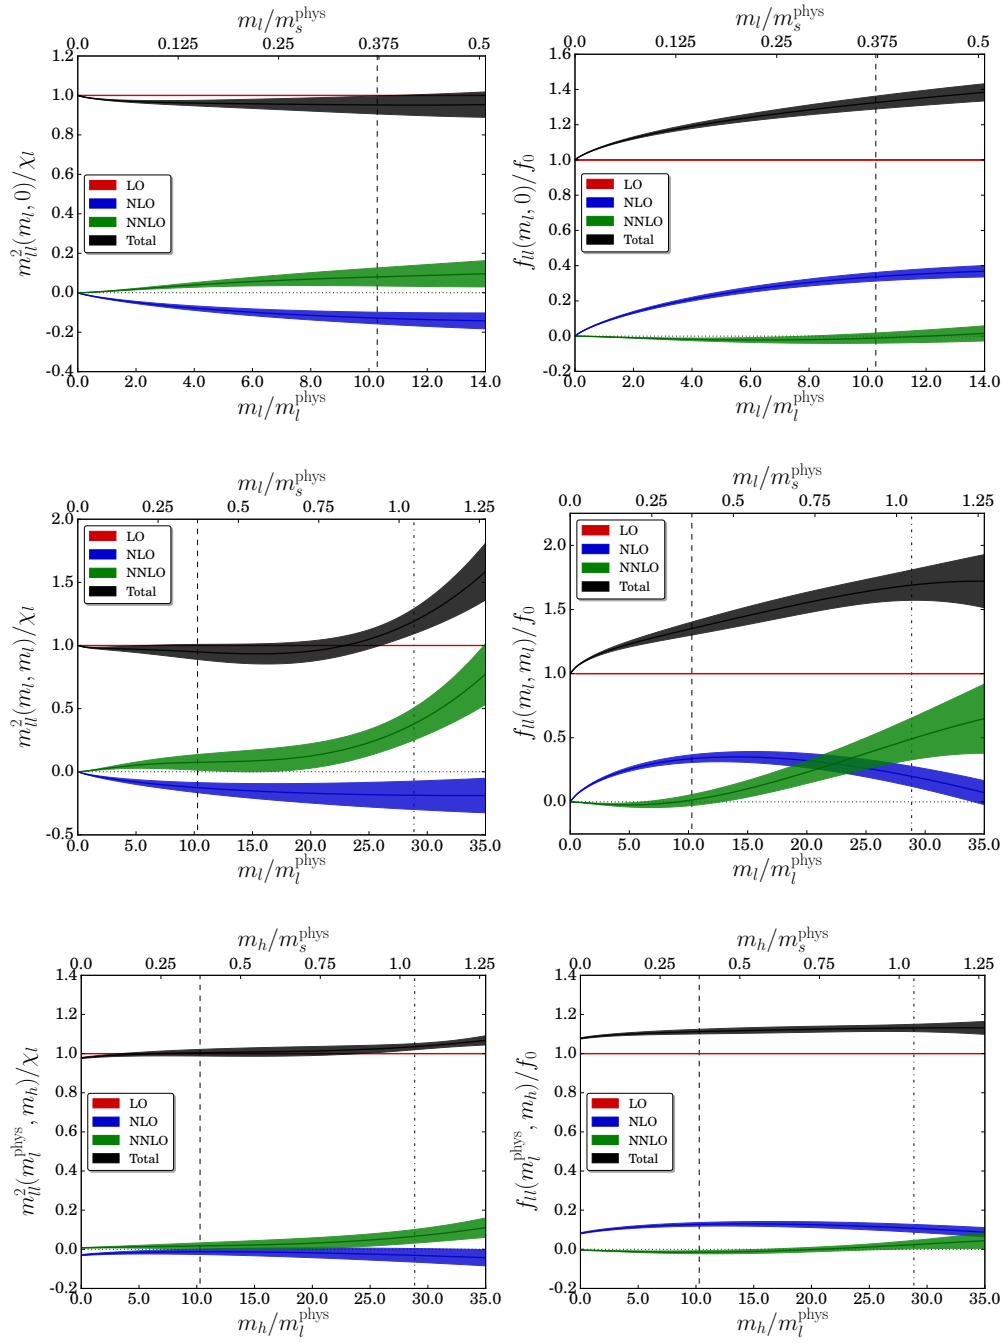

Figure 4.6: Heavy quark chiral limit ($m_h \to 0$, top), degenerate $SU(3)$ limit ($m_l = m_h$, middle), and heavy sea quark mass dependence with $m_l = m_l^{\rm phys}$ fixed (bottom), of the pseudoscalar mass (left) and decay constant (right). The dashed (dash-dotted) vertical line corresponds to the heaviest light (heavy) quark mass constrained by lattice data in the fit.

4.4.6 Comparison with 2007 RBC/UKQCD $SU(3)$ Fits

In Ref. [17] the RBC-UKQCD collaboration studied fits of next-to leading order $SU(3)$ partially quenched ChPT to an earlier version of the 24I ensemble set. While RBC-UKQCD observed that they were able to obtain fits with good χ^2/dof if a cut was placed on the average valence mass — $(m_x+m_y)/2 \leq 0.01$ in lattice units, effectively restricting to partially quenched pseudoscalar masses near the pion-scale of these ensembles — they also observed that their prediction for the chiral decay constant, $f_0 = 93.5\,\text{MeV}$, was very low compared to other lattice and phenomenological predictions. In our current NLO fit with a 370 MeV mass cut, for example, we find $f_0 = 114.4(2.8)\,\text{MeV}$. The fits broke down completely and gave large χ^2/dof when this valence cut was removed and heavy-light data was included. These observations, together with estimates of the size of the NNLO corrections obtained by including NNLO analytic terms, led RBC-UKQCD to conclude that the $SU(3)$ theory was too slowly convergent to obtain reasonable NLO fits to the 24I ensembles, and they have since consistently used $SU(2)$ heavy-meson ChPT to fit the quark mass dependence of heavy-light observables. However, these results were based on a single lattice spacing with heavy input values for the quark masses; even on the lightest 24I ensemble $m_\pi^{\text{unitary}} = 340\,\text{MeV}$ and $m_K^{\text{unitary}} = 593\,\text{MeV}$. In this work we revisit these issues with a much larger data set containing ensembles with several lattice spacings and with quark masses extending down to the physical point.

We have already observed in Table 4.4 that we obtain $f_0 \simeq 115\,\text{MeV}$ from our most recent NLO fits, consistent with other lattice and phenomenological studies of the $SU(3)$ low energy constants. One potential explanation for this discrepancy is that fits which only contain heavy data constrain the NLO expression for the pseudoscalar decay constant in a regime where the chiral logarithms are approximately constant. This makes the fits unreliable, since the logarithms can produce excessive curvature when extrapolated from the heavy region back to the physical point, leading to large systematic uncertainties. To test this hypothesis we have performed two fits: the first ("all ensembles") contains input from our full data set, and determines the quark mass ratios, Z_l and Z_h, and the lattice spacings. For consistency with our earlier work, we include the heaviest 24I and 32I ensembles despite concluding in Section 4.4.4 that NLO $SU(3)$ ChPT fails to describe this

data, and do not include the NLO finite volume corrections. In the second fit ("24I only") we freeze Z_l, Z_h, and the lattice spacings to the values obtained from our "all ensembles" fit, and remove all data except the 24I measurements which were used in Ref. [17]. We note, however, that the "24I only" fit is still not identical to the analysis of Ref. [17] for a number of reasons: since the time of this earlier analysis, RBC-UKQCD has doubled the number of configurations on the 24I ensembles and re-analyzed the spectrum to improve the statistical resolution of this data. In addition, the 24I lattice scale and the physical quark masses are being determined in the "all ensembles" fit by constraining m_π, f_π, and m_Ω to take their experimentally known values at the physical point, where the quark mass dependence of m_{xy} and f_{xy} have been fit to $SU(3)$ PQChPT. In Ref. [17] the 24I lattice scale and the physical quark masses were determined by constraining m_π, m_K, and m_Ω to take their experimentally known values at the physical quark masses in an $SU(2)$ PQChPT chiral fit. The $SU(3)$ chiral fits were deemed unreliable in that work, and were only performed in lattice units without subsequently matching to a continuum scaling trajectory.

The results of this study are summarized in Table 4.5 and Figure 4.7. We observe some tension between the pseudoscalar decay constant measured on the 24I ensemble and the unitary f_π curve in the "all ensembles" fit (panels (b) and (c) of Figure 4.7): this data is systematically $\mathcal{O}(2-3\sigma)$ low on the lighter $m_l^{\text{sea}} = 0.005$ lattice, and $\mathcal{O}(2-4\sigma)$ high on the heavier $m_l^{\text{sea}} = 0.01$ lattice. In this regime the influence of the logarithms is small, and the chiral ansatz is approximately linear in the light quark mass. The 24I data clearly prefers a somewhat steeper slope than the full data set, and, as we observe in the "24I only" fit, this leads to a dramatically lower prediction in the chiral limit, as well as a systematic drift in the values of the leading and next-to leading order LECs. To rule out the possibility that this is entirely a statistical fluke peculiar to the 24I ensembles we have also performed fits where we introduce a lower cut on the pseudoscalar mass — m_{xy}^{min} — and repeat the "all ensembles" fit while varying this lower cut while holding the upper cut fixed. We find that there is a monotonic downward drift in f_0 as the lower cut is increased, and, in particular, we find $f_0 \sim 90$ MeV for a lower cut comparable to the unitary pion mass on the lighter 24I ensemble (339.6(1.2) MeV), consistent with the "24I only" fit. We conclude that extrapolations of very heavy data using the one-loop $SU(3)$ ansatz can be misleading, justifying the cautious approach taken in

Ref. [17]. We note that similar conclusions could be drawn from panels (a) and (b) of Figure 4.3, where tension between the 24I ensembles and our current NLO fits is clearly visible.

Fit	B_0 (GeV)	f_0 (MeV)	$10^3 L_4$	$10^3 L_5$	$10^3 L_6$	$10^3 L_8$
[17]	4.06	93.5(7.3)	0.14(8)	0.87(10)	0.07(6)	0.56(4)
All ensembles	4.05(7)	107.0(1.7)	0.03(3)	0.92(4)	0.03(2)	0.62(2)
24I only	3.96(19)	86.5(4.1)	0.19(3)	0.72(7)	0.11(3)	0.52(3)

Table 4.5: Leading order and next-to leading order low energy constants from the fits discussed in the text, compared to an earlier RBC-UKQCD NLO $SU(3)$ fit. Statistical errors were not explicitly stated for B_0 in physical units in Ref. [17], since the conclusion was that these fits were unreliable.

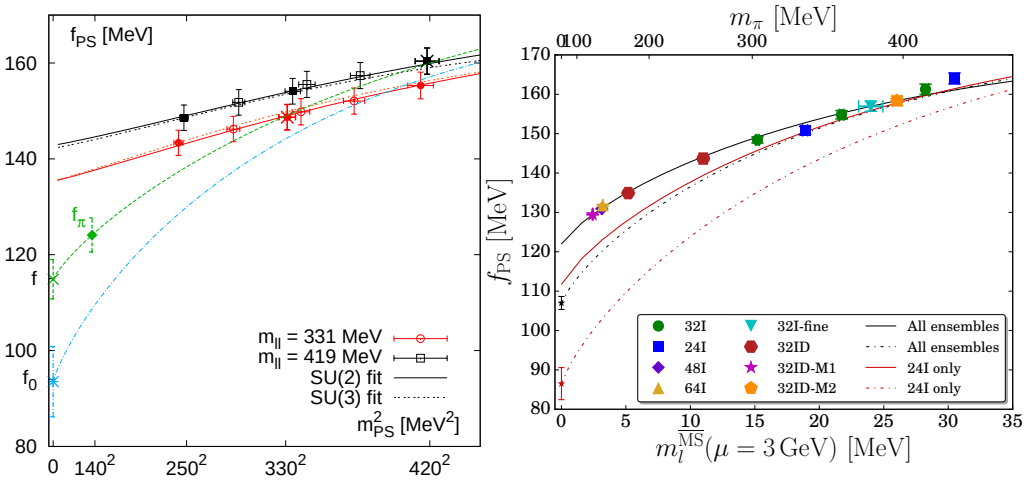

(a) Reprinted from [17]

(b) This work

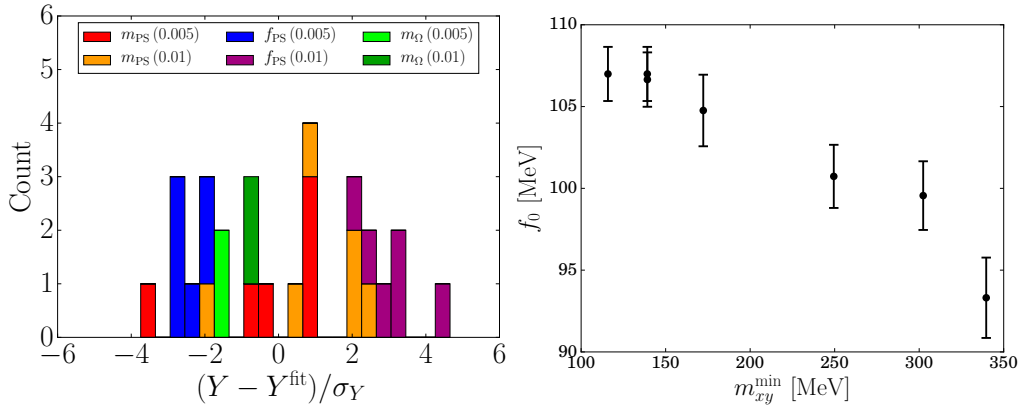

(c) 24I data from the "all ensembles" fit

(d) Dependence of f_0 from "all ensembles" fit on lower mass cut

Figure 4.7: In panel (a) we reprint our summary of $SU(2)$ and $SU(3)$ fits to the pseudoscalar decay constant on the 24I ensemble from an earlier work. Closed (open) [cross] symbols denote measurements with degenerate (nondegenerate) [unitary] quarks. The red and black curves are the partially quenched $SU(2)$ (solid) and $SU(3)$ (dotted) fits to each ensemble, whereas the green and blue curves are the unitary $SU(2)$ extrapolation and the $SU(3)$ extrapolation with three degenerate quarks ($m_l = m_h = m$), respectively. Panel (b) shows two $SU(3)$ fits from this work: the first includes the full data set ("all ensembles"), while the second fit is restricted to the same set of 24I measurements analyzed in the fits from the left panel ("24I only"). In this figure the solid curves show the light quark mass dependence of the light-light pseudoscalar decay constant with $m_h = m_s^{\text{physical}}$ fixed, and the dashed curves show the degenerate $SU(3)$ extrapolation ($m_l = m_h = m$). In panel (c) we show a stacked histogram of the deviation between the "all ensembles" fit and the 24I data in units of the standard deviation of the data. In panel (d) we repeat the "all ensembles" fit, introducing a lower cut on the pseudoscalar mass, m_{xy}^{min}, and plot the dependence of f_0 on m_{xy}^{min}.

4.4.7 Predictions

Unquenched LECs

Table 4.6 summarizes our results for the unquenched $SU(3)$ leading order and next-to leading order low energy constants, computed from the relations in Appendix 4.A.1. We follow the same procedure we used for the $SU(2)$ case in Ref. [1] to compute renormalized values for B_0 and the $SU(3)$ chiral condensate

$$\Sigma_0 = \frac{B_0 f_0^2}{2}. \tag{4.14}$$

B_0 and Σ_0 are renormalized in the $\overline{\text{MS}}$ scheme at $\mu = 2.0\,\text{GeV}$, and include an estimate of the systematic error due to the perturbative matching to $\overline{\text{MS}}$. We obtain this estimate by taking the difference in central value between the value of B_0 or Σ_0 obtained using the RI-SMOM and RI-SMOM$_{\gamma^\mu}$ intermediate schemes; the central value is the value we obtain from the RI-SMOM scheme.

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
$B_0^{\overline{\text{MS}}}(\mu = 2\,\text{GeV})$	2.783(66)(15) GeV	2.581(38)(14) GeV	2.562(60)(14) GeV	2.650(36)(14) GeV	2.665(45)(14) GeV
f_0	114.4(2.8) MeV	112.0(1.6) MeV	128.2(8.3) MeV	127.5(2.4) MeV	115.6(1.8) MeV
$\Sigma_0^{1/3,\overline{\text{MS}}}(\mu = 2\,\text{GeV})$	263.0(5.8)(5) MeV	252.9(3.3)(5) MeV	276.1(10.7)(5) MeV	278.2(3.5)(5) MeV	261.1(3.8)(5) MeV
$10^3 L_1$	—	—	-0.01(24)	-1.46(20)	-0.44(25)
$10^3 L_2$	—	—	-0.94(46)	-2.05(21)	-0.93(69)
$10^3 L_3$	—	—	-1.0(1.8)	3.68(47)	0.5(1.2)
$10^3 L_4$	-0.102(59)	-0.044(34)	-0.26(12)	-0.54(11)	-0.190(55)
$10^3 L_5$	0.934(73)	0.913(32)	0.67(92)	1.01(12)	0.86(10)
$10^3 L_6$	-0.070(40)	0.018(24)	-0.032(52)	-0.239(46)	-0.117(41)
$10^3 L_7$	—	—	-0.23(27)	-0.202(81)	-0.13(10)
$10^3 L_8$	0.639(31)	0.466(11)	0.05(44)	0.469(48)	0.364(62)

Table 4.6: Unquenched $SU(3)$ LECs computed from partially quenched $SU(3)$ fits at the chiral scale $\Lambda_\chi = 770\,\text{MeV}$. Missing entries are not constrained by the fits at a given order. For B_0 and Σ_0 the first error is statistical and the second is a systematic uncertainty in the perturbative matching to $\overline{\text{MS}}$.

In Figures 4.8 and 4.9 we compare our preferred determinations of the leading order and next-to

leading order unquenched $SU(3)$ LECs (blue circles) to other lattice predictions (black squares) and three phenomenological fits (green diamonds). Here "NLO fit" refers to our NLO $SU(3)$ PQChPT fit with a 370 MeV cut and "NNLO fit" refers to our NNLO $SU(3)$ PQChPT fit with a 510 MeV cut and frozen leading order LECs. "Prediction" is the value from Section 4.5 which includes our full error budget. This comparison is different from our treatment of the $SU(2)$ case, where we compared our results to the FLAG averages [23]: since the $SU(3)$ LECs are still relatively poorly determined by lattice calculations, FLAG cites a series of three MILC papers for reference values in lieu of computing a global lattice average. We follow the FLAG nomenclature and refer to these studies as MILC 2009 [24], MILC 2009A [7], and MILC 2010 [8]. All three studies are based on fits of the pseudoscalar masses and decay constants to a series of asqtad improved staggered fermion ensembles, and account for taste-splitting effects using NLO rooted staggered ChPT, but differ in their treatment of NNLO and higher terms and fit constraints. The three phenomenological fits are Gasser and Leutwyler's original determination of the NLO $SU(3)$ LECs from Ref. [4], as well as two updated fits including NNLO terms from Bijnens and Jemos [25] and Bijens and Ecker [26]. The latter two phenonemonological fits rely on experimental and lattice input from many different physical processes, including masses and decay constants of the π, K, and η, values and slopes of the scalar pion charge radius and $K_{\ell 4}$ form factors, $\pi\pi$ and πK scattering lengths, the quark mass ratio $m_s/\hat{m}$, and the known values of the $SU(2)$ LECs.

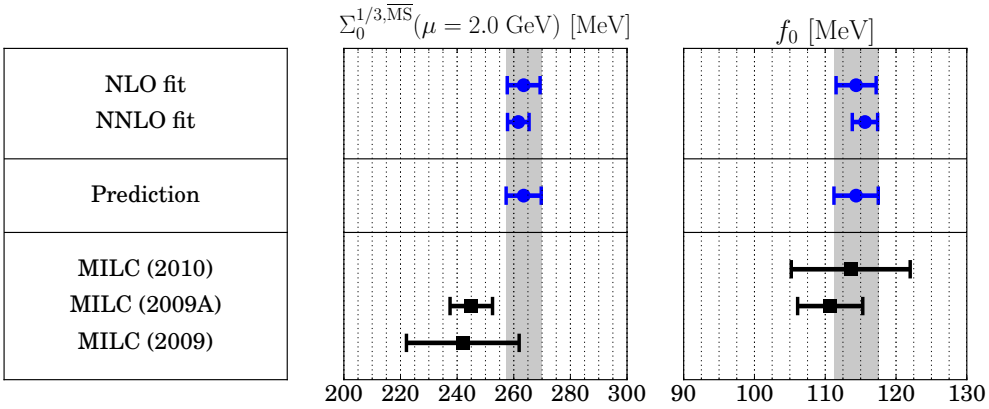

Figure 4.8: Leading order $SU(3)$ ChPT LECs from this work compared to other lattice results [7, 8, 24].

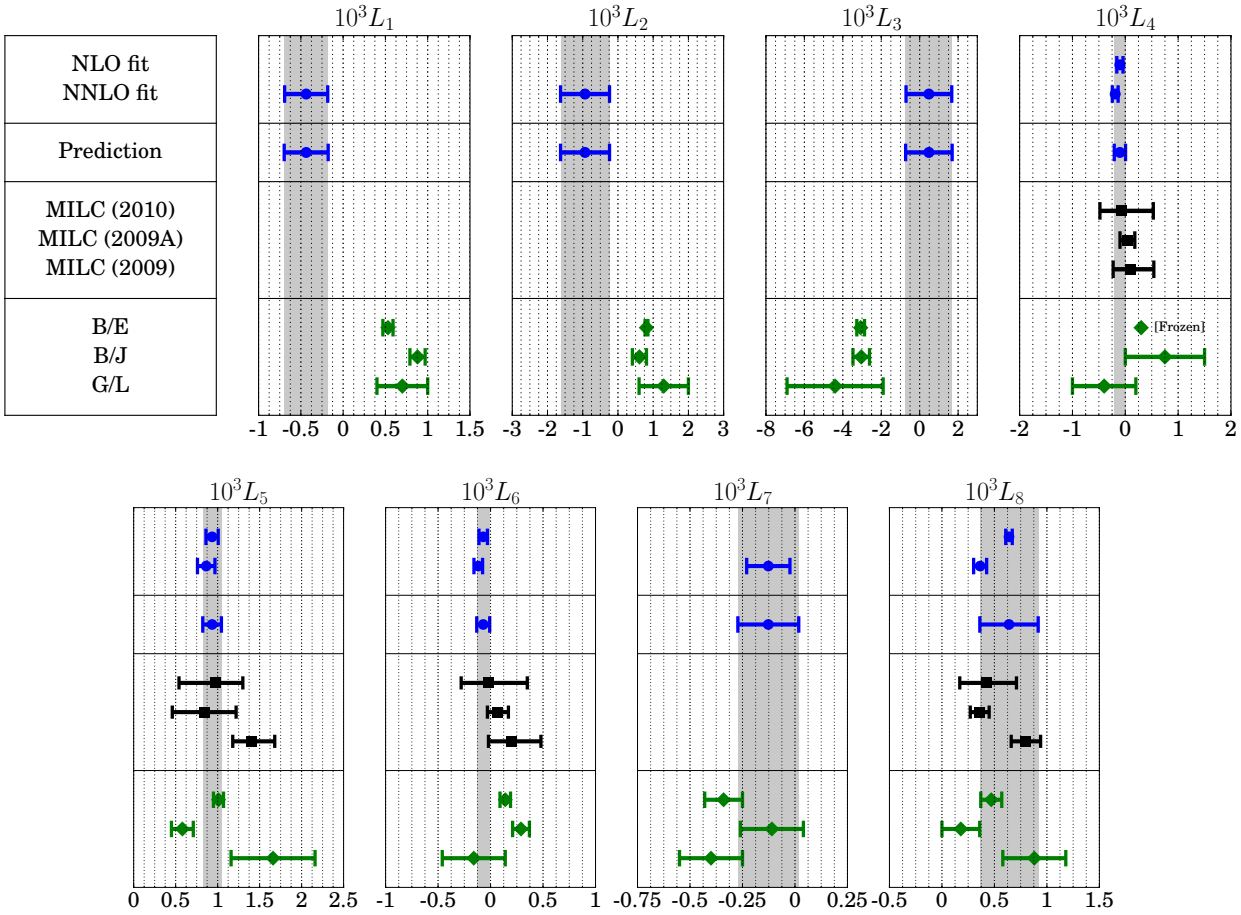

Figure 4.9: Next-to leading order $SU(3)$ ChPT LECs compared to other lattice [7, 8, 24] and phenomenological [4, 25, 26] determinations. The fit by Bijnens and Ecker [26] applies $L_4 \equiv 0.3$ as a constraint.

While our values for the $SU(3)$ low energy constants are generally consistent with other lattice and phenomenological fits, we observe clear tension in two places: the chiral condensate (Σ_0) and three of the unquenched NLO LECs which are only determined in the NNLO fits (L_1, L_2, and L_3). Comparing our results for the chiral condensate to those of the MILC collaboration, we observe that our fits prefer values which are $\sim 6-8\%$ larger than the MILC values; it is unclear if this discrepancy is purely a tension between the fits or if it is also associated with differences in the renormalization procedure. We also observe some tension between our values of L_1, L_2, and L_3 and the three phenomenological determinations. These LECs only enter into the pseudoscalar

mass and decay constant at two-loop order, and thus we expect that they are weakly constrained in our fits. As a result, one might expect that these LECs are relatively free to vary without changing the overall fit quality very much. The phenomenonological fits, in contrast, include $\pi\pi$ and πK scattering data, among many other observables, which strongly constrains these same LECs beginning at one-loop order. The tension we observe may suggest that more data covering a larger range of quark masses or additional quantities beyond the pseudoscalar mass and decay constant are necessary if one wants to reliably extract all of the NLO $SU(3)$ LECs[2]. It is interesting to note, however, that we do not observe such tension in L_7, which also only enters into our fits at two-loop order. Finally, we also observe some weaker systematic tension in the values for L_8: it appears that the fits which only contain chiral logarithms up to one-loop order (ours, MILC 2009, and Gasser/Leutwyler) consistently prefer a somewhat higher value for L_8 than the fits to the full two-loop expressions (ours, MILC 2009A and 2010, Bijnens/Jemos, and Bijnens/Ecker). Given the relatively slow convergence of the $SU(3)$ series at the physical strange quark mass, the $\mathcal{O}(1\sigma)$ sensitivity of the LECs to fit systematics and constraints is not terribly surprising.

By integrating out the strange quark in the $SU(3)$ theory one can also write down explicit relations between the $SU(2)$ and $SU(3)$ LECs, which we collect in Appendix 4.A.3 and use to predict the unquenched $SU(2)$ LECs from our $SU(3)$ fits. The values we obtain for the leading and next-to-leading order $SU(2)$ LECs are summarized in Table 4.7, and are plotted in Figures 4.10 and 4.11 alongside the final results, including our full error budget summed in quadrature, from our recent direct $SU(2)$ fits [1]. We also compare our predictions for the $SU(2)$ LECs to the 2013 $N_f = 2 + 1$ FLAG lattice averages [7, 18, 19, 23, 27–29] and two phenomenological fits: the first is Gasser and Leutwyler's original determination of the $SU(2)$ LECs in Ref. [30], and the second

[2]We observed similar behavior in our NNLO $SU(2)$ fits [1]: the values of the NLO $SU(2)$ LECs $\bar{\ell}_1$ and $\bar{\ell}_2$ — which are the $SU(2)$ analogues of L_1, L_2, and L_3 — disagreed outside statistical errors between an NNLO fit with a mass cut of 370 MeV and an NNLO fit with a mass cut of 450 MeV, while the other LO and NLO LECs, which one would expect to be more strongly constrained, were all consistent. The values of $\bar{\ell}_1$ and $\bar{\ell}_2$ obtained from the fit with the 450 MeV cut were consistent with phenomenological results, suggesting that, at least in the context of our fits, relatively heavy data beyond the range which can be described accurately by NLO ChPT was necessary to determine these LECs accurately.

is Colangelo et al.'s updated fit of experimental pion scattering and scalar charge radius data to NNLO $SU(2)$ ChPT and the Roy equations [31]. We do not attempt to compute the NLO $SU(2)$ LECs $\bar{\ell}_1$ and $\bar{\ell}_2$, since these are related to the $SU(3)$ LECs L_1, L_2, and L_3, which we have argued may not be reliably determined in our fits.

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
$B^{\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	2.804(33)(15) GeV	2.792(28)(15) GeV	2.86(15)(2) GeV	2.801(59)(15) GeV	2.679(86)(15) GeV
f	122.7(5) MeV	122.8(3) MeV	131(12) MeV	121.2(2.2) MeV	121.0(2.2) MeV
$\Sigma^{1/3,\,\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	276.4(1.3)(5) MeV	276.1(1.0)(5) MeV	291.0(13.0)(5) MeV	274.0(2.2)(5) MeV	269.6(4.5)(5) MeV
$\bar{\ell}_3$	2.85(18)	3.07(11)	3.95(53)	3.37(18)	4.09(36)
$\bar{\ell}_4$	3.908(53)	3.986(26)	3.55(43)	3.397(66)	3.76(13)
$10^3 l_7$	—	—	10.3(9.3)	3.7(2.7)	2.5(3.5)

Table 4.7: Unquenched $SU(2)$ LECs computed from partially quenched $SU(3)$ fits and one-loop relations. Missing entries are not constrained by the fits at a given order. For B and Σ the first error is statistical and the second is a systematic uncertainty associated with the perturbative matching to $\overline{\rm MS}$.

We generally observe excellent consistency between the direct $SU(2)$ fits and the converted $SU(3)$ fits, which provides a further check on the sensibility of our results. In particular, we note that there is excellent agreement between our values for the $SU(2)$ chiral condensate Σ computed from the $SU(3)$ fits, our prediction for Σ from direct $SU(2)$ fits reported in Ref. [1], and the FLAG lattice average. We also note that the values for the $SU(2)$ chiral decay constant f we compute from our $SU(3)$ fits, including the free NNLO fits, are consistent within error, even though we observe a substantial difference between the $SU(3)$ chiral decay constants f_0. Comparing the entries in Table 4.4 with Equation (D.19), we note that the shifts in f_0 are compensated by shifts in the NLO LEC L_4 in such a way that the value of f remains consistent between the free and frozen fits. We do not attempt to quantify the systematic error associated with neglecting the two-loop contributions to the expressions which relate the $SU(3)$ LECs to the $SU(2)$ LECs, which could very well be large at the physical strange quark mass. For this reason we prefer values of the $SU(2)$ LECs from direct $SU(2)$ fits to values converted from the $SU(3)$ fits, even in cases where the latter superficially have

similar errors (e.g. $\overline{\ell}_4$ or l_7).

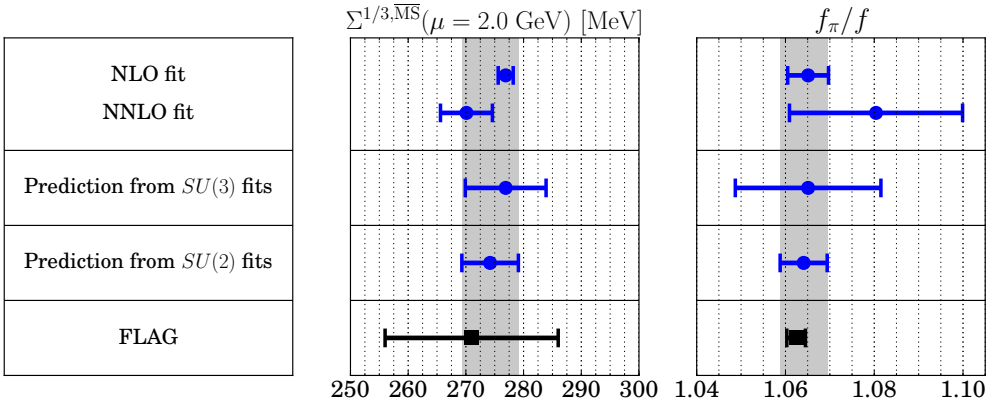

Figure 4.10: Leading order $SU(2)$ ChPT LECs computed from the $SU(3)$ fit results and compared to the 2013 FLAG lattice averages.

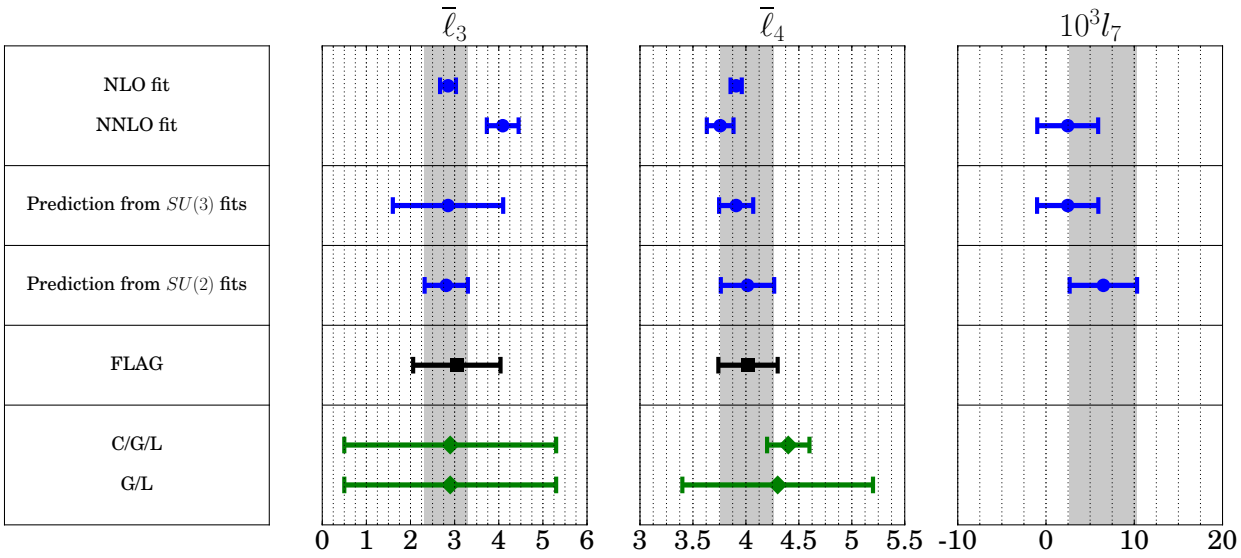

Figure 4.11: Next-to leading order $SU(2)$ ChPT LECs computed from the $SU(3)$ fit results and compared to the 2013 FLAG lattice averages and two phenomenological determinations [30, 31].

Zweig Rule Breaking

Large-N_c arguments suggest that the terms which cause the leading order $SU(2)$ and $SU(3)$ LECs to deviate are suppressed by powers of $1/N_c$; in the limit $N_c \to \infty$ one has exact equality $B = B_0$,

$f = f_0$, and $\Sigma = \Sigma_0$ (Zweig rule). Computing the ratios B/B_0, f/f_0, and Σ/Σ_0 from fits to lattice data provides an interesting first-principles test of the large-N_c approximation to QCD, since the sizes of the deviations from unity are a quantitative test of the validity of this approximation. We compute these ratios two ways: first by directly computing ratios of LECs from two of our NLO fits — the $SU(2)$ fit with $m_\pi^{\rm cut} = 370\,{\rm MeV}$, and the $SU(3)$ fit with $m_{xy}^{\rm cut} = 370\,{\rm MeV}$ — and from two of our NNLO fits — the $SU(2)$ fit with $m_\pi^{\rm cut} = 450\,{\rm MeV}$, and the $SU(3)$ fit with $m_{xy}^{\rm cut} = 510\,{\rm MeV}$ and frozen leading order LECs — under the superjackknife (blue circles), where the $SU(2)$ fits are from Ref. [1]. The NLO fits we use are completely self-consistent in the sense that the chiral fits are constrained by exactly the same data for the pseudoscalar masses and decay constants, whereas, for the NNLO case, the $SU(3)$ fit includes some kaon-scale data containing a heavy quark that is not included in the $SU(2)$ fit. In the second approach we first compute the leading order $SU(2)$ LECs from the $SU(3)$ LECs and then form the same ratios with these converted $SU(2)$ LECs in the numerator (red circles). We find that both methods give consistent values for the ratios, and in Figure 4.12 we compare our results to the MILC 2009 [24] and MILC 2009A [7] studies (black squares).

While our results for the ratio f/f_0 are consistent with the MILC results, our larger value for B_0 leads to an $\mathcal{O}(1-2\sigma)$ tension in the ratios B/B_0 and Σ/Σ_0, implying smaller violations of the Zweig rule.

Other Physical Predictions

Table 4.8 summarizes a number of predictions for physical quantities we make based on our $SU(3)$ fits. We predict m_K, f_K, and the ratio f_K/f_0 directly from the global fit by interpolating/extrapolating to the physical light and strange quark masses in the infinite volume, continuum limit. We also compute the next-to leading order QCD isospin breaking effects in the kaon system and the $I = 1/2$ ($a_0^{1/2}$) and $I = 3/2$ ($a_0^{3/2}$) πK scattering lengths from the relations in Appendix 4.A.4.

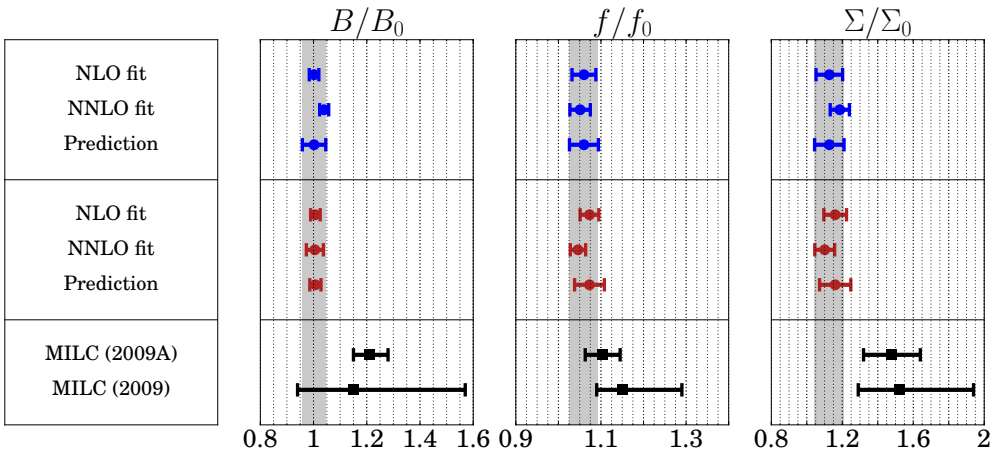

Figure 4.12: Ratios of the leading order $SU(2)$ and $SU(3)$ low energy constants from this work compared to those from the MILC studies [7, 24]. The first three rows (blue circles) are computed by taking ratios between $SU(3)$ LECs from this work and LECs from direct $SU(2)$ fits in Ref. [1]. The second three rows (red circles) are computed by taking ratios between $SU(3)$ LECs from this work and $SU(2)$ LECs obtained from the $SU(3)$ fits and the one-loop conversion formulae in Appendix 4.A.3.

			Free		**Frozen LO LECs**
	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
m_K	0.5171(64) GeV	0.4913(29) GeV	0.479(70) GeV	0.4982(30) GeV	0.4952(41)
f_K	0.15584(97) GeV	0.15566(20) GeV	0.160(42) GeV	0.15562(47) GeV	0.15601(49) GeV
f_K/f_0	1.363(36)	1.390(20)	1.25(39)	1.221(22)	1.349(22)
$[m^2_{K^0} - m^2_{K^\pm}]_{\rm QCD}/\Delta m_{du}$	5.44(24) GeV	3.658(62) GeV	1.75(93) GeV	3.46(28) GeV	2.74(39) GeV
$[\frac{f_{K^0}}{f_{K^\pm}} - 1]_{\rm QCD}/\Delta m_{du}$	3.01(13) GeV^{-1}	3.068(32) GeV^{-1}	1.9(1.9) GeV^{-1}	2.48(19) GeV^{-1}	2.72(27) GeV^{-1}
$m_\pi a_0^{1/2}$	—	—	0.124(18)	0.1435(56)	0.1376(92)
$m_\pi a_0^{3/2}$	—	—	-0.067(14)	-0.0781(47)	-0.0671(84)

Table 4.8: Predictions from NLO and NNLO fits and $SU(3)$ ChPT. $\Delta m_{du} \equiv m_d - m_u$.

The predictions for m_K and f_K are most interesting for the lighter mass cut fits, since the fits with $m_{xy}^{\rm cut} = 510\,\text{MeV}$ contain direct lattice measurements of m_K and f_K on the physical point ensembles. For these fits we are essentially performing a small interpolation to the physical kaon, and we expect that any smooth fit ansatz which matches the lattice data reasonably well in this regime would also accurately predict m_K and f_K. The lighter mass cut fits offer a more

interesting test of $SU(3)$ PQChPT: the LECs are determined entirely by partially quenched data with two light quarks, and m_K and f_K are true predictions obtained by extrapolating the fit up to the physical kaon. While the predictions from the NNLO fit are consistent with the experimental values $m_K^{\text{phys}} = 495.65\,\text{MeV}$ and $f_K^{\text{phys}} = 156.1\,\text{MeV}$, these predictions also have very large statistical uncertainties. The predictions from the NLO fit are consistent with the experimental kaon mass and decay constant to 4(1)% and 0.2(6)%, respectively.

In the remainder of this section we focus on the fit with $m_{xy}^{\text{cut}} = 370\,\text{MeV}$ as our preferred NLO fit, and the fit with $m_{xy}^{\text{cut}} = 510\,\text{MeV}$ and frozen LO LECs as our preferred NNLO fit. We compute the isospin breaking corrections to the kaon masses and decay constants using Equations (D.21) and (D.22). At one-loop these depend only on the LECs L_4, L_5, L_6, and L_8, which are determined in both our NLO and NNLO fits. We observe that the predictions for the decay constant splittings are consistent between these two fits, but the predictions for the mass splittings differ by a factor of two. This discrepancy seems to arise from the difference in L_8, which enters into Equation (D.21) as a term $\propto L_8 m_K$. Assigning our full systematic error budget to these predictions, as described in Section 4.5, we find $[m_{K^0}^2 - m_{K^\pm}^2]_{\text{QCD}}/\Delta m_{du} = 5.4(0.2)(2.7)\,\text{GeV}$ and $[f_{K^0}/f_{K^\pm} - 1]_{\text{QCD}}/\Delta m_{du} = 3.0(0.1)(0.3)\,\text{GeV}^{-1}$, where the first error is statistical and the second is systematic. The RM123 collaboration has performed a direct calculation of the $\mathcal{O}(\Delta m_{du})$ QCD isospin breaking effects in a number of low-energy observables, including the kaon mass and decay constant [32]. They find $[m_{K^0}^2 - m_{K^\pm}^2]_{\text{QCD}}/\Delta m_{du} = 2.57(8)\,\text{GeV}$ and $[f_{K^0}/f_{K^\pm} - 1]_{\text{QCD}}/\Delta m_{du} = 3.3(3)\,\text{GeV}^{-1}$, in good agreement with our ChPT predictions.

Our final prediction is for the $I = 1/2$ ($a_0^{1/2}$) and $I = 3/2$ ($a_0^{3/2}$) πK scattering lengths. While the individual scattering lengths are not known experimentally, the DIRAC collaboration has recently measured the isospin-odd linear combination $a_0^- = (a_0^{1/2} - a_0^{3/2})/3$ and found $m_\pi |a_0^-| = 0.11(^{+9}_{-4})$. Computing the correlated difference from the results in table 4.8 we find $m_\pi a_0^- = 0.068(4)$. We note that a direct calculation of the πK scattering lengths has recently been performed on the 48I and 64I ensembles, and found $m_\pi a_0^{1/2} = 0.16(3)$ and $m_\pi a_0^{3/2} = -0.07(2)$ after extrapolating to the continuum limit [33].

4.5 Error Budget and Final Results for the Unquenched $SU(3)$ LECs

In this section we discuss the error budget for our determination of the leading and next-to leading order unquenched $SU(2)$ and $SU(3)$ low energy constants, and report our final values including all systematics. We assign the following error to each LEC in table 4.9:

- *Influence of heavy data as determined by underweighting correlated data in the fits*: While our global fits are uncorrelated, we know that the partially quenched measurements on a given ensemble are highly correlated since they are computed with the same set of field configurations. If we were fitting to a function which exactly represented our data, as opposed to an expansion with some limited precision, our uncorrelated fits would not introduce any systematic bias into our answers. Since this is not the case, changing the weighting of the heavy mass ensembles, which contain highly correlated partially quenched measurements, gives us an estimate of the systematic effects on our results due to the worsening systematic disagreement betweeen PQChPT and QCD at heavier quark masses. We estimate the impact on our fits by taking the difference in central value between the LECs of an unweighted, uncorrelated fit (Section 4.4) and the LECs of a fit where the contributions to χ^2 from ensembles with multiple partially quenched measurements have been systematically underweighted to capture the dominant effects of correlations (Appendix 4.B).

We also assign additional errors to the LECs which are determined by both NLO and NNLO fits (B_0, f_0, Σ_0, L_4, L_5, L_6, and L_8):

- *Influence of mass cut and truncation of the (continuum) chiral expansion*: We estimate the sensitivity of the LECs to varying the mass cut and to the truncation of N^{3}LO and higher order terms by taking the difference in central value between the NLO fit with a 370 MeV mass cut and the NNLO fit with a 510 MeV mass cut and frozen leading order LECs. We note that this treatment is different from the error budget for our $SU(2)$ fits [1], where we computed two independent systematic error estimates. In the $SU(3)$ case we have argued that

the 510 MeV cut NLO fit and the 370 MeV cut NNLO fit are both likely to be unreliable, and so we choose to estimate these systematics together using our preferred fits.

- *Finite volume effects*: As a conservative bound on the influence of NNLO and higher order FV corrections, as well as neglected cross terms — *e.g.* (NLO continuum ChPT) × (NLO FV correction) — we compute the difference in central value between an NLO PQChPT fit with NLO FV corrections and an NLO PQChPT fit with no FV corrections, both with a pseudoscalar mass cut of 370 MeV.

Finally, for B_0 and Σ_0, which are renormalized in the $\overline{\text{MS}}$ scheme, we include an additional systematic:

- *Renormalization*: We renormalize quantities in $\overline{\text{MS}}$ by first renormalizing in either the RI-SMOM or RI-SMOM$_{\gamma^\mu}$ scheme, and then perturbatively match to $\overline{\text{MS}}$ at a scale $\mu = 2\,\text{GeV}$ where perturbation theory is known to be reliable. We estimate the systematic error associated with this procedure by taking the difference between the results we obtain from the RI-SMOM and RI-SMOM$_{\gamma^\mu}$ intermediate schemes. The central value we report is from the RI-SMOM scheme.

