

Just a Taste

Lectures on Flavor Physics

Lecturer: Yuval Grossman^(a)
L^AT_EX Notes: Flip Tanedo^(b)

*Institute for High Energy Phenomenology,
Newman Laboratory of Elementary Particle Physics,
Cornell University, Ithaca, NY 14853, USA*

E-mail: (a) yg73@cornell.edu, (b) pt267@cornell.edu

This version: September 13, 2012

Abstract

This is a set of L^AT_EX'ed notes from Cornell University's Physics 7661 (special topics in theoretical high energy physics) course by Yuval Grossman in Fall 2010. The lectures as given were flawless, all errors contained herein reflect solely the typist's editorial and/or intellectual deficiencies!

Contents

1	Introduction	1
2	Model building	2
2.1	Example: the Standard Model	3
2.2	Global, accidental, and approximate symmetries	4
2.3	Learning how to count [parameters]	5
2.4	Counting parameters in low-energy QCD	6
2.5	Counting parameters in the Standard Model	8
3	A review of hadrons	10
3.1	What we mean by 'stable'	10
3.2	Hadron quantum numbers	12
3.3	Binding energy	14
3.4	Light quarks, heavy quarks, and the heaviest quark	15
3.5	Masses and mixing in mesons	16
3.6	The pseudoscalar mesons	19
3.7	The vector mesons	20
3.8	Why are the pseudoscalar and vector octets so different?	25
3.9	Hadron names	26

4	The flavor structure of the Standard Model	29
4.1	The CKM matrix	29
4.2	Parameterizations of the CKM matrix	32
4.3	CP violation	33
4.4	The Jarlskog Invariant	35
4.5	Unitarity triangles and <i>the</i> unitarity triangle	36
5	Charged versus neutral currents	38
6	Why FCNCs are so small in the Standard Model	39
6.1	Diagonal versus universal	39
6.2	FCNCs versus gauge invariance	39
6.3	FCNCs versus Yukawa alignment	40
6.4	FCNCs versus broken gauge symmetry reps	41
7	Parameterizing QCD	43
7.1	The decay constant	47
7.2	Remarks on the vector mesons	50
7.3	Form factors	52
7.4	Aside: Goldstones, currents, and pions	53
8	Flavor symmetry and the CKM	57
8.1	Measuring $ V_{ud} $	57
8.2	Measuring $ V_{us} $	61
8.3	Measuring $ V_{cs} $	69
8.4	Measuring $ V_{cd} $	73
9	Intermission: Effective Field Theory	77
9.1	EFT is not a dirty word	78
9.2	A trivial example: muon decay	81
9.3	The trivial example at one loop	83
9.4	Mass-independent schemes	86
9.5	Operator Mixing	86
10	Remarks on Lattice QCD	86
10.1	Motivation and errors	87
10.2	‘Solving’ QCD	87
10.3	The Nielsen-Ninomiya No-Go Theorem	88
10.4	The quenched approximation	90
11	Heavy quark symmetry and the CKM	90
11.1	The hydrogen atom	91
11.2	Heavy quark symmetry: heuristics	92
11.3	Heavy quark symmetry: specifics	93
11.4	HQET	95

11.5	Measuring $ V_{cb} $	97
11.6	Measuring V_{ub}	101
12	Boxes, Penguins, and the CKM	102
12.1	The trouble with top	102
12.2	Loops, FCNCs, and the GIM mechanism	103
12.3	Example: $b \rightarrow s\gamma$	103
12.4	History of the GIM mechanism	104
12.5	Measuring the $b \rightarrow s\gamma$ penguin	105
12.6	Measuring $b \rightarrow s\gamma$ versus $b \rightarrow d\gamma$	106
13	Meson Mixing and Oscillation	108
13.1	Open system Hamiltonian	109
13.2	Time evolution	112
13.3	Flavor tagging	112
13.4	Time scales	114
13.5	Calculating Δm and $\Delta\Gamma$	116
14	Today's lecture: CP violation	120
14.1	CPV in mesons	122
A	Notation and Conventions	124
B	Facts that you should memorize	125
C	Lie groups, Lie algebras, and representation theory	126
C.1	Groups and representations	126
C.2	Lie groups	128
C.3	More formal developments	131
C.4	$SU(3)$	133
D	Homework solutions	134
E	Critical reception of these notes	150
F	Famous Yuval Quotes	151

1 Introduction

“What is the most important symbol in physics? Is it this: $+$? Is it this: \times ? Is it this: $=$? No. I claim that it is this: \sim . Tell me the order of magnitude, the scaling. *That* is the physics.” –Yuval Grossman, 21 August 2008.

These notes are transcriptions of the Physics 7661: Flavor Physics lectures given by Professor Yuval Grossman at Cornell University in the fall of 2010. Professor Grossman also gave introductory week-long lecture courses geared towards beginning graduate students at the 2009 European School of High-Energy Physics and the 2009 Flavianet School on Flavor Physics [44] and the 2010 CERN-Fermilab Hadron Collider Physics Summer School [3].

A course webpage with homework assignments and (eventually) solutions is available at: <http://lepp.cornell.edu/~yuvalg/P7661/>.

There is no required textbook, but students should have ready access to the *Review of Particle Physics* prepared by the Particle Data Group and often referred to as ‘the PDG.’ The lecturer explains that the PDG contains “everything you ever wanted to know about anything.” All of the contents are available online at the PDG webpage, <http://pdg.lbl.gov/>. Physicists may also order a free copy of the large and pocket PDG which is updated every two years. The large version includes several review articles that make very good bed-time reading. The data in the PDG will be necessary for some homework problems.

Additional references that were particularly helpful during the preparation of these notes were the following textbooks,

- *Dynamics of the Standard Model*, by Donoghue, Golowich, and Holstein.
- *Gauge theory of elementary particle physics*, by Cheng and Li.

Both of these are written from a theorist’s point of view but do so in a way that is very closely connected to experiments. (A good litmus test for this is whether or not a textbook teaches chiral perturbation theory.)

Finally, much of the flavor structure of the Standard Model first appeared experimentally in the decays of hadrons. The techniques used to describe these decays are referred to as the current algebra or the partially conserved axial current and have fallen out of modern quantum field theory courses. We will not directly make use of these methods but will occasionally refer to them for completeness. While many reviews exist on the subject, including [80] and [79], perhaps the most accessible and insightful for modern students is the chapter in Coleman’s *Aspects of Symmetry* on soft pions [22].

Problem 1.1. Bibhushan missed part of the first lecture because he had to attend the course that he’s TA’ing, “Why is the sky blue?” As a sample homework problem, what is the color of the sky on Mars? (Solutions problems in these notes appear in Appendix D.)

Finally, an apology. There are several important topics in flavor physics that we have been unable to cover. Among the more glaring omissions are lepton flavor (neutrino physics), soft collinear effective theory for b decays, non-relativistic QCD, chiral symmetries, and current algebra techniques.

2 Model building

In this lecture we will briefly review aspects of the Standard Model to frame our study of its flavor structure. Readers looking for more background material can peruse the first few sections of [44].

The overall goal of high-energy physics can be expressed succinctly in the following form:

$$\mathcal{L} = ? \tag{2.1}$$

That is, our job is to determine the Lagrangian of nature and experimentally determine its parameters.

In order to answer this question we would like to build models. In fact, it is perhaps more accurate to describe a theorist's job not as model *building*, but rather model *designing*. In order to design our Lagrangian, we need to provide three ingredients:

1. The **gauge group** of the model,
2. The **representations** of the fields under this gauge group,
3. The pattern of **spontaneous symmetry breaking**.

The last point is typically represented by a sign, for example the sign of the Higgs mass-squared parameter at the unstable vacuum ($\mu^2 < 0$).

Once we have specified these ingredients, the next step is to write the *most general renormalizable Lagrangian* that is invariant under the gauge symmetry and provides the required spontaneous symmetry breaking pattern. This is far from a trivial statement. The ‘most general’ statement tells us that *all* terms that satisfy the above conditions *must* be present in the Lagrangian, even the terms that may be phenomenologically problematic. For example, even though we might not want to include a term that induces proton decay, we cannot simply omit it from our model without some symmetry principle that forbids it.

On the other hand, renormalizability strongly constrains the form of a Lagrangian and, in fact, limits us to only a finite number of terms. This condition comes to us from the principles of effective field theory and Wilsonian renormalization group. We assume that the UV (more fundamental) theory may generate all possible operators—including non-renormalizable terms—at the UV scale. By dimensional analysis the non-renormalizable operators depend on negative powers of the UV scale Λ . Thus at the ‘low’ energies $\mu \ll \Lambda$ where the Standard Model is valid we expect that these operators are suppressed by powers of $\mu/\Lambda \ll 1$. Thus in principle such non-renormalizable operators exist but they should come with small coefficients. Eventually we will be interested in including these terms and understanding their effect on low-energy observables. For more background on this subject, see reviews by Hollowood [52], Manohar [63], and *The Weak Interactions* by Georgi [40].

At this point a few remarks are in order. One can take these as rules for designing models.

- We must impose Poincaré invariance (the ‘gauge group’ with respect to gravity). Eventually in this course we may get to non-relativistic QCD which breaks Poincaré invariance at the *effective* level, but not the *fundamental* level.
- We use the language of quantum field theory.

- We do not impose global symmetries. We know, for example, that global symmetries are broken by gravity since a black hole might eat a proton and spit out a positron as part of its Hawking radiation. The gauged electromagnetic symmetry is protected, but the global baryon number conservation is broken. Another way of saying this which we will discuss below is that global symmetries are not ‘fundamental.’
- The ‘basic’ fermion representation are two-component chiral (Weyl) spinors. Four component fields end up becoming useful for some calculations, but four dimensional theories are generally chiral. An encyclopedic reference for Weyl spinors is [28].
- A model is not yet a description of nature until one experimentally measures the finite number of physical parameters associated with it. We will clarify what we mean by a ‘physical’ parameter below, but the finite number of such parameters comes from renormalizability. If a Lagrangian contains k physical parameters, then one must first perform k measurements before being able to predict the value of further measurements.

Note in the last point that the idea of the number of parameters required to define a theory is independent of how they are parameterized. In fact, many not-so-obvious ideas in field theory such as the renormalization group become obvious once one understands that they are reparameterizations of these physical parameters. In practice, if one makes $n > k$ measurements of a theory, one doesn’t first do k ‘parameter measurements’ and then $(n - k)$ observations of the theory. Instead one would take all n measurements and do a statistical fit (e.g. χ^2) for the k parameters to check for self-consistency. See James Wells’ lectures for a theorist’s discussion of this process [93, 92]. Philosophically, however, it is useful and important to remember that a model (a Lagrangian) by itself does not make predictions, it must come with measurements of its parameters.

2.1 Example: the Standard Model

Our best and most successful model is the Standard Model. It serves as an excellent example of the above principles because (1) it works and (2) it is familiar. The three ingredients are as follows:

1. GAUGE GROUP: $SU(3)_c \times SU(2)_L \times U(1)_Y$.
2. MATTER REPRESENTATIONS: a useful mnemonic for the fermion fields of the Standard Model is *QUDLE* (“cuddly”), which is short for the left-handed quark doublet Q , right-handed up U , right-handed down D , left-handed lepton doublet L , and the right-handed charged lepton E . Using the convention where all fields are written in terms of left-handed Weyl spinors (taking CP conjugations as necessary), we may write the gauge representations using the notation $(c, L)_Y$:

$$\begin{aligned}
 Q_i & (\mathbf{3}, \mathbf{2})_{1/6} \\
 U_i^c & (\bar{\mathbf{3}}, \mathbf{1})_{-2/3} \\
 D_i^c & (\bar{\mathbf{3}}, \mathbf{1})_{1/3} \\
 L_i & (\mathbf{1}, \mathbf{2})_{-1/2} \\
 E_i & (\mathbf{1}, \mathbf{1})_1,
 \end{aligned}$$

we've written the generation index $i = 1, 2, 3$. Note that this index does *not* imply a global symmetry. In addition to the fermions we also have a complex scalar Higgs field, ϕ , with the representation $(\mathbf{1}, \mathbf{2})_{1/2}$. Note that for $SU(2)$ the ϵ tensor allows us to convert between the fundamental $\mathbf{2}$ and antifundamental $\bar{\mathbf{2}}$ representations. For example, we may write $\tilde{\phi}_a = \epsilon_{ab}\phi^b$.

3. SPONTANEOUS SYMMETRY BREAKING: $SU(2)_L \times U(1)_Y \rightarrow U(1)_{EM}$. We can see this from the sign of the Higgs mass term, $\mathcal{L} \supset \mu^2|\phi|^2$, so that the potential contains a term $-\mu^2|\phi|^2$.

The next step is to write down the most general renormalizable Lagrangian. This can be divided into three parts,

$$\mathcal{L}_{SM} = \mathcal{L}_{kin} + \mathcal{L}_{Higgs} + \mathcal{L}_{Yuk}. \quad (2.2)$$

The kinetic Lagrangian \mathcal{L}_{kin} includes the gauge interactions through the covariant derivative and non-Abelian field strength. The Higgs Lagrangian \mathcal{L}_{Higgs} gives the Mexican hat potential

$$\mathcal{L}_{Higgs} = \mu^2|\phi|^2 - \lambda|\phi|^4. \quad (2.3)$$

The Yukawa part of the Lagrangian \mathcal{L}_{Yuk} contains the terms

$$\mathcal{L}_{Yuk} = y_{ij}^e \bar{L}^i \phi E^j + y_{ij}^d \bar{Q}^i \phi D^j + y_{ij}^u \bar{Q}^j \tilde{\phi} U^i + \text{h.c.} \quad (2.4)$$

In the Standard Model fermions only get masses through these terms. In a general model you can also have bare masses when gauge and spacetime symmetries permit it.

2.2 Global, accidental, and approximate symmetries

One might ask how we can have baryon number conservation if our rules above stated that we may not explicitly impose any global symmetries. The global symmetries of the Standard Model result as *outputs* of the theory rather than as external constraints. In particular, they come from the structure imposed by renormalizability and gauge invariance. Global symmetries that appear only because non-renormalizable terms aren't considered are called **accidental symmetries**. These are broken explicitly by non-renormalizable terms, but since these terms are small one can often make use of these symmetries. When people ask why the proton doesn't decay in the Standard Model, then baryon number conservation is only part of the answer; the 'fundamental' reason is that baryon-number-violating operators in the Standard Model are non-renormalizable.

In addition to accidental symmetries there are other **approximate symmetries** which are parametrically small in the sense that they become exact global symmetries when a parameter (or set of parameters) is set to zero. An example is **isospin symmetry**, which is an approximate global symmetry of the Standard Model. It is broken by the quark masses. Are these parametrically small? We can only say this about dimensionless parameters, and it turns out that the relevant parameter is the ratio of the quark mass splittings to the QCD strong coupling scale,

$$\frac{m_u - m_d}{\Lambda_{QCD}}. \quad (2.5)$$

Isospin symmetry is additionally broken by ‘the most famous small parameter in the world,’ α_{EM} since the charge of the u is different from the charge of the d .

There’s one more small parameter that you might be familiar with. The Higgs potential obeys a **custodial symmetry**. It is broken at one-loop by the Yukawa couplings. While one might argue that the Yukawas themselves are $\mathcal{O}(1)$ and thus not small, since the breaking only occurs at loop-level the relevant parameter is $y^2/16\pi^2$ which *is* small. For more information about custodial symmetry and a nice overview of the symmetries of the Standard Model, see Scott Willenbrock’s TASI 2004 lectures [94].

These accidental and approximate symmetries can be very useful, but we should always remember that they are not as fundamental as the gauge symmetries that we begin with.

2.3 Learning how to count [parameters]

As we mentioned above, a model will never graduate to a ‘description of Nature’ until its physical parameters are measured. What do we mean by *physical* parameters? It is useful to start by reviewing some *unphysical* parameters which cannot be measured. This is an idea that has been with us since Newtonian mechanics, where the absolute value of the energy is not physical (only differences are measurable). In quantum mechanics, a key idea is that the overall phase of a wavefunction is not physical.

We know that the Standard Model has 18 physical parameters. For a reminder of these parameters accessible to a general audience, see [18]. It doesn’t matter *how* we write these 18 parameters. In fact, we may even parameterize the Standard Model in different ways—but the number of physical parameters does not change.

To learn how to count these parameters, let’s consider the leptonic Yukawa term,

$$y_{ij}^e \bar{L}^i \phi E^j. \tag{2.6}$$

How many parameters does the matrix y^e have? Since it is a 3×3 complex matrix, it contains 18 real parameters. We know, however, that the SM lepton sector only has three physical parameters: the masses of the charged leptons. Thus there are 15 unphysical parameters hidden in (and obfuscating) the y^e matrix! Since unphysical parameters are useless for predictions, it’s always nicer to choose parameterization where everything is physical.

Let’s see how we could have figured out that there were 3 physical parameters in this sector of the Standard Model. If $y^e = 0$, then the Lagrangian would enjoy a larger global symmetry: the global symmetry of the kinetic term, $\mathcal{L}_{\text{kin}} \supset \bar{E}_i \not{D} E_i + \bar{L}_i \not{D} L_i$. This is a $U(3)_E \times U(3)_L$ global symmetry associated with the rotation between the three generations of E and L . The Yukawa term breaks the global symmetry down to $U(1)^3$ (corresponding to the phases of the mass eigenstate fields). We may use the now-broken $U(3)_E \times U(3)_L$ symmetry to rotate the symmetry-breaking direction y^e to a convenient direction. Clearly this rotation will use only the broken generators of $U(3)_E \times U(3)_L$. We call y^e a **spurion** of this symmetry. A particularly convenient direction is to rotate y^e to a diagonal matrix proportional to the charged lepton masses where it is now clear that there are three physical parameters. To recap, we have used the *broken* generators of the symmetry group to remove *unphysical* parameters by orienting the spurion y^3 to convenient direction in symmetry space. (It is tautological that the unbroken generators would not affect the spurion direction.)

From this simple example one can see that the number of physical parameters is related to the generators of the unbroken symmetry by

$$\# \text{ physical parameters} = \# \text{ parameters} - \# \text{ broken generators.} \quad (2.7)$$

In the case of the lepton Yukawa there were 15 broken generators out of 18 total generators, confirming that there are 3 physical parameters in this sector. Make sure you understand why this formula makes sense. Unlike unbroken generators which represent symmetries of the system, the generator of a *broken* symmetry will change a system and so allow us to align the system to match our parameterization and—by doing so—eliminate unphysical parameters. This rather abstract statement can be made very concrete with the following example.

A very simple example. To rephrase the above argument in terms of a very simple example from quantum mechanics, recall the Zeeman effect where the energy levels of the Hydrogen atom are split in the presence of a magnetic field. In the absence of such a field the system enjoys an $SO(3)$ rotational symmetry and so has degenerate ground states. The field breaks the symmetry and splits the degeneracy by breaking $SO(3)$ down to $SO(2)$, which is the remaining rotational symmetry perpendicular to the direction of the field. There are two broken generators. The physical parameter in the system is the magnitude of the magnetic field. Generically, the field points in an arbitrary direction $\mathbf{B} = B_x \hat{\mathbf{x}} + B_y \hat{\mathbf{y}} + B_z \hat{\mathbf{z}}$. We can use the two broken generators to rotate the coordinates so that $B_x = B_y = 0$. These were unphysical parameters. This leaves a single physical parameter, B_z , which in these coordinates is precisely the magnitude of \mathbf{B} .

2.4 Counting parameters in low-energy QCD

We may now generalize this picture to more complicated groups. Let us consider the low-energy Lagrangian for the strong force in which only the light quark species (u , d , s) are active. Recall that this is the underlying structure of Murray Gell-Mann’s Eightfold Way for light hadron classification. As before the three ingredients for our model are:

1. GAUGE GROUP: $SU(3)_c$. We ignore $U(1)_{EM}$ as a small perturbation.
2. MATTER REPRESENTATIONS: for each flavor (u , d , s) we have a left-chiral fundamental representation and a right-chiral fundamental representation. Thus we have:

$$\begin{array}{ll} u_L & \mathbf{3} \\ d_L & \mathbf{3} \\ s_L & \mathbf{3} \end{array} \qquad \begin{array}{ll} u_R & \mathbf{3} \\ d_R & \mathbf{3} \\ s_R & \mathbf{3}. \end{array}$$

3. SPONTANEOUS SYMMETRY BREAKING: at *tree level* there is none. At the quantum level $SU(3)_A$ is broken spontaneously by the QCD chiral condensate, $\langle \bar{q}_L q_R + \bar{q}_R q_L \rangle$.

As before we write out the most general renormalizable Lagrangian,

$$\mathcal{L} = \sum_{i=u,d,s} (\bar{q}_L^i i \not{D} q_L^i + \bar{q}_R^i i \not{D} q_R^i) + (m_{ij} \bar{q}_L^i q_R^j + \text{h.c.}). \quad (2.8)$$

Note the important difference between this low-energy QCD Lagrangian and the Standard Model: the flavor symmetry of the kinetic term allows one to rotate up, down, and strange quarks (of a given chirality) between one another! In the Standard Model this is prohibited because this mixes different components of $SU(2)_L$ doublets or, alternately, because this mixes particles of different charge. Do not confuse this $U(3)_L \times U(3)_R$ flavor symmetry with the $U(3)_L \times U(3)_R$ flavor symmetry of the Standard Model's quark sector which mixes the three *generations* between one another. (There are a lot of 3s floating around, make sure you don't mix them up.)

It is useful to write down this flavor symmetry in terms of vector and axial symmetries,

$$U(3)_L \times U(3)_R = U(3)_V \times U(3)_A, \quad (2.9)$$

where the vector (axial) transformation corresponds to

$$q_L^i \rightarrow q_L^i = U^i_j q_L^j \quad q_R^i \rightarrow q_R^i = U^{(\dagger)i}_j q_R^j. \quad (2.10)$$

In other words, the vector symmetry rotates the left- and right-handed quarks in the same way while the axial symmetry rotates them oppositely.

Now let's look at the mass terms of low-energy QCD. We may diagonalize m_{ij} , find the eigenvalues, and write down the three physical mass parameters. Thus for a generic mass matrix we see that flavor is a good quantum number and the $U(3)_V \times U(3)_A$ symmetry of the kinetic term is broken down to $U(1)_V^3$ representing the phases of each flavor mass eigenstate (these are now *Dirac* spinors). If, for some reason, the mass matrix is universal, $m \propto \mathbb{1}$, then the mass terms break $U(3)_L \times U(3)_R$ to $U(3)_V$. We see that the mass terms always breaks the axial part of the kinetic term's flavor symmetry.

One might now ask the following clever question:

Why should we include the mass terms in the QCD Lagrangian at all since we know these come from the Yukawa sector of the Standard Model?

While this is true, we must remind ourselves that our model building rules tell us that we *must* include all renormalizable terms that respect our gauge symmetries. For low-energy QCD this means that we must include these mass terms. Now one might ask an even more clever question:

Fine, we include these terms, but then we go out and measure them and they're very small. Shouldn't we still be able to ignore these terms since they are much smaller than the relevant mass scale, Λ_{QCD} ?

Indeed! We know that $m_q \ll \Lambda_{\text{QCD}}$ for all active quarks at low energies, so the limit $m_{ij} \rightarrow 0$ (or alternately taking the dimensionless parameter $m_{ij}/\Lambda_{\text{QCD}} \rightarrow 0$) should be sensible. However, the point is that in the absence of masses the chiral condensate breaks the $SU(3)_A \subset U(3)_A$ and does so at the scale Λ_{QCD} (recall that $\langle \bar{q}_i q_j \rangle \sim \Lambda_{\text{QCD}}^3 \delta_{ij}$). Thus we really should assume that this symmetry is broken. One more clever retort:

Fine! The $SU(3)_A$ should be broken either by mass terms or by the QCD condensate, but what about the remaining $U(1)_A \subset U(3)_A$?

If we do our parameter counting, we have m_{ij} an arbitrary complex matrix with 18 real parameters. The breaking of $U(3)_V \rightarrow U(1)_V^3$ gives $3^2 - 3 = 6$ broken generators. $SU(3)_A$ is broken and gives us $3^2 - 1 = 8$ broken generators. Thus if we pretend that $U(1)_A$ is unbroken, we would have $18 - 14 = 4$ physical parameters. We know that the correct answer is 3, the three Dirac quark masses. What has happened?

The answer is that $U(1)_A$ is broken: it is anomalous, i.e. broken by quantum effects. The $U(1)_A$ axial anomaly is notorious in the history in particle physics¹. For our purposes we must make the additional rule that at the *quantum level*, anomalous symmetries are no good for counting parameters. If this is the case, then there's one less broken symmetry generator and so we expect there to be one more physical parameter relative to the classical analysis. Does such a parameter exist? Yes; it is precisely the non-perturbative Θ_{YM} term which transforms as a shift with respect to $U(1)_A$!

Thus there are two ways of counting parameters:

1. CLASSICALLY: look only at the most general, renormalizable perturbative Lagrangian (i.e. without Θ terms) and perform the counting using the classical symmetries without any regard as to whether or not they are broken by anomalies.
2. QUANTUM MECHANICALLY: consider the most general, renormalizable Lagrangian *including* non-perturbative terms and only consider non-anomalous symmetries.

Both are consistent as long as you *stay within the regime of the description*, in other words, if you're counting classical (tree-level) parameters, then don't include quantum effects like anomalies and Θ angles. Conversely, if you're counting quantum parameters, then you must include *both* the effect of anomalies and the non-perturbative terms associated with them.

2.5 Counting parameters in the Standard Model

Let's go back to our favorite example, the Standard Model. Where do the Standard Model's 18 parameters come from?

- The kinetic Lagrangian \mathcal{L}_{kin} contains three physical parameters, the couplings g', g, g_3 which determine the strengths of each force.
- The Higgs Lagrangian $\mathcal{L}_{\text{Higgs}}$ contains two physical parameters which we may take to be the vacuum expectation value of its neutral component v and the quartic interaction strength λ . (One could also have chosen the unstable mass $-\mu^2$ or the physical mass m_H^2 .)
- Finally, the Yukawa Lagrangian contains the three charged lepton masses which we discussed above plus ten physical parameters associated with the quark Yukawa sector.

¹See [34] for an excellent presentation with an emphasis on instantons by a reputable, charming, and handsome author.

To explain the last ten parameters, recall that the quark Yukawa sector takes the form

$$\mathcal{L}_{\text{Yuk}} \sim y^u \bar{Q} \phi U + y^d \bar{Q} \phi D. \quad (2.11)$$

There are 36 total real parameters between the y^u and y^d . The generational symmetry of the kinetic term is $U(3)^3$ which gives 27 total generators. How many of these survive after turning on the y^u and y^d interactions? Baryon number $U(1)$ survives with one generator, and so we have 26 broken generators. The number of physical parameters is $36 - 26 = 10$.

In fact, we can separate the physical real parameters from the physical phases. A generic Yukawa matrix has 9 real parameters and 9 phases, one for each element: $y_{ij} = r e^{i\theta}$. How many real parameters and phases are there in a unitary matrix? We know that a unitary $U(N)$ matrix has N total parameters (N^2 complex elements with the constraint equation $U_{ik}(U^\dagger)_{kj} = \delta_{ij}$ for each element) and that an $O(N)$ *orthogonal* matrix has $\frac{1}{2}N(N-1)$ parameters which are all real. Since an orthogonal matrix is just a unitary matrix “without the phases,” we may subtract the latter from the former to find that an $N \times N$ unitary matrix has $\frac{1}{2}N(N+1)$ phases and $\frac{1}{2}N(N-1)$ real parameters. Thus our $U(N)^3$ kinetic term symmetry gives 9 real parameters and 18 phases. The Yukawas break this to $U(1)_B$ so that the number of real physical parameters is

$$2 \text{ Yukawas} \times 9 \text{ } \mathbb{R} \text{ param. each} - 9 \text{ broken } \mathbb{R} \text{ generators} = 9 \text{ physical } \mathbb{R} \text{ parameters.} \quad (2.12)$$

Similarly for the complex phase,

$$2 \text{ Yukawas} \times 9 \text{ phases each} - 17 \text{ broken phase generators} = 1 \text{ physical phase.} \quad (2.13)$$

These physical parameters are the quark masses, the mixing angles, and the CP-violating phase.

The above counting came from a classical analysis. Let us briefly remark on what happens in the quantum picture. The $U(1)_A$ symmetry is anomalous and so should not be counted as having ever been a symmetry of the kinetic term. This gives one less broken generator again giving us the Θ_{QCD} term.

Problem 2.1. Extra generations in the Standard Model. Count the number of physical flavor parameters in the quark sector of the Standard Model with N generations. Show that the quark sector of such a model has $N(N+3)/2$ real parameters and $(N-1)(N-2)/2$ phases. Determine the number of mixing angles.

Problem 2.2. Exotic light quarks, part I. Consider a copy of the Standard Model with a modified quark sector without the c , b , and t quarks. The remaining quark representations are modified to $Q_L = (u_L, d_L)$ forming an $SU(2)_L$ doublet and s_L an $SU(2)_L$ singlet. All the right-handed quarks are singlets (u_R, d_R, s_R) . All color and electric charges are the same as in the Standard Model. Electroweak symmetry follows as in the usual Standard Model. How many physical parameters are in the model? How many of them are real and how many are phases?

Problem 2.3. Parameters of the MSSM. Show that there are 110 physical parameters in the flavor sector of the minimal supersymmetric Standard Model (MSSM). Don’t forget to include general soft SUSY-breaking terms. How do these divide into masses, mixing angles, and phases?

Problem 2.4. The Θ_{Weak} angle for $SU(2)_L$. In Section 2.4 we discussed how the counting of parameters pointed to the appearance of Θ_{YM} as a physical parameter when taking into account quantum corrections. Use a similar argument to show that such an angle is *not* physical for $SU(2)_L$. (**Hint:** one could alternately ask why the anomalous $U(1)_{B+L}$ doesn't affect our Standard Model counting.)

3 A review of hadrons

Peeking at the PDG for the first time may intimidate a student by the many long lists of particles. One worries that this aspect of particle physics reduces to memorizing the periodic table or, even worse, memorizing the botanical names of flowers. This distaste for the hadronic spectra seems more pronounced in students of theory than experiment since the latter very quickly familiarize themselves with the kinds of particles that appear in different parts of a detector. Theorists, on the other hand, tend to want to live in a magical place where Lagrangians are written not in terms of hadrons but in terms of ‘fundamental’ constituents like u and d quarks. A major theme of this course is to be able to pull one’s head out of the UV to be able to make the transition between ‘fundamental’ Lagrangians and the actual measurable physical states at low energies. This transition is far from trivial and a good theorist must be able to go back and forth between the Standard Model Lagrangian, low energy spectra, and—hopefully—a new physics Lagrangian. It will be useful to have the PDG (or the PDG website) nearby while reading this section. A good refresher of the topics in this section are the undergraduate textbooks by Griffiths [42] and Perkins [70].

As you have known since kindergarden, QCD states at low energies are hadrons and we can distinguish between mesons, which are bosons, and baryons, which are fermions. With that out of the way, we can move on to slightly more non-trivial things.

3.1 What we mean by ‘stable’

Next let us highlight an important abuse of language. We will divide hadrons between **stable** particle and **unstable** particles, where the latter are often called **resonances**. These do not refer to stability in the ‘absolute’ sense. Rather, they refer to stability with respect to whether a particle is “stable on detector time scales.” Recall that there are two ways to measure an intermediate particle’s lifetime: (1) one can directly measure its decay width or (2) for sufficiently long-lived particles one can measure a displaced vertex. Stable particles are those whose lifetimes are at least large enough to be measured by the latter method.

A more concrete definition is that a stable particle is one that does not decay through QCD interactions. In this sense a kaon is a stable meson, it certainly decays, but these channels are only through the weak interaction. In fact, let’s make a somewhat fancy definition that we’ll justify later:

Definition 3.1. A **stable particle** is one that is either an eigenstate of the Hamiltonian or that only decays through weak interactions; in other words, it is an eigenstate of the Hamiltonian in the limit $g \rightarrow 0$.

Resonances are then defined to be those particles which are not stable. We can now connect the ‘intuitive’ definition of stability with this semi-technical definition by looking at the characteristic lifetimes of hadrons that decay via the strong versus weak forces.

Example 3.2. Consider the ρ meson. This is a resonance of up and down quarks. What is its approximate mass and width? Since the ρ is composed of light quarks, the dominant contribution to its mass comes from QCD. We thus expect the mass to be on the order of Λ_{QCD} , which is hundreds of MeV. What is its width? The couplings are order 1 so we estimate it to again be Λ_{QCD} . Now go ahead and look it up in the PDG. unless we specify an excited state we will mean the lowest-mass particle by that name, so in this case we mean the $\rho(770)$. We find $m_\rho \approx 775$ MeV and $\Gamma_\rho \approx 150$ MeV. Not bad.

The ρ width is very large, just a factor of five smaller than its mass. In fact, some resonances have a width that is so large (on the order of their mass) that it becomes very hard to determine whether or not it is a bound state at all. Recall that when we do QFT we work with states in the asymptotic past and future. When a particle’s decay width is on the order of its mass, the notion of asymptotic state becomes ill defined. The bottom line is that all of these resonances decay very quickly.

What about the stable particles?

Example 3.3 (This example about useful numbers that you should memorize.). What about the B^0 meson? *Typist’s note: because this is one of Yuval’s famous stories, I have tried to transcribe the following just as it was dictated in lecture. See also [1]*

What is the mass of the B meson? All of you know the story right? No? What is the mass of the B meson? Flip? (*No response.*) Okay, let’s start with a little story. When I was young—like Flip’s age—I went to TASI, just like Flip. And we started hiking, just like Flip². And I saw this sign, and the sign said, “Mile High City.” And it said “mile high” and it said “5280 feet.” You know, I was like—I couldn’t—I couldn’t breathe for a few seconds. That’s *exactly* the mass of the B meson in MeV! Exactly, up to four digits. Okay? So the [moral of] the story is as follows. You always have to remember [it] like this. It’s always much easier to remember two things than [it is to] remember one thing. Since most of you don’t even know how many feet there is in one mile, now you have two things to remember: the number of feet in one mile is exactly mass of the B meson in MeV.

What about the lifetime of the B^0 meson? It decays weakly and turns out to have a lifetime of $\tau_B = 1.5 \times 10^{-12}$ seconds (about a picosecond). A useful conversion is

$$1 \text{ GeV} = 1.5 \times 10^{24} \text{ sec}^{-1}. \quad (3.1)$$

From this one ends up with a decay width of $\Gamma_B = 4 \times 10^{-4}$ eV.

Problem 3.1. Derive (3.1) if you are not already familiar with it.

²Addendum: Flip didn’t actually do any hiking at TASI because he spent the weekends calculating penguin diagrams with Yuval’s student.

Now we can see the difference between a resonance and a stable particle. Consider the order of magnitudes of the widths of the ρ and the B^0 :

$$\Gamma_B \sim 10^{-4} \text{ eV} \tag{3.2}$$

$$\Gamma_\rho \sim 10^8 \text{ eV}. \tag{3.3}$$

This is why we call the B ‘stable.’ You should also have intuited the significance of the term resonance, since these particles tend to be QCD excited states of a given quark content. It is worth remembering that this big difference between the weak and the strong interaction doesn’t come from the $SU(2)_L$ coupling constant being much smaller than the strong coupling (it’s not); it comes from the virtual W that is emitted by a b quark in the decay of the B meson.

A nice analogy is tritium: this is a hydrogen isotope with two neutrons. The excited 2P tritium state will emit a photon to decay to the 1S ground state at a time scale on the order of a nanosecond. The 1S state will eventually decay into helium-3 with a lifetime on the order of 18 years. One can similarly imagine an excited (heavy) B resonance decaying rapidly to the B^0 by emitting a pion and then the B^0 decaying via the weak interaction on a much longer time scale.

One should be comfortable thinking about the plethora of hadrons as excited hydrogen atoms in basic quantum mechanics. In fact, let us turn to one of the basic features of hydrogen in quantum mechanics: the quantum numbers that describe a state.

A PDG hint. In the PDG particles whose name contains their mass in parenthesis, e.g. $\rho(770)$, are resonances while those that do not are stable.

3.2 Hadron quantum numbers

There are three kinds of hadronic quantum numbers:

1. Exact quantum numbers
2. Quantum numbers which are exact under QCD, but not the weak interactions
3. Approximate quantum numbers, even within QCD.

There are two exact quantum numbers: electromagnetic **charge** Q and **spin** J , by which we mean the sum of the constituent particle spins and the orbital angular momenta. Since much of this course will involve particle mixing, these quantum numbers are important because they give strict rules about which particles can mix. If two particles have the same quantum numbers then [in general] they mix, i.e. the state which commutes with the Hamiltonian is generically a linear combination of all such states. However, states with different quantum numbers *cannot* mix.

There is actually one more exact quantum number that we should have guessed from the hydrogen atom analogy, the mass (the ‘ N ’ quantum number). This is certainly conserved, but we won’t use this so much since we’ll often be working with states that are flavor but not mass eigenstates.

Example 3.4. Do the π and ρ mesons mix? No—they have different spin and spin is an *exact* symmetry and *must* be conserved.

Note that in order for particles to mix they needn't be exactly degenerate, but having a small mass difference certainly helps³. One might also ask why we don't list **baryon number** as an exact quantum number. Since baryons have $2J + 1$ even and mesons have $2J + 1$ odd, baryon number conservation automatically comes out of conservation of angular momentum. (We will ignore objects such as glueballs.)

Next let us consider the quantum numbers that are not quite exact but are at least exact with respect to QCD. These quantum numbers are **parity** and **charge conjugation**. We will be particularly interested in their combination, CP . These are 'parities' in the sense that they are \mathbb{Z}_2 discrete symmetries. Mesons in the PDG are denoted by their spin, parity, and charge conjugation parity by J^{PC} . The superscript is a mnemonic that they are only approximate—but *really good*—symmetries.

Example 3.5. Look up the listing for the π^\pm meson in the PDG. Why does it only list J^P instead of J^{PC} ? (Hint: what do we know about degenerate states versus eigenstates?) The π^\pm are degenerate states that are conjugate to one another, i.e. a state π^+ has neither charge parity plus or minus since it's not an eigenstate of charge parity. How can this be? Recall that degenerate states needn't be states of the symmetries of the theory, all that is required is that a superposition of these states is an eigenstate.

While you're looking at the π^\pm , you'll note that there are additional quantum numbers listed: I^G . I is isospin, which we'll discuss below. What is G ? This is G -parity, a generalization of charge quantum number that acts on isospin multiplets to replace the C quantum number for particles like the π^\pm . The strict definition is

$$G = C e^{i\pi I_2}, \tag{3.4}$$

where the second factor is a rotation in isospin space that takes $I_3 \rightarrow -I_3$.

Example 3.6. Consider the π^0 meson. The PDG tells us that it has $J^{PC} = 0^{-+}$. Does this make sense? Why should the pion have negative internal parity? We know that it is a pair of quarks. It's the lightest meson so it wants to have $J = 0$ with as little total and orbital angular momentum as possible. It does this by having two antiparallel spin one-half quarks and so it must pick up a minus sign under parity. Indeed, everything makes sense.

What about T ? We know that CPT is required for any consistent Lorentz invariant quantum field theory. The violation of CP by the weak interactions thus implies a (compensating) violation of T . In neutrino physics, the CP conjugate of the process $\nu_\mu \rightarrow \nu_e$ is $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. On the other hand, the T conjugate is $\nu_e \rightarrow \nu_\mu$. CPT tells us that the amplitudes for the CP conjugate and T conjugate processes must be equal. One particularly interesting search for T violation is through the electric dipole moment (EDM) of the neutron. EDM of an elementary

³This is better stated as a comparison of the diagonal versus off-diagonal elements of the mass matrix. A nice related undergraduate-level quantum mechanics question is to ask the validity of degenerate perturbation theory when the states have a small mass splitting. For a given mass splitting, how big are the corrections to degenerate perturbation theory from non-degenerate perturbation theory and how do these compare to the corrections from the next order in degenerate perturbation theory?

particle points in the direction of its spin \mathbf{s} (by Lorentz invariance there is no other possibility) and is proportional to the electric field \mathbf{E} . Note that

$$\begin{array}{ll} P\mathbf{s} &= \mathbf{s} & P\mathbf{E} &= -\mathbf{E} \\ T\mathbf{s} &= -\mathbf{s} & T\mathbf{E} &= \mathbf{E}, \end{array}$$

Thus a nonzero EDM is a signal of both P and T violation. For more information about this see [70, 75].

Finally, the approximate quantum numbers are those associated with flavor. An example of ‘flavor’ for the light (i.e. ‘flavorless’) mesons composed of only u and d quarks is **isospin**. This is because when we neglect the mass difference between the up and the down quarks ($\Delta m_{ud} \ll \Lambda_{\text{QCD}}$) then u and d form an isodoublet and can mix.

What other approximate flavor quantum numbers might we consider? Charm and strangeness are ‘okay.’ Actually, strangeness is the worst possible example we could have chosen since its mass is so close to Λ_{QCD} . On the other hand, beauty (bottom-ness, if you prefer) is a *good* approximate quantum number because it’s so heavy. Can a B meson and a charmed meson mix? Only by a very small amount since the b quark is so heavy. When you see a stable meson with a mass on the order of 5 GeV you *know* it contains a b quark. The s quark, on the other hand, is not particularly light nor heavy. We’ll get to know its mesons soon.

The punchline of this discussion is that degenerate states with the same quantum numbers generally mix, sometimes with large or small mixing. When they have different *exact* quantum numbers the *cannot* mix, but having different *approximate* numbers means that they can mix a little. How much mixing is allowed for a given approximate quantum number will be a big part of our course.

Problem 3.2. Using the PDG. Use the PDG to answer the following questions:

1. What are the component quarks of the D^+ meson? What is its mass?
2. What are the component quarks of the Λ baryon? What is its spin?
3. What is $\text{Br}(\tau \rightarrow \mu\nu\bar{\nu})$?
4. What is the width (in eV) of the B^+ meson?
5. What is the average distance a B^+ meson will travel if $\gamma = 4$?

3.3 Binding energy

As we already started to discuss above, the masses of the hadrons come from two sources: the ‘bare masses’ of the quarks (coming from electroweak symmetry breaking) and the binding energy coming from QCD interactions. We may continue our analogy with hydrogen since the mass of a hydrogen atom comes from the masses of its constituent proton and electron plus a potential energy $V = -13.6$ eV. Indeed, we can experimentally check that the mass of hydrogen is 13.6 eV less than the sum of the masses of the proton and electron. In this sense this analogy is *too* good,

since this statement comes from being able to physically pull a proton and electron asymptotically apart and measure their masses independent of their mutual electric field.

In QCD the situation is more complicated since confinement tells us that we cannot pull individual quarks apart and measure them independently. What we can do is something a bit counterintuitive and sneaky: we can probe these quantities at high energies where QCD is weak. If we can extract the quark masses at high scales then we can *hope* to run these masses down to the IR, but it's clear that we're cheating a little. The act of running the masses down involves the inclusion of interactions, which is precisely what we wanted to separate from our mass measurement. Of course, even the measurement of the masses in the UV is tricky since UV quantities don't depend strongly on low-scale masses. Another way of saying this is that we usually define masses at low energies.

This is really a problem of the light quarks since we have probed them at energies much larger than their pole masses (to the extent that a pole mass is well defined). The bottom line is that these quark masses are very regularization-scheme dependent. When someone tells you a quark mass, the particular value shouldn't mean much without an accompanying regularization scheme. This is like someone offering to sell you a copy of *Mathematica* without specifying whether its for Linux or Mac.

While we're on the subject of differences between QCD and QED, let us remark that you should have already noticed that unlike atoms, hadrons can have masses *larger* than the sum of their constituent quarks. In other words, the potential energy can be positive. This, of course, is a signature of confinement.

3.4 Light quarks, heavy quarks, and the heaviest quark

So far we've discussed two kinds of quarks and hinted at properties of their associated hadrons. There are the light quarks (u, d, s) which have masses—even though they are renormalization-scheme dependent—much less than Λ_{QCD} . These guys have their relevant mass scales set by Λ_{QCD} and the differences between their masses are negligible enough to talk about approximate symmetries between them. (This is a bit of a stretch for the s quark, but it's certainly true for the u and d .)

On the other hand, there are the heavy quarks (c, b, t) with masses much greater than Λ_{QCD} . Unlike the light quarks, mesons made of heavy quarks have their mass scales set by the bare masses of their valence quarks. Again, the c is a bit of a borderline case, but we don't have to be picky. This is certainly true for the stable B mesons,

$$m_b \approx 4500 \text{ MeV} \tag{3.5}$$

$$m_{B^0} = 5280 \text{ MeV}. \tag{3.6}$$

What about the top quark?

It turns out that because the top quark is *so* heavy ($m_t \gg \Lambda_{\text{QCD}}$) that it decays before it hadronizes. For example, at a collider we might produce a b quark that hadronizes into an excited B meson (resonance) which quickly decays into a stable B meson which must then decay by the weak interaction. We understand this story from our ρ and B examples above: the weak interaction widths are much smaller than strong interaction widths. For the top quark, on the

other hand, the weak decay is *more* important than the strong force. The top width is 2 GeV from a weak decay which is much greater than Λ_{QCD} . Thus the top quark will decay via a W well before it can form any kind of hadron. In fact, because $m_t > m_W$, the top decays into an *on-shell* W which gives additional kinematic enhancement since it decays via a two-body decay rather than a three-body decay required by the b .

We sometimes say that the top “doesn’t form hadrons.” We should be careful with how we phrase this. Because the top decays so quickly we can never identify the hadron that it forms. Arguably it is even an ill-posed question to identify the top hadron in that time scale. In principle, though, we can “turn off” the weak interaction and calculate a spectrum of top hadrons. This is a very strongly coupled problem that is intractable by known techniques, but *in principle* one can calculate the spectrum of top hadrons when the weak force is negligible. In this sense it’s not quite correct to say that there’s ‘no such thing’ as a top hadron, it’s just that we don’t see them.

The reason why this linguistic issue is important is that it puts one in danger of turning poor language into poor physics. For example, one might want to try to explain that the top quark doesn’t form hadrons because such hadrons are not eigenstates of the full Hamiltonian. This is *wrong!* Sure they’re not eigenstates of the Hamiltonian, but *neither are any of the ‘stable’ hadrons that we’ve been discussing*, even those made up of light quarks! Any particle which decays is not an eigenstate of the Hamiltonian. When we discuss particles like the B meson, we really mean objects that are eigenvalues of the Hamiltonian in the limit when the weak interaction is turned off. This now ties back to Definition 3.1 for a stable particle which is defined in the limit when the $SU(2)_L$ coupling $g \rightarrow 0$.

3.5 Masses and mixing in mesons

Why do some mesons have definite quark content while others oscillate?

We’ve discussed the ‘problem’ of QCD where quarks are tied up in strongly-coupled bound states like mesons. There’s a second problem associated with QCD that is of utmost importance to us: not only is the quark tied up in a confined state, but once a quark of a given flavor is bound up in a physical-state meson, the meson doesn’t necessarily preserve that quark’s flavor. As you know a u will want to hadronize into a pion, but the pion will oscillate ($\pi \sim u\bar{u} + d\bar{d}$) and not necessarily want to stay a in bound state u -quarks.

A major theme in this course will be how flavor-mixing is induced by objects that naturally live in different bases. The most famous example is the CKM matrix which rotates between the mass and interaction bases of the Standard Model. Before turning to this, however, we can already see this idea at work in the mixing apparent in meson states.

If we work in a basis where we ‘know’ the masses of different quarks (modulo RG-dependence), then we can say that a given quark pair describes meson state. We can say that such a combination has a definite ‘flavor’ in the sense that the valence quark and anti-quark are well-defined.

Example 3.7. A neutral pion has valence quark structure $\pi^0 = (u\bar{u} - d\bar{d})/\sqrt{2}$. We can say that it is 50% a $u\bar{u}$ state and 50% a $d\bar{d}$ state.

However, these mesons can mix so that the physical meson states are linear combinations of quark pairs of different flavor combinations. On the other hand, we already mentioned above that

B and charmed mesons shouldn't mix appreciably because the b quark makes the B -meson so much heavier. Why is it that the π^0 has a 45° mixing while the B^0 is [nearly] a definite flavor state? Why can't we have a physical meson that is an appreciable admixture of bottom and charm, $(\bar{b}d + \bar{c}u)/\sqrt{2}$? Nothing 'fundamental' (symmetries, for example) prohibits this, so what is it that tells us how mesons mix?

It turns out that the answer is related to the origin of meson masses, which we saw above can predominantly come from either (1) the QCD binding energy (of order Λ_{QCD}), or (2) the valence quark masses. In QCD we know that isospin is a good symmetry in the limit where the quark masses are small compared to Λ_{QCD} . This holds for the light quarks (u, d, s). In this case the dominant contribution to the meson masses comes from the QCD potential energy and the mesons have masses on the order Λ_{QCD} plus, for higher resonances, any additional energy from internal degrees of freedom such as orbital angular momentum. Each light quark has isospin $1/2$ so that the bound state $q\bar{q}$ (for $q = u, d, s$) has isospin $0 \oplus 1$. In other words, we may have iso-singlets and iso-triplet mesons coming from light quark bound states. This is just the same story as hyperfine splitting in the hydrogen atom. Just like the hydrogen atom the singlet and triplet states have different energies (masses), but within the triplet all three states are [nearly] degenerate. We say that for the light quark mesons like the π^0 the QCD potential energy gives mass according to isospin and not flavor. What about mesons such as the B^0 whose masses come from a heavy bare quark mass? In this case the meson knows about flavor because the bare mass contribution from the heavy quark is larger than the Λ_{QCD} -scale potential energy contribution. It is now *flavor* and not isospin that determines the meson mass.

We should think these effects in terms of 2×2 matrices⁴. First consider the mass term generated by the QCD potential. This is a diagonal matrix acting on a space spanned by the iso-singlet $|1\rangle$ and iso-triplet $|3\rangle$ mesons,

$$\mathcal{L}_{\text{mass, iso}} \sim (\langle 1| \quad \langle 3|) \begin{pmatrix} m_1 & 0 \\ 0 & m_3 \end{pmatrix} \begin{pmatrix} |1\rangle \\ |3\rangle \end{pmatrix}. \quad (3.7)$$

Consider the mass terms coming from the valence quarks. For simplicity let's consider only the case where this contribution is dominated by only one of the quarks (mesons with two heavy quarks have shorter lifetimes). We can also write this contribution in terms of a 2×2 matrix, but this time acting on the 'flavor-basis' states defined by the presence of the heavy quark:

$$\mathcal{L}_{\text{mass, flavor}} \sim (\langle q| \quad \langle q'|) \begin{pmatrix} m_q & 0 \\ 0 & m_{q'} \end{pmatrix} \begin{pmatrix} |q\rangle \\ |q'\rangle \end{pmatrix}. \quad (3.8)$$

The point is that we must take the sum of these two mass matrices, but (3.7) and (3.8) are written in *different bases*. Thus shifting one matrix to the other basis will introduce off-diagonal terms which indicate mixing. The basis of mass eigenstates is generally something in between these two bases, the relevant question is whether the mass basis is closer to the isospin basis or the flavor basis. In other words, we must ask which effect is more important: isospin splitting (in this

⁴The best physicists make the most out of simple systems, such as the diagonalization of 2×2 matrices. In fact, the diagonalization of 2×2 matrices were defining moments for Cabbibo (leading to the CKM matrix) and Glashow, Iliopoulos, and Maiani (the GIM mechanism). In this sense the thing that separates great physicists is knowing *which* 2×2 matrices to diagonalize.

case singlet-triplet splitting) or the difference in heavy quark masses. If isospin splitting is more important, then we can use the isospin basis and treat the valence quark masses as a perturbation that introduces off-diagonal terms. This is what happens when the mesons such as the π^0 that are made up of light quarks whose masses are all small.

