

Quantum Field Theory IV - Renormalization

Nic Ford

1 Introduction

This article is part of a series on physics for mathematicians, and the third in a sub-series on quantum field theory. It is a direct sequel to the first, second, and third articles in that series, and it will be very helpful to be familiar with the contents of those articles before tackling this one.

In the previous article, we examined ϕ^4 theory, the interacting scalar field theory with Lagrangian

$$\mathcal{L} = \frac{1}{2}((\partial\phi)^2 - m^2\phi^2) - \frac{\lambda}{4!}\phi^4.$$

We discussed how to express this theory's scattering amplitudes — and in particular the momentum-space time-ordered n -point functions $\tilde{G}^{(n)}$ that appear in the LSZ formula — in terms of power series whose terms are indexed by Feynman diagrams. Using the Källén–Lehmann spectral representation, we also found a useful connection between the two-point function and the quantities μ and Z that appear in the LSZ formula: the quantity multiplying the delta function in $\tilde{G}^{(2)}(p, -q)$ has a pole at $p^2 = \mu^2$, and the residue at that pole is Z .

As pretty as this story was, it had a rather serious problem: almost none of the integrals it produces converge. In the final section of the last article, we discussed why this happened and started to sketch out a method for solving it. It starts with the observation that the parameters m and λ that appear in the Lagrangian are not physically measurable. (The same is in fact true of Z , as we'll discuss when we carry out the computation.) Even if we didn't have these divergent integrals to worry about, we would need to address this: in order to compare our results with experiment, they need to be expressed in terms of quantities whose values we can actually discover!

But we do have some divergent integrals to worry about, and perhaps surprisingly this injunction to write everything in terms of measurable quantities offers a way out of that problem as well. The basic idea is to start by writing all of our diagram integrals as limits of finite quantities as some parameter called a “cutoff” goes to infinity, and then write these finite quantities as functions of our physically measurable quantities. This part of the process is called **regularization**. When we do this, m , λ , and Z will be functions of the cutoff, but the hope is that the scattering amplitudes themselves will — as a function of our new, physically measurable parameters — converge to something finite.

This procedure, where we rewrite the value of each Feynman diagram in terms of our new, more physically relevant parameters, is called **renormalization**, and it's the key final step that we'll need to extract numbers from a quantum field theory that we can compare with experiment.

In this piece, I hope to describe how it works well enough for you to understand the idea, but I won't explain all the computations in as much detail as most physics texts. If you're interested in digging into that side of the story, I'll include references to other sources as we go.

It's also worth mentioning that there are two different points of view one can take on renormalization, often called the "classical" and "Wilsonian" perspectives. I think it's useful to understand both to have a complete picture of what's going on, but this article will focus entirely on the classical perspective. We'll have a bit more to say about the relationship between these two perspectives at the end of this article, and I hope to explore the Wilsonian perspective in depth later in this series.

There were several books I found helpful in the process of putting this article together, especially Peskin and Schroeder's *An Introduction to Quantum Field Theory*, Gerald Folland's *Quantum Field Theory: A Tourist Guide for Mathematicians* and Michel Talagrand's *What Is a Quantum Field Theory? A First Introduction for Mathematicians*. I am very grateful to Jordan Watkins for his helpful comments on an earlier version of this article.

2 Two More Diagram Games

Before we can embark on our journey, we need to say a couple more things about the relationship between Feynman diagrams and scattering amplitudes. Recall that our method for computing scattering amplitudes goes through the LSZ formula. The formula — which is valid as long as all the momenta are distinct — is

$$\langle p_{n+1}, \dots, p_{n+m} | p_1, \dots, p_n \rangle^{\text{in}} = \frac{(-i)^{n+m}}{Z^{(n+m)/2}} \prod_{k=1}^{n+m} (p_k^2 - \mu^2) \tilde{G}^{(n+m)}(p_1, \dots, p_n, -p_{n+1}, \dots, -p_{n+m}),$$

where μ is the mass of the particle and Z is the constant that appeared in the relationship we found between the interpolating field and the one-particle state. (Specifically, this relationship was $\langle p | \phi(0) | \Omega \rangle = Z^{1/2}$, but this won't matter much going forward.) In particular, recall that this means that $\tilde{G}^{(n+m)}$ has a pole at 0 in the variables $p_i^2 - \mu^2$.

