MITOCW | 3. Field Redefinitions

IAN STEWART: So last time, we were midway through this proof that the equations of motion, i.e., field redefinitions, can be used to simplify the theory. And I stated last time that really, all we have to worry about is the change to the Lagrangian. But when we looked at the path integral to do things properly and make a change variable in the path integral, it wasn't just the Lagrangian that changed.

The Lagrangian changed, and we could set up our field redefinition to do what we want, but there were also changes to the Jacobian and the source, and that's what we had started to talk about last time. So when I look at the change the Jacobian, which is this thing, we could write that as a ghost Lagrangian. And the argument for why we don't need to worry about this is as follows.

So the effective field theory is going to be valid for small momentum. There's an expansion and there's some scale introduced by the higher dimensional operators, lamda new, and we're looking at low momentum relative to that. In the particular case that we're dealing with, this parameter eta had dimensions, and so lambda nu was 1 over root eta.

And we put factors of eta in front of our higher dimension operators, so that's like putting 1 over lambda nus in front of those operators. And the reason that we don't have to worry so much about this ghost Lagrangian is that these ghosts are going to get mass that's of the size lambda nu.

So what that means is that the ghost, along with perhaps other particles that are up at the scale lamda new are things that we don't have to worry about, and we can effectively integrate out the ghost from the theory in the same way that we would think about removing some particles that had masses of order lambda nu. So I'll show you why they get masses of order lambda nu in a second.

So let's do that by way of picking a particular example. So far, we've kept things fairly general without specifying what this set of fields it's in this thing that we called T was. We just left it as an arbitrary thing. Let's pick a particular T. The argument here is actually fairly general, but I think it'd be easier to see it for particular example.

So if we pick this to be T, we have this term which has no relation to T. That term is important. And then we have the terms from this. And so we'd have a ghost Lagrangian that would look like that. So this here is going to be the mass term, and so far, it doesn't look like it has the right dimensions. And that's because we've got our kinetic term, if you like, with the wrong dimensions.

So if we want ghost fields of standard dimensions, or canonically normalized kinetic term in this case, we would take c and rescale it to c over root eta. And if we do that, then this does become a lambda nu. So that removes the eta in the interaction term here. It removes it from this kinetic term that happened to be there, and then the only place that the eta's showing up now in this case is in this term.

But now you see explicitly in terms of these ghost fields that they have mass lambda nu. OK, and that's just what I was claiming. So the important point here was that you did need this 1. This 1 was playing a very important role in this argument. If the 1 wasn't there, this wouldn't work.

And that 1 was related to something that was part of one of our starting assumptions, that when we made the field redefinition, that the one-particle states would stay the same. So if you look back at what our field redefinition was and you trace back where that 1 came from, we had a field redefinition that was of this form, and this guy was a function of other fields. But we always had this term, and that term was related to the presence of this 1, and that was important for this argument, OK.

So there has to be a term that's linear in the new field. You want the same one-particle states. So these ghosts, as you know about ghosts, always appear in loops. And so if you want to think about removing this massive particle, it's like a massive particle that occurs in some loops. It's really just like that. So it's just like a heavy particle that you would remove from the theory that would only appear in loops.

And we're going to discuss exactly how that works and how to remove heavy particles that appear in loops in detail a little bit later, so for now, that's enough detail for number two. So are there any questions about that? OK, so the change to the Jacobian is effectively a massive ghost that we could integrate out of the theory.

When we integrate it out of the theory, it would shift the values of the couplings in the Lagrangian, but it doesn't have any impact that can't be absorbed in other operators in the effective theory. So then the final thing we have to worry about is the source. So there is this term with j phi dagger, and when we take-- there was an extra term with j phi dagger that we had last time.

I'm going to remind you what it looked like. So in our path integral, we had an exponential of a bunch of stuff, but then there was one final term that was induced that involved j phi dagger from the field redefinition. So when you take derivatives with respect to j phi dagger, there's sort of the usual term that we want to source, and then there's this new term.

So what about that new term? That's the final thing we have to worry about. And this is where it's important that we're actually considering observables. So we need to consider observables. So let's start by considering some Green's functions or time-ordered products.

So I'm just taking a string of fields. I could have some other fields here as well. I'm going to make my life a little easier by taking the phi to be real, and that's just for notational simplicity. I could write phi dagger and phis, but. So when we make the field redefinition, what happens to that time-ordered product?

So this is one way of thinking about these extra terms in the source. So you take functional derivatives with respect to the other piece that involves j phi, you get phi. So if you take [INAUDIBLE] that piece, you get the extra eta Ts. Or you can just think about making the field redefinition directly in this matrix element, and you'd have a matrix element that looks like that.

So it's not at all obvious that these extra pieces that involve fields can't change the value of this Green's function, and in fact, they will change the value of that Green's function. But the important thing is that we have to consider observables and not just Green's functions. So when we're going to think about observables, we should think about the LSZ formula, which connects Green's functions to S-matrix elements.