For the LECs which first enter into the $SU(3)$ ChPT expressions for the psuedoscalar mass and decay constant at two-loop order — L_1, L_2, L_3, and L_7 — we do not attempt to quantify any systematics other than the first since these LECs typically have $\mathcal{O}(50\%)$ or larger statistical errors, and are more appropriately considered bounds than high-precision determinations. Likewise, we do not attempt to quantify systematic errors for the partially quenched LECs (Section 4.4.2) or for our predictions of the πK scattering lengths (Section 4.4.7), but we could, in principle, follow the same procedure to assign our full error budget to these quantities.

$\boldsymbol{B_0^{\overline{\rm MS}}(\mu = 2\,\mathrm{GeV})}$	$\boldsymbol{2.80(7)(13)\,\mathrm{GeV}}$
$\boldsymbol{f_0}$	$\boldsymbol{114.4(2.8)(1.3)\,\mathrm{MeV}}$
$\boldsymbol{\Sigma_0^{\overline{\rm MS}}(\mu = 2\,\mathrm{GeV})}$	$\boldsymbol{263.5(5.8)(2.4)\,\mathrm{MeV}}$
$10^3 L_1$	$-0.44(25)(5)$
$10^3 L_2$	$-0.93(69)(4)$
$10^3 L_3$	$0.5(1.2)(0.2)$
$\boldsymbol{10^3 L_4}$	$\boldsymbol{-0.102(59)(89)}$
$\boldsymbol{10^3 L_5}$	$\boldsymbol{0.934(73)(83)}$
$\boldsymbol{10^3 L_6}$	$\boldsymbol{-0.070(40)(47)}$
$10^3 L_7$	$-0.13(10)(1)$
$\boldsymbol{10^3 L_8}$	$\boldsymbol{0.64(3)(28)}$

Table 4.9: Final predictions for the unquenched $SU(3)$ LECs including all statistical and systematic errors. The reported errors are the statistical (left) and the total systematic (right) obtained by summing the contributions we discuss in the text in quadrature. Bold entries correspond to LECs which enter into both NLO and NNLO fits, for which we assign the full error budget; for the other entries the mass cut, chiral truncation, and finite volume systematics are assumed to be negligible compared to the statistical error and are not quantified. The central values and statistical errors of B_0, f_0, $\Sigma_0^{1/3}$, L_4, L_5, L_6, and L_8 are from an NLO fit with a 370 MeV cut, while the central values and statistical errors of L_1, L_2, L_3, and L_7 are from an NNLO fit with a 510 MeV cut and frozen leading order LECs. The $\{L_i\}$ are quoted at the chiral scale $\Lambda_\chi = 770\,\mathrm{MeV}$.

4.6 Conclusions

In this work we have performed fits of pseudoscalar masses and decay constants from a series of RBC-UKQCD domain wall fermion ensembles to next-to leading and next-to-next-to leading order $SU(3)$ partially quenched chiral perturbation theory. We reported values for a large set of partially quenched low energy constants, and used these values to compute the unquenched leading and next-to leading order LECs. We also examined the range of quark masses for which NLO and NNLO ChPT accurately describe our lattice data, and used the newly determined LECs from NNLO fits to make one-loop predictions for the size of isospin breaking effects in the kaon sector and for the πK scattering lengths, which we compare to other lattice and experimental results. We have observed that both NLO and NNLO $SU(3)$ PQChPT can accurately describe partially quenched lattice data containing two light valence quarks, while the NNLO terms are necessary to describe data with pseudoscalar masses extending up to the scale of the physical kaon. To emphasize this point we plot histograms of the percent deviation between the data and fit

$$\Delta \equiv \frac{(Y - Y^{\text{fit}})}{(Y + Y^{\text{fit}})/2} \times 100 \tag{4.15}$$

for our preferred fits, NLO PQChPT with a 370 MeV pseudoscalar mass cut and NNLO PQChPT with a 510 MeV pseudsocalar mass cuct and frozen leading order LECs, in Figure 4.13.

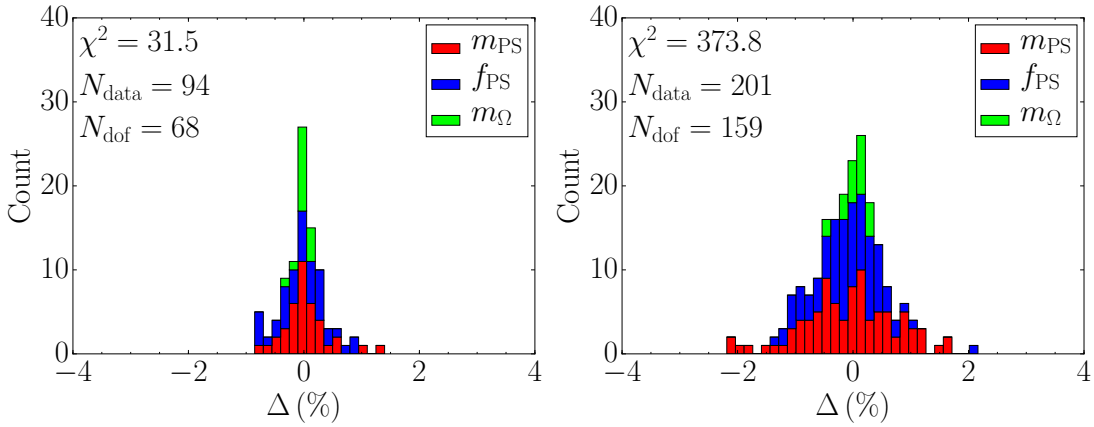

(a) NLO, 370 MeV cut

(b) NNLO, 510 MeV cut with frozen LO LECs

Figure 4.13: Percent deviation between fits and data. We plot stacked histograms of the quantity $\Delta \equiv 200 \times (Y - Y^{\text{fit}})/(Y + Y^{\text{fit}})$.

As we summarize and interpret the various fits we have performed to assess the accuracy with which ChPT formulae match our data, it is important to note the particular features of the (unitary) light-light pseudoscalar mass squared, m_{ll}^2, as a function of the light quark mass, m_l. While leading order ChPT predicts that $m_{ll}^2 \propto m_l$, lattice calculations show that this linearity persists to a good approximation even when m_l is significantly heavier than the physical up and down quarks[3]. This linearity implies that if NLO ChPT is fit to lattice data, the NLO corrections must be small for the fit to accurately represent the data, as we see, for example, in Figure 4.4a for the $SU(3)$ case. Likewise, for an NNLO fit to accurately represent the data either both the NLO and NNLO corrections must be small (as we observed in ref. [1] for $SU(2)$) or they must have the same magnitude and contribute with opposite sign (*e.g.* Figure 4.4c for $SU(3)$). In particular, we observe a substantial change in the NLO contributions between Figure 4.4a and Figure 4.4c. Since we find that all four of our preferred fits represent m_{ll}^2 well, we note that the NNLO fit does not produce a series with hierarchical terms |LO|> |NLO|> |NNLO|, but does produce a series where the corrections (NLO + NNLO) to LO are small. Given this understanding of the fits to

[3]There is visible curvature for very small quark masses, consistent with chiral logarithms, as seen, for example, in Figure 3.2.

the (squared) pseudoscalar masses, we can summarize our fits, paying particular attention to the pseudoscalar decay constants, which are markedly nonlinear functions of the quark masses and thus probe whether the terms in the chiral expansion have a reasonable hierarchy of sizes.

We have found that we obtain good fits of the NLO $SU(3)$ expressions for the pseudoscalar masses and decay constants to partially quenched lattice data containing two light quarks for meson masses ranging from 120 MeV to 370 MeV. Extending the upper mass cut to 510 MeV — slightly heavier than the physical kaon — and including heavy-light data we find a large (uncorrelated) $\chi^2/\text{dof} \approx 6.5$, demonstrating that NLO fits cannot accurately reproduce the data over this fit range at the level of precision of our statistical errors. While a poor χ^2/dof indicates that the ChPT fit curves are many standard deviations away from the lattice data, we also note that, expressed as the percent deviation between the fit curve and the data, this is at worst $\mathcal{O}(5\%)$ for this fit. However, the NLO corrections to the decay constant are 40% of the size of the LO term at the heaviest light-light and heavy-light points included in the fit, suggesting that the $\mathcal{O}(5\%)$ agreement we observe is an accident of curve fitting rather than a reliable agreement between the data and ChPT. We also note that removing lighter mass points from these NLO fits results in a systematic shift in the fitted LECs — in particular, for the chiral decay constant f_0 — which reproduces the difficulty with fitting NLO $SU(3)$ PQChPT to data from the 24I ensemble first observed in Ref. [17]. These single ensemble fits gave $f_0 \simeq 95\,\text{MeV}$, suggesting that the NLO $SU(3)$ expansion is unreliable at the relatively heavy masses constrained by the 24I data.

Repeating the $SU(3)$ fits with two upper mass cuts of 370 and 510 MeV and including NNLO corrections, we observed a substantial shift in the value of f_0 from the NLO value $f_0 \simeq 114\,\text{MeV}$ to $f_0 \simeq 128\,\text{MeV}$ for both NNLO fits. This indicates that, with our current data, the fit is balancing LO and NNLO terms — this could be an indication that we need more data to properly constrain the NNLO terms, an indication that the expansion is breaking down, or a combination of both issues. We find that we can stabilize our fits by constraining the leading-order LECs using the values we obtain from our NLO fits. After doing this, we find that the terms remain hierarchically ordered for f_{ll} with the heavy (dynamical) quark fixed at the physical strange quark mass, *i.e.* we find |LO|> |NLO|> |NNLO|. For f_{ll}, the total correction (NLO + NNLO) at the heaviest

light quark mass included in the fit is about 35% of the LO prediction. Using the fit results and extrapolating to quarks which are 30% heavier, we observe that the NLO and NNLO terms become comparable in size. This evidence indicates that at a pseudoscalar mass of approximately 500 MeV — essentially the kaon mass — we are near the limit of where NNLO $SU(3)$ ChPT provides a sensible approximation to the data. The fits agree with the data points to within a few percent and are self-consistent in the sense that, assuming NNLO ChPT is valid and performing the fits, we obtain LECs which show the series is reasonably convergent up to NNLO. We have also calculated the unquenched leading and next-to leading order $SU(3)$ ChPT LECs and find that our values are generally consistent with other lattice and phenomenological results.

We note that while we find $SU(3)$ NNLO PQChPT represents the data with percent-scale accuracy up to the kaon mass scale, we have used data ranging from slightly below the physical pion mass to the physical kaon mass to determine the terms in the ChPT expansion. We do not have enough accurate data to try and determine the full NNLO expansion from data restricted to a lighter mass range — for example, the 300 to 450 MeV range — which we could then extrapolate up to the physical kaon mass. Furthermore, all of our ensembles have dynamical heavy quark masses near the physical strange quark mass; ideally one would like a series of ensembles extending from the $SU(3)$ chiral limit up to the physical quark masses as a playground to systematically explore $SU(3)$ ChPT. Further lattice simulations, producing ever-more-accurate measurements of the pseudoscalar masses and decay constants, will provide more information about the LECs and the behavior of the expansion.

Acknowledgements

We thank members of the RBC-UKQCD collaboration for helpful discussions and support of this work. We are especially grateful to J. Bijnens, who provided Fortran routines used to compute next-to-next-to leading order partially-quenched $SU(3)$ chiral perturbation theory expressions for the pseudoscalar masses and decay constants. The fits were performed using the Columbia University CUTH cluster.

R.D.M. and D.J.M. are supported in part by U.S. DOE grant #DE-SC0011941. C.K. is sup-

ported by a RIKEN foreign postdoctoral research (FPR) grant.

References

[1] P. A. Boyle et al., "Low Energy Constants of $SU(2)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD", Phys. Rev. **D93**, 054502 (2016).

[2] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[3] Steven Weinberg, "Phenomenological Lagrangians", Physica A: Statistical Mechanics and its Applications **96**, 327–340 (1979).

[4] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory: Expansions in the Mass of the Strange Quark", Nucl.Phys. **B250**, 465 (1985).

[5] Bijnens, Johan and Danielsson, Niclas and Lähde, Timo A., "Three-Flavor Partially Quenched Chiral Perturbation Theory at NNLO for Meson Masses and Decay Constants", Phys.Rev. **D73**, 074509 (2006).

[6] R. Mawhinney, "NLO and NNLO Chiral Fits for 2+1 Flavor DWF Ensembles", PoS **LAT2009**, 081 (2009).

[7] A. Bazavov et al., "MILC Results for Light Pseudoscalars", PoS **CD09**, 007 (2009).

[8] A. Bazavov et al., "Results for Light Pseudoscalar Mesons", PoS **LATTICE2010**, 074 (2010).

[9] Y. Iwasaki and T. Yoshie, "Renormalization Group Improved Action for $SU(3)$ Lattice Gauge Theory and the String Tension", Phys.Lett. **B143**, 449 (1984).

[10] P. M. Vranas, "Gap Domain Wall Fermions", Phys.Rev. **D74**, 034512 (2006).

[11] D. Renfrew, T. Blum, N. Christ, R. Mawhinney, and P. Vranas, "Controlling Residual Chiral Symmetry Breaking in Domain Wall Fermion Simulations", PoS **LATTICE2008**, 048 (2008).

[12] D. B. Kaplan, "A Method for Simulating Chiral Fermions on the Lattice", Phys. Lett. **B288**, 342–347 (1992).

[13] Y. Shamir, "Chiral Fermions from Lattice Boundaries", Nucl. Phys. **B406**, 90–106 (1993).

[14] R. C. Brower, H. Neff, and K. Orginos, "Mobius Fermions: Improved Domain Wall Chiral fermions", Nucl. Phys. Proc. Suppl. **140**, [,686(2004)], 686–688 (2005).

[15] R. Brower, H. Neff, and K. Orginos, "Möbius Fermions", Nucl.Phys.Proc.Suppl. **153**, 191–198 (2006).

[16] R. C. Brower, H. Neff, and K. Orginos, "The Möbius Domain Wall Fermion Algorithm", (2012).

[17] C. Allton et al., "Physical Results from 2+1 Flavor Domain Wall QCD and $SU(2)$ Chiral Perturbation Theory", Phys.Rev. **D78**, 114509 (2008).

[18] Y. Aoki et al., "Continuum Limit Physics from $2 + 1$ Flavor Domain Wall QCD", Phys. Rev. **D83**, 074508 (2011).

[19] R. Arthur et al., "Domain Wall QCD with Near-Physical Pions", Phys.Rev. **D87**, 094514 (2013).

[20] K.A. Olive and Particle Data Group, "Review of Particle Physics", Chinese Physics C **38**, 090001 (2014).

[21] K. Levenberg, "A Method for the Solution of Certain Non-Linear Problems in Least Squares", Quarterly Journal of Applied Mathmatics **II**, 164–168 (1944).

[22] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters", Journal of the Society for Industrial and Applied Mathematics **11**, 431–441 (1963).

[23] S. Aoki, Y. Aoki, C. Bernard, T. Blum, G. Colangelo, et al., "Review of Lattice Results Concerning Low-Energy Particle Physics", Eur.Phys.J. **C74**, 2890 (2014).

[24] A. Bazavov et al., "Full Nonperturbative QCD Simulations with 2+1 Flavors of Improved Staggered Quarks", Rev.Mod.Phys. **82**, 1349–1417 (2010).

[25] J. Bijnens and I. Jemos, "A New Global Fit of the L_i^r at Next-to-Next-to-Leading Order in Chiral Perturbation Theory", Nucl.Phys. **B854**, 631–665 (2012).

[26] J. Bijnens and G. Ecker, "Mesonic Low-Energy Constants", Ann.Rev.Nucl.Part.Sci. **64**, 149–174 (2014).

[27] A. Bazavov et al., "Staggered Chiral Perturbation Theory in the Two-Flavor Case and $SU(2)$ Analysis of the MILC Data", PoS **LATTICE2010**, 083 (2010).

[28] S. Borsanyi, S. Durr, Z. Fodor, S. Krieg, A. Schafer, et al., "$SU(2)$ Chiral Perturbation Theory Low-Energy Constants from 2+1 Flavor Staggered Lattice Simulations", Phys.Rev. **D88**, 014513 (2013).

[29] S. R. Beane, W. Detmold, P. M. Junnarkar, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, A. Torok, and A. Walker-Loud, "$SU(2)$ Low-Energy Constants from Mixed-Action Lattice QCD", Phys. Rev. **D86**, 094509 (2012).

[30] J. Gasser and H. Leutwyler, "Chiral Perturbation Theory to One Loop", Annals Phys. **158**, 142 (1984).

[31] G. Colangelo, J. Gasser, and H. Leutwyler, "$\pi\pi$ Scattering", Nucl.Phys. **B603**, 125–179 (2001).

[32] G. M. de Divitiis et al., "Isospin Breaking Effects Due to the Up-Down Mass Difference in Lattice QCD", JHEP **04**, 124 (2012).

[33] T. Janowski, P. Boyle, A. Jüttner, and C. Sachrajda, "Determination of $K - \pi$ Scattering Lengths at Physical Kinematics", 33rd International Symposium on Lattice Field Theory, 2015.

[34] J. Bijnens, G. Colangelo, and G. Ecker, "Renormalization of Chiral Perturbation Theory to Order p^6", Annals Phys. **280**, 100–139 (2000).

[35] J. Gasser, C. Haefeli, M. A. Ivanov, and M. Schmid, "Integrating out Strange Quarks in ChPT", Phys.Lett. **B652**, 21–26 (2007).

[36] Z. Fu, "Lattice Study on πK Scattering with Moving Wall Source", Phys.Rev. **D85**, 074501 (2012).

4.A χPT Relations

In this appendix we collect various relations used in the analysis in the body of the paper. We do not explicitly reprint the expressions for the pseudoscalar masses and decay constants and the corresponding finite volume corrections used in the chiral fits: instead we refer the reader to the

appendices of Ref. [17]. The NNLO pseudoscalar masses and decay constants were computed using Fortran routines provided by J. Bijnens.

4.A.1 Relations Between PQχPT and χPT LECs at NLO

The $SU(N_f)$ ChPT Lagrangian can be recovered from the more general $SU(N_f)$ PQChPT Lagrangian in the limit of equal sea and valence quark masses. Here we have collected the explicit expressions relating the NLO LECs in this limit from Ref. [34]. The analogous expressions for the NNLO LECs can be found in the same reference, but we do not use them here. For $N_f = 3$, the NLO ChPT LECs $\{L_i\}_{i=1}^{10}$ are related to the NLO PQChPT LECs $\{\hat{L}_i^{(3)}\}_{i=0}^{12}$ by

$$\begin{aligned} L_1 &= \tfrac{1}{2}\hat{L}_0^{(3)} + \hat{L}_1^{(3)} \quad & L_3 &= -2\hat{L}_0^{(3)} + \hat{L}_3^{(3)} \\ L_2 &= \hat{L}_0^{(3)} + \hat{L}_2^{(3)} \quad & L_i &= \hat{L}_i^{(3)}\,, \;\; i = 4, \ldots, 10 \end{aligned} \tag{D.16}$$

and $\hat{L}_{11}^{(3)} = \hat{L}_{12}^{(3)} = 0$.

4.A.2 Scale Independent $SU(2)$ LECs

Conventionally, one quotes values of the scale independent $SU(2)$ ChPT LECs $\{\overline{\ell}_i\}_{i=1}^{6}$ rather than $\{l_i\}_{i=1}^{6}$. These are obtained by running the $\{l_i\}_{i=1}^{6}$ from the energy scale at which they are defined, μ, to the physical pion mass using

$$\overline{\ell}_i = \gamma_i l_i - \log\left(\frac{m_\pi^2}{\mu^2}\right), \tag{D.17}$$

where the coefficients

$$\gamma_1 = 96\pi^2, \quad \gamma_2 = 48\pi^2, \quad \gamma_3 = -64\pi^2, \quad \gamma_4 = 16\pi^2, \quad \gamma_5 = -192\pi^2, \quad \gamma_6 = -96\pi^2, \tag{D.18}$$

were computed in Ref. [30]. The remaining LEC l_7 has no scale dependence.

4.A.3 Relations Between $SU(2)$ and $SU(3)$ LECs

By integrating out the strange quark in the $SU(3)$ theory and matching to the $SU(2)$ theory, one can write down explicit relations between the LECs. Gasser and Leutwyler worked out the one-loop

expressions [4] for the leading order LECs

$$\begin{aligned} B &= B_0 \left[1 - \frac{\chi_s}{72\pi^2 f_0^2} \log\left(\frac{2\chi_s}{3\Lambda_\chi^2}\right) - \frac{16\chi_s}{f_0^2}\left(L_4 - 2L_6\right)\right] \\ f &= f_0 \left[1 - \frac{\chi_s}{32\pi^2 f_0^2} \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right) + \frac{8\chi_s}{f_0^2} L_4\right] \end{aligned} \tag{D.19}$$

and next-to leading order LECs

$$\begin{aligned} l_1 &= 4L_1 + 2L_3 - \frac{1}{768\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right] \\ l_2 &= 4L_2 - \frac{1}{384\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right] \\ l_3 &= -8L_4 - 4L_5 + 16L_6 + 8L_8 - \frac{1}{576\pi^2}\left[1 + \log\left(\frac{2\chi_s}{3\Lambda_\chi^2}\right)\right] \\ l_4 &= 8L_4 + 4L_5 - \frac{1}{64\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right] \\ l_5 &= L_{10} + \frac{1}{384\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right] \\ l_6 &= -2L_9 + \frac{1}{192\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right] \\ l_7 &= \frac{f_0^2}{8\chi_s}\left[1 + \frac{5\chi_s}{36\pi^2 f_0^2}\log\left(\frac{2\chi_s}{3\Lambda_\chi^2}\right)\right] + 4\left(L_4 - L_6 - 9L_7 - 3L_8 + \frac{1}{256\pi^2}\left[1 + \log\left(\frac{\chi_s}{2\Lambda_\chi^2}\right)\right]\right) \end{aligned} \tag{D.20}$$

The full two-loop expressions are also known [35], but we do not make use of them here.

4.A.4 One-Loop $SU(3)$ Predictions

We use the $SU(3)$ LECs determined by our fits to examine the one-loop predictions for isospin breaking [4] in the kaon mass

$$\begin{aligned} \left[\frac{\left(m_{K^0}^2 - m_{K^\pm}^2\right)}{\left(m_d - m_u\right)}\right]_{\text{QCD}} = B_0 \Bigg[1 &+ \frac{\chi_\eta}{24\pi^2 f_0^2}\log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right) + \frac{1}{16\pi^2 f_0^2}\frac{\chi_s + \chi_l}{\chi_s - \chi_l}\left(\chi_\eta \log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right) - \chi_l \log\left(\frac{\chi_l}{\Lambda_\chi^2}\right)\right) \\ &+ \frac{16}{f_0^2}\left(2L_8 - L_5\right)\left(\chi_l + \chi_s\right) + \frac{48}{f_0^2}\left(2L_6 - L_4\right)\overline{\chi}\Bigg] \end{aligned} \tag{D.21}$$

and kaon decay constant

$$\left[\frac{\left(f_{K^0}/f_{K^\pm} - 1\right)}{\left(m_d - m_u\right)}\right]_{\text{QCD}} = B_0 \left[\frac{8}{f_0^2}L_5 - \frac{1}{32\pi^2 f_0^2}\left\{1 + \log\left(\frac{\chi_l + \chi_s}{2\Lambda_\chi^2}\right) + \frac{2}{\chi_\eta - \chi_l}\left(\chi_\eta \log\left(\frac{\chi_\eta}{\Lambda_\chi^2}\right) - \chi_l \log\left(\frac{\chi_l}{\Lambda_\chi^2}\right)\right)\right\}\right]. \tag{D.22}$$

We use the subscript "QCD" to emphasize that this is only the contribution to the mass splitting from QCD isospin breaking, and does not include electromagnetic corrections. We also compute the S-wave πK scattering lengths using the formulae of Ref. [36].

4.B Fits with Weighted χ^2

In this appendix we repeat the fits discussed in the main body of the text, this time minimizing a weighted χ^2 defined by

$$\chi^2_e = \frac{1}{N_e} \sum_{j \in e} \left(\frac{y^j_e - f^j_e}{\sigma^j_e} \right)^2 , \ \chi^2 = \sum_e \chi^2_e \tag{D.23}$$

where N_e is the number of non-degenerate quark mass combinations (m_x, m_y, m_l, m_h) used for pseudoscalar measurements on ensemble e. As we have argued in Appendix D of Ref. [1], this procedure captures the dominant influence of the strong correlations between partially quenched measurements computed on the same ensemble but with different valence quark masses on the global fit. The differences between these fits and the unweighted fits discussed in the main text are factored into our final error budget following the procedure we discuss in section 4.5.

Mass Cut	24I	32I	32ID	32I-fine	48I	64I	32ID-M1	32ID-M2
370 MeV	3	12	20	1	1	1	1	—
510 MeV	19	30	33	1	2	2	2	1

Table 4.10: The value of N_e for each ensemble and mass cut.

4.B.1 Fit Parameters

				Free		Frozen LO LECs
		NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
	χ^2/dof	0.041	0.319	0.028	0.066	0.114
24I	am_l^{phys}	-0.001844(49)	-0.001830(45)	-0.001817(61)	-0.001830(78)	-0.00190(20)
	am_h^{phys}	0.03341(49)	0.03367(32)	0.0356(29)	0.0331(20)	0.0306(30)
	a^{-1}	1.730(19) GeV	1.725(12) GeV	1.715(35) GeV	1.774(45) GeV	1.843(101) GeV
32I	am_l^{phys}	0.000286(14)	0.000283(10)	0.000295(21)	0.000267(23)	0.000245(55)
	am_h^{phys}	0.02496(64)	0.02555(23)	0.02574(42)	0.02551(27)	0.0249(20)
	a^{-1}	2.306(24) GeV	2.321(15) GeV	2.294(34) GeV	2.354(33) GeV	2.397(87) GeV
32ID	am_l^{phys}	-0.000176(38)	-0.000150(36)	-0.00010(13)	-0.00003(18)	-0.00011(83)
	am_h^{phys}	0.0413(28)	0.0439(22)	0.0468(59)	0.0442(77)	0.042(21)
	a^{-1}	1.400(13) GeV	1.388(15) GeV	1.375(23) GeV	1.388(59) GeV	1.426(273) GeV
32I-fine	am_l^{phys}	0.000080(28)	0.000063(27)	0.000058(45)	0.000031(59)	0.00003(14)
	am_h^{phys}	0.0213(34)	0.01814(63)	0.0192(24)	0.0179(17)	0.0181(51)
	a^{-1}	3.094(30) GeV	3.104(31) GeV	3.104(61) GeV	3.138(75) GeV	3.138(202) GeV
48I	am_l^{phys}	0.0007018(63)	0.0007124(54)	0.000729(54)	0.000759(54)	0.00071(20)
	am_h^{phys}	0.03605(31)	0.03622(19)	0.0383(23)	0.0369(16)	0.0350(42)
	a^{-1}	1.725(2) GeV	1.724(2) GeV	1.710(11) GeV	1.714(23) GeV	1.745(69) GeV
64I	am_l^{phys}	0.0006237(74)	0.0006260(49)	0.0006295(100)	0.000623(13)	0.000615(37)
	am_h^{phys}	0.02486(68)	0.02559(22)	0.02553(33)	0.02592(24)	0.0257(21)
	a^{-1}	2.348(6) GeV	2.349(4) GeV	2.345(8) GeV	2.349(8) GeV	2.359(44) GeV
32ID-M1	am_l^{phys}	0.000733(47)	0.000614(22)	0.00078(38)	0.00097(53)	0.0008(20)
	am_h^{phys}	0.097(37)	0.0760(29)	0.092(14)	0.075(13)	0.066(55)
	a^{-1}	1.046(8) GeV	1.042(3) GeV	1.034(23) GeV	1.023(64) GeV	1.054(248) GeV
32ID-M2	am_l^{phys}	—	-0.003360(15)	—	-0.003388(84)	-0.00344(11)
	am_h^{phys}	—	0.02603(61)	—	0.0251(17)	0.0235(55)
	a^{-1}	—	1.961(16) GeV	—	2.028(67) GeV	2.103(148) GeV

Table 4.11: The (uncorrelated) χ^2/dof, unrenormalized physical quark masses in bare lattice units (without m_{res} included), and the values of the inverse lattice spacing a^{-1} in physical units, obtained from fits to $SU(3)$ PQChPT with the stated pion mass cuts.

				Free		Frozen LO LECs
		NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
24I	Z_l	0.9701(56)	0.9659(22)	0.963(46)	0.936(29)	0.95(13)
	Z_h	0.934(26)	0.9583(76)	0.911(54)	0.959(33)	0.987(100)
	R_a	0.7503(29)	0.7429(28)	0.7475(62)	0.754(18)	0.769(42)
32I	Z_l	1.0	1.0	1.0	1.0	1.0
	Z_h	1.0	1.0	1.0	1.0	1.0
	R_a	1.0	1.0	1.0	1.0	1.0
32ID	Z_l	0.942(13)	0.9382(95)	0.919(78)	0.872(57)	0.89(28)
	Z_h	0.978(59)	0.960(36)	0.907(99)	0.96(14)	0.99(34)
	R_a	0.6073(65)	0.5977(76)	0.5993(86)	0.590(23)	0.59(12)
32I-fine	Z_l	1.000(30)	1.024(30)	1.033(52)	1.059(78)	1.05(17)
	Z_h	0.87(15)	1.044(31)	0.99(11)	1.060(98)	1.04(29)
	R_a	1.342(16)	1.337(13)	1.353(41)	1.333(39)	1.309(77)
48I	Z_l	0.9701(56)	0.9659(22)	0.963(46)	0.936(28)	0.95(13)
	Z_h	0.934(26)	0.9583(76)	0.911(54)	0.959(33)	0.987(100)
	R_a	0.7483(78)	0.7428(48)	0.745(10)	0.728(13)	0.728(50)
64I	Z_l	1.0	1.0	1.0	1.0	1.0
	Z_h	1.0	1.0	1.0	1.0	1.0
	R_a	1.0183(86)	1.0122(60)	1.022(12)	0.998(12)	0.984(34)
32ID-M1	Z_l	0.7230(91)	0.7594(28)	0.722(84)	0.684(81)	0.70(31)
	Z_h	0.57(22)	0.747(32)	0.62(11)	0.777(99)	0.85(52)
	R_a	0.4538(80)	0.4488(40)	0.451(16)	0.435(31)	0.44(11)
32ID-M2	Z_l	—	1.0162(86)	—	1.005(40)	1.012(83)
	Z_h	—	1.018(18)	—	1.027(36)	1.04(11)
	R_a	—	0.8446(36)	—	0.862(18)	0.878(46)

Table 4.12: Ratios of lattice spacings (R_a) and light and heavy quark masses (Z_l, Z_h) between each ensemble and the reference 32I ensemble.

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
B_0	4.240(84) GeV	3.924(48) GeV	3.90(16) GeV	4.048(86) GeV	4.05(15) GeV
f_0	0.1147(29) GeV	0.1111(16) GeV	0.1281(44) GeV	0.128(23) GeV	0.1147(43) GeV
$10^3 \hat{L}_0^{(3)}$	—	—	0.66(95)	-0.11(59)	-0.2(2.2)
$10^3 \hat{L}_1^{(3)}$	—	—	-0.160(77)	-0.34(46)	-0.29(62)
$10^3 \hat{L}_2^{(3)}$	—	—	-1.33(31)	-0.64(24)	-0.70(28)
$10^3 \hat{L}_3^{(3)}$	—	—	-0.40(23)	-0.02(74)	-0.1(1.1)
$10^3 \hat{L}_4^{(3)}$	-0.111(62)	-0.022(33)	-0.30(19)	-0.38(69)	-0.19(21)
$10^3 \hat{L}_5^{(3)}$	0.922(87)	0.902(32)	0.99(43)	1.15(79)	0.92(53)
$10^3 \hat{L}_6^{(3)}$	-0.077(40)	0.028(23)	-0.09(28)	-0.23(39)	-0.11(16)
$10^3 \hat{L}_7^{(3)}$	—	—	-0.20(19)	-0.11(29)	-0.14(25)
$10^3 \hat{L}_8^{(3)}$	0.626(35)	0.4599(96)	0.30(79)	0.52(37)	0.41(37)
$10^6 \left(\hat{K}_{17}^{(3)} - \hat{K}_{39}^{(3)}\right)$	—	—	-6(13)	-1.0(2.6)	-1(11)
$10^6 \left(\hat{K}_{18}^{(3)} - \hat{K}_{40}^{(3)}\right)$	—	—	-1(15)	0.2(1.4)	1.8(9.2)
$10^6 \hat{K}_{19}^{(3)}$	—	—	-3(20)	-2.9(1.3)	-2.6(4.4)
$10^6 \hat{K}_{20}^{(3)}$	—	—	-4.0(6.5)	-2.4(5.0)	-1.8(7.3)
$10^6 \hat{K}_{21}^{(3)}$	—	—	-2.7(4.1)	-1.6(4.8)	0.3(4.7)
$10^6 \hat{K}_{22}^{(3)}$	—	—	1.65(40)	2.3(3.5)	1.04(61)
$10^6 \hat{K}_{23}^{(3)}$	—	—	-2.3(5.1)	-1.1(3.0)	-1(11)
$10^6 \hat{K}_{25}^{(3)}$	—	—	0(10)	-0.6(1.1)	-0.5(3.1)
$10^6 \hat{K}_{26}^{(3)}$	—	—	-3.8(5.0)	-1.9(2.2)	-0.1(3.8)
$10^6 \hat{K}_{27}^{(3)}$	—	—	0.62(43)	1.0(1.2)	0.52(44)
$m^{(\Omega)}$	1.6633(33) GeV	1.6676(32) GeV	1.6652(38) GeV	1.667(10) GeV	1.666(36) GeV
c_{m_l,m_Ω}	4.2(1.1)	2.20(51)	3.3(1.1)	2.6(3.5)	2.8(7.5)
c_{m_y,m_Ω}	5.33(17)	5.36(13)	5.56(37)	5.33(75)	5.3(1.5)
c_{m_h,m_Ω}	-2.2(2.1)	0.80(99)	-1.3(2.9)	3.2(5.5)	4(12)

Table 4.13: $SU(3)$ PQChPT LECs fit at the chiral scale $\Lambda_\chi = 770\,\text{MeV}$ in units of the canonical size at a given order in the chiral expansion. The parameters $m^{(\Omega)}$ and c_{m_q,m_Ω} are the constant term and m_q slopes for the (linear) m_Ω ansatz, respectively. The value of B_0 quoted here is unrenormalized.

4.B.2 Predictions

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
$B_0^{\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	2.796(62)(15) GeV	2.588(39)(14) GeV	2.573(108)(14) GeV	2.670(63)(15) GeV	2.673(102)(15) GeV
f_0	114.7(2.9) MeV	111.1(1.6) MeV	128.1(4.4) MeV	128(23) MeV	114.7(4.3) MeV
$\Sigma_0^{1/3,\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	264.0(5.5)(0.5) MeV	251.8(3.4)(0.5) MeV	276.4(7.0)(0.5) MeV	279.1(3.2)(0.5) MeV	260.0(9.7)(0.5) MeV
$10^3 L_1$	—	—	0.17(48)	-0.39(69)	-0.4(1.7)
$10^3 L_2$	—	—	-0.67(66)	-0.75(75)	-0.9(2.2)
$10^3 L_3$	—	—	-1.7(2.0)	0.2(1.7)	0.3(5.3)
$10^3 L_4$	-0.102(59)	-0.044(34)	-0.30(19)	-0.38(69)	-0.19(21)
$10^3 L_5$	0.934(73)	0.913(32)	0.99(43)	1.15(79)	0.92(53)
$10^3 L_6$	-0.070(40)	0.018(24)	-0.09(28)	-0.23(39)	-0.11(16)
$10^3 L_7$	—	—	-0.20(19)	-0.11(29)	-0.14(25)
$10^3 L_8$	0.639(31)	0.466(11)	0.30(79)	0.52(37)	0.41(37)

Table 4.14: Unquenched $SU(3)$ LECs computed from partially quenched $SU(3)$ fits at the chiral scale $\Lambda_\chi = 770\,{\rm MeV}$. Missing entries are not constrained by the fits at a given order. For B_0 and Σ_0 the first error is statistical and the seccond is a systemati uncertainty in the perturbative matching to $\overline{\rm MS}$.

			Free		**Frozen LO LECs**
LEC	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
$B^{\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	2.817(38)(15) GeV	2.804(30)(15) GeV	2.780(426)(15) GeV	2.630(180)(14) GeV	2.695(647)(15) GeV
f	123.1(1.0) MeV	122.0(0.5) MeV	130.1(9.5) MeV	126.4(6.1) MeV	120.0(9.6) MeV
$\Sigma^{1/3,\overline{\rm MS}}(\mu = 2\,{\rm GeV})$	277.4(2.0)(0.5) MeV	275.3(1.3)(0.5) MeV	286.6(27.9)(0.5) MeV	276.0(6.4)(0.5) MeV	268.7(15.3)(0.5) MeV
$\bar{\ell}_3$	2.85(18)	3.07(11)	3.9(1.6)	4.21(88)	4.0(2.7)
$\bar{\ell}_4$	3.909(52)	3.982(26)	3.71(24)	3.70(39)	3.79(50)
$10^3 l_7$	—	—	6.1(6.3)	0(14)	2(11)

Table 4.15: Unquenched $SU(2)$ LECs computed from partially quenched $SU(3)$ fits and one-loop relations. Missing entries are not constrained by the fits at a given order. For B and Σ the first error is statistical and the second is a systematic uncertainty associated with the perturbative matching to $\overline{\rm MS}$.

			Free		**Frozen LO LECs**
	NLO (370 MeV cut)	NLO (510 MeV cut)	NNLO (370 MeV cut)	NNLO (510 MeV cut)	NNLO (510 MeV cut)
m_K	0.5136(97) GeV	0.4951(19) GeV	0.51(17) GeV	0.4980(39) GeV	0.497(15) GeV
f_K	0.1555(12) GeV	0.15582(15) GeV	0.156(27) GeV	0.1563(10) GeV	0.1563(23) GeV
f_K/f_0	1.355(36)	1.403(21)	1.22(24)	1.22(21)	1.349(22)
$[m^2_{K^0} - m^2_{K^\pm}]_{\rm QCD}/\Delta m_{du}$	5.45(26) GeV	3.677(67) GeV	2.3(4.1) GeV	3.02(49) GeV	3.0(2.4) GeV
$[\frac{f_{K^0}}{f_{K^\pm}} - 1]_{\rm QCD}/\Delta m_{du}$	3.02(14) GeV^{-1}	3.079(40) GeV^{-1}	2.46(92) GeV^{-1}	2.78(61) GeV^{-1}	2.9(1.5) GeV^{-1}
$m_\pi a_0^{1/2}$	—	—	0.1245(88)	0.129(17)	0.141(35)
$m_\pi a_0^{3/2}$	—	—	-0.0643(69)	-0.059(11)	-0.069(34)

Table 4.16: Predictions from NLO and NNLO fits and $SU(3)$ ChPT. $\Delta m_{du} \equiv m_d - m_u$.

Chapter 5

Semileptonic Kaon Decays

The Kaon Semileptonic Form Factor from Domain Wall QCD at the Physical Point

C. Kelly[1] and R.D. Mawhinney[1], D.J. Murphy[1], P.A. Boyle[3], N.H. Christ[2], J.M. Flynn[5], N. Garron[4], C. Jung[1], A. Jüttner[5], R.D. Mawhinney[2], D.J. Murphy[2], C.T. Sachrajda[5], F. Sanfilippo[5], and H. Yin[2]

[1]*Department of Physics, Columbia University, New York, NY 10027, USA*[1]*Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA*

[2]*Department of Physics, Columbia University, New York, NY 10027, USA*

[3]*School of Physics and Astronomy, University of Edinburgh, EH9 3JZ, UK*

[4]*School of Computing and Mathematics and Centre for Mathematical Science, Plymouth University, Plymouth PL4 8AA, UK*

[5]*School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK*

Abstract

We present the first calculation of the kaon semileptonic form factor at zero momentum transfer in $N_f = 2 + 1$ domain wall QCD with physical quark masses. We jointly analyze two sets of lattice simulations with different lattice spacings and large physical volumes, computed with pion momenta tuned to achieve zero momentum transfer. We obtain the continuum result $f_+^{K\pi}(0) = 0.9685(34)(14)$, where the first error is statistical and the second error is systematic. Combining this calculation with data from experimental measurements of the $K \to \pi\ell\nu$ decay rate, we predict the CKM matrix element $|V_{us}|= 0.2233(5)(8)$, where the first error is from experiment and the second is from the lattice calculation.

5.1 Introduction

In the Standard Model of particle physics flavor-changing weak decays are parametrized by the 3×3, unitary Cabibbo-Kobayashi-Maskawa (CKM) matrix [1, 2]. The elements of the CKM matrix, V_{ij}, parametrize decays of a quark with flavor i to a quark with flavor j, and can be described in terms of three real parameters and a single complex phase; the complex phase allows for decay mechanisms which violate CP-symmetry. Accurate determinations of the CKM matrix elements are necessary to fix four of the free parameters of the Standard Model, as well as to understand CP-violation in weak decays and to test unitarity[1].

In this work we present a new lattice QCD calculation of the hadronic contribution to the flavor-changing $K \to \pi$ decay, which is encapsulated by the vector form factor $f_+^{K\pi}(q^2)$ evaluated at vanishing momentum transfer $q^2 = 0$. Precision measurements of the $K \to \pi\ell\nu$ ($K_{\ell3}$) decay rate — which constrain the product $|V_{us}|f_+^{K\pi}(0)$ — together with lattice input for the form factor gives the most precise constraint on $|V_{us}|$ currently known. When combined with constraints on the other first-row CKM matrix elements — V_{ud} and V_{ub} — this also enables a precision Standard Model unitarity test through the quantity $\delta_u \equiv 1 - |V_{ud}|^2 - |V_{us}|^2 - |V_{ub}|^2$.

Recent lattice determinations of the vector form factor [3–8] are nicely summarized in the Flavor Lattice Averaging Group (FLAG) reviews [9, 10]. This quantity is currently known with an

[1] In particular, tension between the experimentally determined values of the CKM matrix elements and unitarity may prove to be an important indicator of new physics.

overall uncertainty of 0.3%. Improving the bound on $f_+^{K\pi}(0)$ is necessary since the lattice error is currently the largest source of uncertainty entering into the standard determination of the CKM matrix element V_{us}, and also since the KLOE-2 experiment [11] promises to further tighten the experimental constraint. The lattice error is typically dominated by statistical uncertainty associated with the Monte Carlo sampling of the lattice QCD path integral, but, historically, has also been due to systematic errors associated with using chiral perturbation theory or phenomenological models to extrapolate simulations with unphysically heavy pion masses to the physical point. Advances in algorithmic methods and supercomputer resources now allow simulations to be performed directly at physical quark masses, eliminating this systematic.

Here we present the first prediction of the vector form factor $f_+^{K\pi}(0)$ with physical mass domain wall quarks in the continuum limit of $N_f = 2 + 1$ flavor lattice QCD. The physics described by our simulations corresponds to nature up to isospin breaking corrections in the light quark masses, electromagnetic corrections, and contributions associated with vacuum polarization effects from the neglected heavy quark flavors; these corrections are expected to be sub-leading compared to current uncertainties.

5.2 Measurement Strategy

Experimental measurements of the $K^0 \to \pi^-$ decay rate constrain the product $|V_{us}| f_+^{K\pi}(0) = 0.2163(5)$ [12]. To extract $|V_{us}|$ we compute the QCD matrix element

$$\langle \pi(p_\pi)| V_\mu |K(p_K)\rangle = f_+^{K\pi}(q^2)\,(p_K + p_\pi)_\mu + f_-^{K\pi}(q^2)\,(p_K - p_\pi)_\mu\,, \tag{5.1}$$

where $V_\mu = Z_V \overline{u}\gamma_\mu s$ is the flavor-changing vector current, Z_V is the vector current renormalization coefficient, and $q^\mu = p_K^\mu - p_\pi^\mu$ is the momentum transfer between the kaon and pion. The vector form factor can also be related to a similar matrix element of the scalar density $S = \overline{u}s$ using a Ward identity for the vector current

$$q^\mu \langle V_\mu \rangle = (m_s - m_l)\,\langle \overline{s}u \rangle\,. \tag{5.2}$$

At the kinematical point $q^2 = 0$ the analogue of Equation (5.1) for the scalar density is

$$\langle \pi(p_\pi)| \, S \, |K(p_K)\rangle \Big|_{q^2=0} = \frac{m_K^2 - m_\pi^2}{m_s - m_l} f_+^{K\pi}(0), \tag{5.3}$$

where $m_u = m_d \equiv m_l$ is the (degenerate) light quark mass. Both calculations are summarized by the quark line diagram of figure 5.1. In all of our calculations we consider a kaon at rest and a pion

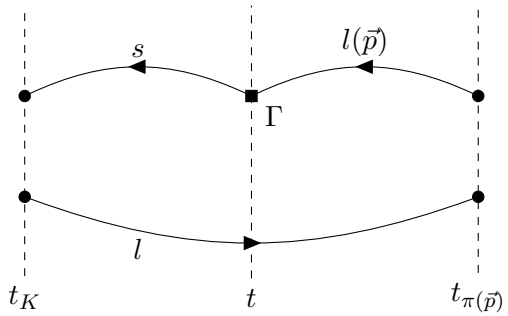

Figure 5.1: Quark line diagram for the $K_{\ell 3}$ three-point functions $\langle \pi(\vec{p})|\bar{s}\Gamma u|K\rangle$. The box ($\Gamma$) denotes an insertion of γ_μ for the vector matrix element, or $\mathbb{1}$ for the scalar matrix element.

with momentum $\vec{p}_\pi$, where twisted boundary conditions [13] are used to tune $\vec{p}_\pi$ such that $q^2 = 0$. We also make use of the all-mode averaging (AMA) technique as described in Ref. [14].

5.2.1 Simulation Parameters

We simulate $N_f = 2 + 1$ domain wall QCD using a series of ensembles with unitary pion masses ranging from 693 MeV down to the physical value of 139 MeV. Our older, unphysical ensemble sets A and C were generated using the Shamir kernel [15, 16] and were analyzed in Ref. [6]. We have generated two new physical pion mass Möbius domain wall fermion [17] ensembles, denoted A_{phys} and C_{phys}, with large volumes, which are the focus of this analysis, and are discussed in detail in Ref. [14]. All ensembles use the Iwasaki gauge action [18]. The full set of ensembles is summarized in Table 5.1.