2.1 Proper Diagrams

The Feynman rules we discussed last time give a method for computing $\tilde{G}^{(n+m)}$. Since we're about to dig into some concrete computations, it will be helpful to turn this into a method for computing the scattering amplitudes directly. This will have a couple of side benefits as well: it will mean we don't have to think about the pole in $\tilde{G}^{(n+m)}$, and it will reduce the number of Feynman diagrams that we need to worry about.

Recall that the value of every connected Feynman diagram contains a delta function forcing the sum of the incoming momenta to be equal to the sum of the outgoing momenta. Since $\tilde{G}^{(n+m)}$ is a sum of diagrams, this delta function will appear in the scattering amplitude as well. It will be nice to have some notation which gets rid of it: let's write

$$\langle p_{n+1}, \dots, p_{n+m} | p_1, \dots, p_n \rangle^{\text{in}} = (2\pi)^4 \delta \left(\sum_{i=1}^n p_i - \sum_{j=1}^m p_{n+j} \right) i\mathcal{M}(p_1, \dots, p_n; p_{n+1}, \dots, p_{n+m}).$$

This function \mathcal{M} is sometimes called the **invariant matrix element**, and our goal will be to rewrite the Feynman rules in a way that computes \mathcal{M} instead of $\tilde{G}^{(n+m)}$. (The i multiplying it in that equation is a convention which is sadly universal enough that I don't think I can drop it.)

It will also be helpful to have some notation for the quantity multiplying the delta function in $\widetilde{G}^{(n+m)}$ itself. Let's write

$$\widetilde{G}^{(n+m)}(p_1, \dots, p_{m+n}) = (2\pi)^4 \delta\left(\sum_{i=1}^{n+m} p_i\right) H^{(n+m)}(p_1, \dots, p_{m+n}).$$

(Unlike \mathcal{M} , this notation is *not* standard and we won't use it after the present discussion is finished.)

Both \mathcal{M} and $H^{(n+m)}$ should be thought of as functions on the subspace of $(\mathbb{R}^4)^{n+m}$ where the relevant sum of momenta vanishes. (Otherwise the preceding equations would not specify the values of the functions off of this subspace.) In particular, we'll always write the parameters to $H^{(2)}$ as p and $-p$.


Recall also that our Feynman rules included a factor of $i/(p_k^2 - m^2 + i\epsilon)$ for every external leg with momentum p_k . (An "external leg" is an edge connected to an external vertex.) At first glance, this seems like it should be the source of the pole in $\widetilde{G}^{(n+m)}$ and exactly cancel the $p_k^2 - \mu^2$ in the LSZ formula, but not quite: one involves the parameter m from the Lagrangian, and the other involves the particle mass μ .

Luckily, there *is* an object in our Feynman diagram story that has a pole at $p^2 = \mu^2$: we learned from our examination of the Källén–Lehmann spectral representation that

$$H^{(2)}(p, -p) = \frac{iZ}{p^2 - \mu^2 + i\epsilon} + (\text{continuous near } p^2 = \mu^2).$$

We can take advantage of this with the following trick.

Consider the following diagram:



Because of momentum conservation, the momentum of the edge with the red dotted line passing through it has to be equal to the external momentum p_1 . I encourage you to use to the Feynman rules to show that this diagram's contribution to $H^{(4)}$ formally splits up as the *product* of the two pieces on either side of the dotted line (where we leave out the $+i\epsilon$'s from the denominators of the propagators to save space):

$$\begin{aligned} & \frac{(-i\lambda)^2}{6} \int \frac{d^4 l_1}{(2\pi)^4} \int \frac{d^4 l_2}{(2\pi)^4} \frac{i}{p_1^2 - m^2} \frac{i}{l_1^2 - m^2} \frac{i}{l_2^2 - m^2} \frac{i}{(p_1 - l_1 - l_2)^2 - m^2} \frac{i}{p_1^2 - m^2} \\ & \cdot \left(\frac{i}{p_2^2 - m^2} \frac{i}{p_3^2 - m^2} \frac{i}{p_4^2 - m^2} \frac{(-i\lambda)^2}{2} \int \frac{d^4 k}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{(p_1 + p_2 - k)^2 - m^2} \right) \end{aligned}$$

We could have done this same thing with *any* two-external-leg diagram to the left of the dotted line, and the part in parentheses on the second line of this expression would be unchanged.