On the other hand, if the quark mass differences are more important, then we should use the flavor basis and treat the isospin-dependent masses as a perturbation that introduce off-diagonal terms. This is the case with our hypothetical mixing between a $\bar{b}b$ and a $\bar{c}c$ (these take the place of the $|q\rangle$ and $|q'\rangle$ states above). A physical mass state is some linear combination

$$\cos \theta \bar{b}b + \sin \theta \bar{c}c, \quad (3.9)$$

where θ is the mixing angle. In this case, however, $\sin \theta$ is suppressed by a factor on the order of magnitude of

$$\frac{\Lambda_{\text{QCD}}^2}{m_b^2 - m_c^2} \ll 1. \quad (3.10)$$

Indeed, the physical states are called the $J/\Psi = \bar{c}c$ and the $\Upsilon = \bar{b}b$. The J/Ψ (“so good they named it twice”) led to a Nobel prize for the discovery of charm while a resonance of the latter—the $\Upsilon(4S)$ —decays to nearly-at-rest B^0 mesons and play a key role in ‘ B -factories’ like the LHCb, BaBar, and Belle. In general states mesons composed of a heavy quark and its antiparticle are called **quarkonia**, the J/Ψ and Υ are examples of charmonium and bottomonium states, respectively. (Panda–anti-panda = *pandamonium*.)

Please take note that the important quantity for mixing is the *difference* between the valence quark masses. The overall scale of these masses do not affect mixing (though they of course set the overall mass scale).

It is also worth observing that the mass difference $|m_d - m_u|$ also splits the triplet. You should also now realize that we were being a bit sneaky in the previous section when we wrote 2×2 matrices. We explore these ideas in the following problem.

Problem 3.3. Meson eigenstates for general quark masses. Consider a model of QCD with only up and down quarks of unspecified masses m_u and m_d , that is we do not know *a priori* whether the isospin or flavor basis is dominant. Write down the meson mass matrix in the flavor basis. (**Hint:** the matrix is larger than 2×2 .)

Finally, we can milk our hydrogen analogy even further to remark on one more contribution to the meson masses⁵. As we will see below this effect is necessary to explain the mass splitting between the (pseudo) vector isotriplet ρ mesons and the pseudoscalar isotriplet pions. The only difference between these mesons is the orientation of the quark spins (they’re different representations an *actual*-spin, rather than isospin), so we might expect an analog to the spin-spin term in the energy states of hydrogen where the energy splitting is proportional to the dot product of the electron and proton spins $\mathbf{s}_e \cdot \mathbf{s}_p/m_e m_p$. By fitting for the proportionality constant by plugging in values quark masses it turns out (rather remarkably) that one can fit the masses of seven of the eight light mesons to percent accuracy. The eighth meson is the η' and is the subject of the $U(1)_A$ problem that plagued particle physics until instanton effects were properly understood [34].

⁵Here the typist has borrowed from Chapter 5.5 of *Introduction to Elementary Particles* by Griffiths.


3.6 The pseudoscalar mesons

Pseudo-Goldstone bosons. Before we discuss the mesons, we should remark that a useful way to think about these mesons is that they are the Goldstone bosons associated with the spontaneous breaking of the $SU(3)_A$ flavor symmetry by in low-energy QCD the chiral condensate. However, since these flavor symmetries are only approximate symmetries, these Goldstone bosons are not exactly massless. Additional contributions to the masses come from the ‘weak gauging’ of $SU(2)_L \times U(1)_Y$ subgroup of $SU(3)_V \times SU(3)_A \times U(1)_B$. Particle that are ‘approximate’ Goldstones are called **pseudo-Goldstone bosons** [91]. The idea that strongly-coupled systems can be described in terms of weakly-coupled degrees of freedom is *en vogue* in model building today and should elicit the AdS/CFT correspondence. It should be no surprise that one particularly interesting avenue along these lines is that the Higgs might be a pseudo-Goldstone boson; see [24] for a recent review.

Now that we have a handle to intuitively understand what determines the physical meson states (i.e. which quark pairs will mix) we can familiarize ourselves with the spectrum of low-lying mesons. Heavier mesons are resonances that are excitations of internal degrees of freedom. For now let’s look at the lightest mesons, the pseudoscalars ($J^P = 1^-$); in the next section we’ll consider the vectors, but these are a bit heavier by the spin-spin energy of the quarks. We know, for example, that the u and d have a very good approximate isospin symmetry because $m_d - m_u \ll \Lambda_{\text{QCD}}$. From the usual addition of angular momentum in quantum mechanics, we know that mesons made out of the u and d can be decomposed into an iso-singlet and an iso-triplet. The singlet is the η ($m_\eta = 550$ MeV) and the triplet is the π ($m_\pi \sim 135$ MeV). Electromagnetism provides a higher-order splitting that distinguishes the neutral pion π^0 and the charged pions π^\pm .

Now let us consider adding the s quark to enlarge isospin to an $SU(3)$ symmetry. A bit of terminology: usually this symmetry is called $SU(3)$ *flavor*, which might be slightly confusing since then the phrase ‘flavor basis’ is somewhat ambiguous (definite quark content or definite $SU(3)$ representation?). Instead in this document we will use the phrase $SU(3)$ isospin, which is a bit of an oxymoron since isospin should be $SU(2)$, but at least makes it clear what basis is relevant. We know that since m_s is quite a bit heavier than $m_{u,d}$ that this is slightly dubious, but the difference is still less than Λ_{QCD} so let’s see how far we can push this. The tensor product of a fundamental and anti-fundamental of $SU(3)$ is $\mathbf{3} \otimes \bar{\mathbf{3}} = \mathbf{1} \oplus \mathbf{8}$, where the $\mathbf{8}$ is the adjoint.

Do we find this structure in the meson spectrum? Yes! The singlet is the η' ($m_{\eta'} = 960$ MeV) and the pseudoscalar octet composed of the pions, the η , and the kaons ($m_K \sim 490$ MeV). We can see that the pions of the $SU(2)$ *pretty-good* iso-triplet and the kaons which flesh out the *just-okay* $SU(3)$ octet indeed has a mass difference on the order of the s mass ($m_s \sim 150$ MeV), which breaks the $SU(3)$ isospin/flavor symmetry. To provide a graphical mnemonic, the weight diagram for the pseudoscalar nonet (octet + singlet) is:



As we saw with the pions, the whole point of isospin symmetry here is that group theory is able to tell us how states of different quark content ought to mix. In the $SU(2)$ case we saw that the interesting mixing occurred in the $I_3 = 0$ states ($\pi^0, \eta \propto u\bar{u} \pm d\bar{d}$). Similarly, the interesting mixing occurs in the $I_3 = 0$ states of $SU(3)$ where group theory predicts that the η (part of the octet) and η' (the singlet) have the content

$$\eta = \frac{1}{\sqrt{6}} (u\bar{u} + d\bar{d} - 2s\bar{s}) \quad (3.11)$$

$$\eta' = \frac{1}{\sqrt{3}} (u\bar{u} + d\bar{d} + s\bar{s}). \quad (3.12)$$

This is, at least, what we expect from $SU(3)$ isospin. It turns out the mass difference $|m_s - m_{u,d}|$ leads to a mixing angle that gives only a 10% correction and we are indeed justified in working in the isospin basis.


For completeness, we mention that using chiral perturbation theory (i.e. the pseudoscalar octet's ‘true identity’ as pseudo-Goldstone bosons) we may represent these fields in an $SU(3)$ adjoint matrix as

$$\begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta & \pi^+ & K^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta & K^0 \\ K^- & \bar{K}^0 & -\frac{2}{\sqrt{6}}\eta \end{pmatrix}. \quad (3.13)$$

3.7 The vector mesons

The vector mesons $J^P = 1^-$, however, are a bit of a different story⁶. We can draw the $\mathbf{8} \oplus \mathbf{1}$ weight diagram to compare with the pseudoscalars:

⁶This section borrows from Chapter 4.4 of [20].



The ρ s ($m_\rho = 770$ MeV) form an $SU(2)$ iso-triplet, the analog of the pions. The $SU(2)$ iso-singlet is the ω ($m_\omega = 783$ MeV), which is the analog of the η . Including the strange quark and enlarging to $SU(3)$ isospin symmetry the K^* s ($m_{K^*} = 892$ MeV) are the analogs of the spin-0 kaons and the ϕ ($m_\phi = 1020$ MeV) is the I_3 analog of the η' .

Do you see the problem? No? Here's a hint.

Problem 3.4. The pseudoscalar Gell-Mann–Okubo mass formula. In the early 1960s Gell-Mann and Okubo were independently searching for predictions of the quark model on the spectrum of mesons. Write down the masses of the different particles (π , K , η) in the pseudoscalar octet \oplus singlet by summing some universal (flavor-independent) QCD-induced mass with the valence quark masses. Assume $m_u = m_d$ but leave m_s independent. Assume that the quark masses and the universal mass are unknown. **Hint:** Write the *squared* meson masses as the sum of the universal part and terms which are *linear* in the quark masses times some overall dimensionful constant. Show that

$$4m_K^2 = m_\pi^2 + 3m_\eta^2. \quad (3.14)$$

Remark: Writing out the *squared* mesons masses *not* the natural thing to do; in general one should write the *linear* meson masses as a universal part plus a linear combination quark masses. We will discuss below, however, why the pseudoscalar octet is special.

Näively we expect the vector meson states to have the analogous quark content as the pseudoscalars; that is the I_3 states ought to have the same mixing of quark flavor states so that

$$\omega \stackrel{?}{=} \frac{1}{\sqrt{6}} (u\bar{u} + d\bar{d} - 2s\bar{s}) \quad (3.15)$$

$$\phi \stackrel{?}{=} \frac{1}{\sqrt{3}} (u\bar{u} + d\bar{d} + s\bar{s}). \quad (3.16)$$

In particular, we expect the Gell-Mann–Okubo formula to hold since the additional spin-spin term in the mass is flavor-universal:

$$4m_{K^*}^2 \stackrel{?}{=} 3m_\omega^2 + m_\rho^2. \quad (3.17)$$

Plugging in the values above we get a prediction that $m_\omega = 926.5$ MeV, which is very different from the experimentally observed $m_\omega^{\text{exp.}} = 783$ MeV. What went wrong?

Ah, this seems like it's on the order of m_s . Shouldn't this just be a correction to our $SU(3)$ isospin approximation that we already expected?

No! The Gell-Mann–Okubo formula *already* treats the valence quark masses as being different and so should *already* account for the $SU(3)$ isospin breaking from $m_s \neq m_{u,d}$. We're still missing something.

Well, we remarked that the quadratic sum of meson masses should only hold for the pseudoscalar mesons. [Why?] The correct form of the Gell-Mann–Okubo formula for the vector mesons, then, should be linear in the meson masses.

Good! This is absolutely true. We'll discuss this in the next section. But meanwhile, have a revised Gell-Mann–Okubo formula for vector (and any higher spin) mesons:

$$4m_{K^*} \stackrel{?}{=} 3m_\omega + m_\rho. \quad (3.18)$$

Unfortunately this predicts $m_\omega = 1446$, which is an even *worse* fit! We're really in trouble now!

Problem 3.5. If you haven't figured you the mystery yet, *stop*. Just stop right there. Don't read beyond this problem. Go to the PDG and look up the quark content of the ω and ρ mesons.

It turns out that our assumption about the quark content of the I_3 states (3.15) and (3.16) is wrong. Given the theme of these lectures, you should already know why: these are two particles with the same quantum numbers, so that—in general—*they mix*! We can thus make the ansatz that the ω and ϕ are not actually the I_3 components of the iso-octet and iso-singlet as we had initially assumed (and that was true for the η_s). Instead, these states are assumed to be linear combinations of the *actual* I_3 components of the iso-octet and iso-singlet, which we shall call ψ_8 and ψ_1 with masses m_8 and m_1 . The quark content is given by

$$\psi_8 = \frac{1}{\sqrt{6}} (u\bar{u} + d\bar{d} - 2s\bar{s}) \quad (3.19)$$

$$\psi_1 = \frac{1}{\sqrt{3}} (u\bar{u} + d\bar{d} + s\bar{s}). \quad (3.20)$$

In other words, the mass matrix for ψ_8 and ψ_1 contains off-diagonal elements $m_{81} = m_{18}$,

$$M = \begin{pmatrix} m_8 & m_{81} \\ m_{18} & m_1 \end{pmatrix}. \quad (3.21)$$

It may not seem like this helps much since it looks like we have three free parameters (m_8 , m_1 , and the mixing angle θ) to fit two experimentally determined values m_ω and m_ϕ . However, we should remember that the Gell-Mann–Okubo formula *must work* for the ψ_8 and ψ_1 states. Thus (3.18) should be modified to

$$4m_{K^*} = 3m_8 + m_\rho. \quad (3.22)$$

where $m_8 = 1446$ MeV is the mass of the $I_3 = 0$ member of the octet. One might also wonder about the off-diagonal elements of the $\psi_{8,1}$ mass matrix; these masses are also fixed by the same valence quark analysis—see the following problem. (In what follows we will always be careful to take differences of meson masses so that the universal mass contributions from angular momentum and the QCD potential can be ignored.)

Problem 3.6. The Schwinger meson mass relation. Shortly after Gell-Mann and Okubo, Schwinger published his own mass sum rule for mesons for $SU(3)$ nonets (octet+singlet) that is robust against mixing [83]. In this problem we will prove Schwinger's relation and see that it works for the vector mesons. First determine the elements of the mass matrix (3.21) in terms of the measurable masses of the other octet mesons (m_K^* , m_ρ). Note that these mesons don't mix (why?) and so their masses are good parameters. Next use the invariance of the trace and determinant to prove the following relations

$$m_\omega + m_\phi = 2m_{K^*} \quad (3.23)$$

$$m_\omega m_\phi = \frac{1}{9} [(4m_{K^*} - m_\rho)(2m_{K^*} + m_\rho) - 8(m_\rho - m_{K^*})^2] \quad (3.24)$$

Finally, combine these equations to write down the **Schwinger meson mass relation**,

$$-(m_8 - m_\omega)(m_8 - m_\phi) = \frac{8}{9}(m_\rho - m_{K^*})^2, \quad (3.25)$$

where m_8 is expressed in terms of measurable masses by the Gell-Mann–Okubo relation (3.22). Check that the vectors satisfy this relation. (Please do this question *especially* if you are confused about what's going on!)

The observation that the vector meson masses obey the Schwinger relation tells us that the model of $SU(3)$ broken by m_s is consistent and *does* explain the ω and ϕ masses once one accounts for mixing. The next thing we need to check is to find out how much the ψ_8 and ψ_1 states actually mix. If you did the previous problem, then you already have the explicit form of the mass matrix elements $m_8, m_1, m_{18} = m_{81}$ in terms of meson masses that one can look up in the PDG. Thus it shouldn't be any trouble for you to do the following problems (we're back to diagonalizing 2×2 matrices!):

Problem 3.7. The ω - ϕ mixing angle, part I. Parameterize the mixing between the $SU(3)$ flavor eigenstates ψ_8 and ψ_1 into the mass eigenstates ω and ϕ as follows:

$$\begin{pmatrix} \phi \\ \omega \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \psi_8 \\ \psi_1 \end{pmatrix}, \quad (3.26)$$

where the flavor eigenstates have quark content (3.19-3.20). In this problem we will show several formulae for the mixing angle. For $\tan^2 \theta$ show that

$$\tan^2 \theta = -\frac{m_8 - m_\phi}{m_8 - m_\omega} = -\frac{m_1 - m_\omega}{m_1 - m_\phi} = \frac{m_8 - m_\phi}{m_1 - m_\phi} = \frac{m_1 - m_\omega}{m_8 - m_\omega}. \quad (3.27)$$

While those formulae do not give the overall sign of θ , prove the following $\tan \theta$ formulae which do:

$$\tan \theta = \frac{-m_{81}}{m_8 - m_\omega} = \frac{m_{81}}{m_1 - m_\phi} = \frac{m_8 - m_\phi}{m_{81}} = \frac{m_1 - m_\omega}{-m_{81}}. \quad (3.28)$$

Finally, for good measure, prove for $\cos 2\theta$ and $\sin 2\theta$,

$$\cos 2\theta = \frac{m_8 - m_1}{m_\phi - m_\omega} \quad \sin 2\theta = \frac{2m_{81}}{m_\omega - m_\phi}. \quad (3.29)$$

Problem 3.8. The ω - ϕ mixing angle, part II. Now that we have the formulae above, rewrite them in terms of physical meson masses so that it is easy to go to the PDG and plug in the appropriate numbers. In particular, prove the relations given in the “Quark Model” review of the PDG:

$$\tan^2 \theta = -\frac{4m_{K^*} - m_\rho - 3m_\phi}{4m_{K^*} - m_\rho - 3m_\omega} \quad (3.30)$$

$$\tan \theta = \frac{4m_{K^*} - m_\rho - 3m_\phi}{2\sqrt{2}(m_\rho - m_{K^*})}. \quad (3.31)$$

Show that the expression for $\tan 2\theta$ obtained from (3.29) is independent of physical measurements and remark on the significance of this value.

These are very important relations and we end up with a mixing angle of $\theta \approx 37^\circ$. Note that the mixing angle is determined by the ratio of mass differences! (Remind yourself that $m_{81} \propto (m_s - m_u)$!) So the reason why the Gell-Mann–Okubo relation failed for the vector mesons is that there is an appreciable mixing between the ψ_8 and ψ_1 . We say ‘appreciable’ because $37^\circ \gg 0^\circ$, but how much mixing is actually contained in 37° ? To ask the question in a different way: is there a way that we could have expected this amount of mixing? Shouldn’t we have expected either zero or 45° ? This would be a good homework problem, but it’s so important (and simple) that we’ll do it together.

Example 3.8. Ideal mixing. We assume that $SU(3)$ is broken only by the difference ($m_s - m_u$) (where $m_d = m_u$ by isospin). Thus it is *intuitively* clear that the physical states in the theory should appear in a way that reflects this breaking. This is just like saying that in the fine structure of Hydrogen one breaks along total angular momentum so that one obtains the singlet and triplet states; then in the hyperfine structure one breaks along the direction of orbital angular momentum so that one obtains the $|1, \pm 1\rangle, |1, \pm 0\rangle$ states. We can use (3.29) to determine the expected angle θ . A quick way to do this is to write down the valence quark masses in each component. You know from the previous problems that

$$m_8 = \frac{2}{3}(m_u + 2m_s) \quad (3.32)$$

$$m_1 = \frac{2}{3}(2m_u + m_s) \quad (3.33)$$

$$m_{81} = \frac{2\sqrt{2}}{3}(m_u - m_s). \quad (3.34)$$

Plugging these in we obtain the **ideal mixing** angle

$$\tan 2\theta_{\text{ideal}} = 2\sqrt{2}, \quad (3.35)$$

which is approximately 35.5° .

Aha! So we see that the value of $\theta = 37^\circ$ that we found indeed is very close to the prediction of **ideal mixing**. More importantly, by now you should *physically* understand what ideal mixing signifies: the $s\bar{s}$ component of the $I_3 = 0$ mesons should separate from the $u\bar{u}$ and $d\bar{d}$ components.

The physical states should mirror the breaking pattern, just like the fine structure of hydrogen. Indeed, writing

$$\omega = \cos \theta \psi_8 - \sin \theta \psi_1 \quad (3.36)$$

$$\phi = \sin \theta \psi_8 + \cos \theta \psi_1 \quad (3.37)$$

and remembering the iso-state quark content

$$\psi_8 = \frac{1}{\sqrt{6}} (u\bar{u} + d\bar{d} - 2s\bar{s}) \quad (3.38)$$

$$\psi_1 = \frac{1}{\sqrt{3}} (u\bar{u} + d\bar{d} + s\bar{s}), \quad (3.39)$$

one finds that for ideal mixing the ϕ is a *pure* $s\bar{s}$ state while the ω contains *no* strangeness,

$$\omega = \frac{1}{\sqrt{2}} (u\bar{u} + d\bar{d}) \quad \phi = s\bar{s}. \quad (3.40)$$

In this sense it is *maximally* sensitive to how much the s quark mass breaks $SU(3)$ isospin symmetry and so we end up with the ω and ϕ being much heavier than their cousins.

So now we've completely turned around the 'problem' that we originally proposed at the beginning of this section. There's nothing odd about the vector mesons—they behave *precisely* as we would expect given the way that $SU(3)$ is broken. Indeed, one finds that the $J = 2$ tensor mesons have a mixing angle that is also very close to ideal mixing [55]. The *real* problem is that the *pseudoscalars* are misbehaving because they somehow conspire to give zero mixing when we expect ideal mixing! It is a deep principle in physics that nothing is ever zero 'by accident.'

3.8 Why are the pseudoscalar and vector octets so different?

Now that we've identified the pseudoscalars as the trouble-makers, why should they behave so differently from the higher-spin mesons? We already mentioned part of the answer: the pseudoscalars are secretly under-cover Goldstone bosons of $SU(3)_A$, i.e. they are pseudo-Goldstone bosons which obtain masses according to chiral perturbation theory. Unfortunately the formalism for chiral perturbation theory is beyond the scope of this course⁷. However, the crux of the matter is that the breaking of $SU(3)$ by quark masses gives relations of the form (using soft pion techniques)

$$f_\pi^2 m_\pi^2 = \frac{m_u + m_d}{2} \langle 0 | u\bar{u} + d\bar{d} | 0 \rangle \quad (3.41)$$

$$f_K^2 m_K^2 = \frac{m_u + m_d}{2} \langle 0 | u\bar{u} + s\bar{s} | 0 \rangle \quad (3.42)$$

$$f_\eta^2 m_\eta^2 = \frac{m_u + m_d}{6} \langle 0 | u\bar{u} + d\bar{d} | 0 \rangle + \frac{4m_s}{3} \langle 0 | s\bar{s} | 0 \rangle. \quad (3.43)$$

⁷For some reviews see [71, 72, 29, 81] or more recent reviews on modern versions of this idea [24, 82]. Additionally, older-style textbooks such as Cheng & Li [20] or Donoghue et al. [58] give very nice introductions.

We can assume that all of the expectation values are at the same QCD condensate scale, $\langle 0|q\bar{q}|0\rangle = \Lambda_{\text{QCD}}^3$ and further that the decay constants are all equivalent $f_\pi = f_K = f_\eta = f$. What we are left with is a relation for the *quadratic* meson masses that is *linear* in the symmetry-breaking terms, the quark masses. This explains why the original Gell-Mann–Okubo formula for the pseudoscalar mesons (3.14) is written with squared meson masses while the formulae for the vectors (3.18) is linear in the meson masses.

So maybe this is the source of the problem. Unfortunately, one can go ahead and try to use all of the mixing angle formulae we derived in Problems 3.6 and 3.8. You can work this out as an exercise, but it turns out that one doesn't get consistent answers! Another way to say this is that the Schwinger mass relation (3.25) doesn't hold. Clearly something is still wrong. Before moving on, though, let us remark that an independent measurement of the pseudoscalar mixing angle can be obtained by looking at various decay rates ('partial widths') as explained in the PDG's "Quark model" review; the angle is found to be between -10° and -20° .

So what is it that is causing the pseudoscalars to mix in such an unexpected and seemingly inconsistent way? It turns out that we've already met this before when we were counting parameters: it is our old friend the $U(1)_A$ axial anomaly. The fact that $U(1)_A$ is broken means that ψ_1 gets an *additional* mass contribution (in principle from the $F\tilde{F}$ term). See, for example, the discussions in the book by Donoghue et al. [58] or the older article by Isgur [55]. The reviews by Feldmann [32] and Gasser and Leutwyler [35] are particularly thorough.

Problem 3.9. A physical interpretation of $\theta = -10^\circ$. Change to the basis where the mesons have quark content following the $SU(3)$ breaking pattern preferred by the symmetry-breaking $(m_s - m_u)$ parameter,

$$\psi_x = s\bar{s} \tag{3.44}$$

$$\psi_y = \frac{1}{\sqrt{2}}(u\bar{u} + d\bar{d}). \tag{3.45}$$

Write the deviation from the ideal mixing angle θ_{ideal} as ϕ so that $\theta = \theta_{\text{ideal}} + \phi$. Show that

$$\tan 2\phi = \frac{m_{xy}}{m_x^2 - m_y^2}. \tag{3.46}$$

Consider the limits $m_{xy} \ll m_x^2 - m_y^2$ and $m_{xy} \gg m_x^2 - m_y^2$ and use these to explain the significance of the mixing angles $\theta = \theta_{\text{ideal}}$ and $\theta = \theta_{\text{ideal}} - 45^\circ$. Comment on the stability of these values in the aforementioned limits.

Problem 3.10. Mixing of the $I_3 \neq 0$ states. Explain why the K^\pm and π^\pm do not mix appreciably. It is not sufficient to invoke $SU(3)$ flavor symmetry since we know that $m_s \gg m_{u,d}$ so that this is not a very good symmetry. **Hint:** Assume that $SU(3)$ is only broken by the quark masses.

3.9 Hadron names

"You can know the name of a bird in all the languages of the world, but when you're finished, you'll know absolutely nothing whatever about the bird. So let's look at the

bird and see what it's doing—that's what counts. I learned very early the difference between knowing the name of something and knowing something.” –Richard Feynman⁸

While Feynman makes a point that is well-known by any physicist who has talked to someone from a ‘taxonomic’ field (like botany), it is important to be familiar with the names of common hadrons. At the very least one should be motivated that experimentalists are all very familiar with different hadrons (since these are the things which fly through the detector), so it's worth knowing what they're talking about when they say things such as, “A hyperon can punch through and mimic a muon.”

We've already talked about the pseudoscalar mesons which have $J^{PC} = 0^{-+}$. We have the π , η , η' , K , D , B . The first three are called ‘unflavored’ by the PDG. By this we really mean that it contains no *heavy* flavor, i.e. no net strangeness, charm, or beauty. The η s contain some admixture of $s\bar{s}$, but this has no net strangeness quantum number. The K , D , and B are all code for flavored mesons. The kaons (K) all have net strangeness, the D s have net charm, and the B s have net beauty.

We've also already met several vector mesons which have $J^{PC} = 1^{--}$. The iso-triplets are the ρ s and we're now very familiar with the ω and ϕ as our famous states which mix isospin representations. The flavored vector mesons are indicated by stars relative to the pseudoscalars: K^* , D^* , B^* .

There are other mesons. The $J^{PC} = 0^{++}$ mesons are formed by quark configurations with orbital angular momentum $\ell = 1$ but fundamental spin $s = 1$ such that the net spin is zero. These have less exotic names, a_0 , a_1 , a_2 , ... What about the 1^{+-} states? These are formed from states with $\ell = 1$ and $s = 0$. Can you guess their names? b_0 , b_1 , ... If you guessed that correctly then you have at least a minimal amount of ‘physics intuition.’

For completeness we should also discuss the names of some popular baryons. The *most* popular baryons are well-known by everybody: the proton and neutron. For most of this course, however, we won't spend much time thinking about the other baryons. It is worth briefly reviewing why this is the case.

Why we don't bother with baryons. We understand how we get mesons: we smash particles together at a collider and the hard scattering processes tend to create objects like $q\bar{q}$ pairs. Confinement tells us that we must pull out a $q\bar{q}$ pair from the vacuum in order to form color-neutral final-states. If we wanted to create final-state baryons, on the other hand, we'd need to pull out *two* $q\bar{q}$ pairs from the vacuum. An undergraduate might then say that this must be a higher-order process since we're borrowing ‘more’ energy from the QCD vacuum, and thus baryon production is rare.

This is *wrong!* QCD is strongly interacting, so the cost of pulling a $q\bar{q}$ pair from the vacuum is order one. Indeed, we know that for high-energy collisions we end up with jets that each have dozens of $q\bar{q}$ pairs. What *does* end up costing you something is combinatorics. When one pulls out two $q\bar{q}$ pairs from the vacuum, it's far more likely (counting color multiplicities) that they will produce three mesons than three baryons. Thus the suppression is not that making

⁸“What is Science?”, presented at the fifteenth annual meeting of the National Science Teachers Association, in New York City (1966) published in *The Physics Teacher* Vol. 7, Issue 6 (1969).

enough pairs to produce baryons is less likely, but rather that once you do this it is still more likely to form mesons.

Still, let's go over the baryon names. These are half-integer spin guys. For three light quarks (u, d combinations), i.e. isospin $I = 1/2$, these are called nucleons N . This is just a fancy name for 'proton or neutron.' For $I = 3/2$ we have the Δ s ($m_\Delta = 1232$ MeV). These come in charges $q = +2, +1, 0, -1$.

For baryons with two light quarks we have the Λ s ($I=0$) and Σ s ($I=1$). The isospin comes only from the light quarks. These particles are named according to their heavy quark. No index implies an s , while a subscript c or b indicates charmed or beautiful baryons, for example Λ_c .

Baryons with just one light quark ($I = 1/2$) are called Ξ , which is too difficult for particle physicists to pronounce consistently so we tend to call them 'cascades.' These are also named according to their heavy quarks with s implied for no index. Thus we would name a baryon with one light quark, a strange quark, and a charm quark the Ξ_c . How do we know which light quark is active? By the charge; for example the Ξ^+ must be a uus combination. (You might say that this depends on electromagnetism to distinguish states; this is exactly correct, since without isospin breaking the question of 'which light quark' is ill-posed.)

Finally, what about the baryons with no light quarks? There is one famous one, the Ω composed of three strange quarks. This is the particle that Ne'eman and Gell-mann predicted on the basis of the decouplet of $SU(3)$. They wrote down its charge, quark content, and mass and the Ω was discovered a few years later.

We will not have many occasions to discuss baryons with multiple heavy quarks. The first such baryon, the Ξ_{cc} was only tentatively observed in the past decade. Recall that while it is easy to produce light quarks from the vacuum, heavy quarks with $m \gg \Lambda_{\text{QCD}}$ are a whole different story. Such production is suppressed by $\exp(-m/\Lambda_{\text{QCD}})$. Thus, for a Ξ_{cb} one can hope to produce the valence b quark from a high-energy parton-parton collision but the accompanying c must be produced from the vacuum at the cost of exponential suppression.

Well, wait a second. The LHC has a really large center-of-mass energy. Can't it produce all of the heavy quarks we need?

Do not be confused by the fact that the LHC is a multi-TeV collider! We would produce these heavy hadrons just as well at HERA as we would at the LHC. The key point (and a deep principle in physics) is *decoupling*. The *hard scattering* event which occurs at high scales is likely to produce any of the quarks since $m_q \ll \text{TeV}$. These quarks propagate away from one another for a distance scale on the order of $1/\Lambda_{\text{QCD}}$ and it is *only* at that distance scale that hadronization from QCD really becomes active. This is a fact that is very well known by anyone who works with Monte Carlo generators: you use MadGraph for the hard scattering, but Pythia for hadronization and showering. You can separate these steps because they occur at two totally different scales and are governed by different physics.

Problem 3.11. An improved Gell-Mann–Okubo formula for the baryon octet. The masses of the baryons in the $SU(3)$ iso-octet do not fit the naïve isospin predictions. Apparently there is a considerable contribution from the spin-spin interactions. Fortunately, the magnetic

moments of the isospin octet baryons are relatively well-measured. Use this to write down predictions for the baryon masses taking into account the constituent quark masses *and* the hyperfine splitting from spin-spin interactions.

4 The flavor structure of the Standard Model

Now that we’ve waded a bit in the hadronic ‘brown muck’,⁹ let us pop our heads back up into the clouds of theory. In fact, let’s go to our beloved friend, the Standard Model. From our parameter counting in Section 2.5 we know that the Standard Model has four physical flavor parameters. In this section we’ll see how these show up in the Standard Model. The experimental goal of flavor physics is to measure these four parameters in as many ways as possible to check for consistency and hope for a signal of new physics.

4.1 The CKM matrix

We should already know that we will be interested in the **CKM matrix**, named after Cabbibo, Kobayashi, and Maskawa. Cabbibo, unfortunately, passed away in 2010 before he had a chance to compete for a Nobel prize after Kobayashi and Maskawa each took part of the 2008 prize. We should hold a special regard for Cabbibo since he was the one to diagonalize the relevant 2×2 matrix (between the d and s quarks), and we know that clever physicists know how to identify the *right* 2×2 matrix to diagonalize.

At the heart of the matter the CKM matrix comes from a basis rotation coming from the fact that one cannot diagonalize all of the matrices in the Standard Model. The CKM matrix takes you from one basis to another and is the source of flavor mixing. The particular flavor matrices that cannot be simultaneously diagonalized are the W -boson couplings, the masses of the d -type quarks, and the masses of the u -type quarks. At most one can diagonalize two of these and one has to choose which—if any—should be diagonal. We will choose to diagonalize the masses and leave the W coupling being non-diagonal. It is important to emphasize over and over again that this is a *choice* that we are free to make. In neutrino physics, for example, we choose the neutrino masses to be not diagonal. In principle we could have even chosen a different basis where none of the matrices are diagonal.

From our principle of writing the most general renormalizable Lagrangian, the only non-trivial (i.e. not from the kinetic terms) quark interactions are Yukawa couplings (2.4),

$$\mathcal{L}_{\text{Yuk}} = y_{ij}^d \bar{Q}^i \phi D^j + y_{ij}^u \bar{Q}^j \tilde{\phi} U^i + (\text{lepton term}) + \text{h.c.} \quad (4.1)$$

Upon electroweak symmetry breaking the Higgs vacuum expectation value (‘vev’) turns these into mass matrices. To get these masses we use our unphysical degrees of freedom (explicitly broken global symmetries of the kinetic term) to rotate the Higgs vev to be real and in the down component of ϕ (and so the up component of $\tilde{\phi}$). Writing the doublet fields out explicitly,

$$Q_L^i = \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix}, \quad (4.2)$$

⁹A term coined by Howard Georgi, [37].

we see that the Higgs vev gives us mass terms of the form

$$\mathcal{L}_{\text{mass},q} = m_{ij}^d \bar{d}_L^i d_R^j + m_{ij}^u \bar{u}_L^i u_R^j, \quad (4.3)$$

where the mass matrices are related to the Yukawas and the Higgs vev $\langle \phi \rangle_2 = v/\sqrt{2}$ by

$$m_{ij}^q = \frac{v}{\sqrt{2}} y_{ij}^q. \quad (4.4)$$

In general these m^q s are not diagonal. So far there are many unphysical parameters. We now move to the mass basis. Again we note that there's nothing *holy* about the mass basis! This is a *choice* that we are free to make, we could have chosen a different basis without affecting the physics.

Is the mass basis special? When we learn QFT we always start out in the mass basis. We should remember, however, our roots in relativistic quantum mechanics where the ‘mass basis’ is just the energy basis, which is just another basis that one is free to choose for one’s states. In QM we’re used to working with states that aren’t in the energy basis, so we should be willing to transfer this intuition to QFT. *However*, there is an important counter point about why the mass basis is ‘special’ in QFT. In the mass basis one can go ahead and completely solve the quadratic part of the path integral without any of the problems associated with the infinite-dimensionality of QFT. This means that it is a natural basis to expand about in perturbation theory. Of course, we are free *not* to use the mass basis and in the limit where the masses are perturbative relative to the energy scales of a process one can work with two-point mass insertion Feynman rules. This is useful, for example, when calculating higher-dimensional penguins [25]. More generally the mass-insertion approximation is often used in the SUSY flavor sector, though there exist computer codes which make this unnecessary [77]. Anyway, the *point* is that it is perfectly valid to do QFT in bases other than the mass basis, though there are good reasons why the mass basis is a default in calculations.

Let’s diagonalize m . How do we diagonalize an arbitrary complex matrix? We use a bi-unitary transformation, i.e. a unitary transformation acting on the left and another unitary transformation acting on the right. By comparison, Hermitian matrices can be diagonalized with the same rotation acting from the left and right. In general,

$$\hat{m}_{ij}^q = V_{L ik}^q m_{k\ell}^q \left(V_R^{q\dagger} \right)_{\ell j}. \quad (4.5)$$

We use the notation that a matrix with a hat \hat{m} is diagonal. In principle these L, R indices have nothing to do with the L, R indices, but one can see from (4.3) that the i index is associated with the left-chiral quark while the j index is associated with the right-chiral quark. Thus rotate the left-chiral and right-chiral fields by V_L and V_R accordingly:

$$q_L^i = (V_L^q)_{ij} q_L^j \quad (4.6)$$

$$q_R^i = (V_R^q)_{ij} q_R^j. \quad (4.7)$$

On the left-hand side we've written the interaction-basis field and on the right-hand side we've written the corresponding linear combination of mass-basis fields. From here on we will assume the mass basis and drop the prime. (Don't you hate it when textbooks do this?)

How does our Lagrangian look in this mass basis? The Yukawa interactions are now diagonal since the mass matrices are proportional to the Yukawa matrices, but now the coupling of the W contains off-diagonal terms. Going from the interaction to the mass basis, we find

$$\mathcal{L}_{Wqq} = \frac{g}{\sqrt{2}} \bar{u}_L i \gamma_\mu d_L W^\mu \rightarrow \frac{g}{\sqrt{2}} \bar{u}_L i \gamma_\mu (V_{uL} V_{dL}^\dagger) d_L W^\mu, \quad (4.8)$$

where the factor of $1/\sqrt{2}$ came from the normalization of the W^\pm states relative to the $W^{1,2}$ states. We can now identify the famous **CKM** matrix,

$$V \equiv V_{uL} V_{dL}^\dagger. \quad (4.9)$$

Let's discuss some properties of the CKM matrix.

- The CKM matrix is unitary. $V^\dagger V = V^\dagger V = \mathbb{1}$. We will see, however, that it is often convenient to parameterize it in a non-unitary way that captures all of the relevant physics.
- The number of physical parameters is always four (and always three mixing angles and a complex phase), but there are many convenient parameterizations. Choosing a convenient parameterization is independent of physics.
- The index structure of the CKM matrix is

$$V = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (4.10)$$

We have *chosen* to order the quarks by increasing masses. We have also *chosen* to name the mass eigenstates by their dominant flavor. This is *not* what we do with neutrino mixing, for example, where we refer to mass states ν^1 and interaction ('flavor') states ν^e .

Before moving on, let's re-emphasize once again for the millionth time that we are making an arbitrary (but convenient) choice about how we choose the physical parameters of V . We are free to choose any parameterization and the physics must be parameterization-invariant. Any quantity which is parameter dependent *cannot* be physical.

Example 4.1. V_{ub} is parameter dependent, but $|V_{ub}|$ is parameter independent. Of the two, only the latter can be a physical observable.

In this sense parameterization-invariance is like gauge invariance. Anything physical must be parameterization invariant, and any 'physical' quantity which you calculate that ends up being parameterization-dependent must be a mistake.

4.2 Parameterizations of the CKM matrix

Our parameter counting taught us to expect that the quark sector of Standard Model contains three mixing angles and one complex phase. Recall our explanation for dividing the parameters this way came from comparing the number of parameters in an orthogonal matrix versus a unitary matrix; review Problem 2.1 and the discussion above it for a refresher. Given that the whole of the flavor structure in the Standard Model boils down to these four numbers (plus the three Dirac masses for the quarks), the general unitary matrix we wrote in (4.10) must contain many unphysical parameters since element V_{ij} is generally a complex number.

Problem 4.1. Simplest unitary parameterization of the CKM. Given that we know that the physical content of the CKM matrix must boil down to three mixing angles and one phase, how minimally can one encode these ingredients into a single unitary matrix? In particular, How many of the elements of the CKM matrix can be made real and how many must be complex? *ANSWER: The most general CKM matrix can be parameterized by at most 4 real parameters that cannot all be in the same row. This means we must have a minimum of five complex elements.*

In light of the above mathematical result, it is customary to choose V_{ud}, V_{us}, V_{cb} , and V_{tb} to be purely real. The remaining elements are complex. One standard parameterization of the mixing angles is: $\theta_{12}, \theta_{13}, \theta_{23}$. In the limit of two generations (e.g. d and s), θ_{12} is the usual Cabibbo angle. The two indices tell us which plane we're rotating about. We will call the phase δ_{KM} after Kobayashi and Maskawa. In terms of these parameters, the CKM matrix takes the form (writing $c_{12} = \cos \theta_{12}$, etc.)

$$V_{\text{CKM}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}. \quad (4.11)$$

This is nice, but it's a little unwieldy. It would be nice to have an approximation that captures the essential physics in a way that makes it more transparent. This is a bit of a shift in paradigm, so let's take a moment to discuss some philosophy. The reason why do we usually use approximations in physics is because we often can't solve things exactly. However, here we have an *exact* parameterization of the CKM matrix, but we want to move away from it to an approximation. This sounds extremely stupid. It will seem even more stupid when you realize that the approximation that we'll make is not even unitary—we *lose* one of the fundamental properties of the matrix! However, part of being a physicist means knowing what you can neglect. Yuval says it this way:

I tell the undergrads that it is not by mistake that [our department] is in [the college of] Arts and Sciences. All of physics is about the art of making the right approximation. Making the right approximation can teach you a lot.

The first person to make such an approximation was Wolfenstein (you might also be familiar with him from the MSW effect in neutrino physics). His key insight was that the orders of magnitude of the CKM matrix seemed to follow a particular pattern:

$$|V| \sim \begin{pmatrix} 1 & \lambda & \lambda^3 \\ \lambda & 1 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}, \quad (4.12)$$

where $\lambda \approx 0.2$. Motivated by this, he defined four *different* parameters to describe the physical content of the CKM matrix: λ, A, ρ, η . One takes λ to be a small parameter worthy of expanding about and the others are $\mathcal{O}(1)$. These parameters are defined relative to the standard parameterization by

$$s_{12} = \lambda = \frac{|V_{us}|}{\sqrt{|V_{ud}|^2 + |V_{us}|^2}} \quad (4.13)$$

$$s_{13} = A\lambda^2 = \lambda \left| \frac{V_{cb}}{V_{us}} \right| \quad (4.14)$$

and finally (this is the somewhat ugly one)

$$s_{13}\delta^{i\delta} = A\lambda^3(\rho + i\eta) = V_{ub}^* = \frac{A\lambda^3(\bar{\rho} + i\bar{\eta})\sqrt{1 - A^2\lambda^4}}{\sqrt{1 - \lambda^2[1 - A^2\lambda^4(\bar{\rho} + i\bar{\eta})]}}, \quad (4.15)$$

where $\bar{\rho} + i\bar{\eta} = -(V_{ud}V_{ub}^*)/(V_{cd}V_{cb}^*)$ is invariant under how one chooses to define the CKM phase. The bars mean that we should treat these as the ‘exact’ values and $\rho + i\eta$ are expansions of this value about small λ . One is free to plug in these relations into (4.11) to get a perfectly unitary and even uglier representation of the CKM matrix. The beauty of the **Wolfenstein parameterization**, however, is that we may use it to write the CKM matrix as a Taylor expansion in λ , so that

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4). \quad (4.16)$$

Looking at the factors of λ we can see that the diagonals are order one, and we get the structure that (4.12) Wolfenstein observed. Now look at the upper left 2×2 . Do you recognize this? (“If you don’t see it, then there’s something wrong with your physics education!”) The elements are just the expansion for sine and cosine. In fact, this is just the 2×2 Cabibbo mixing matrix! We see that to first order in λ the first two generations don’t know about the third. There’s one more feature: complex numbers only show up in the 3-1 mixing element. Let us make a brief detour to highlight the physical significance of this complex number.

4.3 CP violation

Since the punchline of our discussion will be CP violation, let’s remind ourselves of some properties of discrete symmetries¹⁰.

Of course we’re talking about the usual suspects: C, P, CP, and CPT. We know that any local Lorentz-invariant field theory preserves CPT. So far everything we’ve observed agrees with CPT, so we will assume that this is respected by nature. (Note that there *are* interesting cases where effective CPT-violating theories are useful.) What about the other guys? None of them *needs* to be conserved by a theory. Experimentally we know that QCD and QED each conserve both *C* and *P* separately, and so also conserve *CP*. The weak interaction is a different story.

¹⁰Yuval bet 25 cents that he’s taught first semester QFT enough to know the exact chapter in Peskin & Schroder where this is discussed: 3.6. He’s correct, but a better reference for this is Ticiati [89].

Electroweak theory is, by construction, parity violating. This is because it is a chiral theory, meaning it treats left- and right-handed fields differently. Clearly once you write down such a theory interchanging left and right brings you to a different theory. Any chiral theory tautologically violates parity. What about charge conjugation? Applying C to a left-handed field χ indeed transforms it into a right-handed field, $\bar{\psi}$. The particular transformation isn't important to us here, but is explained thoroughly in, e.g. [89, 47]. Thus a chiral theory violates *both* P and C . To repeat things over again but with more refined language, we may say

QED and QCD are vectorial and so preserve P and C separately, but electroweak theory is chiral and so violates both P and C .

What about CP ? Since both P and C transformations take left- and right-handed fields into one another, $\chi \leftrightarrow \bar{\psi}$, at first glance CP isn't *necessarily* violated. This is indeed true at a model-building level. Following our prescription for designing a model, we may specify the key ingredients of a model (gauge group, representations, breaking pattern) and then write down the most general renormalizable Lagrangian. This is a model, but it's not actually a “predictive description of nature” until its physical parameters are determined explicitly. The Standard *Model* as written in terms of undetermined parameters may or may not violate CP . It is only once the parameters of the Standard Model were measured¹¹ that we found that indeed, the values of some of the parameters violates CP .

How does a parameter violate CP ? It turns out that CP -violation reduces to the presence of a non-vanishing physical phase. (This is precisely what we called δ_{KM} above!) You should read these as synonyms:

“physical non-zero phase” = “this theory is CP violating.”

We won't go into details here, though we've already done most of the work in our parameter counting. For details see [14, 47]. Instead, we'll provide a hand-wavy argument an focus on the Yukawa terms. Let us consider the u -type Yukawa term and—for this one time only—explicitly write out its hermitian conjugate,

$$\mathcal{L}_{\text{Yuk},u} = y_{ij} \chi_Q^i \phi \psi_U^j + y_{ij}^* \bar{\chi}_Q^i \phi^\dagger \bar{\psi}_U^j, \quad (4.17)$$

where we've explicitly written out Weyl spinors to emphasize that our fields are chiral. What happens when we apply CP to this? The key point is that fermion bilinear $\bar{\Psi}\Psi$ in Dirac notation¹²) is invariant under CP . In terms of Weyl spinors,

$$\chi_Q^i \phi \psi_U^j \xrightarrow{CP} \bar{\chi}_Q^i \phi^\dagger \bar{\psi}_U^j \quad (4.18)$$

So that the CP conjugated Yukawa terms look like

$$(CP)\mathcal{L}_{\text{Yuk},u} = y_{ij} \bar{\chi}_Q^i \phi^\dagger \bar{\psi}_U^j + y_{ij}^* \chi_Q^i \phi \psi_U^j. \quad (4.19)$$

¹¹This is an anachronistic statement since the Standard Model was developed theoretically hand-in-hand with the experiments that probed its parameters.

¹²It is arguable whether or not the discussion is simpler using Dirac spinors. On the one hand the transformation properties are straightforward, but on the other hand one should technically also write out explicit chiral projection operators which end up not mattering. The original lecture was given in Dirac spinors, but these will be typed with Weyl spinors or else the typist's adviser would disown him in utter shame.

Note that we did *not* act on the coefficient, which is just a number. We conclude that the Lagrangian is CP-invariant *only* when $y = y^*$, i.e. when y is real. We could have of course chosen y to be real by rephasing our fields. The complete argument requires showing that one *cannot* make this phase rotation, and in fact reduces to the fact that we've already shown from parameter counting that there is a physical phase left over. Of course, the final observation is that in the Standard Model there is such a phase *and* experiments have verified that it is non-trivial. The key statement is this:

A physical complex parameter that is measured to be non-trivial implies CP violation.

So far we haven't mentioned flavor at all. Charge and parity are discrete symmetries of our field theory, in a sense coming from Lorentz symmetry. In the Standard Model, however, CP violation *always* comes with flavor violation. We could see this in our parameter counting by noting that we really *needed* to have $N \geq 3$ flavors (i.e. $N \times N$ Yukawa matrices) in order to have a physical phase.

Now let's go back to the Wolfenstein parameterization (4.16). The interesting feature we noted at the end of the last subsection was that to leading order the only complex parameters show up in the 1-3 and 3-1 elements. In other words, to very good approximation CP violation *only* occurs in interactions between the first and third generations! One can go to higher order in the expansion in λ to get CP violation in other elements, but clearly these effects will be suppressed by additional powers of λ .

4.4 The Jarlskog Invariant

Now that we're familiar with the existence of the CP-violating phase, we would like to be able to quantify it in a meaningful way—i.e. in a way which is manifestly basis-independent. What we need is some kind of invariant that identifies CP violation. Such an object exists and it is called the **Jarlskog invariant**, J [57]. It is defined by

$$\text{Im} [V_{ij}V_{kl}V_{il}^*V_{kj}^*] = J \sum_{mn} \epsilon_{ikm}\epsilon_{jln}, \quad (4.20)$$

where there is no sum on the left-hand side. In terms of our CKM parameterizations, this corresponds to

$$J = c_{12}c_{23}c_{13}^2 s_{12}s_{23}s_{13} \sin \delta_{\text{KM}} \quad (4.21)$$

$$\approx \lambda^6 A^2 \eta. \quad (4.22)$$

This is our parameterization-independent quantity that measures the amount of CP violation in our model. The most remarkable observation is that this depends on *every* physical mixing angle! Thus if *any* of the mixing angles are zero, there would be *no* CP violation. This emphasizes the deep connection between flavor and CP. In fact, we can see that the amount of CP violation in the Standard Model is small, but it is *not* small because the CP phase δ_{KM} is small. Quite on the contrary, it is small because of the mixing angles. We can see this in the Wolfenstein parameterization where the Jarlskog invariant comes along with six powers of λ .

The Jarlskog is a geometric object coming from the theory of the $U(3)$ group. Some details can be found in [41].

4.5 Unitarity triangles and *the* unitarity triangle

Using the unitarity of the CKM matrix we can write down equations for the off diagonal elements of VV^\dagger , for example

$$\sum_i V_{id}V_{is}^* = 0. \quad (4.23)$$

We have 6 such relations (three for the rows, three for the columns) and can plot each relation as a triangle in the complex plane. These are called **unitarity triangles**.

Some of these triangles are rather flat and almost linear. Consider the example above:

$$V_{ud}V_{us}^* + V_{cd}V_{cs}^* + V_{td}V_{ts}^* = 0. \quad (4.24)$$

Remembering the Wolfenstein parameterization, we can see that the first two terms are $\mathcal{O}(\lambda)$ while the third term is $\mathcal{O}(\lambda^5)$. The last term is very small and so one ends up with two (normalized) sides that are $\mathcal{O}(1)$ and one that is smaller by a factor of $\sim 10^{-3}$.

Instead, let's look at the most interesting unitarity triangle. It is so interesting that it has a 'special name,' we call it ***the* unitarity triangle**.

$$\sum_i V_{id}V_{ib}^* = 0. \quad (4.25)$$


Doing the same order-of-magnitude estimate from the Wolfenstein parameterization,

$$V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* = 0. \quad (4.26)$$

is a sum of terms that are *each* $\mathcal{O}(\lambda^3)$. Each side is roughly the same order, so we can expect it to look much more triangular. To normalize our triangle, it is customary to normalize the sides. Let us divide by $V_{cd}V_{cb}^*$. The resulting relation is

$$\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*} + 1 + \frac{V_{td}V_{tb}^*}{V_{cd}V_{cb}^*} = 0. \quad (4.27)$$

How do we draw this? Treat each term in the sum as a vector on the complex plane. Recall from preschool that to sum complex numbers we can think of adding vectors by attaching the tail of each vector to the head of the preceding vector. So let's start at zero, add the unit term to get to $(1, 0)$, and sum either of the remaining two terms to get the third point of the triangle. (The third term will bring you back to zero, of course.) You end up with a nice triangle:



The great thing about *the* unitarity triangle is that because each side is $\mathcal{O}(1)$ relative to the others, it is robust against experimental errors.

Problem 4.2. A geometric interpretation of the Jarlskog invariant. Show that the area of all six unitarity triangles are all exactly equal to $J/2$. This makes sense since each of the triangles gives the *same* information about CP violation and we argued that the parameterization-invariant quantity that captures this information is J . *Note: this holds for unitarity triangles before normalizing one side to unit length.*

Problem 4.3. Area of unitarity triangles. Check the claim of the previous problem by writing out the areas of any two different unitarity triangles written with respect to the standard parameterization of the CKM matrix.

This is a very famous triangle. As Yuval says, “There aren’t many famous triangles in the history of mankind, [but] this is one of them.” In the diagram above we gave the angles names, α, β, γ . Note that there is physical significance to the direction with which we measure the angles; for example, if (ρ, η) had been in the lower quadrant the signs would differ. According to popular legend, these were given by Yossi Nir while he was at SLAC. There’s another parameterization that was also invented at SLAC by Bjorken,

$$\phi_1 = \beta \qquad \phi_2 = \alpha \qquad \phi_3 = \gamma. \qquad (4.28)$$

These appear to have been popularized in Japan by Sanda so that now the flavor physicists at Belle use different conventions from those at BaBar and LHC**b**. In fact, at a recent flavor conference in Japan the proceedings were edited so that all of the articles used the ϕ notation.