Ensemble	Action	β	a (fm)	L/a	T/a	am_l	am_s^{sea}	am_s^{val}	θ_u^i / $\frac{2\pi}{L}$	m_π (MeV)	$m_\pi L$
A_3	DWF+I	2.13	0.11	24	64	0.03	0.04	0.04	—	693	9.3
A_2	DWF+I	2.13	0.11	24	64	0.02	0.04	0.04	—	575	7.7
A_1	DWF+I	2.13	0.11	24	64	0.01	0.04	0.04	—	431	5.8
A_5^4	DWF+I	2.13	0.11	24	64	0.005	0.04	0.04	—	341	4.6
A_5^3	DWF+I	2.13	0.11	24	64	0.005	0.04	0.03	—	341	4.6
C_8	DWF+I	2.25	0.08	32	64	0.008	0.03	0.025	—	431	5.5
C_6	DWF+I	2.25	0.08	32	64	0.006	0.03	0.025	—	360	4.8
C_4	DWF+I	2.25	0.08	32	64	0.004	0.03	0.025	—	304	4.1
A_{phys}	MDWF+I	2.13	0.11	48	96	0.00078	0.0362	0.0362	0.5893	139	3.8
C_{phys}	MDWF+I	2.25	0.08	64	128	0.000678	0.02661	0.02661	0.5824	139	3.9

Table 5.1: Summary of ensembles used in this analysis. DWF and MDWF denote domain wall fermions with the Shamir and Möbius kernels, respectively, and I denotes the Iwasaki gauge action. L and T are the size of the lattice in the spatial and temporal directions, respectively. m_π is the unitary pion mass. For the physical point ensembles θ_u^i is the twist angle applied to the up quark field in the three spatial directions, in units of $2\pi/L$.

5.2.2 Twisted Boundary Conditions

In a cubic box with side length L and lattice spacing a the allowed quark momenta in the spatial directions are

$$\vec{p} = \frac{2\pi}{L}\vec{k}, \quad \vec{k} \in \mathbb{Z}^3. \tag{5.4}$$

Assuming the kaon is at rest and the pion carries 3-momentum $\vec{p}_\pi$, the kinematical equation $q^2 = 0$ can be solved to find

$$|\vec{p}_\pi| = \frac{m_K^2 - m_\pi^2}{2m_K} \approx 229 \text{ MeV}. \tag{5.5}$$

Since this momentum is not, in general, a multiple of $2\pi/L$, the lattice calculation must be performed for several nearby values of the pion momentum and interpolated to $q^2 = 0$ — as we have done in our earliest calculation [4] using a phenomenological ansatz — introducing a systematic error associated with the interpolation. Alternatively, *twisted boundary conditions* can be imposed

on the quark fields $\psi(x)$:

$$\psi(x + L\hat{e}_i) = e^{i\theta_i}\psi(x), \tag{5.6}$$

where θ_i is a free parameter, allowing the first allowed momentum state to be tuned such that $q^2 = 0$ is satisfied directly. We also gain a further reduction in statistical noise by imposing the same twist angle in all three spatial directions, allowing the spatial components of the vector current matrix element (Eqn. (5.1)) to be averaged. In Ref. [19] Bernard et al. have demonstrated that, in general, additional form factors $h_\mu^{K\pi}(q)$ enter into the matrix element of Eqn. (5.1) at finite volume and at nonzero twist. However, the authors also demonstrate that these additional terms are associated with cubic symmetry breaking, and vanish when the choice of twist respects cubic symmetry. By twisting equally in all three spatial directions we avoid this complication in our calculation.

5.2.3 All-Mode Averaging (AMA)

The substantial cost of solving for light quark propagators on the large volume, physical pion mass A_{phys} and C_{phys} ensembles required us to make several algorithmic refinements to our measurement strategy. All correlation functions associated with these ensembles were computed using Coulomb gauge-fixed wall source propagators, together with the all-mode averaging (AMA) technique introduced in Ref. [20]. In the AMA formalism one replaces a direct calculation of an expensive lattice observable $\mathcal{O}$ with a less-expensive approximation $\mathcal{O}'$ and a correction term $\Delta\mathcal{O}$. The lattice action and ensemble averages $\langle\mathcal{O}\rangle$, $\langle\mathcal{O}'\rangle$, and $\langle\Delta\mathcal{O}\rangle$ are all assumed to be invariant under a group G of lattice symmetries. We define the AMA estimator by averaging the inexpensive approximation $\mathcal{O}'$ over some number N of transformations $g \in G$, and applying the correction term ΔO:

$$\mathcal{O}_{\text{AMA}} = \frac{1}{N}\sum_{g\in G}\mathcal{O}'_g + \Delta\mathcal{O}, \tag{5.7}$$

where the notation $\mathcal{O}_g$ denotes $\mathcal{O}$ computed after g is applied. We find in practice that the statistical error per unit of computer time can be markedly reduced using AMA with a judicious choice of $\mathcal{O}'$ and $\Delta\mathcal{O}$, relative to computing $\mathcal{O}$ directly.

In the context of this calculation the relevant lattice symmetry is the group of translations in the temporal direction. Quark propagators were computed using a deflated mixed-precision conjugate

gradient (CG) solver, with 600 (1500) single-precision low-mode deflation vectors obtained from the EigCG algorithm applied to a four dimensional volume source on the A_{phys} (C_{phys}) ensemble. We further distinguish between *exact* and *sloppy* light quark propagators. Exact light quark propagators were computed using a tight CG stopping residual $r = 10^{-8}$ for 7 (8) time slices. To avoid bias associated with the even-odd preconditioning used in the CG solves we randomly shifted the time slices used to compute exact propagators on each configuration. Sloppy light quark propagators were computed using a reduced precision $r = 10^{-4}$ and for all time slices. Strange quark propagators were sufficiently inexpensive that exact solves were computed for all time slices. For a given two- or three-point function we then constructed a sloppy estimate ($\mathcal{O}'$) for all time slices with the sloppy light quark propagators, and a correction term ($\Delta\mathcal{O}$) using the exact light quark propagators on time slices for which these are available. We then compute the AMA estimator according to (5.7), after averaging $\mathcal{O}'$ over all time translations. The full measurement package, which also computes observables related to the $K \to (\pi\pi)_{I=2}$ decay [21] and other low-energy QCD observables [14] from the same propagators, took 5.5 days per configuration on the A_{phys} ensemble using 1 rack (1024 nodes) of IBM Blue Gene/Q hardware, and 5.3 hours per measurement on the C_{phys} ensemble using 32 racks of Blue Gene/Q sustaining 1.2 PFlop/s. Additional details of the calculation can be found in [14].

The set of quark propagators described above is sufficient to generate AMA three-point functions, where at least one of the quarks coupling to the external current is a strange quark, for all possible source-sink-separations $\Delta t/a = |t_f - t_i|/a$ up to $T/2a$ (*e.g.* $K \to \pi$ and $K \to K$) and for all T/a possible translations in the temporal direction. Results at constant $\Delta t/a$ but with different t_i/a and t_f/a were binned together into a single, time-translation averaged measurement. For our choice of the 7 (8) source planes for the exact light quark solves on the A_{phys} (C_{phys}) ensemble the $\pi \to \pi$ three-point function entering into *e.g.* the determination of Z_V^π through Equation (**??**) can be computed on every fourth (fifth) source-sink-separation following the AMA prescription.

In all cases we use the bootstrap resampling technique with 500 samples to determine the statistical errors.

5.3 Data Analysis

5.3.1 Methodology

In the following we use the notation $C_{P_i}(t,\vec{p})$ to denote the pion and kaon two-point correlation functions

$$C_{P_i}(t,\vec{p}) \equiv \sum_{\vec{x},\vec{y}} \left\langle \mathcal{O}_{P_i,s_2}(t,\vec{y})\mathcal{O}^{\dagger}_{P_i,s_1}(0,\vec{x}) \right\rangle \overset{t/a\gg 1}{\approx} \frac{Z_{P_i,s_1}Z^{*}_{P_i,s_2}}{2E_{P_i}(\vec{p})}\left(e^{-E_{P_i}(\vec{p})t} + e^{-E_{P_i}(\vec{p})(T-t)}\right), \tag{5.8}$$

where $P_i \in \{\pi, K\}$, and $\mathcal{O}_{P_i,s}$ is an interpolating operator for the state P_i — $\mathcal{O}_{\pi,s} = \overline{u}\omega_s\gamma_5 d$ or $\mathcal{O}_{K,s} = \overline{d}\omega_s\gamma_5 s$ — with smearing kernel ω_s. In this study we have used $\omega_{s_1} = \mathrm{W}$ and $\omega_{s_2} \in \{\mathrm{L}, \mathrm{W}\}$ where "L" is a local source or sink and "W" is Coulomb gauge-fixed wall source or sink. We also use the notation $C_{\Gamma,P_iP_f}(t_i,t,t_f,\vec{p}_i,\vec{p}_f)$ to denote the three-point correlation functions

$$\begin{aligned} C_{\Gamma,P_iP_f}(t_i,t,t_f,\vec{p}_i,\vec{p}_f) &\equiv \sum_{\vec{x}_i,\vec{x},\vec{x}_f} \left\langle \mathcal{O}_{P_f,s_2}(t_f,\vec{x}_f)\Gamma(t,\vec{x})\mathcal{O}^{\dagger}_{P_i,s_1}(t_i,\vec{x}_i) \right\rangle \\ &\overset{t/a\gg 1}{\approx} \frac{Z_{P_i,s_1}Z^{*}_{P_f,s_2}}{4E_{P_i}(\vec{p}_i)E_{P_f}(\vec{p}_f)} \langle P_f(\vec{p}_f)|\,\Gamma\,|P_i(\vec{p}_i)\rangle \\ &\quad \times \left\{\theta(t_f-t)e^{-E_{P_i}(\vec{p}_i)(t-t_i)-E_{P_f}(\vec{p}_f)(t_f-t)} + c_\Gamma\theta(t-t_f)e^{-E_{P_i}(\vec{p}_i)(T+t_i-t)-E_{P_f}(\vec{p}_f)(t-t_f)}\right\} \end{aligned} \tag{5.9}$$

corresponding to the vector matrix element of Equation (5.1) ($\Gamma = V_\mu$) or the scalar matrix element of Equation (5.3) ($\Gamma = S$). In this notation $P_i, P_f \in \{\pi, K\}$ denote the initial and final states, and $\vec{p}_i$ and $\vec{p}_f$ the corresponding momenta. The constant $c_{V_0} = -1$ for the temporal component of the vector current, and $c_{V_i} = c_S = +1$ otherwise. For the three-point functions the final approximate equality holds in the combined limits of large Euclidean time separation between the initial and final states ($|t_f - t_i|/a \gg 1$), as well as insertion times t for the vector current or scalar density that are far from the source and sink time slices t_i and t_f.

While the K_{l3} form factor can be extracted by fitting directly to the asymptotic, large time separation limits of Equations (5.8) and (5.9), it is also possible to form ratios of the Euclidean

two- and three-point functions with simpler asymptotic forms [13]:

$$R_{V_\mu,K\pi}(t_K,t,t_\pi,\vec{p}_K,\vec{p}_\pi) \equiv 2\sqrt{E_\pi(\vec{p}_\pi)E_K(\vec{p}_K)}\left[\frac{C_{V_\mu,K\pi}(t_K,t,t_\pi,\vec{p}_K,\vec{p}_\pi)C_{V_\mu,\pi K}(t_K,t,t_\pi,\vec{p}_\pi,\vec{p}_K)}{\widetilde{C}_\pi(t_\pi-t_K,\vec{p}_\pi)\widetilde{C}_K(t_\pi-t_K,\vec{p}_K)}\right]^{1/2}$$
$$\approx \frac{1}{Z_V}\left((p_K+p_\pi)_\mu f_+^{K\pi}(q^2)+(p_K-p_\pi)_\mu f_-^{K\pi}(q^2)\right) \tag{5.10}$$

for the vector current, and

$$R_{S,K\pi}(t_K,t,t_\pi,\vec{p}_K,\vec{p}_\pi) \equiv 2\sqrt{E_\pi(\vec{p}_\pi)E_K(\vec{p}_K)}\left(\frac{m_s-m_l}{m_K^2-m_\pi^2}\right)\left[\frac{C_{S,K\pi}(t_K,t,t_\pi,\vec{p}_K,\vec{p}_\pi)C_{S,\pi K}(t_K,t,t_\pi,\vec{p}_\pi,\vec{p}_K)}{\widetilde{C}_\pi(t_\pi-t_K,\vec{p}_\pi)\widetilde{C}_K(t_\pi-t_K,\vec{p}_K)}\right]^{1/2}$$
$$\approx f_+^{K\pi}(q^2)+\frac{q^2}{m_K^2-m_\pi^2}f_-^{K\pi}(q^2) \tag{5.11}$$

for the scalar density. These ratios are, again, approximately equal to the final, asymptotic form up to noise and excited state contamination for large separations between the initial kaon and final pion $|t_\pi - t_K|/a \gg 1$ and t far from the source and sink. Here $C_{\Gamma,\pi K}$ is the three-point function for the time-reversed process $\pi \to K$, and

$$\widetilde{C}_{P_i}(t,\vec{p}) \equiv C_\pi(t,\vec{p}) - \frac{1}{2}C_{P_i}(T/2,\vec{p})\,e^{-E_{P_i}(\vec{p})(T/2-t)} \tag{5.12}$$

denotes the pseudoscalar two-point function with the backward propagating around-the-world mode analytically removed using the fitted energy $E_{P_i}(\vec{p})$.

The vector current renormalization factor Z_V is extracted from a similar analysis: we compute the analogue of 5.1 for the electromagnetic current and two pions or two kaons at rest. In this case the form factors are trivial — current conservation implies $f_+^{\pi\pi}(0) = f_+^{KK}(0) = 1$ and $f_-^{\pi\pi}(0) = f_-^{KK}(0) = 0$ — allowing us to extract Z_V directly from the temporal component of Equation (5.9), or from the ratio

$$R_{V_0,P_iP_f}(t_i,t,t_f,\vec{0},\vec{0}) = \frac{\widetilde{C}_{P_i}(t_f-t_i,\vec{0})}{C_{V_0,P_iP_f}(t_i,t,t_f,\vec{0},\vec{0})} \approx Z_V^{P_i}. \tag{5.13}$$

with $P_i = P_f \in \{\pi, K\}$. We observe that Z_V^π and Z_V^K differ by mass dependent cutoff effects at finite lattice spacing, and hence the form factors renormalized with either choice of Z_V follow two independent scaling trajectories which we expect to agree in the continuum limit $a \to 0$ by universality.

5.3.2 Fits on the Physical Point Ensembles

In this work we have considered two methods for extracting the $K_{\ell 3}$ form factor. In the first method, we simultaneously fit:

- The pion and kaon two-point functions — $C_\pi(t, \vec{p})$ and $C_K(t, \vec{0})$,
- The three-point functions determining Z_V^π and Z_V^K — $C_{V_0,\pi\pi}(t_i, t, t_f, \vec{0}, \vec{0})$ and $C_{V_0,KK}(t_i, t, t_f, \vec{0}, \vec{0})$,
- The three-point functions determining $f_+^{K\pi}(0)$ through the vector and scalar current, as well as their time-reversed counterparts — $C_{\Gamma,K\pi}(t_i, t, t_f, \vec{0}, \vec{p}_\pi)$ and $C_{\Gamma,\pi K}(t_i, t, t_f, \vec{p}_\pi, \vec{0})$,

by minimising a single, global χ^2. Two-point functions are fit to the asymptotic limit of Equation (5.8), and three-point functions are fit to the asymptotic limit of Equation (5.9). In the second method we perform fits to the ratio (5.10), for example, by minimizing

$$\chi^2 = \sum_{\Delta t, V_\mu} \left(\frac{R_{V_\mu,K\pi}(\Delta t) - \frac{1}{Z_V}\left((p_K + p_\pi)_\mu \, f_+^{K\pi}(q^2) + (p_K - p_\pi)_\mu \, f_-^{K\pi}(q^2) \right)}{\sigma(\Delta t)} \right)^2 \tag{5.14}$$

using a range of separations Δt for which we observe a good signal. For even values of Δt we take the midpoint $t = \Delta t/2$ to evaluate $R_{V_\mu,K\pi}$, and for odd values of Δt we average the two values of $R_{V_\mu,K\pi}$ straddling the midpoint, since $\Delta t/2$ is not a point on the lattice. The sum over V_μ involves two components — the temporal component, and the average of the three spatial components — allowing us to fit both $f_+^{K\pi}(0)$ and $f_-^{K\pi}(0)$, although only the former is of interest for extracting $|V_{us}|$. One can similarly construct analogues of Equation (5.14) to fit Z_V^π, Z_V^K, and $f_+^{K\pi}(0)$ from the scalar matrix element.

For the fit to the scalar ratio (5.11) we observe a non-negligible sensitivity to the choice of fit range, presumably due to contamination from excited states. This is most pronounced on the C_{phys} ensemble, for which the plateau exhibits a small but clear upward slope. To account for this, we fit a model with includes an exponentially decaying excited state

$$f_+^{K\pi}(0, \Delta t) = f_+^{K\pi}(0) + Ce^{-m\Delta t}. \tag{5.15}$$

We find that this works well in practice: the value of $f_+^{K\pi}(0)$ from this exponential fit is consistent

with the value we get from fitting a constant, but is stable under variations of the fit range. The full set of fits is summarized in Figure 5.2 and Table 5.2.

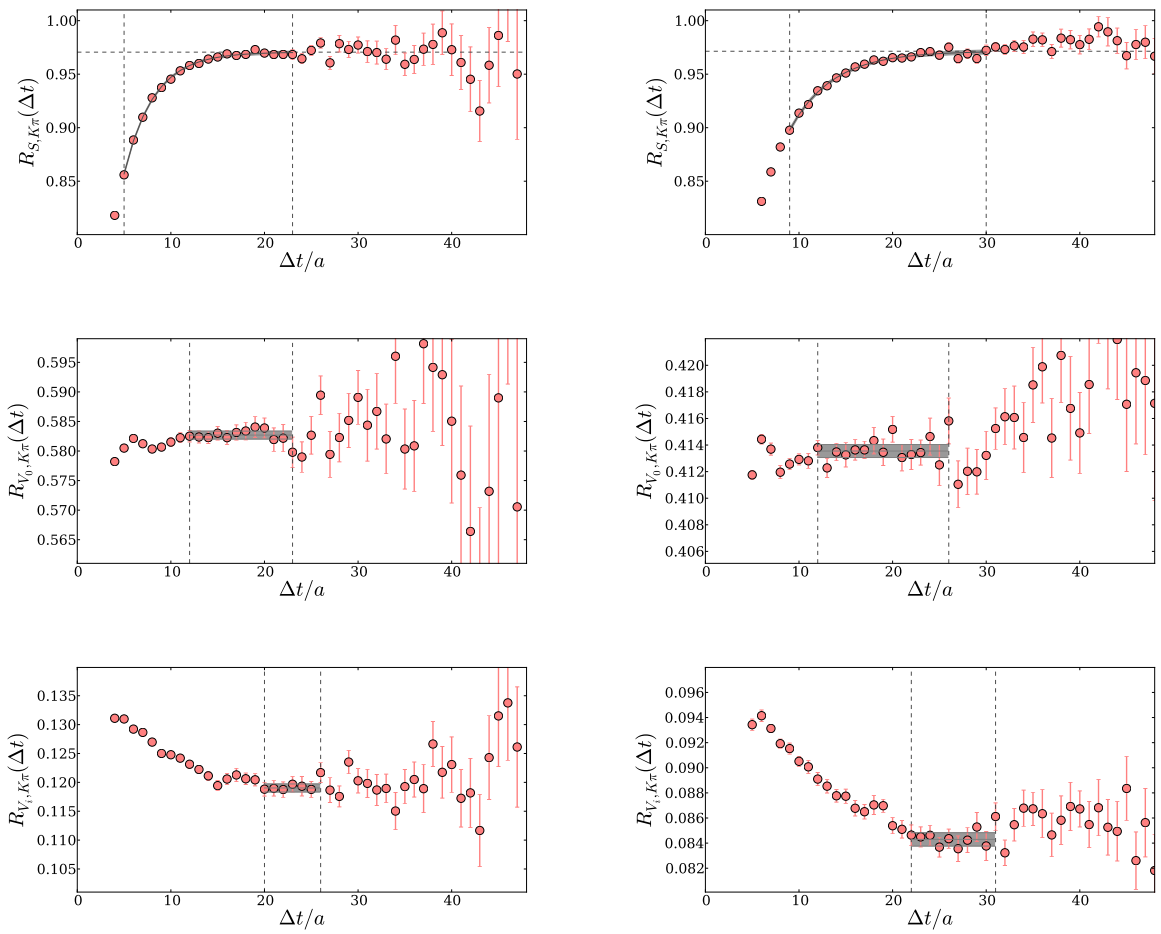

Figure 5.2: Fits to extract $f_+^{K\pi}(0)$ through the scalar density (top), temporal component of the vector current (middle), and average spatial component of the vector current (bottom) on the A_{phys} (left) and C_{phys} (right) ensembles. Vertical lines denote the choice of fit range, and the shaded bands denote the fit result and statistical uncertainty.

We find that we obtain consistent values for the form factor from either of the two analysis methods we have considered. The results we present in the rest of the paper, however, are based on the second method, for the following reasons: we observe as much as a factor of 5 difference in the statistical uncertainty in $Z_V^{\pi,K}$ between the ratio fit approach and the global fit approach; for

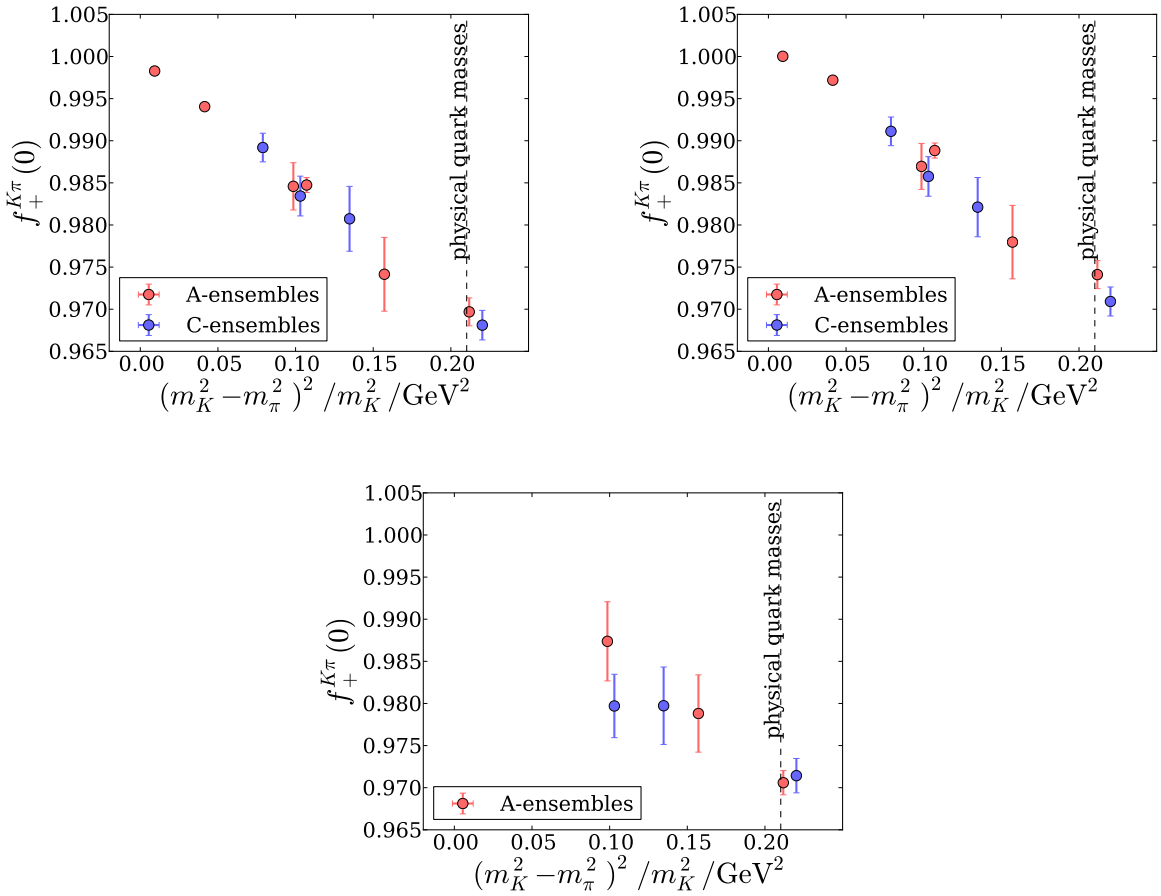

Figure 5.3: Simulation results for $f_+^{K\pi}(0)$ on each ensemble, measured through the vector matrix element renormalized by Z_V^π (upper left) and Z_V^K (upper right), as well as the scalar matrix element (bottom).

this particular quantity the ratio (5.13) is clearly superior. We argue that this can be understood by noting that the measurement of $Z_V^{\pi,K}$ through the ratio is less contaminated by excited states and thus the operator can be placed closer to the source/sink, leading to reduced statistical errors. We also observe that the three-point functions $C_{\Gamma,K\pi}$ and $C_{\Gamma,\pi K}$ are not symmetric between the source and sink walls, since the initial and final states are different, making it *a priori* difficult to decide on sensible fit ranges for extracting the form factors.

Ensemble	am_π	am_K	Z_V^π	Z_V^K	$f_+^{K\pi}(0)$ $Z_V^\pi V_\mu$	$Z_V^K V_\mu$	S
A_3	0.38840(39)	0.41628(39)	0.716106(77)	0.717358(75)	0.998289(79)	1.000033(80)	—
A_2	0.32231(47)	0.38438(46)	0.71499(12)	0.717252(93)	0.99404(29)	0.99719(28)	—
A_1	0.24157(38)	0.35009(39)	0.71408(20)	0.717047(74)	0.98474(89)	0.98884(90)	—
A_5^4	0.19093(46)	0.33197(58)	0.71399(58)	0.71679(13)	0.9746(43)	0.9784(43)	0.9793(46)
A_5^3	0.19093(45)	0.29818(52)	0.71399(58)	0.71570(16)	0.9850(27)	0.9874(27)	0.9878(47)
C_8	0.17249(50)	0.24125(47)	0.74435(40)	0.74580(12)	0.9890(17)	0.9909(17)	—
C_6	0.15104(41)	0.23276(45)	0.74387(56)	0.74563(13)	0.9833(24)	0.9857(24)	0.9796(39)
C_4	0.12775(41)	0.22624(51)	0.74480(94)	0.74585(16)	0.9805(39)	0.9819(35)	0.9796(47)
A_{phys}	0.08046(11)	0.28856(14)	0.71081(14)	0.714051(20)	0.9703(16)	0.9747(16)	0.9712(14)
C_{phys}	0.059010(95)	0.21524(11)	0.742966(81)	0.745121(23)	0.9673(18)	0.9701(17)	0.97097(21)

Table 5.2: Simulation results in lattice units. We report three values for the vector form factor: the first two are obtained from the vector matrix element after renormalizing the vector current with Z_V^π or Z_V^K, and the third is obtained from the scalar matrix element.

5.3.3 Corrections to the Physical Point

The simulation results presented in the previous chapter contain implicit systematic errors: the dominant sources of error can be attributed to the finite lattice spacing, finite ensemble volume, and slight mistunings in the input parameters. In particular, the values for the twist angles and valence quark masses used on the physical point ensembles were based on estimates of the spectrum computed from a small number of configurations early in the data generation run. As a result, there is a small discrepancy between the simulated q^2, m_π, and m_K and the desired kinematical point constrained by experiment. We emphasize, however, that these discrepancies are small, and the corrections described here are smaller than the statistical errors we quote. In this section we discuss our method for performing the q^2 and mass corrections independently on the two ensemble sets A and C, resulting in two values of $f_+^{K\pi}(0)$ at different (finite) lattice spacings. We then perform a continuum extrapolation, which is described in the following section. Our final prediction of the $K_{\ell 3}$ form factor retains a finite volume systematic, which we estimate and include in our final error

budget.

To correct the momentum transfer to $q^2 = 0$ we fit a pole ansatz

$$f_+^{K\pi}(q^2) = \left(1 + \frac{q^2}{M^2}\right)^{-1} f_+^{K\pi}(0) \tag{5.16}$$

to $f_+^{K\pi}(q^2)$ computed at two different kinematical points: the value of $q^2 \approx 0$ corresponding to our choice of twist angle, and $q^2_{\max} = (m_K - m_\pi)^2$, corresponding to a pion and kaon at rest. We then determine $f_+^{K\pi}(0)$ from the fit, and use this corrected value for the form factor in all subsequent steps of the analysis. We know from our previous work that the momentum dependence of the form factor should be well described by a pole ansatz [4]. We have also performed this correction using a linear ansatz

$$f_+^{K\pi}(q^2) = f_+^{K\pi}(0) + \alpha q^2 \tag{5.17}$$

as a cross-check, and find that the change in the extrapolated form factor is much smaller than the statistical error, indicating that the systematic error associated with the choice of q^2 parametrization is negligible in this study.

Likewise, we must make a small correction to $f_+^{K\pi}(0)$ from the simulated values of m_π and m_K to the physical values $m_{\pi^-} = 139.6\,\text{MeV}$ and $m_{K^0} = 497.6\,\text{MeV}$ [22]. We have considered a number of ansätze for performing this correction, which we divide into families $\mathcal{A}$, $\mathcal{B}$, $\mathcal{E}$, and $\mathcal{F}$:

$$\begin{cases} \text{fit } \mathcal{A}: \; f_+^{K\pi}(q^2 = 0, m_\pi^2, m_K^2) = 1 + f_2(f, m_\pi^2, m_K^2) \\ \text{fit } \mathcal{B}: \; f_+^{K\pi}(q^2 = 0, m_\pi^2, m_K^2) = 1 + f_2(f, m_\pi^2, m_K^2) + A_1 \left(m_K^2 + m_\pi^2\right) \left(m_K^2 - m_\pi^2\right)^2 \\ \text{fit } \mathcal{E}: \; f_+^{K\pi}(q^2 = 0, m_\pi^2, m_K^2) = A + A_0 \Delta M^2 \\ \text{fit } \mathcal{F}: \; f_+^{K\pi}(q^2 = 0, m_\pi^2, m_K^2) = A + A_0 \Delta M^2 + A_1 \left(m_K^2 + m_\pi^2\right) \Delta M^2 \end{cases} . \tag{5.18}$$

Here $\Delta M^2 \equiv \left(m_K^2 - m_\pi^2\right)^2 / m_K^2$ is an $SU(3)$-breaking polynomial motivated by the Ademollo-Gatto theorem [23] — which states that, near the $SU(3)$-symmetric limit $m_l = m_s$, the leading corrections to the form factor are $\mathcal{O}\left((m_s - m_l)^2\right)$ — and f_2 is the next-to leading order (NLO) term from $SU(3)$ chiral perturbation theory (χPT) [6]:

$$\begin{cases} f_2(f, m_\pi^2, m_K^2, m_\eta^2) = \dfrac{3}{2} H(f, m_\pi^2, m_K^2) + \dfrac{3}{2} H(f, m_\eta^2, m_K^2) \\ H(f, m_P^2, m_Q^2) = -\dfrac{1}{64\pi^2 f^2} \left(m_P^2 + m_Q^2 + \dfrac{2\, m_P^2\, m_Q^2}{m_P^2 - m_Q^2} \log\left(\dfrac{m_Q^2}{m_P^2}\right) \right) \end{cases} . \tag{5.19}$$

The $SU(3)$ χPT expression is parametrized entirely by the decay constant f at NLO; we use the tree-level relation $m_\eta^2 = (4m_K^2 - m_\pi^2)/3$ to evaluate Equation (5.19), which is consistent with expanding $f_+^{K\pi}(0)$ to NLO. The same set of ansätze have also been studied in our last $K_{\ell 3}$ calculation [6].

We have first considered fits to families $\mathcal{A}$ and $\mathcal{B}$, which are both parametrized by the $SU(3)$ decay constant in the chiral limit f, and, for family $\mathcal{B}$, the coefficient A_1, which models the analytic part of the NNLO corrections. Sample fits of this type to $f_+^{K\pi}(0)$ determined through the vector current matrix element renormalized by Z_V^π on the A series of ensembles are plotted in Figure 6. We observe that next-to leading order $SU(3)$ χPT is a poor description of our data, leading to a

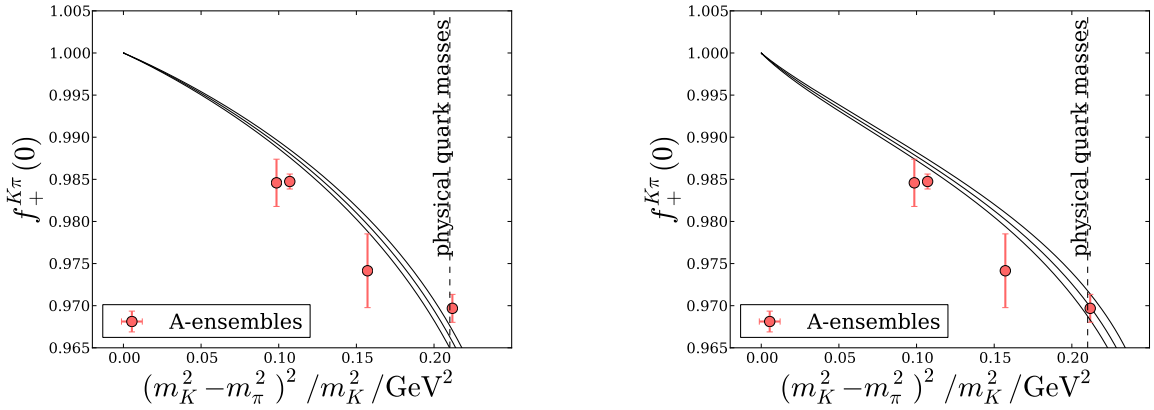

Figure 5.4: Sample fits of $f_+^{K\pi}(0)$, determined through the vector current matrix element renormalized by Z_V^π, to ansatz $\mathcal{A}$ (left) and $\mathcal{B}$ (right) using data from the A ensembles with a mass cut $m_\pi \lesssim 450$ MeV.

fit with $\chi^2/\mathrm{dof} = 4.6$ (left panel). Adding a model for the next-to-next-to leading order analytic terms improves the fit somewhat — $\chi^2/\mathrm{dof} = 1.7$ (right panel) — but is still unacceptably poor. We attribute this to the fact that $SU(3)$ χPT is an effective field theory, valid near the chiral limit $m_l = m_s = 0$, and must break down for sufficiently heavy quark masses. On all but the most recent A_{phys} and C_{phys} ensembles the simulated pion and kaon are heavier than their physical counterparts, and so this data is likely outside the region of validity of the next-to leading order chiral expansion. Since we lack sufficient data to constrain the full NNLO χPT expression for the form factor, we have discarded ansätze $\mathcal{A}$ and $\mathcal{B}$ in favor of $\mathcal{E}$ and $\mathcal{F}$.

Fits $\mathcal{E}$ and $\mathcal{F}$ are motivated by the Ademollo-Gatto theorem, which suggests that ΔM^2 should parametrize the leading $SU(3)$ breaking effects near the $SU(3)$ symmetric limit $m_l = m_s$. We observe in Figure 5.3 that this parametrization actually describes the full range of simulations remarkably well, suggesting that higher order contributions are small even at the physical point. Fit $\mathcal{E}$ is linear in ΔM^2, while fit $\mathcal{F}$ contains an additional polynomial term modeling the next-to leading order $SU(3)$ breaking effects. Representative fits of these types are summarized in Tables 5.3 and 5.4, as well as in Figure 5.5. These fits are generally of very good quality, with small

ME	m_π^{cut} (MeV)	$f_+^{K\pi}(0)^A$	$f_+^{K\pi}(0)^C$	A^A	A^C	A_0^A (GeV^2)	A_0^C (GeV^2)	χ^2/dof
$Z_V^\pi V_\mu$	355	0.9703(16)	0.9689(16)	0.9970(52)	1.001(10)	-0.127(28)	-0.155(51)	0.37
	450	0.9704(16)	0.9687(16)	0.9994(23)	1.0002(26)	-0.138(17)	-0.150(17)	0.28
	600	0.9701(13)	0.9687(16)	0.99990(49)	1.0002(26)	-0.1416(80)	-0.150(17)	0.24
	700	0.97071(99)	0.9687(16)	0.999568(96)	1.0002(26)	-0.1373(49)	-0.150(17)	0.28
$Z_V^K V_\mu$	355	0.9748(16)	0.9715(16)	0.9977(50)	1.0005(97)	-0.109(27)	-0.138(47)	0.28
	450	0.9748(16)	0.9714(16)	1.0027(24)	1.0017(25)	-0.133(17)	-0.144(16)	0.47
	600	0.9748(13)	0.9714(16)	1.00269(48)	1.0017(25)	-0.1327(78)	-0.144(16)	0.38
	700	0.97747(100)	0.9714(16)	1.001120(96)	1.0017(25)	-0.1125(50)	-0.144(16)	2.05
S	355	0.9715(14)	0.9717(19)	1.0022(80)	0.994(13)	-0.146(40)	-0.105(60)	0.00
	450	0.9715(14)	0.9716(19)	1.0022(80)	0.9890(68)	-0.146(40)	-0.083(35)	0.10

Table 5.3: Results for global fit $\mathcal{E}$ on ensembles A and C with a variety of pion mass cuts. The first column indicates the method used to extract the form factor (vector matrix element (ME), renormalized with Z_V^π or Z_V^K, or scalar ME). A superscript A or C in the top line denotes the ensemble set associated with each fit parameter.

values of χ^2/dof. The only evidence of non-negligible corrections beyond the leading ΔM^2 term parametrizing ansatz $\mathcal{E}$ is observed in fits to the vector matrix element data renormalized by Z_V^K which include the heaviest data near the $SU(3)$ symmetric limit. In this limit we observe a sharp rise in χ^2/dof for fit $\mathcal{E}$, and a significant improvement by adding the additional, next-to leading term in ansatz $\mathcal{F}$. We prefer ansatz $\mathcal{E}$ with a reasonable pion mass cut over ansatz $\mathcal{F}$, however, since

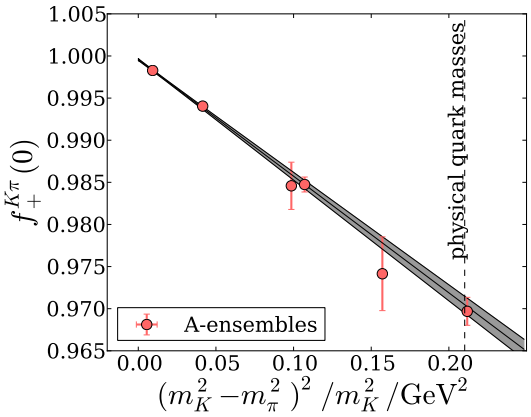

(a) Fit $\mathcal{E}$, A ensembles

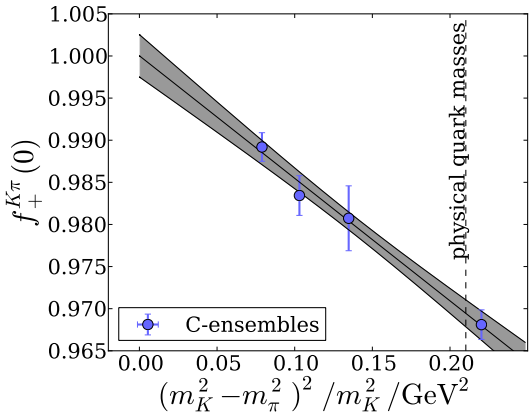

(b) Fit $\mathcal{E}$, C ensembles

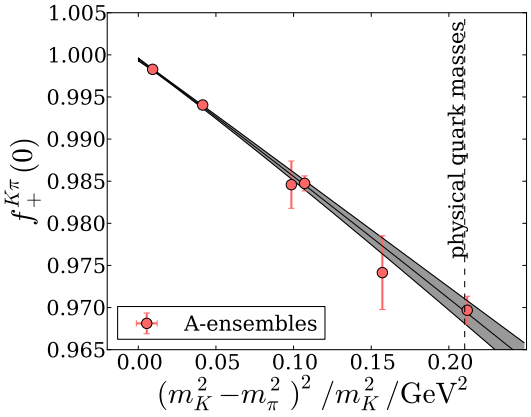

(c) Fit $\mathcal{F}$, A ensembles

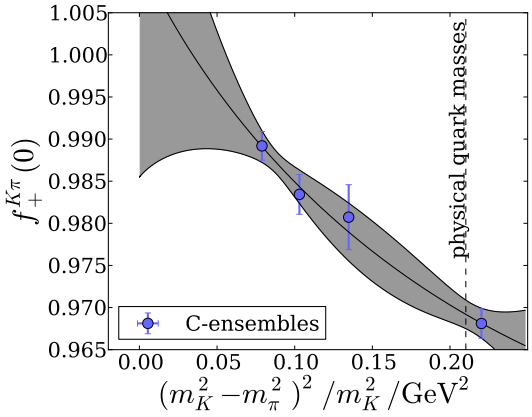

(d) Fit $\mathcal{F}$, C ensembles

Figure 5.5: Representative fits of $f_+^{K\pi}(0)$, measured through the vector matrix element and renormalized by Z_V^π, to ansätze $\mathcal{E}$ and $\mathcal{F}$.

ME	$m_\pi^{\rm cut}$ (MeV)	$f_+^{K\pi}(0)^A$	$f_+^{K\pi}(0)^C$	A^A	A^C	A_0^A (GeV^2)	A_0^C (GeV^2)	A_1^A (GeV^4)	A_1^C (GeV^4)	χ^2/dof
$Z_V^\pi V_\mu$	450	0.9697(16)	0.9710(38)	0.9990(77)	1.012(26)	-0.143(17)	-0.08(14)	0.02(12)	-0.41(92)	0.59
	600	0.9697(16)	0.9710(38)	1.0000(15)	1.012(26)	-0.145(15)	-0.08(14)	0.001(56)	-0.41(92)	0.40
	700	0.9695(15)	0.9710(38)	0.99946(16)	1.012(26)	-0.148(12)	-0.08(14)	0.020(23)	-0.41(92)	0.33
	800	0.9695(15)	0.9710(38)	0.99946(16)	1.012(26)	-0.148(12)	-0.08(14)	0.020(23)	-0.41(92)	0.33
$Z_V^K V_\mu$	450	0.9740(16)	0.9740(37)	0.9986(75)	1.015(25)	-0.138(18)	-0.07(13)	0.08(12)	-0.49(88)	0.80
	600	0.9739(16)	0.9740(37)	1.0021(15)	1.015(25)	-0.142(15)	-0.07(13)	0.027(57)	-0.49(88)	0.61
	700	0.9734(15)	0.9740(37)	1.00067(16)	1.015(25)	-0.151(12)	-0.07(13)	0.079(23)	-0.49(88)	0.72

Table 5.4: Results for global fit $\mathcal{F}$ on ensembles A and C with a variety of pion mass cuts. We do not report results for the form factor determined through the scalar matrix element since we have too little data to constrain ansatz $\mathcal{F}$ reliably.

we lack sufficient data to reliably fit ansatz $\mathcal{F}$ to data extracted from the scalar matrix element. We observe in Table 5.3 that the statistical error on the value of $f_+^{K\pi}(0)$ after interpolating to the physical mass point is essentially constant for $m_\pi^{\rm cut} \lesssim 600$ MeV, indicating that, as long as we cut the heaviest data near the $SU(3)$ symmetric limit, the fits are dominated by the $A_{\rm phys}$ and $C_{\rm phys}$ ensembles, with data from previous $K_{\ell 3}$ calculations simply determining the slope of the small interpolation we perform.

A second observation we note in Table 5.3 is that the values of the slope parameters A_0^A and A_0^C are generally consistent within the quoted statistical errors. We might expect this since domain wall fermions are $\mathcal{O}(a)$ improved, implying that cut-off effects enter as

$$A_0(a) = A_0(0)\left(1 + \alpha\left(a\Lambda_{\rm QCD}\right)^2 + \cdots\right). \tag{5.20}$$

We can estimate using $\Lambda_{\rm QCD} \sim 200 - 300$ MeV that the difference between A_0^A and A_0^C should be $\mathcal{O}(1\%)$, which is indeed smaller than the measured statistical uncertainties. We have also performed fits to determine α directly and verified that it is consistent with zero. Motivated by these observations, we have performed a second series of fits using a modified ansatz $\mathcal{E}$ with the additional constraint $A_0^A = A_0^C \equiv A_0$. These fits are summarized in Table 5.5 and Figure 5.6. We again observe high quality fits with small χ^2/dof for all but the most aggressive pion mass cut

ME	$m_\pi^{\rm cut}$ (MeV)	$f_+^{K\pi}(0)^A$	$f_+^{K\pi}(0)^C$	A^A	A^C	A_0 (GeV^2)	χ^2/dof
$Z_V^\pi V_\mu$	355	0.9701(15)	0.9690(16)	0.9982(45)	0.9970(52)	-0.134(24)	0.31
	450	0.9699(12)	0.9691(14)	1.0002(17)	0.9994(19)	-0.144(12)	0.28
	600	0.9699(12)	0.9692(12)	0.99999(46)	0.9993(15)	-0.1431(72)	0.23
	700	0.97050(95)	0.9696(11)	0.999581(96)	0.9986(13)	-0.1383(48)	0.31
$Z_V^K V_\mu$	355	0.9745(15)	0.9717(15)	0.9991(43)	0.9962(50)	-0.117(23)	0.28
	450	0.9743(12)	0.9718(14)	1.0035(17)	1.0009(20)	-0.139(12)	0.42
	600	0.9745(12)	0.9721(12)	1.00280(44)	1.0004(15)	-0.1348(71)	0.38
	700	0.97696(97)	0.9735(11)	1.001154(96)	0.9977(13)	-0.1151(48)	2.23
S	355	0.9716(14)	0.9716(19)	0.9997(68)	0.9997(72)	-0.133(34)	0.17
	450	0.9719(13)	0.9710(18)	0.9949(53)	0.9940(53)	-0.109(26)	0.55

Table 5.5: Results for global fit $\mathcal{E}$ to a combined data set including both A and C ensembles and the constraint $A_0^A = A_0^C \equiv A_0$.