So, if we consider the sum of *all* diagrams which arise in this way, we simply get

$$H^{(2)}(p_1, -p_1) \cdot \left(\frac{i}{p_2^2 - m^2} \frac{i}{p_3^2 - m^2} \frac{i}{p_4^2 - m^2} \frac{(-i\lambda)^2}{2} \int \frac{d^4 k}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{(p_1 + p_2 - k)^2 - m^2} \right).$$

The second factor is just the value of the diagram to the right of the dotted line, except without the contribution from the external leg with momentum p_1 . (This contribution isn't present in the second factor because it's taken care of by the $H^{(2)}$ factor.)

If we do the same thing with the other three legs, we end up with a factor of $H^{(2)}(p_k, -p_k)$ for each leg, together with a factor that corresponds to the middle part of the diagram. The process of taking a diagram and removing the largest possible two-external-leg subdiagrams from each external leg is called **amputating** the diagram, and the “middle part” that is left over after performing the amputation is called a **proper diagram**.

Let's write $P(p_1, \dots, p_n; p_{n+1}, p_{n+m})$ for the sum of the values of all proper diagrams with n incoming and m outgoing external legs with the given momenta. (Like H , this notation is made up, so don't try to find it in other sources!) By convention, P includes neither the contributions from the external legs nor the delta function. The upshot of the preceding discussion is that

$$\begin{aligned} H^{(n+m)}(p_1, \dots, p_n, -p_{n+1}, -p_{n+m}) \\ = H^{(2)}(p_1, -p_1) \cdots H^{(2)}(p_{n+m}, -p_{n+m}) \cdot P(p_1, \dots, p_n; p_{n+1}, p_{n+m}). \end{aligned}$$

Because we don't include the contributions from the external legs in P , the pole is now entirely contained in the $H^{(2)}$ factors. If we then multiply by the factors of $p_k^2 - \mu^2$ in the LSZ formula, this will exactly cancel the pole in each $H^{(2)}$, and the term in $H^{(2)}$ that's continuous near $p^2 = \mu^2$ will vanish. I encourage you to verify that the final conclusion ends up being that

$$i\mathcal{M}(p_1, \dots, p_n; p_{n+1}, \dots, p_{n+m}) = Z^{(n+m)/2} P(p_1, \dots, p_n; p_{n+1}, p_{n+m}).$$

In other words, we can compute $i\mathcal{M}$ by summing over all proper diagrams — that is, all diagrams that can't be amputated nontrivially — where each diagram is assigned a value according to the following rules:

- Label each external leg with the corresponding incoming or outgoing momentum.
- For each internal vertex, include a factor of $-i\lambda$.
- Introduce new variables k_i for each internal edge, orienting them however you want, and include a factor of $\frac{i}{k^2 - m^2 + i\epsilon}$.
- Using momentum conservation at each vertex, eliminate as many k_i 's as possible, writing them as linear combinations of the other momenta, until the number of free momentum variables is equal to the number of loops in the diagram.
- For each remaining internal momentum variable k , introduce an integral $\int \frac{d^4 k}{(2\pi)^4}$.
- Take the limit $\lim_{\epsilon \rightarrow 0^+}$ and divide by the symmetry factor.

Here are two examples of proper diagrams with the values they're assigned by our new rules (again with the $+i\epsilon$'s omitted):

$$= \frac{(-i\lambda)^2}{2} \int \frac{d^4k}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{(p_1 + p_2 - k)^2 - m^2}$$

$$= \frac{(-i\lambda)^3}{2} \int \frac{d^4k}{(2\pi)^4} \int \frac{d^4l}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{l^2 - m^2} \frac{i}{(p_1 + p_2 - k)^2 - m^2} \frac{i}{(k + l - p_3)^2 - m^2}$$


2.2 One-Particle Irreducible (1PI) Diagrams

The other diagram trick we’re going to go over is about the two-point function itself. Recall yet again that we have

$$\tilde{G}^{(2)}(p, -q) = (2\pi)^4 \delta(p - q) \left[\frac{iZ}{p^2 - \mu^2 + i\epsilon} + (\text{continuous near } p^2 = \mu^2) \right].$$

At least conceptually, this gives us a way to compute Z and μ by summing over all diagrams with two external legs and looking at the reciprocal of the result. While this does work, there is a trick that simplifies this computation quite a bit.