In terms of the CKM matrix parameters, the angles are given by

$$\alpha = \arg \left(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*} \right) \qquad (4.29)$$

$$\beta = \arg \left(\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*} \right) \qquad (4.30)$$

$$\gamma = \arg \left(\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*} \right). \qquad (4.31)$$

For completeness, we can also write out the length of the sides of the unitarity triangle in terms of these parameters,

$$R_u = \left| \frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*} \right| = \sqrt{\rho^2 + \eta^2} \qquad (4.32)$$

$$R_t = \left| \frac{V_{td}V_{tb}^*}{V_{cd}V_{cb}^*} \right| = \sqrt{(1 - \rho)^2 + \eta^2}. \qquad (4.33)$$

The remaining side is, by normalization, set to unit length.

5 Charged versus neutral currents

What did experimentalists see in the charged current and how did this lead to the construction of the Standard Model?

We know from our theoretical construction of the Standard Model that we've chosen a basis for the quarks where the only flavor-off-diagonal interaction comes from the W -boson, which we know is charged. Thus the neutral boson interactions are flavor-conserving. Now let's switch gears and look at this from the experimental side. Suppose we don't know anything about the Standard Model, all we know are the experiments whose data is summarized in the PDG.

Let us 'go through the data' and compare some charged current and neutral current processes.

Example 5.1. Kaon decay. Consider the following two decays:

$$\text{Br}(K^+ \rightarrow \mu^+ \nu) = 64\% \qquad \text{Br}(K_L \rightarrow \mu^+ \mu^-) = 7 \times 10^{-9}. \quad (5.1)$$

Because the left-hand side of each decay is hadronic and the right-hand side is leptonic, we can determine that $K^+ \rightarrow \mu^+ \nu$ has a charged intermediate state (charged current) and that $K_L \rightarrow \mu^+ \mu^-$ has a neutral intermediate state (neutral current).

What do we learn from this? Both processes change flavor, but the charged current has a much *much* larger rate—larger by eight orders of magnitude. Naïvely we would expect both processes to come from some flavor-violating structure so that they should both be of the same order of magnitude. Do we see the same pattern in other decays?

Example 5.2. B and D decays. Making use of the PDG, we find for two particular B decays

$$\text{Br}(B^\pm \rightarrow D^\pm \ell \nu_\ell X) = 11\% \qquad \text{Br}(B^0 \rightarrow K^* \ell^+ \ell^-) = 10^{-6}. \quad (5.2)$$

Similarly, for two particular D decays,

$$\text{Br}(D^\pm \rightarrow K^0 \mu^\pm \nu) = 9\% \qquad \text{Br}(D^0 \rightarrow K^\pm \pi^\mp \mu^\pm \mu^\mp) < 5 \times 10^{-4}. \quad (5.3)$$

Note, as one can see from the PDG, experimentalists have not yet measured flavor-changing neutral currents in the charm sector, this is why we only have upper bounds.

We see the same pattern! The charged current interaction is much larger than the neutral current.

Further, we should note that this *only* shows up for flavor-*changing* processes. For the flavor-conserving processes, the weak-interactions charged and neutral currents indeed occur at roughly the same rates. We leave it to look up demonstrative examples in the PDG.

Example 5.3. Isolating the weak interaction, part I. One set of flavor-conserving processes are $\nu_e X \rightarrow e X'$ and $\nu_e X \rightarrow \nu_e X'$ where X and X' are nuclear states. Why would it be misleading to look at the rates for $e X \rightarrow e X'$? The answer is that this is dominated by the electric coupling and so the overall rate will not tell us about interactions mediated by the *weak* force.

Problem 5.1. Isolating the weak interaction, part II. Experimentalists are clever. Show that you are also clever by explaining how one can experimentally isolate the weak part of the $eX \rightarrow eX'$ rate. **Hint:** a general principle is, “we can see very small things if they break a symmetry.”

The large discrepancy between the charged and neutral flavor-changing currents and the order-of-magnitude agreement in the flavor-conserving currents is a major experimental observation about the structure of the underlying theory. In fact, flavor-changing neutral currents (FCNCs) are always a big issue in designing models beyond the Standard Model.

If we haven't emphasized it enough so far, *familiarize yourself with the PDG*. This is a theorist's best friend when one needs to connect to experiment. When we want to measure something, our 'experiment' is to open up the PDG and look up the actual measured value. Take time to do this; experimentalists spent tens of years extracting and thinking about this information—you can afford to spend ten minutes doing the same thing.

6 Why FCNCs are so small in the Standard Model

6.1 Diagonal versus universal

Let's start with some almost trivial remarks. We know that flavor-changing effects occur due to off-diagonal couplings. Thus one might want to say that couplings do not change flavor if they are **diagonal** in flavor-space. This is generally *not* true since 'diagonal' is a basis-dependent property. One can shift to a different basis and a diagonal matrix will transform into a matrix with both diagonal and off-diagonal parts. It is only **universal** matrix—those which are proportional to the identity—which are diagonal in every basis. Sometimes it is sufficient to be diagonal, but in general flavor-changing effects are only protected against by having couplings that are universal.

6.2 FCNCs versus gauge invariance

In the Standard Model we know that we have flavor-changing in the charged current, namely the W boson. So what about flavor-changing neutral currents? As a caveat in this section we will only discuss *tree-level* FCNCs; the appearance of loop-level FCNCs will be a major topic later in this course. It is sufficient for us to look only at the bosons of the Standard Model since they are the force-mediators.

Example 6.1. No fermion force mediators. Fermions cannot be force-mediators since forces are *classical* effects; in other words, we are interested in the effective potential. Long range forces come from a multi-state background of quantum particles. By their statistics fermions cannot form coherent states and so one cannot construct a classical fermion background. Simply put, fermions are inherently quantum mechanical. Or, in Feynman's words when describing a neutrino model of quantum gravity [76] (see also [53]),

Since the neutrino has half-integral spin, the single exchange does not result in a static force, because after a single exchange the neutrino source is no longer in the same state as it was initially. In order to get a force out of an exchange, it is necessary that the

diagram including the exchange should be able to interfere with the diagonal terms in the scattering amplitude, that is, it should add to the amplitude that nothing happens.

The relevant Standard Model gauge bosons are the Higgs h , the Z , the photon γ , and the gluon g . (Alternately, the neutral Higgs, the pseudoscalar Higgs, W^3 , B , and g .) Why do these particles not have any tree-level FCNCs? The photon and gluon have a simple answer: gauge invariance. Gauge symmetry forces the kinetic terms (where the gauge couplings live) to be universal.

6.3 FCNCs versus Yukawa alignment

The Higgs couplings to the fermions comes from the Yukawas. In the Standard Model these are trivially diagonal since they are proportional to the fermion mass matrices. In other words, they are **aligned** with the mass matrices¹³. For example, consider the down-type Yukawas,

$$y_{ij}^d \bar{Q}^i \phi D^j. \quad (6.1)$$

Expanding the Higgs ϕ about its vev, we obtain

$$y_{ij}^d \bar{Q}^i (v + ih) D^j. \quad (6.2)$$

Thus it is clear that diagonalizing the mass matrix $v y^d$ simultaneously diagonalizes the $\bar{Q} h D$ coupling. This is one of the rare examples when it is sufficient for a matrix to be diagonal but not universal. This is very specific to the fact that the Standard Model only has one Higgs. Consider adding a second Higgs in the most trivial way,

$$y_{ij}^1 \bar{Q}_L^i (v_1 + ih_1) D_R^j + y_{ij}^2 \bar{Q}_L^i (v_2 + ih_2) D_R^j. \quad (6.3)$$

In this model we can give each Higgs a different potential, $V_1(H_1)$ and $V_2(H_2)$ so that h_1 and h_2 have different masses no mixing. The point is that the fermion mass matrix is now $(y^1 v_1 + y^2 v_2)$ while the Higgs couplings are governed by y^1 and y^2 separately. One can see that now the Higgs coupling matrices are (in general) *not* proportional to the fermion mass matrix and so these are flavor-changing.

Example 6.2. Two Higgs Doublet Models. Note that this is *not* the way we usually add a second higgs in two Higgs doublet models (2HDM). In those models the up- and down-type sectors couple to different Higgs bosons which each get an independent vev. In this case the Higgs couplings *are* diagonal because they are aligned with the mass matrices. To provide some context, the most famous 2HDM model is the MSSM since supersymmetry requires at least two Higgses due to constraints from holomorphy and anomaly cancellation. Note that if you have four Higgses in the MSSM you again generally find FCNCs through the Higgs couplings.

Example 6.3. Experimental signature of Higgs FCNCs. This model has explicit flavor-changing neutral currents in the quark sector. Where do you expect to see trouble? We'll narrow it down to a multiple choice question: do you expect more disagreement in the kaon sector or the

¹³This alignment plays a major role in the suppression of certain $\mu \rightarrow e\gamma$ diagrams in a warped extra dimensions[25].

B sector? **Answer:** In this particular model there is some subtlety because the mass matrix can be very different from the couplings, but let us assume that this effect is not particularly perverse. Because the Higgs couples according to the Yukawas, it couples more strongly to massive particles. Thus we end up with a much stronger coupling to B mesons (from the b quark) than kaons (which only have an s quark). This is an important lesson: there are some cases where FCNCs in the kaon sector are negligible and offer model-builders some wiggle room.

6.4 FCNCs versus broken gauge symmetry reps

Now what about the Z ? This is a gauge boson, but it is a gauge boson for a *broken* gauge symmetry so there is no reason to expect gauge invariance to protect against FCNCs. We don't have FCNCs at tree-level in the Standard Model, so something is still protecting the Z .

Because the Z is neutral it only connects fermions with the same electric charge and color. Color is trivially satisfied since $SU(3)_c$ has nothing to do with $SU(2)_L \times U(1)_Y$. Electric charge, on the other hand, is related to the $SU(2)_L$ and $U(1)_Y$ charge by $Q = T_3 + Y$. In the SM all quarks with the same electric charge Q also have the same T_3 and therefore also the same Y . Recall that the Z coupling only depends on these quantities,

$$g^{Zff} = g \cos \theta_W T_3 - g' \sin \theta_W Y. \quad (6.4)$$

Thus particles of the same charge all have the same coupling to the Z , i.e. the Z coupling is universal for each of the up-type and (separately) down-type quarks. This is why the Standard Model Z doesn't give flavor-changing neutral currents. To see this explicitly, consider the terms in the Lagrangian giving the Z coupling to left- and right-handed up quarks,

$$\mathcal{L}_Z = \frac{g}{\cos \theta_W} \left[\bar{u}_L^i \gamma_\mu \left(\frac{1}{2} - \frac{2}{3} \sin^2 \theta_W \right) u_L^i + \bar{u}_R^j \gamma_\mu \left(-\frac{2}{3} \sin^2 \theta_W \right) u_R^j \right] Z^\mu. \quad (6.5)$$

To go to the mass basis, we perform a unitary rotation on the external fields. Let's just write out the u_L term,

$$\mathcal{L}_Z = \frac{g}{\cos \theta_W} \left[\bar{u}_L V_{uL} \gamma_\mu \left(\frac{1}{2} - \frac{2}{3} \sin^2 \theta_Q \right) V_{uL}^\dagger u_L \right] Z^\mu \quad (6.6)$$

What is the structure of this term in flavor space? Everything in the parenthesis is universal. Thus the flavor structure is trivial,

$$V_{uL} V_{uL}^\dagger = \mathbb{1}. \quad (6.7)$$

and we have no FCNCs.

Problem 6.1. A homework problem for Yuval. As a useful exercise, rewrite (6.5) using Weyl spinors. Be sure to get the bars and Pauli matrices correct.

Let us write this in terms of a general principle,

Theorem 6.4. *In order to completely prevent flavor-changing neutral currents in the gauge sector, particles with the same unbroken gauge quantum numbers must also have the same quantum numbers under the broken gauge group.*

So in the Standard Model we have no Z FCNCs because the particles of a given electric charge also have the same $SU(2)_L$ charge (namely T_3). In general this needn't have been the case. Consider a slightly different model where the d_L and s_L quarks had different $SU(2)_L$ representations, but still the same electromagnetic and color charge.

$$\bar{d}_L \left[T_3^{(d)} - Q \sin^2 \theta_W \right] d_L \tag{6.8}$$

$$+ \bar{s}_L \left[T_3^{(s)} - Q \sin^2 \theta_W \right] s_L \tag{6.9}$$

If $T_3^{(d)} \neq T_3^{(s)}$ then we would not be able to avoid FCNCs.

Example 6.5. For example, consider the somewhat contrived case where $Q_{d,s} = 0$, $T_3^{(d)} = 1$, $T_3^{(s)} = 0$. Thus the s doesn't see the Z and the flavor structure of the relevant term in the Z coupling will be

$$V_{dL} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} V_{dL}^\dagger. \tag{6.10}$$

This is trivially not universal.

Example 6.6. Why the charm shouldn't have been surprising. The charm quark was discovered in the J/ψ particle in 1974. The Cabibbo angle for two-generation mixing was understood in the 1960s, but charm wasn't even hypothesized until the 70s. Thus through most of the 60s it was well understood that there was a (u, d) isospin (really $SU(2)_L$) doublet and a weird s quark with the same charge as the d but apparently living in an iso-singlet. This was a missed opportunity! We now see that this structure clearly has flavor-changing neutral currents and physicists should have expected to see $K_L \rightarrow \mu^+ \mu^-$. It wasn't until Glashow, Iliopoulos, and Maiani that physicists realized that the non-observation of this mode implies that there should be a charm quark to complete the doublet. By the time physicists discovered the b physicists had learned their lesson and realized immediately that there should also be a t quark. (Or course they expected a t that was not much heavier than the b ; the story of this hierarchy is one of the great unsolved problems in flavor physics.)

Problem 6.2. Exotic light quarks, part II. Recall the model in Problem (2.2), where we change the Standard Model quark sector by removing the c , b , and t quarks and change the representations such that the light quarks are composed of an $SU(2)_L$ doublet $Q_L = (u_L, d_L)$, and singlets s_L and s_R . The quantum numbers are assigned to maintain the same color and electric charge as the Standard Model. Answer the following questions (it helps to have done the first part of this problem):

1. Write out the gauge interactions of the quarks with the Z boson in both the interaction and the mass bases. (Only write the terms that change when you rotate to the mass basis.) Are there, in general, tree-level Z exchange FCNCs?
2. Are there photon and gluon mediated FCNCs? Support your answer by an argument based on symmetries.

3. Are there Higgs exchange FCNCs?
4. Do you expect this model to be experimentally ruled out or not? Why?

Problem 6.3. Exotic light quarks, part III. Repeat the analysis for Problems 2.2 and 6.2 for a modified exotic light quark model where the u_R and d_R are part of an $SU(2)_L$ doublet.

Problem 6.4. The two Higgs doublet model. Consider the two Higgs doublet model (2HDM) where the Standard Model is extended by an additional Higgs doublet. We label the two Higgses ϕ_1 and ϕ_2 with $SU(2)_L \times U(1)_Y$ representations $\mathbf{2}_{-1/2}$ and $\mathbf{2}_{1/2}$. For simplicity we will work with two generations.

- Write down the most general Yukawa potential for the quarks.
- Carry out the diagonalization procedure for such a model. Show that the Z couplings are still flavor diagonal.
- Show that in general the Higgs mediates FCNCs.
- Is this model experimentally viable?

7 Parameterizing QCD

In our review of hadrons we already presented the general problem of QCD: because of confinement the ‘fundamental’ objects in our theory are not the degrees of freedom that we observe in experiments. Our Lagrangians are written in terms of quarks, but our world is made up of hadrons. Extracting quark interactions from these hadronic interactions is far from trivial.

When doing scattering in quantum field theory we work with asymptotic states, say electrons and positrons. In QCD, however, we can’t have quarks at infinity. Thus, even at a formal level, a lot of what we say about quark interactions in QFT isn’t strictly very well defined. We can talk about pion field theories, but then we’re working with pion couplings and are still removed from our fundamental QCD Lagrangian. Thus in simple words, the fundamental problem of QCD is that our theory is described in one language and our experiments are done in another¹⁴. Fortunately there are some ways around this; partly due to clever parameterization. Before moving forward it is important to introduce some language that lies halfway between theory and experiment (a sort of Yiddish). The vocabulary is as follows:

- **Matrix element.** This is just what we mean by scattering amplitude. You should be familiar with this from non-relativistic quantum mechanics.
- **Form factor.** This should sound familiar from (non-relativistic) scattering such as Rutherford scattering. It is the shape correction to the approximation that a scattering object is point-like (S -wave).
- **Decay constant.** This is a non-perturbative property of the mesons that we met briefly when we discussed chiral perturbation theory.

¹⁴Sometimes this is how it feels to type up these notes...

- **Factorization.** This is the deep idea that we may separate different kinds of physics. For example, physics at different scales ought to decouple. In our case, we will be interested in the factorization of the hadronic matrix elements from the leptonic matrix elements. This factorization is realized because the operators that mix hadronic and leptonic pieces are formed by integrating out an intermediate W at ‘low energy’ scales much lower than M_W .

Let’s familiarize ourselves with the concepts behind these words. Let’s start by remembering how we do calculations and start with a simple purely leptonic example. Consider the process $W \rightarrow \ell\nu$. The amplitude is written as

$$\mathcal{A} = \langle \ell\nu | \mathcal{O} | W \rangle, \quad (7.1)$$

where \mathcal{O} is some operator in the Hamiltonian (or Lagrangian up to a sign). We can then write down the **decay rate** for the W to a lepton and its neutrino by

$$\Gamma \sim \int |\mathcal{A}|^2 \times d(\text{phase space}). \quad (7.2)$$

We know that in this example the amplitude is trivially $\mathcal{A} = g/\sqrt{2}$. One can write down the relevant tree-level term for \mathcal{O} ,

$$\mathcal{O} = \bar{\ell}\gamma_\mu W^\mu \nu \quad (7.3)$$

and do the appropriate Wick contractions with the external states. Note that the point here is that the external states are asymptotic so that our operator can ‘literally’ be written in terms of them.

Moving on, we can see what happens when we look at hadronic decays. Consider $\pi \rightarrow \ell\nu$. The matrix element is simple to write,

$$\mathcal{A} = \langle \ell\nu | \mathcal{O} | \pi^+ \rangle. \quad (7.4)$$

The leptonic part of this matrix element is simple, precisely because we again have asymptotic external states which appear as creation and annihilation operators in \mathcal{O} . What about the π^0 ? Let’s say this is a $u\bar{u}$ combination (ignore the $d\bar{d}$ part for now), then we must create an up–anti-up pair. Because this is a *low energy* process ($m_\pi \ll M_W$), we can integrate out the W boson from our theory. The key point is that our operator \mathcal{O} factorizes into the leptonic and hadronic currents that attached to either end of the W boson:

$$\mathcal{O} = \mathcal{O}_\ell \frac{1}{M_W^2} \mathcal{O}_H = (\bar{\ell}\gamma^\mu \nu) \frac{1}{M_W^2} (\bar{u}\gamma_\mu u), \quad (7.5)$$

where we’ve written quotes to indicate that this is a heuristic form of the operators, we’ll write it out honestly below. (In particular, we’ve been sloppy with factors of, e.g., γ^5 .) This means that our matrix element should also factorize,

$$\langle \ell\nu | \mathcal{O} | \pi^+ \rangle = \langle \ell\nu | \mathcal{O}_\ell | 0 \rangle \frac{1}{M_W^2} \langle 0 | \mathcal{O}_H | \pi^+ \rangle. \quad (7.6)$$

The object on the left-hand side is simple to calculate. Don't get confused by the form of the $\langle \ell\nu | \mathcal{O}_\ell | 0 \rangle$ matrix element, it looks like some weird expectation value for, say, the leptonic current—it doesn't make sense kinematically on its own, the matrix element isn't a kinematic object. (The kinematics are split with the hadronic part and are really only accounted for in the phase space integral.) The factor of $1/M_W^2$ combines with coupling constants (implicitly in the \mathcal{O} s) to give the famous Fermi coupling,

$$\frac{G_F}{\sqrt{2}} \equiv \frac{g^2}{8M_W^2}. \quad (7.7)$$

The numerical factors are historical and are motivated below (B.1).


The $\langle 0 | \mathcal{O}_H | \pi^0 \rangle$ might cause some concern. It looks something like a quark current expectation value, except for the fact that the external state has *no explicit quarks*—it's a pion. Sure, a pion is made of quarks, but it's also made up of glue and is an actual asymptotic state while quarks are not. The easiest way to deal with this is to just define this object to be a parameter of our effective theory. We call it the pion **decay constant**, f_π , and we can *heuristically* define it as

$$\langle \pi | \mathcal{O} | 0 \rangle \sim f_\pi. \quad (7.8)$$

We'll get to a more proper and technical definition below; for now just focus on the ideas. Decay constants capture all of the intractable-to-calculate non-perturbative “brown muck” that keeps the pion together. In recent times lattice QCD calculations have made huge steps in being able to calculate these objects to the percent level, but in principle we should treat these as additional physical parameters whose values should be fixed by experiment.

Now let's look at more complicated hadronic decays. Consider $K^+ \rightarrow \pi^0 \ell^+ \nu$.

Example 7.1. Draw the Feynman diagram for this process. (This should be very basic review!)



We've only drawn the $u\bar{u}$ content of the π^0 , but we know by now that the π^0 is really the linear combination $(u\bar{u} - d\bar{d})/\sqrt{2}$. Get used to drawing Feynman diagrams quickly in your head. Experimentalists and flavor physicists deal with the processes whose diagrams all have similar topologies and so no longer bother to explicitly draw diagrams for each other.

The matrix element is again simple to write,

$$\mathcal{A} = \langle \pi^0 \ell \nu | \mathcal{O} | K^+ \rangle. \quad (7.9)$$

The leptonic part of this matrix element is simple, precisely because we again have asymptotic external states which appear as creation and annihilation operators in \mathcal{O} . What about the K^+

and the π^0 ? We know that $K = \bar{s}u$ and for this decay we only care about the $u\bar{u}$ part of the π^0 . The operator that we need looks like¹⁵

$$\bar{u}\gamma_\mu s = \psi_u \sigma_\mu \bar{\psi}_s + \bar{\chi}_u \bar{\sigma}_\mu \chi_s \quad (7.10)$$

since we want it to annihilate an anti- s quark and create a u quark. Of course, these are statements about quarks. What we *really* want is something that annihilates a kaon and generates a pion. This is where our QCD ‘language’ problem manifests itself: the quark operator coming from the Lagrangian is not directly related to the hadronic external states coming out of an experiment.

We certainly have creation and annihilation operators on the quark level, but just because you annihilate an \bar{s} and produce a u it does *not* mean that you have automatically taken a K^+ and turn it into a π . This could have been the quark-level operator for any number of other processes with different external states, say $\Lambda \rightarrow p\mu\nu$. So how do we calculate this matrix element?

Once again we appeal to factorization. In particular, we know that we can (heuristically) write

$$\mathcal{O} = \mathcal{O}_\ell \frac{1}{M_W^2} \mathcal{O}_H = (\bar{\ell}\gamma^\mu\nu) \frac{1}{M_W^2} (\bar{u}\gamma_\mu s), \quad (7.11)$$

where we’ve explicitly separated the leptonic and hadronic sides of the integrated-out W boson diagram. The quotation marks remind us not to take the actual explicit form of the $\mathcal{O}_{\ell,H}$ operators that we wrote above too seriously yet, we’ll shortly discuss their actual form. The matrix element thus factorizes,

$$\langle \pi^0 \ell\nu | \mathcal{O} | K^+ \rangle = \langle \ell\nu | \mathcal{O}_\ell | 0 \rangle \frac{1}{M_W^2} \langle \pi | \mathcal{O}_H | K^+ \rangle. \quad (7.12)$$

As we know from above, we’re perfectly happy calculating the leptonic piece, $\langle \ell\nu | \mathcal{O} | 0 \rangle$. The hard part is the hadronic piece, $\langle \pi | \mathcal{O} | K \rangle$, which is deeply entrenched in how QCD converts quarks into bound states. In the case of a single hadron we ended up packaging this non-perturbative garbage into a single parameter, the decay constant. What about when we have two hadrons, H_1 and H_2 ? We can define another parameter,

$$\langle H_1 | \mathcal{O} | H_2 \rangle \sim F. \quad (7.13)$$

We call this a **form factor** (where, again, we’ll save a technical definition for later). We choose this name because this looks like scattering: one state H_1 coming in and a different state H_2 coming out. Usually the physical process includes leptons coming out, so form factors are closely related to semileptonic decays, e.g. $B \rightarrow D\ell\nu$.

Matrix elements with more than two hadrons do not have any special names.

Being able to calculate these decay constants and form factors would already be a big step. Peter Lepage, currently the Dean of Cornell’s College of Arts and Sciences, made a career out of making lattice measurements of these non-perturbative constants. While there have been many

¹⁵This is the last time for a while that I’ll explicitly write out Weyl spinors since it will actually be convenient to see explicit factors of γ^5 rather than ambiguous minus signs. This, I think, is one of two good reasons to use Dirac spinors. The other good reason is because one is working in 5D where the Dirac representation is required. (This makes you wonder why Csaba is so attached to Weyl spinors, doesn’t it?)

achievements on this front, such as correctly predicting some values before they were experimentally measured, there's still a long way to go before the QCD 'language problem' is 'solved'. For example, three body hadronic matrix elements are still in the asymptotic future. We need to find a way to get around this so that we can take one non-perturbative result as a parameter and apply it to make predictions of other processes. A more recent approach involves the hope that the AdS/CFT correspondence might be able to provide a perturbative method to calculate these strongly-coupled objects; this is often referred to as AdS/QCD. Recent progress has led to promising results in meson spectroscopy, but predictions of actual matrix elements are still 'undelivered' [26].

To reiterate the big picture: Recent progress with holographic and lattice techniques notwithstanding, we cannot calculate hadronic matrix elements from first principles. The best that we can do is to *parameterize our ignorance* using symmetries. Indeed, a bit of cleverness can go a long way. Sometimes we can just measure the matrix elements in one process and apply the result to other processes—this is the best case, one doesn't even need theory to do this! When this fails, however, we have to be a bit more slick and use approximate symmetries. The program is as follows

1. Use symmetries to determine which matrix elements are relevant.
2. For an unknown hadronic matrix element, write out the most general linear combination of dynamical variables (e.g. momenta) which match the Lorentz structure of the matrix element. (Do not forget the discrete symmetries!)
3. Write the coefficients of these terms as parameters (decay constants and form factors) that have to be determined.
4. Use every trick you can to determine which of these terms are relevant.

7.1 The decay constant

Let's take a closer look at

$$\langle 0 | \mathcal{O} | \pi^+ \rangle. \quad (7.14)$$

The operator \mathcal{O} should annihilate a pion and so must have the form

$$\mathcal{O} \sim \bar{u} \Gamma d \quad (7.15)$$

for some Dirac structure Γ . We can always simplify the Dirac structure to one of S, P, V, A, T (scalar, pseudoscalar, vector, axial vector, tensor). Of these five operators, symmetries force some of them to vanish.

Example 7.2. Which operators vanish? Suppose we knew that the main contribution come from either the $V = \gamma^\mu$ or the $A = \gamma^\mu \gamma^5$ operator, while the other vanishes. Which one is zero?

Solution: We know that π is a pseudoscalar and QCD conserved parity. The vacuum $|0\rangle$ is parity-even, thus \mathcal{O} must be parity-odd, i.e. the vector must vanish and the main contribution must come from the axial operator. Similarly, $S = T = 0$.

Let's look at the axial matrix element $\langle 0 | \bar{u} \gamma^\mu \gamma_5 d | \pi \rangle$ a bit more closely. We don't know how to calculate this object, but we saw above that we can parameterize our ignorance in terms of a **decay constant** f_π , which we now properly define to be

$$\langle 0 | A^\mu | \pi \rangle \equiv -i p^\mu f_\pi. \quad (7.16)$$

There should be nothing surprising about the structure of the right-hand side. Lorentz structure forces us to have a p^μ on the right-hand side since this is the only vectorial quantity available to us. The rest of the right-hand side is an overall coefficient, which we define to be $-i f_\pi$ with mass dimension 1. In principle this *could* have had some dependence on Lorentz scalar quantities, of which we only have m_π , but this is not a *dynamic* variable—it is some fixed value that does not change when we change the kinematic configuration of the decay.


It turns out that f_π can be measured and has the value (whose order of magnitude you should memorize)

$$f_\pi \sim 131 \text{ MeV}. \quad (7.17)$$

Sometimes this will be written in terms of self-adjoint isospin currents and fields as¹⁶,

$$\langle 0 | j_{5a}^\mu | \pi_b \rangle = i \delta_{ab} F_\pi p^\mu \quad F_\pi \equiv \frac{f_\pi}{\sqrt{2}} \approx 93 \text{ MeV}. \quad (7.18)$$

How do we make this measurement? We should look at the dominant pion decay, $\pi^+ \rightarrow \mu^+ \nu$, which has a branching ratio of about 99.99%. The amplitude for this decay is given by the simple s -channel diagram:



Problem 7.1. Draw this diagram using [massless] chiral spinors.

It is straightforward to calculate the amplitude for this diagram,

$$\mathcal{M}(\pi \rightarrow \mu \nu) = -\frac{g^2}{4M_W^2} f_\pi V_{ud} m_\mu \bar{\mu}_R \nu_L. \quad (7.19)$$

Problem 7.2. A basic Feynman diagram calculation. Derive (7.19). Yuval says this is a very important exercise in QFT, so derive it using *both* Dirac and Weyl spinors. Think about the pros and cons of each calculation. Explain the physical significance of the m_μ factor (**Hint:** see *Problem 7.1*).

¹⁶See Georgi chapter 4.3 for a discussion [40].

One can now convert this to a decay width (hint: the PDG has a review section for two- and three-body decay kinematics),

$$\Gamma(\pi \rightarrow \mu\nu) = \frac{G_F^2}{8\pi} |V_{ud}|^2 f_\pi^2 m_\ell^2 m_\pi \left(1 - \frac{m_\mu^2}{m_\pi^2}\right)^2. \quad (7.20)$$

The factor in the parenthesis is the factor coming from the phase space integral; it is conventionally normalized so that in the limit $m_\ell \rightarrow 0$ the phase space factor goes to 1. What about the factor of m_μ^2 ? This is clearly the same factor of m_μ in (7.19).

Problem 7.3. An exercise in calculating decay rates. Derive (7.20) from (7.19). *Hint: If you're spending a lot of time with γ matrix identities then you're doing this the hard way.*

At first sight this factor seems puzzling: this means that the $\pi \rightarrow \mu\nu$ decay is dominant over the $\pi \rightarrow e\nu$ decay, even though the latter has a much bigger phase space. How did this factor appear? It is there to preserve angular momentum. We know that the Standard Model is chiral. The W that mediates this decay only couples to left-handed fermions. However, the pion is a spin-0 scalar. The W that mediates this decay can only directly decay into left-handed fermions. Angular momentum conservation requires a left- and right-handed fermion pair, thus there must be an explicit Dirac mass term converting the left-handed muon $\mu_L = \chi_\mu$ to a right-handed muon $\mu_R = \bar{\psi}_\mu$. (Note that this interpretation is trivial when using Weyl spinors but somewhat subtle when using Dirac spinors.) We say that this decay has a **chiral suppression**.

Chirality versus helicity. It is interesting that experimentalists often refer to this as *helicity* suppression. These are two different—but closely related—ideas that every particle physicist should be comfortable with. Chirality is a property of a *field*, it is the difference between the $SU(2)_L$ doublet versus singlet fields (i.e. χ versus ψ). Helicity, on the other hand, is a property of a *particle* which one can measure (it is an *experimental* quantity). For massless fields chirality and helicity are the same. For massive fields one can flip chirality through mass insertions.

Everything in (7.20) is known except for f_π ; we can thus measure this decay to determine f_π . What is most interesting to us, however, is the *physics* behind this quantity. For example, how does this quantity scale? Can we understand it heuristically?

Example 7.3. The positronium ‘decay constant.’ It is often useful to think of a toy example as a crutch for physical intuition. In this case, an instructive example is positronium. This is an e^+e^- bound state and should have a decay constant, just like the pion. What is it?

The relevant question is how does positronium decay. Semiclassically this occurs when the electron and positron touch, i.e. the probability that they are in the same place. The lifetime of positronium clearly depends on the overlap of the e^+ and e^- wavefunctions. Assuming that the positron is at the origin, the ‘decay constant’ of positronium can thus be heuristically interpreted as something like $|\Psi_e(r=0)|^2$.

Caveat emptor! This is just a toy model to build physical intuition, it is not meant to be an actual representation of the physics of the pion!

From the above example, we can consider limits when one quark becomes very heavy. In this case its wavefunction becomes very narrow and we would expect the decay constant to decrease.

We've now explored the meaning of the axial matrix element and have parameterized the QCD brown muck into a single pion decay constant. Armed with this, what can we say about the pseudoscalar matrix element? It turns out that we don't need to do any more calculations or introduce any more parameters: we can use a slick trick to get the pseudoscalar matrix element from the axial matrix element. This trick falls under the umbrella of the current algebra, something which has fallen out of favor in modern textbooks. For a more formal introduction, see [20, 79, 35]. Here we'll just sketch what we need.

Let us consider the pseudoscalar kaon matrix element,

$$P = \langle K(p) | \bar{s} \gamma_5 u | 0 \rangle. \quad (7.21)$$

This is manifestly a Lorentz scalar quantity. The result that we would like to show is

$$P = i f_K \frac{m_K^2}{m_u + m_s}. \quad (7.22)$$

The trick to derive this result is to take the divergence of the axial current,

$$i \partial_\mu (\bar{s} \gamma_\mu \gamma_5 u) = i (\partial_\mu \bar{s}) \gamma_\mu \gamma_5 u + \bar{s} \gamma_\mu \gamma_5 (i \partial_\mu u) \quad (7.23)$$

$$= (i \not{\partial} \bar{s}) \gamma_5 u - \bar{s} \gamma_5 (i \not{\partial} u) \quad (7.24)$$

$$= -(m_s + m_u) \bar{s} \gamma_5 u. \quad (7.25)$$

Now take this and put it between the kaon and the vacuum. The derivative is just ($\partial_\mu = -i p_\mu$) is just a momentum, so we get

$$\langle K | i \partial_\mu (\bar{s} \gamma_\mu \gamma_5 u) | 0 \rangle = p_\mu (-i f_K p^\mu) = -i f_K m_K^2, \quad (7.26)$$

from this we obtain (7.22) straightforwardly.

Example 7.4. Scalar dominance. Pop quiz: which matrix element is more important, the pseudoscalar or the axial vector? The difference is a factor of $m_K^2 / (m_u + m_s)$. Since the kaon is made up of light quarks, the dominant mass scale is set by λ_{QCD} which is greater than the sum of the valence quark masses. Thus the (pseudo) scalar is enhanced relative to the (axial) vector operator. This is a generic feature for the light quark mesons. Note that this is a purely QCD effect (the decay constants parameterize QCD), electroweak theory only gives axial and vector contributions from the W .

7.2 Remarks on the vector mesons

So far we've discussed the decay constants for pseudoscalars. What about the vector mesons? Consider the ρ matrix element with the vacuum,

$$\langle 0 | \mathcal{O} | \rho \rangle = ? \quad (7.27)$$

Can we measure this object? What we are really asking is whether or not the ρ has a nonzero matrix element with the QCD vacuum, i.e. does it have leptonic decays via the weak interaction that would cause this matrix element to be nonzero? For the ρ the answer is *no*, it only decays via the strong force¹⁷, in other words, the ρ is a resonance; it decays into two pions. The same is also true for the ϕ . These particles do not decay weakly.

What about the (heavy) flavored vectors? The J/ψ and the Υ can both decay to leptons, so it makes sense to write down their decay constants. We see, however, that for the light vector mesons decay constants are not particularly useful. Naïvely we might want to write

$$\langle 0|V^\mu|\rho\rangle \stackrel{?}{\propto} f'_\rho p^\mu, \quad (7.28)$$

but this turns out to be *too* naïve. For two very important reasons. First the parities for the two sides do not match since ρ is a pseudoscalar. More fundamentally, we know from the LSZ theorem in quantum field theory that amplitudes with external vectors are proportional to the polarization vector, ϵ^μ .

A remark on the LSZ theorem. Just as we know that an amplitude with external fermions should have spinor contractions of the form $\bar{u}(p) \cdots u(p)$, an amplitude with an external vector should be proportional to the polarization vector ϵ^μ . We—of course—do *not* have the external plane wave spinors since the whole point of these decay constants and form factors is that we parameterize how these objects hadronize.

In fact, there's a slightly deeper reason why the matrix element must be proportional to ϵ^μ —the LSZ theorem in quantum field theory. Just as we know that an amplitude with external fermions should have spinor contractions of the form $\bar{u}(p) \cdots u(p)$, an amplitude with an external vector should be proportional to the polarization vector ϵ^μ . We—of course—do *not* have the external plane wave spinors since the whole point of these decay constants and form factors is that we parameterize how these objects hadronize.

Thus the correct thing to write is

$$\langle 0|V^\mu|\rho\rangle \sim f_\rho \epsilon^\mu, \quad (7.29)$$

where now it is clear that both sides transform the same way under parity and that the proportionality to ϵ from the LSZ theorem is satisfied.

Example 7.5. Vector parity. Let us check that the two sides indeed transform the same way under parity. ρ is a pseudoscalar and the vector current has parity $(-)^{\mu}$, where this is notation to mean

$$(-)^{\mu} = \begin{cases} + & \text{if } \mu = 0 \\ - & \text{if } \mu = 1, 2, 3 \end{cases} . \quad (7.30)$$

In other words, a vector transforms under parity as $P[V^\mu(x)] = V_\mu(P[x])$. The total parity of the matrix element is then $-(-)^{\mu}$, matching what we expect from an axial vector and only axial vector available on the right-hand side is ϵ^μ .

¹⁷A quick glance at the PDG will show a small branching ratio of the neutral ρ to $\ell^+ \ell^-$, but this is due to mixing with the ω .

Dynamical quantities. In the example of the vector mesons we saw that there are other dynamical quantities other than the external particle momenta. This might lead one to think a bit more about similar quantities that exist in our system. A natural question is whether we should also include the spin vectors s^μ of the quarks. It is a somewhat subtle point that these should *not* be considered to be real physical dynamical quantities! The spin index just identifies the chirality of the external state spinors; in other words, this just identifies the Weyl spinor χ or ψ (or their conjugates).

7.3 Form factors

Let us introduce form factors with a simple example which will be useful later on: the β decay of a neutron into a proton, $n \rightarrow p^+ e \bar{\nu} u$. Let us focus only on the vector current. The relevant matrix element that we cannot calculate from first principle is

$$\langle p^+(p') | \bar{d} \gamma^\mu u | \bar{n}(p) \rangle. \quad (7.31)$$

This is a ‘brown muck’ object that we’d like to parameterize. To do so we must write out the most general linear combination of kinematic variables (and products of those variables),

$$\langle p^+(p') | \bar{d} \gamma^\mu u | \bar{n}(p) \rangle \sim a p^\mu + b p'^\mu, \quad (7.32)$$

where the coefficients a and b are our **form factors**. We know that these can only depend on Lorentz scalars; there are three of these available: p^2 , p'^2 , and $p \cdot p'$. The first two are just masses and are not dynamical, they don’t change when we change the momenta. The third one is a bona-fide dynamical Lorentz scalar. It is conventional to write this in a different momentum basis.

Let’s define the momentum q by

$$q \equiv p - p'. \quad (7.33)$$

$$q^2 = p^2 + p'^2 - 2p \cdot p'. \quad (7.34)$$

This should be a familiar quantity from deep inelastic scattering calculations. We can now use $(p + p')$ and q as our basis of momenta in this problem. We write our form factors as f_\pm which are functions of q^2 and thus write our matrix element as

$$\langle p^+(p') | \bar{d} \gamma^\mu u | \bar{n}(p) \rangle = f_+(q^2) (p + p')^\mu + f_-(q^2) (p - p')^\mu. \quad (7.35)$$

For now we will leave this definition of the f_\pm form factors. Below we will apply it to the determination of the $|V_{ud}|$ CKM matrix element and see how, in specific cases, this structure can be made to simplify even further.

THIS SECTION UNDER CONSTRUCTION: INCLUDE “TIMELIKE/SPACELIKE” ANALYSIS FROM DISCUSSION WITH PIRJOL.

For a more non-trivial example of how to parameterize a form factor with more dynamical variables, see the following problem.

Problem 7.4. The $\langle D^{*+}(p_D, \epsilon) | V^\mu | \bar{B}(p_B) \rangle$ form factor. Consider the semi-leptonic decay $B \rightarrow D^* \ell \nu$. Let us again focus on calculating the hadronic vector matrix element.

1. Before diving into this particular decay, consider $B \rightarrow D^+ \ell \nu$. Convince yourself that the relevant hadronic vector matrix element takes the form of (7.35).

2. Prove that

$$\langle D^{*+}(p_D, \epsilon) | V^\mu | \bar{B}(p_B) \rangle = ig(q^2) \varepsilon^{\mu\nu\alpha\beta} \epsilon_\nu^* (p_D + p_B)_\alpha q_\beta, \quad (7.36)$$

where ε is the totally antisymmetric tensor.

3. Using the fact that T acts as a complex conjugate, show that $g(q^2)$ is real.

4. Parameterize the axial matrix element, $\langle D^*(p_D, \epsilon) | V^\mu | \bar{B}(p_B) \rangle$.

7.4 Aside: Goldstones, currents, and pions

This section aims to provide an informal glimpse of the current algebra techniques originally used to understand hadron decays. More advanced techniques, such as chiral perturbation theory, as well reviewed elsewhere. This section is outside the main narrative of these lectures and may be omitted, but it will hurt my feelings a little bit.

There is a somewhat antiquated way of looking at Goldstone's theorem based on what was called the current algebra that also connect the identity of the pions (and its octet cousins) as pseudo-Goldstone boson. This presentation comes from Cheng and Li Section 5.3 [20] and Coleman's lecture on soft pions [22].

Given a conserved current $\partial_\mu J^\mu(x) = 0$ we may define a conserved charge $Q(t) = \int d^3x J^0(x)$ that acts as the generator of the symmetry. When the symmetry is not spontaneously broken, this charge is also conserved. In the case of spontaneous symmetry breaking, the charge is no longer very well defined since the integral does not converge—but for our purposes this is a technical detail since we will only care about commutators. Current conservation trivially tells us that

$$0 = \int d^3x [\partial_\mu J^\mu(x), \phi(0)] \quad (7.37)$$

$$= -\partial_0 \int d^3x [J^0(x), \phi(0)] + \int d\mathbf{S} \cdot [\mathbf{J}(x), \phi(0)]. \quad (7.38)$$

Here we can identify the conserved charge associated with the symmetry current,

$$Q = \int d^3x J^0(x), \quad (7.39)$$

this also plays the role of the generator of the symmetry group with respect to the fields in the Lagrangian. For a sufficiently large volume the surface integral vanishes and we thus find that for a symmetry-preserving vacuum configuration,

$$\frac{d}{dt} [Q(t), \phi(0)] = 0. \quad (7.40)$$

Spontaneous symmetry breaking can thus be *defined* by the condition that this commutator does not vanish on the vacuum,

$$\langle 0|[Q(t), \phi(0)]|0\rangle \neq 0. \quad (7.41)$$

This is intuitively clear since it tells us that the generator of the symmetry enacts a non-trivial transformation of the vacuum configuration of field $\phi(0)$. Note that (7.40) *still holds* since this came from $\partial_\mu J^\mu(x) = 0$, which is a statement about the Lagrangian *independent* of the vacuum.

To properly understand this spontaneous breaking we may insert a complete set of orthonormal momentum states to write the left-hand side as (dropping an overall normalization),

$$\sum_n \delta^{(3)}(\mathbf{p}_n) (\langle 0|J_0(0)|n\rangle \langle n|\phi(0)|0\rangle e^{-iE_n t} - \langle 0|\phi(0)|n\rangle \langle n|J_0(0)|0\rangle e^{iE_n t}). \quad (7.42)$$

Spontaneous symmetry breaking is the condition where this expression is non-vanishing. Further, (7.40) tells us that this expression must be time-independent so that spontaneous symmetry breaking requires that there exists a state $|n'\rangle$ such that $E_{n'} = 0$ for $\mathbf{p}_{n'} = 0$, i.e. the state is massless. Further, we have the properties that $\langle n'|\phi(0)|0\rangle \neq 0$ and $\langle 0|J_0(0)|n'\rangle \neq 0$.

This turns out to be closely related to the LSZ reduction formula, which gives us a rigorous prescription for extracting off-shell matrix elements from the path integral by extracting the poles in the external particles. Going back to the definition of the pion form factor,

$$\langle 0|A^\mu|\pi\rangle \equiv -ip^\mu f_\pi, \quad (7.43)$$

we find an expression very similar (*unsurprisingly!*) to our manipulations above. Taking a divergence, we find that

$$\langle 0|\partial_\mu A^\mu|\pi\rangle \equiv m_\pi^2 f_\pi. \quad (7.44)$$

This is an interesting statement in light of the LSZ formula, which one can interpret as a statement about what it means to assign a field to a particle—this, of course, is related the whole point of parameterizing our ignorance about QCD matrix elements. In particular, any number of local operators can be used as field for a particle. If ϕ is a ‘good’ field, then so are objects made up of powers and derivatives of ϕ that maintain its canonical normalization. This is fixed from the LSZ formula and Lorentz invariance,

$$\langle p|\phi(x)|0\rangle = (2\pi)^{-3/2}(2E)^{-1/2}e^{ip\cdot x} \quad (7.45)$$

for a scalar field.

Example 7.6. A sketch of the proof. The LSZ reduction formula tells us that to obtain off-shell matrix elements we should take the n -point Green’s function

$$G^{(n)}(k_1, \dots, k_n) = \int \prod_k^i d^4x_i e^{i\sum_j k_j \cdot x_j} \langle 0|\hat{T} [\hat{\phi}(x_1) \cdots \hat{\phi}(x_n)] |0\rangle \quad (7.46)$$

where $\hat{\phi}(x_n)$ is the interaction picture field operator, and multiply it by inverse propagators to identify the on-shell pole in the external particles, e.g.,

$$\mathcal{M}(\text{init} \rightarrow \text{final}) \propto \int \prod_i^n d^4x_i \prod_j^n \mathcal{O}_j e^{ik_j \cdot x_j} \langle 0 | \hat{T} \left[\hat{\phi}(x_1) \cdots \hat{\phi}(x_n) \right] | 0 \rangle, \quad (7.47)$$

where \mathcal{O}_j is the appropriate operator in the free Lagrangian which cancels the external state propagators, in this case it is just the Klein-Gordon operator. We've neglected overall prefactors of the form $(2\pi)^{-3/2}(2E)^{-1/2}$ as well as renormalization factors. Here we have taken all external legs off shell by Fourier transforming in each of them; let us consider the simplified case where we do this for only one field, $\tilde{\phi}$, which may need not be a 'fundamental' field but is canonically normalized according to (7.45). The LSZ prescription tells us that to obtain an on-shell matrix element with $(n - 1)$ external ϕ particles and one external off-shell $\tilde{\phi}$ field, one calculates the necessary n -point functions, multiplies by $(p^2 - m^2)$ for the $\tilde{\phi}$ external line, and then takes the $\tilde{\phi}$ on shell. In other words, one identifies the pole in the $\tilde{\phi}$ squared momentum. The only diagrams with such poles are those with one particle intermediate $\tilde{\phi}$ states on this leg, and these are normalized by (7.45) to contribute a factor of unity. Thus we indeed find that we generate the *same* matrix element (i.e. combinations of couplings and internal propagators) that we would have generated for *any* canonically $\tilde{\phi}$ field. This is further explained in Section 2.1 of Coleman's lectures [22].

From all of this we see that in light of (7.45), what (7.44) is *really* trying to tell us is that

$$\phi_\pi = \frac{\partial_\mu A^\mu}{f_\pi m_\pi^2} \quad (7.48)$$

is a perfectly good field for the pions. Purists can attach $SU(2)$ indices as appropriate, e.g. $f_\pi \rightarrow \delta^{ab} f_\pi$, but—like minus signs—if you complain about it, then it's *your* homework to do it thoroughly¹⁸. This form of the pion field should not be surprising. After all, we know that in the limit where it is a Goldstone boson, it should have a shift symmetry. This is made explicit here by the derivative form of the field which guarantees the 'derivative coupling' that one finds in chiral perturbation theory. The relation of the divergence of the axial current to the pion is called the **partially-conserved axial current** (PCAC) hypothesis.

In fact, more generally, the LSZ statement about writing down an appropriate field to represent a given particle is at the heart of chiral perturbation theory. This was first put on solid footing by Callan, Coleman, Wess, and Zumino in what they called 'phenomenological Lagrangians' [23, 19, 21]. See Chapter 4 of Donoghue [58] for a more modern treatment with explicit demonstrations of the representation independence of such models.

Armed with the insight of (7.48), we can start to get a feel for the current algebra techniques once practiced by ancient physicists. One quick result is called the **Goldberger-Treiman** relation.

Problem 7.5. The Goldberger-Treiman relation. Derive the Goldberger-Treiman relation,

$$f_\pi g_{\pi NN} = m_N g_A(0), \quad (7.49)$$

¹⁸This policy is attributed to Csaba Csáki and his graduate electrodynamics course.

where $g_{\pi NN}$ is the pion-nucleon vertex function evaluated at $q^2 = 0$,

$$\langle p(k_p) | \partial_\mu A^\mu | n(k_n) \rangle = \frac{2f_\pi m_\pi^2}{q^2 - m_\pi^2} g_{\pi NN}(q^2) i \bar{u}_p(k_p) \gamma_5 u_n(k_n), \quad (7.50)$$

and $g_A(0)$ is the the nucleon axial vector coupling in

$$\langle p(k_p) | A_\mu | n(k_n) \rangle = \bar{u}_p(k_p) [g_A(q^2) \gamma_\mu \gamma_5 + q_\mu h_A(q^2) \gamma_5] u_n(k_n). \quad (7.51)$$

We can make more powerful use of the connection between the pion and the divergence of the axial current by invoking the algebra of the symmetry group. Recall that from a Lagrangian one can write down the current associated with a transformation

$$\phi_i \rightarrow \phi'_i = \phi_i + i\epsilon^a t_{ij}^a \phi_j \quad (7.52)$$

in the usual way,

$$J_\mu^a = -i \frac{\delta \mathcal{L}}{\delta(\partial^\mu \phi_i)} t_{ij}^a \phi_j. \quad (7.53)$$

From this we can define the generator/charge of the transformation using (7.39). For a symmetry of the Lagrangian, these generators satisfy the commutation relations of the symmetry group itself,

$$[Q^a(t), Q^b(t)] = i f^{abc} Q^c(t), \quad (7.54)$$

where we note that the charges are time-dependent and the algebra is satisfied at equal times. Charges are nice, but even better quantities are currents, since currents actually contain local (i.e. not integrated) products of the actual fields we care about. From the rigorous derivation of the charge commutation relations above one can also simply write out the commutation relation for charges with the time component of currents,

$$[Q^a(t), J_0^b(t, \mathbf{x})] = i f^{abc} J_0^c(t, \mathbf{x}). \quad (7.55)$$

By Lorentz invariance we can promote this tho

$$[Q^a(t), J_\mu^b(t, \mathbf{x})] = i f^{abc} J_\mu^c(t, \mathbf{x}). \quad (7.56)$$

In fact, it turns out we may go even further and write out current-current commutations relations, yielding the **current algebra**,

$$[J_0^a(t, \mathbf{x}), J_0^b(t, \mathbf{y})] = i f^{abc} J_0^c(t, \mathbf{x}) \delta^{(3)}(\mathbf{x} - \mathbf{y}). \quad (7.57)$$

We neglect the spatial J_i^a commutators since they turn out to include extra terms which usually end up canceling unphysical local singularities ($x \rightarrow y$). For the $SU(N)_L \times SU(N)_R$ ($N = 2, 3$) flavor symmetries of interest to us, the ancient (B.C., “before chiral perturbation theory”) current algebra between the vector and axial currents were used to produce relations between form factors

and pion scattering amplitudes. Commutation relations naturally arose when pulling out partial derivatives from time ordered products, for example

$$\partial^\mu \langle f|T(A_\mu(x)\mathcal{O}(0))|i\rangle = \langle f|T(\partial^\mu A_\mu(x)\mathcal{O}(0))|i\rangle + \delta(x_0)\langle f|T[A_0(x),\mathcal{O}(0)]|i\rangle, \quad (7.58)$$

where typically \mathcal{O} is another current. Thus the current algebra provided a way to simplify the matrix elements in hadronic physics before more sophisticated techniques were available.