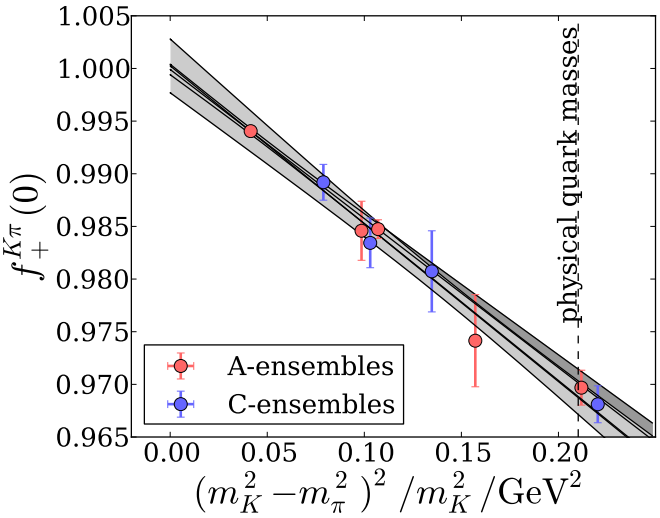

Figure 5.6: Representative fit $\mathcal{E}$ to all data for the form factor measured through the vector matrix element and renormalized with Z_V^π. The slope parameter A_0 is constrained to be equal for the A ensembles and C ensembles.

of 700 MeV. We also observe that the statistical errors in the interpolated $f_+^{K\pi}(0)$ have a stronger dependence on the pion mass cut for these fits. We attribute this to the heaviest data points near the $SU(3)$ symmetric limit, which have smaller statistical errors, and have a more pronounced effect when the slope parameters are constrained than when they are allowed to vary between the two ensemble sets.

5.3.4 Continuum Extrapolation

After performing the corrections described in section 5.3.3 we are left with three values of $f_+^{K\pi}(0)$ — two independent normalizations of the vector current matrix element 5.10 using Z_V^π and Z_V^K, as well as the scalar matrix element 5.11 — at each of the two lattice spacings considered in our analysis. While we could perform three independent linear extrapolations to the continuum, we instead chose to impose universality, and extrapolate all three calculations of the form factor to a common continuum limit. This is summarized in Table 5.6 and Figure 5.7. Repeating the full analysis with different choices of the mass cut does not change the result within the statistical error: we choose $m_\pi \leq 450\,\text{MeV}$ as a reasonable mass cut to quote a final result, and conservatively choose not to constrain the slope parameters A_0 for the A and C ensembles, obtaining a continuum limit of $f_+^{K\pi}(0) = 0.9685(34)$. The quoted error is purely statistical, and includes the uncertainty in the lattice spacings for the A and C ensembles.

m_π^{cut} (MeV)	355	450	600
Global fit $\mathcal{E}$	0.9687(35)	0.9685(34)	0.9685(34)
Global fit $\mathcal{E}$, A_0 constrained	0.9690(33)	0.9689(25)	0.9691(22)
Global fit $\mathcal{F}$, A_1 constrained	—	0.9683(35)	0.9685(34)
Global fit $\mathcal{F}$, A_0, A_1 constrained	0.9694(34)	0.9687(26)	0.9690(22)

Table 5.6: Continuum limit results for the form factor $f_+^{K\pi}(0)$ based on variants of fits $\mathcal{E}$ and $\mathcal{F}$ for the interpolation to physical masses.

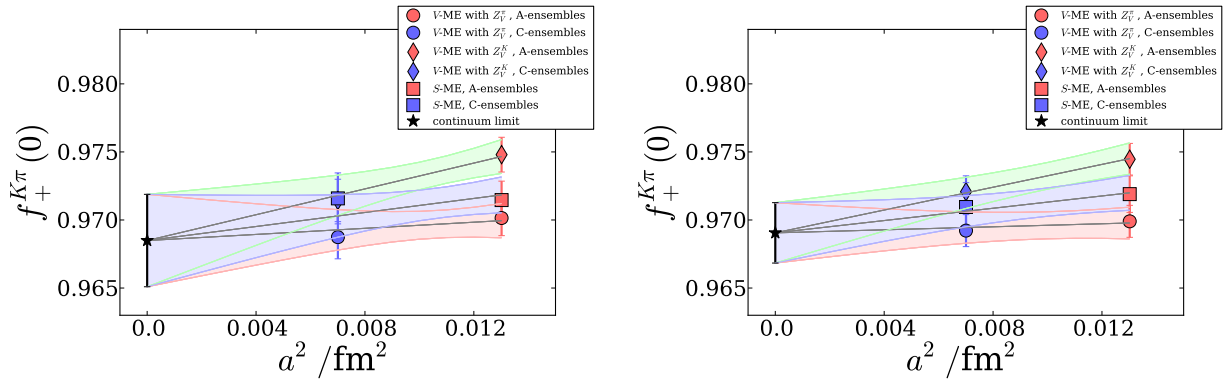

Figure 5.7: Joint extrapolation to a common continuum limit of data obtained from mass interpolation fit $\mathcal{E}$ with a pion mass cut of 600 MeV. In the left plot the slope parameter A_0 is allowed to differ between the A ensembles and C ensembles. In the right plot we have imposed $A_0^A = A_0^C$, obtaining a consistent central value but reduced statistical uncertainty for the final result.

5.3.5 Final Result and Error Budget

After interpolating in q^2, m_π, and m_K to the physical point, and extrapolating to the continuum, we are left with the following systematic errors:

- *Finite Volume (FV)*: Since the $K \to \pi$ matrix element contains only single particle initial and final states, we expected finite volume effects to be exponentially suppressed in $m_\pi L$. We naively estimate these effects to be of order $(1 - f_+^{K\pi}(0))e^{-m_\pi L} = 0.0007$. χPT [24] estimates an error approximately twice as large for the A_{phys} and C_{phys} ensembles, but does not completely describe our calculation since it does not include the effects of imposing twisted boundary conditions on the valence light quark fields. Thus, we quote twice our naive error — 0.0014 — as our value for the FV systematic.

- *Partial Quenching*: The calculations on ensemble A_5^3 and the full C ensemble set were performed with a partially quenched strange quark. We expect any associated systematic errors to be small compared to the other errors we quote, and we have explicitly checked that excluding ensemble A_5^3 from the analysis does not change the result. On the C ensembles the relative difference between the sea and valence strange quark masses are even smaller than for

the A_5^3 ensemble, so we conclude that partial quenching systematics are negligible compared to our other sources of error.

- *Isospin Breaking*: The unitary light quarks in our study are isospin symmetric, unlike the physical up and down quarks. We partially correct for this in our final result by interpolating in the valence sector to the physical π^- and K^0 masses, but this leaves unquantified systematic errors associated with the isospin symmetric light sea quarks. We again expect this systematic to be negligible compared to other sources of error. Techniques to address sea isospin breaking effects in future calculations are currently being developed [25–28].

Taking these into account, we obtain our final result

$$f_+^{K\pi}(0) = 0.9685(34)_{\text{stat}}(14)_{\text{FV}}, \tag{5.21}$$

leading to the prediction

$$|V_{us}| = 0.2233(5)_{\text{experiment}}(9)_{\text{lattice}}. \tag{5.22}$$

Together with $|V_{ud}| = 0.97425(22)$ from super-allowed nuclear β-decay measurements [22], and neglecting $|V_{ub}| \approx 10^{-3}$, we observe a $\sim 1.5\sigma$ tension with first-row CKM unitarity

$$1 - |V_{ud}|^2 - |V_{us}|^2 = 0.0010(4)_{V_{ud}}(2)_{V_{us}^{\text{exp}}}(4)_{V_{us}^{\text{lat}}} = 0.0010(6). \tag{5.23}$$

5.4 Conclusion

In this work we have presented a first-principles calculation of the kaon semileptonic form factor with vanishing momentum transfer and physical light quark masses in domain wall QCD. We have also demonstrated how to utilize our older, heavy pion mass ensembles to correct for slight mistunings and extrapolate to the continuum. We find that the previously dominant systematic errors associated with chiral extrapolations of simulations with heavy quark masses are no longer present, and that the small interpolation to the physical π^- and K^0 masses we perform is indeed dominated by the results on the new A_{phys} and C_{phys} ensembles. While we observed that we could reduce the statistical error in our final prediction by as much as 30% if we include data from heavy

simulations near the $SU(3)$ symmetric point and assume the influence of cut-off effects on the fit parameters is sub-statistical, we have chosen not to make these assumptions in our final results, since they increase the model dependence in ways that are difficult to quantify. We now find that finite volume errors are the dominant source of systematic error.

While a domain wall QCD calculation of the $K_{\ell 3}$ form factor $f_+^{K\pi}(0)$ with physical quark masses and in the continuum limit represents a significant step forward, there are still a number of potential improvements we hope to address in the future. As we have argued above, the dominant source of systematic error is now the finite volume error; a better understanding of the size of these effects in partially twisted chiral perturbation theory, or from a series of lattice simulations which vary the simulation volume while leaving all other details fixed, is desirable. In addition, the analysis in Ref. [12] leading to the experimental constraint $|V_{us}|f_+^{K\pi}(0) = 0.2163(5)$ accounts for electromagnetic and isospin breaking effects in the framework of chiral perturbation theory. Given the current 0.4% overall error in the lattice prediction for the $K_{\ell 3}$ form factor, we have now reached a point where it may be worthwhile to consider how to treat these effects non-perturbatively within the framework of lattice field theory. Progress in this direction is discussed in Refs. [25–28].

References

[1] N. Cabibbo, "Unitary Symmetry and Leptonic Decays", Phys. Rev. Lett. **10**, [,648(1963)], 531–533 (1963).

[2] M. Kobayashi and T. Maskawa, "CP Violation in the Renormalizable Theory of Weak Interaction", Prog. Theor. Phys. **49**, 652–657 (1973).

[3] A. Bazavov et al., "Determination of $|V_{us}|$ from a Lattice-QCD Calculation of the $K \to \pi\ell\nu$ Semileptonic Form Factor with Physical Quark Masses", Phys. Rev. Lett. **112**, 112001 (2014).

[4] P. A. Boyle, A. Juttner, R. D. Kenway, C. T. Sachrajda, S. Sasaki, A. Soni, R. J. Tweedie, and J. M. Zanotti, "K(l3) Semileptonic Form-Factor from 2+1 Flavour Lattice QCD", Phys. Rev. Lett. **100**, 141601 (2008).

[5] P. A. Boyle, J. M. Flynn, A. Juttner, C. Kelly, C. Maynard, H. Pedroso de Lima, C. T. Sachrajda, and J. M. Zanotti, "$K \to \pi$ Form Factors with Reduced Model Dependence", Eur. Phys. J. **C69**, 159–167 (2010).

[6] P. A. Boyle, J. M. Flynn, N. Garron, A. Jüttner, C. T. Sachrajda, K. Sivalingam, and J. M. Zanotti, "The Kaon Semileptonic Form Factor with Near Physical Domain Wall Quarks", JHEP **08**, 132 (2013).

[7] A. Bazavov et al., "Kaon Semileptonic Vector Form Factor and Determination of $|V_{us}|$ Using Staggered Fermions", Phys. Rev. **D87**, 073012 (2013).

[8] V. Lubicz, F. Mescia, S. Simula, and C. Tarantino, "$K \to \pi l \nu$ Semileptonic Form Factors from Two-Flavor Lattice QCD", Phys. Rev. **D80**, 111502 (2009).

[9] G. Colangelo et al., "Review of Lattice Results Concerning Low Energy Particle Physics", Eur. Phys. J. **C71**, 1695 (2011).

[10] S. Aoki et al., "Review of Lattice Results Concerning Low-Energy Particle Physics", Eur. Phys. J. **C74**, 2890 (2014).

[11] G. Amelino-Camelia et al., "Physics with the KLOE-2 Experiment at the Upgraded DAϕNE", Eur. Phys. J. **C68**, 619–681 (2010).

[12] M. Antonelli et al., "An Evaluation of $|V_{us}|$ and Precise Tests of the Standard Model from World Data on Leptonic and Semileptonic Kaon Decays", Eur. Phys. J. **C69**, 399–424 (2010).

[13] J. M. Flynn, A. Juttner, C. T. Sachrajda, P. A. Boyle, and J. M. Zanotti, "Hadronic Form Factors in Lattice QCD at Small and Vanishing Momentum Transfer", JHEP **05**, 016 (2007).

[14] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[15] D. B. Kaplan, "A Method for Simulating Chiral Fermions on the Lattice", Phys. Lett. **B288**, 342–347 (1992).

[16] Y. Shamir, "Chiral Fermions from Lattice Boundaries", Nucl. Phys. **B406**, 90–106 (1993).

[17] R. C. Brower, H. Neff, and K. Orginos, "The Möbius Domain Wall Fermion Algorithm", (2012).

[18] Y. Iwasaki and T. Yoshié, "Renormalization Group Improved Action for $SU(3)$ Lattice Gauge Theory and the String Tension", Physics Letters B **143**, 449–452 (1984).

[19] C. Bernard, J. Bijnens, E. Gámiz, and J. Relefors, "Twisted Finite-Volume Corrections to $K_{\ell 3}$ Decays with Partially-Quenched and Rooted-Staggered Quarks", JHEP **03**, 120 (2017).

[20] T. Blum, T. Izubuchi, and E. Shintani, "New Class of Variance-Reduction Techniques using Lattice Symmetries", Phys. Rev. **D88**, 094503 (2013).

[21] T. Blum et al., "$K \to \pi\pi$ $\Delta I = 3/2$ Decay Amplitude in the Continuum Limit", Phys. Rev. **D91**, 074502 (2015).

[22] K. A. Olive et al., "Review of Particle Physics", Chin. Phys. **C38**, 090001 (2014).

[23] J. F. Donoghue, E. Golowich, and B. R. Holstein, *Dynamics of the standard model*, Cambridge Books Online (Cambridge University Press, 1992).

[24] K. Ghorbani and H. Ghorbani, "Kaon Semi-Leptonic Form Factor at Zero Momentum Transfer in Finite Volume", Eur. Phys. J. **A49**, 134 (2013).

[25] G. M. de Divitiis et al., "Isospin Breaking Effects Due to the Up-Down Mass Difference in Lattice QCD", JHEP **04**, 124 (2012).

[26] N. Tantalo, "Lattice Calculation of Isospin Corrections to K_{l2} and K_{l3} Decays", in 7th International Workshop on the CKM Unitarity Triangle (CKM 2012) Cincinnati, Ohio, USA, September 28-October 2, 2012 (2013).

[27] A. Portelli, "Review on the Inclusion of Isospin Breaking Effects in Lattice Calculations", PoS **KAON13**, 023 (2013).

[28] N. Carrasco, V. Lubicz, G. Martinelli, C. T. Sachrajda, N. Tantalo, C. Tarantino, and M. Testa, "QED Corrections to Hadronic Processes in Lattice QCD", Phys. Rev. **D91**, 074506 (2015).

Chapter 6

The Exact One Flavor Algorithm

Domain Wall Fermion QCD with the Exact One Flavor Algorithm

C. Jung[1], C. Kelly[2], R.D. Mawhinney[2], and D.J. Murphy[2]

[1]Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA

[2]Department of Physics, Columbia University, New York, NY 10027, USA

Abstract

Lattice QCD calculations including the effects of one or more non-degenerate sea quark flavors are conventionally performed using the Rational Hybrid Monte Carlo (RHMC) algorithm, which computes the square root of the determinant of $\mathscr{D}^\dagger\mathscr{D}$, where $\mathscr{D}$ is the Dirac operator. The special case of two degenerate quark flavors with the same mass is described directly by the determinant of $\mathscr{D}^\dagger\mathscr{D}$ — in particular, no square root is necessary — enabling a variety of algorithmic developments, which have driven down the cost of simulating the light (up and down) quarks in the isospin-symmetric limit of equal masses. As a result, the relative cost of single quark flavors — such as the strange or charm — computed with RHMC has become more expensive. This problem is even more severe in the context of our measurements of the $\Delta I = 1/2$ $K \to \pi\pi$ matrix elements on lattice ensembles with G-parity boundary conditions,

since G-parity is associated with a doubling of the number of quark flavors described by $\mathscr{D}$, and thus RHMC is needed for the isospin-symmetric light quarks as well. In this paper we report on our implementation of the exact one flavor algorithm (EOFA) introduced by the TWQCD collaboration for simulations including single flavors of domain wall quarks. We have developed a new preconditioner for the EOFA Dirac equation, which both reduces the cost of solving the Dirac equation and allows us to re-use the bulk of our existing high-performance code. Coupling these improvements with careful tuning of our integrator, the time per accepted trajectory in the production of our 2+1 flavor G-parity ensembles with physical pion and kaon masses has been decreased by a factor of 4.2.

6.1 Introduction

Lattice QCD simulations are typically performed using variants of the hybrid Monte Carlo (HMC) algorithm, which includes the effects of dynamical sea quarks through the determinant of a fermion matrix evaluated by stochastically sampling a discretized QCD path integral. Conventional simulations choose the Hermitian fermion matrix $\mathcal{M} = \mathscr{D}^\dagger\mathscr{D}$ rather than the lattice Dirac operator $\mathcal{M} = \mathscr{D}$, since the latter, in general, has a complex spectrum, and is thus less amenable to standard numerical algorithms. While $\mathscr{D}$ describes a single quark flavor, $\mathscr{D}^\dagger\mathscr{D}$ describes two degenerate quark flavors with the same mass. As a result the standard HMC algorithm naturally describes the light (up and down) quarks in the isospin-symmetric limit $m_u = m_d$ considered in most lattice calculations. Simulations including single quark flavors (such as the strange or charm) are typically performed by taking an overall square root of the determinant of $\mathcal{M} = \mathscr{D}^\dagger\mathscr{D}$, leading to the rational hybrid Monte Carlo (RHMC) algorithm. While RHMC has found widespread usage in the lattice QCD community, RHMC calculations are typically more expensive than HMC calculations for the same input quark mass, in part because many of the techniques which have been developed to accelerate HMC simulations of degenerate quark flavor pairs are not applicable to RHMC.

A number of recent developments in the HMC algorithm used by the RBC/UKQCD collaboration have driven down the cost of simulating degenerate pairs of isospin-symmetric quarks with the same mass. These developments include: extensive force tuning via Hasenbush mass preconditioning [1], the zMöbius domain wall fermion action [2], reduced L_s approximations to the light quark

determinant [2], and the use of implicitly restarted, mixed-precision defect correction methods in the conjugate gradient algorithm [1]. In Table 6.1 we list timings for a recent large-scale calculation which utilizes these techniques. We now find that the single-flavor strange and charm quark deter-

Action Component	**Timings**	
Gauge	5970 s	12.0%
Light Quarks	19600 s	39.4%
Strange and Charm Quarks	24200 s	48.6%
Total	49770 s	—

Table 6.1: Timings for one HMC trajectory of RBC/UKQCD's $80^2 \times 96 \times 192 \times 32$ $N_f = 2+1+1$ ensemble with physical quark masses and $a^{-1} \approx 3$ GeV on a 12,288-node Blue Gene/Q partition [2].

minants, which we simulate using the RHMC algorithm, are collectively the most expensive part of the calculation. To address this, we have turned to exploring TWQCD's recently proposed exact one flavor algorithm (EOFA), which allows for simulating single quark flavors without the need for RHMC [3]. Preliminary results have suggested that EOFA simulations can outperform RHMC simulations, both in terms of computer time and a reduced memory footprint, while producing exactly the same physics [4, 5].

The RBC/UKQCD collaboration's ongoing efforts to probe direct CP-violation in $K \to \pi\pi$ decays provide a second motivation for exploring EOFA. The collaboration has recently reported the first calculation of the $\Delta I = 1/2$ $K \to \pi\pi$ decay amplitude with physical kinematics in Ref. [6], which, when combined with previous results for the $\Delta I = 3/2$ amplitude [7] determines the Standard Model CP-violating parameters ϵ and ϵ' entirely from first principles. An important ingredient in this calculation was the introduction of G-parity boundary conditions for the quark fields [6, 8]: since the pion is G-parity odd, the pion momenta are quantized along G-parity directions as

$$p_\pi^i = \frac{(2n_i + 1)\,\pi}{L}, \quad n_i \in \mathbb{Z}, \tag{6.1}$$

[1] We elaborate on the details of our defect correction solver in Section 6.6.

allowing the ensemble parameters to be tuned such that the $K \to \pi\pi$ decay has both physical kinematics and the final pions in the ground state. Since the G-parity transformation $G = Ce^{i\pi I_y}$ is the product of charge conjugation and a 180° isospin rotation about the y-axis — at the lattice boundary the light quark doublet transforms as $(u, d) \mapsto (\overline{d}, -\overline{u})$ — the G-parity Dirac operator inherently describes two quark flavors. The standard lattice technique for obtaining a Hermitian, positive-definite fermion matrix — by taking the square of the Dirac operator, $\mathcal{M} = \mathscr{D}^\dagger \mathscr{D}$ — results in a theory with four degenerate quark flavors on a G-parity ensemble, and a square root is required to reduce to a two-flavor simulation. Describing the light quark pair on a G-parity ensemble is a particularly attractive target for EOFA, since many of the techniques we use to accelerate the calculation of the light quark determinant for ensembles with periodic boundary conditions — including defect correction solvers, the forecasted force gradient integrator [9], and Hasenbusch mass preconditioning — are not applicable or of limited utility for RHMC simulations, but are expected to perform well in the context of EOFA. More generally, since there is no straightforward way to start the multishift conjugate gradient solver used for RHMC with a nonzero initial guess, techniques which rely on forecasting or restarting the solver are not applicable.

In this work we discuss the RBC/UKQCD collaboration's implementation and tests of the exact one flavor algorithm, as well as the use of EOFA in generating gauge field configurations for our ongoing first-principles calculation of the ratio of Standard Model parameters ϵ'/ϵ from $\Delta I = 1/2$ $K \to \pi\pi$ decays with G-parity boundary conditions. We have independently implemented EOFA in the Columbia Physics System (CPS), BAGEL fermion sparse matrix library (BFM), and the Grid data parallel C++ QCD library (Grid), for Shamir and Möbius domain wall fermions, with periodic, anti-periodic, and G-parity boundary conditions. We will demonstrate in the following sections that a significant improvement over the RHMC algorithm in terms of wall clock time is indeed possible with EOFA after introducing a variety of preconditioning and tuning techniques. Early work in this direction was presented at the 34th International Symposium on Lattice Field Theory [5]; here we will elaborate on the details and discuss our first large-scale EOFA calculation.

6.2 The Exact One Flavor Algorithm

The exact one flavor algorithm was developed by the TWQCD collaboration and used to enable efficient simulations of single quark flavors on GPU clusters, where memory usage is a significant constraint. In Ref. [10] the authors discuss their construction of a positive-definite pseudofermion action describing a single flavor of Wilson or domain wall quark, and elaborate on the details of this construction in Ref. [3]. The key is their observation that a ratio of determinants of domain wall fermion (DWF) Dirac operators can be factorized as

$$\det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m_2)}\right) = \frac{1}{\det\left(\mathcal{M}_L\right)} \cdot \frac{1}{\det\left(\mathcal{M}_R\right)}, \tag{6.2}$$

with $\mathcal{M}_L$ and $\mathcal{M}_R$ Hermitian and positive-definite. In a subsequent paper the authors benchmark EOFA against RHMC for $N_f = 1$ and $N_f = 2 + 1$ lattice QCD simulations, and demonstrate a number of advantages of the EOFA formalism [4]. These include substantial reductions in the pseudofermion force and in the memory footprint of the algorithm, since, in the context of EOFA, inversions of the Dirac operator can be performed using the ordinary conjugate gradient (CG) algorithm rather than the multishift CG used for RHMC. They ultimately find that they are able to generate HMC trajectories 15-20% faster using EOFA rather than RHMC after retuning their integration scheme to take advantage of these properties.

We note that the construction of the exact one flavor pseudofermion action has been detailed by TWQCD in Ref. [3, 10] and summarized in our own formalism in Ref. [5]. We will not repeat this discussion here, other than to give a brief overview and to introduce the notation used in this work. We write the 5D Möbius domain wall fermion (MDWF) operator $\mathscr{D}_{\text{DWF}}$ in terms of the 4D Wilson Dirac operator D_W and 5D hopping matrix $L_{ss'}$ as

$$\begin{aligned}
\left(\mathscr{D}_{\text{DWF}}\right)_{xx',ss'} &= \left(\left(c+d\right)\left(D_W\right)_{xx'} + \delta_{xx'}\right)\delta_{ss'} + \left(\left(c-d\right)\left(D_W\right)_{xx'} - \delta_{xx'}\right)L_{ss'} \\
\left(D_W\right)_{xx'} &= \left(4 + M_5\right)\delta_{xx'} - \frac{1}{2}\sum_{\mu}\left[\left(1-\gamma_\mu\right)U_\mu(x)\delta_{x+\hat{\mu},x'} + \left(1+\gamma_\mu\right)U_\mu^\dagger(x')\delta_{x-\hat{\mu},x'}\right] \\
L_{ss'} &= \left(L_+\right)_{ss'}P_+ + \left(L_-\right)_{ss'}P_-
\end{aligned} \tag{6.3}$$

with

$$\left(L_+\right)_{ss'} = \left(L_-\right)_{s's} = \begin{cases} -m\delta_{L_s-1,s'}, & s = 0 \\ \delta_{s-1,s'}, & 1 \le s \le L_s - 1 \end{cases}. \tag{6.4}$$

Here x and s are spacetime indices in the 4D bulk and along the fifth dimension, respectively, with L_s denoting the total number of s sites, $P_\pm = (1 \pm \gamma_5)/2$ denoting the chiral projection operators, and $(R_5)_{ss'} \equiv \delta_{s,Ls-1-s'}$ denoting the operator which performs a reflection in the fifth dimension. We recover four-dimensional quark fields q and $\overline{q}$ with definite chiralities from the five-dimensional quark fields ψ and $\overline{\psi}$ described by $\mathscr{D}_{\text{DWF}}$ at the boundaries of the fifth dimension

$$\begin{aligned} q_R &= P_+ \psi_{L_s-1} \qquad & q_L &= P_- \psi_0 \\ \overline{q}_R &= \overline{\psi}_{L_s-1} P_- \qquad & \overline{q}_L &= \overline{\psi}_0 P_+ \end{aligned} \quad . \tag{6.5}$$

Green's functions constructed from q and $\overline{q}$ approximate continuum QCD arbitrarily well in the limit of vanishing lattice spacing and infinite 5D spacetime volume. The tunable parameters in Eqn. (6.3) are the domain wall parameter M_5, the bare quark mass m, and the Möbius scale $\alpha = 2c$; the parameter d is fixed at $d = 1/2$. DWF with the Shamir kernel is recovered from the more general Möbius operator in the limit $\alpha \to 1$. For more detail regarding our MDWF formalism we direct the reader to Ref. [11].

The construction of the exact one flavor action for domain wall fermions begins by factorizing the MDWF Dirac operator as

$$\mathscr{D}_{\text{DWF}} = \mathscr{D}_{\text{EOFA}} \cdot \widetilde{\mathscr{D}}, \tag{6.6}$$

with

$$\begin{aligned} (\mathscr{D}_{\text{EOFA}})_{xx',ss'} &\equiv (D_W)_{xx'} \, \delta_{ss'} + \delta_{xx'} \, (M_+)_{ss'} \, P_+ + \delta_{xx'} \, (M_-)_{ss'} \, P_- \\ (\widetilde{\mathscr{D}})_{ss'} &\equiv d \, (\delta_{ss'} - L_{ss'}) + c \, (\delta_{ss'} + L_{ss'}) \end{aligned} \quad . \tag{6.7}$$

The operator $\widetilde{\mathscr{D}}$ relating $\mathscr{D}_{\text{DWF}}$ and $\mathscr{D}_{\text{EOFA}}$ has no dependence on the gauge field, so we are free to replace $\mathscr{D}_{\text{DWF}}$ with $\mathscr{D}_{\text{EOFA}}$ in Eqn. (6.2) without modifying physical observables described by a properly normalized path integral. In fact, it can be shown analytically using the explicit form of $\widetilde{\mathscr{D}}$ listed in Appendix 6.B that

$$\det (\widetilde{\mathscr{D}}) = \left((c+d)^{L_s} + m \, (c-d)^{L_s} \right)^{12V}, \tag{6.8}$$

where $V = L^3 T$ is the 4D spacetime volume. This substitution facilitates the construction of a proper action since the operator $\gamma_5 R_5 \mathscr{D}_{\text{EOFA}}$ is manifestly Hermitian for any choice of the Möbius scale α, whereas $\mathscr{D}_{\text{DWF}}$ satisfies a less trivial γ_5-Hermiticity condition when $\alpha \neq 1$ [12]. However,

this comes at the cost of substantially more expensive inversions, since $\mathscr{D}_{\text{EOFA}}$ is dense in ss' whereas $\mathscr{D}_{\text{DWF}}$ has a well-known tridiagonal block structure.

After introducing $\mathscr{D}_{\text{EOFA}}$, TWQCD's construction proceeds by applying the Schur identity

$$\det\left[\begin{pmatrix} A & B \\ C & D \end{pmatrix}\right] = \det(A)\det\left(D - CA^{-1}B\right) = \det(D)\det\left(A - BD^{-1}C\right) \tag{6.9}$$

to $\mathscr{D}_{\text{EOFA}}$, treated as a 2×2 block matrix in its spinor indices, and rearranging terms to arrive at the right-hand side of Eqn. (6.2). Crucially, factors of $\gamma_5 R_5$ can be freely inserted under the determinant to replace $\mathscr{D}_{\text{EOFA}}$ with the Hermitian operator $H \equiv \gamma_5 R_5 \mathscr{D}_{\text{EOFA}}$, since $\det(\gamma_5) = \det(R_5) = 1$. The final form of the exact one flavor pseudofermion action is $S_{\text{EOFA}} = \phi^\dagger \mathcal{M}_{\text{EOFA}} \phi$, with

$$\mathcal{M}_{\text{EOFA}} \equiv 1 - kP_-\Omega_-^\dagger \left[H(m_1)\right]^{-1}\Omega_- P_- + kP_+\Omega_+^\dagger \left[H(m_2) - \Delta_+(m_1,m_2)P_+\right]^{-1}\Omega_+ P_+. \tag{6.10}$$

In Appendix 6.B we collect explicit expressions for k, $\Omega_\pm$, $\Delta_\pm$, $\mathscr{D}_{\text{EOFA}}$, and $\widetilde{\mathscr{D}}$ for Shamir and Möbius DWF, since, to the authors' knowledge, these expressions have not previously appeared in the literature. In Ref. [3] these operators are constructed recursively for the more general case of Zolotarev-type DWF with weights $\rho_s = c\omega_s + d$ and $\sigma_s = c\omega_s - d$ that are allowed to vary along the fifth dimension, subject to the constraint that ω_s is reflection-symmetric in s.

6.3 Summary of Ensembles Used in This Work

The properties of the lattices used in this work are summarized in Tables 6.2 and 6.3. In all cases we use the Iwasaki gauge action (I) [13], and on some ensembles supplement this with the dislocation suppressing determinant ratio (DSDR) [14, 15]; we abbreviate the combined action including both terms as "ID". The additional DSDR term is designed to suppress the dislocations of the gauge field associated with tunneling between topological sectors, thereby reducing the degree of residual chiral symmetry breaking. For strong coupling simulations, where these dislocations occur frequently, the DSDR term reduces the costs associated with light quark masses while still maintaining good topological sampling. We simulate $N_f = 2+1$ quark flavors using domain wall fermions, with either the Shamir (DWF) [16, 17] or Möbius (MDWF) [18–20] kernel. Finally, on ensembles marked "-G" we use G-parity boundary conditions in one or more of the spatial directions.

Ensemble	Action	β	$L^3 \times T \times L_s$	Möbius Scale	G-Parity B.C.	am_l	am_h
16I	DWF + I	2.13	$16^3 \times 32 \times 16$	—	—	0.01	0.032
16I-G	DWF + I	2.13	$16^3 \times 32 \times 16$	—	x	0.01	0.032
16ID-G	MDWF + ID	1.75	$16^3 \times 32 \times 8$	2.00	x,y,z	0.01	0.045
24ID	MDWF + ID	1.633	$24^3 \times 64 \times 24$	4.00	—	0.00789	0.085
32ID-G	MDWF + ID	1.75	$32^3 \times 64 \times 12$	2.67	x,y,z	0.0001	0.045

Table 6.2: Summary of ensembles and input parameters used in this work. Here β is the gauge coupling, $L^3 \times T \times L_s$ is the lattice volume decomposed into the length of the spatial (L), temporal (T), and fifth (L_s) dimensions, and am_l and am_h are the bare, input light and heavy quark masses. On the 16I-G, 16ID-G, and 32ID-G ensembles G-parity boundary conditions are applied to the fermion fields at one or more of the spatial boundaries of the lattice; otherwise periodic boundary conditions are applied, and in all cases antiperiodic boundary conditions are used along the temporal direction.

The 16I ensemble was first generated and used to study light meson spectroscopy with domain wall fermions in Ref. [21]. The 16I-G ensemble is identical to the 16I ensemble except for the boundary conditions along the x-direction, which have been changed from periodic to G-parity. Likewise, the parameters of the 16ID-G ensemble have been chosen based on a series of $\beta = 1.75$ DSDR ensembles generated in Ref. [22], but have G-parity boundary conditions in all three spatial directions. Collectively, these three lattices are used as inexpensive, small-volume test ensembles with unphysical, heavy pion masses to perform cross-checks of the EOFA algorithm and its implementation in the BFM and CPS code libraries. The larger 24ID [23] and 32ID-G [24] ensembles have physical pion masses and are currently being generated as part of production RBC/UKQCD calculations.

Ensemble	L (fm)	a^{-1} (GeV)	m_π (MeV)
16I	1.95(5)	1.62(4)	400(11)
16I-G	1.95(5)	1.62(4)	388(14)
16ID-G	2.29(1)	1.378(7)	575(11)
24ID	4.82(19)	0.981(39)	137.1(5.5)
32ID-G	4.57(2)	1.378(7)	143.1(2.0)

Table 6.3: Summary of spatial volumes, lattice cutoffs, and pion masses in physical units for the ensembles used in this work. All values for the 16I and 32ID-G ensembles are from Refs. [21] and [6], respectively. On the 16I-G (16ID-G) ensemble we assume the lattice cutoff is the same as the 16I (32ID-G) ensemble since the same action and value of β has been used. The pion masses on the 16I-G and 16ID-G ensembles have been extracted using the fitted value of the lowest energy pion states from Table 6.8 and the continuum dispersion relation. Finally, the determination of the lattice scale for the 24ID ensemble was performed in Ref. [25], and the determination of the pion mass in Ref. [23].

6.4 Hybrid Monte Carlo with EOFA

In lattice QCD correlation functions are computed in terms of a discretized Euclidean path integral

$$\langle \mathscr{O}_1 \cdots \mathscr{O}_n \rangle = \frac{1}{\mathscr{Z}} \int \mathscr{D}U \left(\prod_f \mathscr{D}\psi_f \mathscr{D}\overline{\psi}_f \right) (\mathscr{O}_1[U] \cdots \mathscr{O}_n[U])\, e^{-S[U,\overline{\psi}_f,\psi_f]}. \tag{6.11}$$

Here U is the gauge field, ψ_f is the quark field associated with flavor f, and $S[U,\overline{\psi}_f,\psi_f]$ is the action, which decomposes into a sum of contributions from the gauge field, fermions, and possibly other terms (e.g. the dislocation suppressing determinant ratio). To avoid having to deal with anticommuting Grassman variables in a computer, dynamical fermion flavors are integrated out and then reintroduced in terms of bosonic "pseudofermion" fields ϕ as

$$\frac{1}{\mathscr{Z}} \int \mathscr{D}\psi \mathscr{D}\overline{\psi}\, e^{-\overline{\psi} M \psi} = \det(M) = \frac{1}{\det(M^{-1})} = \frac{1}{\mathscr{Z}} \int \mathscr{D}\phi \mathscr{D}\phi^\dagger\, e^{-\phi^\dagger M^{-1} \phi}, \tag{6.12}$$

provided M is positive-definite. While pseudofermions can be represented straightforwardly in a computer, they come at the cost of applications of M^{-1} rather than M, which is not typically

available in an explicit form. Even after discretization the integration in Eqn. (6.11) is far too expensive to perform directly due to the enormous number of degrees of freedom on a typical lattice. Instead, Monte Carlo techniques are used to ergodically sample a sequence of representative configurations of the gauge field $\{U_i\}$, for which

$$\langle \mathscr{O}_1 \cdots \mathscr{O}_n \rangle \approx \frac{1}{N} \sum_{i=1}^{N} \mathscr{O}_1(U_i) \cdots \mathscr{O}_n(U_i). \tag{6.13}$$

The standard Monte Carlo technique used in modern lattice QCD calculations is known as the Hybrid Monte Carlo (HMC) algorithm.

HMC generates a Markov chain of gauge field configurations $\{U_i\}$ by evolving a Hamiltonian system in unphysical Molecular Dynamics (MD) "time". This Hamiltonian system is constructed by treating $U_\mu(x)$ as a generalized coordinate, introducing an $\mathfrak{su}(3)$-valued conjugate momentum $\pi_\mu(x)$, and forming the standard Hamiltonian

$$H = \frac{1}{2}\pi^2 + S(U). \tag{6.14}$$

The associated equations of motion

$$\begin{cases} \partial_\tau U_\mu(x) = \pi_\mu(x) U_\mu(x) \\ \partial_\tau \pi_\mu(x) = -T^a \partial^a_{x,\mu} S(U) \end{cases} \tag{6.15}$$

can then be integrated using numerical integration techniques. The integration is performed over intervals of length $\Delta\tau$ — referred to as a single MD trajectory — as a sequence of N small steps $\delta\tau$, with $N = \Delta\tau/\delta\tau$. Finite precision integration errors are corrected stochastically with a Metropolis accept/reject step: after every N integration steps by $\delta\tau$ the total change in the Hamiltonian ΔH is computed, and the current gauge field $U'_\mu(x)$ is accepted as the next configuration in the Markov chain with probability

$$P_{\text{accept}} = \min\left(1, e^{-\Delta H}\right). \tag{6.16}$$

One can show that the resulting algorithm satisfies detailed balance provided the scheme used to numerically integrate Eqn. (6.15) is reversible [26]. Ergodicity is achieved by performing a heatbath step each time the integration is restarted to pick a new conjugate momentum $\pi_\mu(x)$, and thus a new trajectory in the phase space $\{(U, \pi)\}$. HMC generates a sequence of gauge field configurations

whose statistical independence is governed by the length of each MD trajectory, $\Delta\tau$. The number of MD trajectories separating statistically independent gauge field configurations is typically determined *ex post facto* by examining the integrated autocorrelation times of representative physical observables.

The fermionic contribution to the Hamiltonian in Eqn. (6.14) introduces a technical obstacle for the HMC algorithm since the lattice Dirac operator $\mathscr{D}$ has a complex spectrum. Replacing $\mathscr{D}$ with the Hermitian fermion matrix $M = \mathscr{D}^\dagger\mathscr{D}$ in Eqn. (6.12) has a number of advantages. Most importantly, it allows M^{-1} to be applied to pseudofermion vectors using the conjugate gradient algorithm, and it allows for a straightforward pseudofermion heatbath step: at the beginning of each MD trajectory a random Gaussian vector η is drawn according to $P(\eta) \propto \exp(-\eta^\dagger\eta/2)$ and the initial pseudofermion field is seeded as $\phi = \mathscr{D}\eta$, ensuring that ϕ is correctly sampled as $P(\phi) \propto \exp(-\phi^\dagger M^{-1}\phi/2)$. However, the fermion matrix $M = \mathscr{D}^\dagger\mathscr{D}$ describes two degenerate quark flavors with the same mass. Single flavor simulations are typically performed by taking an overall square root of the fermion determinant,

$$\det(\mathscr{D}) = \left[\det\left(\mathscr{D}^\dagger\mathscr{D}\right)\right]^{1/2}. \tag{6.17}$$

In the pseudofermion formalism applications of the operator $(\mathscr{D}^\dagger\mathscr{D})^{-1/2}$ are approximated by a matrix-valued function $f(\mathscr{D}^\dagger\mathscr{D})$,

$$\left[\det\left(\mathscr{D}^\dagger\mathscr{D}\right)\right]^{1/2} = \frac{1}{\mathcal{Z}}\int \mathscr{D}\phi\mathscr{D}\phi^\dagger e^{-\phi^\dagger(\mathscr{D}^\dagger\mathscr{D})^{-1/2}\phi} \simeq \frac{1}{\mathcal{Z}}\int \mathscr{D}\phi\mathscr{D}\phi^\dagger e^{-\phi^\dagger f(\mathscr{D}^\dagger\mathscr{D})\phi} \tag{6.18}$$

where $f(x)$ is a suitably constructed approximation to the inverse square root, valid over the spectral range of $\mathscr{D}^\dagger\mathscr{D}$. Variants of the HMC algorithm which construct f from different classes of functions have been proposed and used in the literature; the most common is the rational HMC (RHMC) algorithm [27], where

$$f(x) = \alpha_0 + \sum_{k=1}^{N}\frac{\alpha_k}{\beta_k + x} \tag{6.19}$$

is a rational function. While rational functions are in many ways a good choice — they are economical in the sense that the inverse square root can usually be well-approximated by a modest number of terms, and the multishift CG algorithm can be used to efficiently invert $(\mathscr{D}^\dagger\mathscr{D} + \beta_k)$

for all k simultaneously — the additional complexity of evaluating $f(\mathscr{D}^\dagger\mathscr{D})$ and the associated molecular dynamics pseudofermion force makes single flavor RHMC simulations significantly more costly than degenerate two flavor HMC simulations at the same bare quark mass. This additional cost can be largely attributed to the significant linear algebra overhead associated with multishift CG.

EOFA provides an alternative construction of a single-flavor pseudofermion action through Eqn. (6.2): a ratio of fermion determinants can be factorized as a product of two determinants, each of which involves an operator which is Hermitian and positive-definite. This product can then be represented as a path integral over a bosonic pseudofermion field with a two-term action (Eqn. (6.10))

$$\det\left(\frac{\mathscr{D}_{\text{EOFA}}(m_1)}{\mathscr{D}_{\text{EOFA}}(m_2)}\right) = \frac{1}{\mathcal{Z}}\int \mathscr{D}\phi\mathscr{D}\phi^\dagger e^{-\phi^\dagger\mathcal{M}_{\text{EOFA}}\phi}, \tag{6.20}$$

leading to an algorithm which is "exact" in the sense that it avoids the numerical approximations required to implement the square root in RHMC (Eqn. (6.18)) and related HMC variants. EOFA is also expected to be somewhat faster than RHMC, since there is no rational approximation entering into evaluations of the Hamiltonian or the pseudofermion force, eliminating the overhead associated with multishift CG. In the remainder of this section we elaborate on the details of the action, heatbath step, and pseudofermion force entering into the Hamiltonian equations of motion (Eqn. (6.15)) for HMC with EOFA.

6.4.1 Action

The EOFA action (Eqn. (6.10)) computes a ratio of determinants of $\mathscr{D}_{\text{EOFA}}$ upon integrating out the pseudofermion fields (Eqn. (6.20)). This ratio can be related to the conventional determinant ratio computed by the RHMC algorithm through Eqns. (6.6) and (6.8), leading to the relationship

$$\det\left(\frac{\mathscr{D}_{\text{DWF}}(m_1)}{\mathscr{D}_{\text{DWF}}(m_2)}\right) = \left(\frac{(c+d)^{L_s}+m_1\,(c-d)^{L_s}}{(c+d)^{L_s}+m_2\,(c-d)^{L_s}}\right)^{12V}\det\left(\frac{\mathscr{D}_{\text{EOFA}}(m_1)}{\mathscr{D}_{\text{EOFA}}(m_2)}\right). \tag{6.21}$$

We use this relationship as a test of the equivalence of RHMC and EOFA, as well as our implementation of the EOFA action, by stochastically computing the left side of Eqn. (6.21) with the

RHMC action

$$\mathcal{M}_{\text{RHMC}} = \left[\mathscr{D}^\dagger_{\text{DWF}}\mathscr{D}_{\text{DWF}}(m_2)\right]^{1/4}\left[\mathscr{D}^\dagger_{\text{DWF}}\mathscr{D}_{\text{DWF}}(m_1)\right]^{-1/2}\left[\mathscr{D}^\dagger_{\text{DWF}}\mathscr{D}_{\text{DWF}}(m_2)\right]^{1/4} \tag{6.22}$$

and the right side with the EOFA action (Eqn. (6.10)) on the same gauge field configuration. Observing that we can, in general, rewrite a determinant as

$$\det\left(\mathcal{M}^{-1}\right) = \frac{1}{\mathcal{Z}}\int \mathscr{D}\phi\mathscr{D}\phi^\dagger e^{-\phi^\dagger\mathcal{M}\phi} = \frac{1}{\mathcal{Z}}\int \mathscr{D}\phi\mathscr{D}\phi^\dagger e^{-\frac{1}{2}\phi^\dagger\Sigma^{-1}\phi}e^{\phi^\dagger(\frac{1}{2}\Sigma^{-1}-\mathcal{M})\phi} \tag{6.23}$$

suggests the following simple Monte Carlo integration scheme: we draw random pseudofermion vectors by independently sampling the real and imaginary parts of each component from the standard normal distribution $\mathcal{N}(\mu=0,\sigma=1)$, and compute the expectation value

$$-\operatorname{logdet}\left(\mathcal{M}^{-1}\right) \approx \left\langle \phi_i^\dagger\left(\mathcal{M}-\frac{1}{2}\Sigma^{-1}\right)\phi_i\right\rangle_i, \tag{6.24}$$

where the average is computed using the jackknife resampling technique. This will accurately approximate the true log determinant for finite, realistically calculable values of N provided m_1 and m_2 are sufficiently close that the integrand is well-approximated by a Gaussian with unit variance. To address this latter systematic, we consider splitting Eqn. (6.21) as a product of determinants

$$\det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m_2)}\right) = \det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m_1')}\right)\left[\prod_{i=1}^{N_m}\det\left(\frac{\mathscr{D}(m_i')}{\mathscr{D}(m_{i+1}')}\right)\right]\det\left(\frac{\mathscr{D}(m_{N_m}')}{\mathscr{D}(m_2)}\right) \tag{6.25}$$

with equally-spaced intermediate masses

$$m_i' = m_1 + \frac{m_2-m_1}{N_m+1}i, \quad i=1,\ldots,N_m, \tag{6.26}$$

and study the dependence of the result on N_m (this procedure is identical to the method introduced in Ref. [28] for computing quark mass reweighting factors). In the upper panel of Figure 6.1 we plot the log determinants of $\mathcal{M}^{-1}_{\text{RHMC}}$ and $\mathcal{M}^{-1}_{\text{EOFA}}$ as a function of N_m, with $N=10$ stochastic evaluations, computed using a single thermalized trajectory of the 16I, 16I-G, and 16ID-G ensembles. For the case of the 16ID-G ensemble, which uses the Möbius DWF fermion action, we also include the overall constant multiplying the right side of Eqn. (6.21) so that in all cases we are computing the same determinant ratio of $\mathscr{D}_{\text{DWF}}$ using either action.