Consider the following diagram:



This diagram can clearly be divided into three pieces, each separated from its neighbors by a single edge. The key observation is that, by momentum conservation, the momentum attached to each of these connecting edges (marked in red in the diagram) has to be the same as the momentum on each of the outer edges. This means that, just like with the proper diagrams we just examined, the value of this diagram splits up as a *product* consisting of one factor for each subdiagram and an $i/(p^2 - m^2 + i\epsilon)$ for each connecting edge.

If you split up any diagram in this way as far as possible, the pieces will be what physicists call **one-particle irreducible** or **1PI** diagrams, that is, diagrams that can’t be disconnected by deleting a single edge. (By convention, the diagram consisting of just a single edge is not 1PI.) This inspires the following definition. We’ll write $-i\Sigma(p^2)$ for the sum of the values of all 1PI diagrams with two external legs, where we don’t include the external propagators or the factor of $(2\pi)^4 \delta(p - q)$. Physicists refer to Σ as the **self-energy**. I encourage you to convince yourself that the preceding analysis implies that

$$\tilde{G}^{(2)}(p, -q) = (2\pi)^4 \delta(p - q) \left[\Delta + \Delta(-i\Sigma(p^2)\Delta) + \Delta(-i\Sigma(p^2)\Delta)^2 + \dots \right],$$

where $\Delta = i/(p^2 - m^2 + i\epsilon)$ is the propagator.

What's nice about this expression is that it takes the form of a geometric series: we can rewrite it as

$$\begin{aligned}\tilde{G}^{(2)}(p, -q) &= (2\pi)^4 \delta(p - q) \frac{i}{p^2 - m^2 + i\epsilon} \cdot \frac{1}{1 - \Sigma(p^2) \frac{1}{p^2 - m^2 + i\epsilon}} \\ &= (2\pi)^4 \delta(p - q) \frac{i}{p^2 - m^2 - \Sigma(p^2) + i\epsilon}.\end{aligned}$$

(Does this geometric series converge? It doesn't actually matter for our purposes: we will only ever be using this in perturbation theory, where we discard all terms which contain a sufficiently high power of λ . I sometimes find it helpful to adopt the perspective that all equations involving diagrams are equations of formal power series in λ rather than of series of numbers, which makes questions of convergence much less important. More on this at the end.)

Comparing this to our earlier expression gives us a way to relate Σ directly to Z and μ : equating the quantities multiplying the delta functions in our two expressions for $\tilde{G}^{(2)}$ tells us that

$$\frac{i}{p^2 - m^2 - \Sigma(p^2) + i\epsilon} = \frac{iZ}{p^2 - \mu^2 + i\epsilon} + (\text{continuous near } p^2 = \mu^2),$$

so if we take the reciprocal we learn that (discarding the $+i\epsilon$'s from the propagator)

$$-i(p^2 - m^2 - \Sigma(p^2)) = \frac{p^2 - \mu^2}{iZ + (p^2 - \mu^2)(\text{continuous near } p^2 = \mu^2)}.$$

By plugging in $p^2 = \mu^2$ to Σ and its derivative, we can get rid of the need to know anything about the mysterious second term in the denominator: we simply learn that

$$\begin{aligned}\Sigma(\mu^2) &= \mu^2 - m^2 \\ 1 - \Sigma'(\mu^2) &= \frac{1}{Z}.\end{aligned}$$

This gives us a procedure for relating μ and Z to diagrams that is quite a bit simpler than the one we started with! When we dig into the computation later in this article, this is what we'll use.

3 Wick Rotation and Regularization Schemes

The program we sketched above for extracting finite values from our Feynman diagrams had two steps: first we will write each divergent integral as a limit of finite quantities, and then we will express these finite quantities in terms of new, more physically meaningful parameters with the hope that these new expressions will converge to something finite.

Let's start by tackling the first of these steps. This part of the process is called **regularization**, and we actually have a considerable amount of freedom in how exactly to carry it out. Such a choice is called a **regularization scheme**, and we'll go through several such schemes in a moment to give you a sense of the diversity of options that are available here.