So let me remind you about that. And I'll remind you what it looks like. First one is fixed set of fields, which are scalar fields, but we could put fermions in it as well. So what this says is, if I look at this-- so it's an integral over the spacetime. I'm sticking in particular momentum for the particular fields. And if I look at it, if I look at the leading term, the leading pole term as the particles are taken on shell, then that's an observable known as the S-matrix.

So some number of these particles are incoming. Some number of them are outgoing. That affects the sign that I put in this plus minus. I'm not trying to be too detailed about which ones are outgoing, which ones are incoming, but some of them are incoming, some of them are outgoing. And for each one, I want to strip off the leading pole. I want to look at the coefficient, and it has to have a pole, so this is a product over all particles.

We really need to pull in all the different cases and all for all the external particles, OK. So if I wanted to make this an equality, I'd say plus dot dot dot. There's other terms, and the thing that's observable is the coefficient of these poles, not the Z but this thing. So if I make changes to my theory and they effect this thing but get cancelled by this thing, as long as they don't affect this thing, we're OK.

And the claim is that we make this change to the source and it won't affect the matrix on the S, the S-matrix. So again, we can do some examples, and I think the examples are fairly quickly convincing that this is the case. So we'll do three different examples, first a rather trivial one.

What if T was just by itself? So T is just phi, so this is just 1 plus eta phi. So if you think back, when we have T, this was the form of the terminal Lagrangian, so this is just eta phi del squared phi, so it's just changing the kinetic term for phi. If we look at our matrix on it, we're just getting a factor of 1 plus eta for each of the fields.

So let's say we had 4 of them just so we don't have to-- so say I had 4 of them. If I had n of them, it would just be to the nth power, but if I have power of them, they'd be to the fourth power, and I'd just get this extra prefactor. And it looks like it's changed G.

And indeed, it has changed the left-hand side of this equation, but when we calculate the z factor and we canonically normalize-- if you want to think about it that way-- then we would exactly cancel off these 1 plus etas. So the root Z here is going to be 1 plus eta as well.

So when we look at the-- this is just the residue of the free propagator. The residue gets changed when you change the kinetic term, so that doesn't do anything to the S-matrix. So that's one way we can be protected. We change the left-hand side. It's compensated by a change to the Z and leaves S invariant.

Let's do another example, a little more non-trivial. So let's just take some cubic term in our field redefinition. So this is like having a 5 cubed del squared phi term. So this extra term will give rise to extra terms in the primary product. And let me just write one of them, again thinking of it as a 4-point function.

So let's just imagine I look at the change that comes from changing the 4th guy. So there's other terms from changing 5x1, 5x2 and 5x3, but if I work to order eta, then I'd change one of them at a time, and we're only working to order eta here. So the claim is actually that this matrix element has no effect on this structure.

It can affect the dots, but it's not going to affect the leading term, and that's because having a 5 cubed means that you're not having a one-particle state. So if you try to draw it as a [INAUDIBLE] diagram, in the position space, you'd label these external points and you inject momenta there.

And then we label point 4, but there's three fields there, so maybe we have to tie it up like this or something. And when you look at asymptotically what's going to happen from [INAUDIBLE] a 5 cubed, you do not get a single-particle pole from a 5 cubed. So this guy is less singular. There's no single-particle pole, and hence, gives no contribution to scatter. Yeah.

AUDIENCE: As far as the theorem goes, this step that you're showing, it proves my hypothesis, right?

IAN STEWART: In what sense?

AUDIENCE: The hypothesis of the theorem is that [INAUDIBLE] does not change a one-particle state.

IAN STEWART: That's right. That was an assumption.

AUDIENCE: I mean, is that the same as what you're--

IAN STEWART: It's the same, yeah. It's effectively the same, though it's being more careful about what the statement of that means, right, because I mean, you're changing the residue but you're not changing the S-matrix.

AUDIENCE: Right. This is what I was thinking when you said that.

IAN STEWART: Yeah, all right. So really, this statement of-- yeah, I could have been a little more careful. So instead of saying that the hypothesis of the theorem is that it wouldn't change one-particle states, what I should have really said is that there's a linear term in the field redefinition. Yeah, which now I'm showing you is equivalent to not changing the one-particle states.

And let's do one final example just to rule out all possible things we could think of that are different. So what if we had a derivative [INAUDIBLE] squared phi prime? So then we would do the following. We could add and subtract a mass term. The reason we would want to do that is, if we looked at a term like this, it's, again, less singular than this term. If this term is giving the one particle state, this term has got a factor of the propagator upstairs.

So if you look for a pole that would come from that term, it's canceling out. You get p squared minus m squared over p squared minus m squared. So there's no pole from this term. And this term is, again, just of the type from our first example. It's just shifting 5 by some constant. So probably to get dimensions right, I should put some factors of eta in here. So again, something like that doesn't change anything, because it can be decomposed into something that has no pole and then something that's just a shift, OK?