We observe, as expected, that both formalisms agree for sufficiently large N_m. Likewise, we observe that at sufficiently small N_m the calculation generally becomes unreliable since we do not attempt to account for the systematic error associated with approximating the integrand of the path integral by a Gaussian with unit variance (*i.e.* setting $\Sigma = 1$ in Eqns. (6.23) and (6.24).). In both cases "sufficiently" small or large N_m is controlled by the size of the splitting between m_1 and m_2. We also observe that, for a given choice of N_m and N, both the statistical and systematic errors of the determinant ratio computed via EOFA are suppressed relative to the errors of the determinant ratio computed via RHMC. We argue that the observed error suppression can be explained by comparing the spectrum of $\mathcal{M}_{\text{RHMC}}$ to the spectrum of $\mathcal{M}_{\text{EOFA}}$, which we plot in the lower panels of Figure 6.1 for a very small lattice volume (4^5) where the complete spectrum can be computed directly. While both operators have similar condition numbers, we find that most of the spectrum of the EOFA action is concentrated into a small interval $[1, 1 + \Delta]$ with $\Delta \sim \mathcal{O}(0.1)$, leading to an action which is easier to estimate stochastically.

We propose that TWQCD's EOFA construction can be thought of as a kind of preconditioning which computes the same determinant ratio as RHMC but modifies the operator inside the determinant ($\mathcal{M}_{\text{RHMC}}$), mapping its spectrum onto a more compact interval. This suggests an additional application of the EOFA formalism: quark mass reweighting factors can be computed substantially more cheaply using the EOFA action than using the RHMC action, especially at light quark masses, even if the ensemble was generated using RHMC. This could be useful, for example, to include the dynamical effects of isospin breaking in ensembles generated with isospin-symmetric up and down quarks.

6.4.2 Heatbath

At the beginning of each HMC trajectory we wish to draw a random pseudofermion field ϕ according to the distribution $P(\phi) \propto \exp(-\phi^\dagger \mathcal{M}_{\text{EOFA}} \phi)$. To do this, we first draw a random vector η by independently sampling the real and imaginary parts of each component from the normal distribution with $\mu = 0$ and $\sigma^2 = 1/2$, and then compute $\phi = \mathcal{M}_{\text{EOFA}}^{-1/2} \eta$. As with the RHMC algorithm we approximate the inverse square root by an appropriately constructed rational function, but we

stress that in the context of EOFA this rational approximation enters only into the heatbath and is not necessary to compute the EOFA action itself or the associated pseudofermion force. Naively applying a rational approximation with the form of Eqn. (6.19) to the operator $\mathcal{M}_{\text{EOFA}}$ results in

$$\mathcal{M}_{\text{EOFA}}^{-1/2} \simeq \alpha_0 + \sum_{k=1}^{N_p} \alpha_l \left[\frac{1}{\gamma_l} - kP_- \Omega_-^\dagger \left[H(m_1) \right]^{-1} \Omega_- P_- \right. \\ \left. + kP_+ \Omega_+^\dagger \left[H(m_2) - \Delta_+(m_1, m_2) P_+ \right]^{-1} \Omega_+ P_+ \right]^{-1}, \quad (6.27)$$

where we have defined $\gamma_l \equiv (1 + \beta_l)^{-1}$. In this form, the nested inversions required to seed the heatbath would make EOFA prohibitively expensive. However, the Woodbury matrix identity

$$(A + BCD)^{-1} = A^{-1} - A^{-1} B \left(C^{-1} + DA^{-1} B \right)^{-1} DA^{-1} \quad (6.28)$$

and the cancellation between cross-terms involving products of the chiral projection operators can be used to manipulate this expression into the equivalent form

$$\mathcal{M}_{\text{EOFA}}^{-1/2} \simeq \alpha_0 + \sum_{k=1}^{N_p} \alpha_l \gamma_l \Big\{ 1 + k\gamma_l P_- \Omega_-^\dagger \left[H(m_1) - \gamma_l \Delta_-(m_1, m_2) P_- \right]^{-1} \Omega_- P_- \\ - k\gamma_l P_+ \Omega_+^\dagger \left[H(m_2) - \beta_l \gamma_l \Delta_+(m_1, m_2) P_+ \right]^{-1} \Omega_+ P_+ \Big\}. \quad (6.29)$$

With this expression the EOFA heatbath step can be performed at the cost of $2N_p$ CG inversions using a rational approximation with N_p poles. Unlike the case of RHMC, multishift CG algorithms are not applicable to the EOFA heatbath since each of the $2N_p$ operators in Eqn. (6.29) generates a different Krylov space. Furthermore, since the operators $\Delta_\pm P_\pm$ have a large number of zero modes and are therefore not invertible, there is no simple transformation by which this system can be recast into a form amenable to multishift CG.

In the left panel of Figure 6.2 we test Eqn. (6.29) on a single thermalized configuration of the 16I ensemble by computing the quantity

$$\varepsilon \equiv \frac{\left| \eta^\dagger \eta - \phi^\dagger \mathcal{M}_{\text{EOFA}} \phi \right|}{\eta^\dagger \eta} \quad (6.30)$$

after seeding the pseudofermion field ϕ with a random Gaussian vector η. In exact arithmetic $\varepsilon = 0$; in practice it measures the relative error in the heatbath step arising from the choice of

CG stopping conditions and rational approximation to the inverse square root. We repeat this calculation, varying the number of poles in the rational approximation but keeping the stopping conditions fixed, and observe that ε reaches the limits of double-precision arithmetic even with a relatively modest number of poles compared to what is typically required to compute non-integer powers of $\mathscr{D}^\dagger\mathscr{D}$ accurately in the context of RHMC. In the right panel of Figure 6.2 we demonstrate this explicitly by computing the condition numbers $\kappa = \lambda_{\max}/\lambda_{\min}$ of both operators as a function of the bare input quark mass. In Section 6.6 we show how aggressive tuning of the rational approximation and stopping conditions, together with forecasting techniques for the initial CG guesses, can be combined to ameliorate the cost of the $2N_p$ inversions required to apply $\mathcal{M}_{\text{EOFA}}^{-1/2}$.

6.4.3 Pseudofermion Force

The pseudofermion force

$$T^a \partial^a_{x,\mu} S(U) \equiv T^a \frac{d}{ds} S\left(e^{sT^a} U_\mu(x)\right)\Big|_{s=0} \tag{6.31}$$

measures the back-reaction of the pseudofermions on the HMC system (Eqn. (6.15)) under an infinitesimal variation of the gauge field. In our notation $\{T^a\}$ is a basis for the Lie algebra $\mathfrak{su}(3)$, with the CPS normalization convention

$$\text{Tr}\left(T^a T^b\right) = -\frac{1}{2}\delta^{ab}. \tag{6.32}$$

The EOFA pseudofermion force can be worked out explicitly by differentiating the EOFA action (Eqn. (6.10)) and applying the matrix identity

$$\partial_x M^{-1} = -M^{-1}\left(\partial_x M\right) M^{-1}, \tag{6.33}$$

resulting in

$$T^a \partial^a_{x,\mu} S(U) = kT^a \left(\gamma_5 R_5 \chi_1\right)^\dagger \left(\partial^a_{x,\mu} D_W\right) \chi_1 - kT^a \left(\gamma_5 R_5 \chi_2\right)^\dagger \left(\partial^a_{x,\mu} D_W\right) \chi_2, \tag{6.34}$$

with

$$\chi_1 \equiv \left[H(m_1)\right]^{-1} \Omega_- P_- \phi \tag{6.35}$$

and

$$\chi_2 \equiv \left[H(m_2) - \Delta_+(m_1, m_2) P_+\right]^{-1} \Omega_+ P_+ \phi. \tag{6.36}$$

Standard manipulations can be used to write

$$\begin{aligned} a^\dagger \left(\partial^a_{x,\mu} D_W\right) b &= -\frac{k}{2} \sum_{x,s,\mu} \Big[a^\dagger(x,s) T^a U_\mu(x) \left(1-\gamma_\mu\right) b(x+\hat{\mu},s) \\ &\qquad\qquad - a^\dagger(x+\hat{\mu},s) U^\dagger_\mu(x) T^a \left(1+\gamma_\mu\right) b(x) \Big] \\ &= -\frac{k}{2} \sum_{x,s,\mu} \Bigg[U_\mu(x) \Bigg(\operatorname*{Tr}_{\text{spin}} \Big[\left(1+\gamma_\mu\right) a(x+\hat{\mu},s) b^\dagger(x,s) \Big] + \\ &\qquad\qquad \operatorname*{Tr}_{\text{spin}} \Big[\left(1-\gamma_\mu\right) b(x+\hat{\mu},s) a^\dagger(x,s) \Big] \Bigg) \Bigg], \end{aligned} \tag{6.37}$$

allowing Eqn. (6.34) to be efficiently computed locally in terms of a trace over spinor indices, at the cost of the two inversions required to form χ_1 and χ_2. Since Dirac bilinears of the form $a^\dagger \left(\partial^a_{x,\mu} D_W\right) b$ enter into the pseudofermion forces associated with many of the standard pseudofermion actions for Wilson and domain wall fermions — including the RHMC action — implementing the EOFA pseudofermion force requires little new code beyond what is required to implement the EOFA Hamiltonian.

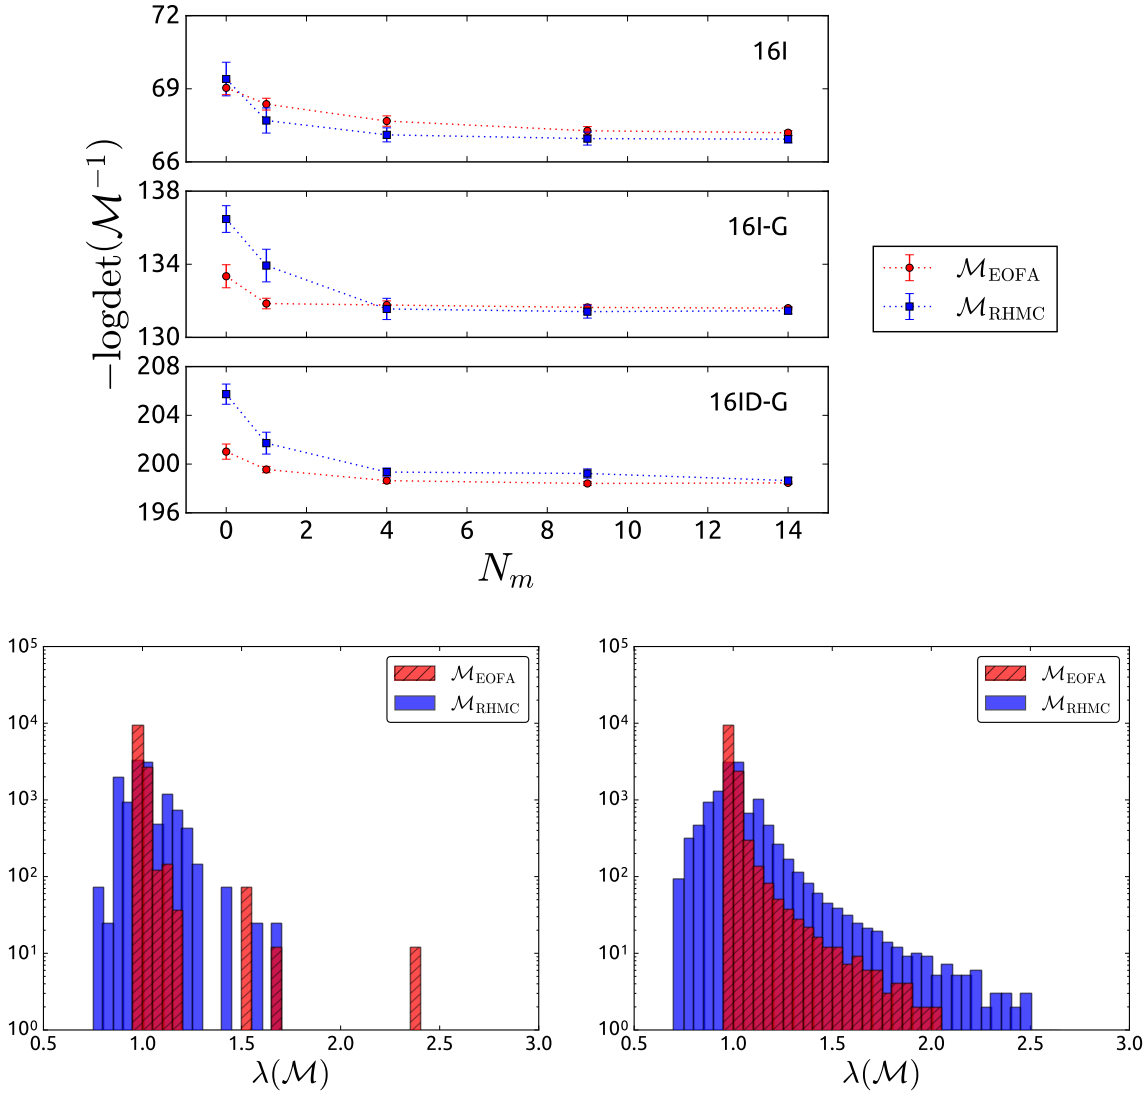

Figure 6.1: Top: log determinants of the EOFA and RHMC actions as a function of the number of intermediate masses (N_m) used to compute Eqn. (6.25), computed on a single, thermalized configuration of the 16I, 16I-G, and 16ID-G ensembles. We set (am_1, am_2) to (0.032, 0.042), (0.032,0.042), and (0.045,0.055) on the 16I, 16I-G, and 16ID-G ensemble, respectively. We note that the error bars are purely statistical; for small N_m there is a large, unaccounted systematic error associated with setting $\Sigma = 1$ in Eqns. (6.23) and (6.24). Bottom: eigenvalue spectra of $\mathcal{M}_{\text{EOFA}}$ and $\mathcal{M}_{\text{RHMC}}$ on a 4^5 lattice with $am_1 = 0.01$, $am_2 = 1.0$, and $aM_5 = 1.8$. In the bottom left plot all of the gauge links are set to $U_\mu(x) = 1$ (*i.e.* the free field limit); in the bottom right plot each gauge link is set to an independent, random $SU(3)$ matrix.

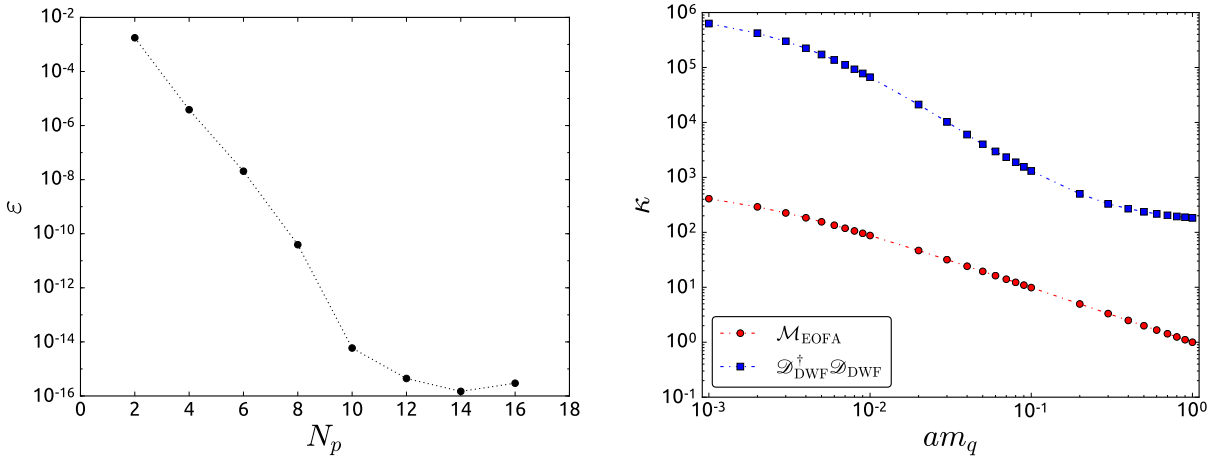

Figure 6.2: Left: relative error — ε, defined by Eqn. (6.30) — in seeding the pseudofermion heatbath as a function of the number of poles in the rational approximation to the inverse square root (N_p), with $am_1 = 0.032$ set to the dynamical heavy quark mass, and a stopping residual of 10^{-10} for all CG inversions. Right: condition numbers of $\mathcal{M}_{\mathrm{EOFA}}$ and $\mathscr{D}^{\dagger}_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}$ as a function of the bare input quark mass (am_q); for $\mathcal{M}_{\mathrm{EOFA}}$ this is the numerator mass ($am_1 = am_q$), while the denominator mass is fixed at $am_2 \equiv 1$. Both calculations were performed on a single, thermalized configuration of the 16I ensemble.

6.5 Small Volume Reproduction Tests

To further test our implementation of EOFA we have reproduced the 16I (16I-G, 16ID-G) ensemble using EOFA for the strange quark (light quarks) in place of RHMC. For these tests we have made no serious effort to tune EOFA for performance; we have simply checked that replacing RHMC with EOFA, but leaving all other details of the simulation fixed, has no discernible impact on physical observables such as the average plaquette, topological susceptibility, and low energy spectrum.

6.5.1 Ensemble Generation

The details of the integrator parameters and nesting are summarized in Tables 6.4 and 6.5, respectively. We use the abbreviations

$$\mathrm{Quo}(m_1, m_2) \equiv \det\left(\frac{\mathscr{D}^\dagger_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}(m_1)}{\mathscr{D}^\dagger_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}(m_2)}\right) \tag{6.38}$$

and

$$\mathrm{RatQuo}_{1/n}(m_1, m_2) \equiv \left[\det\left(\frac{\mathscr{D}^\dagger_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}(m_1)}{\mathscr{D}^\dagger_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}(m_2)}\right)\right]^{1/n} \tag{6.39}$$

to denote the quotient and rational quotient actions, and on the EOFA reproduction ensembles replace each instance of $\mathrm{RatQuo}_{1/2}(m_1, m_2)$ with the EOFA action (Eqn. (6.10)) using the same mass parameters. The 16I and 16I-G EOFA reproduction runs were seeded with an ordered start — *i.e.* all gauge links were initially set to the unit matrix — and evolved for 1500 and 2500 MD time units, respectively. For the 16ID-G ensemble the last RHMC configuration (MD trajectory 908) was used to seed the start of the EOFA reproduction run, and then evolved for an additional 500 MD time units.

6.5.2 Basic Observables

In Table 6.6 we summarize results for the average plaquette $\langle P\rangle$, light quark and strange quark chiral $\langle\overline{\psi}\psi\rangle$ and pseudoscalar $\langle\overline{\psi}\gamma_5\psi\rangle$ condensates, and the topological susceptibility $\chi_t \equiv \langle Q^2\rangle/V$, computed on each ensemble; we observe statistically consistent results between the RHMC and EOFA ensembles in each case. Accompanying plots of the time evolution can be found in Section 6.D. The topological charge Q has been measured using the 5Li discretization introduced in Ref. [29]

Ensemble	Integrator	$\delta\tau$	r_{FG}	r_{MD}	r_{MC}
16I	Force Gradient QPQPQ	0.1000	10^{-7}	10^{-8}	10^{-10}
16I-G	Omelyan ($\lambda = 0.2$)	0.2000	—	10^{-8}	10^{-10}
16ID-G	Force Gradient QPQPQ	0.1667	10^{-6}	10^{-7}	10^{-10}

Table 6.4: Basic integrator and HMC details for the generation of the 16I, 16I-G, and 16ID-G ensembles. We use nested Sexton-Weingarten integration schemes, detailed in Table 6.5, with $\delta\tau$ the coarsest time step used to evolve the outermost level. We denote the CG stopping tolerances used for the force gradient forecasting, molecular dynamics, and Monte Carlo steps by r_{FG}, r_{MD}, and r_{MC}, respectively.

Ensemble	Level	Action	Update
16I	1	Quo(0.01,0.2) + Quo(0.2,1.0) + $\mathrm{RatQuo}_{1/2}$(0.032,1.0)	4:1
	2	Gauge	1:1
16I-G	1	$\mathrm{RatQuo}_{1/2}$(0.01,0.032)	1:1
	2	$\mathrm{RatQuo}_{1/4}$(0.032,1.0) + $\mathrm{RatQuo}_{1/4}$(0.032,1.0) + $\mathrm{RatQuo}_{1/4}$(0.032,1.0)	8:1
	3	Gauge	1:1
16ID-G	1	$\mathrm{RatQuo}_{1/2}$(0.01,0.05) + $\mathrm{RatQuo}_{1/2}$(0.05,1.0) + $\mathrm{RatQuo}_{1/4}$(0.045,1.0)	1:1
	2	DSDR	8:1
	3	Gauge	1:1

Table 6.5: Integrator layouts for the original RHMC runs. Here "Quo" is an abbreviation for the quotient action (Eqn. (6.38)) and "$\mathrm{RatQuo}_{1/n}$" is an abbreviation for the rational quotient action (Eqn. (6.39)), with a rational function approximation used to apply $(\mathscr{D}^\dagger\mathscr{D})^{1/n}$ and its inverse. For the EOFA reproduction runs each instance of $\mathrm{RatQuo}_{1/2}$ is replaced by an EOFA determinant with the same masses (Eqn. (6.10)), while all other ensemble and integrator details are left fixed. The notation A:B for the update scheme denotes the number of steps of the next innermost integrator level (A) per step of the current level (B).

after cooling the gauge fields with 20 steps of APE smearing [30] using a smearing coefficient of 0.45. The ensemble averages were computed after binning over 50 (25) successive MD time units on the 16I and 16I-G (16ID-G) ensembles, where the bin size has been conservatively chosen based on the integrated autocorrelation times measured in Ref. [21] for the 16I ensemble and Ref. [22] for a series of $\beta = 1.75$ DSDR ensembles. We expect that the runs produced for this study are too short to reliably compute integrated autocorrelation times directly, but note that there is no evidence of a difference in an integrated autocorrelation time between the EOFA and RHMC ensembles in the time evolution plots of Section 6.D.

	16I		16I-G		16ID-G	
Observable	RHMC	EOFA	RHMC	EOFA	RHMC	EOFA
$\langle P \rangle$	0.588087(22)	0.588106(26)	0.588033(24)	0.588039(16)	0.514251(43)	0.514200(48)
$\langle \overline{\psi}_l \psi_l \rangle$	0.001697(5)	0.001698(11)	0.0017151(72)	0.0017130(52)	0.005543(11)	0.005563(8)
$\langle \overline{\psi}_s \psi_s \rangle$	0.0037450(31)	0.0037435(74)	0.0037541(51)	0.0037529(34)	0.0085729(82)	0.0085895(69)
$\langle \overline{\psi}_l \gamma_5 \psi_l \rangle$	-0.000015(14)	-0.000012(19)	-0.000003(15)	-0.000006(12)	0.000033(13)	-0.000001(11)
$\langle \overline{\psi}_s \gamma_5 \psi_s \rangle$	-0.000001(8)	-0.000007(12)	-0.0000004(92)	-0.0000034(81)	0.000017(10)	-0.000002(8)
χ_t	$1.03(19) \times 10^{-5}$	$1.81(42) \times 10^{-5}$	$2.16(47) \times 10^{-5}$	$1.53(27) \times 10^{-5}$	—	—

Table 6.6: Average plaquettes, quark condensates, and topological susceptibilities (χ_t) computed on the 16I, 16I-G and 16ID-G lattices and their corresponding EOFA reproduction ensembles. The ensemble averages on the 16I (16I-G) lattices were computed using MD trajectories 500-1500 (500-2500) after binning over 50 successive MD time units. The ensemble averages on the 16ID-G lattices were computed using MD trajectories 500:900 for the RHMC ensemble, and MD trajectories 960:1360 for the EOFA ensemble, after binning over 25 successive MD time units. We do not compute χ_t on the 16ID-G ensemble since the short 400 MD time unit measurement runs are insufficient to adequately sample the topological charge, as evidenced by the time evolutions plotted in Appendix 6.D.

6.5.3 Low Energy Spectra

In Table 6.7 we list results for the pion, kaon, Omega baryon, and residual masses, computed on the 16I ensemble. These calculations were performed using a measurement package previously introduced in Ref. [11], and based on the all-mode averaging (AMA) technique of Ref. [31]. Five *exact* light quark propagators were computed per trajectory using a deflated, mixed-precision CG solver with 600 low-mode deflation vectors and a tight stopping residual $r = 10^{-8}$, while *sloppy* propagators were computed for all time slices using a reduced stopping residual $r = 10^{-4}$. Strange quark propagators were computed with the tight residual $r = 10^{-8}$ for all time slices using ordinary CG with no deflation. AMA correlation functions were then computed by time-translational averaging of the sloppy propagators, using the available exact propagators to correct for bias. The light quark propagators were computed using Coulomb gauge-fixed wall (W) sources, with either local (L) or wall sinks; the strange quark propagators were computed using Coulomb gauge-fixed wall or Z_3 box (Z_3B) sources, and in both cases local sinks.

The pion and kaon masses were extracted by fitting to the asymptotic, large Euclidean time limit of the respective two-point correlation function,

$$\langle 0|\overline{\mathscr{O}}(t)\mathscr{O}(0)|0\rangle \overset{t\to\infty}{\simeq} \frac{\langle 0|\overline{\mathscr{O}}(t)|X\rangle\langle X|\mathscr{O}(0)|0\rangle}{2m_X V}\left(e^{-m_X t} \pm e^{-m_X(T-t)}\right), \tag{6.40}$$

where $\mathscr{O}$ denotes the choice of interpolating operator, $X \in \{\pi, K\}$ is the ground state to which $\mathscr{O}$ couples, and V and T are the spatial volume and temporal extent of the lattice, respectively. In particular, we performed simultaneous fits to the $\langle PP^{LW}\rangle$, $\langle PP^{WW}\rangle$, and $\langle AP^{LW}\rangle$ correlators, with $P(x) = \overline{\psi}(x)\gamma_5\psi(x)$ and $A(x) = \overline{\psi}(x)\gamma_5\gamma_4\psi(x)$, and the first (second) superscript denotes the sink (source) type. The Omega baryon mass was extracted from the two-point correlation function

$$C^{s_1 s_2}_{\Omega\Omega}(t) = \sum_{i=1}^{3}\sum_{\vec{x}}\langle 0|\overline{\mathscr{O}}^{s_1}_{\Omega}(\vec{x},t)_i\mathscr{O}^{s_2}_{\Omega}(0)_i|0\rangle \tag{6.41}$$

with the interpolating operator

$$\mathscr{O}_{\Omega}(x)_i = \varepsilon_{abc}\left(s^{\dagger}_a(x)C\gamma_i s_b(x)\right)s_c(x), \tag{6.42}$$

$s_1 = L$ and $s_2 \in \{W, Z_3B\}$. The correlators were then projected onto the positive parity component

$$P_+ C^{s_1 s_2}_{\Omega\Omega} = \frac{1}{4}\mathrm{Tr}\left[\frac{1}{2}\left(1+\gamma_4\right)C^{s_1 s_2}_{\Omega\Omega}\right] \tag{6.43}$$

and simultaneously fit to double exponential ansätze with common mass terms

$$C^{s_1 s_2}_{\Omega\Omega}(t) = (Z_1)^{s_1 s_2}_{\Omega\Omega} e^{-m_\Omega t} + (Z_2)^{s_1 s_2}_{\Omega\Omega} e^{-m'_\Omega t}. \tag{6.44}$$

Finally, the residual mass was determined by fitting the ratio

$$R(t) = \frac{\langle 0 | \sum_{\vec{x}} j^a_{5q}(\vec{x}, t) | \pi \rangle}{\langle 0 | \sum_{\vec{x}} j^a_5(\vec{x}, t) | \pi \rangle} \tag{6.45}$$

to a constant, where j^a_{5q} is the five-dimensional pseudoscalar density evaluated at the midpoint of the fifth dimension, and j^5_a is the physical pseudoscalar density constructed from the surface fields.

	16I	
Observable	RHMC	EOFA
am_π	0.2424(11)	0.2425(8)
am_K	0.3252(11)	0.3253(7)
am_Ω	1.003(15)	0.994(11)
$am'_{\rm res}(m_l)$	0.0030558(80)	0.0030523(78)

Table 6.7: Low energy spectrum on the 16I ensemble computed from 100 independent measurements beginning with MD trajectory 500 and separated by 10 MD time units. Prior to fitting the correlation functions were binned over groups of 5 measurements. Corresponding effective mass plots can be found in Appendix 6.D.

In addition, we have also measured the ground state pion energy, kaon mass, and residual mass on the 16I-G and 16ID-G ensembles. While the ground state of the kaon is at rest, the ground state of the pion has nonzero momentum $\vec{p}_{100} = (\pm\pi/L, 0, 0)$ on the 16I-G ensemble and $\vec{p}_{111} = (\pm\pi/L, \pm\pi/L, \pm\pi/L)$ on the 16ID-G ensemble due to the boundary conditions. These calculations make use of an extension of the AMA measurement package described above to G-parity ensembles; as discussed in Ref. [8], this requires the inclusion of additional diagrams that are generated by the mixing of quark flavors at the lattice boundaries through the G-parity operation. We measure on 51 configurations of the 16I-G ensemble, beginning with trajectory 500 and with a separation of 40 MD time units, and use sloppy and exact CG stopping tolerances of 10^{-4} and 10^{-10},

respectively, with a single exact solve per trajectory. We likewise measure on 21 configurations of the 16ID-G ensemble, beginning with trajectory 500 (960) for the RHMC (EOFA) ensemble and separated by 20 MD time units, and use the same AMA setup. We perform no additional binning for either ensemble since the separations between consecutive measurements are already comparable to the bin sizes used to compute the plaquette and quark condensates.

	16I-G		16ID-G	
Observable	RHMC	EOFA	RHMC	EOFA
aE_π	0.3175(43)	0.3097(48)	0.4457(101)	0.4614(72)
aE_π^{pred}	0.31197(4)	0.31207(4)	—	—
am_K	0.3271(22)	0.3272(28)	0.4343(34)	0.4382(24)
$am'_{\text{res}}(m_l)$	0.003140(90)	0.003054(86)	0.00919(14)	0.00952(13)

Table 6.8: Low energy spectra on the 16I-G and 16ID-G ensembles computed from 51 and 21 measurements, respectively. On the 16I-G ensemble we also predict the ground state pion energy using the fitted am_π on the 16I ensemble and the continuum dispersion relation $aE_\pi^{\text{pred}} = \sqrt{(am_\pi)^2 + (a\pi/L)^2}$. Corresponding effective mass plots can be found in Appendix 6.D.

6.5.4 Pseudofermion Forces on the 16I Ensemble

TWQCD has observed that the average EOFA pseudofermion force is roughly half the size of the corresponding average RHMC pseudofermion force for a particular dynamical $N_f = 1$ QCD simulation with domain wall quarks performed in Ref. [3]. Following this observation, we examine the forces on the RHMC and EOFA variants of the 16I ensemble. We define a norm on the space of $\mathfrak{su}(3)$-valued pseudofermion force matrices $F_\mu^a(x) \equiv \partial_{x,\mu}^a S(U)$ by

$$\|F_\mu(x)\| \equiv \left[\sum_a F_\mu^a(x) F_\mu^a(x)\right]^{1/2}, \tag{6.46}$$

and consider two measures of the force associated with a given configuration of the gauge field: the first is the RMS force

$$F_{\mathrm{RMS}} \equiv \frac{1}{4V}\left[\sum_{x,\mu} \|F_\mu(x)\|^2\right]^{1/2} \tag{6.47}$$

and the second is the maximum force

$$F_{\mathrm{max}} \equiv \max_{x,\mu} \|F_\mu(x)\|, \tag{6.48}$$

in both cases taken over all lattice sites and link directions. While we expect Equation (6.48) to be a more pertinent definition in the context of HMC simulations — we have empirically found that acceptance is controlled by the size of F_{max} — both F_{RMS} and F_{max} are, *a priori*, reasonable measures of the pseudofermion force.

In Figure 6.3 we compare histograms of F_{RMS} and F_{max} between the RHMC and EOFA 16I ensembles. Each data point corresponds to a single evaluation of the pseudofermion force falling between MD trajectories 500 and 1500. We find that comparing the relative sizes of the RHMC and EOFA forces is highly dependent on whether one chooses F_{RMS} or F_{max}; the mean EOFA F_{RMS} is roughly 30% smaller than the mean RHMC F_{RMS}, but the distributions of F_{max} are nearly indistinguishable. This observation suggests that while the EOFA force distribution may have a smaller mean than the RHMC force distribution, the EOFA distribution also likely has longer tails, such that the largest forces have similar magnitudes. Since we expect the magnitude of the largest forces to correlate more strongly with the efficiency of the integrator than the magnitude of the average forces, as we have argued above, we interpret these results as suggesting that the optimal step size for an EOFA evolution should be similar to that of an RHMC simulation with the same mass parameters, even if the average force is somewhat smaller.

TWQCD has also observed a large hierarchy of scales in the pseudofermion forces associated with each of the two terms in Eqn. (6.34); in Ref. [4] they find that the average force associated with the first term — involving the left-handed component of the pseudofermion field — is more than an order of magnitude smaller than the average force associated with the second term — involving the right-handed component — for two different dynamical QCD simulations. They exploit this observation with a Sexton-Weingarten integration scheme, integrating the first term

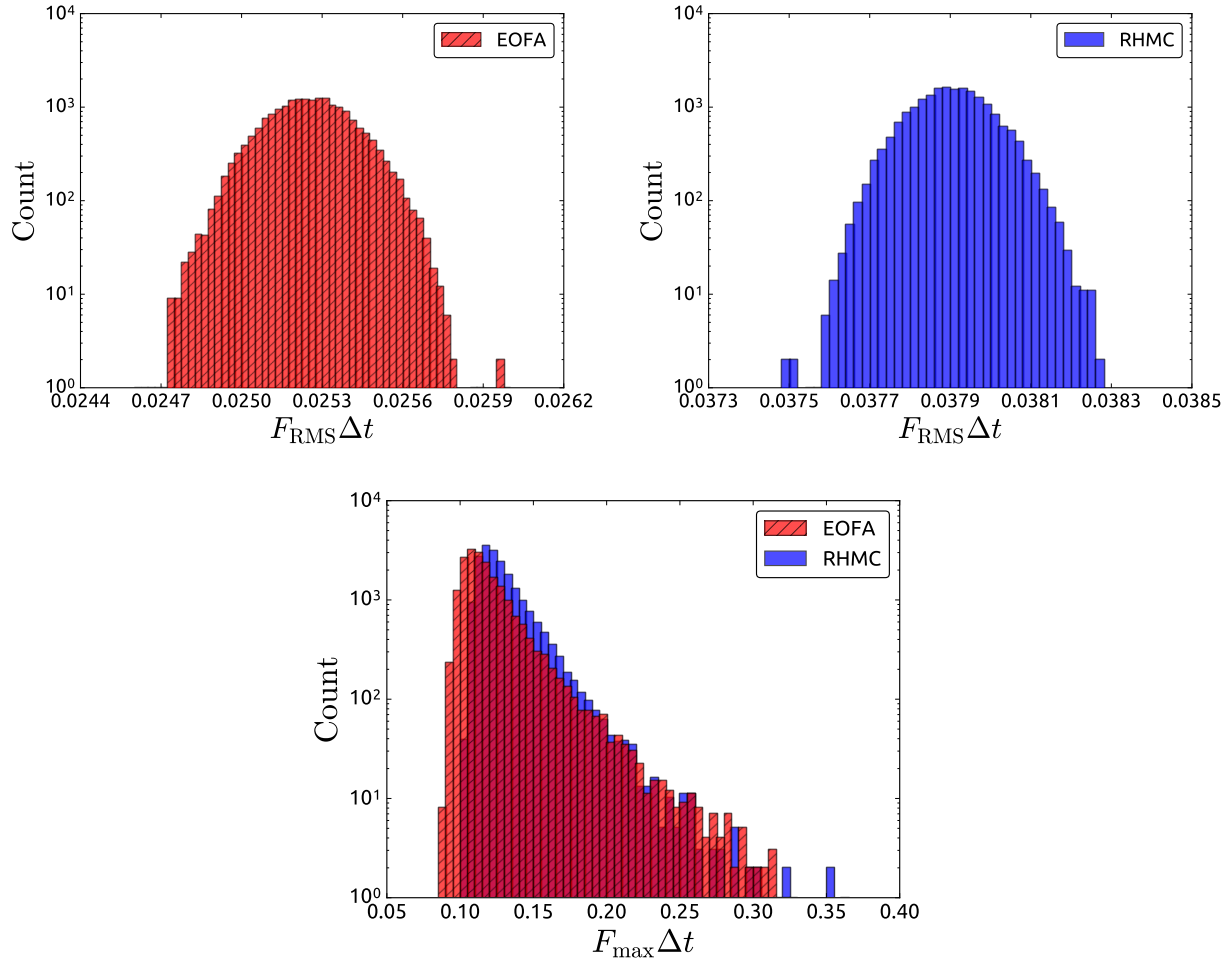

Figure 6.3: Histograms of the RMS and maximum pseudofermion forces associated with force evaluations falling between trajectories 500 and 1500 of the 16I HMC evolutions. F_{RMS} and F_{max} are defined by Equations (6.47) and (6.48), respectively. Δt is the step size used to integrate the pseudofermion force contributions to the HMC evolution.

with a larger time step than the second, and find increased efficiency in their simulations. In Figure 6.4 we compare histograms of the RMS and maximum left-handed and right-handed forces from 1000 thermalized configurations of the 16I EOFA ensemble. Our conclusions are analogous to the comparison between the EOFA and RHMC forces: if one considers F_{RMS} the left-handed force contribution is indeed substantially smaller than the right-handed force contribution, but if one instead considers F_{max} the force distributions are very similar in both magnitude and shape. Based

on the latter observation we leave both terms in Equation (6.34) on the same time step in our large-scale EOFA simulations.

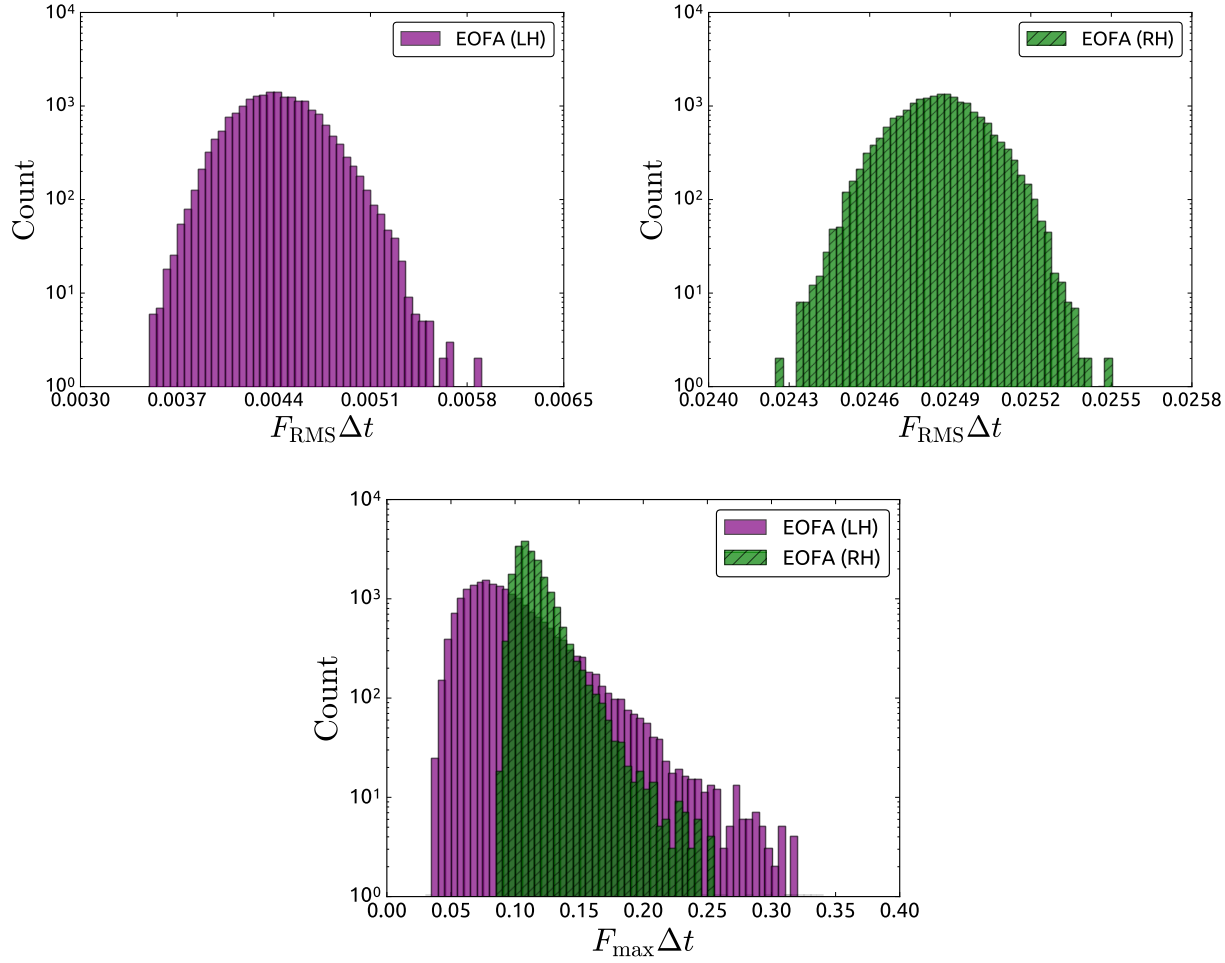

Figure 6.4: Histograms of the RMS and maximum pseudofermion forces associated with the left-handed and right-handed components of the pseudofermion field in Eqn. (6.34).

We also note that neither these small volume test runs, nor TWQCD's simulations in Refs. [3, 4, 10], have considered applying the Hasenbusch mass preconditioning technique [1] to the EOFA formalism. Introducing a set of Hasenbusch masses $\{m'_i\}_{i=1}^{N}$, with $m_1 < m'_i < m_2$, we can write the fermion determinant as

$$\det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m_2)}\right) = \det\left(\frac{\mathscr{D}(m_1)}{\mathscr{D}(m'_1)}\right)\left[\prod_{i=1}^{N-1}\det\left(\frac{\mathscr{D}(m'_i)}{\mathscr{D}(m'_{i+1})}\right)\right]\det\left(\frac{\mathscr{D}(m'_N)}{\mathscr{D}(m_2)}\right). \tag{6.49}$$

While the left-hand side can be simulated using a single pseudofermion field, the associated forces can be large if $m_1 \ll m_2$, requiring a small step size to maintain reasonable acceptance. The right-hand side, in contrast, involves $N+1$ independent pseudofermion fields, but with possibly substantially reduced forces, allowing larger step sizes to be used. For light m_1 one typically observes that the gain from increasing the step size offsets the cost of simulating extra heavy flavors, leading to a more efficient simulation. In Section 6.7 we demonstrate that Hasenbusch preconditioning allows for a substantial speed-up in the context of the 32ID-G ensemble. We also note that in addition to reducing the size of the pseudofermion forces, the Hasenbusch technique preconditions the EOFA force in the sense that the size hierarchy between the left-handed and right-handed force contributions to a single determinant disappears in the limit $m'_i \to m'_{i+1}$. In practice, we find that the mass preconditioned simulation has comparable left-handed and right-handed force contributions even in the RMS sense.

6.6 Optimization and Tuning

In this section we discuss preconditioning and algorithmic techniques which reduce the cost of EOFA simulations. In some cases these are extensions of well-known lattice techniques to the EOFA formalism, while in other cases they are specific to EOFA. We illustrate these techniques using bechmark tests computed with the physical quark mass, Möbius DWF 24ID ensemble, and report timing results for code written in the Columbia Physics System (CPS) and running on 256-node or 512-node Blue Gene/Q partitions.

6.6.1 Inversions of $\mathscr{D}_{\text{EOFA}}$

Since the majority of the computational effort in an HMC simulation is associated with repeatedly inverting the Dirac operator, techniques to more efficiently apply the Dirac operator or to otherwise accelerate these inversions can have a dramatic impact on the overall efficiency of the integrator. To address the former, we make use of the BAGEL assembler generation library [32] to produce highly optimized kernels and fermion solvers for the Blue Gene/Q hardware. To address the latter, we make use of multiple preconditioning techniques, as well as a mixed precision defect correction

CG solver.

The first preconditioning technique we apply — "even-odd" or "red-black" preconditioning — is well-known in the lattice QCD community. Lattice sites are labeled as even if $(x+y+z+t) \equiv 0 \pmod 2$, or odd if $(x+y+z+t) \equiv 1 \pmod 2$, inducing a 2×2 block structure on fermion operators

$$M = \begin{pmatrix} M_{ee} & M_{eo} \\ M_{oe} & M_{oo} \end{pmatrix}. \tag{6.50}$$

Standard tricks can then be used to relate the linear system $M\psi = \phi$ to a better conditioned linear system involving only the odd sub-lattice; this preconditioned system is substantially cheaper to invert since the size of the problem has been halved. After inverting on the odd sub-lattice, the even component of ψ can also be recovered at modest cost, without ever needing to explicitly invert on the even sub-lattice. The details of this construction, and its extension to EOFA, are described in Appendix 6.C.

The second preconditioning technique we apply — Cayley-form preconditioning — is unique to EOFA, and was introduced in Ref. [5]. The generic linear system one needs to solve in the context of EOFA has the form

$$\Big(H(m_1) + \beta\Delta_\pm(m_2, m_3)P_\pm\Big)\psi = \phi, \tag{6.51}$$

where $H = \gamma_5 R_5 \mathscr{D}_{\text{EOFA}}$. For Möbius domain wall fermions $\mathscr{D}_{\text{EOFA}}$ is dense in ss', and thus considerably more expensive to invert than $\mathscr{D}_{\text{DWF}}$, which has a tridiagonal ss' stencil, in terms of wall clock time. However, Eqn. (6.6) suggests that Eqn. (6.51) can be related to an equivalent system in terms of $\mathscr{D}_{\text{DWF}}$ by using $\widetilde{\mathscr{D}}^{-1}$ as a preconditioner. We elaborate on the mathematical details in Appendix 6.C.2, and, in particular, demonstrate that $\Delta_\pm\widetilde{\mathscr{D}}$ has a relatively simple, rank-one form, allowing for substantially more efficient EOFA inversions — even when $\beta \neq 0$ — by working with the preconditioned system. This technique also has the advantage that it allows for EOFA simulations which re-use existing high-performance code for applying $\mathscr{D}_{\text{DWF}}$ with little modification.