3.1 Wick Rotation


There is one feature that almost all regularization schemes share, though, so it's worth describing that first. All our integrals take place in \mathbb{R}^4 , and the integrands are all invariant under Lorentz

transformations. It would be considerably nicer to have integrands which are invariant under the action of $SO(4)$, that is, under ordinary Euclidean rotations in \mathbb{R}^4 .

This is mainly because, if we write k^2 for the Lorentz norm of $k \in \mathbb{R}^4$ and $|k|^2$ for the Euclidean norm, then $\{k : |k|^2 \leq r^2\}$ is compact while $\{k : k^2 \leq r^2\}$ is not. An integral of the form $\int_{\mathbb{R}^4} d^4k f(k)$ where f is a continuous $SO(4)$ -invariant function of k can be split into a one-dimensional integral over $|k|$ and a three-dimensional integral over the 3-sphere of radius $|k|$, and the fact that the 3-sphere is compact means that that part of the integral will always be finite.

We can transform our Lorentz-invariant integrands into $SO(4)$ -invariant ones by means of a clever trick called **Wick rotation**. The idea is fairly simple. Writing k_0 for the time component of some momentum variable k , we take advantage of the fact that the integrand is holomorphic to rotate the contour we're integrating k_0 along from the real axis to the imaginary axis.

Specifically, because the propagators all look like $i/(k^2 - m^2 + i\epsilon)$ and ϵ is going to zero from the positive direction, they have one pole in k_0 just below the positive real axis and one just above the negative real axis. The values being plugged in for k are all homogeneous linear functions of the momenta, so if we rotate all the k_0 's counterclockwise at the same time, we'll avoid the poles. You could formalize this using the fact that the integral along the following contour is zero, together with an argument that the integrand decays fast enough to make the two arc-shaped pieces on the outside go to zero as the radius goes to infinity:



This counterclockwise rotation amounts to introducing, for each momentum variable k , a new variable k^E where

$$(k_0^E, k_1^E, k_2^E, k_3^E) = (-ik_0, k_1, k_2, k_3),$$

and integrating over all $k^E \in \mathbb{R}^4$. (You'll also get an extra factor of i from the fact that $d^4k = i d^4k^E$.) The propagator becomes

$$\frac{i}{k_0^2 - k_1^2 - k_2^2 - k_3^2 - m^2 + i\epsilon} = \frac{i}{-(k_0^E)^2 - (k_1^E)^2 - (k_2^E)^2 - (k_3^E)^2 - m^2 + i\epsilon} = \frac{-i}{|k^E|^2 + m^2 - i\epsilon},$$

and since the denominator now doesn't vanish anywhere near the contour of integration we are free to just let ϵ go to zero and forget about it. We'll see an example of all this being carried out momentarily.

The propagators will include external as well as internal momenta, so in order for this program to work we'll have to Wick-rotate the external momenta as well as the internal momenta.

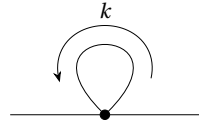
This means that, after we compute the value of the integral, we'll have to analytically continue it back to the original values of the external momentum variables with real rather than imaginary time components. In practice, this usually just amounts to replacing every occurrence of $|p^E|^2$ with $-p^2$, and occasionally using the fact that we are rotating p_0 clockwise to pick the right branch of a log or an n 'th root. For more details, I recommend the discussion that starts on p. 195 of Folland's book mentioned in the introduction.

3.2 Regularizing with a Hard Cutoff

Probably the simplest regularization scheme is to introduce a so-called **hard cutoff**. This means, after Wick-rotating, that we restrict each of our momentum integrals to a ball in \mathbb{R}^4 of radius Λ . This leaves us with an integral of a continuous function on a compact set, so the result is finite. The resulting function of Λ will then blow up as Λ goes to infinity. Physicists refer to this as an **ultraviolet divergence**; Λ is referred to as an **ultraviolet cutoff**, and taking Λ to infinity is called "taking the **ultraviolet limit**." The name comes from the fact that large momenta correspond to high frequencies, and ultraviolet light is at the high-frequency end of the visible spectrum.

While this doesn't happen for ϕ^4 theory, it's also possible for Feynman integrals to blow up for *small* momenta, resulting in what physicists call an **infrared divergence**. (This tends to happen in theories with massless particles, and if we discuss such theories later in this series we will talk about why.) If your theory has infrared divergences and you're doing hard-cutoff regularization, you will also have to introduce a *lower* bound on the absolute value of your momenta, that is, introduce an **infrared cutoff** ϵ and restrict the integrals to the set $\{k : \epsilon < |k| < \Lambda\}$. Since ϕ^4 theory doesn't have any infrared divergences, we'll only worry about ultraviolet cutoffs from now on.