So as long as you have this linear term in your field redefinition, and it doesn't have to even have a trivial-- it doesn't have to even have a trivial prefactor. You could have 2 times phi or whatever you like, because that'll cancel out when you normalize things correctly.

As long as you have that linear term, you're fine. And you don't have to worry about changes to the Jacobian or changes from the source term. So you just have the changes to the Lagrangian to worry about. You don't have to think about the path integral. Just make field redefinitions on the Lagrangian. And that's what you'll get some practice with on the problem set.

So as I said at the beginning before everybody was here, there's a problem set number 1 that's posted. Everyone should make sure that they can actually access it. And if they can't access it, they should let me know-- if you can't get access to the web page. Any final questions about this before we move on? Yeah?

AUDIENCE: I get nervous about trying to cancel a pole by adding an m squared term, because then, I guess, at some point, you start to worry that masses get shifted when you start [? normalizing these. ?]

IAN STEWART: Oh, yeah. Yeah. Yeah, you should be careful about that, too. But you don't have to worry about it.

Yeah.

You could think of it as doing the mass renormalization first and then worry about this, or-- so that you're using a renormalized mass here.

All right. So we'll start a new section. And the goal of this section is basically to be more careful about loop diagrams and, in particular, to show you what matching is between two effective field theories. We'll still be thinking in the context of mass of particles. That's the simplest place to describe this that we'll probably do an example later on of a case that doesn't just involve mass of particles.

So let's start out with a very trivial example just to see what we're talking about. So we'll take two particles. One of them is a heavy scalar. It has mass capital M. Actually, I put a little underline so you can tell my capitals from my lowercase, because the light fermion that we're also going to have has a lowercase mass, little m. OK. So heavy scalar, light fermion.

And so what we want to talk about here is really this picture that we described to you earlier where we've got two theories, 1 and 2, and we want to pass from theory 1 to 2 by removing something from theory 1. So theory 1 here will be just the theory of these things. And we'll make it a renormalizable theory, in the traditional sense, just so we know where to stop. And this guy should be capital. And then there's some [INAUDIBLE] interaction between them. So we'll think of this as being the theory 1, where we're in a situation where capital M is much bigger than little m.

So then we want to think about describing psi at low energies. And that means we can get rid of the scalar as an explicit degree of freedom. So low energy is relative to this mass scale, which is heavy-- M, capital M. We're going to remove the scalar from the theory.

So you could just say remove it, or you could say integrate it out. When you say integrate out, the words you're using have a path-integral connotation, where you had this [INAUDIBLE] in the path integral, and you think about just doing the path integral over that field and removing it. But the words of "removing it" or "integrating it out" are synonymous.

So what is theory 2 going to look like? Well, we'll still have the kinetic term for our fermion. And then removing the scalar will generate some new operators. In particular, there'll be a dimension-six operator like this. And then there could well be other terms. And so if we wanted to figure out what this dimension-six operator is at tree level, that's a pretty straightforward exercise. We would simply think about the Feynman diagram in theory 1.

So here's a Feynman diagram in theory 1 with [INAUDIBLE] couplings here. We would calculate this Feynman diagram. And we would assume that all the momentum of the external particles are small. And that means that this is small. And we would just expand.

And then we would just ask that, when I take the Feynman rule from this-- whoops, should have been a square there-- when I take the Feynman rule from that, then I should get the same thing as the Feynman rule from that. And that would fix this coefficient a, which is, so far, arbitrary. But we can determine it by using theory 1. And that's the idea of doing a matching calculation that you have introduced the theory 2. It has some parameters in front of operators. In this case, I called it a.

And we want to determine those parameters by doing calculations in theory 1. And in particular, what you do is you make sure that S-matrix elements in the two theories agree. But at tree level, that's just matching up diagrams.

So a is simply g squared. So we just have to match up this guy with that guy. And so taking the leading-order term, a is just g squared. Very simple.

So the place where this gets more involved and more-- where you have to think a little bit more-- this seems almost automatic, that you just do calculations in here and here, and just match them up. And that's the basic idea. It's only a little more complicated when you have to take into account loops. So that's what we'll spend most of this section discussing, since this part is easier.

So what about loops? What are some of the issues that come up there? Well, the first one is that the Feynman diagrams diverge, so you have to regulate them. And the thing that can be confusing about thinking about effective field theories and thinking about divergence diagrams is separating the ideas of regularization and the mass scale lambda nu that we've been talking about.

So that's something I want to talk about in some detail, because they're not the same thing. So we would need to cut off ultraviolet divergences to obtain finite results. And that means we introduce cutoff parameters into our results.

So examples would be taking the Minkowski momentum, continuing it to be Euclidean, and then just putting a hard cutoff on it-- that's one example. Or you could use dim reg. That's another example. Or you could use a lattice spacing. That's another example. So these are all examples of how you might cut off the theory to remove ultraviolet divergences.