Finally, we use a restarted, mixed precision defect correction solver to perform the conjugate gradient inversions of the fully preconditioned EOFA system. For memory bandwidth-limited calculations — such as applying the Dirac operator — single precision computations can be performed at

approximately half the cost of full double precision computations. In the defect correction approach to mixed precision CG, the following algorithm is used:

1. Solve the Dirac equation in single precision arithmetic using a reduced stopping tolerance (typically 10^{-4} or 10^{-5}).

2. Compute the current residual using the (single precision) solution in full double precision arithmetic.

3. If the desired final tolerance (typically 10^{-8} or smaller) has been reached, stop. Otherwise, return to step 1, using the residual vector computed in step 2 as the new CG source.

We observe that this algorithm outperforms straight double precision CG by approximately a factor of 2 — as one would expect if the calculation is truly memory bandwidth-limited — provided the local lattice volume on each node is sufficiently large to avoid communications bottlenecks.

In Figure 6.5 we plot the CG residual as a function of the wall clock running time of the inverter for a series of benchmark inversions of Equation (6.51) on the 24ID ensemble. These benchmarks show the inverter performance as we sequentially introduce even-odd preconditioning, Cayley-form preconditioning, and finally, mixed precision CG. We also plot the time required to solve the family of linear systems

$$\left(\mathscr{D}_{\mathrm{DWF}}^{\dagger}\mathscr{D}_{\mathrm{DWF}} + \beta_k\right)\psi = \phi \tag{6.52}$$

using multishift CG for the same set of poles $\{\beta_k\}$ used in the rational approximation to $x^{-1/2}$ in the RHMC evolution that generated the 24ID ensemble. This allows a baseline estimate of the cost of evaluating the EOFA Hamiltonian or pseudofermion force against the cost of evaluating the RHMC Hamiltonian or pseudofermion force at the same quark mass. We observe a factor of 3.9 speed-up for fully preconditioned EOFA over the even-odd preconditioned RHMC system. In both cases the underlying operator being inverted is $\mathscr{D}_{\mathrm{DWF}}$; the slower RHMC benchmark demonstrates the overhead associated with multishift CG relative to solving a single system with standard CG, both due to the inability to fully utilize mixed precision methods and due to the additional linear algebra required at each iteration.

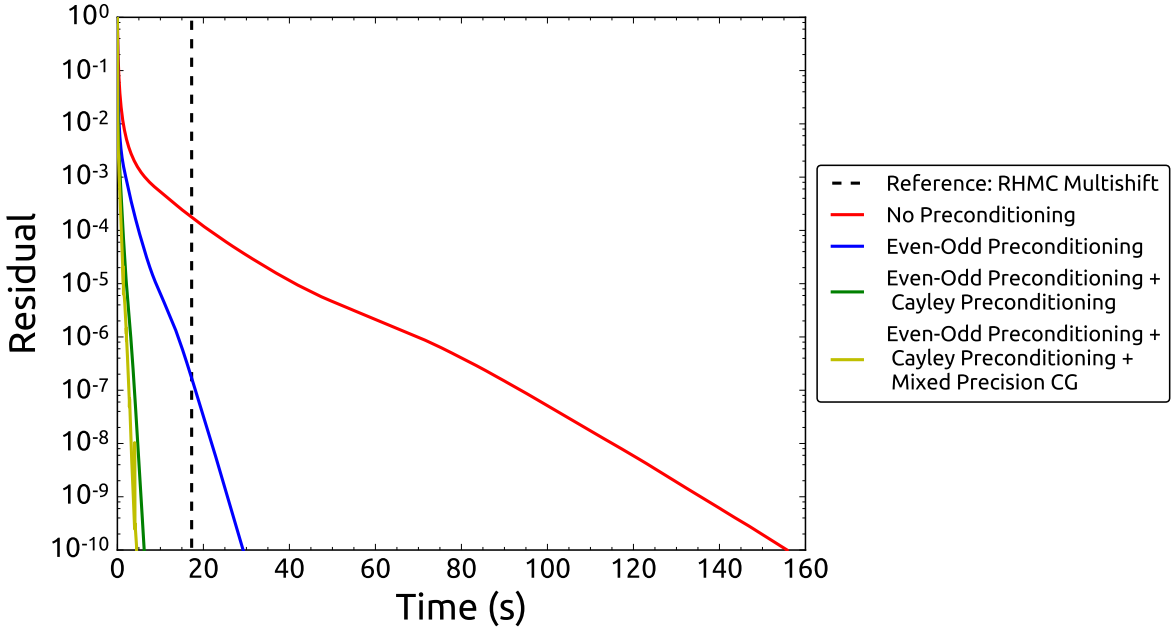

Figure 6.5: Wall clock time required to solve Eqn. (6.51) to a stopping tolerance of 10^{-10} at the physical strange quark mass on the 24ID ensemble, as the preconditioning and algorithmic refinements discussed in the text are introduced sequentially. The dashed vertical line corresponds to the time required to apply $(\mathscr{D}^{\dagger}_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}})^{-1/2}$ by solving Equation (6.52) using the high-performance implementations of even-odd preconditioned $\mathscr{D}_{\mathrm{DWF}}$ and multishift CG in the BAGEL library.

6.6.2 Heatbath

Achieving the full performance improvement suggested by the inversion benchmarks in Section 6.6.1 is complicated by the form of the EOFA heatbath, which is expected to be more expensive than the RHMC heatbath, even with efficient EOFA code. Applying $\mathcal{M}^{-1/2}_{\mathrm{EOFA}}$ (Eqn. (6.29)) requires two independent CG inversions per pole used in the rational approximation to $x^{-1/2}$, since multishift CG is not applicable: we use two algorithmic techniques to reduce this cost. The first is a forecasting technique initially proposed by Brower et al. [33] in the context of more general HMC simulations, and later used successfully by TWQCD in the context of the EOFA heatbath [3]. The idea is the following: given a set of solutions $\{\psi_k\}_{k=1}^{N}$ to Equation (6.51) for N different poles $\{\beta_k\}_{k=1}^{N}$, one

can use the linear combination

$$\psi_{N+1} = \sum_{k=1}^{N} c_k \psi_k \tag{6.53}$$

minimizing the functional

$$\Phi\left[\psi\right] = \psi^\dagger \Big(H + \beta_{N+1} \Delta_\pm P_\pm \Big) \psi - \phi^\dagger \psi - \psi^\dagger \phi \tag{6.54}$$

as the initial CG guess for the next inversion with pole β_{N+1}. The coefficients c_k satisfy

$$\sum_{k=1}^{N} c_k \psi_l^\dagger \Big(H + \beta_{N+1} \Delta_\pm P_\pm \Big) \psi_k = \psi_l^\dagger \phi, \tag{6.55}$$

and can be computed explicitly using *e.g.* Gauss-Jordan elimination. Since Equation (6.54) is the same functional minimized by the conjugate gradient algorithm itself, accurate initial guesses can be computed for modest N provided the $\{\beta_k\}_{k=1}^{N+1}$ are similar in magnitude. In Figure 6.6 we test this forecasting technique using the 24ID ensemble and a rational approximation with 8 poles, and find that the iteration count required to solve Eqn. (6.51) to a tolerance of 10^{-10} is more than halved for the last few poles.

The second technique we have used to accelerate the heatbath is motivated by observing that the coefficients entering into Equation (6.29) span several orders of magnitude for a typical rational approximation to $x^{-1/2}$. We find typical values $k\alpha_l\gamma_l^2/\alpha_0 \sim \mathcal{O}(10^{-3} - 10^{-5})$, suggesting that the inversions can be performed with reduced stopping tolerances relative to the desired accuracy of $\mathcal{M}_{\text{EOFA}}^{-1/2}\psi$, since the solution vectors are ultimately multiplied by small coefficients when the result is formed. We have explored the following simple optimization scheme to relax the stopping conditions for each pole:

1. Choose a desired tolerance for the heatbath, ε_{tol}, where ε is defined by Equation (6.30).

2. Choose one of the inversions required to compute $\mathcal{M}_{\text{EOFA}}^{-1/2}$ according to Equation (6.29), and relax the stopping tolerance until the overall error in the heatbath ε reaches ε_{tol}.

3. Iterate over each inversion until all stopping conditions have been tuned.

We report results for the 24ID ensemble in Table 6.9. Using a rational approximation with 6 poles, and $\varepsilon_{\text{tol}} = 10^{-10}$, we observe that the total heatbath time is more than halved while only slightly

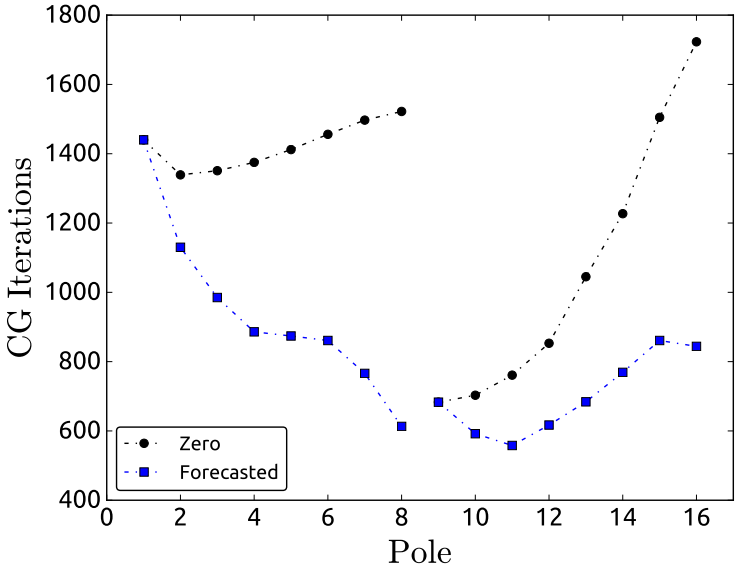

Figure 6.6: CG iterations required to invert Equation (6.51) for each of the 16 values of β entering into a rational approximation of $\mathcal{M}_{\text{EOFA}}^{-1/2}$ with 8 poles on the 24ID ensemble. The first 8 poles ($\beta = -\gamma_l$) are associated with the first (LH) term in Equation (6.29), while the second 8 poles ($\beta = -\beta_l\gamma_l$) are associated with the second (RH) term. We find no improvement from using solutions to the LH system to forecast solutions to the RH system and vice-versa, since the Dirac operator being inverted in either case is evaluated with a different quark mass.

increasing the error. We have also checked that the final error and heatbath running time after tuning is insensitive to the exact order in which the stopping tolerances are tuned.

	ε	Total Heatbath Time
Untuned	1.52×10^{-11}	129.5 s
Tuned	7.79×10^{-11}	68.9 s

Table 6.9: The relative error (ε) and total running time for the EOFA heatbath on the 24ID ensemble before and after applying the tuning algorithm discussed in the text.

6.7 Large-Scale EOFA Calculations

In this section we turn to two ongoing ensemble generation calculations currently being performed by the RBC/UKQCD collaboration. The first is a strong-coupling $N_f = 2 + 1$ $24^3 \times 64 \times 24$ Iwasaki+DSDR lattice (24ID) intended for exploratory studies and calculations requiring high statistics [23]. The second (32ID) has been used for a first-principles calculation of the ratio of Standard Model CP-violation parameters ϵ'/ϵ from $\Delta I = 1/2$ $K \to \pi\pi$ decays in Ref. [6]. RBC/UKQCD is currently generating more gauge field configurations to reduce the statistical errors in the $\Delta I = 1/2$ decay amplitudes. Both ensembles have physical quark masses and large volumes, allowing for tests of the performance of EOFA in the context of state-of-the-art domain wall fermion calculations.

Tables 6.10 and 6.11 summarize the details of the integrator parameters and nesting for these evolutions. The ensembles labeled RHMC correspond to the evolutions of Ref. [23] (24ID) and Ref. [6] (32ID-G). For the ensembles marked EOFA, we have changed the strange quark (light quark) action to EOFA for the 24ID (32ID-G) ensemble and retuned the details of the integrator as described in the remainder of the section. For the 32ID-G ensemble — where, due to the G-parity flavor doubling, the EOFA action naturally describes the degenerate light quark pair — we have also switched from an Omelyan integrator to a force gradient integrator, and inserted additional Hasenbusch preconditioning determinants.

6.7.1 24ID Ensemble

We use the 24ID ensemble as a straightforward benchmark of RHMC against an equivalent EOFA simulation to describe a physical heavy quark flavor. Here this is the strange quark, but $N_f = 2+1+1$ simulations with dynamical strange and charm quarks are another obvious target of EOFA. We make no serious attempt to retune the integrator after switching to EOFA beyond tuning the heatbath step with the following procedure:

1. Compute the largest and smallest eigenvalues of $\mathcal{M}_{\text{EOFA}}$ (Eqn. (6.10)) for a few thermalized configurations of the gauge field, and use these measurements to inform the bounds of the

Ensemble	Integrator	$\delta\tau$	r_{FG}	r_{MD}	r_{MC}
24ID (RHMC)	Force Gradient QPQPQ	0.0833	10^{-5}	10^{-7}	10^{-10}
24ID (EOFA)	Force Gradient QPQPQ	0.0833	10^{-5}	10^{-7}	10^{-10}
32ID-G (RHMC)	Omelyan ($\lambda = 0.22$)	0.0625	—	10^{-7}	10^{-10}
32ID-G (EOFA)	Force Gradient QPQPQ	0.1667	10^{-5}	10^{-7}	10^{-10}

Table 6.10: Basic integrator and HMC details for the generation of the 24ID and 32ID-G ensembles. We denote the coarsest time step used to evolve the outermost level by $\delta\tau$, and the CG stopping tolerances used for the force gradient forecasting, molecular dynamics, and Monte Carlo steps by r_{FG}, r_{MD}, and r_{MC}, respectively. We elaborate on the details of the integrator nesting in Table 6.11.

rational approximations to $x^{-1/2}$ constructed via the Remez algorithm.

2. Add poles to the rational approximation, with all CG stopping tolerances set to r_{MC}, until $\varepsilon < r_{\mathrm{MC}}$ (Eqn. (6.30)) is reached.

3. With the rational approximation now fixed from step 2, tune the CG stopping tolerances corresponding to each pole, following the procedure described in Section 6.6.2, and keeping $\varepsilon < r_{\mathrm{MC}}$.

After tuning the heatbath, we then ran a single trajectory of the RHMC evolution and the EOFA evolution on a 256-node Blue Gene/Q partition. For the EOFA ensemble, we compare two schemes. The first ("dense") is a straightforward implementation of Möbius DWF as proposed in Ref. [3]: we invert Equation (6.51) directly, where $H = \gamma_5 R_5 \mathscr{D}_{\mathrm{EOFA}}$ and the other dense 5D operators appearing in the EOFA action are listed explicitly in Appendix 6.B.2. We also do not apply the final step in our heatbath tuning procedure, leaving all CG stopping tolerances in the heatbath fixed at $r_{\mathrm{MC}} = 10^{-10}$. In the second EOFA scheme ("preconditioned") we fully tune the heatbath step and apply the Cayley-form preconditioning detailed in Appendix 6.C.2 to inversions of Equation (6.51). Timing breakdowns for the strange quark part of the evolution are reported in Table 6.12.

We observe that the dense EOFA formalism is actually somewhat slower than RHMC: the

Ensemble	Level	Action	Update
24ID (RHMC)	1	$\mathrm{RatQuo}_{1/2}(0.085, 1.0)$	1:1
	2	$\mathrm{Quo}(0.00107, 0.00789) + \mathrm{Quo}(0.00789, 0.0291) + \mathrm{Quo}(0.0291, 0.095) + \mathrm{Quo}(0.095, 0.3) + \mathrm{Quo}(0.3, 0.548) + \mathrm{Quo}(0.548, 1.0)$	1:1
	3	Gauge + DSDR	1:1
24ID (EOFA)	1	$\mathrm{EOFA}(0.085, 1.0)$	1:1
	2	$\mathrm{Quo}(0.00107, 0.00789) + \mathrm{Quo}(0.00789, 0.0291) + \mathrm{Quo}(0.0291, 0.095) + \mathrm{Quo}(0.095, 0.3) + \mathrm{Quo}(0.3, 0.548) + \mathrm{Quo}(0.548, 1.0)$	1:1
	3	Gauge + DSDR	1:1
32ID-G (RHMC)	1	$\mathrm{RatQuo}_{1/2}(0.0001, 0.007)$	1:1
	2	$\mathrm{RatQuo}_{1/2}(0.007, 1.0) + \mathrm{RatQuo}_{1/4}(0.045, 1.0)$	1:2
	3	DSDR	1:2
	4	Gauge	1:1
32ID-G (EOFA)	1	$\mathrm{EOFA}(0.0001, 0.0058) + \mathrm{EOFA}(0.0058, 0.0149) + \mathrm{EOFA}(0.0149, 0.059) + \mathrm{EOFA}(0.059, 0.177) + \mathrm{EOFA}(0.177, 0.45) + \mathrm{EOFA}(0.45, 1.0) + \mathrm{RatQuo}_{1/4}(0.045, 1.0)$	5:1
	2	DSDR	1:2
	3	Gauge	1:1

Table 6.11: Integrator layouts for the 24ID and 32ID-G ensembles. The notation A:B for the update scheme denotes the number of steps of the next innermost integrator level (A) per step of the current level (B).

additional complexity of the EOFA heatbath, together with the more expensive inversions of the dense 5D operator $\mathscr{D}_{\mathrm{EOFA}}$, negate the expected performance gains from the simpler forms of the Hamiltonian and force evaluations. We emphasize, however, that we have made no attempt to retune the integrator details to optimize for EOFA; TWQCD has shown in Ref. [4] that dense EOFA simulations can outperform RHMC simulations after optimizing the integrator layout for EOFA. After introducing Cayley-form preconditioning — so that we are inverting the tridiagonal operator $\mathscr{D}_{\mathrm{DWF}}$ rather than $\mathscr{D}_{\mathrm{EOFA}}$ when we solve Equation (6.51) — we find that EOFA outperforms RHMC by a significant factor of 3.5.

Step	**RHMC** Time (s)	%	**EOFA** (Dense) Time (s)	%	**EOFA** (Preconditioned) Time (s)	%
Heatbath	42.6	2.7	340.6	15.1	68.9	15.5
Force gradient integration (total)	1485.6	94.8	1840.6	81.8	355.9	80.1
Final Hamiltonian evaluation	39.4	2.5	68.8	3.1	19.8	4.4
Total	1567.6	—	2250.0	—	444.6	—
(Total RHMC) / Total	1.0	—	0.7	—	3.5	—

Table 6.12: Strange quark timings for a single MD trajectory of the 24ID ensemble on a 256-node Blue Gene/Q partition. We compare RHMC to EOFA with ("preconditioned") and without ("dense") Cayley-form preconditioning.

6.7.2 32ID-G Ensemble

One particularly promising feature of EOFA in the context of G-parity ensembles is the potential for aggressive Hasenbusch mass preconditioning of the light quark determinant; this makes the 32ID-G ensemble a particularly interesting case study since the EOFA formalism is used to describe a physical mass light quark pair. In Ref. [6] the RBC/UKQCD collaboration observed that mass preconditioning is not particularly effective for the RHMC light quark determinant, since each molecular dynamics step requires one multishift inversion of $\mathscr{D}^\dagger\mathscr{D}$ evaluated at the numerator quark mass and two multishift inversions of $\mathscr{D}^\dagger\mathscr{D}$ evaluated at the denominator quark mass. The latter two solves become prohibitively expensive if many intermediate masses are introduced, negating the expected gain from integrating the preconditioned pseudofermion forces with larger step sizes. The EOFA force, on the other hand, is no more expensive to evaluate than the force associated with the standard quotient action (Eqn. (6.38)), so it is natural to expect better performance from Hasenbusch preconditioning.

In Table 6.13 we list details of the tuning runs we have used to explore potential schemes for evolving the 32ID-G ensemble with EOFA light quarks. We started by switching from an Omelyan integrator, for which the leading errors are $\mathcal{O}(\delta\tau^2)$, to a force gradient integrator, for which the

leading errors are $\mathcal{O}(\delta\tau^4)$, and studied the effects of inserting mass preconditioning determinants one at a time (runs 1-7). We then identified two promising mass preconditioning schemes — one with four intermediate masses (runs 8-10), and the other with five intermediate masses (runs 11-14) — and continued tuning the step size, CG stopping conditions, and heatbath, to optimize the job time per trajectory and Monte Carlo acceptance. The initial RHMC scheme used in Ref. [6] corresponds to run 1, and the final EOFA scheme we have adopted for our continuing ensemble generation corresponds to run 12.

Run	Integrator Type	Light Hasenbusch Masses	$\Delta\tau$	r_{MD}	N_{traj}	Acceptance	Efficiency
1	**O**	**0.007**	**0.0625**	$\mathbf{10^{-8}}$	**850**	**88%**	—
2	O	—	0.0625	10^{-8}	10	40%	1.2
3	FG	0.043	0.0625	10^{-8}	10	100%	2.0
4	FG	0.018, 0.12	0.0625	10^{-8}	10	100%	1.8
5	FG	0.0118, 0.0412, 0.23	0.0625	10^{-8}	10	100%	1.7
6	FG	0.0075, 0.023, 0.11, 0.4	0.0625	10^{-8}	10	100%	1.7
7	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.0625	10^{-8}	10	100%	1.5
8	FG	0.0103, 0.029, 0.12, 0.41	0.1000	10^{-6}	15	67%	4.0
9	FG	0.0103, 0.029, 0.12, 0.41	0.1000	10^{-7}	20	95%	3.0
10	FG	0.0103, 0.029, 0.12, 0.41	0.1667	10^{-7}	20	75%	4.5
11	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.1000	10^{-6}	40	80%	3.0
12	**FG**	**0.0058, 0.0149, 0.059, 0.177, 0.45**	**0.1667**	$\mathbf{10^{-7}}$	**850**	**93%**	**4.2**
13	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.2000	10^{-7}	60	65%	4.5
14	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.2000	10^{-8}	25	72%	3.9

Table 6.13: HMC details for the production ensemble generation run (1) of Ref. [6], as well as 13 tuning runs after switching to EOFA light quarks (2-14). We use the following notation: "O" denotes the Omelyan integrator, "FG" denotes the force gradient integrator, "N_{traj}" is the number of trajectories generated for the timing run, "acceptance" is the fraction of gauge field configurations which were accepted in the final Monte Carlo step, and "efficiency" is the ratio of the total job time per trajectory for the specified integration scheme to the total job time per trajectory of the scheme used in run 1. Entries in bold correspond to the original RHMC scheme (1) and the final, fully tuned EOFA scheme (12).

We find, in practice, that Hasenbusch mass preconditioning is extremely effective for the EOFA light quark determinant. In addition to reducing the size of the pseudofermion force, we also observe that the largest eigenvalue of the EOFA action, Equation (6.10), decreases rapidly as $m_2 \to m_1$. As a consequence, the heatbath is also less expensive with Hasenbusch preconditioning, since, as we increase the number of intermediate masses, we can simultaneously decrease the range and number of poles entering into the rational approximation used for each determinant. Table 6.14 summarizes the measured spectral range, the heatbath error, and the total heatbath cost for each of the runs 2-7. For this ensemble the first Hasenbusch mass reduces the cost of the heatbath by more than a factor of two, and subsequent Hasenbusch masses essentially leave the cost fixed.

For each of the runs 2-7 we generated ten trajectories, beginning from the same seed configuration, and analyzed the resulting distributions of F_{RMS} and F_{max}. In panel (a) of Figure 6.7 we plot distributions of F_{max} from 850 trajectories of the production RHMC ensemble generation calculation (run 1). Since we are using exactly the same RHMC action for the strange quark on the RHMC and EOFA ensembles, we tune by adjusting the number and magnitude of the intermediate light Hasenbusch masses such that the forces associated with each of the light quark determinants are comparable to the strange quark force. This allows us to simplify the integrator layout to a three-level scheme, with the light and strange quark determinants updated on the same level. We find that four intermediate Hasenbusch masses are sufficient to ensure that the strange quark force is dominant in the sense of F_{RMS}, and that five intermediate Hasenbusch masses are sufficient in the sense of F_{max}. Panel (b) shows an analogous force distribution for the latter mass preconditioning scheme.

In runs 8-10 we explore further tuning of a scheme with four light Hasenbusch masses, and in runs 11-14 we explore further tuning of a scheme with five light Hasenbusch masses. We note that the Monte Carlo acceptance is relatively poor in runs 8-10 — as we argued in Section 6.5, this is consistent with the view that the acceptance should be controlled by the largest integration errors accrued during the trajectory, which are proportional to F_{max} rather than F_{RMS} — and thus have abandoned this mass preconditioning scheme in favor of the scheme used in runs 11-14. We have then tuned the step size of the outermost integrator level ($\delta\tau$) and the CG stopping tolerance used

N_{LHSB}	Mass Ratio	$\lambda_{\min}$	$\lambda_{\max}$	N_{poles}	ε	Δt_{HB} (s)
0	0.0001/1.0	1.0	1150	11	6.91×10^{-11}	5263.2
1	0.0001/0.043	1.0	33.3	7	3.50×10^{-11}	2226.6
	0.043/1.0	1.0	22.8	7	6.82×10^{-12}	
2	0.0001/0.018	1.0	13.5	6	2.13×10^{-11}	2043.8
	0.018/0.12	1.0	6.4	5	1.11×10^{-11}	
	0.12/1.0	1.0	8.3	6	6.18×10^{-12}	
3	0.0001/0.0118	1.0	8.9	6	6.08×10^{-12}	2307.8
	0.0118/0.0412	1.0	3.3	4	4.09×10^{-11}	
	0.0412/0.23	1.0	5.5	5	1.29×10^{-11}	
	0.23/1.0	1.0	4.3	5	1.11×10^{-11}	
4	0.0001/0.0075	1.0	5.9	5	9.63×10^{-12}	2080.7
	0.0075/0.023	1.0	2.8	4	3.55×10^{-11}	
	0.023/0.11	1.0	4.6	5	1.00×10^{-11}	
	0.11/0.4	1.0	3.6	4	1.98×10^{-11}	
	0.4/1.0	1.0	2.5	4	2.34×10^{-11}	
5	0.0001/0.0058	1.0	4.7	5	1.11×10^{-11}	2289.0
	0.0058/0.0149	1.0	2.3	4	1.64×10^{-11}	
	0.0149/0.059	1.0	3.7	4	9.65×10^{-11}	
	0.059/0.177	1.0	3.0	4	4.14×10^{-11}	
	0.177/0.45	1.0	2.5	4	2.71×10^{-11}	
	0.45/1.0	1.0	2.2	4	1.64×10^{-11}	

Table 6.14: Measured spectral range of $\mathcal{M}_{\text{EOFA}}$, heatbath relative error (ε), and total time for the heatbath step (Δt_{HB}), using N_{LHSB} intermediate mass preconditioning steps and an order N_{poles} rational approximation to $x^{-1/2}$, with all CG stopping tolerances set to $r_{\text{MC}} = 10^{-10}$. Timings are reported for a 512-node Blue Gene/Q partition.

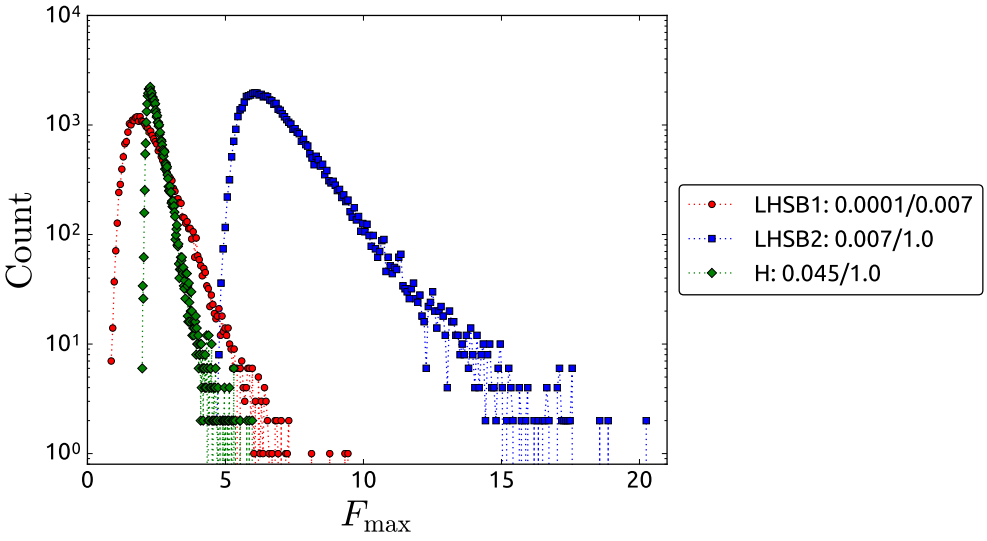

(a) RHMC Ensemble

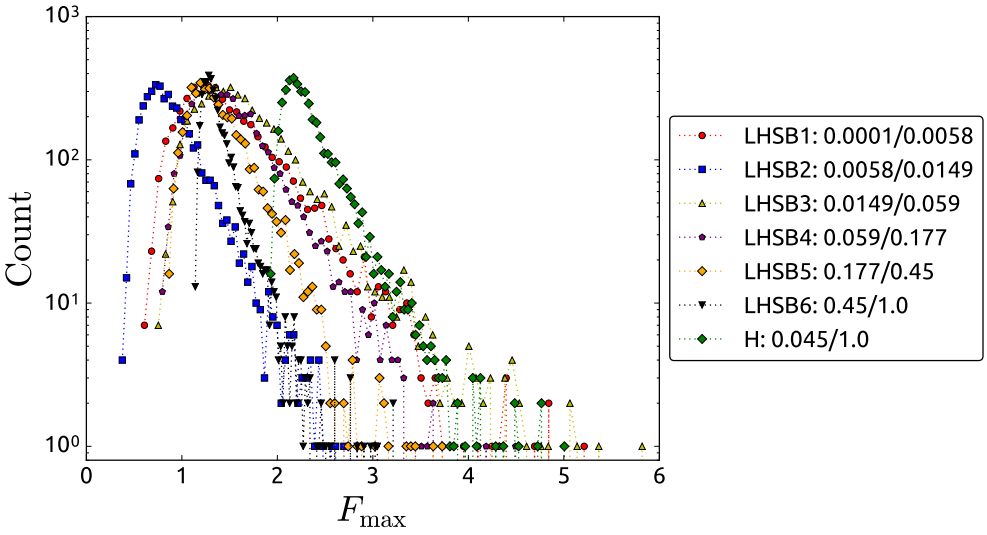

(b) EOFA Ensemble

Figure 6.7: Histograms of the maximum force, defined by Equation (6.48), measured between trajectories 500 and 1350 on the 32ID-G RHMC ensemble and measured between trajectories 1350 and 2200 on the 32ID-G EOFA ensemble. We use the abbreviation "LHSB" in the legends to denote the various mass ratios entering into our mass preconditioning scheme for the light quark determinant, and "H" to denote the strange quark determinant.

in the molecular dynamics evolution (r_{MD}) to minimize the mean time required to generate an accepted gauge field configuration, resulting in the scheme of run 12. In addition, we have applied the heatbath tuning procedure described in Section 6.6.2 in all of the runs 8-14, allowing us to relax CG stopping tolerances for the individual solves in the heatbath, while keeping the overall error bounded by $r_{\mathrm{MC}} = 10^{-10}$. For the final scheme (12) this optimization further reduced the cost of the light quark heatbath from approximately 2300 s, as reported in Table 6.14, to approximately 850 s after tuning.

Comparing the fully tuned EOFA scheme (12) to the original RHMC scheme (1) in Table 6.13, we find that we are able to generate EOFA trajectories a factor of 4.2 times faster than RHMC trajectories, while maintaining a slightly higher acceptance rate of 93%. We emphasize, however, that this improved performance is only partially attributable to the simpler form of the EOFA Hamiltonian and force evaluations: we have also switched from an Omelyan integrator to a force gradient integrator, retuned the step sizes and integrator layout, and, in some cases, applied optimizations to the EOFA simulation that are not applicable to RHMC simulations (*e.g.* mixed precision CG). Figure 6.8 briefly summarizes the respective techniques used in the RHMC and EOFA evolution schemes. We have now adopted the EOFA scheme tested in run 12 for ensemble generation in our ongoing $\Delta I = 1/2$ $K \to \pi\pi$ calculation [34]. We expect the resulting performance gain to enable up to four times as many measurements in our current production run as we would have been able to generate using the initial RHMC evolution scheme, enabling a significantly more precise first-principles determination of the Standard Model ratio ϵ'/ϵ.

RHMC	**EOFA**
• Omelyan integrator ($\delta\tau = 0.0625$)	• Force gradient integrator ($\delta\tau = 0.1667$)
• One light quark Hasenbusch mass	• Five light quark Hasenbusch masses
• Multishift CG with single precision $\not{D}$ but accumulating solution and search vectors in double precision, coupled with reliable update to correct residual	• Mixed precision defect correction CG
	• Even-odd preconditioning
	• Cayley-form preconditioning
	• Force gradient forecasting [9]
• Even-odd preconditioning	• Heatbath forecasting
	• Heatbath stopping tolerance tuning

Figure 6.8: Comparison of optimizations used in the RHMC 32ID-G simulation to the optimizations used in the EOFA 32ID-G simulation.

6.8 Conclusion

In this work we have explored the viability of the exact one flavor algorithm (EOFA) as an alternative to the rational Hybrid Monte Carlo (RHMC) algorithm in molecular dynamics simulations of lattice QCD with domain wall fermions and periodic or G-parity boundary conditions. We have verified the formal equivalence of EOFA to RHMC through statistical tests of the EOFA action (Section 6.4), and checked, using a series of inexpensive, small volume ensembles with heavy pions, that physical observables such as the plaquette, quark condensates, topological susceptibility, and low energy spectrum are consistent between ensembles generated using EOFA and ensembles generated using RHMC (Section 6.5). We have then discussed preconditioning and tuning techniques for EOFA simulations (Section 6.6 and Appendix 6.C), and finally, demonstrated that EOFA can substantially outperform RHMC for state-of-the-art lattice QCD simulations with large volumes

and physical quark masses (Section 6.7). In particular, we find that we are able to generate gauge field configurations for the ongoing RBC/UKQCD calculation of the $\Delta I = 1/2$ $K \to \pi\pi$ decay amplitudes a factor of 4.2 times faster with EOFA. The keys to this dramatic speed-up are a novel preconditioning technique which relates inversions of the EOFA Dirac operator ($\mathscr{D}_{\text{EOFA}}$) to cheaper inversions of the standard domain wall fermion Dirac operator ($\mathscr{D}_{\text{DWF}}$), and the ability to apply mixed precision defect correction solvers and extensive Hasenbusch mass preconditioning in the context of EOFA.

Future work will explore further physics applications of EOFA. We intend to generate variants of the 24ID ensemble with non-degenerate up and down quark masses in the near future. These ensembles will enable exploratory studies of isospin breaking effects in the meson and baryon spectra, as well as in other precision lattice calculations such as the extraction of the CKM matrix element V_{us} from semileptonic kaon decays [35]. Other potential applications include domain wall QCD simulations with dynamical charm quarks in the sea, and simulations with light, $SU(3)$-symmetric quarks. The latter simulations could be used, for example, to better constrain the strange quark dependence of our $SU(3)$ chiral perturbation theory studies [36], or to probe the location of the critical point separating the crossover and first-order phase transition regions in three-flavor domain wall QCD at finite temperature.

Acknowledgments

We thank members of the RBC/UKQCD Collaboration for helpful discussions and support of this work, and are particularly grateful to Peter Boyle and Norman Christ. Calculations were performed using the Blue Gene/Q computers of the RIKEN-BNL Research Center and Brookhaven National Lab. The software used includes the CPS QCD code (`https://github.com/RBC-UKQCD/CPS`) [37], supported in part by the USDOE SciDAC program, the BAGEL (`http://www2.ph.ed.ac.uk/~paboyle/bagel/Bagel.html`) assembler kernel generator for high-performance optimized kernels and fermion solvers [32], and the Grid data parallel C++ QCD library (`https://github.com/paboyle/Grid`) [38]. R.D.M. and D.J.M. are supported in part by U.S. DOE grant #DE-SC0011941. C.J. is supported in part by U.S. DOE Contract #AC-02-98CH10886 (BNL).

References

[1] M. Hasenbusch and K. Jansen, "Speeding up Lattice QCD Simulations with Clover Improved Wilson Fermions", Nucl. Phys. **B659**, 299–320 (2003).

[2] G. McGlynn, "Advances in Lattice Quantum Chromodynamics", PhD thesis (Columbia University, 2016).

[3] Y.-C. Chen and T.-W. Chiu, "Exact Pseudofermion Action for Monte Carlo Simulation of Domain-Wall Fermion", Phys. Lett. **B738**, 55–60 (2014).

[4] Y.-C. Chen and T.-W. Chiu, "One-Flavor Algorithms for Simulation of Lattice QCD with Domain-Wall Fermion: EOFA versus RHMC", PoS **IWCSE2013**, 059 (2014).

[5] D. J. Murphy, "Domain Wall Fermion Simulations with the Exact One-Flavor Algorithm", PoS **LAT2016**, 272 (2016).

[6] Z. Bai et al., "Standard Model Prediction for Direct CP Violation in $K \to \pi\pi$ Decay", Phys. Rev. Lett. **115**, 212001 (2015).

[7] T. Blum et al., "$K \to \pi\pi$ $\Delta I = 3/2$ Decay Amplitude in the Continuum Limit", Phys. Rev. **D91**, 074502 (2015).

[8] C. Kelly, "Progress Towards an Ab Initio, Standard Model Calculation of Direct CP-Violation in K-Decays", in Meeting of the APS Division of Particles and Fields (DPF 2013) Santa Cruz, California, USA, August 13-17, 2013 (2013).

[9] H. Yin and R. D. Mawhinney, "Improving DWF Simulations: the Force Gradient Integrator and the Möbius Accelerated DWF Solver", PoS **LATTICE2011**, 051 (2011).

[10] K. Ogawa, T.-W. Chiu, and T.-H. Hsieh, "One-Flavor Algorithm for Wilson and Domain-Wall Fermions", PoS **LAT2009**, 033 (2009).

[11] T. Blum et al., "Domain Wall QCD with Physical Quark Masses", Phys. Rev. **D93**, 074505 (2016).

[12] R. C. Brower, H. Neff, and K. Orginos, "Mobius Fermions: Improved Domain Wall Chiral fermions", Nucl. Phys. Proc. Suppl. **140**, [,686(2004)], 686–688 (2005).

[13] Y. Iwasaki and T. Yoshié, "Renormalization Group Improved Action for $SU(3)$ Lattice Gauge Theory and the String Tension", Physics Letters B **143**, 449–452 (1984).

[14] P. M. Vranas, "Gap Domain Wall Fermions", Phys. Rev. D **74**, 034512 (2006).

[15] D. Renfrew, T. Blum, N. Christ, R. Mawhinney, and P. Vranas, "Controlling Residual Chiral Symmetry Breaking in Domain Wall Fermion Simulations", PoS **LATTICE2008**, 048 (2008).

[16] D. B. Kaplan, "A Method for Simulating Chiral Fermions on the Lattice", Physics Letters B **288**, 342–347 (1992).

[17] Y. Shamir, "Chiral Fermions from Lattice Boundaries", Nuclear Physics B **406**, 90–106 (1993).

[18] R. Brower, H. Neff, and K. Orginos, "Möbius Fermions: Improved Domain Wall Chiral Fermions", Nuclear Physics B - Proceedings Supplements **140**, 686–688 (2005).

[19] R. Brower, H. Neff, and K. Orginos, "Möbius Fermions", Nuclear Physics B - Proceedings Supplements **153**, 191–198 (2006).

[20] R. C. Brower, H. Neff, and K. Orginos, "The Möbius Domain Wall Fermion Algorithm", (2012).

[21] C. Allton et al., "2+1 Flavor Domain Wall QCD on a $(2\text{ fm})^3$ Lattice: Light Meson Spectroscopy with $L_s = 16$", Phys. Rev. **D76**, 014504 (2007).

[22] R. Arthur et al., "Domain Wall QCD with Near-Physical Pions", Phys. Rev. **D87**, 094514 (2013).

[23] R.D. Mawhinney, "Scaling and Properties of $1/a = 1$ GeV, $2 + 1$ Flavor Möbius Domain Wall Fermion Ensembles", PoS **LATTICE2016**, 136 (2016).

[24] Z. Bai, T. Blum, P. A. Boyle, N. H. Christ, J. Frison, N. Garron, T. Izubuchi, C. Jung, C. Kelly, C. Lehner, R. D. Mawhinney, C. T. Sachrajda, A. Soni, and D. Zhang, "Standard Model Prediction for Direct CP Violation in $K \to \pi\pi$ Decay", Phys. Rev. Lett. **115**, 212001 (2015).

[25] P. A. Boyle et al., "Low Energy Constants of $SU(2)$ Partially Quenched Chiral Perturbation Theory from $N_f = 2 + 1$ Domain Wall QCD", Phys. Rev. **D93**, 054502 (2016).

[26] H. J. Rothe, *Lattice Gauge Theories: An Introduction*, World Scientific Lecture Notes in Physics (World Scientific, Singapore, 2012).

[27] M. A. Clark, "The Rational Hybrid Monte Carlo Algorithm", PoS **LAT2006**, 004 (2006).

[28] Y. Aoki et al., "Continuum Limit Physics from $2 + 1$ Flavor Domain Wall QCD", Phys. Rev. **D83**, 074508 (2011).

[29] P. de Forcrand, M. G. Pérez, and I.-O. Stamatescu, "Topology of the $SU(2)$ Vacuum: A Lattice Study Using Improved Cooling", Nuclear Physics B **499**, 409–449 (1997).

[30] M. Albanese et al., "Glueball Masses and String Tension in Lattice QCD", Physics Letters B **192**, 163–169 (1987).

[31] T. Blum, T. Izubuchi, and E. Shintani, "New class of variance-reduction techniques using lattice symmetries", Phys. Rev. D **88**, 094503 (2013).

[32] P. Boyle, "The BAGEL Assembler Generation Library", Comput.Phys.Commun. **180**, 2739–2748 (2009).

[33] R. C. Brower, T. Ivanenko, A. R. Levi, and K. N. Orginos, "Chronological Inversion Method for the Dirac Matrix in Hybrid Monte Carlo", Nucl. Phys. **B484**, 353–374 (1997).

[34] C. Kelly, "Progress in the Calculation of ϵ' on the Lattice", PoS **LAT2016**, 308 (2016).

[35] P. A. Boyle et al., "The Kaon Semileptonic Form Factor in $N_f = 2 + 1$ Domain Wall Lattice QCD with Physical Light Quark Masses", JHEP **06**, 164 (2015).

[36] R. D. Mawhinney and D. J. Murphy, "NLO and NNLO Low Energy Constants for $SU(3)$ Chiral Perturbation Theory", PoS **LATTICE2015**, 062 (2016).

[37] C. Jung (RBC and UKQCD Collaborations), "Overview of Columbia Physics System", PoS **LAT2013**, 417 (2014).

[38] P. Boyle, A. Yamaguchi, G. Cossu, and A. Portelli, "Grid: A Next Generation Data Parallel C++ QCD Library", (2015).

[39] B. N. Datta, *Numerical linear algebra and applications, second edition*, 2nd (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2010).

6.A Derivation of the Exact One Flavor Action

6.A.1 Preliminaries

Schur Complement

Theorem 6.A.1. (Schur Decomposition)

Let

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \tag{F.56}$$

with A invertible. Then

$$\det(M) = \det(A)\det(D - CA^{-1}B). \tag{F.57}$$

*$W_A \equiv D - CA^{-1}B$ is the **Schur complement of** $\mathbf{A}$.*

Likewise, if D is invertible, then

$$\det(M) = \det(D)\det(A - BD^{-1}C) \tag{F.58}$$

*with $W_D \equiv A - BD^{-1}C$ the **Schur complement of** $\mathbf{D}$.*

Proof. If A is invertible, we can factor

$$\begin{aligned} \begin{pmatrix} A & B \\ C & D \end{pmatrix} &= \begin{pmatrix} A & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & A^{-1}B \\ C & D \end{pmatrix} \\ &= \begin{pmatrix} A & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ C & 1 \end{pmatrix} \begin{pmatrix} 1 & A^{-1}B \\ 0 & D - CA^{-1}B \end{pmatrix}. \end{aligned} \tag{F.59}$$

Eqn. (F.57) follows immediately by taking the determinant. The proof for the second form — in terms of the Schur complement of D — is analogous. □

One way to think about this theorem is that it provides the correct generalization of the determinant of a 2×2 matrix

$$\det\left(\begin{pmatrix} a & b \\ c & d \end{pmatrix}\right) = ad - bc = a\left(d - ca^{-1}b\right) = d\left(a - bd^{-1}c\right) \tag{F.60}$$

to block matrices where the blocks have arbitrary dimension. The Schur complements also appear when writing the inverse of a 2×2 block matrix in terms of the inverses of its blocks

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} \left(A - BD^{-1}C\right)^{-1} & -A^{-1}B\left(D - CA^{-1}B\right)^{-1} \\ -D^{-1}C\left(A - BD^{-1}C\right)^{-1} & \left(D - CA^{-1}B\right)^{-1} \end{pmatrix} \tag{F.61}$$

as one can check explicitly.

Woodbury Matrix Identity

Theorem 6.A.2. (Woodbury Matrix Identity)

Let A, C, U, and V be matrices such that A, C, and $(C^{-1} + VA^{-1}U)$ are all invertible. Then

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U\left(C^{-1} + VA^{-1}U\right)^{-1}VA^{-1}. \tag{F.62}$$

Proof. We can simply check that multiplying by $(A + UCV)$ on the left indeed gives the identity:

$$\begin{aligned} (A + UCV)\left[A^{-1} - A^{-1}U\left(C^{-1} + VA^{-1}U\right)^{-1}VA^{-1}\right] & \\ = 1 + UCVA^{-1} - \left(U + UCVA^{-1}U\right)\left(C^{-1} + VA^{-1}U\right)^{-1}VA^{-1} & \\ = 1 + UCVA^{-1} - UC\left(C^{-1} + VA^{-1}U\right)\left(C^{-1} + VA^{-1}U\right)^{-1}VA^{-1} & \\ = 1 + UCVA^{-1} - UCVA^{-1} & \\ = 1. & \end{aligned} \tag{F.63}$$

□

Corollary 6.A.3. (Sherman-Morrison Formula).

$$\begin{aligned} (A + B)^{-1} &= A^{-1} - A^{-1}\left(1 + BA^{-1}\right)^{-1}BA^{-1} \\ (A + u \otimes v)^{-1} &= A^{-1} - \frac{A^{-1}(u \otimes v)A^{-1}}{1 + \langle v, A^{-1}u\rangle} \end{aligned} \quad . \tag{F.64}$$

Proof. Both are immediate consequences of the Woodbury identity obtained by setting $U = V = 1$ and $C = (1)$, respectively. □

Note: $(u \otimes v)_{ij} \equiv u_i v_j$ (vector outer product).