Archaeologist–Physicists may learn more about these ancient techniques in Coleman’s lecture on soft pions [21], Cheng & Li [20], the review by Scadron [80], or the book by Treiman, Jackiw, and Gross [79].

8 Flavor symmetry and the CKM

Now that we’ve bounced back and forth between theory and experiment a few times, we see that experimental physics has robust flavor structure that is explained very well by the Standard Model. We must remind ourselves that the Standard Model (and any other model) is predictive after all of its parameters are measured. More practically, one makes several measurements and performs a global fit to determine the parameters. It is only when there are more measurements than parameters that the global fit becomes statistically meaningful¹⁹.

In order to ‘really probe’ the Standard Model’s flavor sector, then, we must measure and re-measure the four parameters of the CKM matrix. This will be our focus for some time. The motivation here is very important, since it is easy to feel like we’re making a big fuss about measuring particular each element of the CKM matrix. Nobody really cares about any particular element, say V_{cb} . The interesting part of determining its value is to add it to the global fit of the Standard Model parameters to actually test the theory. One might complain that as we go through each element of the CKM matrix we’re only measuring quantities which we could have extracted from the other elements; this is, in fact, the entire point! Each measurement only becomes meaningful when the parameter it is measuring has been measured independently in many other ways (e.g. through other elements of the CKM matrix).

8.1 Measuring $|V_{ud}|$


The key process to measure $d \rightarrow u$ transitions is already very familiar: β decay. This, however, comes in different forms:

- **Nuclear** β decay, for example tritium ${}^3\text{H}$ to ${}^3\text{He}$ or ${}^{14}\text{C}$ to ${}^{14}\text{N}$.
- **Neutron** β decay $n \rightarrow p e \nu$ involving a *free* neutron decaying to a *free* proton.
- **Pion** β decay, $\pi^+ \rightarrow \pi^0 e \nu$ or $\pi^+ \rightarrow \mu \nu$.

¹⁹There are some cases where one can make meaningful measurements without the entire fit, for example when one measures a positive-definite parameter to indeed be positive when it needn’t have been from the experiment. Statistically, however, this does not have the same weight and predictive power as an over-constrained global fit.

Even though these all probe the $d \rightarrow u$ transition, these are all different. Even though these are all mediated by the W boson and hence the same *operators* these processes differ by the initial and final states. In particular, they differ according to how the quarks are “dressed” by QCD.

The basic approach for determining V_{ud} in each of these processes is the same. We take the amplitude and compare it to muon decay.



How do these diagrams differ? The external states and the phase space integrals are different, but the actual Feynman rules differ only by a factor of V_{ud} .

8.1.1 Nuclear β decay

Let’s start with nuclear β decay. Nuclei are complicated objects and this seems hard to calculate from first principles. Fortunately, symmetries make our lives much easier. (In case you haven’t noticed, this is a major theme in this class!) In general we would like to determine the matrix element of an operator \mathcal{O} sandwiched between two nucleon states N and N' ,

$$\langle N | \mathcal{O} | N' \rangle. \tag{8.1}$$

In the Standard Model we know that \mathcal{O} can be vector or axial. In fact, we can pare this down further by considering the non-relativistic limit. The energy emitted by the electron in nuclear β decay is much smaller than the rest energy of the nucleus; for example, a tritium nucleus at rest hardly moves after undergoing β decay to ${}^3\text{He}$, thus such an approximation is well within reason. In this limit the axial current differs from the vector by a **spin flip**. This can be seen, for example, from the non-relativistic Dirac equation or by considering the effect of the γ^5 in the Dirac fermion basis.

Now we’d like to use this observation to further simplify our analysis. Independent of the non-relativistic limit, we know that we can appeal to cases where either the A or V operator vanishes. For example, $A = 0$ when the initial and final states are spin zero since a spin-less particle vanishes when acted upon by a spin-changing operator. Said differently, the vector and axial vector are parity odd and even respectively. The external states have a definite parity and since QCD respects parity, only one term can be non-zero. Fewer matrix elements mean fewer calculations, which makes us happy.

In nuclear physics these $N \rightarrow N'$ transitions between $J = 0$ states are called a **superallowed** β decays. In such a transition one of the neutrons within the nucleus N converts into a proton to produce nucleus N' . By definition nothing else has changed, certainly not spin. Now we can be very slick and invoke an approximate symmetry: in the limit of exact isospin symmetry, *nothing changes* when we make this $n \rightarrow p$ transition in the nucleus. The n and p are just states of isospin $+1/2$ and $-1/2$ respectively. This limit involves ignoring the mass differences and electric charges of these particles, but as far as QCD is concerned, it is almost as if nothing happened. The QCD transition matrix element should thus be unity.

This is a deep and useful tool for us: initially we have no right to expect to be able to say anything about this matrix element, but since we have some approximate symmetry which becomes exact in a well defined limit, we can sneakily exert some control over what's going on. In this case, isospin (as far as QCD is concerned) is broken by $\delta m = m_d - m_u \approx 1.4 \text{ MeV} \ll \Lambda_{\text{QCD}}$. We thus expect isospin breaking to give one or two percent corrections. In practice one can use a model to estimate the size of these corrections. Before going into some details of how isospin helps us do this calculation, the result turns out to be

$$|V_{ud}| = 0.97425(22) \implies \lambda = 0.2255(10). \quad (8.2)$$

Look at how many significant digits we have! Of course, here we don't care at all about those digits—but this is certainly an experimental triumph that we have such accuracy despite QCD. Let's sketch how we can get this result.

The most general decay amplitude $A \rightarrow B$ takes the form:

$$\mathcal{M}^\mu = \langle N'(p') | A^\mu + V^\mu | N(p) \rangle. \quad (8.3)$$

We shall choose spin-less initial and final states $J_N = J_{N'} = 0$ so that $\langle N' | A^\mu | N \rangle = 0$. Thus this reduces to the form factors that we introduced in (7.35),

$$M^\mu = \langle N'(p') | V^\mu | N(p) \rangle = f_+(q^2)(p + p')^\mu + f_-(q^2)q^\mu. \quad (8.4)$$

Now we invoke the isospin limit where $m_N = m_{N'}$. Note that in this limit the decay cannot occur kinematically! This is okay. In the isospin limit I can calculate the hadronic matrix element—the matrix element doesn't care about kinematics. Since the real world is very close to the symmetry limit, the matrix element at $q^2 = 0$ should be very close to its true value. In other words, q^2 is very small relative to Λ_{QCD}^2 so that

$$f_\pm(q^2) \approx f_\pm(0). \quad (8.5)$$

Example 8.1. We can test this assumption by looking at the spectrum of this decay with respect to q^2 . If f_\pm is constant then the spectrum is a straight line. This is precisely the dominant feature in the so-called Kurie plot of tritium β decay.

If we wanted, we can go on and calculate the corrections to this approximation. For the moment we will only focus on the 'big picture' of this process; later on we will actually calculate the corrections to a matrix element.

In the isospin limit we have yet another very helpful simplification: $f_-(q^2 \ll \Lambda_{\text{QCD}}^2) = 0$! To see this we invoke the Ward identity,

$$0 = q_\mu \mathcal{M}^\mu = f_+(q^2)(p^2 - p'^2) + f_-(q^2)q^2 = f_+(q^2)(m_A^2 - m_B^2) + f_-(q^2)q^2. \quad (8.6)$$

Isospin tells us that $m_N = m_{N'}$ so that the f_+ term vanishes. This means that $f_-(q^2)q^2 = 0$. Note that we are *not* treating $q^2 = 0$ here, even though we are implicitly assuming that $q^2 \rightarrow 0$ to get to this equation; we know that q^2 isn't physically *exactly* zero in the way that m_A and m_B are *exactly* equal in the isospin limit. Thus we conclude that $f_-(q^2) = 0$.

Now we've reduced everything to a single form factor, $f_+(0)$. We should be very proud of ourselves. Like any good infomercial, however, we have one last “*but wait, there's more!*” trick to pull out. It turns out that while we're working in the isospin limit, $f_+(0)$ doesn't even need to be determined experimentally—we can figure it out simply using group theory!

The vector operator giving us the non-relativistic contribution to this matrix element changes I_3 by one unit; this is precisely the difference between N and N' as far as isospin and QCD is concerned. In terms of quarks, this is because $V^\mu = \bar{u}\gamma^\mu d$ and each of the quarks changes isospin by $1/2$. Since this is a $\Delta I_3 = 1$ operator, we know that it is proportional to the usual $SU(2)$ raising operator acting on isospin space,

$$V|j, m\rangle \propto \sqrt{(j-m)(j+m+1)}|j, m+1\rangle. \quad (8.7)$$

The overall coefficient can be determined, but the point for us is that *we don't care* since we're just want to take the ratio with $\mu \rightarrow e\nu\bar{\nu}$ anyway. The coefficient above is just a Clebsch-Gordan coefficient and tells us that, finally,

$$\langle N'(p')|V^\mu|N(p)\rangle \propto \sqrt{(j-m)(j+m+1)}(p+p')^\mu, \quad (8.8)$$

where we stress that the j and m are *isospin* numbers. This is an incredible result: we started with some hadronic matrix element that we knew nothing about. Using Lorentz invariance, specific decay channels, and approximate isospin symmetry we were able to get rid of *all* QCD unknowns and end up with expressing everything in terms of a single number that comes from group theory. After taking into account corrections to the approximations we've made, one can extract the value for V_{us} in (8.2).

8.1.2 Neutron β decay

Let us briefly comment on the other measurements of V_{ud} . Consider the free neutron decay process $n \rightarrow pe\bar{\nu}$. While naively this appears to be very similar to the nucleon calculation, the external states are different—the hadronic states are now necessarily spin-1/2 so we can't get away with using the $J_N = J_{N'} = 0$ trick to get rid of the axial contribution. This ends up giving six matrix elements to calculate which nontrivially reduce to

$$G_V = \langle n|V^\mu|p\rangle(p+p')_\mu \quad (8.9)$$

$$G_A = \langle n|A^\mu|p\rangle(p+p')_\mu. \quad (8.10)$$

The relevant quantity turns out to be the ratio of these matrix elements, $g_A \equiv G_V/G_A$. This is a ratio between two hadronic matrix elements with no symmetries forcing any obvious hierarchies between the two terms so that we expect $g_A \sim \mathcal{O}(1)$. Indeed, it turns out that $g_A = 1.27$.

Example 8.2. How does one measure g_A ? We can distinguish the V and A matrix elements by their angular distributions so that the angular dependence of β decay is a probe of g_A .

V_{ud} depends on the neutron lifetime and g_A as

$$|V_{ud}|^2 \sim \tau_n \times (1 + 3g_A^2)(1 + \text{radiative corrections}). \quad (8.11)$$

We can calculate the right-hand side to determine V_{ud} ; it turns out to have roughly the same value and precision as the nuclear matrix element method. This should be very reassuring since this method carries different hadronic assumptions. (Dimensionless?) For two recent reviews, see [51] and [9].

8.1.3 Pion β decay

Finally, for pion β decay, $\pi^+ \rightarrow \pi^0 e \nu$, we can see immediately that a very nice feature is that the hadronic matrix element is between two spin-0 states.

$$\langle \pi^0 | \mathcal{O} | \pi^\pm \rangle. \quad (8.12)$$

The pions live in an $I = 1$ isotriplet and are much simpler objects than nuclei. It even turns out that the calculation of the corrections to the isospin limit are much easier for pions. In fact, this is the cleanest way to measure V_{ud} . Why, then, are we only mentioning this decay as a remark after the nucleon and neutron decays; why isn't this the *best* way to measure $|V_{ud}|$? It is suppressed by phase space! The π^\pm and π^0 are nearly degenerate so that the branching ratio is on the order of 10^{-8} . Experimentally we are limited by the total number of these decays that we measure.

Example 8.3. Note that the neutral pions are *not* difficult to detect experimentally. Due to the axial anomaly $\pi^0 \rightarrow \gamma\gamma$ is the dominant decay (98.8%) and it's easy to tag two photons with an invariant mass of $m_\pi = 135$ MeV.


8.2 Measuring $|V_{us}|$

8.2.1 $K^{0,+} \rightarrow \pi^{-,0} \ell^+ \nu$

Moving on, how would we propose to measure V_{us} ? There's an obvious choice: kaon decay. The dominant semileptonic decay is $K^{0,+} \rightarrow \pi^{-,0} \ell^+ \nu$, with a decay rate of 8.4%. We would like to apply the same slick maneuvers to determine $|V_{us}|$ that we used for $|V_{ud}|$; to do this we need an approximate symmetry to give us a handle on the solution in the unbroken-symmetry limit. Here, again, we have an obvious choice: the generalization of isospin to include the strange quark is just $SU(3)$ flavor. In fact, if we stare at this long enough we see that this process with respect to $SU(3)$ flavor symmetry is really just like pion β decay, $\pi^+ \rightarrow \pi^0 \ell^+ \nu$, with respect to isospin.

However, we do have reasons to be very skeptical. For example, we know that while $SU(2)$ is broken by effects of the size $(m_d - m_u)/\Lambda_{\text{QCD}} \sim 1\%$. Thus, with some exceptions (such as the difference in the proton versus neutron lifetimes), we expect isospin breaking effects to be small corrections. For $SU(3)$ flavor, however, we have a breaking on the scale of $m_s/\Lambda_{\text{QCD}} \sim 20\%$. While this is still a valid symmetry limit to expand about, we now expect the precision to be rather poor compared to the precision of our experiments. There is a second reason why we should be skeptical that our previous bag of tricks will work: in our measurement of $|V_{ud}|$ we assumed that $f(q^2) \approx f(0)$. We justified this because $q^2 \ll \Lambda_{\text{QCD}}^2$. In this case, however, because the kaon and pion masses are appreciably different, the momentum transfer is not small. It seems like this approximation should also be very poor.

The $SU(2)$ subgroups of $SU(3)$. Let's pause for some group theory. We know that $SU(2)$ isospin lives inside of the bigger $SU(3)$ flavor symmetry. In fact, we know that there are other $SU(2)$ s that live inside $SU(3)$. We can see these by looking at the root diagram of the octet, which we reproduce below. Each dotted line corresponds to an $SU(2)$ subgroup of $SU(3)$. We can identify these $SU(2)$ s as the familiar isospin ($u \leftrightarrow d$), v -spin ($u \leftrightarrow s$), and u -spin ($d \leftrightarrow s$).



These were named by Harry Lipkin—author of an excellent representation theory text for physicists [61]—as a pun: ‘I-spin, u -spin, v - (we) all spin for iso-spin.’ While we know that $SU(2)$ isospin is a much better symmetry than $SU(3)$ flavor, both $SU(2)$ u -spin and $SU(2)$ v -spin are as badly broken as $SU(3)$ flavor since they both suffer from a breaking on the order of m_s/Λ_{QCD} . Technically, u -spin has a *small* advantage over v -spin since it is not broken by electromagnetism, but since $\alpha \sim 1\%$, this is a small concession compared to its m_s/Λ_{QCD} breaking. Finally let us remark on the shift in mindset that we’ve been espousing: as theorists we like to think about the symmetry limit in such a way that we only care about the symmetry; this is wrong. The more constructive point of view is to think as *phenomenologists*, where the main aspect that we care about the symmetry limit is how the symmetry is broken.

Now that we’ve vociferously noted our objections to $SU(3)$ flavor due to the size of its breaking and the small $q^2/\Lambda_{\text{QCD}}^2$ limit, we can proceed with the calculation guilt-free and see what happens. We are interested in the two decays,

$$K^0 \rightarrow \pi^- e^+ \nu_e \quad (8.13)$$

$$K^+ \rightarrow \pi^0 e^+ \nu_e. \quad (8.14)$$

Actually, we should be honest: we don’t actually see K^0 s, the neutral kaon mass eigenstates are denoted by $K_{S,L}$ (“K-short” and “K-long”) due to mixing. This is an entire rich and fascinating story in itself, but we’ll leave this for later. If you are very serious, then you should replace K^0 with K_L , but for our current purposes we can be naïve.

What is the difference between these two decays? They have different charges. This comes from the fact that the K^0 decay has a spectator \bar{d} quark while the K^+ decay has a spectator \bar{u} . Thus the two decays are related by *isospin*! Thus to the extent that isospin is approximately true, it is sufficient to determine the form factor for only one of these process since the other will be related by symmetry.

Let start by writing out the matrix element,

$$\langle \pi(p_\pi) | \bar{s} \gamma^\mu u | K(p_K) \rangle = f_+(q^2)(p_\pi + p_K)^\mu + f_-(q^2)q^\mu, \quad (8.15)$$

where we haven’t yet specified K^0 or K^+ decay.

Example 8.4. Relating the kaon decay form factors. We know that K^0 is an isospin $|\frac{1}{2}, -\frac{1}{2}\rangle$ state. This is related to K^+ by

$$J_-|j, m\rangle = \frac{1}{\sqrt{2}}\sqrt{(j+m)(j-m+1)}|j, m-1\rangle. \quad (8.16)$$

Thus we see that

$$\frac{f_+^{K^+}(0)}{f_+^{K^0}(0)} = \frac{1}{\sqrt{2}}. \quad (8.17)$$

What about the other decay constant, $f_-(q^2)$? In (8.6) we invoked the isospin symmetry limit to explain why this term should vanish. In the case of kaon decay, we replace isospin with $SU(3)$ flavor this argument becomes rather fishy: after normalizing f_+ to unity, we would at best be able to say that f_- is on the order of 20%. Fortunately, there's another argument available to allow us to neglect f_- : the contributions of the f_- form factor to semileptonic kaon decays are proportional to the lepton mass. Thus for decays with electrons, these decays are suppressed by factors of m_e/m_K . This is essentially the same chiral suppression we encountered in Problem 7.2, though here we have a one-to-three body decay. The point is that the amplitude associated with this term goes like

$$\mathcal{M}(f_-) \sim f_-(q^2) (p_e + p_\nu)_\mu \bar{u}_e \gamma^\mu u_\nu, \quad (8.18)$$

where we've used $p_\pi - p_K = p_e + p_\nu$. Contracting the Lorentz indices we obtain

$$\mathcal{M}(f_-) \sim f_-(q^2) \bar{u}_e (\not{p}_e + \not{p}_\nu) u_\nu \quad (8.19)$$

$$\sim f_-(q^2) m_e \bar{u}_e u_\nu, \quad (8.20)$$

where we've used the equations of motion for the leptons, $\not{p}_e u(p_e) = m_e u(p_e)$ and $\not{p}_\nu u(p_\nu) = 0$.

Example 8.5. That's a nice argument, why didn't we use it before? Let's recall what just happened: in pion β decay we had a nice symmetry argument for why $f_-(q^2) = 0$. We then moved to kaon decay, and found that this argument is no longer so nice since $SU(3)$ flavor isn't as precise as $SU(2)$ isospin. So instead, we cooked up a different reason why the contribution to the amplitude mediated by this term is suppressed anyway. If this 'dynamical' reason was there all along, why didn't we apply it to pion decay?

At first glance one might say that pions are lighter than kaons, and so the mass suppression shouldn't be as dramatic. However, pions are only a factor of about 4 lighter than kaons, which should not be a big difference compared to the size of m_e . This line of thought, however, is a red herring! The point isn't that $m_e \ll m_K$, but rather that $m_e \ll (m_K - m_\pi)$. The reason why this argument doesn't hold for pion decay is that m_e is not *much* smaller than $(m_{\pi^+} - m_{\pi^0})$.

Example 8.6. What about muons? The next obvious question is whether $f_-(q^2)$ can still be neglected when considering semileptonic decays into muons, e.g. $K \rightarrow \pi\mu\nu$. The answer is, as you would expect, no since $m_\mu \approx 100$ MeV is certainly not very small compared to the meson mass difference. This is why we were very specific in (8.13) and (8.14) to specify the outgoing electron rather than a general lepton. For semileptonic kaon decays into muons, we say that we are "sensitive to $f_-(q^2)$." This is a useful process not so much for measuring $|V_{us}|$, but for checking predictions of the form factors from lattice calculations or chiral Lagrangians.

Now it's time for us to face the music. To get to this point we assumed that $q^2 \ll \Lambda_{\text{QCD}}$ and that $SU(3)$ breaking is small. These manifest themselves into the statements

$$f(q^2) = f(0) \tag{8.21}$$

$$f(0) = 1, \tag{8.22}$$

respectively. While the analogous assumptions were completely sensible for pion decay, we're now well outside the regime where we would expect these to give reasonable results. It turns out, somewhat magically, that these *a priori* poor approximations work out much better than we ever had a right to expect.

First consider $f(q^2) = f(0)$. One can check this by plotting the decay spectrum with respect to energy and observing that that it is a straight line. It 'turns out' that the expansion in q^2 takes the form

$$f(q^2) = f(0) + \lambda q^2 + \dots \tag{8.23}$$

where $\lambda \sim 0.04 \text{ fm}^2$ and the higher order terms are not measured. Why should the coefficient be so small? **Work in progress: need to write this in a dimensionless way! This is probably something like q^2/f_π^2 ?** A not-so-satisfactory hand-waving answer is that to good approximation the kaon and pion are point-line. Since the form factors really probe the structure of the particles, the small coefficient is telling us that the process is not sensitive to any meson substructure. As long as $q^2 \ll \Lambda_{\text{QCD}}$ this is what we would expect; form factors should be very small when one probing something which is 'smaller' than the probe.

Next let us consider the approximation $f(0) = 1$. It turns out that the value we find is

$$f(0) = 0.961 \pm .008. \tag{8.24}$$

Naïvely we would have expected $f(0) = 1 \pm 20\%$, how did we get something which is more like $1 \pm 4\%$? There's a hint here if we play with some numerology: how is 4% related to 20%? It is the square!

$$(20\%)^2 = 4\%. \tag{8.25}$$

Besides writing a trivial numerical identity, this is a big hint: it seems like the actual correction to our symmetry-limit isn't linear (20%), but rather second-order. Thus our expansion about the symmetry limit is somehow slightly different from the naïve expansion that we would have written. Somehow $f(0) = 1$ up to *second-order* in $SU(3)$ breaking.

This has a fancy name, the **Ademollo-Gatto theorem** [4]. The argument doesn't quite work for $f_-(q^2)$, but as we argued above, we don't care since we already have a nice m_e/m_K suppression on this term. In the spirit of this course, we present a heuristic argument²⁰ for why the Ademollo-Gatto result should be intuitive. We start by noting that the $f_+(q^2)$ form factor of the $(p+p')^\mu$ term in the matrix element should be unity for $q^2 = 0$ and in the limit where $SU(3)$ is restored. This is because $f_+(0)$ should be understood to be proportional the overlap of the s and u wavefunctions which approaches unity in the unbroken $SU(3)$ limit. Clearly this is the maximum

²⁰This was beautifully explained to us by Shmuel Nussinov over dinner.

value of such a quantity. The heart of the Ademollo-Gatto theorem—and the reason why f_+ ends up being so well behaved despite the breaking of $SU(3)$ —is that at the deviation of a function expanded about a critical point (such as a maximum) is always *second* order, tautologically. Note that the observation that the form factor should be interpreted as a wavefunction overlap is already something that we met with the ‘decay constant’ of positronium in Example 7.3. A similar theorem, Luke’s theorem [62, 11], exists for the limit of large symmetry breaking where the critical point is a minimum.

Finally, getting back to our main task, we can use the above techniques (and include corrections) to obtain

$$V_{us} = 0.2257(21). \quad (8.26)$$

Ademollo-Gatto from the current algebra. A cute and slightly more rigorous way of proving the Ademollo-Gatto theorem for the vector form factors can be presented in light of the current algebra mentioned in Section 7.4. A very rough sketch of the procedure is to note the commutator of the $SU(3)$ generators

$$[Q^{4+i5}, Q^{4-i5}] = Q^3 + \sqrt{3}Q^8 = Q^{\text{EM}} + Y, \quad (8.27)$$

where $Q^{4\pm i5} = (Q^4 \pm iQ^5)/2$. One can then consider taking the expectation value of both sides with respect to an octet state, from which we obtain

$$(Q^{\text{EM}} + Y)_H = -|\langle H' | Q^{4+i5} | H \rangle|^2 + \sum_m |\langle m | Q^{4-i5} | H \rangle|^2 - \sum_n |\langle n | Q^{4+i5} | H \rangle|^2, \quad (8.28)$$

where H' is another state in the same octet. For example, for $H = \Sigma^-$, $H' = \Lambda$ or Σ^0 . The states $\langle m |$ and $\langle n |$ are not in the same octet. In the appropriate limit, it turns out that $\langle H' | Q^{4+i5} | B \rangle \rightarrow f_1(0)$. Taking such a limit and rearranging terms we obtain

$$f_1(0)^2 = -(Q^{\text{EM}} + Y)_B + \mathcal{O}(\delta^2), \quad (8.29)$$

where $\delta = (M_H - M_{H'})/M_H$ is a parameter for $SU(3)$ breaking. This is further explained in the review article by Cabbibo et al. [15].

8.2.2 $K^+ \rightarrow \pi^+\pi^-e\nu$

Now that we’ve shown how to determine V_{us} , we can go on to see how we can cross-check this value. Recall that we don’t actually *care* about the precise values of these CKM matrix elements! For our purposes we can use the Wolfenstein parameterization for a good estimate and we can look up the value in the PDG when any ‘honest’ precision is required. Our goal is to see how physics is done, and this requires checking and cross checking our parameters through orthogonal experiments to *really* test the Standard Model.

The next $|V_{us}|$ -sensitive process we can look at is the four-body decay

$$K^+ \rightarrow \pi^+\pi^-e\nu. \quad (8.30)$$

Back in ancient times like the 1960s, this decay was known as²¹ $K_{\ell 4}$. Since this is a one-to-two hadron decay with two additional leptons, the whole story becomes more complicated. On the plus side, we may take some experimental solace that the two charged pions are at least easy to identify in a detector.

We may parameterize the matrix elements in terms of three *independent* momenta. We chose these to be the momenta p_{\pm} of the charged pions and k of the kaon so that the vector and axial matrix elements are

$$\langle \pi^+ \pi^- | A^\mu | K \rangle = C [f_1(p_+ + p_-)^\mu + f_2(p_+ - p_-)^\mu + f_3(k - p_+ - p_-)^\mu] \quad (8.31)$$

$$\langle \pi^+ \pi^- | V^\mu | K \rangle = C' g \epsilon^{\mu\alpha\beta\gamma} k_\alpha p_{+\beta} p_{-\gamma}. \quad (8.32)$$

The axial matrix element is a sum of three terms that are linear in linearly independent combinations of the basis momenta. The vector matrix element, on the other hand, is composed of a contraction of all three momenta with the totally antisymmetric $\epsilon^{\mu\alpha\beta\gamma}$ tensor as required by parity. This is clear since the vector matrix element has the opposite parity of the axial matrix element and the only way to construct such a one-Lorentz-index quantity out of four-vectors necessarily requires the intrinsic parity of the ϵ tensor.

Example 8.7. We can immediately identify one of the f_i factors to discard, can you see which one? By the same argument that we've seen several times now, we know that the f_3 form factor is proportional to the electron mass upon contraction with the leptonic part of the full matrix element.

The prefactors C and C' are overall normalizations

$$C = -\frac{i}{f_\pi} \quad (8.33)$$

$$C' = \frac{1}{4\pi^2 f_\pi^3} \quad (8.34)$$

which we pull out to normalize $f_1 = f_2 = g = 1$. The factor of $4\pi^2$ is very familiar—in fact, it's *almost* ‘dimensionful’ since it carries dimension of loop number. Indeed, this really is a loop effect since the vector matrix element is related by $SU(3)$ to the matrix element for $\gamma \rightarrow 3\pi$ by an $SU(3)$ flavor rotation. Note that $\gamma \rightarrow 3\pi$ is of course a nonsensical physical amplitude, but it is a nonzero matrix element coming from the axial anomaly.

The normalization of f_1 and f_2 is more subtle and comes from the current algebra techniques mentioned in Section 7.4 (they are *not* simply Clebsch-Gordan coefficients). These come from soft pion techniques where one uses LSZ reduction and current commutation relations to relate the amplitude for a soft pion emission to an amplitude without the pion. In the case of $K \rightarrow \pi\pi e\nu$ there are *two* pions which one may take to be soft. One limit gives $f_1 = f_2$ while the other gives $f_1 + f_2 = 2C$. Note that these relations hold in the limit of *soft* pions, i.e. the limit $p_{\pm} \rightarrow 0$. In principle f_1 and f_2 are functions of *three* dynamical scalars²²: $k \cdot p_+$, $k \cdot p_-$, and $p_+ \cdot p_-$. However,

²¹Similarly, $K \rightarrow \pi e\nu$ was known as $K_{\ell 3}$. Naming conventions in physics have gotten more colorful, but aren't necessarily any better. At least these days we like to believe that we have more interesting objects to name than scattering amplitudes.


²²To be precise, in the soft pion manipulations one takes the pion off shell so that p_{\pm}^2 becomes a ‘dynamical’ variable as well.

in the soft pion limit these are all small. Experiments confirm that there is no strong dependence on the external momenta.

The mystery of f_3 . One of the mysteries of the ancient world was the failure of naïve current algebra techniques to sensibly predict the f_3 form factor (even though we already know f_3 is m_e/m_K suppressed). While the two constraints on f_1 and f_2 were consistent, the f_3 constraints coming from taking each pion to its soft limit gave what appeared to be contradictory predictions:

$$f_3 = 0 = f_+ + f_-, \quad (8.35)$$

where the f_{\pm} come from the $K \rightarrow \pi e \nu$ decays. In other words, f_3 seems to vary strongly with the momenta. In 1966 Weinberg solved the puzzle by identifying the p_{\pm} dependence of f_3 as coming from a kaon pole coming from the following diagram, [90]



This diagram *only* contributes to f_3 so that one ultimately finds

$$f_3 = (f_+ + f_-) \frac{k \cdot p_-}{k \cdot (p_- + p_+)}, \quad (8.36)$$

which indeed gives the correct interpolating behavior between the two single-soft pion limits.

Having set up the problem thusly, one can go on to do the same analysis that we presented for three body semileptonic kaon decay ($K_{\ell 3}$). The procedure is the same, though there are more coefficients and the entire ordeal is generally more complicated and not as fun. The punchline is that one can indeed make an independent measurement to confirm the value of the previous method. As we said, in this section we will care less about the actual values of the Standard Model parameters but rather we will focus on the way in which they are measured and cross-checked. (These are ultimately experimental matters but, as theorists, looking up a number in the PDG is our way of ‘doing an experiment.’)

8.2.3 Hyperon decay

Moving on from the mesons, another way to measure $|V_{us}|$ is to look at the decays of mesons containing strange quarks and no heavier quarks, the **hyperons**. These include the Λ , Σ , Ξ , and Ω baryons. Just as we looked at neutron β decay $n \rightarrow p e \nu$ to measure $|V_{ud}|$, we can also measure $\Lambda \rightarrow p e \nu$ in the analogous way to measure $|V_{us}|$.

Example 8.8. Λ lifetime versus n lifetime. Why does the Λ have a much smaller lifetime than the neutron? If you answer that $|V_{us}| \ll |V_{ud}|$, then *wake up*. This is a trick question. The amplitude for the decay is slightly smaller because $|V_{us}| = \lambda|V_{ud}|$, where $\lambda \approx 0.2$ is the Wolfenstein expansion parameter. The correct answer is that n decay is very phase space suppressed.

The upshot of looking at baryons is that we have many more baryons to work with and that we can measure all of them, at least in principle. The $SU(3)$ baryon analysis depends on two parameters that encode the two irreducible matrix elements for an octet current between two octets. In particular, given two $SU(3)$ flavor octet states B_n and B_m and some octet operator \mathcal{O}_k ,

$$\langle B_n | \mathcal{O}_k | B_m \rangle = F_{\mathcal{O}} f_{knm} + D_{\mathcal{O}} d_{knm}, \quad (8.37)$$

where f_{knm} are the $SU(3)$ structure constants and d_{knm} are group theory factors defined by

$$\{\lambda^k, \lambda^n\} = 2\delta_{kn} + 2d_{knm}. \quad (8.38)$$

This is just the same program of defining form factors to parameterize our ignorance about QCD matrix elements. Note that this expansion implicitly includes the assumption of $SU(3)$ flavor, thus we are neglecting the mass difference between the light quarks.

Thirty years ago when you say ‘D term’ and ‘F term,’ everybody knew what you were referring to. You were referring to the $SU(3)$ invariants in hyperon decay. And now if you say ‘D term’ or ‘F term,’ everybody thinks you’re talking about supersymmetry. Why do they use the same name? Why would they use D and F? Everything comes from the 60s –Yuval

So this is a pretty good deal, we have two form factors in the $SU(3)$ limit and several baryons to make measurements with. One can again turn the crank by following through the calculations and then doing the experiments. There are lots of tricks one can further use to keep a handle on the breaking terms, and at the end of the day one finds that one can measure $|V_{us}|$ with the same precision as the kaons. For a nice review of the current state-of-the-art for semileptonic hyperon decays, see the review by Cabibbo, Swallow, and Winston [15].

Until about 10 years ago, there was a 3σ discrepancy in these measurements, but it has since disappeared. To quote Yuval,

3σ s disappear. Most of them. It’s an amazing fact of life. Only 1% disappear because of statistics, 99% disappear because of systematics. 0.0000007 survive as real effects.

8.2.4 $\tau \rightarrow K^- \nu$

Yet another measurement of $|V_{us}|$ comes from the decay of the τ . Recall that the τ is so heavy that its decays prefer to hadronize. In this sense, the τ is *almost* a hadron. Our strategy is to divide out the hadronic uncertainties by comparing the ratio of two decay rats,

$$\tau^- \rightarrow K^- \nu_\tau \qquad \tau^- \rightarrow \pi^- \nu. \quad (8.39)$$

In the limit of $SU(3)$ flavor, the ratio simply gives

$$R = \frac{\Gamma(\tau^- \rightarrow K^- \nu)}{\Gamma(\tau^- \rightarrow \pi^- \nu)} \stackrel{SU(3)}{=} \frac{|V_{us}|^2}{|V_{ud}|^2}. \quad (8.40)$$

This is too rough for us, let us at least quantify the dominant corrections from the $SU(3)$ limit. One correction is phase space, since the K has mass approximately equal to 500 MeV while the pions are around 140 MeV. Yuval also interjects another mnemonic,

What is the mass of the τ ? 1,777 MeV. That's 1 and then 777, like the casino.

On top of this, we also have a correction from hadronic effects that we know to high precision. What is the hadronic matrix element here? It's just K/π goes to vacuum, thus there are factors of decay constants! Hence another breaking effect is f_K versus f_π . We end up with a revised ratio,

$$R = \frac{\Gamma(\tau^- \rightarrow K^- \nu)}{\Gamma(\tau^- \rightarrow \pi^- \nu)} = \frac{|V_{us}|^2}{|V_{ud}|^2} \times (\text{Phase Space}) \left| \frac{f_K}{f_\pi} \right|^2. \quad (8.41)$$

The powers of two should be clear since we are comparing branching ratios, not amplitudes. Let us make a remark about the measurement of f_K and f_π . To use the above method to determine $|V_{us}|$, we need a measurement of f_K that is somehow ‘orthogonal’ to the $f_K|V_{us}|$ combination above. One would want to measure f_K through leptonic decays such as $K \rightarrow \mu\nu$, but this gives something proportional to the same combination $f_K|V_{us}|$. This is also true for the pions, which is why we focused on β decay rather than $\pi \rightarrow \ell\nu$ to measure $|V_{ud}|$. It turns out, then, that for f_K we really depend on a lattice QCD calculation of f_K as an independent input.

The best way to use the τ is to look at **inclusive** decay. We’ll avoid the details, but the main idea is that the τ is heavy enough that at m_τ one can really treat QCD perturbatively, $\alpha_s(m_\tau) \sim 0.3$. You can calculate how many light and heavy quarks you produce and really trust your perturbative calculation.

The best way to use τ decay: **inclusive** tau decay. The tau is heavy enough such that when it decays we can really treat QCD perturbatively $\alpha_s(m_\tau) = 0.3$. We can calculate the number of light and heavy quarks. The inclusive decay has no form factors, which makes it theoretically very nice, but currently our limiting factor is the number of experimentally observed τ decays.

8.3 Measuring $|V_{cs}|$

The values for $|V_{us}|$ and $|V_{cd}|$ are not measured very well. The errors are much much larger than those for V_{ud} and V_{us} by factors of ten. The reason is clear: the charm quark mass lives in the barren no-man’s-land where it’s too heavy for $SU(3)$ flavor symmetry but not heavy enough to treat in the heavy quark limit. Or, in the timeless words of Briney Spears,

*Feels like I'm caught in the middle
That's when I realized—
I'm not a girl, Not yet a woman.*

Later in the song Spears sings the line “*I’m not a girl, there’s no need to protect me...*,” from, I can only assume, the large corrections associated with quarks which cannot be reasonably treated in either the $SU(3)$ flavor or heavy quark limit.

It turns out, actually, that the best way to work with charm decays is to use lattice QCD techniques. We will present a brief introduction to lattice techniques in Section 10. For now, let us simply state that lattice QCD is ideal for effects that are on the order of Λ_{QCD} and the charm mass is about a factor of few larger than this. The approach is to latticize spacetime so that one immediately has UV and IR cutoffs associated with the lattice spacing and the lattice size, respectively. Thus one expects that such techniques should not be able to capture effects that are much smaller or larger than Λ_{QCD} due to these cutoffs. In other words, the QCD effects that are precisely *at* Λ_{QCD} are those which should be best captured by lattice techniques.

8.3.1 $D \rightarrow K e \nu$

So what process do we want to look at to probe $|V_{cs}|$? The obvious choice is $D \rightarrow K e \nu$. Just as the case for $K \rightarrow \pi e \nu$ we pick the electron final state in favor of the muon final state since this gives us one less form factor to worry about; f_- gives an effect that is suppressed by m_ℓ/m_D . The remaining problem, as we mentioned above, is that we really have no idea what $f_+(0)$ is. The lattice gives us $f_+(0) = 0.7$, but we even don’t know if $q^2 = 0$ is a good approximation. Note that this has indeed shifted quite a bit from the previous symmetry limit where $f_+(0) = 1$. Going through the same rigamarole, we end up with a measurement

$$V_{cs} = 0.98 \pm 0.01_{\text{exp}} \pm 0.10_{\text{thy}} \quad (8.42)$$

This is indeed approximately equal to one. There’s a pretty small experimental error, but by comparison the theory error is huge. Compare this to the case for V_{ud} where the error was in the fourth digit! Here we give our best effort with lattice QCD, but we’re still left with 10% error. This is not systematic error, this is just theory error.

8.3.2 $D_s^+ \rightarrow \ell^+ \nu$

Now we can hope that alternate methods to measure $|V_{cs}|$ might give more insight. In that spirit, let’s now look at the pure leptonic decays of the D_s . It should be clear why we use D_s , since this is the meson which will give a factor of $|V_{cs}|$ when it decays. We are thus interested in the decays

$$D_s^+ \rightarrow \ell^+ \nu, \quad (8.43)$$

where the ℓ can be any of e, μ, τ . Each choice has its own pros and cons. The electron has the smallest branching ratio due to chirality suppression, so we can forget about it. Now we have to choose between the muon and τ . On the one hand, chirality suppression prefers the τ decay, but on the other hand the τ is constrained by phase space since $m_\tau \lesssim M_{D_s}$, where $M_{D_s} = 1970$ MeV. The muon gets more chiral suppression, but enjoys a much bigger phase space. Theoretically these effects even out and both are useful. Experimentally, on the other hand, the τ is a pain to measure since it hadronizes while the muon is a very clean signal. Enterprising experimentalists, however, have been able to use both modes with success.

Before we can go on to extract $|V_{cs}|$, however, we need to know f_{D_s} . This is comforting, since if there's *one* thing that lattice people really know how to calculate accurately, it should be f_{D_s} : the c is not too heavy, and the s is not too light relative to Λ_{QCD} . Indeed, the simplest things for the lattice to calculate are masses and decay constants, and f_{D_s} should be ideal—it's a number that lattice techniques were *born* to predict. Because there's no such thing as an 'easy' lattice calculation, there was of course a 3σ disagreement. Like nearly all 3σ effects, it ended up being pushed down to 1σ and forgotten; but in this case the reconciliation was only done one month prior to this lecture and was performed by the dean of Cornell's College of Arts and Sciences [66].

The lattice gives a prediction for the ratio of f_{D_s}/f_D , from which can reconstruct and even better measurement for $|V_{cs}|$ than we had for the semileptonic decay,

$$|V_{cs}| = 1.030 \pm 0.038, \quad (8.44)$$

where the error is basically theoretical. Does this look strange to you? $|V_{cs}| > 1$, but clearly a 'sensible' this is within error²³. Such a thing happens all the time—we measure something which is supposed to be the sine of an angle, but it ends up also being larger than unity. The point is that it boils down to how we define what we measure. Ultimately we're not necessarily measuring that sine, it is usually the case that we are measuring this indirectly with all sorts of errors propagating.

8.3.3 $W \rightarrow c\bar{s}$

There is yet another way to measure $|V_{cs}|$ is the W boson decay channel $W \rightarrow c\bar{s}$. This is the most direct measurement of $|V_{cs}|$. Can we actually measure this?

Example 8.9. Quick calculation: the branching ratio of $W \rightarrow c\bar{s}$. Neglecting higher order terms in the Wolfenstein expansion parameter λ , a W in its rest frame can decay into any of the three lepton–neutrino pairs or any of the two *active* generations of quarks, u - d or c - s . Note that it does not have the rest energy to decay into third generation quarks. The quarks have an additional factor of three due to color so that there are nine possible decays. We can approximate each as having the same rate so that the second generation quark decays occur 1/3 of the time. This is an appreciable rate and so we expect it to be tractable experimentally.

There's an immediate problem. How are we going to actually *see* $W \rightarrow c\bar{s}$? What we really see is W going to two jets. And as we know from West Side Story,

*When you're a Jet,
you're a Jet all the way
From your first cigarette
To your last dying day.*

The problem is that unlike charmed—and especially bottom—jets we cannot really tag a strange jet. (Tony was a strange jet because he fell in love with Maria, the sister of the leader of the rival gang, the Sharks.)

Now an experimental aside. We can identify jets containing a b or c quark by measuring displaced vertices. This boils down to the lifetime of the b and c mesons.

²³To be frank, even if it were not within error, then we would say that at some point someone underestimated the error.

Example 8.10. The lifetime of the B^0 and D^0 mesons. Which has a longer lifetime, the charmed or the bottom mesons? We can see that the D^0 decay goes like $|V_{cs}|^2$ while the bottom goes like $|V_{cb}|^2$. We know from our Wolfenstein parameterization that the latter is suppressed by $(\lambda^2)^2$. On the other hand, we should also account for the phase space. From naïve dimensional analysis (NDA) we can write down

$$\Gamma \sim m^n G_F^2, \quad (8.45)$$

where m is the meson mass, n is some integer, and G_F contains the necessary M_W^{-2} factors from the W propagator. In order to get the correct overall dimension we must have $n = 5$, so that the ratio of the two branching fractions scale like

$$R = \frac{\Gamma(b)}{\Gamma(c)} = \left(\frac{m_B}{m_D}\right)^5 \lambda^4. \quad (8.46)$$

It turns out that the phase space factors nearly cancels the Wolfenstein suppression leaving $R \sim 0.3$, so that they lifetimes are actually comparable. Note that since $\tau \sim 1/\Gamma$, this means that the B has a longer lifetime by a factor of about three.

What we end up with are B mesons which give a nice, observable displaced vertex while the D meson gives a slightly harder-to-measure displaced vertex. For W decays the b quark isn't an active generation, so *any* displaced vertex is a signal for charm.

The $W \rightarrow c\bar{s}$ measurement, as done by the Delphi experiment, one of the four LEP experiments that ran at CERN. The measurement yielded,

$$V_{cs} = 0.94 \pm 0.29_{\text{stat}} \pm 0.13_{\text{sys}}. \quad (8.47)$$

This number is constrained by the number of W bosons that they measured.

Remaining questions: *How do we get around the problem of not being able to tag the strange quark? Why was Delphi the only experiment to do this measurement? (I think we answered these but I forgot.)*

8.3.4 $W \rightarrow$ hadrons

Another $|V_{cs}|$ measurement which all of the LEP experiments were able to was to look at $W \rightarrow$ hadrons. We can normalize this to the rate for $W \rightarrow$ leptons,

$$R = \frac{\Gamma(W \rightarrow h)}{\Gamma(W \rightarrow \ell)} = N_c^{\text{eff}} \sum_{u,c,d,s,b} (V_{ij})^2 \times (\text{Phase Space}) \quad (8.48)$$

Using the same approximation that we did in Example 8.9, we find that $R \approx 2$. We can write out the branching ratio of the W to leptons as

$$\frac{1}{\text{Br}(W \rightarrow \ell)} = \frac{\Gamma(W \rightarrow h)\Gamma(W \rightarrow \ell)}{\Gamma(W \rightarrow \ell)} = R + 1. \quad (8.49)$$

Thus measuring the branching ratio to leptons very precisely gives R so that the theoretical situation is very nice. In order to measure this precisely we need to include QCD corrections, which we are able to do since we are in a regime where we are able to do perturbative QCD. For example, we need to write down the **effective number of colors** $N_c^{\text{eff}} \approx 3 + \mathcal{O}(\alpha_s) \approx 3.3$.

Please provide more information about this!! *This is a very strange thing to define and seems to violate $SU(3)_{\text{color}}$. Please provide more information on this—references?*

The point is that this measurement ultimately gives the sum

$$|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 + |V_{cd}|^2 + |V_{cs}|^2 + |V_{cb}|^2. \quad (8.50)$$

Of these, the only terms that are not small are $|V_{ud}|$ and $|V_{cs}|$. Since the former value is well known to very good precision, we can take this to be a measurement of $|V_{cs}|$. Experimentally the number is quite precise

$$\sum_{u,c,d,s,b} |V_{ij}|^2 \approx 2.002 \pm 0.027. \quad (8.51)$$

If you take it as a measurement of V_{cs} , then this measurement has the best error, but clearly is an indirect measurement since it really probes the above sum.

8.4 Measuring $|V_{cd}|$

Now let's quickly see how to measure $|V_{cd}|$.

8.4.1 Peculiarities of D^* decays.

Before reviewing the D meson decays that tell us about $|V_{cd}|$, let us pause to make a remark about an interesting fact about the vector meson D^* . How does this particle decay? It should certainly decay to the pseudoscalar, $D^* \rightarrow D(\text{stuff})$. The 'stuff' can be either a pion or a photon,

$$D^* \rightarrow D\pi \quad (8.52)$$


$$D^* \rightarrow D\gamma \quad (8.53)$$

Which one dominates? At first glance, we would say that the pion emission dominates since it is a strong process while the photon emission is electromagnetic. In other words, the D^* is a resonance in the sense of Section 3.1. The strong decay should then dominate by a factor of $\alpha_s/\alpha \gg 1$. This is indeed what happens the kaons, where $K^* \rightarrow K\pi$ is the dominant decay mode.

However, there is an amazing coincidence in nature that the mass difference between the D^* and D is just a *little* larger than the pion mass

$$(m_{D^*}^2 - m_D^2) \gtrsim m_\pi^2, \quad (8.54)$$

such that the phase space suppression of $D^* \rightarrow D\pi$ is almost the same as the ratio of coupling constants so that the rate for is only a factor of two larger than $D^* \rightarrow D\gamma$. The ratio α_s/α is very large, but recall that the phase spaces closes extremely quickly. Heuristically,



so that the decay rate lies roughly halfway down the steep precipice because the mass splitting between the vector and pseudoscalar D mesons, $\Delta m_D^2 \equiv m_{D^*}^2 - m_D^2$, is finely tuned relative to the pion mass. There's no deep reason for this, it is a complete accident.

We can use isospin symmetry to relate this decay to trivially relate this process to

$$D^{*0} \rightarrow D^+ \pi^- \quad (8.55)$$

$$D^{*0} \rightarrow D^0 \pi^0 \quad (8.56)$$

$$D^{*+} \rightarrow D^0 \pi^+ \quad (8.57)$$

$$D^{*+} \rightarrow D^+ \pi^0. \quad (8.58)$$

The relations between these processes in the isospin limit is broken by the mass difference between D^* and D^0 , which ordinarily is negligible. However, because of the delicate phase space tuning, the isospin breaking effects are order one. For example, $D^{*0} \rightarrow D^+ \pi^-$ is completely forbidden because the isospin-breaking mass differences barely pushes it out of the kinematically allowed region. On the other hand, the D^{*+} decays are pushed further into the kinematically allowed region and away from the steep drop so that pion decays indeed dominate over the photon decay: $\text{Br}(D^{*+} \rightarrow D^0 \pi^+) \approx 67.7\%$ and $\text{Br}(D^{*+} \rightarrow D^+ \pi^0) \approx 30.7\%$. The upshot with these decay modes is that they can teach us a lot about the strong interaction since we can reliably calculate the photon decays.


Now let's pause to remind ourselves that the story of isospin splitting is something that we already know very well from our physics childhood: it's simply hyperfine structure of the hydrogen atom²⁴ with 'hydrogen' replaced by 'meson'. The mass difference between the D^* and D is just the analog of hyperfine splitting due to the magnetic interaction between the proton and electron in hydrogen. The pseudoscalar D is the singlet and the vector D^* is triplet of the spin-spin interaction.

Example 8.11. B^* decay. Based on the analogy to the hyperfine splitting of hydrogen, what is the dominant B^* decay? Recall in the hydrogen atom that the spin-spin coupling goes like $1/m_p$ due to the dipole moment of the proton. Thus we expect the splitting to go like $1/m_{B^*}$. Given that the B meson masses are roughly a factor of three larger the D masses—coming from the relative c and b quark masses—we expect the $\Delta m_B \equiv m_{B^*} - m_B$ splitting should be a factor of three smaller than the Δm_D splitting. Since $\Delta m_D \sim m_\pi \sim 140$ MeV, we expect $\Delta m_B \approx 50$ MeV, which is smaller than the pion mass and hence the dominant B^* decay is $B^* \rightarrow B\gamma$.


²⁴To zeroth order all of physics reduces to the harmonic oscillator. The first order correction to this statement is the hydrogen atom. You should know everything about the hydrogen atom.

8.4.2 $D^\pm \rightarrow \ell\nu$

A natural process to start with is $D^- \rightarrow \ell\nu$. This is just the *same* semileptonic pseudoscalar meson decay that we've been used many times now. Let's repeat the same spiel. We can again ask which particular lepton we should consider. We know that there is a chiral suppression, so the electron branching ratio will be very small. This leaves us with the muon and tau. The tau, however, is hard to measure experimentally and suffers from phase space suppression. The D^\pm are made up of a d and a c quark so that this probes V_{cd} . The problem, of course, is that the factor of $|V_{cd}|$ always comes along with the decay constant, which we have to get from the lattice. Combining the phase space and chirality (m_μ) suppression, this process is very small. Another way to say this is that as its mass gets heavier, it becomes easier for the heavy quark to decay by itself without the contact interaction with the light quark. The only way for a pion to decay is through an interaction between its valence quarks and the W , but the charm quark in the D can decay by itself with the other quark simply acting as a spectator. The natural mass scale for charm decay is m_c . The decay factor, $f_D \sim \Lambda_{\text{QCD}}$. When the mesons are light, $m_{\text{quark}}f \sim 1$. However, as the quark mass becomes heavier than the QCD scale, the spectator decays are enhanced since the rate goes like $\Gamma \sim m^5$. In particular, instead of decays of the form



one has to consider



8.4.3 $D \rightarrow \pi\ell\nu$

The best way to determine V_{cd} is through the semileptonic decays,

$$D^\pm \rightarrow \pi^0 \ell^\pm \nu \quad (8.59)$$

$$D^0 \rightarrow \pi^\pm \ell^\mp \nu. \quad (8.60)$$

The two decays are related by isospin and so give the same information. Isospin breaking effects are small since there is no phase space miracle, unlike the D^* decay we mentioned above. There is still a problem: while the D^\pm and D^0 are related by isospin, the D s and π s are not related by any symmetry. This is troubling since it takes away one more symmetry limit that we had to simplify our analysis. Compare this to the kaon decay $K \rightarrow \pi\ell\nu$ where we had $SU(3)$ flavor to relate the left- and right-hand side of the decay equation.