And then there's a second step, renormalization, distinct from regularization. So some of the things I'm teaching you here should be familiar. And the only real generalization that we're going to do is we're going to be able to apply some of the things that you learned about renormalization and regularization to any operators that you might have in the theory, whether they're renormalizable in the dimension-four traditional sense or in a higher-dimensional sense.

So here, what we're doing when we talk about renormalization is we're picking a scheme that gives precise or definite meaning to the parameters in the effective theory. Or, even more explicitly, each coefficient in the Lagrangian, as well as the operators in the Lagrangian, are given a meaning by this procedure.

If we didn't have a renormalization procedure, then, because of the ultraviolet divergences, there'd be ambiguities in how we define the coefficients and the operators. Or, even if you don't have ultraviolet divergences, you still have freedom to pick different schemes for the definitions of the coefficients.

And when you do this, you also can introduce parameters. So you can get parameters from regularization. You can also get parameters from this. So some examples that are familiar-- there's this scale mu that shows up when you do MS bar. If you do something that's a little bit different, if you take Green's functions, and you go to some offshell point, that's another renormalization scheme called offshell momentum subtraction.

And that scheme also has a parameter that shows up. And if you did a Wilsonian renormalization, there would also be a cutoff associated to that. And this Wilsonian cutoff here doesn't have to be the same as this lambda uv cutoff.

OK. So every coefficient, every operator, no matter the dimension, no matter where it turns up in the series, we're going to have to think about whether there's diagrams that generate ultraviolet divergences. And if those diagrams look like these operators, then you're going to get renormalization of those operators. And you're going to have to be careful about taking care of those divergences and also the scheme that you define the operators in.

So you can think about this, just as far as the coefficients are concerned, as starting out with some bare coefficients that depend on your ultraviolet regulator, and switching to some renormalized coefficients which don't depend on the ultraviolet regulator but do depend on the scheme. And then, also, you have some counterterms that depend on both things.

So that's how it would look if you wrote that down for a cutoff, in a Wilsonian sense. If you wrote it down in dimensional regularization, it would look like this. Same idea, different names for the parameters. OK. So in dimensional regularization, epsilon is the ultraviolet regulator. It shows up in the bare coefficients. The renormalized ones don't depend on epsilon anymore. And these are your 1-over-epsilon poles.

All right. So one of the things that makes this tricky is when you start thinking about your power counting. And so I want to spend a few minutes talking about that. So let's consider this example that I wrote down with the four-fermion operator. And if we just loop up the four-fermion operator, then we can get mass renormalization. So this is psi dagger psi squared. And I'm sticking it in and looping it up.

So we start out with a tree-level mass, little m. But there's a one-loop correction that involves one of our higher-dimension operators. And there's some connection to the mass, which I'll call delta m. I'm not going to worry too much about the overall prefactor.

This operator came with an a over M squared. And then there's a loop integral. There's a fermionic propagator, so that's a k slash plus m over k squared minus m squared.

This guy here drops away. And really, we just have a correction that's a scalar. And that's a correction to the mass scalar in the spin space. So it's a, little m, over capital M, and then just this integral. And if we continue it to be Euclidean, then it looks like that if I continue from Minkowski to Euclidean. That's why I put the i there.

So if we just assume that this integral that's sitting there-- there's only one mass scale that seems explicit there-- the little m. So if we just assume that that integral is dominated by the scale m-- little m-- then you could do a power counting for this loop integral.

You just assume that all the factors of k are of order little m. You've got an explicit little m, four powers in the numerator, and two powers downstairs. So you would say that this integral scales like m squared. And then you would put that together with your delta m and say that delta m scales like a, little m cubed, capital M squared.

And so this is something that would then be suppressed. Relative to the original tree-level contribution, it would be suppressed by, perhaps, some [? 4 pi ?] [? is in ?] the loop, but also little m squared over big M squared. So it would be a small correction, which is really what we would like to be the case when we're talking about some higher-dimension operator that was supposed to be suppressed. We'd like it to be a small correction. All right.

So does anyone have any idea what can go wrong with this argument? I've been glib about it. I just said, if the integral is dominated by that. So the thing is that we have to be careful about our choice of regulator, because regulators also can introduce scales. So let's do this more carefully.

Well, let's consider two different choices for the regulator. So we'll start out with just a cutoff, which seems very natural, from an effective field theory point of view-- that your theory is supposed to be only valid up to some scale. So why not just take and cut off the momentum explicitly above some scale?

And you can think of that scale as being of order big M. So we just don't include any momentum that are higher than that scale in our loop integrals. So we're only integrating over the region in momentum space where the theory is supposed to be valid. It seems like a perfectly reasonable thing to do.

So in this case, we can take this integral that has some angular parts to it as well as a radial part, decompose it into the radial piece and the angular piece. I'll give you a handout, or I'll post a page of lecture notes that have all those fun formulas that are useful to remember but not fun to talk about for decomposing integrals in arbitrary dimensions into radial and angular pieces.