Chiu et al. write their single flavor action as

$$\begin{aligned} S_{\text{EOFA}} = &\begin{pmatrix} 0 & \phi_1^\dagger \end{pmatrix} \left[1 - k\Omega_-^\dagger \Big(H(m) \Big)^{-1} \Omega_- \right] \begin{pmatrix} 0 \\ \phi_1 \end{pmatrix} \\ &+ \begin{pmatrix} \phi_2^\dagger & 0 \end{pmatrix} \left[1 + k\Omega_+^\dagger \Big(H(1) - \Delta_+(m) P_+ \Big)^{-1} \Omega_+ \right] \begin{pmatrix} \phi_2 \\ 0 \end{pmatrix} \end{aligned} \tag{F.65}$$

where $H(m) \equiv \gamma_5 R_5 \mathscr{D}_{\text{EOFA}}(m)$ is closely related to the DWF Dirac operator, and the other factors are explained in the derivation below. The action is implicitly written in the chiral representation of the gamma matrices; the ϕ_1 term projects onto the left-handed $(-)$ component and the ϕ_2 term projects onto the right-handed $(+)$ component. The fields ϕ_1 and ϕ_2 each have two spinor components.

6.A.2 Outline

The idea is to manipulate the DWF Dirac operator $\mathscr{D}_{\text{DWF}}(m)$ such that we end up with a positive-definite action describing a single quark flavor which still gives $\det(\mathscr{D}_{\text{DWF}})$ upon integrating out the pseudofermion fields. There are three major steps:

1. Factor out a gauge field independent piece from $\mathscr{D}_{\text{DWF}}(m)$, resulting in a rescaled Dirac operator $(\mathscr{D}_{\text{EOFA}}(m))_{xx';ss'} \equiv (D_W)_{xx'}\delta_{ss'} + \delta_{xx'}(\mathscr{D}^\perp_{\text{EOFA}})_{ss'}$, where D_W is the Wilson Dirac operator. We can then replace $\mathscr{D}_{\text{DWF}}(m)$ with $\mathscr{D}_{\text{EOFA}}(m)$ in the fermion action since this overall factor cancels in a properly normalized path integral.

2. Use the Sherman-Morrison formula, Eqn. (6.A.3), to work out the matrix elements of $\mathscr{D}^\perp_{\text{EOFA}}$ explicitly.

3. Use the relationship between the determinant of a matrix and the determinant of its Schur complements, Eqn. (F.57), to massage the determinant of $\mathscr{D}_{\text{EOFA}}$ into the form

$$\det\left(\frac{\mathscr{D}_{\text{EOFA}}(m_1)}{\mathscr{D}_{\text{EOFA}}(m_2)} \right) = \frac{1}{\det(H_1)} \cdot \frac{1}{\det(H_2)} \tag{F.66}$$

with H_1 and H_2 Hermitian and positive-definite. This is done by treating $\mathscr{D}_{\text{EOFA}}$ as a block matrix in Dirac space and working explicitly in the chiral representation of the gamma matrices. We can then simulate one flavor of DWF with the action $S_{\text{pf}} = \phi_1^\dagger H_1 \phi_1 + \phi_2^\dagger H_2 \phi_2$, which is equivalent to the above form.

6.A.3 Derivation

We start by factoring the DWF Dirac operator:

$$\begin{aligned}
\left[\mathscr{D}_{\text{DWF}}(m)\right]_{xx';ss'} &= \Big((c\omega_s + d) D_W + 1 \Big)_{xx'} \delta_{ss'} + \Big((c\omega_s - d) D_W - 1 \Big)_{xx'} L_{ss'} \\
&= (D_W)_{xx'} \Big[(c\omega_s + d) \delta_{ss'} + (c\omega_s - d) L_{ss'} \Big] + \delta_{xx'} (1 - L)_{ss'} \\
&= \left\{ (D_W)_{xx'} \delta_{ss'} + \delta_{xx'} \left(\Big[d + c\omega_s (1 + L) (1 - L)^{-1} \Big]_{ss'} \right)^{-1} \right\} \\
&\qquad \times \Big[d + c\omega_s (1 + L) (1 - L)^{-1} \Big]_{ss'} .
\end{aligned} \tag{F.67}$$

Since the gauge field only enters into the $\{\cdots\}$ factor through D_W, we can safely drop the $[\cdots]$ terms; their contributions to a properly normalized path integral will cancel. Observing that

$$\begin{aligned}
\frac{1}{d + c\omega (1 + L) (1 - L)^{-1}} &= \frac{P_+ + P_-}{d + c\omega (1 + L) (1 - L)^{-1}} \\
&= P_+ \omega^{-1/2} \frac{1}{\omega^{-1} d + c (1 + L_+) (1 - L_+)^{-1}} \omega^{-1/2} \\
&\qquad + P_- \omega^{-1/2} \frac{1}{\omega^{-1} d + c (1 + L_-) (1 - L_-)^{-1}} \omega^{-1/2},
\end{aligned} \tag{F.68}$$

where we have repeatedly used the orthogonality of the projection operators, for example to manipulate

$$\begin{aligned}
\frac{1 + L}{1 - L} &= \frac{1 + P_+ L_+ + P_- L_-}{1 - P_+ L_+ - P_- L_-} \\
&= \frac{(P_+ + P_-) (1 + P_+ L_+ + P_- L_-)}{1 - P_+ L_+ - P_- L_-} \\
&= P_+ \left(\frac{1 + L_+}{1 - L_+} \right) + P_- \left(\frac{1 + L_-}{1 - L_-} \right),
\end{aligned} \tag{F.69}$$

we end up with Chiu et al.'s rescaled Dirac operator (eqns. (3)-(5) of [3]):

$$\begin{aligned}
\left[\mathscr{D}_{\text{EOFA}}(m)\right]_{xx';ss'} &\equiv (D_W)_{xx'} \delta_{ss'} + \delta_{xx'} (P_+ M_+)_{ss'} + \delta_{xx'} (P_- M_-)_{ss'} \\
M_\pm &= \omega^{-1/2} \left[\omega^{-1} d + c N_\pm \right]^{-1} \omega^{-1/2} \\
N_\pm &= (1 + L_\pm) (1 - L_\pm)^{-1}
\end{aligned} \quad . \tag{F.70}$$

Next, we manipulate $M_\pm$ into a more manageable form. Chiu et al. start by claiming

$$N_\pm(m) = N_\pm(0) - \frac{2m}{1+m}\,(u \otimes u)\,, \tag{F.71}$$

where $u = (1;1;1;\ldots)$, and $u \otimes u \equiv uu^\mathsf{T}$ is the vector outer product. This can be checked by explicit computation:

$$\begin{aligned}
N_+(m) &= (1 + L_+(m))\,(1 - L_+(m))^{-1} \\
&= \begin{pmatrix} 1 & 0 & 0 & \cdots & -m \\ 1 & 1 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & \cdots & m \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}^{-1} \\
&= \begin{pmatrix} 1 & 0 & 0 & \cdots & -m \\ 1 & 1 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} \cdot \frac{1}{1+m} \begin{pmatrix} 1 & -m & -m & \cdots & -m \\ 1 & 1 & -m & \cdots & -m \\ 1 & 1 & 1 & \cdots & -m \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix}, \\
&= \frac{1}{1+m} \begin{pmatrix} 1-m & -2m & -2m & \cdots & -2m \\ 2 & 1-m & -2m & \cdots & -2m \\ 2 & 2 & 1-m & \cdots & -2m \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 2 & 2 & 2 & \cdots & 1-m \end{pmatrix} \\
&= \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 2 & 1 & 0 & \cdots & 0 \\ 2 & 2 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 2 & 2 & 2 & \cdots & 1 \end{pmatrix} - \frac{2m}{1+m} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix} \\
&= N_+(0) - \frac{2m}{1+m}\,(u \otimes u)\,,
\end{aligned} \tag{F.72}$$

and similarly for $N_-(m) = N_+(m)^\mathsf{T}$. This allows us to apply the Sherman-Morrison formula, Eqn. (6.A.3), to the rescaled Dirac operator, Eqn. (F.70). Defining

$$A_\pm \equiv \omega^{-1} d + cN_\pm(0), \quad \lambda_\pm \equiv \langle u, A_\pm^{-1} u \rangle \tag{F.73}$$

we have

$$\begin{aligned}
\left[\omega^{-1} d + cN_\pm(m)\right]^{-1} &= \left[\omega^{-1} d + cN_\pm(0) - \frac{2cm}{1+m}(u \otimes u)\right]^{-1} \\
&= \left[A_\pm - \frac{2cm}{1+m}(u \otimes u)\right]^{-1} \\
&= A_\pm^{-1} + \frac{2cm}{1+m} \frac{A_\pm^{-1}(u \otimes u) A_\pm^{-1}}{1 - \frac{2cm}{1+m}\lambda_\pm} \\
&= A_\pm^{-1} + \frac{2cm}{1+m-2cm\lambda_\pm} A_\pm^{-1}(u \otimes u) A_\pm^{-1}.
\end{aligned} \tag{F.74}$$

We can plug this result into Eqn. (F.70) to obtain

$$M_\pm(m) = \omega^{-1/2} A_\pm^{-1} \omega^{-1/2} + \frac{2cm}{1+m-2cm\lambda_\pm} \omega^{-1/2} A_\pm^{-1}(u \otimes u) A_\pm^{-1} \omega^{-1/2}. \tag{F.75}$$

Chiu et al. further manipulate this by noting that since ω is symmetric under a reflection in the s-coordinate — i.e. under the operator $(R_5)_{ss'} = \delta_{s',L_s-s}$ — we have $\omega = R_5 \omega R_5$, and introduce

$$v_\pm \equiv R_5 A_\pm^{-1} u, \tag{F.76}$$

whose components are given explicitly by[2]

$$\begin{aligned}
(v_+)_{L_s-1} &= (v_-)_0 = \alpha_{L_s-1} \\
(v_+)_s &= (v_-)_{L_s-s+1} = \alpha_s \beta_{s+1} (v_+)_{s+1}, \;\; 1 \le s \le L_s - 2,
\end{aligned} \tag{F.77}$$

where $\alpha_s \equiv 1/(d/\omega_s + c)$ and $\beta_s \equiv d/\omega_s - c$. For Möbius DWF these simplify, and we have

$$(v_+)_{L_s-s-1} = (v_-)_s = \frac{(d-c)^s}{(d+c)^{s+1}}, \;\; \lambda_+ = \lambda_- \equiv \lambda = \sum_s \frac{(d-c)^s}{(d+c)^{s+1}}. \tag{F.78}$$

[2]This is worked out by noting that since $A_\pm$ is a triangular matrix, we can use forward/back substitution to solve $A_\pm R_5 v_\pm = u$ explicitly, which gives this recursive formula for $v_\pm$.

Putting this together, and using $\left(A_\pm^{-1}\right)^\mathsf{T} = A_\mp^{-1}$, $R_5 A_\pm^{-1} R_5 = A_\mp^{-1}$, and $R_5 u = u$:

$$\begin{aligned}
\omega^{-1/2} A_\pm^{-1} u u^\mathsf{T} A_\pm^{-1} \omega^{-1/2} &= R_5 \omega^{-1/2} R_5 A_\pm^{-1} u u^\mathsf{T} A_\pm^{-1} R_5 \omega^{-1/2} R_5 \\
&= R_5 \omega^{-1/2} \left(R_5 A_\pm^{-1} u\right) \left(u^\mathsf{T} R_5 A_\mp^{-1} R_5\right) R_5 \omega^{-1/2} R_5 \\
&= R_5 \omega^{-1/2} \left(R_5 A_\pm^{-1} u\right) \left(R_5^\mathsf{T} A_\pm^{-1} R_5 u\right)^\mathsf{T} \omega^{-1/2} \\
&= R_5 \omega^{-1/2} \left(R_5 A_\pm^{-1} u\right) \left(R_5 A_\pm^{-1} u\right)^\mathsf{T} \omega^{-1/2} \\
&= R_5 \omega^{-1/2} v_\pm v_\pm^\mathsf{T} \omega^{-1/2}.
\end{aligned} \tag{F.79}$$

We can now rewrite Eqn. (F.70) in a more explicit form as

$$\begin{aligned}
\left[\mathscr{D}_{\mathrm{EOFA}}(m)\right]_{xx';ss'} &= (D_W)_{xx'} \delta_{ss'} + \delta_{xx'} \left(P_+ M_+\right)_{ss'} + \delta_{xx'} \left(P_- M_-\right)_{ss'} \\
M_\pm(m) &= \omega^{-1/2} A_\pm^{-1} \omega^{-1/2} + \frac{2cm}{1 + m - 2cm\lambda_\pm} R_5 \omega^{-1/2} \left(v_\pm \otimes v_\pm\right) \omega^{-1/2},
\end{aligned} \tag{F.80}$$

where $A_\pm$, $\lambda_\pm$, and $v_\pm$ are defined explicitly above. In Appendix 6.B we collect explicit expressions for $M_\pm$ for the special cases of Shamir and Möbius DWF, for which $\omega = 1$.

For the last step, we manipulate the determinant of Eqn. (F.80) using the Schur decomposition, Eqn. (6.A.1), to arrive at the action we want. We work explicitly in the chiral representation of the (Euclidean) gamma matrices:

$$\gamma_\mu = \begin{pmatrix} 0 & \sigma_\mu \\ \sigma_\mu^\dagger & 0 \end{pmatrix}, \quad \sigma_\mu = (\vec{\sigma}, i1), \quad \{\gamma_\mu, \gamma_\nu\} = 2\delta_{\mu\nu} 1. \tag{F.81}$$

In this representation

$$\gamma_5 \equiv \gamma_1 \gamma_2 \gamma_3 \gamma_4 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{F.82}$$

and the chiral projection operators are

$$P_+ = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad P_- = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \tag{F.83}$$

$\mathscr{D}_{\mathrm{EOFA}}(m)$ is a 2×2 block matrix in spinor space[3]:

$$\mathscr{D}_{\mathrm{EOFA}}(m) = \begin{pmatrix} W - m_0 + M_+(m) & (\sigma \cdot t) \\ -(\sigma \cdot t)^\dagger & W - m_0 + M_-(m) \end{pmatrix} \tag{F.84}$$

[3]The lower-left component picks up a minus sign since t_μ is anti-Hermitian: $\sigma^\dagger \cdot t = -(\sigma \cdot t)^\dagger$.

Taking the determinant, and applying the Schur decomposition (Eqn. (6.A.1)):

$$\begin{aligned}\det\left(\mathscr{D}_{\mathrm{EOFA}}(m)\right) &= \det\left(W - m_0 + M_+(m)\right) \\ &\quad \times \det\left(W - m_0 + M_-(m) + (\sigma\cdot t)^\dagger \left(W - m_0 + M_+(m)\right)^{-1} (\sigma\cdot t)\right) \\ &= \det\left(W - m_0 + M_-(m)\right) \\ &\quad \times \det\left(W - m_0 + M_+(m) + (\sigma\cdot t) \left(W - m_0 + M_-(m)\right)^{-1} (\sigma\cdot t)^\dagger\right).\end{aligned} \tag{F.85}$$

The fermion determinant, after introducing a Pauli-Villars field with $m = 1$, can be written as

$$\frac{\det\left(\mathscr{D}_{\mathrm{EOFA}}(m)\right)}{\det\left(\mathscr{D}_{\mathrm{EOFA}}(1)\right)} = \frac{\det\left(W - m_0 + M_+(m)\right)\cdot\det\left(H_-(m)\right)}{\det\left(W - m_0 + M_-(1)\right)\cdot\det\left(H_+(1)\right)}, \tag{F.86}$$

where[4]

$$\begin{aligned}H_+(m) &\equiv R_5\left[W - m_0 + M_+(m) + (\sigma\cdot t)\left(W - m_0 + M_-(m)\right)^{-1}(\sigma\cdot t)^\dagger\right] \\ H_-(m) &\equiv R_5\left[W - m_0 + M_-(m) + (\sigma\cdot t)^\dagger\left(W - m_0 + M_+(m)\right)^{-1}(\sigma\cdot t)\right].\end{aligned} \tag{F.87}$$

We can use a trick to eliminate the $\det(W - m_0 + M_\pm)$ factors: consider a slightly generalized Dirac operator

$$\mathscr{D}_{\mathrm{EOFA}}(m_1, m_2) = \begin{pmatrix} W - m_0 + M_+(m_1) & (\sigma\cdot t) \\ -(\sigma\cdot t)^\dagger & W - m_0 + M_-(m_2)\end{pmatrix}. \tag{F.88}$$

Applying the Schur decomposition to $\mathscr{D}_{\mathrm{EOFA}}(m, 1)$ gives

$$\begin{aligned}\det\left(\mathscr{D}_{\mathrm{EOFA}}(m,1)\right) &= \det\left(W - m_0 + M_+(m)\right) \\ &\quad \times \det\left(W - m_0 + M_-(1) + (\sigma\cdot t)^\dagger \left(W - m_0 + M_+(m)\right)^{-1} (\sigma\cdot t)\right) \\ &= \det\left(W - m_0 + M_-(1)\right) \\ &\quad \times \det\left(W - m_0 + M_+(m) + (\sigma\cdot t) \left(W - m_0 + M_-(1)\right)^{-1} (\sigma\cdot t)^\dagger\right).\end{aligned} \tag{F.89}$$

[4]We can stick in an overall factor of R_5 for free since $\det(R_5) = 1$.

Taking the ratio of these two equalities

$$
\begin{aligned}
1 &= \frac{\det\left(\mathscr{D}_{\mathrm{EOFA}}(m,1)\right)}{\det\left(\mathscr{D}_{\mathrm{EOFA}}(m,1)\right)} \\
&= \frac{\det\left(W-m_0+M_+(m)\right)\cdot\det\left(W-m_0+M_-(1)+(\sigma\cdot t)^\dagger\left(W-m_0+M_+(m)\right)^{-1}(\sigma\cdot t)\right)}{\det\left(W-m_0+M_-(1)\right)\cdot\det\left(W-m_0+M_+(m)+(\sigma\cdot t)\left(W-m_0+M_-(1)\right)^{-1}(\sigma\cdot t)^\dagger\right)}
\end{aligned}
\tag{F.90}
$$

and rearranging

$$
\begin{aligned}
\frac{\det\left(W-m_0+M_+(m)\right)}{\det\left(W-m_0+M_-(1)\right)} &= \frac{\det\left(W-m_0+M_+(m)+(\sigma\cdot t)\left(W-m_0+M_-(1)\right)^{-1}(\sigma\cdot t)^\dagger\right)}{\det\left(W-m_0+M_-(1)+(\sigma\cdot t)^\dagger\left(W-m_0+M_+(m)\right)^{-1}(\sigma\cdot t)\right)} \\
&= \frac{\det\left(R_5\left[W-m_0+M_+(m)+(\sigma\cdot t)\left(W-m_0+M_-(1)\right)^{-1}(\sigma\cdot t)^\dagger+(M_+(1)-M_+(1))\right]\right)}{\det\left(R_5\left[W-m_0+M_-(1)+(\sigma\cdot t)^\dagger\left(W-m_0+M_+(m)\right)^{-1}(\sigma\cdot t)+(M_-(m)-M_-(m))\right]\right)} \\
&= \frac{\det\left(R_5\left[W-m_0+M_+(1)+(\sigma\cdot t)\left(W-m_0+M_-(1)\right)^{-1}(\sigma\cdot t)^\dagger-(M_+(1)-M_+(m))\right]\right)}{\det\left(R_5\left[W-m_0+M_-(m)+(\sigma\cdot t)^\dagger\left(W-m_0+M_+(m)\right)^{-1}(\sigma\cdot t)+(M_-(1)-M_-(m))\right]\right)} \\
&= \frac{\det\left(H_+(1)-\Delta_+(m)\right)}{\det\left(H_-(m)+\Delta_-(m)\right)},
\end{aligned}
\tag{F.91}
$$

where we have introduced

$$
\Delta_\pm(m) \equiv R_5\left(M_\pm(1)-M_\pm(m)\right). \tag{F.92}
$$

Using Eqn. (F.80) we can write $\Delta_\pm(m)$ explicitly:

$$
\begin{aligned}
\Delta_\pm(m) &= \left(\frac{c}{1-c\lambda}-\frac{2cm}{1+m-2cm\lambda}\right)\omega^{-1/2}\left(v_\pm\otimes v_\pm\right)\omega^{-1/2} \equiv k\Omega_\pm\Omega_\pm^\dagger \\
k &\equiv \frac{c}{1-c\lambda}\cdot\frac{1-m}{1+m-2cm\lambda}, \quad \left(\Omega_\pm\right)_{ss'} \equiv \omega_s^{-1/2}\left(v_\pm\right)_s\delta_{s',0}.
\end{aligned}
\tag{F.93}
$$

Substituting the last expression into Eqn. (F.86)

$$\begin{aligned}\frac{\det\left(\mathscr{D}_{\mathrm{EOFA}}(m)\right)}{\det\left(\mathscr{D}_{\mathrm{EOFA}}(1)\right)} &= \frac{\det\left(H_+(1)-\Delta_+(m)\right)}{\det\left(H_-(m)+\Delta_-(m)\right)}\cdot\frac{\det\left(H_-(m)\right)}{\det\left(H_+(1)\right)} \\ &= \frac{1}{\det\left(1+\Delta_-(m)H_-(m)^{-1}\right)\cdot\det\left(H_+(1)\left(H_+(1)-\Delta_+(m)\right)^{-1}\right)} \\ &= \frac{1}{\det\left(1+k\Omega_-^\dagger\Big(H_-(m)\Big)^{-1}\Omega_-\right)}\cdot\frac{1}{\det\left(1+k\Omega_+^\dagger\Big(H_+(1)-\Delta_+(m)\Big)^{-1}\Omega_+\right)} \\ &\equiv \frac{1}{\det\left(H_1(m)\right)}\cdot\frac{1}{\det\left(H_2(m)\right)},\end{aligned} \tag{F.94}$$

where we have used $\det(1+AB)=\det(1+BA)$ (Sylvester's determinant theorem) to rearrange factors in the $\Delta_\pm H_\pm$ terms.

This last form suggests an action

$$S_{\mathrm{EOFA}} = \phi_1^\dagger H_1(m)\phi_1 + \phi_2^\dagger H_2(m)\phi_2, \tag{F.95}$$

where ϕ_1 and ϕ_2 are two independent pseudofermion fields, each of which contains two spinor components. It turns out to be more useful to re-write this action in block form, however, since H_1 and H_2 contain implicit matrix inverses and thus would be difficult to invert. Define $H(m)\equiv\gamma_5 R_5\mathscr{D}_{\mathrm{EOFA}}(m)$. Explicitly,

$$\begin{aligned}H(m) &= \begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix}\begin{pmatrix}R_5\left(W-m_0+M_+(m)\right) & R_5\left(\sigma\cdot t\right)\\ -R_5\left(\sigma\cdot t\right)^\dagger & R_5\left(W-m_0+M_-(m)\right)\end{pmatrix} \\ &= \begin{pmatrix}R_5\left(W-m_0+M_+(m)\right) & R_5\left(\sigma\cdot t\right)\\ R_5\left(\sigma\cdot t\right)^\dagger & -R_5\left(W-m_0+M_-(m)\right)\end{pmatrix}.\end{aligned} \tag{F.96}$$

Using Eqn. (F.61), the lower right $(--)$ entry of $H(m)^{-1}$ is

$$\begin{aligned}\left[H(m)^{-1}\right]_{--} &= \left(-R_5\left[W-m_0+M_-(m)+(\sigma\cdot t)^\dagger\left(W-m_0+M_+(m)\right)^{-1}(\sigma\cdot t)\right]\right)^{-1} \\ &= -\frac{1}{H_-(m)}.\end{aligned} \tag{F.97}$$

Thus,

$$\begin{pmatrix} 0 & \phi_1^\dagger \end{pmatrix} \left[1 - k\Omega_-^\dagger \Big(H(m)\Big)^{-1} \Omega_-\right] \begin{pmatrix} 0 \\ \phi_1 \end{pmatrix} = \phi_1^\dagger \left[1 + k\Omega_-^\dagger \Big(H_-(m)\Big)^{-1} \Omega_-\right] \phi_1 = \phi_1^\dagger H_1 \phi_1. \tag{F.98}$$

Similarly,

$$H(1) - \Delta_+(m) P_+ = \begin{pmatrix} R_5 \left(W - m_0 + M_+(1)\right) - \Delta_+(m) & R_5 \left(\sigma \cdot t\right) \\ R_5 \left(\sigma \cdot t\right)^\dagger & -R_5 \left(W - m_0 + M_-(1)\right) \end{pmatrix}, \tag{F.99}$$

so, again using Eqn. (F.61),

$$\begin{aligned} \Big[(H(1) - \Delta_+(m) P_+)^{-1}\Big]_{++} &= \left(R_5 \left[W - m_0 + M_+(1) + (\sigma \cdot t) \Big(W - m_0 + M_-(1)\Big)^{-1} (\sigma \cdot t)^\dagger - \Delta_+(m)\right]\right)^{-1} \\ &= \frac{1}{H_+(1) - \Delta_+(m)}, \end{aligned} \tag{F.100}$$

and we have

$$\begin{aligned} \begin{pmatrix} \phi_2^\dagger & 0 \end{pmatrix} \left[1 + k\Omega_+^\dagger \Big(H(1) - \Delta_+(m) P_+\Big)^{-1} \Omega_+\right] \begin{pmatrix} \phi_2 \\ 0 \end{pmatrix} &= \phi_2^\dagger \left[1 + k\Omega_+^\dagger \Big(H_+(1) - \Delta_+(m)\Big)^{-1} \Omega_+\right] \phi_2 \\ &= \phi_2^\dagger H_2 \phi_2. \end{aligned} \tag{F.101}$$

Thus, the final form of the exact one-flavor action is

$$\begin{aligned} S_{\text{EOFA}} = &\begin{pmatrix} 0 & \phi_1^\dagger \end{pmatrix} \left[1 - k\Omega_-^\dagger \Big(H(m)\Big)^{-1} \Omega_-\right] \begin{pmatrix} 0 \\ \phi_1 \end{pmatrix} \\ &+ \begin{pmatrix} \phi_2^\dagger & 0 \end{pmatrix} \left[1 + k\Omega_+^\dagger \Big(H(1) - \Delta_+(m) P_+\Big)^{-1} \Omega_+\right] \begin{pmatrix} \phi_2 \\ 0 \end{pmatrix}. \end{aligned} \tag{F.102}$$

This reformulation will turn out to be useful for the implementation.

6.A.4 Generalization for Hasenbusch Mass Splitting

More generally, we want to use the Hasenbusch trick: we introduce a series of intermediate masses $\{m_i\}_{i=1}^N$ with $m < m_1 < \cdots < m_N < 1$ and break up the fermion determinant into a product

$$\frac{\det\left(\mathscr{D}_{\text{EOFA}}(m)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(1)\right)} = \frac{\det\left(\mathscr{D}_{\text{EOFA}}(m)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(m_1)\right)} \left(\prod_{i=1}^{N-1} \frac{\det\left(\mathscr{D}_{\text{EOFA}}(m_i)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(m_{i+1})\right)}\right) \frac{\det\left(\mathscr{D}_{\text{EOFA}}(m_N)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(1)\right)}. \tag{F.103}$$

In practice one can achieve a speed-up after some tuning since smaller mass ratios (generally) produce smaller forces.

Consider a single determinant ratio

$$\frac{\det\left(\mathscr{D}_{\text{EOFA}}(m_1)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(m_2)\right)} \tag{F.104}$$

with masses $m_1 < m_2$. The derivation in Section 6.A.3 is mostly unchanged: in particular, if we define slightly more general $\Delta_\pm$ matrices

$$\Delta_\pm(m_1, m_2) \equiv R_5 \left(M_\pm(m_2) - M_\pm(m_1)\right), \tag{F.105}$$

which reduce to Eqn. (F.92) when $m_2 = 1$, we have

$$\frac{\det\left(\mathscr{D}_{\text{EOFA}}(m_1)\right)}{\det\left(\mathscr{D}_{\text{EOFA}}(m_2)\right)} = \frac{\det\left(H_+(m_2) - \Delta_+(m_1, m_2)\right)}{\det\left(H_-(m_1) + \Delta_-(m_1, m_2)\right)} \cdot \frac{\det\left(H_-(m_1)\right)}{\det\left(H_+(m_2)\right)} \tag{F.106}$$

with $H_\pm(m)$ still defined by Eqn. (F.87). The decomposition of the $\Delta_\pm(m)$ matrices we wrote down in Eqn. (F.93) is easy to generalize since the dependence on m enters only into the normalization constant k:

$$\begin{aligned} \Delta_\pm(m_1, m_2) &= \left(\frac{2cm_2}{1 + m_2 - 2cm_2\lambda} - \frac{2cm_1}{1 + m_1 - 2cm_1\lambda}\right) \omega^{-1/2} \left(v_\pm \otimes v_\pm\right) \omega^{-1/2} \equiv k \Omega_\pm \Omega_\pm^\dagger \\ k &\equiv \frac{2c\left(m_2 - m_1\right)}{\left(1 + m_1 - 2cm_1\lambda\right)\left(1 + m_2 - 2cm_2\lambda\right)}, \quad \left(\Omega_\pm\right)_{ss'} \equiv \omega_s^{-1/2} \left(v_\pm\right)_s \delta_{s',0}. \end{aligned} \tag{F.107}$$

Thus, the appropriate generalization of the one-flavor action to Eqn. (F.104) is

$$\begin{aligned} S_{\text{EOFA}}(m_1, m_2) = &\begin{pmatrix} 0 & \phi_1^\dagger \end{pmatrix} \left[1 - k\Omega_-^\dagger \Big(H(m_1)\Big)^{-1} \Omega_-\right] \begin{pmatrix} 0 \\ \phi_1 \end{pmatrix} + \\ &\begin{pmatrix} \phi_2^\dagger & 0 \end{pmatrix} \left[1 + k\Omega_+^\dagger \Big(H(m_2) - \Delta_+(m_1, m_2) P_+\Big)^{-1} \Omega_+\right] \begin{pmatrix} \phi_2 \\ 0 \end{pmatrix}, \end{aligned} \tag{F.108}$$

and we can calculate the full determinant, Eqn. (F.103), with a sum

$$S_{\text{EOFA}} = S_{\text{EOFA}}(m, m_1) + \sum_{i=1}^{N-1} S_{\text{EOFA}}(m_i, m_{i+1}) + S_{\text{EOFA}}(m_N, 1). \tag{F.109}$$

With N intermediate Hasenbusch masses we need to simulate $N+1$ pairs of pseudofermion fields $\{(\phi_1^i, \phi_2^i)\}_{i=0}^N$ according to Eqn. (F.108).

6.B EOFA Operators for Shamir and Möbius DWF

In this appendix we list the operators which enter into $\mathscr{D}_{\text{EOFA}}$ (Eqn. (6.6)) and the EOFA action (Eqn. (6.10)). The more general case of DWF with Zolotarev-type domain wall fermions is constructed implicitly in Ref. [3]; we explicitly list these operators for the more restrictive cases of Shamir and Möbius DWF used in our simulations. We use Θ_s to denote the discrete Heaviside theta function

$$\Theta_s = \begin{cases} 0, & s < 0 \\ 1, & s \geq 0 \end{cases} \tag{F.110}$$

and assume even L_s. The operators $\Omega_\pm$ and $\Delta_\pm$ are related by the identity

$$\Delta_\pm = k\Omega_\pm\Omega_\pm^\dagger, \tag{F.111}$$

and the Möbius operators reduce to the corresponding Shamir operators in the limit $c = d = 1/2$. We note that the dense Möbius expressions listed here are not used inside the inverter; we instead invert the preconditioned system discussed in Appedix 6.C.2.

6.B.1 Shamir Kernel

$$k = m_2 - m_1 \tag{F.112}$$

$$[\Omega_+]_{ss'} = \delta_{s,L_s-1}\,\delta_{s',0} \tag{F.113}$$

$$[\Omega_-]_{ss'} = \delta_{s,0}\,\delta_{s',0} \tag{F.114}$$

$$[\Delta_+(m_1,m_2)]_{ss'} = (m_2-m_1)\,\delta_{s,L_s-1}\,\delta_{s',L_s-1} \tag{F.115}$$

$$[\Delta_-(m_1,m_2)]_{ss'} = (m_2-m_1)\,\delta_{s,0}\,\delta_{s',0} \tag{F.116}$$

$$[M_+(m)]_{ss'} = \delta_{ss'} - \delta_{s,s'+1} + m\delta_{s,L_s-1}\,\delta_{s',0} \tag{F.117}$$

$$[M_-(m)]_{ss'} = \delta_{ss'} - \delta_{s,s'-1} + m\delta_{s,0}\,\delta_{s',L_s-1} \tag{F.118}$$

$$[\widetilde{\mathscr{D}}(m)]_{ss'} = \delta_{ss'} \tag{F.119}$$

$$[\widetilde{\mathscr{D}}(m)^{-1}]_{ss'} = \delta_{ss'} \tag{F.120}$$

6.B.2 Möbius Kernel

$$k = \frac{2c\,(m_2-m_1)\,(c+d)^{2L_s}}{\left[(c+d)^{L_s}+m_1\,(c-d)^{Ls}\right]\left[(c+d)^{L_s}+m_2\,(c-d)^{Ls}\right]} \tag{F.121}$$

$$(\Omega_+)_{ss'} = (-1)^{s+1}\frac{(c-d)^{L_s-s-1}}{(c+d)^{L_s-s}}\,\delta_{s',0} \tag{F.122}$$

$$(\Omega_-)_{ss'} = (-1)^{s}\frac{(c-d)^{s}}{(c+d)^{s+1}}\,\delta_{s',0} \tag{F.123}$$

$$[\Delta_+(m_1,m_2)]_{ss'} = \frac{(-1)^{s+s'}2c\,(m_2-m_1)\,(c+d)^{s+s'}\,(c-d)^{2(L_s-1)-s-s'}}{\left[(c+d)^{L_s}+m_1\,(c-d)^{Ls}\right]\left[(c+d)^{L_s}+m_2\,(c-d)^{Ls}\right]} \tag{F.124}$$

$$[\Delta_-(m_1,m_2)]_{ss'} = \frac{(-1)^{s+s'}2c\,(m_2-m_1)\,(c+d)^{2(L_s-1)-s-s'}\,(c-d)^{s+s'}}{\left[(c+d)^{L_s}+m_1\,(c-d)^{Ls}\right]\left[(c+d)^{L_s}+m_2\,(c-d)^{Ls}\right]} \tag{F.125}$$

$$[M_+(m)]_{ss'} = \frac{(-1)^{s-s'}\,2c\,(c+d)^{Ls-s+s'-1}\,(c-d)^{s-s'-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\Theta_{s-s'-1} + \frac{(c+d)^{Ls-1}-m\,(c-d)^{Ls-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\delta_{ss'} + \frac{(-1)^{s-s'+1}\,2cm\,(c+d)^{s'-s-1}\,(c-d)^{Ls+s-s'-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\Theta_{s'-s-1} \quad \text{(F.126)}$$

$$[M_-(m)]_{ss'} = \frac{(-1)^{s'-s+1}\,2cm\,(c+d)^{s-s'-1}\,(c-d)^{Ls-s+s'-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\Theta_{s-s'-1} + \frac{(c+d)^{Ls-1}-m\,(c-d)^{Ls-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\delta_{ss'} + \frac{(-1)^{s-s'}\,2c\,(c+d)^{Ls+s-s'-1}\,(c-d)^{s'-s-1}}{(c+d)^{Ls}+m\,(c-d)^{Ls}}\,\Theta_{s'-s-1} \quad \text{(F.127)}$$

$$[\widetilde{\mathscr{D}}(m)]_{ss'} = (c+d)\,\delta_{ss'} + (c-d)\,P_+\delta_{s,s'+1} + (c-d)\,P_-\delta_{s,s'-1} - m\,(c-d)\,P_+\delta_{s,0}\,\delta_{s',L_s-1} - m\,(c-d)\,P_-\delta_{s,L_s-1}\,\delta_{s',0} \quad \text{(F.128)}$$

$$[\widetilde{\mathscr{D}}(m)^{-1}]_{ss'} = \left[\frac{m\,(-1)^{s-s'+1}\,(c+d)^{s'-s-1}\,(c-d)^{L_s+s-s'}}{(c+d)^{L_s}+m\,(c-d)^{L_s}} + \frac{(-1)^{s-s'}\,(c-d)^{s-s'}}{(c+d)^{s-s'+1}}\,\Theta_{s-s'}\right]P_+ + \left[\frac{m\,(-1)^{s'-s+1}\,(c+d)^{s-s'-1}\,(c-d)^{L_s+s'-s}}{(c+d)^{L_s}+m\,(c-d)^{L_s}} + \frac{(-1)^{s'-s}\,(c-d)^{s'-s}}{(c+d)^{s'-s+1}}\,\Theta_{s'-s}\right]P_- \quad \text{(F.129)}$$

6.C Four-Dimensional Even-Odd Preconditioning

The inversions required to compute the exact one flavor Hamiltonian can be accelerated using a standard checkerboarding technique: we label lattice sites as "even" if $(x+y+z+t) \equiv 0 \pmod 2$ or "odd" if $(x+y+z+t) \equiv 1 \pmod 2$. This naturally induces a block structure in the Dirac

operator $\mathscr{D}$, which can be LDU decomposed as

$$\underbrace{\begin{pmatrix} \mathscr{D}_{ee} & \mathscr{D}_{eo} \\ \mathscr{D}_{oe} & \mathscr{D}_{oo} \end{pmatrix}}_{\mathscr{D}} = \underbrace{\begin{pmatrix} 1 & 0 \\ \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1} & 1 \end{pmatrix}}_{L} \underbrace{\begin{pmatrix} \mathscr{D}_{ee} & 0 \\ 0 & \mathscr{D}_{oo} - \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1}\mathscr{D}_{eo} \end{pmatrix}}_{D} \underbrace{\begin{pmatrix} 1 & \mathscr{D}_{ee}^{-1}\mathscr{D}_{eo} \\ 0 & 1 \end{pmatrix}}_{U}. \tag{F.130}$$

Left-multiplying the linear system $\mathscr{D}\psi = \phi$ by

$$L^{-1} = \begin{pmatrix} 1 & 0 \\ -\mathscr{D}_{oe}\mathscr{D}_{ee}^{-1} & 1 \end{pmatrix} \tag{F.131}$$

results in the equivalent system

$$\begin{pmatrix} \mathscr{D}_{ee}\psi_e + \mathscr{D}_{eo}\psi_o \\ \left(\mathscr{D}_{oo} - \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1}\mathscr{D}_{eo}\right)\psi_o \end{pmatrix} = \begin{pmatrix} \phi_e \\ \phi_o - \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1}\phi_e \end{pmatrix}, \tag{F.132}$$

leading to the following trick: assuming $\mathscr{D}_{ee}^{-1}$ is available in an explicit form, it suffices to invert

$$\left(\mathscr{D}_{oo} - \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1}\mathscr{D}_{eo}\right)\psi_o = \widetilde{\phi}_o, \tag{F.133}$$

with $\widetilde{\phi}_o \equiv \phi_o - \mathscr{D}_{oe}\mathscr{D}_{ee}^{-1}\phi_e$. This system only involves the odd sublattice, and is thus substantially cheaper to invert than $\mathscr{D}$ using an iterative algorithm like CG. The solution on the even sublattice can then be reconstructed for a trivial additional cost as

$$\psi_e = \mathscr{D}_{ee}^{-1}\left(\phi_e - \mathscr{D}_{eo}\psi_o\right). \tag{F.134}$$

This technique is already well understood in the context of RHMC with Shamir or Möbius DWF; in this appendix we describe how to generalize the method to the exact one flavor algorithm.

In the context of EOFA, the generic linear system one needs to invert takes the form

$$\Big(H(m_1) + \beta\Delta_{\pm}(m_2, m_3)P_{\pm}\Big)\psi = \phi. \tag{F.135}$$

We choose to multiply by an overall factor of $\gamma_5 R_5$, rewriting the system as

$$\Big(\mathscr{D}_{\text{EOFA}}(m_1) + \beta\gamma_5 R_5\Delta_{\pm}(m_2, m_3)P_{\pm}\Big)\psi = \gamma_5 R_5\phi, \tag{F.136}$$

for the following reasons: first, we wish to re-use the existing high-performance implementation of the Wilson $\not{D}$ kernel in the BAGEL library without modification, and second, overall factors of $\gamma_5 R_5$

will cancel inside the inverter since we use CG applied to the normal equations and $(\gamma_5 R_5)^\dagger(\gamma_5 R_5) = 1$. Since $(\mathscr{D}_{\text{EOFA}})_{eo} = (\mathscr{D}_{\text{DWF}})_{eo}$ and $\Delta_\pm \propto \delta_{xx'}$ in the 4D bulk, only the operators coupling sites of the same parity need to be modified to implement even-odd preconditioned EOFA. We take somewhat different approaches for the Shamir and Möbius cases.

6.C.1 Shamir Kernel

Recall that for the Shamir kernel $\mathscr{D}_{\text{DWF}} = \mathscr{D}_{\text{EOFA}}$, so the extension of an inverter for the even-odd preconditioned $\mathscr{D}_{\text{DWF}}$ operator to instead solve Eqn. (F.136) is straightforward. With $\mathscr{D} = \mathscr{D}_{\text{DWF}}$, the same parity fermion matrix has the tridiagonal block structure

$$(\mathscr{D}_{\text{DWF}})_{ee} = (\mathscr{D}_{\text{DWF}})_{oo} = \delta_{xx'}\Big\{ (5 - M_5)\,\delta_{ss'} - P_+\delta_{s,s'+1} - P_-\delta_{s,s'-1} + m_1 P_+\delta_{s,0}\delta_{s',L_s-1} + m_1 P_-\delta_{s,L_s-1}\delta_{s',0}\Big\}. \quad \text{(F.137)}$$

One can check by explicit calculation that the shift operators have the form

$$\begin{cases} \beta\gamma_5 R_5 \Delta_+(m_2, m_3)P_+ = \beta\,(m_3 - m_2)\,P_+\delta_{xx'}\delta_{s,0}\delta_{s',L_s-1} \\ \beta\gamma_5 R_5 \Delta_-(m_2, m_3)P_- = -\beta\,(m_3 - m_2)\,P_-\delta_{xx'}\delta_{s,L_s-1}\delta_{s',0} \end{cases}, \quad \text{(F.138)}$$

so one can consider the operator appearing in Eqn. (F.136) as a slight generalization of Eqn. (F.137) to

$$\mathscr{D}_{ee} = \mathscr{D}_{oo} = \delta_{xx'}\Big\{ (5 - M_5)\,\delta_{ss'} - P_+\delta_{s,s'+1} - P_-\delta_{s,s'-1} + d_+ P_+\delta_{s,0}\delta_{s',L_s-1} + d_- P_-\delta_{s,L_s-1}\delta_{s',0}\Big\}, \quad \text{(F.139)}$$

with

$$d_- = m_1 - \beta\,(m_3 - m_2)\,\delta_{i,-} \quad \text{(F.140)}$$

and

$$d_+ = m_1 + \beta\,(m_3 - m_2)\,\delta_{i,+}, \quad \text{(F.141)}$$

where the index i denotes the chirality of the shift operator. $\mathscr{D}_{ee}^{-1}$ can be efficiently applied using the LDU decomposition of $\mathscr{D}_{ee}$, again as a slight generalization of the standard Shamir DWF case.

6.C.2 Möbius Kernel and Cayley-Form Preconditioning

Using Eqn. (6.7) we can write $\mathscr{D}_{\text{EOFA}}$ in the form

$$(\mathscr{D}_{\text{EOFA}})_{xx',ss'} = (D_W)_{xx'}\,\delta_{ss'} + \delta_{xx'}\left(\mathscr{D}^{\perp}\right)_{ss'}. \tag{F.142}$$

The action of the operator appearing in Eqn. (F.136) on lattice sites of the same parity, then, is given by

$$\mathscr{D}_{ee} = \mathscr{D}_{oo} = \delta_{xx'}\Big\{ (4 - M_5)\,\delta_{ss'} + (M_+(m_1))_{ss'}\,P_+ + (M_-(m_1))_{ss'}\,P_- \\ + \beta\gamma_5 R_5\,(\Delta_\pm(m_2,m_3))_{ss'}\,P_\pm\Big\}, \tag{F.143}$$

with M_+, M_-, and $\Delta_\pm$ as defined in equations (F.124)-(F.127). The matrix elements of $\mathscr{D}_{ee}^{-1} = \mathscr{D}_{oo}^{-1}$ can be found by explicit numerical inversion as part of the setup cost; this is a trivial overhead since it suffices to invert only the ss' subblock of $\mathscr{D}_{ee}$. In this form, the exact factorization of the fermion determinant in Eqn. (6.2) comes at the cost of dense $L_s \times L_s$ matrix operations. While TWQCD has shown that EOFA can still be faster than RHMC after retuning the integrator layout [4], we argue that it is possible to do significantly better by introducing an additional preconditioning step.

We note that the system defined by Eqn. (F.136) can be more efficiently inverted for the case of Möbius DWF by using the operator $\widetilde{\mathscr{D}}^{-1}$ as a right preconditioner, resulting in an equivalent system in terms of $\mathscr{D}_{\text{DWF}}$. For the special case $\beta = 0$ this is straightforward: observing that the relationship between $\mathscr{D}_{\text{EOFA}}$ and $\mathscr{D}_{\text{DWF}}$ (Eqn. (6.6)) can be used to manipulate

$$\mathscr{D}_{\text{EOFA}}\psi = \mathscr{D}_{\text{EOFA}} \cdot \widetilde{\mathscr{D}} \cdot \underbrace{\widetilde{\mathscr{D}}^{-1}\psi}_{\equiv\psi'} = \mathscr{D}_{\text{DWF}}\psi', \tag{F.144}$$

it suffices to solve $\mathscr{D}_{\text{DWF}}\psi' = \gamma_5 R_5 \phi$, from which $\psi = \widetilde{\mathscr{D}}\psi'$ can be recovered at the cost of a single additional matrix multiplication. While we observe that $\mathscr{D}_{\text{DWF}}^\dagger \mathscr{D}_{\text{DWF}}$ has a slightly larger condition number than $\mathscr{D}_{\text{EOFA}}^\dagger \mathscr{D}_{\text{EOFA}}$, leading to a modest increase in the total number of CG iterations required to invert the system, $\mathscr{D}_{\text{DWF}}$ also has a tridiagonal stencil in the fifth dimension, and can thus be applied in $\mathcal{O}(L_s)$ operations — unlike the $\mathcal{O}(L_s^2)$ operations required for the dense $\mathscr{D}_{\text{EOFA}}$ — leading to a substantial reduction in wall clock time for the inversion.