Again drawing from our bag of tricks, we can still write down a ratio of independent decay rates in which the final states are related by $SU(3)$ flavor,

$$\frac{\Gamma(D \rightarrow K \ell \nu)}{\Gamma(D \rightarrow \pi \ell \nu)} \stackrel{SU(3)}{=} \frac{|V_{cs}|^2}{|V_{cd}|^2}. \quad (8.61)$$

We can estimate the error from phase space in the usual way. The correction for the form factor f_+ is measured from the lattice. The important point about the lattice measurement is that this is done by plotting $f_+(0)$ over an unphysical mass strange quark mass \tilde{m}_s and then extrapolating to the physical value. There are many papers on whether one should use linear versus quadratic extrapolations (giving a systematic error), but nice thing is that the *ratio* of these values between $\tilde{m}_s = m_s$ and $\tilde{m}_s = m_c$ is much less sensitive to this error. Of course, one could have avoided using the ratio altogether by directly calculating $D \rightarrow \pi \ell \nu$ using the lattice from the very beginning. In this case one has to extrapolate the mass of the kaon down to the pion mass scale, but this is within the regime in which lattice techniques are effective.

8.4.4 Neutrino scattering

A surprising method for measuring $|V_{cd}|$ which claims to have the smallest uncertainty is neutrino scattering against a fixed target. We know that the fixed target is composed of matter: u and d quarks. However, as any LHC physicist will tell you, it is also composed of sea quarks²⁵. In particular, we expect the light (u, d, s) quarks to pop out of the vacuum with some parton distribution function. This is a huge hadronic uncertainty.

For now let's be naïve and only assume u and d quarks in the target. We can tag a charm final state by looking for a pair of muons since this must come from the process

$$\nu + d \rightarrow \mu^- + c \rightarrow \mu^- + \mu^+ + \nu. \quad (8.62)$$

The cross section goes like $\sigma \propto |V_{cd}|^2$ and one can set cuts for hard 30 GeV muon to isolate the charm decays against backgrounds with $\nu + u \rightarrow \mu^+ + s$. (In other words we take the data as a function of Q^2 and perform nice cuts.) What is nice about this is that we are looking at *inclusive* decays, unlike our previous exclusive decays. We don't care about how the remnant of the charm hadronizes, we only tag on the muons. This gives us a different set of hadronic uncertainties which we will discuss later.

The problem, however, is $\nu + s \rightarrow c + \mu$ coming from interactions with strange sea quarks. Note that the ratio of these cross sections is

$$\frac{|V_{cs}|^2}{|V_{cd}|^2} \sim 20, \quad (8.63)$$

i.e. the Cabbibo angle squared is 5%. The background rate is twenty times larger, but is suppressed by the density of s versus d in the target. This latter suppression is only on the order of hundreds, not billions and so limits our accuracy.

²⁵When typing up lecture notes, the phrase 'sea quark' is very problematic. I suspect this is why this phrase has fallen out of favor.

One thing we can do is to also look at anti-neutrino beams while tagging on a single muon,

$$\bar{\nu} + u \rightarrow \mu + d. \tag{8.64}$$

The sea quark background here comes from a sea c . Fortunately, the c quark is not much of a sea quark because it's heavy. (Can you imagine what it was like typing this?) Thus this anti-neutrino process effectively only measures the constituent quarks so that this cross section can be used to subtract the sea s contribution in the ν scattering.

The hadronic physics involved in this measurement is completely orthogonal to that of the semileptonic $D \rightarrow \pi$ decays. Having these two processes agree gives much more confidence in the result.

9 Intermission: Effective Field Theory

This material is based on a guest lecture by Flip Tanedo and is somewhat out of the main narrative of the course. It should be treated as an intermission in which some useful background information is presented idiosyncratically to better support subsequent lectures.

Our main goals in this section is to introduce the ‘philosophy’ of effective field theory, go through an illustrative example, and highlight a few topics which seem to be glossed over by most QFT textbooks. There is a very long list of very good review articles on these topics, a few that I have found to be very influential are Hollowood [52], Manohar [63], Burgess [13], Georgi [36], Polchinski [73], Rothstein [78], Skiba [86], Strassler [88], and the book by Manohar and Wise [64].

Why is the sky blue? (From Manohar [63].) The game-show answer is a phenomenon called Rayleigh scattering. The idea behind this phrase is that the scattering cross section of light against the atmosphere depends on frequency in some positive power so that higher-frequency (e.g. blue) light dominates in these interactions. This is high-school physics. To get to university physics, we should write down an effective theory. Let us describe the nitrogen molecules in the atmosphere by some field Ψ . This field has the usual non-relativistic Schrodinger kinetic Lagrangian but also interacts with photons. The interactions of interest are those at low energies such that the photon does not probe the Nitrogen substructure—it is non-ionizing and doesn't otherwise change the state of the Nitrogen molecule. The interaction Lagrangian thus looks like

$$\mathcal{L}_{\text{int.}} = \frac{c}{\Lambda^3} \bar{\Psi} \Psi F_{\mu\nu} F^{\mu\nu} \tag{9.1}$$

$$= \frac{c}{\Lambda^3} \bar{\Psi} \Psi (E^2 - B^2). \tag{9.2}$$

Here we've just written out the lowest-order Lorentz- and gauge-invariant coupling between the Ψ and the photon. Since we know that the electric and magnetic fields are gradients, they are proportional to the frequency $E \sim \omega$. Thus the amplitude for scattering induced by this term goes like ω^2 , from which we deduce that the cross section goes like $\sigma \sim \omega^4$. This is indeed the famous frequency dependence in Rayleigh scattering.

9.1 EFT is not a dirty word

Thus far we’ve made significant progress in our study of flavor physics. We’ve come to understand the flavor structure of the Standard Model and have slogged through the use of approximate flavor symmetries to allow us to subject elements of the CKM matrix to multiple experimental cross-checks. Underlying everything we have done, however, is the notion of **effective field theory** (EFT).

In words, EFT is the separation of ‘relevant’ infrared degrees of freedom from ‘irrelevant’ UV information. The choice of words is certainly loaded from what we know about the renormalization flow of (ir)relevant operators, but the nomenclature is also unfortunate since it suggests that the UV information is somehow being discarded. This is wrong and leads to a somewhat pejorative connotation when we would like to debunk. To start, let’s provide a definition in terms of the following equation,

$$\mathcal{L}_{\text{EFT}} = \sum_i C_i \mathcal{O}_i. \quad (9.3)$$

An effective field theory is one where we have **factorized** information about possibly-but-not-necessarily unknown UV physics from low-energy ‘active’ degrees of freedom. The UV information lives in C_i , the **Wilson coefficient**, which is simply a number—a coupling which generally depends on the scale at which it is probed. The IR information lives in \mathcal{O}_i , the **effective operator**. These are composed of all of the ‘low energy’ excitations that are physically accessible at the scales where the EFT is valid.

We’ll formalize these notions as we get to use them more, but let us highlight a few important points:

- We have not “thrown away” *any* information, the point is that we’ve just separated information into C_i or \mathcal{O}_i . (It is true, however, that the map sending UV information into C_i is not invertible.)
- The sum over Wilson coefficient should run over all allowed operators, including non-renormalizable operators (unlike our model-building rules in Section 2).
- Implicitly we have introduced a scale Λ , below which the EFT is ‘effective.’ In other words, we have chosen an energy scale so that physics above that scale goes into the C_i while physics below that scale goes into the \mathcal{O}_i .
- By construction, the EFT is only useful (but we’ll see it’s very useful) for answering questions about physics below Λ .
- This scale also serves to balance the dimension of the effective operators \mathcal{O}_i so that very non-renormalizable (large dimension) operators are suppressed by large powers of $1/\Lambda$.
- The UV physics that we factorize into the Wilson coefficients need not be known.

There is a very *naïve* and almost dismissive belief that effective theories are somehow lesser than ‘fundamental’ theories. This belief is somehow planted deep into our consciousness as physics

students. We can approximate the gravitational acceleration as constant, but only until we start thinking about *more interesting* problems involving orbital motion. We major *physics* rather than chemistry because the latter is just an application of the former, therefore the former must be more pure²⁶. We learn that physics is about Taylor expansions and approximations, which can—*at best*—reproduce ‘full’ calculations. All of these things leave the impression that something which is ‘effective’ is just a poor man’s version of the *correct* tool.

It is important at an ideological level to address this bias and to explain the value of effective field theory in particle physics. After all, one of the great paradigm shifts in our field came from Ken Wilson, who taught us that *all* models of nature should be understood as effective theories. There are three major ways in which we use EFT.

1. **Bottom-up physics**, or “phenomenology” (another word that is sometimes used pejoratively). Here one does *not* know the physics that UV completes one’s EFT, but the Wilson coefficients contain dimensionful factors which point to the scale by which the EFT must be completed. The canonical example is Fermi effective theory, where


$$\mathcal{L} = -\frac{c}{\Lambda^2} J_\mu J^{\dagger\mu}, \tag{9.4}$$

writing $J_\mu = \bar{\Psi}\gamma^\mu(1 - \gamma^5)\Psi$ and c is an $\mathcal{O}(1)$ number. Experiment tells us that we should expect ‘new physics’ by the scale $\Lambda \sim G_F^{-1/2} \approx 300$ GeV. Of course, we now know that the W boson appears at 80 GeV and ‘UV completes’ this theory; as is often the case Nature saves the day with some wiggle room left to spare.

Note that *experimentally* this is where we are in particle physics: we have discovered and measured the properties of fermions and *massive* gauge bosons. We can write down a Lagrangian for the observed particles containing terms such as $\Delta L = \frac{1}{2}m_W W^\mu W_\mu$. Clearly such a Lagrangian suffers from non-unitarity at high scales and is *not* the Standard Model, in fact, it is a non-linear sigma model. The Standard Model one possible UV completion that introduces a Higgs turn this theory into a *linear* sigma model. For more on this point of view, see the excellent lectures by Nima Arkani-Hamed on low energy supersymmetry [6].

2. **Simplify calculations**. On the other hand, it may be the case that we know the UV physics perfectly well. For example, beyond the Standard Model (BSM) model builders focus on the structure of physics at the TeV scale.

²⁶<http://xkcd.com/435/>



TeV-scale model building is a great enterprise, but ultimately all respectable models have to be matched against experiments, and—for the time being—the most constraining experiments live not at the TeV scale, but at the weak scale in the form of electroweak precision tests (EWPT) and at the hadronic scale (Λ_{QCD}) in the form of flavor constraints.

We are wise enough to know that we cannot just brazenly compare a theory at the TeV scale with constraints at different scales; at least in principle one has to run the theory down to those scales via the renormalization group. Generally going from the TeV scale down to M_W is no problem since all of the effective couplings (even $SU(3)_c$) are perturbative and the scales are only one order of magnitude apart. *However*, running down to Λ_{QCD} is a different story completely! This is the domain of Georgi’s brown muck and adventurers are required to slog through with multi-loop calculations to maintain any control over anything. This is hard work!

Fortunately, these RG effects really are dominated by the renormalization of low energy degrees of freedom that we already know and are familiar with: the usual quarks (at least those with masses lighter than the energy scale being probed) and gluons. Sure, there are some small contributions from gluinos or KK quarks or whatever exotica might be charged under $SU(3)_c$; but these contributions to the RG *must* be small because this physics has been effectively integrated out way up at the TeV scale! Thus all we need are some very technically capable groups to actually do the multi-loop RG calculations from M_W to Λ_{QCD} for an *effective theory* composed of gluons and quarks whose masses are below the energy threshold as we run the RG scale down. (There are ‘threshold corrections’ as one flows past each quark mass.) At the hadronic scale we can sensibly place constraints on the Wilson coefficients and use this to compare to the parameter space *back at M_W* (and ideally back at the TeV scale).

Thus one finds several flavor physics updates where theorists and experimentalists convert experimental bounds to constraints on Wilson coefficients at, say M_W . Then any TeV scale model builder can take their theory and write down the M_W -scale Wilson coefficients in terms of their TeV scale parameters. One can then interpret the bounds on the Wilson


coefficients from low energies (which enterprising flavor physicists have done for us) as bounds on combinations of the ‘more fundamental’ parameters of the UV theory.

3. **Strong coupling.** In this case the UV physics may or may not be known, but is intractable by perturbative methods. Such is the familiar case with QCD: we’ve already waxed poetic about how our scattering amplitudes are well defined for asymptotic states, but because quarks are strongly coupled, single-quark asymptotic states do not exist. Instead, we must write down our scattering amplitudes in terms of hadrons, which are the effective degrees of freedom in the regime below $\sim M_W$. A more sophisticated example is the description of light mesons and baryons using chiral perturbation theory.

However, this is in fact a much more general phenomenon in theoretical physics and one that is much deeper than sometimes your low energy degrees of freedom are dramatically different from those of your UV degrees of freedom. Two excellent examples in modern theoretical physics—but certainly not the only two—are the AdS/CFT correspondence and Seiberg duality. In both cases, one theory becomes strongly coupled and intractable as you flow into the IR, but we have strong evidence to believe that there are *dual* theories which are weakly coupled precisely in the regime where the original theories are strongly coupled. These dual theories *a priori* look nothing like one another and have different degrees of freedom. The duality is a prescription for taking a UV model and assigning Wilson coefficients to the effective operators of the IR dual²⁷.

9.2 A trivial example: muon decay

Let us now consider the four-Fermi effective theory and use it to describe muon decay. We would like to approach this EFT as a tool to simplify calculations. We already know the Standard Model and we know that the effect is described by the following diagram:



Using the Standard Model, we can straightforwardly calculate this,


$$\mathcal{M}_{\text{SM}} = \left(\frac{ig}{\sqrt{2}} \right)^2 \bar{u} \gamma^\mu P_L u \frac{-i}{p^2 - M_W^2} \bar{u} \gamma_\mu P_L v. \quad (9.5)$$

We now want to use this as a toy example to see how four-Fermi theory can simplify our calculations given a *known* UV theory. Let us recall that our $(V - A)$ four-fermi Lagrangian takes the form

$$\mathcal{L}_{4F} = C J_\mu J^{\dagger\mu} \equiv C (\bar{\Psi} \gamma_\mu (1 - \gamma^5) \Psi) (\bar{\Psi} \gamma^\mu (1 - \gamma^5) \Psi). \quad (9.6)$$

²⁷There are some differences. For example, with dualities one actually has a one-to-one map between the two models, whereas for an EFT one cannot reconstruct the UV Lagrangian from an EFT Lagrangian.

The $\gamma_\mu(1 - \gamma^5)$ structure is typically called $(V - A)$ and made sense back in ancient times when people used four-Fermi theory. These days we know better that these factors should be thought of as chiral projectors. This term in the Lagrangian gives us a Feynman rule for the insertion of this operator:



Here each black dot is understood to represent a J_μ fermion bilinear, but the two black dots together should be understood to mean a *single* vertex coming from the $J_\mu J^{\dagger\mu}$ operator.

Note that the amplitude is now *trivial* to read off, it's *easier* than (9.5)—it's *not even* tree level! It's just a single-vertex with no internal propagators,

$$\mathcal{M}_{4F} = C (\bar{\Psi} \gamma_\mu (1 - \gamma^5) \Psi) (\bar{\Psi} \gamma^\mu (1 - \gamma^5) \Psi). \quad (9.7)$$

So far the grand promise of EFT seems to pan out, this calculation is indeed ‘simpler,’ just don’t expect trumpets to sound for such a meager victory. The next step is to use this to determine C . Technically what we’re doing is an **operator product expansion** and taking the lowest order term in p^2 , but we needn’t quibble over fancy words. Intuitively we are just saying that the energy scale of the process, $p^2 = m_\mu^2$, is much smaller than the W mass so that $p^2 \ll M_W^2$ and we can neglect the p^2 term. Thus we can determine the four-Fermi Wilson coefficient,

$$C = \frac{g^2}{8M_W^2} \equiv \frac{G_F}{\sqrt{2}}. \quad (9.8)$$

Now we have a predictive theory of low energy fermion scattering that includes $\mu \rightarrow e\nu\bar{\nu}$. This theory holds for energies below M_W . Note what’s happened here: when we were doing bottom-up EFT, four-Fermi theory gave us a cutoff $\Lambda_{4F} \approx 300$ GeV. Now, however, we *know* what the UV completion is, so we can say that \mathcal{L}_{4F} holds for $s^2 < M_W^2$. Note that the UV cutoff lives comfortably below the scale at which the theory actually loses perturbativity.

Let’s compare the two calculations. The Standard Model calculation

- is tree-level
- is valid—in principle—at any energy scale
- is renormalizable, by construction.

On the other hand, the four-Fermi theory

- is ‘not even tree-level’
- is only valid in the range $s^2 < M_W^2$

- is non-renormalizable, and that’s okay.

One way of addressing the non-renormalizability is to say that “this means that four-Fermi theory is an EFT,” but this is a tautological statement. In fact, the Standard Model is *also* an EFT in the sense that it is a theory whose cutoff is taken to infinity so that all irrelevant (i.e. non-renormalizable) operators vanish. This is the Wilsonian perspective for EFT and is well-documented in Hollowood’s lectures [52].

So far there’s been a missing link. We’ve done the matching (9.8) which defines the Wilson coefficient in terms of the UV parameters. But there’s a question we could have asked:

At what scale do we match, (9.8)?


This is a silly question. (9.8) is a tree-level relation. It doesn’t really matter at what scale we do the matching! Well, not quite. There are in fact two important reasons why it does matter:

1. The scale at which the EFT is defined to match the full theory is the factorization scale. Below this scale we expect our EFT to capture the dominant effects of our theory, while above it we make no promises.
2. Naïvely, this relation *does* change as one flows to lower scales due to RG effects.

9.3 The trivial example at one loop

Let us now move on, then, and see how this theory behaves at loop level. Our goal will be to understand how the full theory and EFT behave as we run down to lower scales. In particular, we would like to understand the nature of the factorization scale and the deviations (if any) of our EFT from what we would obtain if we had just used the full theory. The overarching motivation for all this is that we already know that a key part of flavor physics is being able to do multi-loop RG evolution from M_W down to Λ_{QCD} , it is thus rather important to understand the way this process occurs.


At scales below M_W , the primary renormalization comes from QED. This is, of course, rather weak at these scales, but we should think of this as practice for QCD loops when we replace muon decay by a hadronic process. We begin by drawing representative loop diagrams in the full theory:



The first diagram is manifestly finite by power counting. The second diagram is not finite: it has a log divergence. Naïve dimensional analysis (NDA) tells us that the bosons go like $1/k^2$ (in a convenient gauge) and the fermion goes like $1/k$, but recall that the three gauge boson vertex has a power of momentum as well. This is okay: such divergences not problematic and we already knew so much when we said that the Standard Model is renormalizable. In particular, upon

wavefunction renormalization of the active fields, the divergence in this diagram is cancelled and there is no need to introduce any counter term.

Let us compare this to the analogous loops that we might draw in the effective theory:



The diagram on the left corresponds to the box diagram in the Standard Model. We can note immediately that this diagram is now divergent because we've shrunk the W propagator into a single vertex. This is *not* removed by wavefunction renormalization. To the extent that there are actually any divergences in EFT, this is an actual divergence. We knew this was bound to happen since our four-Fermi EFT is non-renormalizable. The point, however, is that there never really are any divergences in an EFT because an EFT is *always* defined with a cutoff and a renormalization prescription. For an EFT which is renormalizable (like the Standard Model), this cutoff can be taken to infinity and forgotten. For non-renormalizable theories, we absolutely must define a UV cutoff *and* a prescription for making physical predictions. Thus for this diagram we must introduce a new counter term δ_{JJ} so that the effective Lagrangian contains the term


$$\mathcal{L}_{4F} = (1 + \delta_{JJ})J_\mu J^{\dagger\mu} \equiv Z J_\mu J^{\dagger\mu}. \quad (9.9)$$

This is fine, it's no cause for worry; this is just part of life working with a non-renormalizable EFT. Note that the counter term implicitly demands that it be accompanied by a prescription for its definition: you must subtract off the divergence, but you are also free to tune the finite parts as well. Thus this diagram can be *made* to match the analogous SM contribution. This is not a test of the validity of EFT, anymore, it is just fixing a free parameter that the EFT had hidden at one-loop. On the other hand, one will note that this does have predictive merit, since fixing this parameter so that the $\mu \rightarrow e\nu\bar{\nu}$ diagram matches at one-loop also trivially makes sure that, for example, the analogous $\mu\bar{e} \rightarrow \nu\bar{\nu}$ diagrams match as well.

Now let us move on to the EFT diagram on the right. This clearly matches the second diagram in the Standard Model with a photon insertion on the W propagator. One might argue that such an effect *cannot* ever hope to be captured in the EFT since the EFT does not know about the W boson. Indeed, if we *only* accounted for the $\mathcal{O} = J_\mu J^{\dagger\mu}$ operator in our EFT, then this is an effect that we would miss. However, this effect is suppressed by the W boson propagators which prevent such diagrams from being appreciable for $s \ll M_W^2$. In fact, in this case we could look at this as the statement that our EFT breaks down when this diagram becomes appreciable, i.e. when $s \sim M_W^2$. This is a valid point.

On the other hand, we can also *extend* our EFT by including a higher order operator. There's nothing wrong with this; we know that higher-order operators are generated at loop-level and so within the EFT we *should* include these effects. Of course, the slippery slope of non-renormalizable theories is that each higher-order operator begets even higher order operators that each require their own counter terms so that in principle there are an infinite number of new parameters that we should tune. (This is the *real* issue with non-renormalizable theories, not that there are any


‘actual’ divergences.) In practice, we avoid this by only taking a finite number of operators. For example, we could have just taken the leading terms in $1/\Lambda$, in which case we would only have the four-Fermi operator, $\mathcal{O} = JJ$. If we want to include the effect of a photon coming off of an internal W in our EFT, we have to include a new operator which induces a vertex of the form:



We already know how to write down this operator: we just write down the lowest-order gauge-invariant combination that gives us four fermions and a gauge boson. The correct operator is

$$\Delta\mathcal{L}_{4F} = C' (\bar{\Psi}\gamma^\mu (1 - \gamma^5) \Psi) (\bar{\Psi}\gamma^\nu (1 - \gamma^5) \Psi) F_{\mu\nu}. \quad (9.10)$$

We remind ourselves that $F_{\mu\nu}$ is gauge invariant for a $U(1)$ gauge field (otherwise one would need to take the trace of the square). One might also have thought to write down objects where the photon appears through a covariant derivative, but these vanish through the use of the equations of motion²⁸. Very naïve dimensional analysis of the external fields (just looking at the 5-point diagram) us that the operator is dimension 5, which might have led us to [very] naïvely think that $C' \sim O(1/M_W^3)$. This should cause us to pause since we know that the corresponding graph in UV theory takes the form



and thus to leading order we expect the EFT term to go like $1/M_W^4$. The resolution is clear: gauge invariance forced (9.10) to contain a factor of $F_{\mu\nu}$, which makes the operator dimension *six* so that $C' \sim 1/M_W^4$. Note that in both the EFT and UV theory there is a linear dependence on the external momentum coming from the derivative in $F_{\mu\nu}$ or alternately the momentum factor appearing in the three gauge boson vertex.

We can interpret this dimension 6 term as an estimate for the ‘error’ of the low energy EFT where we only include the four-Fermi interaction; other corrections include those that occur from taking higher derivative terms in the operator product expansion. We can see that indeed, our EFT becomes unreliable above energies on the order of $\mathcal{O}(M_W)$.

Before closing this discussion, let us qualitatively discuss the loop integrals involved in these processes. The one-loop correction to the effective four-fermion vertex, for example, has an integral

²⁸That the *classical* equations of motion should hold in as an operator relation in the *quantum* theory is a highly non-trivial statement, but turns out to be true. This fact is used extensively when writing out a complete basis of electroweak precision operators [8, 16, 49]. For discussions of this fact, see Section 12 of Politzer [74], Chapter 7.7 of Weinberg volume I, Georgi [38], or Arzt [7]. We thank Witek Skiba for explaining this point to us; see his TASI lectures [86] for further discussion of the topics in this footnote.

that takes the form

$$\int_p^{M_W} d^4k \frac{1}{k^4}. \quad (9.11)$$

As we mentioned above, the logarithmic divergence is not removed by wavefunction renormalization; we'll have more to say about this below. What we would like to emphasize, however, is that this integral can be separated into two pieces:

$$\left(\int_p^\mu + \int_\mu^{M_W} \right) d^4k \frac{1}{k^4}. \quad (9.12)$$

This is a manifestation of the factorization between UV and IR scales, where μ —which was hitherto an arbitrary parameter coming from dim reg—plays the role of separating the high and low scales. This is largely because of the magic of logarithms. This is because the renormalization factor Z in (9.9) takes the form

$$Z = \left(1 + \alpha \log \frac{M_W^2}{p^2} \right) = \left(1 + \alpha \log \frac{M_W^2}{\mu^2} \right) \left(1 + \alpha \log \frac{\mu^2}{p^2} \right), \quad (9.13)$$

so that to leading order in α , (9.12) takes the form

$$\log \frac{M_W^2}{p^2} = \log \frac{M_W^2}{\mu^2} + \log \frac{\mu^2}{p^2}. \quad (9.14)$$

Still being extremely qualitative, the point of this factorization is that the low-energy logarithm ($\log \mu^2/p^2$) goes into the renormalization matrix elements of the effective operator while the high-energy logarithm ($\log M_W^2/\mu^2$) goes into the renormalization of the Wilson coefficient. In other words, the Callan-Symanzik equation for an amplitude can be written as

$$0 = \mu \frac{d}{d\mu} \langle \text{out} | C_i \mathcal{O}_i | \text{in} \rangle = \left(\mu \frac{d}{d\mu} C_i \right) \langle \mathcal{O}_i \rangle + C_i \left(\mu \frac{d}{d\mu} \langle \mathcal{O}_i \rangle \right), \quad (9.15)$$

where $\langle \mathcal{O}_i \rangle$ is the appropriate matrix element. We'll say much more quantitative things about this in Section 9.5.

9.4 Mass-independent schemes

THIS SECTION IS STILL BEING WRITTEN!

9.5 Operator Mixing

10 Remarks on Lattice QCD

We've mentioned lattice QCD a few times now, so let us pause to make a few remarks about what is actually done. We will be necessarily brief, but excellent references for theory/phenomenology students interested in learning more are the lectures by Sharpe [84], LePage [60], and Gupta [46]. The lattice approach was first proposed by Ken Wilson, who was convinced that the approach would solve QCD within half a year. Now, forty years later, lattice calculations are *just* starting to grow into maturity.

10.1 Motivation and errors

The basic idea is to discretize spacetime to numerically solve the partition function of QCD,

$$Z = \int D\phi e^{-iS[\phi]}, \quad (10.1)$$

where ϕ is shorthand for the quark and gluon fields. In order to calculate this we Wick rotate and evaluate the action on paths over the discretized spacetime. The Wick rotation is essential to convert the highly oscillatory integral to one that is exponentially damped and hence numerically tractable.

Ideally we would like to take the lattice size to infinity and the lattice spacing to zero; this corresponds to the continuum limit of an infinite spacetime. The fundamental problem (but not the ‘big’ problem, as we’ll discuss below) with the lattice is that we cannot take these limits and are restricted both in the IR by finite volume and in the UV by finite lattice spacing. This introduces a systematic error for any calculation and part of the goal of lattice QCD calculations is to beat down this error.

On top of this, however, this is also a statistical error coming from the Monte Carlo scans over a finite number of random field configurations. Thus even if one runs for an arbitrarily long time, one will eventually run into this statistical error. **PROBLEM: I’M CONFUSED: SHOULDN’T THE STATISTICAL ERROR DECREASE WITH LONGER RUNNING TIME SO THAT IT EVENTUALLY RUNS INTO A SYSTEMATIC ERROR?**

A third source of error comes from numerical relics coming from doing the Wick rotation on a discrete space. The point here is that while the Wick rotation is perfectly well defined in the continuum space, it is no longer strictly kosher in discrete space.

The current state of the art calculates make use of a spacetime with dimensions $32^3 \times 64$, where the time direction has twice the resolution of any space direction. While this may seem relatively modest, most of the systematic errors come from the particular algorithms used in the numerical evaluation so that there are more gains from using more sophisticated techniques than from using a finer lattice.

10.2 ‘Solving’ QCD

Since we are interested in accounting for the non-perturbative effects at Λ_{QCD} , want the IR cutoff set by the lattice size L and the UV cutoff set by the lattice spacing a to be far from Λ_{QCD} ,

$$L^{-1} \ll \Lambda_{\text{QCD}} \ll a. \quad (10.2)$$

In practice, the current state of the art values are

$$L^{-1} \approx 30 \text{ MeV} \quad (10.3)$$

$$a \approx 3 \text{ GeV}. \quad (10.4)$$

While we remarked earlier that this is fine for transitions like $c \rightarrow s$ or describing the $\phi = s\bar{s}$ meson, it seems like the IR limit is still hopelessly larger than the up and down quark masses and the UV limit is much smaller than the bottom quark mass (and only marginally within the

range of the charm). Physically, the lattice is unable to take into account the up quark because at $m_u \sim 5$ MeV, it is bigger than the entire lattice! Similarly, there's no way for the lattice to currently take into account the effects of a 'real' bottom quark in its path integral.

Thus in order to make use of the lattice beyond the narrow range of mesons in the 'sweet spot' near Λ_{QCD} , physicists have to use as many analytical tricks that they can before passing the problem to a lattice calculation. For example, for pions we must first make full use of chiral QCD before submitting a job to a supercomputer (or, in modern days, a grid computing farm). Similarly, for heavy quark mesons one must extract everything we can with heavy quark effective theory before going to the lattice. The point is that we don't need to pass the full theory to the lattice, only an effective theory.

10.3 The Nielsen-Ninomiya No-Go Theorem

While all of the above considerations are generic problems for lattice calculations, there is one particular 'big problem' in lattice QCD associated with fermions: the Nielsen-Ninomiya no-go theorem [69]. It turns out that chiral fermions are very difficult to generate on a discretized space time. For new physics model builders this is almost a perverse joke: in supersymmetry one introduces quantum extra dimensions to double the particle spectrum. Meanwhile, it has been known for over three decades that by quantizing spacetime itself (i.e. latticizing Minkowski space) one multiplies the fermion spectrum by a factor of *sixteen* (2^4 , where four is the number of discretized dimensions).

Naïvely one can just divide by this factor of 16, but—just like in theories of an ordinary extra dimension—this makes it tricky to generate the usual *chiral* fermions of the Standard Model. It is generally very difficult to restore chiral symmetry on the lattice, though several ideas exist.

The best idea so far, by David B. Kaplan, is very computationally demanding. The idea is to introduce an unphysical fifth lattice dimension. By adding a domain wall in the extra dimension one can set boundary conditions to generate a chiral fermion. This has a nice analogy to the Brillouin zone in solid state physics. For a nice review, see [56]. Contemporary model builders may recognize this as simply an early manifestation of the split-fermion idea.

Domain wall fermions appear in many contexts. One modern application is the **split fermion** scenario, which is reviewed in [45]. The basic idea is to try to explain the hierarchies in the fermion masses and mixings by assuming an extra dimensional scenario where the fermion wavefunctions are localized in different parts of the extra dimension. Assume, for example, that one has a 5D scalar that obtains a vev,


$$\mathcal{L}_\phi = (\partial_M \phi)^2 - \frac{\lambda}{4} (\phi^2 - v^2). \quad (10.5)$$

This gives kink solutions such that $\phi \rightarrow \pm v$ as $x_5 \rightarrow \pm\infty$. Cosmologically as the universe cools down to temperatures below $T < v$ one can imagine that the field settles into regions with $\langle \phi \rangle = \pm v$. Eventually neighboring kink/anti-kink regions will cancel each other out. If, however, the points at $x_5 = \pm\infty$ live in different vacua, then one will be left with a stable

solitonic kink,

$$V(x_5) = N \tanh(ax_5). \quad (10.6)$$

In principle this profile can be sharp or smooth depending on a . The kink is what we know as a brane in field theory, and a determines the brane's size.



Of course in realistic models one has to compactify the extra dimension so that the $x_5 = \pm\infty$ points have to be identified. Practically this means that one has to **orbifold** the extra dimension; but we won't worry about these details. The

Now let's introduce fermions into the game. Unlike free fermions, we will give these fermions a background field,

$$\mathcal{L}_\Psi = \bar{\Psi}(i\not{D} - y\phi)\Psi. \quad (10.7)$$


Doing the usual Kaluza-Klein decomposition and determining the eigenvalue equation for the zero mode profile $f(x_5)$ gives,

$$(\partial_5 + y\phi(x)) f(x_5) = 0. \quad (10.8)$$

The solution is $f = e^{-S}$ where

$$S(x_5) = y \int_0^{x_5} \langle\phi(x_5)\rangle. \quad (10.9)$$

This is a peaked profile that is stuck to the kink.



We can understand this intuitively. At values far away from the origin, $\langle\phi(x_5)\rangle$ is large relative to the temperature of the universe, $T \sim 0$, so that the fermions want to stay near the kink where $\langle\phi(x_5)\rangle \approx 0$.

In the split fermion scenario, we add flavor-dependent terms M_i to the fermion Lagrangian,

$$\mathcal{L}_\Psi = \bar{\Psi}(i\not{D} - y\phi + M_i)\Psi \quad (10.10)$$

which localizes the fermions at different points on the extra dimension. The overlap of wavefunctions, say e_L and e_R , give the masses and mixings. For example, one can have the leptons and quarks localized apart from one another in order to prevent proton decay in GUT models. These days people play similar games with Randall-Sundrum models. (Do I even need to provide references for this anymore?).

This is a nice idea, but it makes the computation more difficult since one has gone from a four- to a five-dimensional lattice.

10.4 The quenched approximation

One idea that is often mentioned when discussing lattice calculations of form factors and decay constants is the **quenched approximation**. This is the approximation in which one ignores the fermion contribution to the path integral. The reason for this approximation is purely computational restrictions: including this contribution using standard techniques leads to a factor of 10^3 increase in the required processor-hours to perform calculations.

Let us remind ourselves that part of the art of our science is being able to use *controlled* approximations to elucidate physics apart from technical mathematics. This has been one of the central themes to these lectures. The quenched approximation, however, is *not* a controlled approximation: there is no continuous approximation parameter to quantify the error. It's simply an approximation made out of necessity and justified post-facto with heuristic arguments. There is no physics in the quenched approximation²⁹.

Like a childhood blanket, the quench approximation is a relic of a time when the lattice QCD field was in its infancy. In the past ten years there has been remarkable progress leading to real unquenched calculations allowing lattice physicists to reduce the error on, for example, quarkonia spectra from 10% to 1%. The point isn't that the error has decreased, however, but that the error is under control and quantifiable.

11 Heavy quark symmetry and the CKM

Recall our 'big picture' motivation in Section 3.5 where we explained that for the heavy quark mesons the dominant contribution to the mass comes from the mass of the heavy quark rather than Λ_{QCD} and that this affects the mixing of these mesons relative to the light mesons. Now that we've squeezed everything that we could out of approximate flavor symmetries to determine CKM matrix elements associated with the light mesons in Section 8, let us now proceed to lay the groundwork for attacking the heavy quark mesons. Our strategy will be to make use of a different kind of handle on these systems, **heavy quark symmetry** (HQS).

Heavy quark symmetry is different from the symmetries which we are already used to. For example, in chiral symmetry we know that there is a parameter in the Lagrangian which we can set to zero to yield a larger symmetry. In heavy quark symmetry we take the limit where the heavy quark mass goes to infinity. This is not a symmetry which is manifest in the Lagrangian, but rather one that is specific to particular meson systems. This is why the subject only blossomed in the 1990s. There is a nice set of review literature:

- A short review by Mark Wise³⁰ [95]

²⁹By now you should get the feeling that 'quench' is a very dirty word in our field. The September 2008 magnet quench in the LHC delayed the LHC time table to design-energy collisions by several years. <http://public.web.cern.ch/press/pressreleases/Releases2008/PR14.08E.html>

³⁰Mark Wise's middle name is Brian. At some point there was a typo on his driver's license which spelled it as 'Brain' so that his official state ID read 'Mark Brain Wise.'

- Technical lectures by Benjamin Grinstein [43]
- Neubert’s lecture notes [68] and a detailed review [67]
- The yellow book by Manohar and Wise [64]³¹
- Finally, the insightful and idiosyncratic TASI lectures by Howard Georgi, which appear to only be available on his website³².

11.1 The hydrogen atom

All of the essential physics of heavy quark effective theory is contained in the undergraduate quantum mechanics of the hydrogen atom. Let’s start with the following question:

What is the difference between hydrogen and deuterium?

Deuterium, you’ll remember, is a hydrogen atom with an additional neutron. Chemists—who are, after all, effective theorists of atomic physics—call these states isotopes and usually treat them as ‘basically the same.’ We know that technically hydrogen and deuterium differ by their quantum numbers under the Lorentz group: they have different mass and spin. Why, then, is it that these two ‘technically different particles’ are ‘basically the same’ in chemistry? To put this on more concrete footing, if you were suddenly transported to a parallel universe where all of the oxygen on Earth were replaced by oxygen-17, we would survive just fine because for the most part chemistry doesn’t care about isotopes³³.

The reason why hydrogen and deuterium are chemically ‘basically the same’ is that the electrons really don’t care about the nucleus. Or, to put it in more general terms that will become the mantra of this section,

The light degrees of freedom are insensitive to the heavy degrees of freedom that source the potential.

In the hydrogen atom the nucleus just plays the role of the electromagnetic potential about which we can do non-relativistic quantum mechanics and calculate energy spectra. To leading order the potential is insensitive to the mass and spin of the actual nuclear source of the electromagnetic field. What do the corrections look like?

Problem 11.1. *An undergraduate quantum mechanics course should be 30% harmonic oscillator, 60% hydrogen atom, and 10% “other important ideas.”* If you are rusty with the quantum mechanics of the hydrogen atom, pick up a self-respecting textbook and review it now. The ‘big picture’ intuition that we’ve been developing should help you better appreciate the physics being done.

³¹An anonymous faculty member in our group once implied that Yuval ‘borrowed’ his copy of this textbook. When that faculty member eventually found his copy, he gave Yuval the extra copy he bought. What a friendly group that we have.

³²<http://www.people.fas.harvard.edu/~hgeorgi/tasi.pdf>

³³We can get away with glib statements like this about chemistry because chemists wouldn’t be reading these notes. <http://xkcd.com/435/>

Recall that in the hydrogen atom we can calculate not only the leading order spectra, E_n , we can also calculate the energy splittings coming from fine structure (e.g. spin-orbit coupling) and hyperfine structure (e.g. spin-spin coupling),


$$E_n \sim m_e \frac{\alpha^2}{n^2} \tag{11.1}$$

$$\Delta E^{\text{fs}} \sim m_e \alpha^4 \tag{11.2}$$

$$\Delta E^{\text{hf}} \sim m_e \alpha^4 \frac{m_e}{m_p}. \tag{11.3}$$

Note the extra m_e/m_p suppression on the hyperfine splitting. Physically this is telling us that in the limit of infinite mass one cannot rotate the proton to reverse its dipole moment. The additional spin coming from the neutron in deuterium clearly changes the hyperfine splitting.

The other difference coming from the larger nuclear mass leads to a different reduced mass. Recall that the ‘reduced’ in ‘reduced mass’ means we are reducing a two-body problem to a one-body problem such as the Kepler system. In the rest frame of the Earth-Sun Kepler system *both* the Earth *and* the sun rotate about a common center of mass,



naturally the heavier object orbits with a much smaller radius. In other words, the whole system is at rest, but the heavy guy has some small velocity. Thus let us use suggestive language and say that deuterium also differs from hydrogen by the velocity of the nucleus.

The two differences between hydrogen and deuterium lead to energy splitting, but both effects decouple in the limit where the nucleus is taken to infinite mass since both effects scale like $1/m_{\text{nuc}}$.

11.2 Heavy quark symmetry: heuristics

The infinite nuclear mass limit is the precise analog to the heavy quark limit that we’ll be using. The heavy quark mesons which we’ll consider—primarily B mesons³⁴—are just like hydrogen systems where the binding comes not from electromagnetism but from QCD. In fact, we may consider a naïve picture where the light quark orbits the b .

We might argue that this is being too naïve. There’s a very big difference between the electromagnetic potential and the QCD potential: while the former is perturbative, the latter is hopelessly non-perturbative. Since virtual partons are $\mathcal{O}(1)$ effects, we shouldn’t even be able to say that the B meson is composed of a b and a light quark, say \bar{u} .

This is precisely the issue of the **brown muck**, the dirty nonperturbative physics that should make the heavy mesons intractable. Our goal is to get around the brown muck without getting ourselves too dirty; we want to find ways to calculate *the things that we care about* without having to understand the intractable physics of the muck. Should this be possible? Yes! In fact, this

³⁴Recall that tops ‘decay before they can hadronize’ and that the D mesons are only borderline ‘heavy.’

should now sound familiar: this is precisely what we meant when we said that chemists don't care about isotopes, and why we don't have to re-derive all of the details of deuterium.

When the bottom quark becomes extremely heavy, the B meson spectrum becomes independent of its mass and the spin. Mass, for us, is really an index for flavor (e.g. b or c quark). We can, of course, completely neglect electromagnetism. The calculation of the 'energy levels' of the B —that is, the spectrum of mesons—follows precisely as the hydrogen atom. For example, the B and B^* states represent the hyperfine splitting, while the higher spin resonances $B_1, B_2 \dots$ are the p -wave states and so on.

Now assume that $m_c, m_b \gg \Lambda_{\text{QCD}}$ and that the spectrum of D and quarks have been measured in nature along with the mass of the lightest B meson. Just as we could figure out everything about deuterium once we understood hydrogen, we can similarly determine the rest of the B spectrum simply based on the D spectrum and the lowest B mass as a reference point. Further, we are able to do this without knowing the physics of the brown muck just as we didn't have to re-derive the deuterium spectrum from scratch. The only difference between the D s and B s are the mass of the heavy quark and, as we said before,

the light degrees of freedom just don't care.

In the heavy-quark limit, the splitting between these states is the same.

Example 11.1. For example, in the $m_b \rightarrow \infty$ limit, we have the relation

$$\frac{m_{B^*} - m_B}{m_{B^*} + m_B} = 0. \quad (11.4)$$

The B and B^* are degenerate in the heavy quark limit. In the language of symmetry, we say that these two states form an $SU(2)$ doublet with respect to the spin orientation of the heavy quark. (Recall that these are just the singlet and triplet states of the hyperfine structure.)

Further, up to some normalization we know that

$$\langle B | \mathcal{O} | D \rangle = 1. \quad (11.5)$$

What we mean by this is that in this transition—once again—*the light degrees of freedom just don't care*. The $b \rightarrow c$ transition hasn't changed the QCD potential which the light quark feels. In fact, it doesn't matter what the operator \mathcal{O} is that enacts the $b \rightarrow c$ decay. This is just a statement about the wavefunction overlap. We say that the B and D form a doublet under a flavor group $SU(2N_f)$, where there are N_f flavors and the factor of two comes from the spin orientation.

11.3 Heavy quark symmetry: specifics

We can now see that heavy quark symmetry is very different from the usual symmetries that we work with in effective field theory. In 'normal' effective field theories integrate out some heavy degrees of freedom. In the chiral Lagrangian, for example, we can take things to zero and see how we recover symmetries. Heavy quark symmetry is different in a very fundamental way. The $M \rightarrow \infty$ limit does *not* increase the symmetry of the Lagrangian. In fact, we will only 'integrate out' *part* of the degrees of freedom.

When we do a regular EFT, we expand about the vacuum and integrate out high frequency modes. In heavy quark effective theory, we will expand about the background of a single heavy quark, say a b . In such a background one *assumes* that the heavy quark is there classically; we get it ‘for free’ without having to consider how it might pop out of the vacuum. If we want any *additional* heavy quarks, however, we have to honestly pay the cost of introducing an additional high frequency mode to the background.

Note that this background breaks CP since it prefers a particle versus its antiparticle, but that’s fine; this is like a ‘spontaneous breaking’ of CP in the region around a heavy quark. Eventually we will see that we should go to the non-relativistic limit, but let us remark that this is different from so-called non-relativistic QCD (NRQCD)³⁵.

Let’s now consider the spectrum of mesons containing a heavy quark. This turns out to be very simple and can be done even before dipping into heavy quark effective theory. The main idea is to consider an expansion in powers of Λ_{QCD}/m_Q , where m_Q is the mass of the quark. Due to the ambiguities in defining a quark mass, we should really say that it is some effective mass—a pole mass, \overline{MS} mass, whatever—which is identified with a physical mass by some prescription. The mass of a hadron containing this heavy quark, m_H , is

$$m_H = m_Q + \bar{\Lambda} + a/m_Q + \dots . \quad (11.6)$$

This is a trivial parameterization where we’ve explicitly written out the zeroth, first, and second order terms in our expansion. The leading term just reminds us that in the heavy quark limit the mass of the hadron and the mass of the quark are basically the same. The sub-leading terms all depend on the light degrees of freedom; the $\bar{\Lambda}$ is just the mass of the light degrees of freedom and the a term parameterizes the interaction between the light degrees of freedom and the heavy quark. Note that $\bar{\Lambda}$ doesn’t know about the mass or spin of the heavy quark, it is a parameter that *only* deals with the light degrees of freedom *independent* of the heavy degree of freedom sourcing the QCD potential. Only the interaction term takes these into account; this is just the hyperfine splitting in hydrogen.

Example 11.2. Consider the case when the hadron H is a B or B^* . One only observes a difference between the B and B^* mass at $\mathcal{O}(1/m_Q)$, i.e. in the a term. On the other hand, the difference between the B and the B_s masses occurs at $\mathcal{O}(m_Q^0)$, i.e. in the $\bar{\Lambda}$ term.

By now you should have guessed what kind of information is contained in a : (1) reduced mass or kinetic energy of heavy guy in the meson rest frame and (2) the spin. In fact, we can further parameterize a as

$$a = -\lambda_1 + 2 \left[J(J+1) - \frac{3}{2} \right] \lambda_2. \quad (11.7)$$

Here the λ_1 term is universal and the λ_2 term is associated with the spin, J .

Example 11.3. Which spin does J correspond to: the spin of the heavy quark or the spin of the meson? J corresponds to the total spin of the meson, just as in the case of the hydrogen atom.

³⁵For a review see, for example, [12]. For an unrelated example of how non-relativistic effective theories can be used in the theory of direct dark matter detection, see [31].

By dimensional analysis $\lambda_{1,2}$ have mass dimension two. In fact, by the method of “*there’s only one other scale lying around,*” we can guess that $\lambda_{1,2} \approx \Lambda_{\text{QCD}}^2$ and $\bar{\Lambda} \approx \Lambda_{\text{QCD}}$.

Example 11.4. For example, consider the ratio of the B_s – B_d mass splitting to the D_s – D_d splitting. Heavy quark symmetry predicts that these splittings are the same up to effects on the order of $1/m_c$. Applying (11.6) gives

$$r_1 \equiv \frac{m(B_s) - m(B_d)}{m(D_s) - m(D_d)} = \frac{\bar{\Lambda}_s - \bar{\Lambda}_d + \mathcal{O}(\frac{1}{m_b})}{\bar{\Lambda}_s - \bar{\Lambda}_d + \mathcal{O}(\frac{1}{m_c})} = 1 + \mathcal{O}(1/m_c). \quad (11.8)$$

As another example, one can show that

$$m^2(B^*) - m^2(B) = 4\lambda_2. \quad (11.9)$$

It is, of course, not surprising that this splitting only comes from the non-universal part of the a term. Comparing to actual measurements, we find

$$\lambda_2(m_B) \approx 0.12 \text{ GeV}^2. \quad (11.10)$$

We can do the same calculation the charmed mesons and we would also get 0.12, note that the bottom gives preferable results since the error goes like $1/m_Q$. In fact, looking at the differences in the squared masses of the vector and pseudoscalar mesons for the B and the D ,

$$m^2(B^*) - m^2(B) = 0.47 \text{ GeV}^2 \quad (11.11)$$

$$m^2(D^*) - m^2(D) = 0.55 \text{ GeV}^2. \quad (11.12)$$

What about the kaon? Of course, we don’t expect this heavy quark symmetry to hold, but it turns out that $m^2(K^*) - m^2(K) = 0.55 \text{ GeV}^2$. What do we get for the ρ and pion? We can neglect the pion mass, and we recall that $m_\rho = 770 \text{ MeV}$ so that the difference of the squared masses is 0.57 GeV^2 . There’s no reason why these should hold, but it is a notable observation that they do.

11.4 HQET

Now let’s get to the heart of heavy quark effective theory (HQET), where we make use of the idea that we will expand about the background of a classical heavy quark. We now have to be a little more formal. Suppose we have a b quark in an on-shell hadron with momentum

$$p_H^\mu = m_h v^\mu. \quad (11.13)$$

Here we have introduced the relativistic four-velocity, v^μ , in anticipation that we will eventually take a non-relativistic limit. If the hadron is on-shell and if the quark is indeed heavy, we can say that the b quark itself is also ‘almost on-shell’ in the sense that to good approximation we know the momentum of the heavy quark,

$$p_q^\mu = m_Q v^\mu + k^\mu, \quad (11.14)$$

where k^μ is a small correction such that $k^\mu \ll m_Q v^\mu$. Let us define a four-velocity for the heavy quark,

$$v_Q^\mu = \frac{p_Q^\mu}{m_Q} = v^\mu + \frac{k^\mu}{m_Q}. \quad (11.15)$$

The second term is now manifestly a small correction and in the heavy quark limit, $v_Q = v_H$.

Now we introduce a trick. Let us define a projection operator

$$P_\pm = \frac{1 \pm \not{v}}{2}. \quad (11.16)$$

where $\not{v} = v_\mu \gamma^\mu$. This projects between particle and anti-particle states. Thus let us define the projected states,

$$h_v(x) = e^{im_Q v \cdot x} P_+ Q(x) \quad (11.17)$$

$$H_v(x) = e^{im_Q v \cdot x} P_- Q(x). \quad (11.18)$$

The overall phase is like the unphysical phase of the Schrodinger equation where the energy is arbitrary (unless we include gravity). All we're doing is pulling out the quickly oscillating part of the Q . These are projection operators in the sense that

$$\not{v} h = h \quad (11.19)$$

$$\not{v} H = -H. \quad (11.20)$$

Let's go to the rest frame of the particle $v_\mu = (1, \mathbf{0})$, then these h and H are just the upper and lower two components of the four component Dirac spinor. These are just the electron and the positron. Note that the h and H particle/antiparticle decomposition is *different* from the χ and $\bar{\psi}$ chiral decomposition.