So if we do that, in this case, then the radial integral-- we can cut off the radial integral with a hard cutoff lambda uv. And just by dimensions, the d 4 k becomes kE times kE cubed. So this is a radial k now. And this integral, we can do exactly.

There's a loop factor, 4 pi squared. There's a lambda uv squared. And then there's a logarithm. It goes like, lambda uv squared over little m squared. There's two scales that are showing up-- lambda uv and little m. The answer depends on both of them.

And then we can start expanding, because little m is supposed to be much smaller than lambda uv. Let me pull out an a m over 4 pi squared. So there would be a logarithmic term. There's this lambda uv squared over capital M squared if I put that factor in. And then there's some-- so everything I'm writing in the square brackets here is dimensionless. This has the right dimensions of a mass.

OK. So you can see that, if I do this, it's not satisfying what I said over here. I'd like the correction to be m cubed over M squared. There are corrections that go like m cubed over M squared, like this one. These ones are even higher. But there's this term. And that's not a small correction.

So for that particular piece, what you're finding is that your power counting-- your naive power counting-- of the loop was wrong, because there's a piece from k of order-- the cutoff that's contributing for that part of it. And that's why your naive scaling argument didn't work.

Now, this is the bare result. And we have to go through the renormalization procedure. And so if we go through the renormalization procedure, you can think of that as taking a piece of the integral-- so in a Wilsonian sense, you would take a piece of the integral and absorb it into the counterterm. So as promised, the counterterm depends on both lambda uv and lambda. And those are just explicitly-- I'm cutting off the radial integral here.

And that does improve things, because then, I can lower this cutoff, make it smaller. And what is left would just be lambda squared over capital M squared and log of m squared over capital lambda squared. So I change the lambda uv's up here into lambdas by that procedure.

So when I renormalize, the psi bar psi squared matrix element-- there's a correction from that to the mass. And the renormalized thing would depend on this Wilsonian cutoff lambda, OK? But I'm not fully getting around the issue that I'm generating this type of term.

So let's do it with a different regulator, which is dimensional regularization. Let's see what happens using the MS-bar scheme. Same calculation. So again, split it into radial and angular parts.

And again, I'll give you formulas for doing something like that. But I won't write them down in lecture. Actually, I think I will write one down a little later. But I'll give you a more complete set as a handout.

So leaving over the radial integral but regulating it dimensionally just means that, instead of having this guy to the cubed power, it's to the d minus 1. And there's some pieces here that depend on epsilon. And my convention for the entire course is that d is 4 minus 2 epsilon, not d minus epsilon.

So again, this is an integral we can do exactly. We get something like that. And then we can expand. We get a 1-over-epsilon pole. We get a logarithm. There is some constant as well. And then there's order-epsilon pieces.

So you should contrast this result here with the result we had over here. The order-epsilon pieces are like these terms. They're the terms that would go away if I took the cutoff here to infinity-- lambda uv to infinity or epsilon to 0. There's terms here that go like m cubed over M squared. That's like this term. And this term is not there.

This 1 over epsilon is related to the log divergence. And this thing here is a power divergence. So the 1 over epsilon is related to this log of lambda uv. You should think of log of lambda uv as going to 1 over epsilon plus log of mu squared. But we don't see the power divergence in dimensional regularization in MS bar.

OK. So what's happening is, from this result, we are getting something that's the size that we expected by our power-counting argument. The regularized result is the right size by power counting. The logarithm here is the same log as in a with, if you like, a correspondence between mu and lambda.

So we're seeing that logarithm there. We're just not seeing this term. And if we wanted to write down the MS-bar counterterm. Some correction to the mass, it would go like, a, m squared. The diagram would look like just the 1-over-epsilon pole in MS bar.

OK. So the two regulators seem to give us similar results but not exactly-identical results. And the question is, how should I think about this? What's the right way of thinking about this extra term? So there's a language that people use in effective field theory and dimensional counting for the difference between thinking about doing a regularization of type A and type B.

And what you say is that, in type B, your regularization of the problem does not break the power counting, which means that you can do power counting for regulated graphs prior to renormalization them. But you can see that, in case A, that would be problematic, because our naive power counting wouldn't have given us the right result.

So what do we say about case A? Well, if case B didn't break the power counting, then case A does. But you can say something a little bit more positive about case A. And that is that, when you think about case A, you can set up the theory with renormalization conditions such that you can restore power counting in the renormalized graphs and renormalized couplings and renormalized operators. But you don't have power counting-- explicit power counting-- in just the regularized results.

So you can think that I just take the Wilsonian cutoff small enough that I make that annoying term small. And order by order, I can do that-- order by order in my loop expansion, order by order in my calculations. OK. So in some sense, there's nothing wrong with doing A. It's just that you have to work a little harder to think about it. And when you think about the power counting, you have to do that for the renormalized quantities, not for the bare quantities.

So if you like, what you're doing here is you're adding counterterms. Another way you can say this is that you're adding counterterms to restore the power counting that you want. And that's not too different than-- you could always do that, actually. I made this analogy that you should think about power counting as being like gauge symmetry.