The $\beta \neq 0$ case is more involved, but can be treated in a similar manner. Right preconditioning Eqn. (F.136) with $\widetilde{\mathscr{D}}^{-1}$ leads to

$$\left(\mathscr{D}_{\mathrm{DWF}}(m_1) \pm \beta R_5 \Delta_\pm(m_2, m_3)\widetilde{\mathscr{D}}P_\pm\right)\psi' = \gamma_5 R_5 \phi, \tag{F.145}$$

where we have used $\gamma_5 P_\pm = \pm P_\pm$. We define a new, preconditioned, shift operator $\widetilde{\Delta}_\pm$ by

$$\widetilde{\Delta}_\pm(m_1, m_2) \equiv R_5 \Delta_\pm(m_1, m_2)\widetilde{\mathscr{D}}P_\pm, \tag{F.146}$$

and note that since $(\widetilde{\Delta})_{eo} = (\widetilde{\Delta})_{oe} = 0$, Eqn. (F.145) can be inverted efficiently even with $\beta \neq 0$ provided we can apply the operator $(\mathscr{D}_{\mathrm{DWF}})_{ee} \pm \beta\widetilde{\Delta}_\pm$ and its inverse in $\mathcal{O}(L_s)$ operations. This turns out to be possible after observing that $\widetilde{\Delta}_\pm$ is rank-one, *i.e.* it can be written as a vector outer product

$$\widetilde{\Delta}_\pm = u_\pm \otimes v_\pm. \tag{F.147}$$

To see this, we start by decomposing $\widetilde{\mathscr{D}}$ into its chiral components — $\widetilde{\mathscr{D}} = \widetilde{\mathscr{D}}_+ P_+ + \widetilde{\mathscr{D}}_- P_-$ — in terms of which we can also decompose

$$\widetilde{\Delta}_\pm = R_5 \Delta_\pm \widetilde{\mathscr{D}}_\pm P_\pm. \tag{F.148}$$

The 5D structure of these operators can be worked out by direct calculation, leading to Eqn. (F.147), with

$$\begin{cases} (u_+)_s = (-1)^s \dfrac{(c-d)^s}{(c+d)^{L_s+s+1}} \left((c+d)^{L_s} + m_1 (c-d)^{L_s}\right) \\ (v_+)_s = k\, \delta_{s,L_s-1} \\ (u_-)_s = (-1)^{s+1} \dfrac{(c-d)^{L_s-1-s}}{(c+d)^{2L_s-s}} \left((c+d)^{L_s} + m_1 (c-d)^{L_s}\right) \\ (v_-)_s = k\, \delta_{s,0} \end{cases}. \tag{F.149}$$

Matrix-vector products involving the preconditioned shift operator and a pseudofermion field can be computed from this decomposition as

$$\begin{cases} \left(\widetilde{\Delta}_+ \psi\right)_s = k\, (u_+)_s\, P_+ \psi_{L_s-1} \\ \left(\widetilde{\Delta}_+^\dagger \psi\right)_s = k\, \delta_{s,L_s-1}\, P_+ \left[\displaystyle\sum_{s'=0}^{L_s-1} (u_+)_{s'}\, \psi_{s'}\right] \end{cases} \tag{F.150}$$

and

$$\begin{cases} \left(\widetilde{\Delta}_{-}\psi\right)_s = k\,(u_-)_s\,P_-\psi_0 \\ \left(\widetilde{\Delta}_{-}^{\dagger}\psi\right)_s = k\,\delta_{s,0}\,P_-\left[\sum_{s'=0}^{L_s-1}(u_-)_{s'}\,\psi_{s'}\right] \end{cases}. \tag{F.151}$$

The inverses can be applied using the Sherman-Morrison formula:

$$\left((\mathscr{D}_{\mathrm{DWF}})_{ee} \pm \beta\,(u_\pm \otimes v_\pm)\right)^{-1} = (\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1} \mp \beta \frac{(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\,(u_\pm \otimes v_\pm)\,(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}}{1 \pm \beta\langle v_\pm, (\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1} u_\pm\rangle}. \tag{F.152}$$

In terms of

$$x_\pm \equiv (\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1} u_\pm, \tag{F.153}$$

which can be constructed numerically using the tridiagonal matrix algorithm [39], the necessary factors can be written as

$$\begin{cases} 1 + \beta\langle v_+, (\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1} u_+\rangle = 1 + \beta k\,(x_+)_{L_s-1} \\ \left(\left[(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\,(u_\pm \otimes v_\pm)\,(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\right]\psi\right)_s = \dfrac{k\,(x_+)_s}{(c+d)^{L_s} + m_1\,(c-d)^{L_s}}\,P_+\left[\sum_{s'=0}^{L_s-1}(c+d)^{s'}\,(c-d)^{L_s-1-s'}\,\psi_{s'}\right] \\ \left(\left[(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\,(u_+ \otimes v_+)\,(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\right]^{\dagger}\psi\right)_s = \dfrac{k\,(c+d)^{s}\,(c-d)^{L_s-1-s}}{(c+d)^{L_s} + m_1\,(c-d)^{L_s}}\,P_+\left[\sum_{s'=0}^{L_s-1}(x_+)_{s'}\,\psi_{s'}\right] \end{cases} \tag{F.154}$$

and

$$\begin{cases} 1 - \beta\langle v_-, (\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1} u_-\rangle = 1 - \beta k\,(x_-)_0 \\ \left(\left[(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\,(u_- \otimes v_-)\,(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\right]\psi\right)_s = \dfrac{k\,(x_-)_s}{(c+d)^{L_s} + m_1\,(c-d)^{L_s}}\,P_-\left[\sum_{s'=0}^{L_s-1}(c+d)^{L_s-1-s'}\,(c-d)^{s'}\,\psi_{s'}\right], \\ \left(\left[(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\,(u_- \otimes v_-)\,(\mathscr{D}_{\mathrm{DWF}})_{ee}^{-1}\right]^{\dagger}\psi\right)_s = \dfrac{k\,(c+d)^{L_s-1-s}\,(c-d)^{s}}{(c+d)^{L_s} + m_1\,(c-d)^{L_s}}\,P_-\left[\sum_{s'=0}^{L_s-1}(x_-)_{s'}\,\psi_{s'}\right] \end{cases} \tag{F.155}$$

which allow Eqn. (F.152) to be applied to a pseudofermion vector in $\mathcal{O}(L_s)$ operations.

In Figure 6.9 we benchmark representative even-odd preconditioned inversions of Eqn. (F.136) on the 24ID ensemble, with and without additional preconditioning by $\widetilde{\mathscr{D}}^{-1}$, at the physical strange quark mass. In addition to observing a substantial improvement in terms of wall clock time for the inversion, we note that this preconditioning scheme also has the advantage that it requires little new code — assuming an existing high-performance implementation of $\mathscr{D}_{\mathrm{DWF}}$ — since $\mathscr{D}_{\mathrm{EOFA}}$ is never applied directly in the preconditioned formalism.

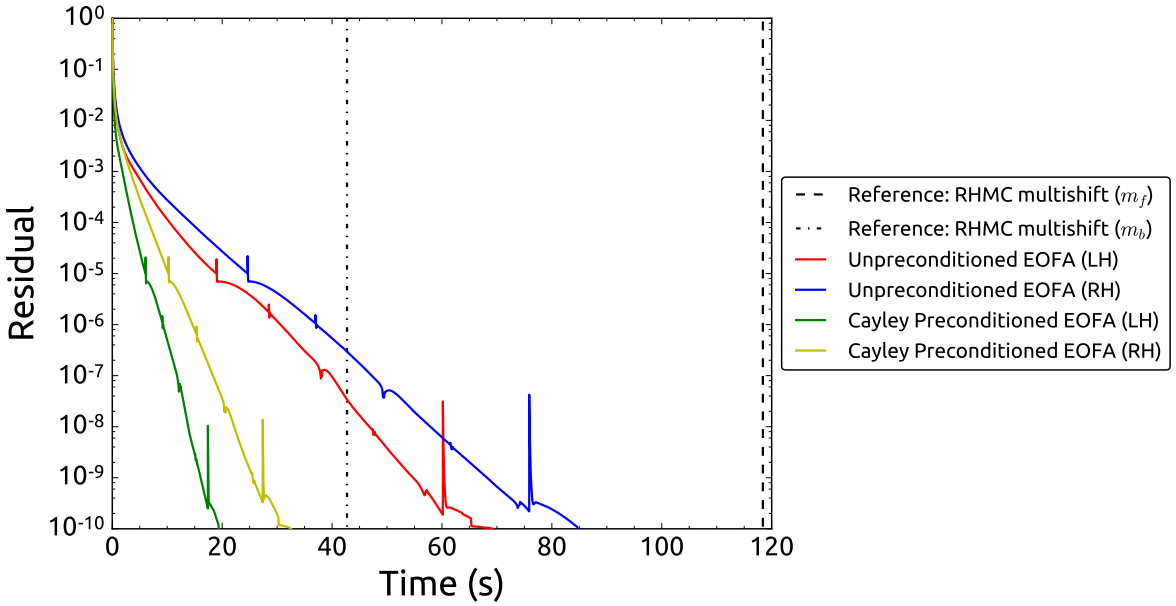

Figure 6.9: Comparison of wall clock inversion times for the two solves required to evaluate the EOFA Hamiltonian or pseudofermion force with and without Cayley-form preconditioning for the strange quark determinant on the 24ID ensemble. The dashed vertical lines show the corresponding total cost of the multishift inversions of $\mathscr{D}^{\dagger}_{\mathrm{DWF}}\mathscr{D}_{\mathrm{DWF}}$ needed to evaluate the RHMC Hamiltonian or pseudofermion force on the same ensemble.

6.D Additional Plots for Small Volume Reproduction Tests

6.D.1 Evolution of the Plaquette, Quark Condensates, and Topological Charge

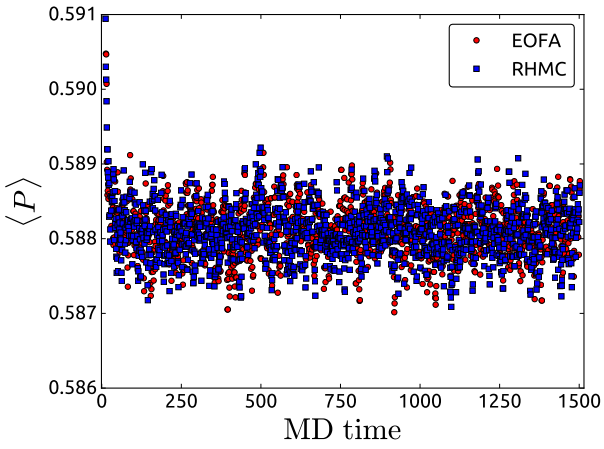

(a) Plaquette

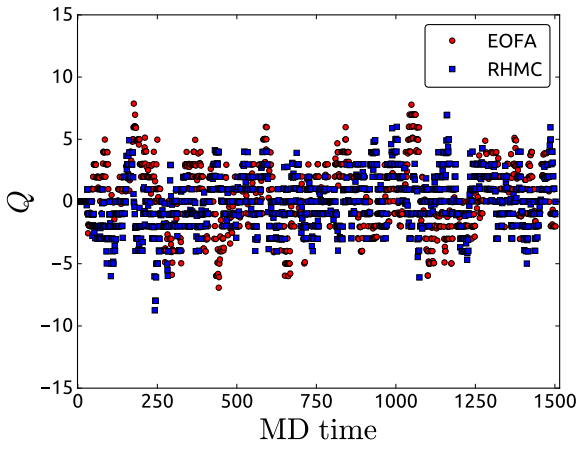

(b) Topological Charge

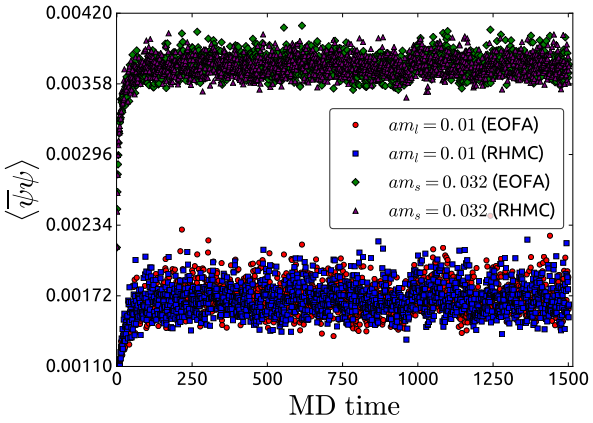

(c) Chiral Condensates

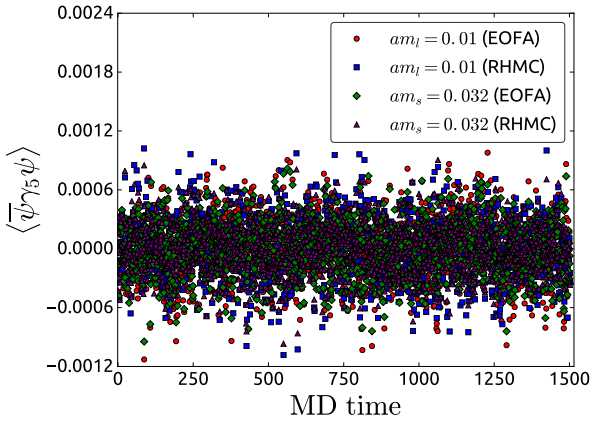

(d) Pseudoscalar Condensates

Figure 6.10: Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16I ensembles.

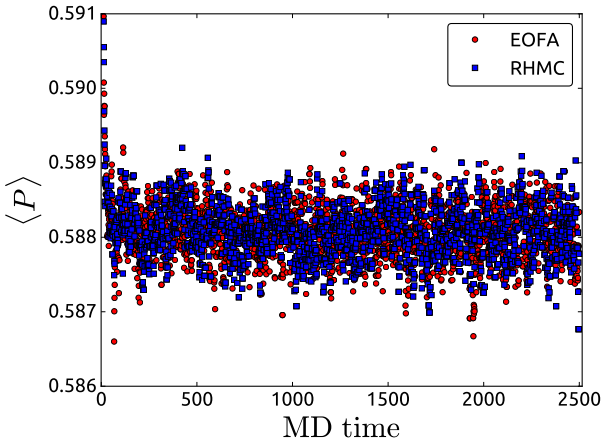

(a) Plaquette

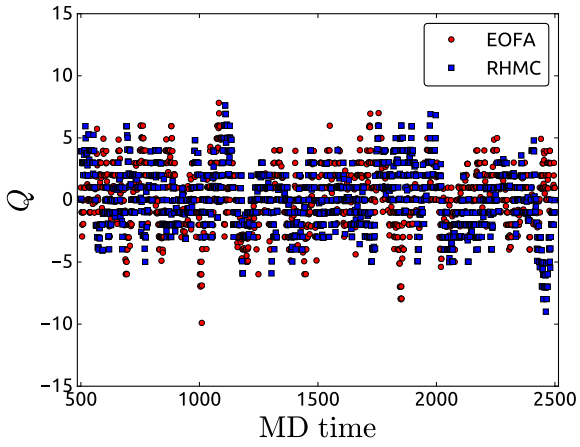

(b) Topological Charge

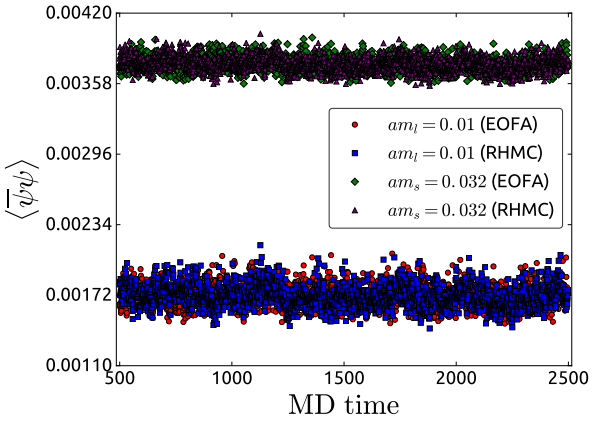

(c) Chiral Condensates

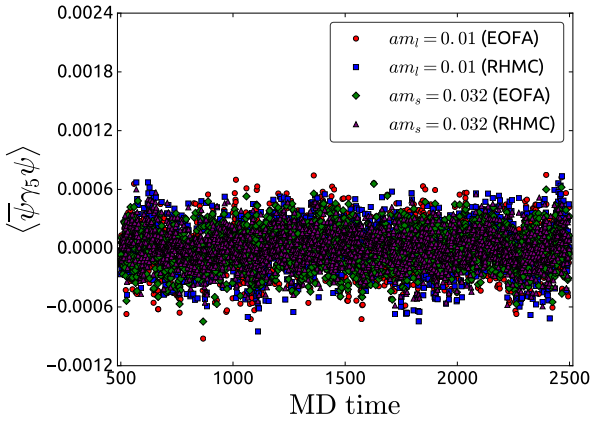

(d) Pseudoscalar Condensates

Figure 6.11: Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16I-G ensembles.

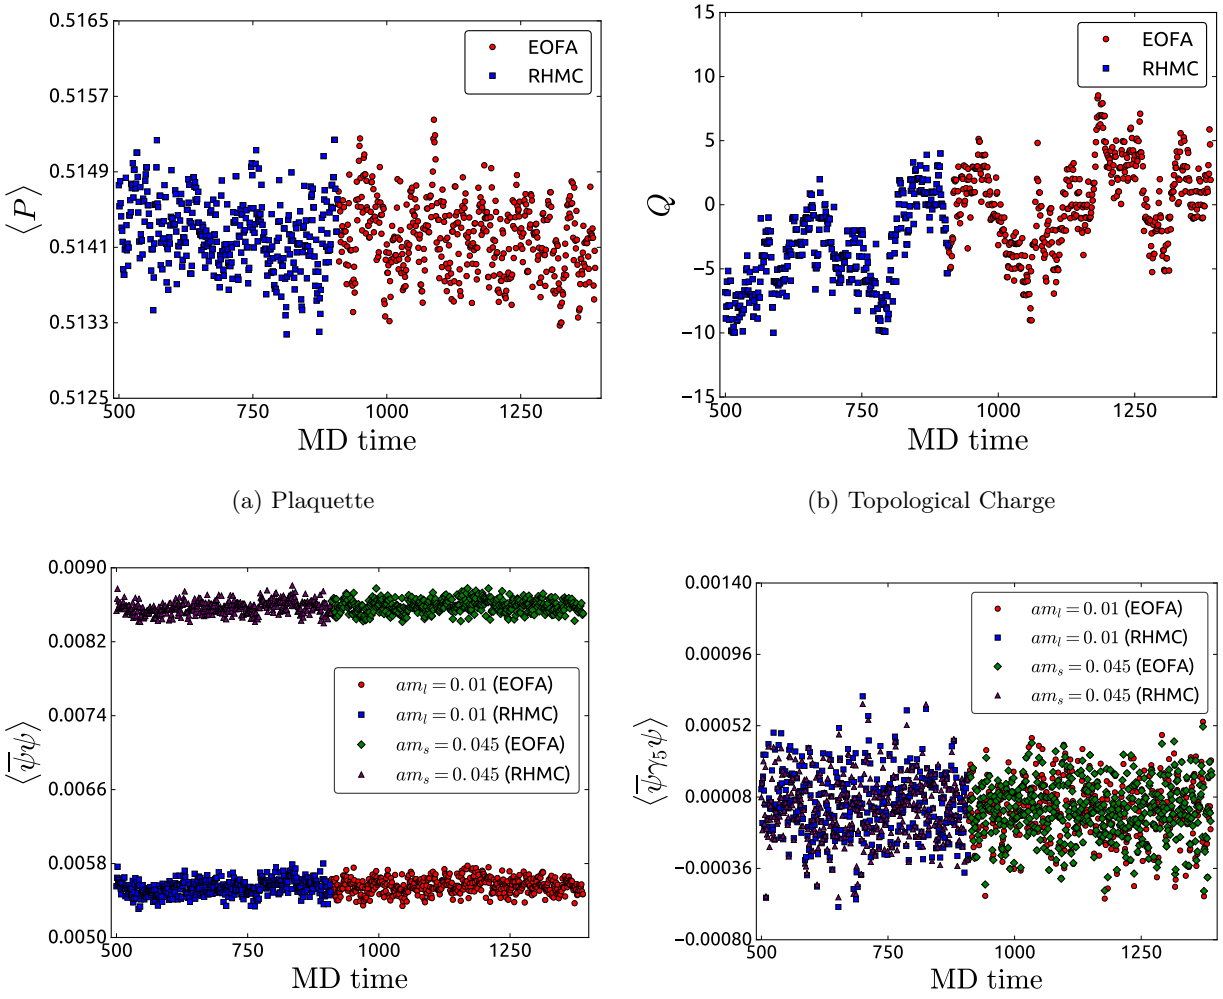

(a) Plaquette

(b) Topological Charge

(c) Chiral Condensates

(d) Pseudoscalar Condensates

Figure 6.12: Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16ID-G ensembles.

6.D.2 Effective Mass Plots

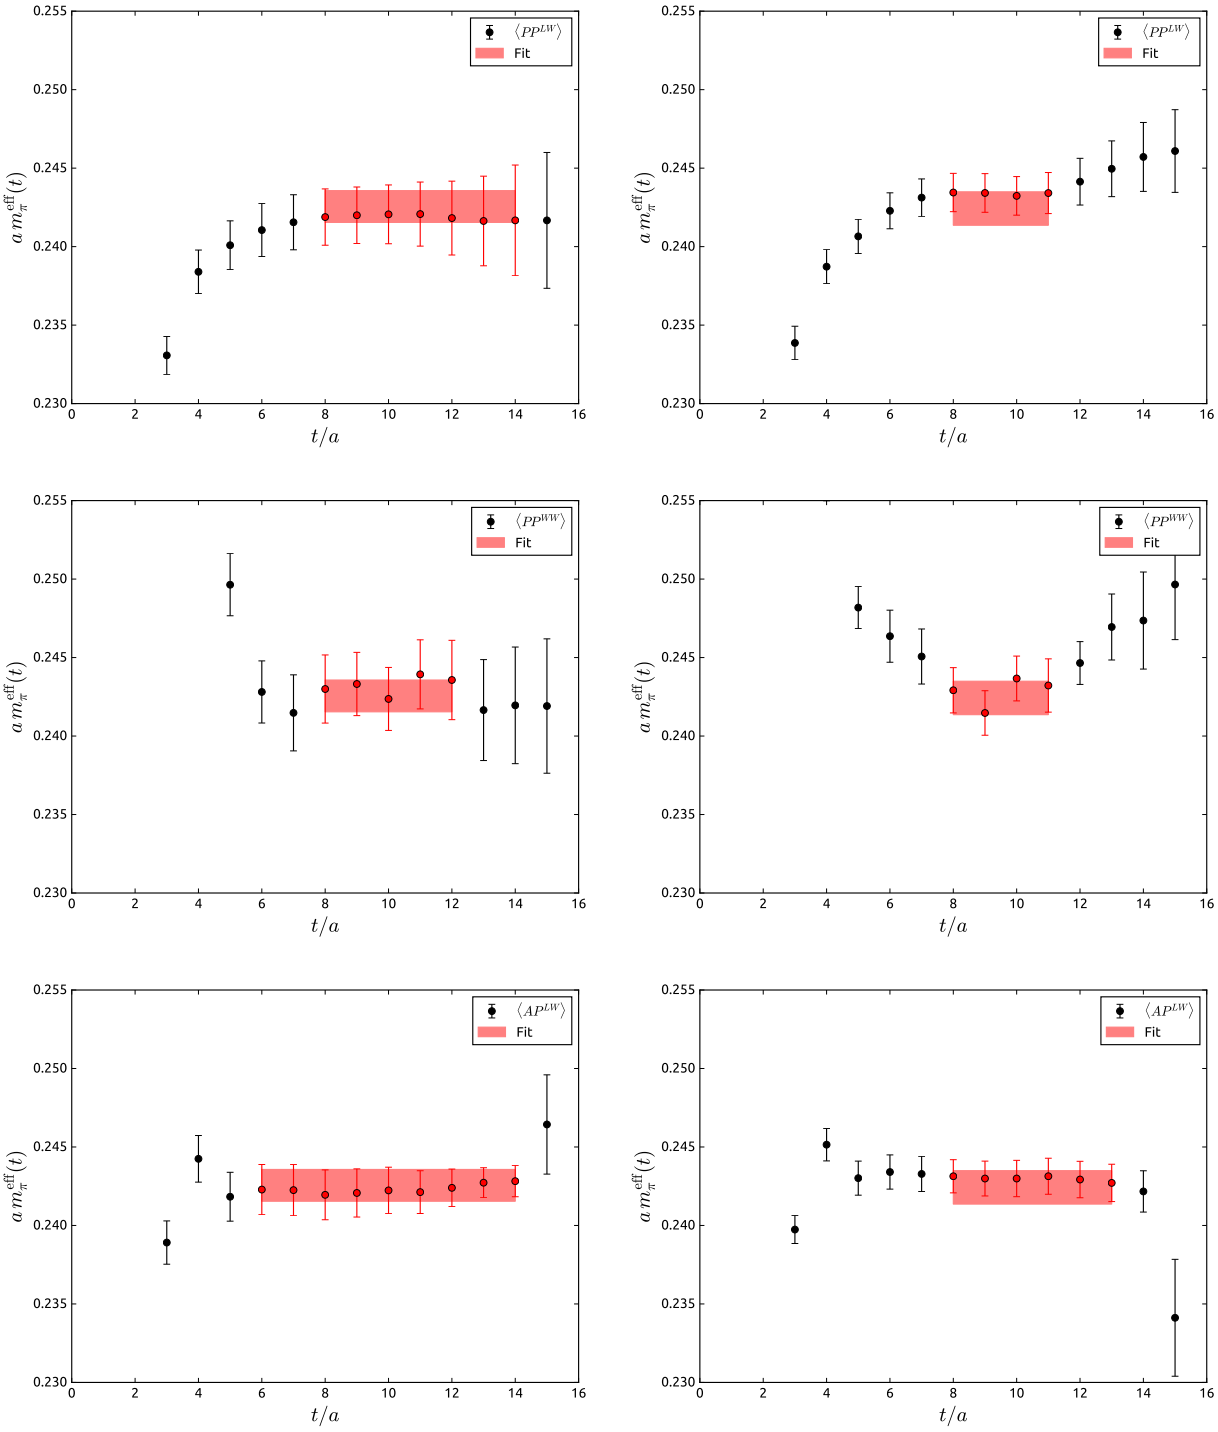

Figure 6.13: Effective pion mass from a simultaneous fit to the $\langle PP^{LW} \rangle$ (top), $\langle PP^{WW} \rangle$ (middle), and $\langle AP^{LW} \rangle$ (bottom) correlation functions, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

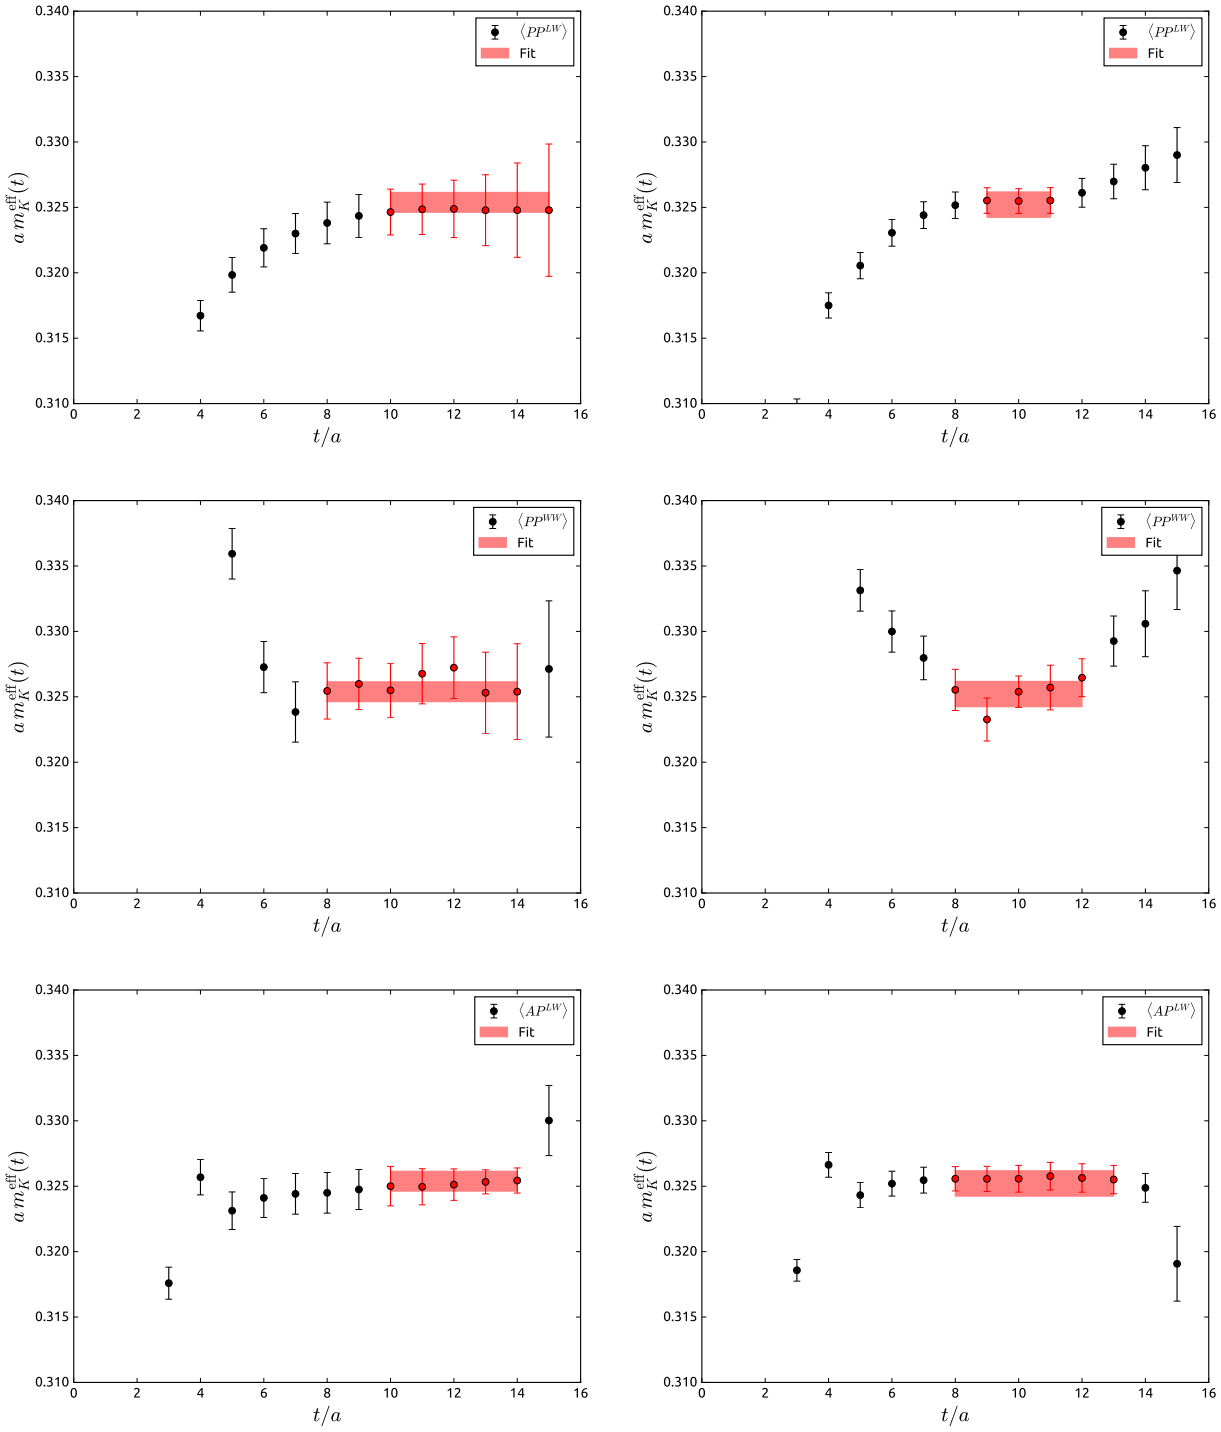

Figure 6.14: Effective kaon mass from a simultaneous fit to the $\langle PP^{LW} \rangle$ (top), $\langle PP^{WW} \rangle$ (middle), and $\langle AP^{LW} \rangle$ (bottom) correlation functions, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

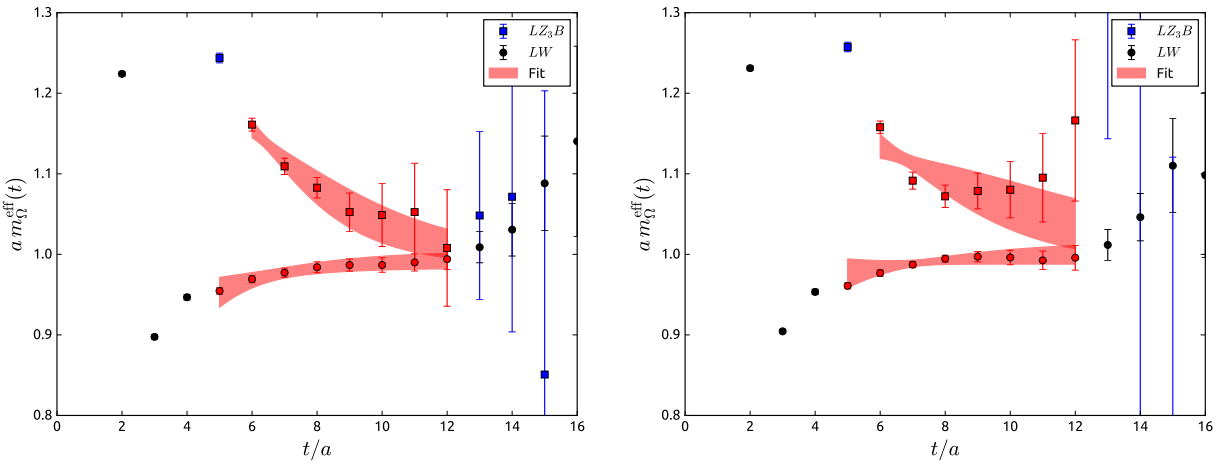

Figure 6.15: Effective Ω baryon mass from a simultaneous two-state fit to wall and Z_3 noise sources, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

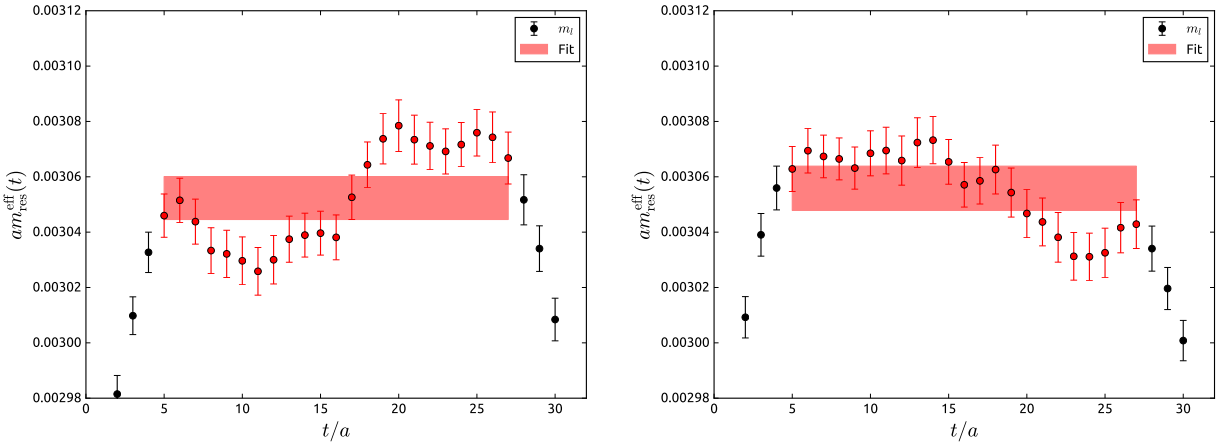

Figure 6.16: Effective $am'_{\rm res}(m_l)$, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

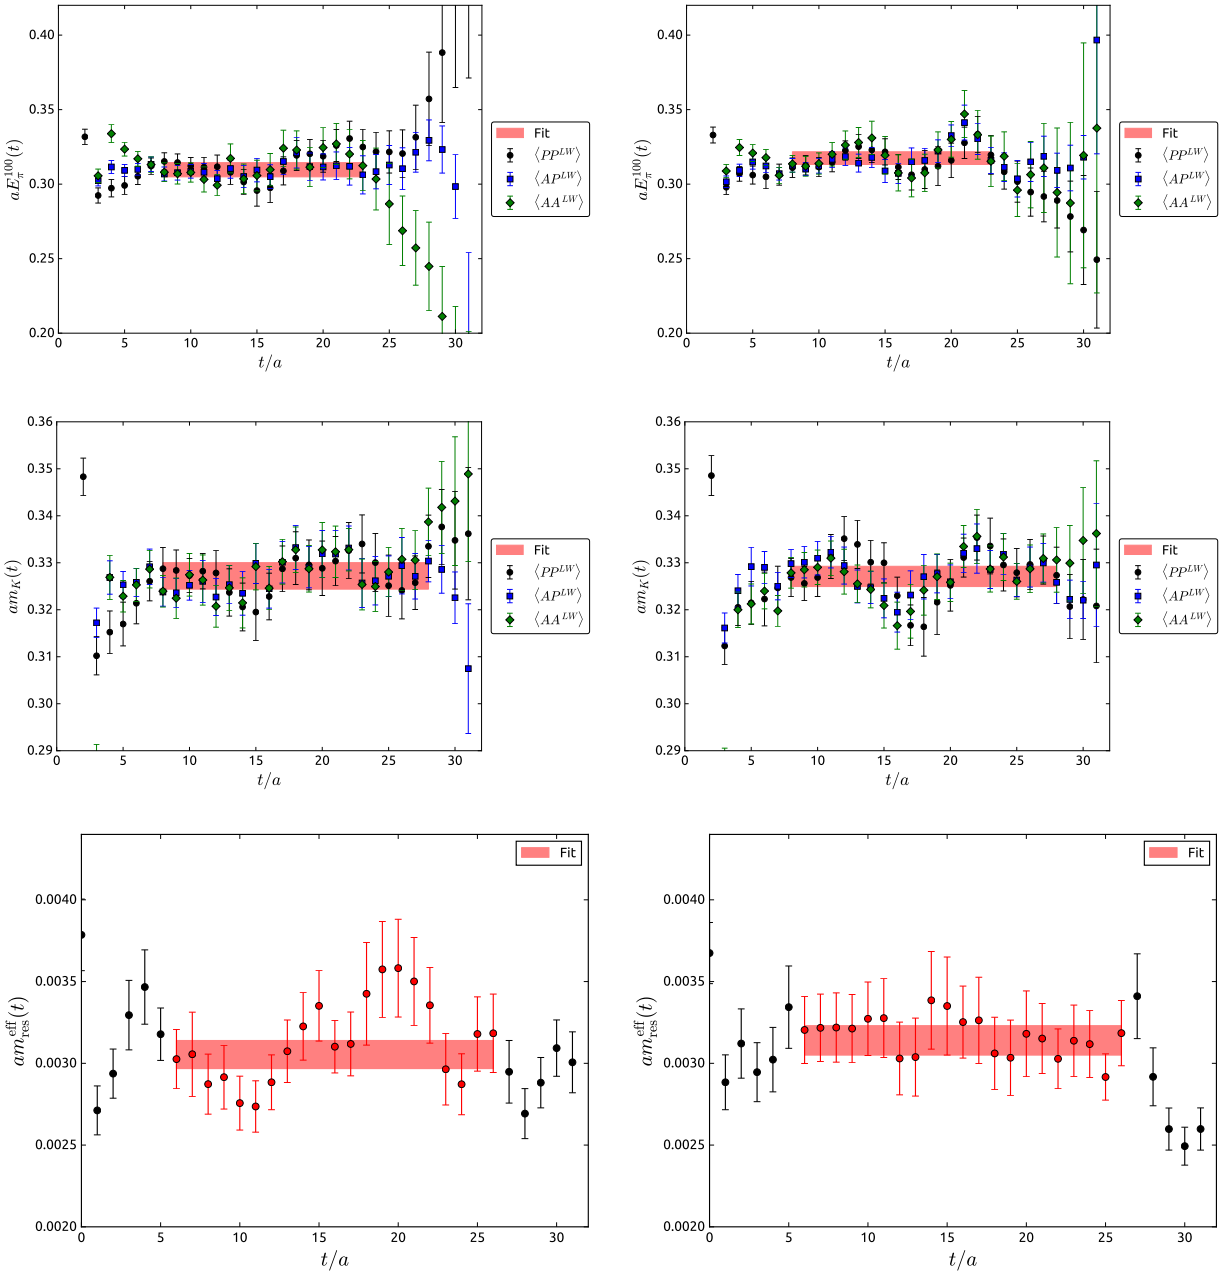

Figure 6.17: Effective ground state pion energy (top), kaon mass (middle), and $am'_{\rm res}(m_l)$ evaluated at the bare light quark mass, as measured on the EOFA (left) and RHMC (right) 16I-G ensembles.

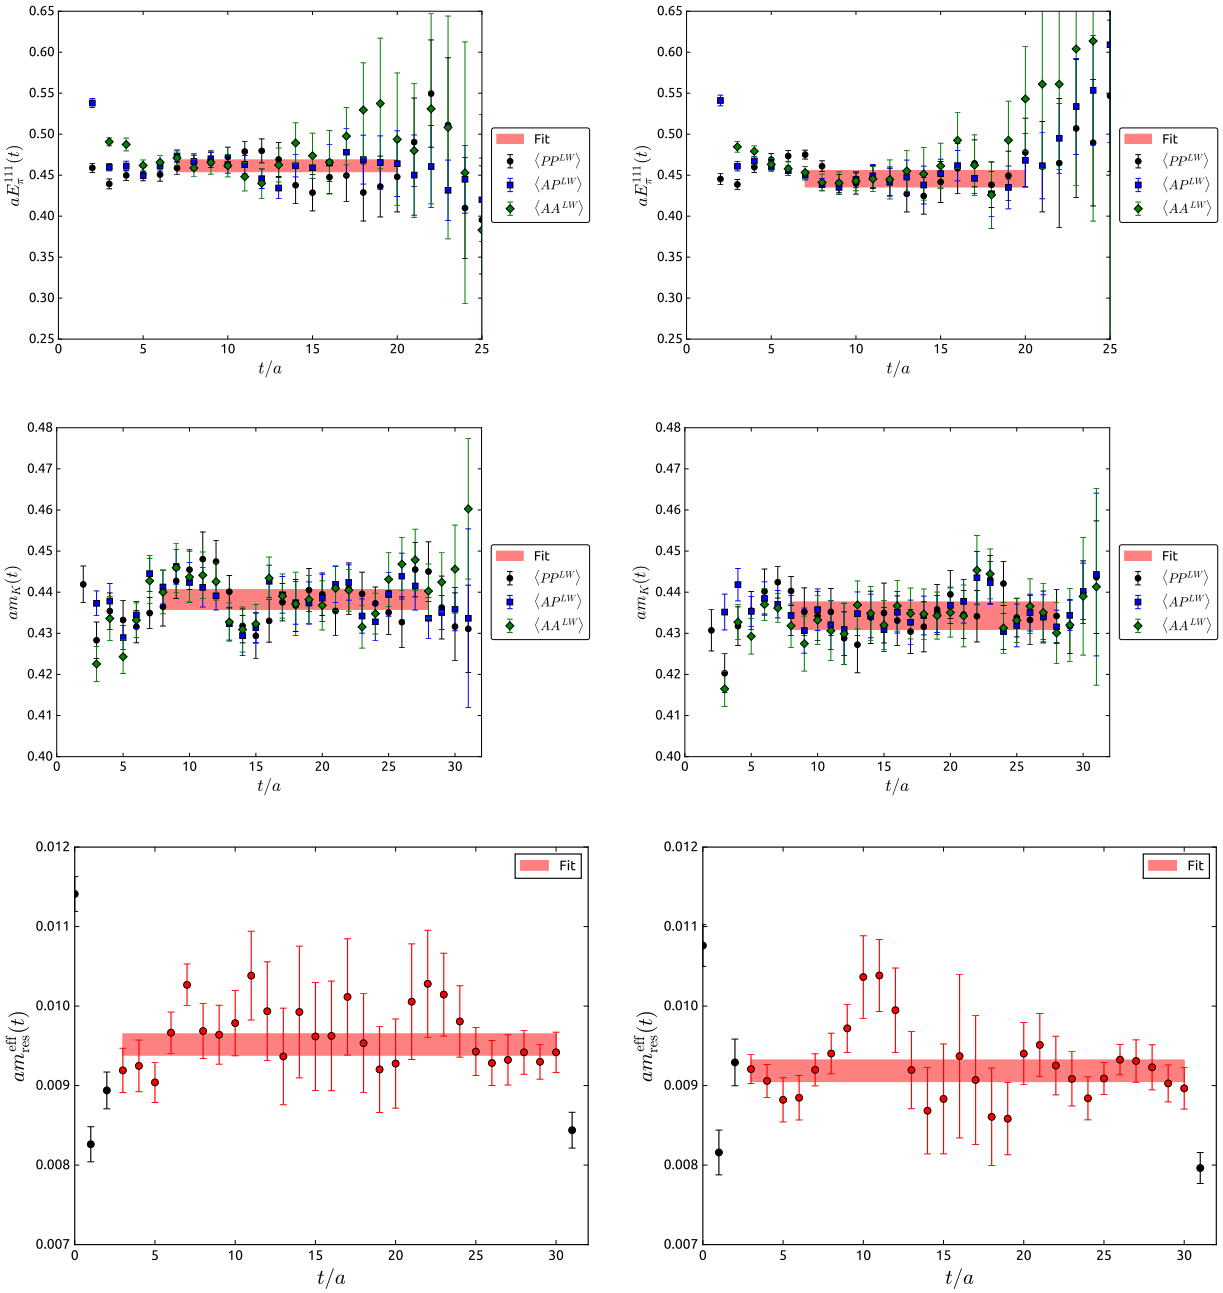

Figure 6.18: Effective ground state pion energy (top), kaon mass (middle), and $am_{\rm res}$ evaluated at the bare light quark mass, as measured on the EOFA (left) and RHMC (right) 16ID-G ensembles.