Let's go ahead and write the heavy quark field as

$$Q(X) = e^{im_Q v \cdot x} (h_v(x) + H_v(x)). \quad (11.21)$$

We can now plug it into the Lagrangian to obtain

$$\mathcal{L} = \bar{Q}(i\not{D} - m_Q)Q \quad (11.22)$$

$$= \bar{h}_v i v \cdot D h_v - \bar{H}_v (i v \cdot D + 2m_Q) H_v + \bar{h}_v i \not{D}_\perp H_v + \bar{H}_v i \not{D}_\perp h_v, \quad (11.23)$$

where we've defined

$$D_\perp \equiv D^\mu - v^\mu (v \cdot D). \quad (11.24)$$

In the rest frame, $v \cdot D_\perp = 0$ so that $D_\perp = (0, \mathbf{D})$. Now let's integrate out the 'heavy degrees of freedom,' H_v . To do this we solve the Euler-Lagrange equation for H_v which gives

$$H_v = \frac{1}{2m_Q + i v \cdot D} \cdot i D_\perp h_v. \quad (11.25)$$

We can use this to eliminate H_v from the Lagrangian,

$$\mathcal{L} = \bar{h}_v i v \cdot D h_v + \bar{h}_v i \not{D}_\perp \frac{1}{2m_q + i v \cdot D} i \not{D}_\perp h_v. \quad (11.26)$$

Great! Mission accomplished. Note, however, that the second term is rather funny-looking; it has a derivative in the denominator. Intuitively this is the remnant of the H_v propagator. It is a non-local term. In order to convert this into a local term—as required for an effective Lagrangian—we do an operator product expansion in $v \cdot D/m_Q \ll 1$, since the derivative acts on the light field and picks up a momentum on the order of Λ_{QCD} . Expanding we get

$$\mathcal{L} = \mathcal{L}_{\text{kin}} + \frac{1}{2m_Q} \sum \bar{h}_v i \not{D}_\perp \left(\frac{i - v \cdot D}{2m_q} \right)^n i \not{D}_\perp j_v \quad (11.27)$$

$$= \mathcal{L}_{\text{kin}} + \frac{1}{2m_Q} \bar{h}_v (i D_\perp)^2 h_v + \frac{g_s}{4m_Q} \bar{h}_v \sigma_{\mu\nu} G^{\mu\nu} h_v + \dots \quad (11.28)$$

So now we've written out the leading-order higher-dimensional operators in our heavy quark EFT. We should actually look at these two terms with great fondness. The first term is just a factor on the order of $p^2/m_Q \sim \Lambda_{\text{QCD}}^2/m_Q$ while the second term is a chromo-magnetic operator. These are precisely the λ_1 and λ_2 terms in (11.7), which we had identified from very physical intuition. Let us define effective operators

$$\mathcal{O}_1 + \mathcal{O}_2 = \frac{1}{2m_Q} \bar{h}_v (i D_\perp)^2 h_v + \frac{g_s}{4m_Q} \bar{h}_v G_{\mu\nu} G^{\mu\nu} h_v \quad (11.29)$$

It just jumps out at you that these terms correspond to kinetic energy of the heavy quark in the rest frame of them meson and the hyperfine interaction between the heavy quark and the brown muck. This allows us to identify the physical meaning of the $\lambda_{1,2}$ parameters,

$$\langle H(v) | \mathcal{O}_{1,2} | H(v) \rangle \sim \lambda_{1,2}. \quad (11.30)$$

The ‘perpendicular’ covariant derivative D_\perp is kind of annoying, but fortunately the deviation between D_\perp and D is third order so that when we work only to second order it can be replaced by the full covariant derivative, D .

11.5 Measuring $|V_{cb}|$

Finally, let's use heavy quark symmetry to obtain something useful, $|V_{cb}|$. As was the case with $SU(3)$ flavor, we want to invoke heavy quark symmetry to simplify our treatment of the form factors.

11.5.1 $B \rightarrow D$ decays

The natural decays to consider are

$$B \rightarrow D \ell \bar{\nu} \quad (11.31)$$

$$B^* \rightarrow D \ell \bar{\nu} \quad (11.32)$$

$$B \rightarrow D^* \ell \bar{\nu}. \quad (11.33)$$

Note immediately that (11.32) is highly suppressed. Under heavy quark symmetry the rates for (11.32) and (11.33) should be the same since the B and B^* (as well as the D and D^*) are doublets under spin symmetry. However, (11.32) is a lost cause since B^* decays via $B^* \rightarrow B\gamma$ electromagnetically, which is much stronger than weak decay, c.f. Example 8.11. Of course, the B subsequently decays according to (11.31), but now we cannot make use of the form factors.

We are thus left with (11.31) and (11.32). If we didn't know anything about heavy quark symmetry we would say that (11.31) contains two form factors, while (11.33) contains four (one for the vector and three for the axial), giving six for these two decays. In addition, we can also use the baryon decay,

$$\Lambda_1 \rightarrow \Lambda_2 \ell \bar{\nu}, \quad (11.34)$$

which is essentially the same as neutron decay. This also has six form factors, giving a grand total of twelve form factors between the three decays. Yuck!

Now for a big surprise: In the heavy quark limit, *all* of these form factors are either zero or the *same* non-zero value, which we can normalize to unity. This allows us to write everything in terms of a single form factor. Even better, it is easy to understand how actual QCD differs from the heavy quark limit in terms of an object called the **Isgur-Wise function** which is universal and depends only on $v_b \cdot v_c$.

Let us work in the specific limit

$$\frac{m_b}{m_c} = \text{const} (\sim 3) \quad (11.35)$$

while simultaneously taking $m_b \rightarrow \infty$. Imagine a B meson to be composed of a b quark surrounded by brown muck. Consider the velocity of the b to be the same as the velocity of the entire meson, $v_b = v$ and suppose the b decays into a c and let us write the velocity of the c as $v_c = v'$.

Now we would like to ask what would happen if we replaced the B meson with a B^* or if we changed the D to a D^* ? Would the brown muck also be excited? No, nothing would happen. We have a theory of QCD where we treat electroweak currents as external sources. If Harry Potter waved his magic wand and turned the b into a c quark, the brown muck would more or less stay the same. This is just like our atomic physics analogy where adding another neutron to an atom won't lead to a big change in its chemical properties; nothing happens. And when nothing happens, the intuitive value of the form factor is unity (suitably normalized, of course).

In the limit of $v \rightarrow v'$, we don't care about b versus c , D versus D^* , or any of these 'magic wand' modifications that don't significantly affect the brown muck. In fact, we don't even care *how* the $b \rightarrow c$ transition occurs, as long as it is short distance. This transition could just as well have come from a vector current, an axial current, some $V \pm A$ combination, or Harry Potter's magic wand.

Now consider $v \neq v'$. In the decay $B \rightarrow D$ this corresponds to the c quark picking up a non-zero velocity in the meson rest frame. Note that we are *not* allowed to assume the non-relativistic limit for the velocity of the c quark, v' . The only thing we know is that the b is non-relativistic in the meson rest frame. (It is a fallacy to say that 'everything' is non-relativistic in the heavy-quark limit.) The c can be relativistic in the rest frame of the B . In this case, the b decays to c with some generically-not-small velocity and suddenly brown muck sees the color flow moving. Then one of two things happen,

1. The brown muck can follow the color flow. This corresponds to $B \rightarrow D$ decay.
2. The brown muck can pop things out of the vacuum to produce other things.

This brings us back to the meaning of the form factor: the probability for $B \rightarrow D$ versus other decay processes when things are not at rest. The analogous statement in the hydrogen/deuterium atom is that not only do we take away a neutron, but we give the proton a kick. The form factor can be interpreted as the overlap between the electron wavefunction before and after the kick and represents the probability of finding an electron in a given shell after the proton kick. The case with the brown muck is completely analogous except that we do not know the explicit wavefunction.

There's one simple option for generating a Lorentz invariant out of v and v' : $w \equiv v \cdot v'$, where w is just what we usually call γ , the boost. The form factor is given by some universal function of only w called the **Isgur-Wise function**, $\xi(w)$, such that $\xi(1) = 1$, i.e. when $v = v'$. In general $w \geq 1$. We don't know much more about $\xi(w)$, but we can see how far we can go with these properties.

Let us choose a somewhat different normalization of so that we right the form factors,

$$\frac{1}{m_B} \langle B(v) | V^\mu | B(v') \rangle = \xi(w) \cdot (v \cdot v')^\mu. \quad (11.36)$$

We can compare this to the 'usual' way of parameterizing the form factors,

$$\langle B | V^\mu | B \rangle = F(q^2)(p + p')_\mu. \quad (11.37)$$

Example 11.5. By the way, why is there no $(p - p')$ term? It is zero in the isospin limit.

Now we would like to calculate the form factors based on the Isgur-Wise function. Let's look at the decay of $B \rightarrow D$ so that we would like

$$\frac{1}{\sqrt{m_B m_D}} \langle B(v) | V^\mu | D(v') \rangle = \xi(w) \cdot (v \cdot v')^\mu. \quad (11.38)$$

In the 'usual' notation,

$$\langle B | V^\mu | D \rangle = F_+(q^2)(p + p')_\mu + F_-(q^2)q_\mu. \quad (11.39)$$

Our hope is to relate the form factors F_\pm to the Isgur-Wise function $\xi(w)$. First consider the case $v = v'$ ($w = 1$). We find

$$F_\pm(q^2)|_{w=1} = \frac{m_B \pm m_D}{2\sqrt{m_B m_D}}. \quad (11.40)$$

HELP: please specify the $\xi(w)$ dependence. In principle there are two form factors, but in the heavy quark limit they are related so that we only have to determine one of them before we can start making predictions.

Three remarks are in order:

1. Does this work for $B \rightarrow D^*$? Yes, it's more complicated because the form factors include more terms, e.g. contractions with ϵ tensors, but at the end of the day everything is indeed expressed in terms of $\xi(w)$.

2. What kind of corrections are there in the heavy quark limit, $\xi(1) = 1$? There's a correction from the finiteness of the heavy quark mass, $1/m_Q^2$ as well as *perturbative* α_S corrections from soft gluons.
3. **Luke's theorem.** If you expand about an extremum, the correction is always second order. This is why the leading order correction to $\xi(w) = \xi(1) + \dots$ starts at $\mathcal{O}(1/m_W^2)$. The point $\xi(1)$ is an extremum because we're looking at a situation with 'maximum wavefunction overlap.'

Now what do we want to do with all of this? We would like to measure $B \rightarrow D$ and use the spectrum to obtain $|V_{cb}|$. Ideally we would like to look at the case where $w = 1$, i.e. where the charm is at rest, but the phase space—and hence statistics—for this is very small. We must thus look at all events, plotting the total number events as a function of w . We can then use the data to extrapolate the curve to the $w = 0$ case. [INSERT PLOT.] Since we know $\xi(1) = 1$ we can expand about this point,

$$\xi(w) = 1 - \rho^2(w - 1), \quad (11.41)$$

and fit for the parameter ρ . There is a small unremovable theoretical uncertainty associated with this. The number that we get is something like 0.44. (HELP: CHECK THIS)

11.5.2 Inclusive decays

There is another way to get information about $|V_{cb}|$. Consider the decay $B \rightarrow X_c \ell \nu$, where X_c means *any* state containing a c quark. This is called an **inclusive decay**. There is a principle called **quark-hadron duality**³⁶ which says that information about inclusive hadronic decays tell us about the parton-level processes. For $B \rightarrow X_c$, for example, this duality tells us that

$$B \rightarrow X_c \ell \nu \approx b \rightarrow c \ell \nu. \quad (11.42)$$

The important symbol here is the ' \approx .' The B decays into something charmed and potentially a lot of other stuff. When we sum over all of this 'other stuff,' we say that this amplitude should be approximately the same as that of the quark-level $b \rightarrow c$ amplitude. Said in another way, one can predict the rates for inclusive hadronic processes by calculating the quark process.

It is important to stress that despite it's seemingly innocuous appearance, this quark-hadron duality is far from trivial. Consider the measurement of the famous ' R ratio',

$$R = \frac{\sigma(e^+e^- \rightarrow \text{hadron})}{\sigma(e^+e^- \rightarrow \mu^+\mu^-)}. \quad (11.43)$$

Recall that the plot of R over q^2 is one of the famous checks for the existence of quarks. Quark-hadron duality tells us that this plot should look like a series of step functions with a step at each quark's mass threshold, with some smoothing due to phase space. [INCLUDE PLOT]

If we compare this to actual experimental plots of the R ratio, we can immediately see the problem. [INCLUDE PLOT] There are lots of peaks associated with hadronic resonances which

³⁶This is a duality from the pre-string theory days, when we [phenomenologists] had more dualities than they did.

clearly do not appear in the quark-level analysis. For example, at the ρ resonance, $e^+e^- \rightarrow \rho \rightarrow \pi\pi$ has a huge cross section—much larger than one would predict from the naive quark-level diagrams.

Clearly there’s a subtlety in the quark-hadron duality principle. The subtlety is that we must *smear* out the data. For example, in *Mathematica* we can smooth out the data. What is the scale of the smoothing? Λ_{QCD} , of course! When we smear out features on this order and smaller—that is when we integrate over these features—we begin to follow the quark-level $b \rightarrow c\ell\nu$ plot.

To measure $|V_{cb}|$ we just have to plot the spectrum of $B \rightarrow X_c\ell\nu$, integrate over some region of q^2 , and pretend that we’re looking at a plot of $b \rightarrow c\ell\nu$. In the heavy quark limit it’s clear that the B decay really is the same as the b decay, so that to leading order we’re done. Of course, after developing HQET we can go further and discuss the corrections coming from λ_1 and λ_2 . [SEE LONG HOMEWORK PROBLEM ON THIS.]

Example 11.6. Operator dependence. Manohar and Wise originally calculated the $1/m^2$ corrections to $B \rightarrow X_c e \nu$ from the vector operator. Suppose you wanted to calculate the scalar operator, as Yuval and Zoltan Ligeti did when they were graduate students. How do you expect the λ_1 and λ_2 corrections to relate to those calculated by Manohar and Wise? The λ_1 prefactor is the same while λ_2 is different. This is because λ_1 is the correction coming from the kinetic energy of the heavy quark, which has nothing to do with the operator.

11.6 Measuring V_{ub}

Let us briefly discuss $|V_{ub}|$. The natural decays to consider are

$$B \rightarrow \pi\ell\nu \tag{11.44}$$

$$B \rightarrow \rho\ell\nu. \tag{11.45}$$

There’s an obvious problem: we cannot use HQET. In fact, there are problems with statistics since this process is CKM suppressed since $|V_{ub}| \sim 0.01|V_{cb}|$. The key here is to use **soft collinear effective theory**, which we do not have time to discuss further.

One idea we might have is to look at inclusive decays,

$$B \rightarrow X_u\ell\nu. \tag{11.46}$$

However, this doesn’t help either, since ‘ X_u ’ is not a well-defined signal experimentally—this would just be ‘hadronic junk.’ For example, an experimentalist can’t tell it X_u apart from X_c : if you see a charm then you know it goes in X_c , but if you don’t see a charm then it’s not necessarily true that there wasn’t one. (The X_u is a small theoretical error on the X_c inclusive decay.)

Fortunately, there’s a trick. We can use phase space. There’s a kinematic endpoint in the energy of the electron which is different for the X_c and X_u decays. The X_u decays allow higher-energy electrons. [SEE PLOT OF $d\Gamma/dE_\ell$, THE X_u LINE IS REALLY SMALL.] The point is that there’s a little area where the up overcomes the charm in the endpoint region. Experimentally this turns out to be do-able, but it’s still very hard.

12 Boxes, Penguins, and the CKM

We’ve taken a grand tour of the CKM matrix elements by looking at flavor-changing tree-level processes. One set of matrix elements have eluded us: $|V_{tb}|$, $|V_{ts}|$, and $|V_{td}|$. This brings us to the top.


12.1 The trouble with top

We know that the top is a special case. For one, we know that it doesn’t hadronize—but this is a *good* thing from the point of view of measuring CKM elements since it suggests that we can avoid the worrying about dreaded form factors. On the other hand, we also know that tops are hard to produce. To date (i.e. in the pre-LHC era) physicists have only produced about a hundred or so tops in the lab. Compare this to maybe billions of *bs* and trillions of *cs*. Experimentally tops are identified by tagging the associated *b* (looking at the invariant mass of the *W* and *b* jet) so that it seems hard to measure anything other than V_{tb} . Further, even if we could get around this—for example by looking at the mass of decay products—there would still be no way to tell V_{td} from V_{ts} .

What a mess. There are still a few things we can do. For example, we can compare measurements of the top with and without *b* tagging,

$$\frac{|V_{tb}|^2}{\sum |V_{ti}|^2} = 0.94 \pm 0.30. \quad (12.1)$$

Another way to do a CKM measurement is through **single top production**. Tops are more likely to be produced from gluons than *W*s, but gluon vertices don’t carry CKM elements. Fortunately, tops produced from a *W* can be distinguished from those produced by a gluon since gluon-borne tops come in particle–anti-particle pairs. The relevant diagram is



Example 12.1. Tevatron versus LHC. Given the initial states for the above diagram, what can we say about what kind of collider we should use to probe this? The Tevatron is a proton anti-proton collider so that there are *actual* up and anti-down valence quarks. Indeed, the Tevatron has produced tens of single tops to date. The LHC, which is a proton–proton collider, requires the anti-down to appear as a sea quark. Actually, this was a trick question: this process occurs at small *x* where the parton distribution functions become large so that there is actually an additional enhancement that makes the LHC not-so-bad at single top production.

This is all we have to say about top decay at tree-level. Unsatisfied? You should be. Fortunately, there’s a sneaky way to get to the top matrix elements: we can look at loop-level processes.

12.2 Loops, FCNCs, and the GIM mechanism

Now that we're passing from tree-level to loop-level processes, Yuval makes an obligatory reference to the circus trees in Gilroy which are grown in such a way that their trunks form loops³⁷.


Now that we've resigned ourselves to loop-level processes and the associated loop suppression factors, a natural set of processes to look for are flavor-changing neutral currents (FCNCs) since we already know that these vanish at tree-level in the Standard Model. A loop correction to a tree-level process is hard to observe, while a small effect which is the only source of signal has at least a chance of being measured well.

One of the well-known morals of flavor physics at loop level is the **GIM mechanism**, named after Glashow, Iliopoulos, and Maiani. The underlying principle is that for a unitary matrix, any pair of columns (or rows) are orthogonal. Most readers will already be familiar with the GIM mechanism from their courses and reading in particle physics. The first non-trivial point is that loop-level FCNCs are most sensitive to the *heavy* (!) quarks running in the loops, so that we can be sensitive to the top.

In loop-level processes particles go off shell, where we would expect the **Appelquist-Carazzone decoupling theorem** to hold [5]. This tells that when we take the mass of the internal particle to infinity, its effects on physics at much lower scales must vanish. But now we are saying that the heavy particles in FCNC loops do *not* decouple. In fact, the result of the amplitude is a function of the form $f(\frac{m_i^2}{m_W^2})$, which goes to a constant for large values of its argument. What's going on?

12.3 Example: $b \rightarrow s\gamma$

Now a classic example, $b \rightarrow s\gamma$, first observed at Cornell, about halfway through CLEO's lifetime. This was when Yuval was in grad school and his first paper was on this so-called **penguin diagram**. The curious etymology of this process is best explained in John Ellis' own words, as cited in Shifman's introduction to the ITEP Lectures in Particle Physics [85]. Here's an example of a penguin diagram,



The dependence of $b \rightarrow s\gamma$ on the CKM matrix takes the form

$$\mathcal{M} \propto \sum_i V_{ib}^* V_{is} = 0, \quad (12.2)$$

where we sum over the internal quarks, i . This product vanishes identically by unitarity. This tells us that any term independent of the internal quark mass must necessarily vanish. (We've

³⁷Yuval once spent fifteen minutes looking at photos of these trees in his office saying 'Wow! Come look at this!' to any student who passed by. Another story Yuval has about Gilroy is the city's annual garlic festival. When he was at SLAC Yuval took his family to the garlic festival and said that his son smelled like garlic for a week afterward and that "it was amazing."

drawn this as a chirality-flipping mass insertion in the diagram above.) We can conclude that

$$\mathcal{M} = \sum_i V_{ib}^* V_{is} f(m_i). \quad (12.3)$$

We can expand f in a power series. If $m_i \ll m_W$ then $f(m_i \ll m_W) \propto m_i^2/m_W^2$. The W indeed satisfies the decoupling limit. When we take m_i large (say $m_i = m_t$), however, then we claim that argument is not valid. Thus we know that $f(m_i^2/m_W^2)$ is linear at small values and constant at large values. The function f is called the Inami-Lim function, because they did all of these diagrams in the 1980s³⁸ [54].

Finiteness of penguins. The penguin amplitudes are manifestly finite in 4D. One naïve argument for this is that any divergences are independent of the internal quark mass (say, using dimensional regularization) so that unitarity kills any divergences. This, however, obfuscates deeper reasons why the amplitude is finite. One heuristic argument that the loop level must be finite since symmetries preventing FCNCs prevent any tree-level counter-term. (Note that this argument does not depend on renormalizability.) However, an even deeper appreciation for the finiteness can be found in an excellent paper [25]. The summary for 4D penguins is that there are two sources of suppression that reduce the superficial degree of divergence for these diagrams:

1. Gauge invariance, in the form of the Ward identity, requires the amplitude to depend explicitly on the external momentum. This is explained very clearly in [59].
2. Lorentz invariance prevents divergences which are odd in the loop momentum, k . In other words, $\int d^4k \not{k}/k^{2n} = 0$. Since the leading order contribution to $b \rightarrow s\gamma$ is even in k , the next-to-leading term is odd and vanishes.

As a final remark on this, one can also argue that the *chiral structure* of this process provides a suppression mechanism. The penguin is mediated by a $\sigma^{\mu\nu} F_{\mu\nu}$ operator that requires an explicit mass insertion. This turns out to be equivalent to gauge invariance.

The reason why decoupling is violated is that they are essentially mediated by the longitudinal part of the W by the Goldstone Equivalence theorem. Thus the coupling to the fermions is essentially a Yukawa coupling, which goes like the fermion mass. When you have a particle whose coupling is proportional mass, then it is clear that decoupling fails. Thus the GIM mechanism tells us that one-loop diagrams carry factors of m_i^2/m_W^2 , where i is summed over the internal quarks.

12.4 History of the GIM mechanism

In the early days of the quark model, physicists thought there were only three quarks u, d, s . $SU(3)$ flavor symmetry seemed to describe the light hadronic states well and there wasn't any motivation for the fourth quark. The GIM mechanism gave a motivation: it is a way to avoid FCNCs at

³⁸It wasn't until the 80s that anyone really believed that the top quark might be so heavy!

tree-level. Recall that if someone with no knowledge of the Standard Model stared at the PDG for a long enough time, that person might realize that it was very curious that neutral current processes like $K_L \rightarrow \mu^+ \mu^-$ were heavily suppressed relative to the charged current processes. In the 60s people couldn't reach the precision to see this decay while it's charged current cousin was readily measurable. GIM gives a way to understand this: the neutral current process is suppressed by a loop factor and by a GIM factor,

$$\mathcal{M} \sim \frac{g^2}{16\pi^2} \frac{m_c^2}{m_W^2}. \quad (12.4)$$

Next Gaillard and T.D. Lee calculated at $K\bar{K}$ mixing, keeping and dropping factors of two somewhat haphazardly as theorists are wont to do. Their calculations gave a prediction the mass of the c , 1.5 GeV which is remarkably close. It turns out that (like many good theorists) Gaillard and Lee had luck on their side, since if one follows their calculation somewhat more honestly, one obtains a value between 0.5 and 10. The charm was finally observed in 1974 at SLAC and Brookhaven.

An interesting question is which quark dominates the $b \rightarrow s\gamma$ penguin? It seems we have already spoiled the answer since we said that heavy quarks dominate in penguin loops, but we should be careful since there's also a CKM suppression associated with the top so that there is a competition between the large mass and the small CKM elements. Comparing the loops with an internal top and charm,

$$\mathcal{M}_t \sim m_t^2 V_{td} V_{ts} \quad (12.5)$$

$$\mathcal{M}_c \sim m_c^2 V_{cd} V_{cs}. \quad (12.6)$$

It turns out that the charm wins but that the top amplitude is in the same ballpark. So in this sense Gaillard and Lee got a little bit lucky once again, since back then nobody would have believed that the top was so heavy that it might challenge the charm contribution.

Example 12.2. Up-type versus down-type quarks. What can we say about $D \rightarrow \pi\gamma$? What about $t \rightarrow c\gamma$? These should all be small. The processes with external down-type quarks (and hence internal up-type quarks) carry factors of m_t^2/m_W^2 which are not small at all. For processes with external up-type quarks, on the other hand, the GIM mechanism is very efficient because $m_{d_i}^2/m_W^2$ is small for any down-type quark d_i . In other words, the fact that the s and b are much lighter than the c and t means that FCNCs are more visible in the down quark sector.

Example 12.3. FCNCs in the lepton sector. GIM is *most* effective in the lepton sector. Given that neutrinos have mass, we would expect processes like $\mu \rightarrow e\gamma$ (the leptonic analog to $b \rightarrow s\gamma$). In the Standard Model rate goes like $m_\nu^4/m_W^4 \sim (10^{-1}/10^{11})^4 \sim 10^{-48}$. Loop factors and couplings give a few more powers, leading to around 10^{-52} . GIM effectively kills leptonic FCNC at one loop.

12.5 Measuring the $b \rightarrow s\gamma$ penguin

One interesting question in $b \rightarrow s\gamma$ is the energy spectrum of the emitted photon. In $\mu \rightarrow e\gamma$ the spectrum is trivial, it is just a line at the only kinematically allowed energy. For $b \rightarrow s\gamma$, on the other hand, hadronization allows an actual shape since the actual measurement is $B \rightarrow X_s\gamma$.

Quark hadron duality tells me that to a good approximation I can just say that I have an s quark and the photon spectrum is just a line. But this approximation relies on smearing. When you cannot smear then the approximation is not good. In fact, the spectrum is far from the delta function. The spectrum is called a **shape function**. This is a truly non-perturbative object that we do not know from first principles. (Even the lattice doesn't come close right now—it's inclusive, and lattice at best can do one or two hadrons.) The one thing that *is* good—remember V_{ub} which relied at the end of the spectrum—is that the shape function in $b \rightarrow s\gamma$ can be related to V_{ub} . These hadronic unknowns can at least be related to one another so that one measurement can be used as input into other measurements.

Given the distribution in photon energies, the reality of experiments requires us to put a cut on the photon energy, bounding it from below. These cuts introduce problems. If we want to measure the $b \rightarrow s\gamma$ branching ratio and we can only measure part of the spectrum, then we have to extrapolate to get the rest of the spectrum. Since the spectrum is pure QCD this is very hard to do with any confidence. Thus the calculation must also be done with a cut; and the lower the cut the more we trust the calculation. Current lower limits from BaBar and Belle go down to about 1.6 GeV. As we push the cut lower and lower we increase the statistics and reduce theoretical error, but also we also pick up more background and lose purity. Eventually we get a number for the inclusive decay $B \rightarrow X_s\gamma$ which is published with an explanation for how this cut is taken.

12.6 Measuring $b \rightarrow s\gamma$ versus $b \rightarrow d\gamma$

How can we distinguish $b \rightarrow s\gamma$ from $b \rightarrow d\gamma$? Consider the ratio

$$\frac{\text{Br}(B \rightarrow D\gamma)}{\text{Br}(B \rightarrow X_s\gamma)} \sim \left| \frac{V_{td}}{V_{ts}} \right|^2 \sim \lambda^2. \quad (12.7)$$

It would be very nice if we could measure this ratio of CKM elements. However, we cannot experimentally distinguish between X_s and X_d . It's the same old story: “things which are very nice theoretically tend to be hard to do experimentally, and if you can measure something experimentally, you don't really know what you're doing theoretically.”

Another approach to go is to go make exclusive measurements. We could, for example, look only for

$$B \rightarrow K^*\gamma. \quad (12.8)$$

Note that angular momentum conservation prevents us from considering $B \rightarrow K\gamma$. The problem with $B \rightarrow K^*\gamma$ is that the branching ratio is much smaller: it turns out to that only roughly 20% of the s quarks hadronize into K^* s. This is nicer, however, because we once again have a definite γ energy from kinematics. The theoretical problem is calculating the hadronic part? We need form factors,

$$\langle B | \mathcal{O} | K^* \rangle. \quad (12.9)$$

At least we can compare this to $B \rightarrow \rho$, allowing us to determine the ratio

$$\left| \frac{V_{td}}{V_{ts}} \right|^2 \sim \lambda^2. \quad (12.10)$$

Comparing $B \rightarrow \rho\gamma$ and $B \rightarrow K^*\gamma$ gives this ratio to $SU(3)$ flavor symmetry breaking effects since the form factors (and hence a big chunk of the hadronic uncertainties) cancel in the $SU(3)$ limit. Ultimately small statistics limit the usefulness of $B \rightarrow \rho\gamma$ (the ρ is miserable to identify because it is so wide), but at least it's an independent measurement.

Temporary: *these will eventually show up in the operator mixing section of these notes.*

Five more minutes: the theory of $b \rightarrow s\gamma$. One of the hardest calculations that's ever been done "really" using perturbation theory. $g - 2$ is one of them, $b \rightarrow s\gamma$ is the other one. The point is that naively all you have to do is calculate these trivial diagrams. This only gives \mathcal{O}_7 . When we run to the b scale we need other operators, e.g. \mathcal{O}_8 . From the W scale down to the b scale these mix. The effect is huge. You'll be off by a factor of six if you miss the mixing (factor of six enhancement from mixing).

Simple way to think about it. What does it mean to do the running? Just take into account higher order loop. What happens if you add a gluon to \mathcal{O}_7 and a QED correction to \mathcal{O}_8 . The two are the same. Just heuristic to show that they really should mix.

The point is: once you have an effect that, at one loop, changes the result by a factor of six, then you start to worry about two loop. In fact, you have to start going to higher loops. (In fact, the 'tree level' was already one-loop.) At three level there's a 10×10 matrix of anomalous dimensions. The Swiss do it even better than the Germans. It takes a long time to do the calculation, end up with some un-removable uncertainty associated with setting the scale. Brings you down to 5% error and probably cannot do much better theoretically. This is a whole industry, people spend their entire careers doing this.

Example 12.4. Chiral enhancement of $b \rightarrow s\gamma$. This is one of the most promising low-energy processes in the search for new physics. This is because $b \rightarrow s\gamma$ is particularly sensitive to charged scalars; let's consider the charged Higgs in a 2HDM as a concrete example. The coupling of the charged Higgs is proportional to the Yukawa so that it's small in b decays. In $b \rightarrow s\gamma$, however, there is an interesting enhancement that makes it dominate over many other b decays. Consider, for example, $b \rightarrow c\tau\nu$ ³⁹. [Draw SM diagram with vector and with charged Higgs.]

How relevant is the charged scalar relative to the vector? The bcH^\pm and $\tau\nu H^\pm$ couplings go like $m_b/v \tan\beta$. For the W -mediated decays one has couplings which go like $g/\sqrt{2}$. In the approximation where m_W and m_H are the same the ratio of the vector to scalar amplitudes is

$$\frac{g^2}{y_b y_t \tan^2 \beta}, \quad (12.11)$$

so that it's not usually easy to probe the charged Higgs. The point, however, is that the charged Higgs diagram it is a *tensor* operator,

$$m_b \sigma_{\mu\nu} F^{\mu\nu}, \quad (12.12)$$

where the 'mass insertion' m_b is required because the $\sigma^{\mu\nu}$ connects opposite-chirality states. There's no way to generate vector operator since $\bar{s}\gamma_\mu b A^\mu$ is not gauge invariant. In the charged

³⁹Yuval calculated the QCD correction to this process—to good approximation it "got him his first postdoc."

Higgs diagram we get the m_b insertion ‘automatically’ from the Yukawa coupling. The vector diagrams, on the other hand, require an additional mass insertion giving a m_b/m_W suppression and enhancing the relative importance of the charged Higgs diagrams that is largely independent of $\tan\beta$. (**I don’t understand this.**) This allows us obtain a bound on the mass of the charged Higgs, $m_{H^\pm} > 300$ GeV in the limit where we can neglect other diagrams. (In SUSY the other penguin diagrams weaken this bound.)

13 Meson Mixing and Oscillation


Mixing and oscillation are two of the key ideas in flavor physics and always come hand-in-hand. As a point of nomenclature, it is useful to distinguish between the two very closely related ideas. **Mixing** is a theoretical statement that tells us that the flavor eigenstates are different from the mass eigenstates. For example, in SUSY we know about the mixing of the gauge-eigenstate bino and winos into the charginos. In mesons we have the flavor and the mass states. **Oscillation**, on the other hand, is the time evolution of a flavor state into another flavor state because they do not match the energy eigenstates.

Given a superposition of two Hamiltonian eigenstates that mix, the oscillation frequency is $\omega = \Delta E$ (where \hbar is one); in the rest frame this is just $\Delta E = \Delta m$. Thus oscillations are a way to use quantum mechanics to measure this Δm , even when the absolute mass scales are hard to probe directly.

The most famous example is kaon mixing between the $K = (\bar{s}d)$ and $\bar{K} = (\bar{d}s)$. In the absence of the weak interactions, these states have degenerate masses. The same thing is true about the $K^+ = (\bar{s}u)$ and $K^- = (\bar{u}s)$. In QCD the K and K^\pm states are protected by $U(1)$ flavor symmetries, a K does not mix with \bar{K} nor does a K^+ mix with a K^- .

The story changes once we turn on the weak interactions. The charged kaons still cannot mix since $U(1)_{\text{EM}}$ steps in place of the broken flavor symmetry to protect against the mixing. However, the K and \bar{K} are left unprotected and the *can* mix into one another, even in the Standard Model. Those that have been paying attention (and don’t already know the punch line) should be perturbed by this: while there’s no explicit flavor symmetry protecting against $K\bar{K}$ mixing, we also know that the Standard Model Lagrangian doesn’t have flavor-changing neutral currents. Since both the K and \bar{K} are neutral states, shouldn’t this flavor-changing neutral effect be prohibited?

No, though it is suppressed. At his TASI 2008 lectures, Herbi Dreiner once likened SUSY model-building to ‘sex in Scandinavia’ because “anything that isn’t explicitly forbidden is tacitly allowed.” This is true for flavor-changing effects: $K\bar{K}$ mixing isn’t *explicitly* forbidden by symmetries, and so we should expect them to happen. In this case, however, FCNCs are *so* scandalous that it is not practiced openly—that is, it isn’t advertised in the Lagrangian and is hidden at tree-level. The flavor-changing $K\bar{K}$ mixing occurs only at loop level due to box diagrams of the form



Marginally inappropriate references aside, what does this mean? Is it mysterious or magical that such an oscillation occurs? No, not at all. This is clear if we forget about the parton substructure of the kaons and just treat them as quantum fields in their own right. The K and \bar{K} are just states of a single complex scalar field. The box diagram above is then a two point amplitude between the field and its conjugate—but we know exactly what this is, this is just a loop-induced mass term. Once again, you can admonish such a thing and have it hidden from tree level, but unless you *explicitly* prohibit it, you can expect it to still happen.

The bottom line is that everything can mix if they have the same unbroken quantum numbers: $SU(3)_c$, $U(1)_{EM}$, and Poincaré. Though in the latter case spin-0 and spin-1 states can *effectively* mix when placed in a system that breaks Lorentz invariance, for example axion-photon mixing in a magnetic field.

13.1 Open system Hamiltonian

The game that we must now play is to start with a state which is a superposition of the Hamiltonian eigenstates and let it evolve in time. There are a lot of subtleties when doing this in quantum field theory which are well-described in the literature. Following Einstein’s famous adage that physics should be as simple as possible—but *no simpler*—we’ll sidestep many of the subtleties and present the most straightforward treatment which accurately represents the physics.

It is convenient to distinguish between states which are **asymptotic** and those which are **resonances**. In QFT we usually work with asymptotic states, which to zeroth order in perturbative interactions have don’t decay and go off to infinity. All decays of these asymptotic states come from interactions which we treat in perturbation theory (it is crucial that the relevant physics is perturbative). Alternately, we can describe the *same* physics with states which have finite lifetime even at zeroth order in perturbation theory. In other words, we include some of the interactions into our definition of the field. A very rough analogy is the difference between the Schrödinger, Heisenberg, and interaction picture in quantum mechanics, where different amounts of time evolution are packaged in the definition of an operator. As we shall see, for meson oscillations it is much easier to work with resonances rather than asymptotic states.

In regular quantum field theory the Hamiltonian is basically the mass of the particle, $H = M$. For a resonance, the Hamiltonian is

$$H = M - \frac{i\Gamma}{2}. \tag{13.1}$$

This is not Hermitian! Are you worried? We never worked with non-Hermitian matrices in quantum mechanics because probabilities are conserved, everything we measure is real, and Hermitian matrices have real eigenvalues. Of course (13.1) is perfectly fine. If a particle can decay then we shouldn’t probability to be conserved. For example in addition to all of the diagrams that

contribute to an S matrix element for given incoming and outgoing states, we must also include the processes where those specified incoming and outgoing states decay during their free evolution. In other words, we are working with an **open system** in which the ‘interaction’ with the system allows the particle to decay. We have taken all of the decay processes out of the system so that the ‘effective’ system that we are work with is no longer Hermitian.

When we have a system with more than one eigenstate then the Hamiltonian is a matrix. Let us thus consider the two state $K\bar{K}$ system where the Hamiltonian is a 2×2 matrix; in the flavor basis it takes the form

$$H \sim \begin{pmatrix} m_K & h_{\text{FCNC}} \\ h_{\text{FCNC}} & m_K \end{pmatrix}. \quad (13.2)$$

We can diagonalize this matrix to obtain the masses and the mass eigenstates, which we call the K_L (‘ K -long’) and K_S (‘ K -short’). The names describe the different lifetimes, so you can already guess that the two states not only have different masses, but different widths. We could have alternately named the states based on their masses, which is what is done for the B system: B_H (‘ B -heavy’) and B_L (‘ B -light’). Actually, for the B meson the relevant sign in the mass splitting has not been measured. **Reminder: Yuval should update this paragraph, how do they measure $\Delta m \times \Delta\Gamma$ in the K ?**

We’ve seen that weak interactions generate a mass term between the K and \bar{K} . CPT tells us that the flavor-conserving Hamiltonian must be the same: $H_{11} = H_{22}$. Further, Hermiticity tells us that

$$M_{12} = M_{21}^* \quad (13.3)$$

$$\Gamma_{12} = \Gamma_{21}^*. \quad (13.4)$$

Diagonalizing H gives the two mass eigenvalues and the two widths. Before doing any work, lets estimate the important quantities, Δm and $\Delta\Gamma$. One should be able too look at the matrix and very intuitively identify

$$\Delta M \sim 2h_{\text{FCNC}}. \quad (13.5)$$

Further, you should see right away that the mixing is 45° .

Example 13.1. Diagonalizing 2×2 matrices. Given a symmetric 2×2 matrix,

$$\begin{pmatrix} a & b \\ b & c \end{pmatrix} \quad (13.6)$$

one should remember that the mixing angle θ is given by

$$\tan 2\theta = \frac{b}{a - c}. \quad (13.7)$$

Note that we get a 45° mixing even when the m_{FCNC} is very small. Doing a full calculation gives, e.g. for the B system,

$$|B_{L,H}\rangle = p|B^0\rangle \pm q|\bar{B}^0\rangle \quad (13.8)$$

subject to the constraint

$$|p|^2 + |q|^2 = 1. \quad (13.9)$$

Look very carefully at (13.8). Usually in a rotation between states one has

$$|A_1\rangle = \cos\theta|B_1\rangle + \sin\theta|B_2\rangle \quad (13.10)$$

$$|A_2\rangle = -\sin\theta|B_1\rangle + \cos\theta|B_2\rangle. \quad (13.11)$$

Our expression for $|B_{L,H}\rangle$ is very different! Usually a rotation from one basis to another preserves the vector length. This is no longer true here since the Hamiltonian is not Hermitian. Another consequence of working in an open system is that the overlap of the two eigenstates is no longer zero, but

$$\langle B_L|B_H\rangle = |p|^2 - |q|^2. \quad (13.12)$$

The two eigenstates are not really orthogonal! Remembering linear algebra from kindergarden, we know that we only need *independent* vectors to span a space, not necessarily orthogonal. This non-orthogonality is called **vacuum regeneration**, a *mass* eigenstate can propagate for a while and then turn into another mass eigenstate! This is different from oscillation, where you start not with something that is not a mass eigenstate.

Let us now consider the case where CP is conserved. CP guarantees that $|p| = |q|$ (though this condition can hold even when CP is violated). In this case $\langle B_L|B_H\rangle = 0$. The Hamiltonian is *still* not Hermitian. The eigenvalues are

$$\mu_\alpha = M_\alpha + \frac{i}{2}\Gamma_\alpha \quad (13.13)$$

where $\alpha = 1, 2$. Note that we must diagonalize the entire Hamiltonian; this is *not* generally the same as separately diagonalizing M and Γ . (The diagonalizations are equivalent only in the CP-conserving limit.) Defining $2\Gamma = \Gamma_1 + \Gamma_2$, the CP-conserving limit gives us

$$\Delta m \stackrel{\text{CP}}{=} 2|M_{12}| \quad (13.14)$$

$$\Delta\Gamma \stackrel{\text{CP}}{=} 2|\Gamma_{12}|. \quad (13.15)$$

We can define three physical parameters out of this:

$$x = \frac{\Delta m}{\Gamma}, \quad y = \frac{\Delta\Gamma}{2\Gamma}, \quad \phi = \arg(\Gamma_{12}^* M_{12}). \quad (13.16)$$

Note the range of values for the x and y parameters: $x \in [0, \infty)$, $y \in [-1, 1]$. In the CP-conserving limit $\phi = 0$.

13.2 Time evolution

Consider the wavefunction $B^0(t)$ which, at time $t = 0$, is some linear combination of flavor eigenstates $|B\rangle$ and $|\bar{B}\rangle$,

$$B^0(0) = a(0)|B\rangle + b(0)|\bar{B}\rangle. \quad (13.17)$$

To evolve this in time we write B and \bar{B} in terms of energy eigenstates B_H and B_L and independently evolve each. We can then project back into the B and \bar{B} basis at any later time.

Let us write $B^0(t)$ to mean a state which is a pure B at $t = 0$. At $t \neq 0$ it is a mixture of B and \bar{B} . Still assuming CP for simplicity, we obtain

$$B^0(t) = \cos\left(\frac{\Delta E t}{2}\right) |B\rangle + i \sin\left(\frac{\Delta E t}{2}\right) |\bar{B}\rangle. \quad (13.18)$$

We'll generalize this to the CP violating case later in this course. ΔE is complex when we have an open system, but let us further assume that $\Gamma = 0$ so that we really do have oscillation with no decay. The probability of $B(t)$ to remain a B at time t is

$$P[B(t) = B] = |\langle B^0(t) | B^0 \rangle|^2 \quad (13.19)$$

$$= \frac{1 + \cos(\Delta E t)}{2} e^{-\Gamma t}, \quad (13.20)$$

where we've tacked on the decay rate post-facto. We see that the B system is indeed oscillating—there's the cosine! The frequency is ΔE , as advertised, so that a measurement of this oscillation frequency tells us Δm . The experiment that we want to do is to produce a B (flavor state) at $t = 0$ and then measure its flavor at some later time.

What is the t ? In neutrino experiments t depends on the distance from your source to the detector, but this is because neutrinos are so long lived. Let's focus on B s which have much shorter lifetimes. In this case t is on the order the lifetime of the B . What kind of Δm s can we probe? We can probe things on the order of Γ , the width of the particle. The width of the particle has to do with the three-body decay of the B meson with some CKM matrix elements. The lifetime δm has to do with some box diagram with some completely different structure. It's very lucky the lifetime is of the same order as the width so that $x \approx 0.7$. This is *completely* luck.

Example 13.2. If, for example, the B mass were doubled, then its lifetime would decrease by 2^5 (recall that $\Gamma \sim m^5$) and $x \ll 1$.

If t is something we can control experimentally, then we can probe the width. But for the decays of mesons, t is something which is given to us by nature. This is why the B system is so ideal for nature: it has *just* the right lifetime so that we can probe the width. [Show nice BaBar figure from Pat Burchat showing the oscillation versus decay amplitude.]

13.3 Flavor tagging

So we want to measure Δm . We have to produce an initial state which we *know* is a B at $t = 0$ and then determine the flavor of that state at some later detection time t . In other words, we need to **flavor tag** at both the production and decay times.

Nomenclature. The phrase ‘*b*-tagging’ is used differently in different collider communities. In supersymmetry this refers to whether or not a jet contains a *b* quark, while in flavor physics we are talking about actual *B* and \bar{B} mesons.


How do we perform these tags? Let’s discuss the easier tag: once we have some unknown state propagating, how can we determine what it is at time *t*, when it decays? The easiest way to tag a *b* quark decay is through the semileptonic decays,

$$b \rightarrow c\mu^-\bar{\nu} \quad (13.21)$$


$$\bar{b} \rightarrow \bar{c}\mu^+\nu. \quad (13.22)$$

By the way, the *B* meson nomenclature is very silly: a *B* meson contains an *anti-b* quark. Sometimes a *B* will decay into final states that are also accessible to the \bar{B} , making it useless for tagging whether it contained a *b* or \bar{b} , for example $B \rightarrow \pi^+\pi^-$. In this case the decaying object is really a superposition.

Example 13.3. Can you use $B \rightarrow D^+\pi^-$ for *b* tagging? Sure, here’s the diagram.



Can we thus conclude that has to be a \bar{B} meson? On the other hand, we could also draw a diagram with an initial *B* meson,



Note the arrows, this second diagram mediates $\bar{B} \rightarrow D^+\pi^-$. The difference is CKM suppression. So this observation isn’t a pure *B*, it’s an almost-pure *B* since there is some interference at the percent level. There is a second way to tag this process; the *D* can be ‘kaon’ tagged; at the inclusive level the charge of the kaon determines the flavor of the *B*.

Sticking to the semileptonic decays where the relatively easy-to-measure muon versus anti-muon *b* tags the final state, we must now find a clever way to *b* tag the initial state. The trick is to use EPR quantum entanglement. At the *B*-factories an electron-positron collider is operated at the $\Upsilon(4S)$ resonance. These mesons have *just* the right rest energy to decay into a $B\bar{B}$ pair nearly at rest. Since these particles are entangled upon production, we may use the decay of one of them as an initial-state *b* tag for the other.

Example 13.4. Without looking at the PDG, what is the spin of the $\Upsilon(4S)$? It has to be a spin-1 (vector) meson since it is produced through $e^+e^- \rightarrow \gamma \rightarrow \Upsilon(4S)$.

As soon as one B decays semileptonically, then at that instant we know precisely the b content of its sibling. That is the moment which we have been calling $t = 0$ when we know the exact flavor content of the propagating state. From that moment on the remaining state oscillates and then identifying its decay allows us to probe the oscillation. If, for example, we tag a state that is initially a B and we subsequently measure a semileptonic decay that could only have come from a \bar{B} , then we know that the system oscillated.

BaBar and Belle are *asymmetric* B factories. This means that the center of mass frame is boosted relative to the lab frame so that the $B\bar{B}$ pair is not produced at rest. Suppose that at some time $t = 0$ one meson decays into a μ^+ , telling us that it was a B meson. Thus at $t = 0$ we know that the other meson is a \bar{B} . Then we can look at the decay to flavor tag the final state and then use the decay position (through vertex detectors) as a measurement of the decay time. Note that the boosted system is important for being able to determine the decay time since it allows experimentalists to measure the physical distance between the vertices. The smoking gun signal for oscillations are same-sign leptons.

The last decade was fantastic for B physics, since in addition to the B factories, the Tevatron also made contributions to B physics. Flavor tagging at high energies is a different story since the $B\bar{B}$ pairs are uncorrelated. At the B factories, it was crucial that we knew the angular momentum of the state; the fact that the $\Upsilon(4S)$ has $L = 1$ gave us the correlation. At high energies we can have many different angular momenta which mix and causes decoherence. Thus in order to tag the B meson (as opposed to other hadronic junk) one has to fully reconstruct the invariant mass of the hadronic shower to identify the B mass. To further tag the b versus \bar{b} quark, one must look at events where the entangled b quark hadronizes into a *charged* B meson since we know that $U(1)_{EM}$ prevents these guys from oscillating. This plays the role of our initial-state flavor tag since the flavor indicated by the B^\pm decay must match the flavor of the B^\pm at its production. The decay of the oscillating B can then be subsequently tagged by looking at our favorite semileptonic decays. While this process may sound like a small portion of the total number of events, hadron machines produce so many bs that we can afford to be very picky about using only nice events.

13.4 Time scales

All of physics is based on the fact that there are different scales. In this case, it is important to ask about the relevant time scales for oscillations. Rewriting (13.20) in terms of $x \equiv \Delta m/\Gamma$

$$P(B^0(t) = B) = \frac{1}{2} (1 + \cos \Delta mt) \quad (13.23)$$

$$\frac{1}{2} (1 + \cos x\Gamma t). \quad (13.24)$$

What is t ? In neutrino or atomic experiments, you can change the propagation time t by placing the detector somewhere. In mesons, on the other hand, t is given to you by nature: it's particle's decay time, $\langle t \rangle \sim 1/\Gamma$. This is, of course, a statement about the average t ; each individual meson decays probabilistically. We see, however, that the dimensionless parameter controlling the

oscillation time scale is x . As always, we can learn about the system by considering asymptotic states.

First consider $x \ll 1$. In this case there are effectively no oscillations. Once the state is produced it simply doesn't have enough time to oscillate. It is produced as a flavor state and—for its short lifetime—it remains a flavor state. An example of this kind of state is the D meson.

Now consider the opposite extreme, $x \gg 1$. In this case the oscillations are washed out since the cosine averages out to zero. What is happening physically? Any measurement of the meson system is done over a small characteristic time scale. In the limit $x \gg 1$, the meson is able to oscillate many times over this measurement period so that the effect of the oscillation is washed out to zero. The natural objects to discuss in this scenario are thus mass eigenstates. This is the case when we produce kaons; we say that the kaon we produce is 50% K_L and 50% K_S .

We are used to $x \gg 1$. This scenario is what we are most familiar with since it's often an implicit assumption. In supersymmetry we can produce a bino and say that it is really some probability of $\chi_{1,\dots,4}$ mass eigenstates. This statement already contains the assumption that $x \gg 1$ so that we integrate over the quick oscillation time. Even in the Standard Model, we can discuss $W \rightarrow cs$ versus $W \rightarrow cd$ decays. The probability of one decay versus another is the square of the relative CKM matrix element, but this is also an assumption about the oscillation frequency.

Finally, consider the case $x \sim 1$; this is the most interesting case since oscillations are actually observable. The state is neither a flavor nor a mass eigenstate, but we really have a state which evolves with time. We can go and measure the oscillation of B and K mesons experimentally and find that

$$x_B \sim x_K \sim 1. \tag{13.25}$$

In fact, $x_B \approx .73$ and $x_K \approx .94$. As we mentioned before, this is complete luck since x depends on the CKM and the masses in some non-trivial way. Other similar mesons are not so fortunate. For example,

$$x_s \sim 25 \tag{13.26}$$

$$x_D \sim 10^{-2}, \tag{13.27}$$

where x_s is standard notation for x_{B_s} . For these guys the width of the meson's weak decay is much smaller than its mass,

$$\Gamma \sim G_F^2 m \sim 10^{-14} m. \tag{13.28}$$

Again, the 'practical' use of these oscillations is to make precise measurements of Δm using quantum mechanics, similar to the Mössbauer effect.

Let's return to our other dimensionless oscillation parameter

$$y = \frac{\Delta\Gamma}{2\Gamma}. \tag{13.29}$$

Taking Γ back into account, our oscillation probability formula becomes

$$P(B^0(t) = B) = \frac{e^{-\Gamma t}}{2} [\cosh(y\Gamma t) + \cos(x\Gamma t)]. \quad (13.30)$$

One can easily check that this is an extension of the previous formula.

Let us consider the some limits with respect to x and y .

- $|y| \ll 1$ and $y \ll x$. In this the effects of y are irrelevant. This is the case in the B system.
- $y \sim x$. Both x and y are important. For the kaon system $x, y \sim 1$ so that we can probe this scale experimentally.
- $y \sim 1$ and $x \gg 1$. Here there are two mass eigenstates with different decay rates (lifetimes) and we can forget about oscillations.

The following table shows the order of magnitude for the x and y values of some our favorite mesons,

	x	y
K	1	1
B	1	10^{-2} (*)
B_s	10	10^{-1} (*)
D	10^{-2}	10^{-2}

The stars indicate when a quantity is not well measured.

We explained above that the trick at B factories is to flavor tag the initial and final states. Because the collision rest frame is boosted relative to the lab frame, one can use the displaced vertices to determine the decay time. One can then fit with respect to the decay time to learn about x . What if you didn't have such an asymmetric B factory? That is, suppose you have an experiment—like CLEO—where you can't measure the decay time and all you can do is flavor tag for particle production. Effectively such an experiment is integrating over time. For simplicity, let's set $y = 0$ and integrate from $t = 0$ to $t = \infty$. In this limit we can easily perform the integral over the oscillation probability—by 'easily' we mean it is easy to plug into *Mathematica*—to obtain

$$\int_0^\infty P(B \rightarrow B) = \frac{1}{2} \frac{2 + x^2}{1 + x^2}. \quad (13.31)$$

This gives us a function that we may invert to determine x . As good theorists, we should check that this expression makes sense by looking at limiting cases. Indeed, our above discussion we know that in the limit $x \rightarrow \infty$ we expect $\int P = 1/2$ while in the limit $x \rightarrow 0$ we expect $\int P = 1$.

13.5 Calculating Δm and $\Delta\Gamma$

Now we should describe how to actually calculate ΔM and $\Delta\Gamma$. This really boils down to determining M_{12} and Γ_{12} . What about M_{11} and Γ_{11} ? These are just directly measured numbers that we don't care because they don't contribute to the mass or width difference.