So let's say you had a theory, and it had nice gauge symmetry, but you picked some crazy regulator that broke gauge symmetry. Well, you could always put counterterms in that would break gauge symmetry and restore it in the renormalized quantity. That would be OK. It would be more work. We don't like to do that. We avoid it at all costs. But if we had to, we could do it.

If you do supersymmetry, you might try to use dimensional regularization and supersymmetry, standard dimensional regularization and supersymmetry, break supersymmetry. If you do that, you have to introduce counterterms that restore supersymmetry.

So the same language of symmetry is being applied here, except now to power counting, where we say that if your regulator messes up your power counting, you can restore it in the renormalized couplings. But you may be smart enough to think up a regulator-- in this case, dimensional regularization-- where you don't have to deal with that complication. OK. So let me write that.

So in an effective theory, you should think about regulating to preserve symmetries as well as to preserve power counting, if you can. And one way, in a more formal language, that you could say what happens with the generation of that term that we talked about with the cutoff is that you'd mix up different orders in the expansion, and it looks like your naively higher-order term is mixing back to a term of lower dimension. And so if you can get away with taking a regulator that doesn't have that mixing back to more relevant operators, then you could preserve your power counting, make it simpler.

This is true irrespective of what the power counting is in. This is a general statement. In the context of what we've been talking about, which is dimensional power counting, there's a particular phrase that goes along with this. And that is that we talk about using a mass independent regulator, like dim reg. If you like, it has a mass scale, mu, but it's put in softly in a way that doesn't mess up power counting. We call that using a mass independent regulator.

So we want to avoid having different orders in the expansion mix up with each other. In general, I should comment, then, that terms that are the same order will definitely typically mix up with each other under renormalization.

So even if you thought you were smart, and you enumerated all the operators, but you missed one, and then you started calculating, that operator might just pop out at you, because you could have some calculation with another operator mix into that operator. Or you could have an operator that you did a matching at tree level, and you didn't generate, but then you start renormalizing that tree-level operator, and another operator pops out at you.

So it's important that you include all the operators that have the same dimension and the same quantum numbers, because if you don't include them, you're bound to get them from loops anyway. And you want to be complete. So if you like, in a matrix [INAUDIBLE] notation, you could say that the bare operators mix up with the set of possible renormalized operators, and there's some matrix of counterterms that would correct them-- connect them.

OK. So any questions so far? So sometimes, in the literature, you'll see that these kinds of things-- regulator discussions in effective field theory-- generate all sorts of papers. Keep this in mind if you ever run into that. You should be able to think about the physics that you're after with different regulators and come to the same conclusion. And it just may be easier with one regulator versus another.

So we've said, in these kind of theories where we have dimensional power counting, that dimensional regularization is special. So I want to talk a bit more about dimensional regularization. So sometimes you'll hear people say that you should always use dimensional regularization for doing the power counting.

But that's not quite true in the way that I've told you. It's not that you have to, it's just that it makes things simpler. But given that we want to make things as simple as possible, let's take dimensional regularization seriously.

So you can actually derive dimensional regularization by just imposing axioms. If you say that you want a loop integration that's linear-- I should have said this earlier. So my notation with dimensional regularization is, I put a little cross on the d. And that means dividing by the 2 pi. So that means d d p over 2 pi to the d.

So linearity means that if I'm integrating some function that can be decomposed into a sum of two pieces, a and b being constants, f and g being functions, then I can write that out as an integral over f plus an integral over g, which really is something that almost every reasonable definition of the integration will satisfy.

The second one is translations, which is more restricting. So that says, if you have some integral over f, but it's a function of p plus q-- q is some external-- I can always shift away the q. It just goes-- p goes to p minus q. And then I just have an integral over p. And along with translations, you can think about having rotations. My whole notation is covariance, so we won't worry so much about rotations. and Lorentz group.

And then the final one that's obviously a little bit special to dim reg is a scaling. So let's say we have a scalar s multiplying our momentum p. Then we can rescale the momentum p and get rid of-- pull the s outside by just taking p goes to p over s. So that changes this to a p. We get an s to the minus d. It pulls out front.

And if we demand that, then that's special to dimensional regularization, because you can see that this depends on d. Even if I call this measure some abstract thing, now there's a d showing up. It's outside the measure. And these three together actually give a unique definition to the integration up to the overall normalization. And that unique thing is dim reg.

So I'm going to refer you to reading, I have posted a chapter from Collins' book on regularization. And around page 65, he talks about how you prove that. It's not too hard. The standard definition of the normalization, which is something you have to specify, is that you let, say, this Gaussian integral be pi to the d over 2. So then, from that, you have some measure in some space that you can then just use.

So one formula that I used earlier on was the ability to split that into pieces which were a radial piece and then an angular piece. And in general, this is a property that this integration [? measure obeys, ?] that you could split it, and you could split it further. You could pull out another angle, for example, and get one less dimension in the angular parts.

And the uv divergences, if we're talking about us divergences-- they're occurring in this radial part-- Euclidean radial part. So by thinking about this kind of decomposition, you're moving the uv divergences to a one-dimensional integration, at least at one loop. And you can always do that.

So in general, in dimensional regularization, there's many ways you could evaluate the integral. You're not used to using this one. You're used to keeping things covariant, using some Feynman parameters, combining propagators together, and then doing the integral. But you could also do it this way, and you get the same answer. So it's really a well-defined measure in the sense that you can manipulate it in different ways. And they should all lead to the same answer for your loop integrals. And that's part of what I'm trying to emphasize by saying that you could derive it by considering axioms.

So d was equal to 4 minus 2 epsilon. Epsilon greater than 0 is what you need to lower the powers of p, and therefore tame the uv. Epsilon less than 0 can be used to regulate infrared singularities. There's some counterintuitive facts about dimensional regularization, and I want to mention a couple of them to you. One of them is that, if I have p to an arbitrary power-- think of it as Euclidean-- that's 0.

So Collins constructs a proof of this on page 71, which is actually a little more involved, in general. I'll just give you an idea of how you can see that, from using our axioms, that something like this better be true. So let's consider a special example that won't be enough to prove it for arbitrary alpha. This is any alpha. Let's consider a special example that at least will be enough to prove it for integers. So we'll consider k's that are integers. And we'll think of k's that are greater than 0.

So if I just expand out this p plus q squared, then the first term is p to the 2k. Then I get some coefficient, p to the 2k minus 2, q squared, some coefficient, p to the 2k minus 4, q to the fourth, et cetera. In general, there's p dot q terms as well. But then I could do integral-- angular average and combine those together with these terms. And that's why I'm not being very explicit about what the coefficients are. But they're some positive numbers.

Now, I could also take this integral, and I could shift. That was one of our axioms. And that is p to the 2k. So that means that all these terms here better be 0. And they have to be 0 for arbitrary q and arbitrary k, or any k under the assumptions that we used, which are that it's an integer and it's positive. So I could expand it in this way. And so therefore, we have all these integrals over p to the powers. And they better be 0. And that's enough to prove this for integer alphas-- positive integer alphas.

OK. And so if you want to fill in between the integers and you want to do the negative cases, then you have to look at Collins. But you could do that, too. Then it requires a little more heavy lifting.

There's one fact about this-- fact number one-- which is a little bit subtle, and you have to be careful. And it's worth noting. So let me do another example, which is by way of warning you that this can sometimes be dangerous. So let's think of a scalar-field theory and a simple loop diagram like this. But let's take 0 momentum and 0 mass.

So if you do that, you'll encounter an integral that looks like this. There's two propagators, so I get a p to the fourth downstairs, and I get d d p. And that integral is 0, but it's 0 in a special way. It's 0 due to a cancellation between ultraviolet and infrared physics.

I said that epsilon could be regulating both infrared divergences as well as ultraviolet divergences. If I only used epsilon to regulate ultraviolet divergences, I'd get a 1-over-epsilon uv. But in this integral, I'm actually using it to do both. It's regulating an infrared divergence as well. And it just comes in with the opposite sign. And since epsilon uv is equal to epsilon IR-- they're just notation to signify what region of physics is giving the divergence-- you get 0.

But even though that's true, that doesn't mean you don't have to add a counterterm for this diagram, because counterterms are supposed to cancel ultraviolet divergences, not infrared ones, OK? So even though it's 0, you still need to add a counterterm, because the 0 is actually a cancellation between ultraviolet and infrared physics. So there's some counterterm. And it would be exactly of this sort, because this is the epsilon uv. And then if you add the bare diagram plus the counterterm, the answer is non-zero. You've canceled, if you like, the uv pole, and you've left over the IR pole.

OK. So you have to be a bit careful about using dimensional regularization, because if you encounter scaleless integrals, it could be that they actually are affecting counterterm. And if you want to do some renormalization-group improvement of the theory or something like that, you have to be aware of this.

If you know that all the infrared poles are going to cancel because you're looking at some infrared safe quantity, then you can be a little bit glib about this, because if they're going to cancel in the end, that means that any of these corresponding uv poles will also cancel. But it's not always the case that you're renormalizing operators that have no 1-over-epsilon IR's [INAUDIBLE] have to be careful about this. So this is a subtlety that sometimes people get wrong when they write papers. So dim reg is beautiful, but there are some things about it that are a little bit tricky.

Another thing that can be confusing about dim reg is that it does this, that it regulates both uv and IR poles. And even though it's doing that, and even though you need different values of epsilon, if you want to do that, it's actually still a well-defined procedure, even in the presence of uv and IR poles, even if they're both in the same integral. And basically, you're using analytic continuation here.

So let me give you a little example which is not exactly related to this but will allow me to show you both how you use analytic continuation and how you could think about separating uv and IR poles. So we'll start out just by thinking about analytic continuation.