We know that

$$M_{12} \sim \langle B | \mathcal{O} | \bar{B} \rangle. \quad (13.32)$$

The only non-trivial thing is the normalization, which is

$$M_{12} = \frac{1}{2m_B} \langle B | \mathcal{O} | \bar{B} \rangle. \quad (13.33)$$


The operator \mathcal{O} is some $(V - A)$ current representing weak decays. Ignoring the Dirac structure, the field content takes the form

$$\mathcal{O} \sim (b\bar{d})(b\bar{d}). \quad (13.34)$$

Check that equation one more time! Does it look weird? Indeed, one might have wanted to write down $(b\bar{d})(\bar{b}d)$ since, after all, we're used to mass operators formed out of a field and its conjugate. This is not an ordinary mass term, it represents a flavor-changing coupling between the B and \bar{B} , and the operator that we might have wanted to write down is manifestly flavor conserving. The 'weird' operator, on the other hand, *is* flavor-violating: $\Delta B = 2$ and $\Delta D = 2$. This is because the mass term indeed connects a particle and an anti-particle, but now because we're mixing B and \bar{B} s, the 'particle' is the B and the 'anti-particle' is an anti- \bar{B} , i.e. also a B .

Example 13.5. A toy example. Before going to the Standard Model, let's work with a simpler model with vector-like quarks that induce tree-level FCNC through a Z boson. Call this coupling κ . The diagram for M_{12} is just the usual s -channel tree diagram and goes like $\frac{\kappa^2}{m_Z^2}$. Stop and check this again: should we have written $\kappa^\dagger \kappa$ instead of κ^2 ? No! Check the diagram to make sure that *both* vertices are given by the *same* term in the Lagrangian, not the term and its Hermitian conjugate.

In the Standard Model, the $B\bar{B}$ mixing diagram occurs at one-loop level as a box diagram. At two-loop order one may also have 'twin-penguin' diagrams which are subdominant (though this is somewhat subtle). The two box diagrams are:



The calculation for each diagram is identical in the limit where all internal particles are heavy, but they're independent diagrams; practically we can just calculate one of them and multiply by two. These diagrams are manifestly finite by power counting. We can parameterize the amplitude straightforwardly just by using power-counting and the GIM mechanism,

$$\mathcal{M} \sim \frac{g^4}{m_W^2} \langle B | \mathcal{O}_{v-A} | B \rangle \sum V_{id} V_{ib}^* V_{jd} V_{jb}^* F(x_1, x_j) \quad (13.35)$$

where

$$x_i = \left(\frac{m_u}{m_W} \right)^2. \quad (13.36)$$

For the B meson calculation the dominant contribution comes from an internal top quark. For the kaons there's a competition between the CKM matrix elements and the mass insertions required by chirality and the GIM mechanism. It turns out that for the real part of the amplitude the internal charm is the winner, though the imaginary part is dominated by the top. Recall that the optical theorem tells us that the imaginary part of the amplitude is related to the decay rate.

Now all we need is the hadronic matrix element, our old friend. We are faced with what seems like a rather intractable hadronic object, $\langle B|(b\bar{d})(b\bar{d})|B\rangle$. We have a trick up our sleeve—but like all tricks up someone sleeve, it's somewhat shady. We start by inserting a sum over a complete set of states,

$$\langle B|(b\bar{d})(b\bar{d})|B\rangle = \sum_n \langle Bb\bar{d}|n\rangle \langle n|(b\bar{d})|\bar{B}\rangle. \quad (13.37)$$

The trick is to assume that the vacuum is the most important contribution in the sum and hence approximate the sum with only the vacuum state, $\sim |n\rangle \langle n| \approx |0\rangle \langle 0|$. Thus in the so-called **vacuum insertion approximation** (VIA) we have

$$\langle B|(b\bar{d})(b\bar{d})|B\rangle \approx \langle B|(b\bar{d})|0\rangle \langle 0|(b\bar{d})|\bar{B}\rangle. \quad (13.38)$$

The great thing about this approximation is that we already know what $\langle 0|(b\bar{d})|\bar{B}\rangle$ is: it is just the B decay constant,

$$\langle B|(b\bar{d})(b\bar{d})|B\rangle \approx f_B^2. \quad (13.39)$$

We've been characteristically hand-wavy since we neglected the Dirac structure of the operator. To do this properly one should invoke Fierz identities. When we do this properly we should use the Fierz identities which give us some additional factors,

$$\langle B|(b\bar{d})(b\bar{d})|B\rangle = \frac{8}{3} f_B^2 B_B. \quad (13.40)$$

The prefactor includes a factor of two from the sum over the two box diagrams and an additional $1 + \frac{1}{3}$ color factor. We've included a parameterization of the correction to the vacuum insertion approximation, B_B . It turns out, however, that for the B meson B_B is only at the order of a few percent. To be clear, there is no *a priori* reason for this approximation to work very well and the agreement is rather surprising. Now that we are armed with this, we can proceed to cross-check the CKM matrix.

The unbelievably heavy top quark. Wagner once wrote a paper in the 80s that assumed the top mass to be around 50 GeV, for which it was promptly rejected by the editor as being unreasonable. When you put the 50 GeV top into the above calculation you predict that B

mixing is very small. In the early 80s flavor physics found that B mixing is, in fact, order one. The natural explanation was that the top was heavy, and indeed, flavor measurements in 1981 suggested $m_t \sim 150$ GeV. People didn't believe this because it was so ridiculously large. It wasn't until much later that electroweak precision tests predicted the same value. Historically people often say that electroweak precision experiments predicted a heavy top, but it was in fact $B\bar{B}$ mixing that was the first avatar of a heavy top—people just weren't ready to believe it!

The next object to calculate is Γ_{12} . The off-diagonal mass M_{12} came from $\langle B|\mathcal{O}|\bar{B}\rangle$. In particular, it came from the *off-shell* contribution. When this matrix element goes on shell it gives a contribution to the width. This is because the real part of the two point function tells us about the mass while the imaginary part tells us about its decay. Intuitively, the Hermitian part of the Hamiltonian is off shell while the anti-Hermitian part goes on shell. For example, $K \rightarrow \pi\pi \rightarrow \bar{K}$ contributes to the width while $K \rightarrow B \rightarrow K$ with an off-shell B contributes to the mass. The intuition here is basically the **optical theorem**; if you go on-shell then you are really decaying and coming back, whereas if you stay off-shell then you're undergoing a two-point contact interaction.


Can we trust such a calculation? There's something fishy. We essentially perform a quark-level calculation. In the calculation of M_{12} we tacked on hadronic effects in f_B . We were off-shell with large q^2 so that our interactions are local and we could use effective operators. In the case of the decay width the hadronic final states make a difference: the actual processes we care about are $K \rightarrow$ hadrons and $\bar{K} \rightarrow$ hadrons. These things occur on shell, the quarks *really* hadronize and we have to take this into account in our kinematics. By doing the quark level calculation, we are implicitly invoking quark-hadron duality. However, the condition for quark-hadron duality is that we can average over a large q^2 , i.e. we can smear over resonances. Here we have $q^2 \rightarrow 0$; we only have *one* kinematic point—the mass of the decaying particle. Unlike deep inelastic scattering, there is no smearing. *A priori* it is hard to predict the validity of our calculation given these conditions. We can hope that it might still be applicable to the B , but we already know that it is not so good for the kaons.

A lesson: flavor physics is imprecise. This summer there was a lot of excitement about a large $\Delta\Gamma$ observation from D0 that seemed to suggest new sources of CP violation. One key lesson in this section on loop processes is that one has to be very careful about the theory assumptions going into the Standard Model calculation for such claims. As we shall see for the kaons, it is entirely possible to inadvertently carry an innocent assumption well beyond its region of validity.

Consider the kaon box diagram where only the u can go on-shell. Phase space tells us that we can only really produce no more than three pions. It's now trivial to see that quark-hadron duality doesn't work since a phase space which permits 'up to three pions' is clearly *not* and inclusive decay. In fact, the K_S mass eigenstate can go only to two pions, while the K_L can go to three. So y is basically one from phase space. $\Delta\Gamma$ is basically the difference between these rates. The kaon calculation is off by a large factor. In the B case it is probably correct within factor of 2.

In the D box diagram, which internal quark dominates? A quick back-of-the-envelope calculation invoking the ratios of quark masses and CKM elements in the Wolfenstein parameterization should convince you that the strange quark. Naïvely we could proceed to calculate this diagram just as we did for the B and K , but it turns out that this would indeed be too naïve [30]. In the B box we had a heavy t quark running in the loop. Now for the D box we have an s quark in the loop, which is *not* heavy and so it is really *not* okay to integrate it out into a local operator. If one does this calculation as if the s were much heavier than the W , then one obtains a result which is four orders of magnitude too small. It is *completely* wrong to make it a local operator; you can never integrate out an s when you're talking about c decay.

The correct thing to calculate is a diagram where only the W s are integrated out, giving the so-called **fish diagram** (because doing anything else would be ‘fishy’),



From this one obtains $x \sim 10^{-2}$. The lesson here is that phase space is not part of the operator product expansion and can hence lead to a violation of quark-hadron duality.

14 Today’s lecture: CP violation

A priori CP violation seems to have nothing to do with flavor physics. CP is a discrete group associated with the Lorentz group, while flavor has to do with internal symmetries. It just turns out that in nature all CP violating observables always come with flavor violation. By now we know that the only source of known CP violation is this one phase of the CKM matrix.

We look for some process that violates CP, which we write as \mathcal{CP} . We look for processes where the probability for one process is different from the probability for the CP conjugate of this process

$$P(A \rightarrow B) \neq P(\bar{A} \rightarrow \bar{B}). \quad (14.1)$$

Note you should distinguish this from T and CPT violation.

$$P(A \rightarrow B) \neq P(B \rightarrow A) \quad (14.2)$$

$$P(A \rightarrow B) \neq P(\bar{B} \rightarrow \bar{A}). \quad (14.3)$$

Since we believe CPT is a good symmetry, we often really distinguish between \mathcal{X} and \mathcal{CP} .

So what’s the big deal? If \mathcal{CP} in the Lagrangian level, it should show up all over the place. However, CP was only first observed experimentally in 1964 and 1967 in the kaon system, and not until 2000 in the B system. Parity violation, on the other hand, is trivial—we see it all the time. To date, we only have a handful of \mathcal{CP} examples.

The point is that if we have $\Gamma(A \rightarrow B) = 10\%$ and $\Gamma(\bar{A} \rightarrow \bar{B}) = 9\%$, then there must be another process such that, for example, $\Gamma(A \rightarrow C) = 50\%$ and $\Gamma(\bar{A} \rightarrow \bar{C}) = 51\%$. Probability is a ‘zero sum’ game. Thus in order to see CP violation, something has to be compensation. The reason for this is CPT, it forces the mismatch to show up elsewhere.

Open question (a free project to anyone who wants to do it): what is the CP violation in the Hydrogen atom? Given a hydrogen and an anti-hydrogen, what is the width of the 21 cm? Suspect it is something like 10^{-100} . The interesting question is where the CP violation enters.

At the end of the day we must have the interference between two amplitudes (diagrams) with a difference in the **strong** and **weak** phases. Usually for a process there is just a single diagram where the higher-loop corrections are extremely small.

The **weak phase** is a CP-odd phase; it changes sign for a CP conjugate process. This phase shows up in the Lagrangian. Given a coupling $V_{ub}W^+\bar{u}b$, then the Hermitian conjugate coupling is $V_{ub}^*W^-\bar{b}u$. The **strong phase**, on the other hand, is a CP-even phase; it does not change sign between a process and its CP conjugate. All the phases in the Lagrangian change sign, so where does this come from? Strong phases come from time evolution. For example, a free particle has a trivial time evolution, $\exp(iEt)$, independent of whether it is a particle or anti-particle. For a more complicated strongly-interacting system this becomes much less trivial. There is some complicated non-perturbative dynamics which, at the end of the day, give you some phase. In order for this to be important you really need a strong coupling: if you have two pions, what is the probability to scatter out to two pion. Another way to see it is going near a resonance; for example, two pion scattering near the ρ resonance. There’s a phase across in the Breit-Wigner resonance in classical scattering (same phenomena as a forced oscillation with friction). We will not discuss **unphysical phases** which can be rotated away by the field definitions.

In mesons we have two processes of primary interest in flavor physics: decays and oscillations. Given an amplitude and its CP conjugate, we may write the amplitudes as

$$A_f = A(B \rightarrow f) = a_1 e^{i(\delta_1 + \phi_1)} + a_2 e^{i(\delta_2 + \phi_2)} \quad (14.4)$$

$$\bar{A}_f = A(\bar{B} \rightarrow \bar{f}) = a_1 e^{i(\delta_1 - \phi_1)} + a_2 e^{i(\delta_2 - \phi_2)}, \quad (14.5)$$

where we write δ as the strong phase and ϕ as the weak phase. It’s very bad notation because the δ in our CKM parameterization is the weak phase. This partitioning into two terms is required by enforcing that the terms have the same coefficients and that the amplitude and CP-conjugate amplitude are related by simply flipping the sign of the weak phase. Similarly,

$$M_{12} = |M_{12}| e^{i\phi_M} \quad (14.6)$$

$$\Gamma_{12} = |\Gamma_{12}| e^{i\phi_\Gamma}. \quad (14.7)$$

It should be trivial that if $a_2 = 0$ then $A_f = \bar{A}_f$, similarly if $\delta_1 = 0$. Thus the conditions to see \mathcal{CP} are that you have two terms such that there are different weak phase and strong phase. At this stage these are still two *terms*, but we will see soon that these terms can be closely related to a tree-level diagram and a penguin.

Our goal experimentally is to find processes where $a_1 \approx a_2$ and there are different δ and ϕ . Note that a s and δ s are related to the strong interaction, only ϕ is related to the weak interaction. We return to the usual story: we have to deal with hadronic matrix elements. Here there are two good cases: either the hadronic matrix elements cancel or they can be directly measured. This was not the case for CP-conserving processes.

A nice question: if $A_f \neq \bar{A}_f$ we demonstrate CP violation. Is the converse true? Can we show that there is CP violation without actually directly measuring it? The analog is measuring the angles of a triangle without ever measuring any particular angle. In principle one can measure CP violation just by measuring the lengths of the sides of the unitarity triangle. By measuring a few decay rates and a few complex conjugate decay rates—such that they are all the same—and we could still demonstrate that there is a phase in the Lagrangian.

We’ve talked very generally about CP violation in mesons. Let us briefly mention two other measurements of CP violation outside of flavor physics. The first is the **triple product**, where in some decay one takes three vectors (e.g. p_1 , p_2 , and spin in a three-body decay) and generate some kind of triple product $(\mathbf{p}_a \times \mathbf{p}_b) \cdot \mathbf{s}$. These types of quantities turn out to be CP odd and sensitive to the sum of the strong and weak phase. At the LHC one can study CP violation with these sorts of triple products. The other possible measurement for CP violation are **electric dipole moments** (EDM). These are CP violating for elementary particles. When you take the dipole moment of such a particle then the CP conjugate of such a state must change since angular momentum doesn’t conserve CP. Note that the magnetic moment is CP-even. The argument only holds for elementary particles. When you have molecules you can have an EDM and move to different states. Measurement of the EDM for the electron is a measurement of \mathcal{CP} . What about the neutron? People say that the neutron EDM is the best bound on SUSY CP violation. Certainly the neutron is not elementary, but there is no excited neutron state. Thus the full statement is that the EDM of any particle that doesn’t have an excited state is a signal of \mathcal{CP} . (Actually, this boils down to having a degenerate ground state.)

So far CPV has only been unambiguously observed in K_L , B^0 , and B^\pm .

14.1 CPV in mesons

There are basically two processes: decay and oscillation. To be sensitive to the CP phases we need two amplitudes to interfere. There are three options

1. CPV in decay. This is an interference in the decay amplitude. For example the B^\pm doesn’t oscillate, so it has to be of this type.
2. CPV in mixing. This is an interference between M_{12} and Γ_{12} , the two ways to mix form B into a \bar{B} .
3. CPV in interference between mixing and decay.

Let's start with CPV in decay. The observable that we will define is

$$a_{\text{CP}}[f] \equiv \frac{\Gamma(\bar{B} \rightarrow \bar{f}') - \Gamma(B \rightarrow f)}{\Gamma(\bar{B} \rightarrow \bar{f}') + \Gamma(B \rightarrow f)} \quad (14.8)$$

$$= \frac{|\bar{A}/A|^2 - 1}{|\bar{A}/A|^2 + 1}. \quad (14.9)$$

We have used

$$A \equiv A(B \rightarrow f) \quad (14.10)$$

$$\bar{A} \equiv A(\bar{B} \rightarrow \bar{f}). \quad (14.11)$$

We can always parameterize the amplitude as

$$A(B \rightarrow f) = \tilde{A} (1 + r e^{i(\phi+\delta)}) \quad (14.12)$$

$$A(\bar{B} \rightarrow \bar{f}) = \tilde{A} (1 + r e^{i(-\phi+\delta)}). \quad (14.13)$$


Plugging this in we get

$$a_{\text{CP}} = r \sin \phi \sin \delta. \quad (14.14)$$


Once again we get two amplitudes so that $r \neq 0$ and we also need ϕ and δ not equal to zero.

Now some examples. Up to now we haven't explained how these values are related to our fundamental parameters. ϕ will be related to the CKM. We can use symmetries to get information about the hadronic matrix element, which shows up in r . In δ , however, the whole effect has to do with symmetry breaking so the leading effect is zero. This is why CPV is hard to measure.

Let us consider is direct CPV in $B \rightarrow K\pi$. Can we draw an amplitude for this? Yes.



We also have the penguin diagram



We can estimate these.

$$\mathcal{M}_{\text{penguin}} \sim \alpha_W \alpha_s \frac{\lambda^2}{16\pi^2} \quad (14.15)$$

$$\mathcal{M}_{\text{tree}} \sim \alpha_w \lambda^4. \quad (14.16)$$

The tree dominates by a factor of five by our naïve analysis, however experimentally it turns out that it's the other way around, the penguin dominates. It is very important that these turn out to be roughly the same size. Note that the isospin is different. The tree level is a combination of isospin 0 and isospin 1, while the penguin is isospin zero. See homework.

Now that we have the two amplitudes, we must identify the phases. We assume that the top dominates the $b \rightarrow s$ penguin (though the charm also contributes). [This was a bit difficult to type up.] Which phase is in that penguin then? None! You could have seen this—the only phase shows up in the 1-3 elements. What about the tree diagram? Look back at our unitarity triangle. The $b\bar{u}$ vertex comes with the γ angle. ... what is the strong phase? We don't know. r is the absolute value of the ratio of these two amplitudes.

There are four different $B \rightarrow K\pi$ decays. We can use isospin and SU(3) to relate these. The B^+ was the first system where we saw CPV without flavor mixing.

Acknowledgements

Y.G. would like to thank the students of Cornell University's Physics 7661 course for their questions which reflected their enthusiasm for the subject. P.T. would also like to thank them for their questions which gave him enough time for him to T_EX these notes in real-time during the lectures. P.T. would additionally like to thank all who read these lectures and offered corrections and suggestions, including Panagiotis Athanasopoulos, Hardik Panjwani, Yuhsin Tsai, and Yariv Yanay. Special thanks are in order for Andreas Kronfeld for very helpful explanations about lattice QCD. This work is supported in part by the NSF grant number PHY-0355005. P.T. is additionally supported by a National Science Foundation graduate research fellowship and a Paul & Daisy Soros Fellowship For New Americans. The contents of this article do not necessarily represent the views of either institution.

A Notation and Conventions

4D Minkowski indices are written with lower-case Greek letters from the middle of the alphabet, μ, ν, \dots . 5D indices are written in capital Roman letters from the middle of the alphabet, M, N, \dots . Tangent space indices are written in lower-case Roman letters from the beginning of the alphabet, a, b, \dots . Flavor indices are written in lower-case Roman letters near the beginning of the alphabet, i, j, \dots . We will use the notation that a matrix with a hat, e.g. \hat{m} , is diagonal.

We use the particle physics ('West Coast,' mostly-minus) metric for Minkowski space, $(+, -, -, -)$. Dirac spinors Ψ are related to left- and right-chiral Weyl spinors $(\chi, \bar{\psi})$ respectively via

$$\Psi = \begin{pmatrix} \chi \\ \bar{\psi} \end{pmatrix}. \quad (A.1)$$

Our convention for σ^0 and the three Pauli matrices $\vec{\sigma}$ is

$$\sigma^0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \sigma^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (\text{A.2})$$

with the flat-space γ matrices given by

$$\gamma^\mu = \begin{pmatrix} 0 & \sigma^\mu \\ \bar{\sigma}^\mu & 0 \end{pmatrix}, \quad (\text{A.3})$$

where $\bar{\sigma}^\mu = (\sigma^0, -\vec{\sigma})$. We normalize γ^5 so that the left-chiral projection operator is $P_L = \frac{1}{2}(1 - \gamma^5)$, i.e. $(\gamma^5)^2 = 1$ rather than the normalization that satisfies the 5D Clifford algebra.

We will write the CKM matrix as V and the Jarlskog invariant as J .

B Facts that you should memorize

During the course there were a few facts that were deemed useful enough to commit to memory.

- The Fermi constant in terms of the $SU(2)_L$ coupling,

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8M_W^2}. \quad (\text{B.1})$$


The $\sqrt{2}$ in the denominator on the left-hand side is usually absorbed by f_π to become F_π . The factor of 8 on the right-hand side comes from the W coupling constant being $g/\sqrt{2}$ and the factors of 1/2 that come from, for example, only picking up the axial part of a chiral interaction (e.g. only the second term of $\frac{1}{2}\gamma^\mu(1 - \gamma^5)$).

- The mass of the B meson is 5280 MeV, which is precisely the number of feet in a mile. See also the *Spiked Math* comic [1].
- $1 \text{ GeV} = 1.5 \times 10^{24} \text{ sec}^{-1}$
- $1 \text{ GeV} = 5 \times 10^{13} \text{ cm}^{-1}$
- The **Wolfenstein parameterization** of the CKM matrix...

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4), \quad (\text{B.2})$$

where one should note that $\lambda \approx 0.2$ is a good expansion parameter and all other parameters are $\mathcal{O}(1)$. Note that the top-left 2×2 has the structure of the Cabbibo mixing matrix and that CP violation *only* appears with a λ^3 in the 3-1 mixing terms.

- The unitarity triangle:



- The pion decay constant is $f_\pi \sim 131$ MeV.
- The τ mass is 1777 MeV. According to Yuval, “this is just one, and then 777, like the casino.”

C Lie groups, Lie algebras, and representation theory

This appendix is based on notes by Y. Grossman and Y. Nir.

Symmetries play a crucial role in model building. You are already familiar with some symmetries and their consequences. For example, the Poincaré group implies the conservation of energy, momentum, and angular momentum. The mathematical structures that describe symmetries are called **Lie groups**. In this appendix we will briefly review how these objects allow us to relate symmetries to physical interactions. Three particularly useful texts on this subject are

- *Lie Algebras in Particle Physics*, Georgi [39]
- *Semi-simple Lie Algebras and their Representations*, Cahn [17]
- “Group theory for unified model building”, Slansky [87]

C.1 Groups and representations

We begin with a series of definitions.

Definition C.1. A **group** G is a set $\{x_i\}$ (finite or infinite) with a multiplication law \cdot subject to the following four requirements:

- **Closure:**

$$x_i \cdot x_j \in G \quad \forall x_i \tag{C.1}$$

- **Associativity:**

$$x_i \cdot (x_j \cdot x_k) = (x_i \cdot x_j) \cdot x_k. \tag{C.2}$$

- There exists an **identity** element, I (or e) such that

$$I \cdot x_i = x_i \cdot I = x_i \quad \forall x_i. \quad (\text{C.3})$$

- For each element $x_i \in G$ there exists an **inverse** element x_i^{-1} such that

$$x_i \cdot x_i^{-1} = x_i^{-1} \cdot x_i = I. \quad (\text{C.4})$$

Definition C.2. A group is **Abelian** if all its elements commute:

$$x_i \cdot x_j = x_j \cdot x_i \quad \forall x_i, x_j. \quad (\text{C.5})$$

A group is **non-Abelian** if it is... not Abelian; in other words, at least one pair of elements does not commute.

Let us give a few examples,

- \mathbb{Z}_2 , also known as parity, is a group with two elements, I and P , such that I is the identity and $P^{-1} = P$. This completely specifies the multiplication table. This group is finite and Abelian.
- \mathbb{Z}_N , for integer N , is a generalization of \mathbb{Z}_2 . It contains N elements labeled from zero to $N - 1$. The multiplication law is the same as addition modulo N : $x_i x_j = x_{(i+j) \bmod N}$. The identity element is x_0 , and the inverse element is given by $x_i^{-1} = x_{N-i}$. This group is also finite and Abelian.
- Positive real numbers with standard multiplication is an infinite Abelian group.
- S_3 , permutation group of 3 objects. It contains 6 elements and is non-Abelian. Work for yourself the 6 elements and the multiplication table.

Definition C.3. A **representation** is a realization of the group multiplication law on a linear vector space so that group elements are identified with matrices.

Definition C.4. Two representations are **equivalent** if they are related by a similarity transformation.

Definition C.5. A representation is **reducible** if it is equivalent to a representation that is block diagonal.

Definition C.6. An **irreducible** representation (‘irrep’) is a representation that is not reducible.

Definition C.7. An irrep that contains matrices of size $n \times n$ is said to be of **dimension** n .

Fact C.8. Any reducible representation can be written as a direct sum of irreps, e.g. $D = D_1 + D_2$.

Fact C.9. The dimension of all irreps of an Abelian group is one.

Fact C.10. Any finite group has a finite number of irreps R_i . If N is the number of elements in the group, the irreps satisfy

$$\sum_{R_i} [\dim(R_i)]^2 = N. \quad (\text{C.6})$$

Definition C.11. For any group there exist a **trivial** representation such that all the matrices are just the number 1. This representation is also called the **singlet** representation and, as we see later, is of particular importance for us.

Let us give some examples for the above statements.

- \mathbb{Z}_2 : has a trivial irrep with $I = 1, P = 1$. The other irrep is $I = 1, P = -1$. These two irreps clearly satisfy (C.6).
- \mathbb{Z}_N : An example of a non-trivial irrep is $x_k = \exp(i2\pi k/N)$.
- S_3 : We will explore this group further in the homework.

The groups that we are particularly interested in are the *transformation groups of physical systems*. Such transformations are associated with **unitary operators** in the Hilbert space. We will often describe the elements of the group by the way that they transform physical states. When we refer to representations of the group, we mean either the appropriate set of unitary operators, or, equivalently, by the matrices that act on the states of the Hilbert space.

C.2 Lie groups

While finite groups are very important, particle physicists are particularly interested in continuous groups which are infinite-dimensional with cardinality \aleph_1 . These groups are called **Lie groups**.

Definition C.12. A **Lie group** is an infinite group whose elements are labeled by a finite set of N continuous real parameters α_ℓ , and whose multiplication law depends smoothly on the α_ℓ 's. The number N is called the dimension of the group.

Fact C.13. An Abelian Lie group has $N = 1$. A non-Abelian Lie group has $N > 1$.

Our first example of a Lie group is $U(1)$. It represents addition of real numbers modulo 2π , that is, rotation on a circle. Such a group has an infinite number of elements that are labeled by a single continuous parameter α . We can write the group elements as $M = \exp(i\alpha)$. We can also represent it by $M = \exp(2i\alpha)$ or, more generally, as $M = \exp(iX\alpha)$ with X real. Each X generates an irrep of the group.

We are mainly interested in **compact** Lie groups. We will not give a formal definition, but we can use the $U(1)$ example to give an intuitive explanation of compactness. A group whose multiplication rule is modular addition is compact, while non-modular addition would be non-compact. In the first case, if you repeat the same addition a number of times, you may return to your starting point. In the latter this would never happen. In other words, in a compact Lie group, the parameters have a finite range, while in a non-compact group, their range is infinite. Do not confuse that with the number of elements, which is infinite in either case. Another familiar example is the Lorentz group: Rotations are compact while boosts are not.

Fact C.14. *The elements of any compact Lie group can be written as*

$$M_i = \exp(i\alpha_\ell X_\ell) \tag{C.7}$$

such that X_ℓ are Hermitian matrices that are called **generators**. (We use the standard summation convention: $\alpha_\ell X_\ell \equiv \sum_\ell \alpha_\ell X_\ell$.)

Let us perform some algebra before moving on to our next definition. Consider two elements of a Lie group, A and B , such that A is the exponentiation of a the generator X_a (i.e. only $\alpha_a \neq 0$) and is the exponentiation of a the generator X_b (i.e. only $\alpha_b \neq 0$). One may of course parameterize the generators such that this is true for any A and B . Let us further set $\alpha_a = \alpha_b = \lambda$ so that

$$A \equiv \exp(i\lambda X_a), \quad B \equiv \exp(i\lambda X_b). \tag{C.8}$$

Since A and B are in the group, each of them has an inverse. Thus,

$$C = BAB^{-1}A^{-1} \equiv \exp(i\beta_c X_c) \tag{C.9}$$

is also in the group. Let us take λ to be a small parameter and expand around the identity. Clearly in this limit all of the the β_c are also small. Keeping the leading order terms, we get

$$C = \exp(i\beta_c X_c) \approx I + i\beta_c X_c, \quad C = BAB^{-1}A^{-1} \approx I + \lambda^2[X_a, X_b]. \tag{C.10}$$

In the $\lambda \rightarrow 0$ limit, we have

$$[X_a, X_b] = i\frac{\beta_c}{\lambda^2} X_c. \tag{C.11}$$

Clearly, the combinations

$$f_{abc} \equiv \lambda^{-2}\beta_c \tag{C.12}$$

are independent of λ . Furthermore, while λ and β_c are infinitesimal, the f_{abc} -constants do not diverge. This brings us to a new set of definitions.

Definition C.15. The commutation relations of the generators of a Lie group [see (C.11)]

$$[X_a, X_b] = if_{abc}X_c \tag{C.13}$$

are the called the **Lie algebra**. The constants f_{abc} are called the **structure constants** of the algebra.

Note the following points regarding the Lie Algebra:

- The algebra defines the local properties of the group but not its global properties. Usually, this is all we care about⁴⁰.

⁴⁰A notable exception in high-energy physics is the Poincaré group whose nontrivial global structure requires that particles of a given spin should be representations of the *universal cover* of the Poincaré group rather than the Poincaré group itself. See, for example, [33] for a pedagogical exposition.

- The Lie algebra is closed under the commutation operator.
- Just like Lie groups, one can define representations of the Lie algebra, i.e. matrix representations of X_ℓ . In particular, each representation has its own dimension. (Do not confuse the dimension of the representation with the dimension of the group!)
- The generators satisfy the Jacobi identity

$$[X_a, [X_b, X_c]] + [X_b, [X_c, X_a]] + [X_c, [X_a, X_b]] = 0. \quad (\text{C.14})$$

- For each Lie algebra there is the trivial (**singlet**) representation which is $X_\ell = 0$ for all ℓ . The trivial representation of the algebra generates the trivial representation of the group.
- Since an Abelian Lie group has only one generator, its algebra is always trivial. Thus, the algebra of $U(1)$ is the only Abelian Lie algebra.
- Non-Abelian Lie groups have non-trivial algebras.

The example of $SU(2)$ algebra is well-known from introductory quantum mechanics:

$$[X_a, X_b] = i\varepsilon_{abc}X_c. \quad (\text{C.15})$$

Usually, in QM, X is called L or S or J . The $SU(2)$ group represents non-trivial rotations in a two-dimensional complex space. Its algebra is the same as the algebra of the $SO(3)$ group, which represents rotations in the three-dimensional real space.

Let us explain what we mean when we say that “the group represents rotations in a space.” This is clear in the context of quantum mechanics. Consider a finite Hilbert space of, say, a particle with spin S . The matrices that rotate the direction of the spin are written in terms of exponent of the S_i operators. For a spin one-half particle, the S_i operators are written in terms of the Pauli matrices. For particles with spin different from $1/2$, the S_i operators will be written in terms of different matrices. We learn that the group represents rotations in some space, while the various representations correspond to different objects that can “live” in that space.

There are three important irreps that have special names. The first one is the trivial—or singlet—representation that we already mentioned. Its importance stems from the fact that it corresponds to something that is invariant under rotations.

The second important irrep is the **fundamental** representation. This is the smallest non-trivial irrep. For $SU(2)$, this is the spinor representation. An important property of the fundamental representation is that it can be used to get all other representations. We will return to this point later. Here we just remind you that this statement is well familiar from QM. One can get spin-1 by combining two spin-1/2, and you can get spin-3/2 by combining three spin-1/2. Every Lie group has a fundamental irrep.

The third important irrep is the **adjoint** representation. It is made out of the structure constants themselves. Think of a matrix representation of the generators. Each entry, T_{ij}^c is labelled by three indices. One is the index of the generator itself c that runs from 1 to N , such that N is the dimension of the algebra. The other two indices, i and j , are the matrix indices that run from 1 to the dimension of the representation.

Caution. Do not become confused by the different dimensions we've mentioned. The dimension of the *algebra*, which is finite, is different from the dimension of the *group*, which is infinite for Lie groups. These are both intrinsic properties of the group, for example $SU(2)$ is an infinite-dimensional group which has a three-dimensional algebra whose elements are the Pauli matrices. This algebra has different representations, each of which has a different dimension. In particle physics we are always interested in finite-dimensional representations. The only exception is the Poincaré group, for which we use a trick to work with finite-dimensional representations; see [33].

One can show that each Lie group has one representation where the dimension of the representation is the same as the dimension of the group; this is the adjoint representation. This representation is obtained by defining

$$(X_c)_{ab} \equiv -if_{abc}. \tag{C.16}$$

In other words, the structure constants themselves satisfy the algebra of their own group. In $SU(2)$, the adjoint representation corresponds to spin-1. It is easy to see that the $SU(2)$ structure constants ε_{ijk} are just the set of the three 3×3 representations of spin-1.

C.3 More formal developments

Definition C.16. A **subalgebra** M is a set of generators that are closed under commutation.

Definition C.17. Consider an algebra L with a subalgebra M . M is an **ideal** if for any $x \in M$ and $y \in L$ $[x, y] \in M$. (For a subalgebra that is not ideal we still have $[x, y] \in L$.)

Definition C.18. A **simple** Lie algebra is an algebra without a non-trivial ideal. (Any algebra has a trivial ideal, the algebra itself.)

Definition C.19. A **semi-simple** Lie algebra is an algebra without a $U(1)$ ideal.

Any Lie algebra can be written as a direct product of simple Lie algebras. Thus we can think about each of the simple algebras separately. You are already familiar with this! For example, consider the hydrogen atom. We can think about the Hilbert space as a direct product of the spin of the electron and that of the spin of the proton.

A useful example is that of the $U(2)$ group, which is not semi-simple:

$$U(2) = SU(2) \times U(1). \tag{C.17}$$

A $U(2)$ transformation corresponds to a rotation in two-dimensional complex space. Think, for example, about the rotation of a spinor. It can be separated into two: The trivial rotation is just a $U(1)$ transformation, that is, a phase multiplication of the spinor. The non-trivial rotation is the $SU(2)$ transformation, that is, an internal rotation between the two spin components.

Definition C.20. The **Cartan subalgebra** is the largest subset of generators whose matrix representations can all be diagonalized at once.

Obviously, these generators all commute with each other and thus constitute a subalgebra.

Definition C.21. The number of generators in the Cartan subalgebra is called the **rank** of the algebra.

Let us consider a few examples. Since the $U(1)$ algebra only has a single generator so it has rank one. $SU(2)$ is also rank one. You can make one of its three generators, say S_z , diagonal, but not two of them simultaneously. $SU(3)$ is rank two. Much of this course will depend on the structure of $SU(3)$ so we will explore it in much more detail below.

Our next step is to introduce **roots** and **weights**. We do this by first presenting an example.

Example C.22. Consider the $SU(2)$ algebra. It has three generators. We usually choose S_3 to be in the Cartan subalgebra, and we can combine the two other generators, S_1 and S_2 , into raising and lowering operators, $S^\pm = S_1 \pm iS_2$. Any representation can be defined by the eigenvalues under the operation of the generators in the Cartan subalgebra, in this case S_3 . For example, for the spin-1/2 representation, the eigenvalues are $-1/2$ and $+1/2$; For the spin-1 representation, the eigenvalues are -1 , 0 , and $+1$. Under the operation of the raising (S^+) and lowering (S^-) generators, we “move” from one eigenstate of S_3 to another. For example, for a spin-1 representation, we have $S^+|-1\rangle \propto |0\rangle$.

Let us now consider a general Lie group of rank n . Any representation is characterized by the possible eigenvalues of its eigenstates under the operation of the Cartan subalgebra: $|e_1, e_2, \dots, e_n\rangle$. We can assemble all the operators that are not in the Cartan subalgebra into “raising” and “lowering” operators. That is, when they act on an eigenstate they either move it to another eigenstate or annihilate it.

Definition C.23. The **weight vectors** or **weights** of a representation are the possible eigenvalues of the generators in the Cartan subalgebra.

Definition C.24. The **roots** of the algebra are the various ways in which the generators move a state between the possible weights.

Fact C.25. *The weights completely describe the representation.*

Fact C.26. *The roots completely describe the Lie algebra.*

Note that both roots and weights live in an n -dimensional vector space, where n is the rank of the group. The number of roots is the dimension of the group. The number of weights is the dimension of the irrep.

Example C.27. Let us return to our $SU(2)$ example. The vector space of roots and weights is one-dimensional. The three roots are $0, \pm 1$. The trivial representation has only one weight, zero; The fundamental has two, $\pm 1/2$; The adjoint has three, $0, \pm 1$ (the weights of the adjoint representations are just the roots); and so on.

C.4 $SU(3)$

In this section we discuss the $SU(3)$ group. It is more complicated than $SU(2)$ but it allows us to demonstrate a few aspects of Lie groups that cannot be demonstrated with $SU(2)$. Of course, it is also particularly relevant to particle physics.

$SU(3)$ is a generalization of $SU(2)$. It may be useful to think about it as rotations in three-dimensional complex space. Similar to $SU(2)$, the full symmetry of the rotations is called $U(3)$, and it can be written as a direct product of simple groups, $U(3) = SU(3) \times U(1)$. The $SU(3)$ algebra has eight generators. (There are nine independent Hermitian 3×3 matrices. They can be separated to a unit matrix, which corresponds to the $U(1)$ part, and eight traceless matrices, which correspond to the $SU(3)$ part.)

Similar to the use of the Pauli matrices for the fundamental representation of $SU(2)$, the fundamental representation of $SU(3)$ is usually written in terms of the **Gell-Mann matrices**,

$$X_a = \lambda_a/2, \tag{C.18}$$

with

$$\begin{aligned} \lambda_1 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, & \lambda_2 &= \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \\ \lambda_3 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, & \lambda_4 &= \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \\ \lambda_5 &= \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, & \lambda_6 &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \\ \lambda_7 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, & \lambda_8 &= \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}. \end{aligned} \tag{C.19}$$

We would like to emphasize the following points:

1. The Gell-Mann matrices are traceless, as we expect for the generators of a unitary group. One can check that the condition of unitarity of group elements imposes tracelessness on the elements of the algebra.
2. The Lie algebra of $SU(3)$ contains three $SU(2)$ subalgebras. One of them is manifest and it is given by λ_1 , λ_2 and λ_3 . Can you find the other two?
3. It is manifest that $SU(3)$ is of rank two: λ_3 and λ_8 are in the Cartan subalgebra.

Having explicit expressions of fundamental representation at our disposal, we can now draw a **weight diagram**—a graphical way to visualize the roots and weights of a Lie algebra. Let us first sketch how this works for the fundamental (spinor) representation of $SU(2)$. We have two basis vectors (spin-up and spin-down) and apply S_z on them and obtain the two weights, $+1/2$ and $-1/2$.

For $SU(3)$ we follow the same steps. We take the three vectors,

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad (\text{C.20})$$

and on each of them apply the two generators in the Cartan subalgebra, X_3 and X_8 . We find the three weights

$$\left(+\frac{1}{2}, +\frac{1}{2\sqrt{3}} \right), \quad \left(-\frac{1}{2}, +\frac{1}{2\sqrt{3}} \right), \quad \left(0, -\frac{1}{\sqrt{3}} \right). \quad (\text{C.21})$$

We can plot this in a weight diagram in the $X_3 - X_8$ plane. Please do it.

Once we have the weights we can obtain the roots. They are just the combination of generators that move us between the weights. Clearly, the two roots that are in the Cartan are at the origin. The other six are those that move us between the three weights. It is easy to find that they are

$$\left(\pm\frac{1}{2}, \pm\frac{\sqrt{3}}{2} \right), \quad (\pm 1, 0). \quad (\text{C.22})$$

Again, it is a good idea to plot it. This root diagram is also the weight diagram of the adjoint representation.

D Homework solutions

1.1 According to *Wikipedia*, the Martian sky is pinkish-red and becomes blue around the horizon during sunset [2]. Rayleigh scattering, which causes the blue color of the Earth sky, is a small effect on Mars. Instead, it is believed that dust containing magnetite causes the reddish hue of the Martian sky.

2.1 Let us review our counting for real parameters and phases. Each Yukawa matrix has N^2 elements which can each be written as $y_{ij} = r_{ij}e^{i\theta_{ij}}$. Thus each Yukawa contains N real parameters and N phases. We would now like to identify the number of broken generators. An element of $U(N)$ is an $N \times N$ complex matrix satisfying

$$UU^\dagger = \mathbb{1}. \quad (\text{D.1})$$

This gives one equation of constraint for each complex component for a total of N^2 parameters. How many of these parameters are real? This is precisely the number of parameters of an element of $O(N)$, namely $\frac{1}{2}N(N-1)$. The remaining $\frac{1}{2}N(N+1)$ parameters are phases.

In the quark sector of our N generation Standard Model, we have $2N^2$ real parameters and $2N^2$ phases coming from the y^u and y^d matrices. How many of these are physical? We subtract the number of broken generators. For generic Yukawa matrices with no underlying structure, the $U(N)^3$ flavor symmetry is broken down to $U(1)_B$. Thus we have $3 \times \frac{1}{2}N(N-1)$ broken generators

associated with real parameters and $3 \times \frac{1}{2}N(N+1) - 1$ broken generators associated with phases. This gives a total of

$$2N^2 - \frac{3}{2}N(N-1) = \frac{1}{2}N(N+3) \quad \text{physical real parameters} \quad (\text{D.2})$$

$$2N^2 - \frac{3}{2}N(N+1) + 1 = \frac{1}{2}(N-2)(N-1) \quad \text{physical phases.} \quad (\text{D.3})$$

Of the real parameters, N of them are Dirac masses and so $\frac{1}{2}N(N+1)$ must be mixing angles. Let us remark that for the full N -generation Standard Model we must also add N real parameters for the lepton masses, the 3 real parameters for the gauge couplings, and the two real parameters governing the Higgs sector.

2.2 The problem already defined the three main ingredients of a model. It's up to us to now write the most general Lagrangian and identify the number of physical parameters. We will only consider the quark sector since all other sectors are the same as the Standard Model. There are five matter fields in the quark sector $(Q_L, u_R, d_R, s_L, s_R)$. The only difference from the Standard Model is that the s_L is an $SU(2)_L$ singlet so that it must have $Y = -1/3$ to obtain the correct electric charge. The kinetic terms have a $U(2) \times U(1)^3$ symmetry where the $U(2)$ corresponds to mixing the d_R and s_R fields. Note that the s_L cannot mix with these fields even though it seems to have the same quantum numbers as the s_R ; the s_L and s_R are different Lorentz representations. In other words, s_L is a left-chiral spinor χ_α while s_R is a right-chiral spinor $\bar{\psi}^{\dot{\alpha}}$. As one can see from the indices, these transform differently under rotations. Using the counting introduced in the previous problem for the number of real parameters and phases in unitary matrices, we see that the kinetic terms have a symmetry group with 1 real parameter generator and 6 phase generators.

Let us now identify how many total parameters are in the Lagrangian and how many of the above symmetries are broken. The most general quark Yukawa sector is

$$\mathcal{L}_{\text{Yuk.,}Q} = y_u \bar{Q}_L \phi u_R + y_d \bar{Q}_L \tilde{\phi} d_R + y_s \bar{Q}_L \tilde{\phi} s_R + \text{h.c.}, \quad (\text{D.4})$$

where the y_i are complex numbers so that we have 3 real parameters and 3 phases. In the Standard Model the Yukawa sector was the only source of additional quark interactions. In our exotic model, however, the fact that the s_L is a singlet allows us to write additional bare mass terms,

$$\mathcal{L}_{\text{mass}} = m_s \bar{s}_L s_R + m_d \bar{s}_L d_R + \text{h.c.} \quad (\text{D.5})$$

Each m_a is a complex number so that this gives 2 additional real parameters and 2 additional phases. We now have a total of 5 real parameters and 5 phases. What symmetries are preserved by these Lagrangian terms? We are left with an analog of $U(1)_B$ where the left-handed and right-handed fields all transform with the same phase. This means that of the symmetries in the kinetic sector, the one real parameter and 5 of the 6 phases are broken. The number of physical parameters is equal to the number of total parameters minus broken generators, so this gives us a total of four real parameters and no phases.

Does this make sense? Three real parameters are associated to the Dirac masses of the particles. We have one left over real parameter, but we can see that when the Higgs gets its electroweak

symmetry-breaking vev there is a mass term mixing the s and d quarks so that the mass term after EWSB looks like

$$\mathcal{L} \supset (\bar{d}_L \quad \bar{s}_L) \begin{pmatrix} y_d \frac{v}{\sqrt{2}} & y_s \frac{v}{\sqrt{2}} \\ m_d & m_s \end{pmatrix} \begin{pmatrix} d_R \\ s_R \end{pmatrix} + \text{h.c.} \quad (\text{D.6})$$

so that we can see that the left over real parameter gives the $d - s$ mixing angle.

2.3 A careful counting by Dimopoulos and Sutter gives 110 parameters: 30 masses, 29 mixing angles, and 41 phases [27]. You are referred to their paper for a very nice analysis. See also a nice presentation by Haber [48] for some variants of the MSSM.

2.4 The conventional explanation for why the $SU(2)_L$ Θ angle is not physical comes from the use of the anomalous $U(1)_{B+L}$ symmetry to rotate it away. (One cannot do this in QCD with $U(1)_A$ since this would give the quark masses an imaginary component.) A cute way to check this is to count the number of physical parameters in the $SU(2)_L$ sector at a *quantum* level and check that they are all accounted for *without* having to introduce a Θ_L term. We can cast this argument in a different way in light of our parameter counting. We saw that the $U(1)_A$ anomaly of QCD led to the appearance of another physical parameter at the quantum level, Θ_{QCD} . If we look at the $SU(2)_L$ sector of the Standard Model, we might also expect an additional parameter to appear, Θ_{weak} , because of the $U(1)_{B+L}$ anomaly. However, the key point is that $U(1)_{B+L}$ isn't broken by the Yukawas, and so it was never counted as a 'broken symmetry' in our tree-level analysis. Thus the $U(1)_{B+L}$ anomaly doesn't lead to one less broken symmetry and so does not introduce an additional physical parameter. In other words, Θ_{Weak} is not physical. By the way, a useful way to remember which global symmetries are broken anomalies is to look at the instanton that induces the appropriate Θ term. In this case the sphaleron violates $B + L$ while preserving $B - L$.

3.1 A quick-and-dirty way to get this is to use the Heisenberg relation $\Delta E \Delta t \sim \hbar$.

3.2 These are trivial to look up in the PDG. The paper version is slightly more satisfying to thumb through, while the online version is kinder to the environment.

1. The easiest place to look is the Meson Summary Table. The D^+ is a $c\bar{d}$ bound state with mass $m_{D^+} \approx 1870$ MeV.
2. The easiest place to look is the Baryon Summary Table. The Λ is a uds bound state with spin $1/2$.
3. The easiest place to look is the Lepton Summary Table. (Do you see a theme?) The τ will decay to $\mu\nu\bar{\nu}$ 17% of the time.
4. The B^+ meson has a lifetime of $\tau_{B^\pm} = 1.6 \times 10^{-12}$ s. This translates to a width of

$$\Gamma_{B^+} = (1.6 \times 10^{-12} \text{ s})^{-1} \times \frac{1 \text{ GeV}}{1.5 \times 10^{24} \text{ s}^{-1}} = 4.4 \times 10^{-13} \text{ GeV} = 4.4 \times 10^{-4} \text{ eV}. \quad (\text{D.7})$$

5. The PDG tells us that $c\tau = 491 \mu\text{m}$. With a boost factor $\gamma = 4$ this gives us $\gamma c\tau = 1.9$ mm.

3.3 A convenient basis is $\{|u\bar{u}\rangle, |u\bar{d}\rangle, |d\bar{u}\rangle, |d\bar{d}\rangle\}$. We can convert to the isospin basis using

$$|0, 0\rangle = \frac{1}{\sqrt{2}} (|u\bar{u}\rangle - |d\bar{d}\rangle) \quad (\text{D.8})$$

$$|1, -1\rangle = |d\bar{u}\rangle \quad (\text{D.9})$$

$$|1, 0\rangle = \frac{1}{\sqrt{2}} (|u\bar{u}\rangle + |d\bar{d}\rangle) \quad (\text{D.10})$$

$$|1, 1\rangle = |u\bar{d}\rangle \quad (\text{D.11})$$

where we've written states according to $|I, I_3\rangle$. Note that u and \bar{d} are analogous to $|\uparrow\rangle$ while \bar{u} and d are analogous to $|\downarrow\rangle$ in the usual $SU(2)$ spin notation. (To see why just remember that the kinetic term must be isospin invariant and that the $SU(2)$ metric is the antisymmetric tensor.) Let us say that the singlet gets a mass m_1 and the triplet states get mass m_3 . The mass matrix thus takes the form

$$\begin{pmatrix} 2m_u^2 + \frac{1}{2}m_3^2 + \frac{1}{2}m_1^2 & & & \frac{1}{2}m_1^2 - \frac{1}{2}m_3^2 \\ & m_u^2 + m_d^2 + \frac{1}{2}m_3^2 & & \\ & & m_u^2 + m_d^2 + \frac{1}{2}m_3^2 & \\ \frac{1}{2}m_1^2 - \frac{1}{2}m_3^2 & & & 2m_d^2 + \frac{1}{2}m_3^2 + \frac{1}{2}m_1^2 \end{pmatrix}. \quad (\text{D.12})$$

From here one can happily plug into *Mathematica* to solve for the exact eigenstates and eigenvalues. As a sanity-check note that in the limit $m_{u,d} \rightarrow 0$ the $u\bar{u}$ and $d\bar{d}$ states mix with a 45° angle giving the π^0 and η . The $u\bar{d}$ and $d\bar{u}$ states do not mix because both bases conserve I_3 . Note further that electromagnetism also preserves the I_3 quantum number.

3.4 Writing the universal mass contribution as m_Λ and using the quark content predicted by $SU(3)$, e.g. by using (3.11) and (3.12), we find

$$m_\pi^2 = m_\Lambda^2 + 2\mu m_u \quad (\text{D.13})$$

$$m_K^2 = m_\Lambda^2 + \mu(m_u + m_s) \quad (\text{D.14})$$

$$m_\eta^2 = m_\Lambda^2 + \frac{2\mu}{3}(m_u + 2m_s). \quad (\text{D.15})$$

We can 'solve' these equations by eliminating the unknown quantities m_Λ , m_u , and m_s to obtain the **Gell-Mann–Okubo mass formula** for pseudoscalar mesons,

$$4m_K^2 = m_\pi^2 + 3m_\eta^2. \quad (\text{D.16})$$

Note that the meson masses in this expression are summed in quadrature. For general $SU(3)$ octets (vector mesons) the meson masses would be summed linearly. We will explain this subtle difference in Section 3.8.

3.5 This is actually a bit of a trick question because it's hard to find the quark content from the summary table. Textbooks list the quark content as being approximately

$$\omega = \frac{1}{\sqrt{2}} (u\bar{u} + d\bar{d}) \quad (\text{D.17})$$

$$\phi = s\bar{s}. \quad (\text{D.18})$$

3.6 Representation theory for $SU(3)$ tells us that the $I_3 = 0$ octet and singlet states have the quark content

$$\psi_8 = \frac{1}{\sqrt{6}} (u\bar{u} + d\bar{d} - 2s\bar{s}) \quad (\text{D.19})$$

$$\psi_1 = \frac{1}{\sqrt{3}} (u\bar{u} + d\bar{d} + s\bar{s}). \quad (\text{D.20})$$

As we mentioned before we will be working with differences of masses, so we can ignore all universal contributions to meson masses coming from angular momentum and the QCD potential. Further, we will assume that $SU(2)$ isospin is a good symmetry such that $m_d = m_u$. Then the mass of the ψ_8 comes from averaging over the probabilities that ψ_8 will contain a given quark pair. In other words, one should look at

$$\langle \psi_8 | \psi_8 \rangle = \frac{1}{6} (\langle u\bar{u} | u\bar{u} \rangle + \langle d\bar{d} | d\bar{d} \rangle + 4\langle s\bar{s} | s\bar{s} \rangle). \quad (\text{D.21})$$

to determine that

$$m_8 = \frac{1}{6} (2 \times 2m_u + 2 \times 4m_s) = \frac{2}{3}(m_u + 2m_s), \quad (\text{D.22})$$

where we remember an overall factor of two because the state $u\bar{u}$ has valence quark mass contribution $\Delta m = 2m_u$. Note that we look at the square of the state $|psi\rangle$ because we want the *probability* that the meson contains a given quark content; this has *nothing* to do with the idea that we should be summing the masses linearly or quadratically. It is “more obvious than obvious⁴¹” that we should sum the masses linearly. (The pseudoscalars are special because they are pseudo-Goldstone bosons that need to be treated in chiral perturbation theory where the *squared* meson masses are linear in the symmetry-breaking parameters, the quark masses.) Using the same analysis we may write out the valence quark mass contribution for the vector kaon and ρ ,

$$m_{K^*} = m_u + m_s \quad (\text{D.23})$$

$$m_\rho = 2m_u. \quad (\text{D.24})$$

This gives us a basis to re-express quark masses in terms of the non-mixing octet meson masses. (These mesons don’t mix because the quark masses break $SU(3) \rightarrow U(1)^3$ so that strangeness is conserved.) Thus we readily obtain

$$m_8 = \frac{1}{3}(4m_{K^*} - m_\rho) \quad (\text{D.25})$$

$$m_1 = \frac{1}{3}(2m_{K^*} + m_\rho). \quad (\text{D.26})$$

You might have worried that we do not know what the mixing term $m_{18} = m_{81}$ should look like, but the point is that the mixing term comes precisely from the fact that when $SU(3)$ is broken, $\langle \psi_8 | \psi_1 \rangle \neq 0$. In particular,

$$\langle \psi_8 | \psi_1 \rangle = \frac{1}{3\sqrt{2}} (\langle u\bar{u} | u\bar{u} \rangle + \langle d\bar{d} | d\bar{d} \rangle - 2\langle s\bar{s} | s\bar{s} \rangle), \quad (\text{D.27})$$

⁴¹This is a delightful phrase I borrowed from Tony Zee.

note the minus sign! We can thus read off

$$m_{81} = m_{18} = \frac{2}{3\sqrt{2}}(2m_u - 2m_s) = \frac{2\sqrt{2}}{3}(m_\rho - m_{K^*}). \quad (\text{D.28})$$

Good. Now we know that the flavor basis matrix is made up of these elements, but upon rotation to the mass basis the matrix is diagonalized to

$$R \begin{pmatrix} m_8 & m_{81} \\ m_{18} & m_1 \end{pmatrix} R^T = \begin{pmatrix} m_\phi & \\ & m_\omega \end{pmatrix}. \quad (\text{D.29})$$

We can relate the physical masses $m_{\phi,\omega}$ to the expressions above by considering invariants under rotations, namely the trace and determinant. The trace relation is easy and gives us

$$m_\omega + m_\phi = 2m_{K^*}. \quad (\text{D.30})$$

The determinant relation gives us

$$m_\omega m_\phi = \frac{1}{9} [(4m_{K^*} - m_\rho)(2m_{K^*} + m_\rho) - 8(m_\rho - m_{K^*})^2], \quad (\text{D.31})$$

where we recognize $m_8 = 3(4m_{K^*} - m_\rho)$ in the first term on the right-hand side. We would like to combine these equations into the Schwinger relation. This is most easily obtained by starting with the determinant relation and, in the first term on the left-hand side, writing

$$(2m_{K^*} + m_\rho) = 6m_{K^*} + m_\rho - 4m_{K^*} = 3(m_\omega + m_\phi) - 3m_8. \quad (\text{D.32})$$

This leads us to

$$3m_8 [3(m_\omega + m_\phi) - 3m_8] - 8(m_\rho - m_{K^*}) = 9m_\omega m_\phi. \quad (\text{D.33})$$

Rearranging terms and then factorizing we finally obtain

$$-(m_8 - m_\omega)(m_8 - m_\phi) = \frac{8}{9}(m_a - m_{K^*})^2. \quad (\text{D.34})$$

This is satisfied up to an error of $\mathcal{O}(160 \text{ MeV}^2)$, which is quite good considering that it is *quadratic* in mass.

3.7 The rotation from the ψ_8, ψ_1 basis to the physical basis can be written as

$$R^T M R = \begin{pmatrix} m_\phi & \\ & m_\omega \end{pmatrix}, \quad R \equiv \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (\text{D.35})$$

Inverting this relation one finds

$$\begin{pmatrix} m_8 & m_{81} \\ m_{18} & m_1 \end{pmatrix} = \begin{pmatrix} c^2 m_\phi + s^2 m_\omega & cs(m_\omega - m_\phi) \\ cs(m_\omega - m_\phi) & s^2 m_\phi + c^2 m_\omega \end{pmatrix}, \quad (\text{D.36})$$

where we've written $\sin \theta = s$ and $\cos \theta = c$ for short-hand. Now note the following relations:

$$m_8 - m_\omega = c^2(m_\phi - m_\omega) \quad (\text{D.37})$$

$$m_8 - m_\phi = s^2(m_\omega - m_\phi) \quad (\text{D.38})$$

$$m_1 - m_\omega = s^2(m_\phi - m_\omega) \quad (\text{D.39})$$

$$m_1 - m_\phi = c^2(m_\omega - m_\phi), \quad (\text{D.40})$$

where all we've used is $\cos^2 \theta + \sin^2 \theta = 1$. From this it is easy to take ratios and derive the relations in (3.27). For the $\tan \theta$ expressions use the observation that

$$m_{81} = c s(m_\omega - m_\phi) \quad (\text{D.41})$$

and take the appropriate ratios with the $\sin^2 \theta$ and $\cos^2 \theta$ expressions above. Finally, for the double angle formulae use

$$m_8 - m_1 = (c^2 - s^2)(m_\phi - m_\omega) \quad (\text{D.42})$$

and

$$\cos 2\theta = \cos^2 \theta - \sin^2 \theta \quad (\text{D.43})$$

$$\sin 2\theta = 2 \cos \theta \sin \theta. \quad (\text{D.44})$$

3.8 Writing out only the valence quark mass contribution (and taking $m_u = m_d$) we have

$$m_{K^*} = m_u + m_s \quad (\text{D.45})$$

$$m_\rho = 2m_u. \quad (\text{D.46})$$

From this we may express the diagonal flavor-eigenstate masses

$$m_8 = \frac{1}{3}(4m_{K^*} - m_\rho) \quad (\text{D.47})$$

$$m_1 = \frac{1}{3}(2m_{K^*} + m_\rho). \quad (\text{D.48})$$

Taking the inner product $\langle \phi_8 | \phi_1 \rangle$ we may also extract the off-diagonal mass term

$$m_{81} = m_{18} = \frac{2\sqrt{2}}{3}(m_u - m_s) = \frac{2\sqrt{2}}{3}(m_\rho - m_{K^*}). \quad (\text{D.49})$$

Armed with these it is straightforward to plug into (3.27), (3.28), and (3.29). To give a complete list for each expression,

$$\tan^2 \theta = -\frac{4m_{K^*} - m_\rho - 3m_\phi}{4m_{K^*} - m_\rho - 3m_\omega} \quad (\text{D.50})$$

$$= -\frac{2m_{K^*} + m_\rho - 3m_\omega}{2m_{K^*} + m_\rho - 3m_\phi} \quad (\text{D.51})$$

$$= \frac{4m_{K^*} - m_\rho - 3m_\phi}{2m_{K^*} + m_\rho - 3m_\phi} \quad (\text{D.52})$$

$$= \frac{2m_{K^*} + m_\rho - 3m_\omega}{4m_{K^*} - m_\rho - 3m_\omega}. \quad (\text{D.53})$$

$$\tan \theta = \frac{-2\sqrt{2}(m_\rho - m_{K^*})}{4m_{K^*} - m_\rho - 3m_\omega} \quad (\text{D.54})$$

$$= \frac{-2\sqrt{2}(m_\rho - m_{K^*})}{2m_{K^*} + m_\rho - 3m_\phi} \quad (\text{D.55})$$

$$= \frac{4m_{K^*} - m_\rho - 3m_\phi}{2\sqrt{2}(m_\rho - m_{K^*})} \quad (\text{D.56})$$

$$= \frac{2m_{K^*} + m_\rho - 3m_\omega}{-2\sqrt{2}(m_\rho - m_{K^*})}. \quad (\text{D.57})$$

$$\cos 2\theta = \frac{2}{3} \frac{m_{K^*} - m_\rho}{m_\phi - m_\omega} \quad (\text{D.58})$$

$$\sin 2\theta = \frac{4\sqrt{2}}{3} \frac{m_\rho - m_{K^*}}{m_\omega - m_\phi}. \quad (\text{D.59})$$

Note that the value of $\tan 2\theta$ obtained from the last two equations is independent of any meson masses. This is because (3.29) was written such that the m_ϕ and m_ω dependence cancels and we're left with two quantities which are calculated using valence quark masses. The upshot is that the value of θ derived this way is precisely what is predicted by *ideal mixing*.

3.9 When the mixing term $m_{xy} \ll m_x^2 - m_y^2$ then $\phi = 0$ and we have the ideal mixing condition. On the other hand, if the splitting $m_x^2 - m_y^2 \ll m_{xy}$ then $\phi \rightarrow 45^\circ$. This value is an asymptotic and so is relatively stable in this regime. For more discussion see Isgur [55].

3.10 These states do not mix because the light quark masses break $SU(3)$ to $U(1)_u \times U(1)_d \times U(1)_s$ such that strangeness is conserved.

3.11 See Griffiths *Introduction to Elementary Particles*, Chapter 5.6.2–5.6.3 (Chapter 5.10 in the first edition).

4.1 See [47]. (Joao Silva was attending the lecture where this problem was posed and only nodded that this is a problem in his book.)

4.2 This is straightforward: use (area) = (base) \times (height). To be somewhat pedantic, take the following steps

1. Normalize the triangle so that one side has unit length in the real direction.
2. The height of the triangle is given by the imaginary part of one of the other (normalized) terms in the sum.
3. Use the formula for the area of a triangle.
4. Return to the previous normalization by multiplying this area by the absolute value of the side that we scaled and rotated, squared (because this is an area).

4.3 You are free to pick any two unitary triangles you want, but two of them are particularly easy. Looking at the standard CKM parameterization (4.11) we see that V_{ud} , V_{us} , V_{cb} , and V_{tb} are purely real. Thus there are two triangles that have one leg parallel to the real axis:

$$\sum_i V_{id} V_{is}^* = 0 \quad (\text{D.60})$$

$$\sum_i V_{cj} V_{tj}^* = 0. \quad (\text{D.61})$$

Thus the base of the triangle (D.60) is

$$V_{ud} V_{us} = c_{12} c_{13} s_{12} c_{13}, \quad (\text{D.62})$$

while the base of the triangle (D.61) is

$$V_{cb} V_{tb} = s_{23} c_{13} c_{23} c_{13}. \quad (\text{D.63})$$

Great. That's the easy part. To calculate the height it is sufficient to take the imaginary part of either of the remaining terms in the sum. For triangle (D.60) we'll take $V_{cd} V_{cs}^*$

$$|\text{Im}(V_{cd} V_{cs}^*)| = |\text{Im}(s_{12}^2 s_{13} c_{23} s_{23} e^{-i\delta} - c_{12}^2 s_{13} c_{23} s_{23} e^{i\delta})| \quad (\text{D.64})$$

$$= \sin \delta s_{13} c_{23} s_{23}. \quad (\text{D.65})$$

For triangle (D.60) we'll take $V_{cs} V_{ts}^*$,

$$|\text{Im}(V_{cs} V_{ts}^*)| = |\text{Im}(-c_{12} s_{12} s_{13} c_{23}^2 e^{-i\delta} + c_{12} s_{12} s_{13} s_{23}^2 e^{i\delta})| \quad (\text{D.66})$$

$$= \sin \delta c_{12} s_{12} s_{13}. \quad (\text{D.67})$$

Comparing the (base) \times (height) of each triangle, we find that both triangles have area equal to one half of

$$c_{12} s_{12} s_{13} c_{13}^2 c_{23} s_{23} \sin \delta. \quad (\text{D.68})$$

This is indeed our definition of the Jarlskog, (4.21).

5.1 Look up references on ‘atomic parity violation.’ See, for example, [10, 50] for reviews. The main point is that part of the Z coupling is pseudovector-like which means its coupling is proportional to a fermion’s helicity. This is parity-odd unlike the electric force. The goal this, is to devise experiments that are sensitive to this left-right asymmetry. This is rather non-trivial, and a naïve estimate of the size of such an asymmetry is [10]

$$A_{\text{LR}} = \frac{P_L - P_R}{P_L + P_R} \quad (\text{D.69})$$

where $P_{L,R} = |A_{\text{EM}} \pm A_W^{\text{odd}}|^2$. The weak amplitude $A_W \sim g^2/(q^2 + M_Z^2)$ where the characteristic momentum scale is given by the Bohr radius,

$$q \sim \frac{1}{\alpha m_e}. \quad (\text{D.70})$$

This gives us

$$A_{\text{LR}} \approx \alpha^2 \frac{m_e^2}{M_Z^2} \approx 10^{-15}, \quad (\text{D.71})$$

which is hopelessly small. Fortunately, there are various enhancement mechanisms which make the measurement of this quantity experimentally tractable. See Section 2.4 of [10] for an excellent discussion.

6.1 In terms of Weyl spinors, (6.5) reads

$$\mathcal{L}_Z = \frac{g}{\cos \theta_W} \left[\bar{\chi}_u^i \bar{\sigma}_\mu \left(\frac{1}{2} - \frac{2}{3} \sin^2 \theta_W \right) \chi_u^i + \psi_u^j \sigma_\mu \left(-\frac{2}{3} \sin^2 \theta_W \right) \bar{\psi}_u^j + \text{h.c.} \right] Z^\mu. \quad (\text{D.72})$$

6.2 Refer to the solution to Problem (2.2) for background.

- In the interaction basis the couplings to the Z are given by the usual formula

$$g_z = g \cos \theta_W T^3 - g' \sin \theta_W Y. \quad (\text{D.73})$$

From the particle content we see that the singlets d_R and s_L have the same T^3 and Y quantum numbers. The other particles all have different quantum numbers and hence different couplings to the Z . From (D.6) we see that the bare mass terms cause the mass matrix for the d and s quarks to be different from the Yukawa basis. Let us say that this mass matrix M is diagonalized by $\hat{M} = W_L M W_R^\dagger$. Then the rotations

$$\begin{pmatrix} d'_{L,R} \\ s'_{L,R} \end{pmatrix} = W_{L,R} \begin{pmatrix} d_{L,R} \\ s_{L,R} \end{pmatrix} \quad (\text{D.74})$$

shift the interaction basis fields (d, s) to the mass basis fields d', s' . In terms of these fields, the coupling to the Z in the kinetic term is written as

$$(\bar{d}'_{L,R} \quad \bar{s}'_{L,R}) W_L^\dagger \begin{pmatrix} g_Z^{d_{L,R}} & \\ & g_Z^{s_{L,R}} \end{pmatrix} W_R i \gamma^\mu Z_\mu \begin{pmatrix} d'_{L,R} \\ s'_{L,R} \end{pmatrix}. \quad (\text{D.75})$$

The key point is that while $g_Z^{d_R} = g_Z^{s_R}$, i.e. the right-chiral coupling matrix is diagonal, $g_Z^{d_L} \neq g_Z^{s_L}$ so that the left-chiral couplings become non-diagonal after the bi-unitary rotation by W_L and W_R . Thus there *are* FCNCs in the Z coupling to the left-chiral down quark sector.

- There are no photon or gluon FCNCs. The photon and gluon couplings are all universal with respect to particles within a given flavor representation.
- Because the Yukawa couplings and the mass matrix (D.6) are not proportional to one another (i.e. not aligned) there are FCNCs from the Higgs. Note that these can be easy to ‘hide’ since the Higgs couplings to the light quarks is small.

- We now know that this model has tree-level FCNCs. As a quick and dirty approximation we can expect that the magnitude of the flavor-changing neutral currents and charged currents should be of the same order of magnitude. However, looking at the PDG, we can consider characteristic kaon decays (kaons because the new flavor structure is only in the s - d sector),

$$\text{Br}(K^+ \rightarrow \mu^+ \nu_\mu) = 64\% \quad (\text{D.76})$$

$$\text{Br}(K_L^0 \rightarrow \mu^+ \mu^-) = 7 \times 10^{-9}. \quad (\text{D.77})$$

We know that the leptons only couple to the bosons through the $SU(2)_L$ gauge bosons, so first decay is through the charged current while the second is through the neutral current. It would be very difficult to explain this discrepancy if a model has tree-level FCNCs and hence we expect this model to be ruled out.

6.3 The Yukawa and mass terms in the Lagrangian for this model takes the form

$$\mathcal{L}_{\text{Yuk.}+\text{mass}} = y_L \bar{Q}_L \phi s_R + y_R \bar{Q}_R \phi s_L + m_Q \bar{Q}_L Q_R + m_s \bar{s}_L s_R + \text{h.c.} \quad (\text{D.78})$$

These terms break $U(1)^4$ flavor symmetry in the kinetic terms to a $U(1)$ overall phase rotation on each field. We thus expect a total of $4 - 0 = 4$ physical real parameters and $4 - 3 = 1$ phases. The Dirac mass matrix takes the form

$$\begin{pmatrix} m_Q & y_L \frac{v}{\sqrt{2}} \\ y_R \frac{v}{\sqrt{2}} & m_s \end{pmatrix}, \quad (\text{D.79})$$

while the up-quarks have a single mass term $m_Q \bar{u}_L + \bar{u}_R + \text{h.c.}$. We may choose the real physical parameters to be the Dirac masses for each generation and the d - s mixing angle and the physical phase to be that of the up quark mass term. The coupling to the Z boson in the interaction and mass bases follow as it did in the previous problem. In particular, in the mass basis one must perform a rotation (D.74) so that the Z couplings take the form (D.75). Now both the left- and right-handed d and s quarks have different $SU(2)_L \times U(1)_Y$ quantum numbers so that $g_Z^d \neq g_Z^s$ and so the rotation introduces FCNCs through the Z for both chiralities.

As before the photon and gluon still do not mediate FCNCs because their couplings are all universal with respect to particles within a given flavor representation. The Higgs again induces FCNCs because the bare mass terms prevent the mass matrix and Yukawa matrix from being proportional to one another. The appearance of FCNCs in the Z and Higgs sectors lead us to expect this model to be ruled out experimentally.

6.4 For a nice history and collection of references about the 2HDM, see [48].

- The most general Yukawa potential takes the form

$$\mathcal{L}_{\text{Yuk.}} = (y_{ij}^{1u} \phi_1 + y_{ij}^{2u} \tilde{\phi}_2) \bar{Q}^i u^j + (y_{ij}^{1d} \tilde{\phi}_1 + y_{ij}^{2d} \phi_2) \bar{Q}^i d^j + \text{h.c.} \quad (\text{D.80})$$


- The diagonalization procedure proceeds as usual for the Standard Model except that the mass matrices are now

$$m_{ij}^u = y_{ij}^{1u} \frac{v_1}{\sqrt{2}} + y_{ij}^{2u} \frac{v_2}{\sqrt{2}} \quad m_{ij}^d = y_{ij}^{1d} \frac{v_1}{\sqrt{2}} + y_{ij}^{2d} \frac{v_2}{\sqrt{2}}. \quad (\text{D.81})$$

These are diagonalized by some biunitary transformation. As far as the kinetic terms are concerned, this rotation on the quark fields has the same effect as the rotation in the Standard Model. In particular, the Z couplings are still universal and remain flavor-diagonal.

- It is straightforward to see that there are FCNCs in the Higgs sector. For example, the neutral scalar Higgs components couple to the fermions by replacing $v_i/\sqrt{2}$ in the mass matrices with h_a for $a = 1, 2$ labeling each the two Higgs doublets. It is clear that in general the mass matrix is different from the Yukawa matrix for either Higgs since, e.g. $m_{ij}^u \neq y_{ij}^{au}$ for either $a = 1, 2$. There are, of course special cases; for example if these Higgses are allowed to mix in such a way that the mass eigenstate couples to combinations of Yukawas proportional to m^u and m^d . We can also note that in the limits $\tan \beta \equiv v_2/v_1 \rightarrow 0, \infty$ the Yukawas align with the mass matrices because one of the vevs vanishes. It is generally true in any multi-Higgs model that FCNCs are avoided in the Higgs sector so long as all the fermions of a given electric charge only couple to a single Higgs.
- Higgs FCNCs are typically stronger for heavier quarks, though in this 2HDM one can arrange for cancellations between large couplings in the Yukawas. A natural place to rule out this type of model is in the FCNC constraints of the B mesons.

7.1 The chiral diagram takes the following form:



The key point is that the W boson only couples to left-chiral fermions in the Standard Model. Thus in order to preserve angular momentum, there must be a Dirac mass insertion to convert the left-chiral fermion to a right-chiral fermion.

7.2 Method 1: Dirac spinors. For this calculation one needs to write in explicit projection operators. The amplitude takes the form

$$\mathcal{M} = \frac{g^2}{2} V_{ud} \bar{v}_d \gamma^\mu P_L u_u \frac{-i}{M_W^2} \bar{u}_\mu \gamma_\mu P_L v_\nu. \quad (\text{D.82})$$

The first spinor contraction is just $\frac{1}{2} \langle \pi^+ | A^\mu | 0 \rangle$, where A^μ is the axial current. Note the factor of $1/2$ coming from the fact that we're only taking the axial part of the projection operator, $\gamma^\mu P_L = \frac{1}{2} \gamma^\mu (1 - \gamma^5)$. Plugging in the expression for the pion decay constant,

$$\mathcal{M} = -\frac{g^2}{4M_W^2} V_{ud} f_\pi \bar{u} \not{p} P_L v \quad (\text{D.83})$$

$$= -\frac{g^2}{4M_W^2} V_{ud} f_\pi m_\mu \bar{u} P_L v, \quad (\text{D.84})$$

where in the second line we used the equation of motion for the outgoing charged lepton⁴². The benefit of this calculation is that one can use all of the usual Feynman rules from Peskin. However, things become a bit more clunky with factors of γ^5 . These can become distracting for more complicated problems.

Method 2: Weyl spinors. We don't have to worry about projection operators anymore since the external states are chiral. The amplitude is

$$\mathcal{M} = \frac{g^2}{2} \bar{\chi}_d \bar{\sigma}^\mu \chi_d \frac{-i}{M_W^2} \bar{\chi}_\mu \bar{\sigma}_\mu \chi_\nu \quad (\text{D.85})$$

$$= -\frac{g^2}{4M_W^2} V_{ud} f_\pi \bar{\chi} \not{p} \chi \quad (\text{D.86})$$

$$= -\frac{g^2}{4M_W^2} V_{ud} f_\pi m_\mu \psi \chi, \quad (\text{D.87})$$

where in the second line we have remembered to insert a factor of $1/2$ because we're only taking the axial part. The benefit of this calculation is that one doesn't have to deal with γ^5 . On the other hand, this also means that those factors are not there to remind you which piece is the axial part. In Weyl notation γ^5 simply shows up as a relative sign between different chiralities, which is somewhat subtle. **Remark:** Weyl spinors can really go a long way to simplify amplitudes with some chiral protection; for example $\mu \rightarrow e\gamma$ in an extra dimension [25].

Finally, the significance of the mass insertion should be clear from Problem 7.1: it signifies the chirality flip in this decay. The pion is spin zero and the W mediating the decay only couples to left-handed particles. In order to preserve angular momentum, one requires an explicit mass insertion on the final state fermions. Note that the use of the equation of motion is equivalent to a mass insertion on an external leg when using massless chiral fermions.

7.3 The formula for the decay rate is

$$d\Gamma = \frac{1}{32\pi^2} |\mathcal{M}|^2 \frac{|\mathbf{p}_\ell|}{m_\pi^2} d\Omega, \quad (\text{D.88})$$

where we've written the muon subscript as ℓ to avoid confusion with Lorentz indices. In case you didn't feel inclined to re-derive this formula and your favorite quantum field theory textbook isn't nearby, then you can always look this up in the kinematics review of the pocket PDG—which should *always* be nearby.

The first step is to square the matrix element and sum over spins. This matrix element should be particularly simple because we only have one spinor bilinear. This is worked out in most QFT textbooks so we won't belabor the four lines of work required to extract the relevant expression. Writing p_ℓ for the lepton four-vector and p_n for the neutrino four-vector (to avoid confusion with Lorentz indices) we obtain from (7.19),

$$\sum_s |\mathcal{M}|^2 = \left(\frac{g^2}{4M_W^2} \right)^2 f_\pi^2 |V_{ud}|^2 m_\ell^2 2p_\ell \cdot p_n. \quad (\text{D.89})$$

⁴²We've been a little sloppy here, but there is some elegance to being able to do a sloppy-but-accurate calculation. The pion momentum p should really be written in terms of $p = p_\mu + p_\nu$ and for each term the equation of motion for the appropriate lepton should be used. We already know, however, that $\not{p}_\nu v_\nu = 0$ since the neutrino is massless in the Standard Model.

The factor of $2p_\ell \cdot p_n$ is simply the result of summing over the spins of the spinor structure $|\bar{\mu}_R \nu_L|^2$. Note that this is *trivial* and just comes from the completeness relations of the plane wave spinors.

Some strategy. It is important to get the right answer, but it is also important to do so in a way that doesn't make your life difficult. This toy calculation is an important example. We have made use of the fact that we are only looking at the axial current so that our amplitude takes the form

$$\mathcal{M} \propto f_\pi (p_\ell + p_n)_\alpha \bar{\mu}_L \gamma^\alpha \nu_L. \quad (\text{D.90})$$

One could have naïvely taken the square of the spinorial part as written. If one were to work with Dirac spinors, this would be a terrible mess,

$$\bar{u}(p_\ell) \gamma^\alpha \frac{1}{2} (1 - \gamma^5) v(p_n). \quad (\text{D.91})$$

Squaring and summing over spins requires taking a trace over four γ matrices times $(1 - \gamma^5)$ and one obtains a funny relation in terms of two-index bilinears in p_ℓ and p_n including an ugly term with an $\varepsilon^{\alpha\beta\mu\nu}$. See, e.g. equation (5.19) in Peskin and Schroeder [65]. One will eventually find that the ε term cancels because it is contracted with $(p_\ell + p_n)_\mu (p_\ell + p_n)_\nu$. This is *a lot of work!*

We can do a little better by using Weyl spinors. Nima Arkani-Hamed was once quited as saying that he used to think that fermions were very difficult to work with because of all of the γ matrix manipulations required; but then he discovered Weyl spinors and everything was easy because the only structure comes from the 2×2 Pauli matrices. In this case the simplification is obvious: the ε term drops out from the very beginning because we never have to deal with the annoying γ^5 matrix. The relevant relations to perform the analogous calculation in Weyl space can be found in [28]; or in a pinch you can extract them from the Dirac spinor relations. This is indeed a meaningful simplification and justifies the use of Weyl spinors in 'actual calculations' (rather than just abstractly to refer to chiral fields); however, we can do better!

What we did above in (D.89) was to go a step further and get rid of the Dirac/Pauli matrix structure altogether. This was easy: we just used the fact that the γ^μ (alternately $\bar{\sigma}^\mu$) contracted a $(p_\ell + p_n)_\mu$. This means we can use the equation of motion to explicitly pull out the mass factor m_ℓ , where we have used $m_n = 0$. Now we get an expression of the form $\bar{u}(p_\ell) v(p_n)$ in Dirac notation, or in Weyl notation $\psi(p_\ell) \chi(p_n)$. Squaring and summing over spins is trivial for these since it is just the completeness relation for the plane wave spinors.

It is trivial to determine the contraction $2p_\ell \cdot p_n$ since this is constrained by the kinematics,

$$m_\pi^2 = (p_\ell + p_n)^2 = m_\ell^2 + 2p_\ell \cdot p_n. \quad (\text{D.92})$$

Thus we are led to (we drop the \sum_s symbol)

$$|\mathcal{M}|^2 = \left(\frac{2}{\sqrt{2}} G_F \right)^2 f_\pi^2 |V_{ud}|^2 m_\ell^2 m_\pi^2 \left(1 - \frac{m_\ell^2}{m_\pi^2} \right), \quad (\text{D.93})$$

where we've also replaced the $g^2/8M_W^2$ with $G_F/\sqrt{2}$ according to the definition of the Fermi constant (B.1). Let's simplify the other factors in (D.88). The magnitude of the muon three-momentum can be derived trivially from kinematics; conservation of four-momentum gives us

$$|\mathbf{p}_n| = |\mathbf{p}_\ell| \tag{D.94}$$

$$m_\pi^2 = |\mathbf{p}_n| + \sqrt{|\mathbf{p}_\ell|^2 + m_\ell^2}. \tag{D.95}$$

From this we obtain

$$|\mathbf{p}_\ell| = \frac{m_\pi}{2} \left(1 - \frac{m_\ell^2}{m_\pi^2} \right). \tag{D.96}$$

The angular integral is trivial

$$d\Omega = d(\cos\theta)d\phi = 4\pi, \tag{D.97}$$

and we can now plug in all of these factors to obtain the result we wanted

$$\Gamma = \frac{1}{8\pi} G_F^2 f_\pi^2 |V_{ud}|^2 m_\ell^2 m_\pi \left(1 - \frac{m_\ell^2}{m_\pi^2} \right)^2. \tag{D.98}$$

7.4 This problem is an excellent example of a form factor multiplying a non-trivial combination of dynamical variables.

1. The key difference between the D and D^* is that the latter has a polarization vector. Without this the parameterization of the matrix element reduces to (7.35).
2. The dynamical variables available to us are the D^* polarization ϵ and the two meson momenta p_B and p_D . Eventually we will repackage the momenta into $(p_B + p_D)$ and $q \equiv (p_B - p_D)$. Now we would like to use spacetime symmetries to determine the structure of the parameterization. The V^μ operator contains a Lorentz index, so we know the right-hand side must be a vector or axial vector. The matrix element on the left-hand side is necessarily parity-even since QCD respects P , so the right-hand side must be vectorial. In the next part we will use T to determine that $g(q^2)$ is real, but note that we *never* need to consider C since QCD also respects charge conjugation and the right-hand side is composed of dynamical variables which are all trivially even under C . Thus all we have to worry about is constructing a Lorentz vector out of ϵ , P_B , and p_D .

There is a subtlety in the parity of vectorial objects. The parity of a pure vector V^μ can be written as $P[V^\mu] = (-)^\mu$, by which we mean the parity is even (+) for $\mu = 0$ and odd (-) for $\mu = 1, 2, 3$. Note that the μ on $(-)^mu$ isn't a Lorentz index, it's just there to tell us whether $(-) = +$ or $-$. Thus a $J^P = 1^-$ meson has parity $-(-)^\mu$. We see that the parity of the terms in the matrix element are

$$P[V^\mu] = (-)^\mu \quad P[D^*] = - \quad P[B] = -. \tag{D.99}$$

We confirm that the matrix element is P -even with respect to ‘overall parities’ – and has the correct left-over ‘vector parity’ $(-)^{\mu}$ that is required for a single-index object. The dynamical objects that we have, however, do not seem to have the correct parities:

$$P[\epsilon^{\mu}] = -(-)^{\mu} \qquad P[p_{B,D}] = (-)^{\mu}. \qquad (\text{D.100})$$

In fact, we can see that the only non-zero vector bilinear $\epsilon_{\mu} p_B^{\mu}$ also does not have the correct parity. One might be led to believe that this matrix element must vanish. However, we have one more trick up our sleeves. In d -dimensional space we have a d -index totally antisymmetric $\varepsilon^{\mu_1 \dots \mu_d}$ tensor with which we can construct Lorentz contractions. In more formal language, we have the additional operation of taking a Hodge dual to convert p -forms into $(d-p)$ forms. Our our present case this enables us to consider triple products of ϵ , p_D , and p_B to construct a Lorentz vector. The natural object to write down is

$$\varepsilon^{\mu\nu\alpha\beta} \epsilon_{\nu}^* p_{D\alpha} p_{B\beta}. \qquad (\text{D.101})$$

The key thing to recall is that ε has parity

$$P[\varepsilon^{\mu\nu\alpha\beta}] = -(-)^{\mu} (-)^{\nu} (-)^{\alpha} (-)^{\beta}, \qquad (\text{D.102})$$

where we’ve written out the usual ‘vector parity’ for a four-tensor-like object. Because of the additional overall minus sign, we see that ε is a **pseudotensor**. One should already be familiar with this in the coordinate definition of the volume form in which the ε tensor carries information about orientation. Armed with this we now see that the $(\varepsilon \epsilon^* p_D p_B)^{\mu}$ contraction is indeed a parity-even vector to which the hadronic matrix element $\langle D^{*+}(p_D, \epsilon) | V^{\mu} | \bar{B}(p_B) \rangle$ may be proportional. To complete the analysis, we remark that we are free to change momentum variables to $(p_D + p_B)$ and $q \equiv (p_D - p_B)$. It is clear that we are free to make the replacement

$$p_{D\alpha} p_{B\beta} \longrightarrow (p_D + p_B)_{\alpha} q_{\beta} \qquad (\text{D.103})$$

since these are contracted with the ε antisymmetric tensor so that only the $p_{D\alpha} p_{B\beta} - p_{D\beta} p_{B\alpha}$ term is picked out. Finally, the overall coefficient $g(q^2)$ can only be a function of q^2 since this is the only dynamical Lorentz scalar quantity.

3. The matrix element transforms as a complex conjugate under time reversal. The only complex element from our dynamical variables on the right-hand side is ϵ^* , which contains an imaginary element. The sign of this element represents the transverse polarization of the D^* and we expect it to flip under complex conjugation. Otherwise, the rest of the element is real and so $g(q^2)$ must also be real.
4. The axial current is not pure QCD and comes from the chiral nature of the weak interactions and hence we are allowed to have $\langle D^{*+}(p_D, \epsilon) | A^{\mu} | \bar{B}(p_B) \rangle$ which is ‘overall parity’ odd (i.e. $P = -(-)^{\mu}$). One immediate choice is a term proportional to $\epsilon^{*\mu}$ since this has precisely the correct parity and Lorentz structure. In addition to this, we can also form objects out of the three dynamical variables since $\epsilon^*(p_D + p_B)q$ also has the same correct parity and Lorentz structure; we just have to insert the Lorentz index at each place. The $\epsilon^{*\mu}(p_D + p_B) \cdot q$ term

is the same as the $\epsilon^{*\mu}$ since $(p_D + p_B \cdot q)$ is not dynamical (it's a sum of masses). We are left with terms of the form

$$\langle D^{*+}(p_D, \epsilon) | A^\mu | \bar{B}(p_B) \rangle = g_{A1}(q^2)\epsilon^{*\mu} + g_{A2}\epsilon^* \cdot (p_D + p_B)q^\mu + g_{A3}(\epsilon^* \cdot q)(p_D + p_B)^\mu. \quad (\text{D.104})$$

7.5 Let us write the hadronic matrix element for neutron β decay $n \rightarrow p$ through the axial current,

$$\langle p(k_p) | A_\mu | n(k_n) \rangle = \bar{u}_p(k_p) [g_A(q^2)\gamma_\mu\gamma_5 + q_\mu h_A(q^2)\gamma_5] u_n(k_n). \quad (\text{D.105})$$

Here g_A and h_A are functions which depend on the momentum transfer $q \equiv k_n - k_p$. Note that we are *not* ‘parameterizing our ignorance’ about QCD (we still have cumbersome spinor structure on the right-hand side); instead, we are saying that the neutron decays into the proton *literally* through the tree-level coupling to the axial current. If you want to be precise about isospin indices, one ought to write $A_\mu \rightarrow A_\mu^+ = A_\mu^1 + iA_\mu^2$, but we’ll try to keep the equations as simple as possible. We can now take the divergence of this equation to obtain

$$\langle p(k_p) | \partial_\mu A^\mu | n(k_n) \rangle = i [g_A(q^2)m_N^2 + q^2 h_A(q^2)] \bar{u}_p(k_p)\gamma_5 u_n(k_n), \quad (\text{D.106})$$

where m_N is the nucleon mass. The left-hand side of this equation, however, is simply what we called the pion field in (7.48), up to overall factors:

$$\langle p(k_p) | \partial_\mu A^\mu | n(k_n) \rangle = f_\pi m_\pi^2 \langle p(k_p) | \phi_\pi | n(k_n) \rangle \quad (\text{D.107})$$

$$= \frac{2f_\pi m_\pi^2}{q^2 - m_\pi^2} g_{\pi NN}(q^2) i \bar{u}_p(k_p)\gamma_5 u_n(k_n), \quad (\text{D.108})$$

where in the second line we have defined the pion-nucleon vertex function $g_{\pi NN}(q^2)$ and pulled out the pion pole. This gives the πNN coupling for $q^2 = m_\pi^2$. Comparing (D.106) and (D.108) we obtain,

$$f_\pi g_{\pi NN}(0) = m_N g_A(0). \quad (\text{D.109})$$

This is nearly the Goldberger-Treiman relation, except that the pion-nucleon vertex function is off-mass shell. We must make the assumption that $g_{\pi NN}(q^2)$ is slowly varying such that $g_{\pi NN}(0) \approx g_{\pi NN}(m_\pi^2)$. This gives us a relation which holds within 10%. Note that this relation is actually rather trivial from the point of view of chiral perturbation theory, see e.g. Donoghue Section 12.3 [58].

E Critical reception of these notes

Thanks for writing these notes. Keep it up—or else I’ll have to take my own notes during lecture. P.A.

I’ve learned a lot from your flavor notes. It’s the only way I can learn flavor physics from my adviser. Y.T.

You're still taking classes? Aren't you a little old for that? D.S. (comment via instant message during lecture)

I would suggest putting more details to derive this equation. I do so little algebra in class that I feel bad to leave the little I that do out... Y.G.

That's too hard to remember, I prefer to just remember the number. P.A. (after Yuval explained the *Spiked Math* mnemonic [1] to remember the number of feet in a mile and, hence, the mass of the B meson in MeV)

These are important stories, not like the usual stories that I tell you. Y.G.

F Famous Yuval Quotes

There's one thing you can do, it's called a charge thing. It has something to do with charge. [29 Oct 2010, after being asked how to do a certain measurement.]

You've never seen penguins that are twins? [2 Nov 2010, describing the penguin-mediated $D\bar{D}$ diagram.]

Calculation is not for us—only approximation. [2 Nov 2010]

References

- [1] Spiked Math Comics: 5 tomato feet.
- [2] Wikipedia, the free encyclopedia: Astronomy on Mars.
- [3] Fifth CERN-Fermilab Hadron Collider Physics Summer School, 2010. Streaming video of lectures available online.
- [4] M. Ademollo and Raoul Gatto. Nonrenormalization Theorem for the Strangeness Violating Vector Currents. *Phys. Rev. Lett.*, 13:264–265, 1964.
- [5] Thomas Appelquist and J. Carazzone. Infrared Singularities and Massive Fields. *Phys. Rev.*, D11:2856, 1975.
- [6] Nima Arkani-Hamed. Introduction to Low Energy Supersymmetry, 2007. Lectures at the Summer School on Particle Physics, Cosmology and Strings.
- [7] Christopher Arzt. Reduced effective Lagrangians. *Phys. Lett.*, B342:189–195, 1995.
- [8] Riccardo Barbieri, Alex Pomarol, Riccardo Rattazzi, and Alessandro Strumia. Electroweak symmetry breaking after LEP-1 and LEP-2. *Nucl. Phys.*, B703:127–146, 2004.
- [9] E. Blucher et al. Status of the Cabibbo angle. 2005.

- [10] M.-A. Bouchiat and C. Bouchiat. Parity violation in atoms. *Reports on Progress in Physics*, 60:1351–1396, November 1997.
- [11] C. Glenn Boyd and David E. Brahm. Vanishing of $1/m$ corrections at threshold. *Phys. Lett.*, B257:393–398, 1991.
- [12] Nora Brambilla, Antonio Pineda, Joan Soto, and Antonio Vairo. Effective field theories for heavy quarkonium. *Rev. Mod. Phys.*, 77:1423, 2005.
- [13] C. P. Burgess. Introduction to effective field theory. *Ann. Rev. Nucl. Part. Sci.*, 57:329–362, 2007.
- [14] Cliff Burgess and Guy Moore. *The Standard Model: A Primer*. Cambridge University Press, 2006.
- [15] Nicola Cabibbo, Earl C. Swallow, and Roland Winston. Semileptonic hyperon decays. *Ann. Rev. Nucl. Part. Sci.*, 53:39–75, 2003.
- [16] G. Cacciapaglia, C. Csaki, G. Marandella, and A. Strumia. The minimal set of electroweak precision parameters. *Phys. Rev.*, D74:033011, 2006.
- [17] Robert Cahn. *Semi-Simple Lie Algebras and Their Representations*. Dover Publications, 2006.
- [18] Robert N. Cahn. The Eighteen arbitrary parameters of the standard model in your everyday life. *Rev. Mod. Phys.*, 68:951–960, 1996.
- [19] Curtis G. Callan, Jr., Sidney R. Coleman, J. Wess, and Bruno Zumino. Structure of phenomenological Lagrangians. 2. *Phys. Rev.*, 177:2247–2250, 1969.
- [20] Ta-Pei Cheng and Ling-Fong Li. *Gauge Theory of elementary particle physics*. Oxford University Press, 1988.
- [21] S. Coleman. How to build phenomenological lagrangians. In *Erice 1968, Theory and Phenomenology In Particle Physics. International School*, New York 1969, 648-657.
- [22] Sidney Coleman. *Aspects of Symmetry: Selected Erice Lectures*. Cambridge University Press, 1988.
- [23] Sidney R. Coleman, J. Wess, and Bruno Zumino. Structure of phenomenological Lagrangians. 1. *Phys. Rev.*, 177:2239–2247, 1969.
- [24] Roberto Contino. Tasi 2009 lectures: The Higgs as a Composite Nambu- Goldstone Boson. 2010.
- [25] Csaba Csaki, Yuval Grossman, Philip Tanedo, and Yuhsin Tsai. Warped Penguins. 2010.
- [26] Csaba Csaki, Matthew Reece, and John Terning. The AdS/QCD Correspondence: Still Undelivered. *JHEP*, 05:067, 2009.

- [27] Savas Dimopoulos and David W. Sutter. The Supersymmetric flavor problem. *Nucl. Phys.*, B452:496–512, 1995.
- [28] Herbi K. Dreiner, Howard E. Haber, and Stephen P. Martin. Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry. 2008.
- [29] G. Ecker. Chiral perturbation theory. *Prog. Part. Nucl. Phys.*, 35:1–80, 1995.
- [30] Adam F. Falk, Yuval Grossman, Zoltan Ligeti, and Alexey A. Petrov. SU(3) breaking and D0 - anti-D0 mixing. *Phys. Rev.*, D65:054034, 2002.
- [31] JiJi Fan, Matthew Reece, and Lian-Tao Wang. Non-relativistic effective theory of dark matter direct detection. 2010.
- [32] Thorsten Feldmann. Quark structure of pseudoscalar mesons. *Int. J. Mod. Phys.*, A15:159–207, 2000.
- [33] Flip Tanedo. Fun and Supersymmetry: but mostly the latter. <http://www.lepp.cornell.edu/~pt267/files/notes/FlipSUSY.pdf>.
- [34] Flip Tanedo. 't Hooft and η' : Instantons and their applications. http://www.lepp.cornell.edu/~pt267/files/documents/A_instanton.pdf.
- [35] J. Gasser and H. Leutwyler. Quark Masses. *Phys. Rept.*, 87:77–169, 1982.
- [36] H. Georgi. Effective field theory. *Ann. Rev. Nucl. Part. Sci.*, 43:209–252, 1993.
- [37] Howard Georgi. Heavy quark effective field theory. HUTP-91-A039, available at <http://www.people.fas.harvard.edu/~hgeorgi/tasi.pdf>.
- [38] Howard Georgi. On-shell effective field theory. *Nucl. Phys.*, B361:339–350, 1991.
- [39] Howard Georgi. *Lie Algebras In Particle Physics: from Isospin To Unified Theories (Frontiers in Physics)*. Westview Press, 2 edition, 1999.
- [40] Howard Georgi. *Weak Interactions and Modern Particle Theory*. Dover Publications, 2009.
- [41] Gary W. Gibbons, Steffen Gielen, C. N. Pope, and Neil Turok. Measures on Mixing Angles. *Phys. Rev.*, D79:013009, 2009.
- [42] David Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, 2nd edition, 2008.
- [43] Benjamin Grinstein. An Introduction to heavy mesons. 1995.
- [44] Yuval Grossman. Introduction to flavor physics. 2010.
- [45] Yuval Grossman and Gilad Perez. Realistic construction of split fermion models. *Phys. Rev.*, D67:015011, 2003.
- [46] Rajan Gupta. Introduction to lattice QCD. 1997.

- [47] João Paulo Silva Gustavo C. Branco, Luís Lavoura. *CP violation CP violation CP Violation*. International Series of Monographs on Physics. Oxford University Publications, 1999.
- [48] Howard E. Haber. Low-energy supersymmetry and its phenomenology. *Nucl. Phys. Proc. Suppl.*, 101:217–236, 2001.
- [49] Zhenyu Han and Witold Skiba. Effective theory analysis of precision electroweak data. *Phys. Rev.*, D71:075009, 2005.
- [50] W. C. Haxton and C. E. Wieman. Atomic parity nonconservation and nuclear anapole moments. *Annual Review of Nuclear and Particle Science*, 51(1):261–293, 2001.
- [51] Andreas Hocker and Zoltan Ligeti. CP violation and the CKM matrix. *Ann. Rev. Nucl. Part. Sci.*, 56:501–567, 2006.
- [52] Timothy J. Hollowood. 6 Lectures on QFT, RG and SUSY. 2009.
- [53] Stephen D. H. Hsu and Pierre Sikivie. Long range forces from two neutrino exchange revisited. *Phys. Rev.*, D49:4951–4953, 1994.
- [54] T. Inami and C. S. Lim. Effects of Superheavy Quarks and Leptons in Low-Energy Weak Processes $k(L) \rightarrow \mu \text{ anti-}\mu$, $K^+ \rightarrow \pi^+$ Neutrino anti-neutrino and $K^0 \text{ } \bar{\nu}_i \text{ } \nu_j$ anti- K^0 . *Prog. Theor. Phys.*, 65:297, 1981.
- [55] Nathan Isgur. Why the Pseudoscalar Meson Mixing Angle Is 10-Degrees. *Phys. Rev.*, D12:3770–3774, 1975.
- [56] Karl Jansen. Domain wall fermions and chiral gauge theories. *Phys. Rept.*, 273:1–54, 1996.
- [57] C. Jarlskog. Commutator of the quark mass matrices in the standard electroweak model and a measure of maximal cp nonconservation. *Phys. Rev. Lett.*, 55(10):1039–1042, Sep 1985.
- [58] Barry R. Holstein John F. Donoghue, Eugene Golowich. *Dynamics of the Standard Model*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 1994.
- [59] L. Lavoura. General formulae for $f_1 - i f_2 \gamma$. *Eur. Phys. J.*, C29:191–195, 2003.
- [60] G. P. Lepage. Lattice QCD for Novices. 1998.
- [61] Harry Lipkin. *Lie Groups for Pedestrians*. North-Holland Publishing Company, 1966.
- [62] Michael E. Luke. Effects of subleading operators in the heavy quark effective theory. *Phys. Lett.*, B252:447–455, 1990.
- [63] Aneesh V. Manohar. Effective field theories. 1996.
- [64] A.V. Manohar and M.B. Wise. *Heavy Quark Physics*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 2004.

- [65] Daniel Schroeder Michael Peskin. *An Introduction To Quantum Field Theory*. Frontiers in Physics. Westview Press, 1995.
- [66] Heechang Na, Christine T. H. Davies, Eduardo Follana, G. Peter Lepage, and Junko Shigemitsu. The $D \rightarrow K, l\nu$ Semileptonic Decay Scalar Form Factor and $|V_{cs}|$ from Lattice QCD. 2010.
- [67] Matthias Neubert. Heavy quark symmetry. *Phys. Rept.*, 245:259–396, 1994.
- [68] Matthias Neubert. Heavy-quark effective theory. 1996.
- [69] Holger Bech Nielsen and M. Ninomiya. No Go Theorem for Regularizing Chiral Fermions. *Phys. Lett.*, B105:219, 1981.
- [70] Donald Perkins. *Introduction to High Energy Physics*. Cambridge University Press, 2000.
- [71] A. Pich. Introduction to chiral perturbation theory. *AIP Conf. Proc.*, 317:95–140, 1994.
- [72] A. Pich. Chiral perturbation theory. *Rept. Prog. Phys.*, 58:563–610, 1995.
- [73] Joseph Polchinski. Renormalization and Effective Lagrangians. *Nucl. Phys.*, B231:269–295, 1984.
- [74] H. David Politzer. Power Corrections at Short Distances. *Nucl. Phys.*, B172:349, 1980.
- [75] N. F. Ramsey. ELECTRIC DIPOLE MOMENTS OF ELEMENTARY PARTICLES. *Rept. Prog. Phys.*, 45:95–113, 1982.
- [76] William Wagner David Pines Richard Feynman, Fernando Moringo. *Feynman Lectures On Gravitation*. Frontiers in Physics. Westview Press, 1995.
- [77] Janusz Rosiek, Piotr Chankowski, Athanasios Dedes, Sebastian Jager, and Philip Tanedo. *SUSY_FLAVOR : acomputationaltoolforFCNCandCP – violatingprocessesintheMSSM*. 2010.
- [78] Ira Z. Rothstein. TASI lectures on effective field theories. 2003.
- [79] David J. Gross Sam B. Treiman, Roman Jackiw. *Lectures on current algebra and its applications*. Princeton Series in Physics. Princeton University Press, 1972.
- [80] M. D. Scadron. Current Algebra, PCAC and the Quark Model. *Rept. Prog. Phys.*, 44:213–292, 1981.
- [81] Stefan Scherer. Introduction to chiral perturbation theory. *Adv. Nucl. Phys.*, 27:277, 2003.
- [82] Martin Schmaltz and David Tucker-Smith. Little Higgs Review. *Ann. Rev. Nucl. Part. Sci.*, 55:229–270, 2005.
- [83] Julian Schwinger. A Ninth Baryon? *Phys. Rev. Lett.*, 12:237–239, 1964.

- [84] Stephen R. Sharpe. Phenomenology from the lattice. 1994.
- [85] Mikhail A. Shifman. ITEP Lectures in Particle Physics. 1995.
- [86] Witold Skiba. TASI Lectures on Effective Field Theory and Precision Electroweak Measurements. 2010.
- [87] R. Slansky. Group Theory for Unified Model Building. *Phys. Rept.*, 79:1–128, 1981.
- [88] Matthew J. Strassler. An unorthodox introduction to supersymmetric gauge theory. 2003.
- [89] Robin Ticciati. *Quantum Field Theory for Mathematicians*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1999.
- [90] Steven Weinberg. Current-commutator calculation of the k_{l4} form factors. *Phys. Rev. Lett.*, 17(6):336–340, Aug 1966.
- [91] Steven Weinberg. Approximate symmetries and pseudoGoldstone bosons. *Phys. Rev. Lett.*, 29:1698–1701, 1972.
- [92] James Wells. Status of the Standard Model. Lectures at the Particle Physics, Cosmology and Strings Summer School at Perimeter Institute, August 2007.
- [93] James D. Wells. TASI lecture notes: Introduction to precision electroweak analysis. 2005.
- [94] Scott Willenbrock. Symmetries of the standard model. 2004.
- [95] Mark B. Wise. Heavy quark theory. 1994.