So suppose I had some integral. So what does analytic continuation mean? Or, how should I construct it? So let me, again, write it in a way where I've separated out the radial integral. And let's suppose that this integral here is perfectly well-defined for d in some range. So this is well-defined for some range of positive d. And then let's say that we wanted to continue that integral to negative d.

Now, the problem with negative d may be that, when you get to negative d, you're getting some infrared divergences, and you have to figure out how to deal with them. And you can do this, if you'd like, step by step. So if we wanted to extend the lower limit down to minus 2 from 0, then we would do the following.

We take our integral, write it out in this angular/radial separation, split up, in the radial variable, the piece that's ultraviolet, which is the high-momentum piece, from the piece that's infrared. And in the piece that's infrared, we could also just do some addition and some subtraction to make it more convergent as p goes to 0.

So for example, we could subtract f at 0. This thing would fall off faster and hence, give more powers of p. And we could make it more convergent at 0. And then we just integrate the subtraction up to the cutoff c.

And the idea of introducing this cutoff c is that we split the uv piece and the ultraviolet piece. And so we can do one kind of continuation for d up here, making epsilon positive, one kind of continuation down here. And the result, when we put these back together, is independent of c, OK?

So that's the sense, actually, in which what I said up here is-- that you can use it for both uv and IR divergences, because you could always introduce some parameter c to split-- they're occurring in different regions of space base, so you could always split them up, regulate each one with different values of d, and then put them back together. And the answer, when you put them back together, is independent of c. OK.

Now, if you wanted to define-- that was one of our goals. The other one was just to show you what we mean by analytic continuation. So since it's independent of c, then you could do the following. So for minus 2 less than d less than 0, let's take c goes to infinity. And so for that particular range, what you find, then, if you take that limit--

Well, because of the c to the minus d, the c is downstairs if d is negative. So that term goes away, and you're just left with this term. This term here goes away, too, because c approaches the upper limit, and everything is regulated. And so I just would be left with that.

OK. Making d negative is-- OK. So that would be the definition. So you can see some of the kind of tricks that you could use with dimensional regularization or adding and subtracting terms. And these are the things that are valid things to do. Any questions about that? OK.

So when we do dimensional regularization in MS bar, you're used to doing that for a gauge theory. That's what you've learned about. But you can also do it for any effective field theory. And the logic is the same. So let me remind you of the gauge-theory logic and then just tell you how you would define MS bar precisely for the fermionic effective theory with the "psi bar psi squared" operator that we had. So we talked about dimensional regularization. Let's talk about the MS scheme.

So if we've set up our effective theory in a way where we've made the mass scale explicit, which is often a nice thing to do-- and we did that when we set up the effective theory where you have the capital M showing up explicitly. If you do that, then the couplings start out dimensionless. And that's a nice thing, to have dimensionless coupling. And the MS scheme is simply the scheme where you want to introduce a scale to keep the renormalized couplings dimensionless.

So the example you're familiar with is just having a gauge coupling. And if you go through the dimensions of the fields here, which I do in my notes, but I'm going to assume that you've got some familiarity with this, you find that the bare coupling has dimension epsilon. And so you define a renormalized coupling as dimensionless. And you introduce a factor of mu to the epsilon. So you say g bare, which has dimensions, some dimensionless z factor, some mu to the epsilon to make up those dimensions. And then left over is the renormalized coupling.

And the idea and the strategy for any other coupling in the effective theory is the same, so let's do one other example. So take our [? dimension-six ?] a bare over capital M squared, psi bar psi squared. Do dimension counting on this guy, which I do, again, in the notes.

But if you go through that dimension counting, and you remember that the dimensions for the fermions are assigned by the kinetic term-- so we have a dimension counting for the fermions. We know that this is dimension minus 2. The whole thing has to add up to d. So that tells us what the a bare is. And we get 4 minus d, in this case. And so then we can write down a formula analogous to that one but for the a coefficient. And it's mu to the 2 epsilon. OK. So it's as simple as that.

So in looking at the action where this is a term in the Lagrangian, we want to ensure that, when we go to the-- that we figure out the dimensions of this guy in dimensional regularization. That's 4 minus d, which we determine from knowing the other pieces here. And then we just make a redefinition to give a dimensionless coupling.

So that's how MS bar will work. Well, this is MS, but this is how minimal subtraction works for defining all the operators that you may have. And then minimal subtraction is simply a rescaling. And that's the same as it is in gauge theory, where we get rid of some annoying factors, and [? g is ?] a slightly different definition. So this was MS, and this is MS bar.

OK. So it really works in a very similar way to gauge theory. And you figure out the factors of mu to the epsilon to include in your calculation this way. Sometimes you see books do it by saying the loop measure is continued within mu to the epsilon. That's not right. This is right. If you do that, you'll get into trouble-- not always. That's why the books can get away with it. But in general, you'll get into trouble.

All right. So we should stop there.