As a quick example, let's use a hard cutoff to compute a regularized value for this diagram, sometimes called the "tadpole":



$$= \frac{-i\lambda}{2} \int \frac{d^4k}{(2\pi)^4} \frac{i}{k^2 - m^2 + i\epsilon}$$

Our first step is to Wick-rotate. As explained above, this shakes out to replacing k^2 with $-|k^E|^2$ and d^4k with id^4k^E . All together, our integral becomes

$$\frac{-i\lambda}{2} \int \frac{d^4k^E}{(2\pi)^4} \frac{1}{|k^E|^2 + m^2}.$$

Now that the integrand is spherically symmetric, we can introduce the cutoff. If we restrict the integral to the ball of radius Λ and switch to spherical coordinates, I encourage you to verify that the result is

$$-i\lambda\pi^2 \int_0^\Lambda dr \frac{r^3}{r^2 + m^2} = \frac{-i\lambda\pi^2}{2} \left(\Lambda^2 + m^2 \log \frac{m^2}{m^2 + \Lambda^2} \right).$$

(Note that $2\pi^2$ is the surface area of the unit 3-sphere in \mathbb{R}^4 .) Because the value of this diagram doesn't happen to depend on the external momenta, we can skip the final step of analytically continuing the external momenta back to values with real time coordinates.

This of course blows up as Λ goes to infinity, which shouldn't be surprising since the original integral was divergent. Remember, though, that our goal is not to make the value of each individual diagram finite, but to make the scattering amplitudes finite *when written as functions of more physically meaningful parameters*. As we'll see when we carry out the program in more detail below, this creates opportunities for the divergent parts of different diagrams to cancel, leaving a quantity that approaches a finite limit as the cutoff is removed.

The hard cutoff is probably the simplest regularization scheme to think about, but it has a pretty severe disadvantage from a computational perspective: once there are at least two momentum variables in the integral, it's basically impossible to write the result in closed form as a function of Λ . Because of this, for all but the simplest diagrams, hard-cutoff regularization is basically never used to perform actual computations.

Other regularization schemes naturally involve other tradeoffs. We'll go through a couple other possibilities to give you a sense of the options before settling on the scheme we're ultimately going to use.

3.3 Pauli–Villars Regularization

Pauli–Villars regularization is a regularization scheme where we replace the propagator with a different rational function. We start by Wick-rotating, which replaces each propagator (up to factors of i we won't bother keeping track of here) with $1/(|k|^2 + m^2)$. Then, for some very large Λ , we replace this propagator with

$$\frac{1}{|k|^2 + m^2} - \frac{1}{|k|^2 + \Lambda^2} = \frac{\Lambda^2 - m^2}{(|k|^2 + m^2)(|k|^2 + \Lambda^2)}.$$

When $|k| \ll \Lambda$, this expression is very close to the original propagator, but it decays like $|k|^4$ rather than $|k|^2$, so it stands a better chance of producing a convergent integral.

Because the Pauli–Villars expression for the propagator is very close to the original one for small momenta, it can be useful to think of this as a variant of the hard-cutoff scheme we described earlier where the contributions of large momenta are being suppressed gradually rather than all at once. Unlike the hard-cutoff scheme, it's possible in the Pauli–Villars scheme to find closed forms for the regularized values of all the diagrams we're about to consider.

But it's still somewhat cumbersome from a computational perspective. One unfortunate feature is that performing this transformation of the propagator might not be enough to make every diagram converge on its own, and it may be necessary to introduce a second Pauli–Villars parameter Λ_2 and transform the propagator again. Because of this (along with some other subtleties involving gauge theories that won't come up for us here) it's become less popular than the dimensional regularization scheme we describe below.

3.4 Lattice Regularization

Lattice regularization involves replacing space with the lattice $(\epsilon\mathbb{Z})^3$ for some small distance ϵ and then taking the limit as ϵ goes to 0. This ϵ turns out to serve as an ultraviolet cutoff — in general, a great rule of thumb is that problems at short distances correspond to problems at high momenta, and vice versa. Our field operators are then thought of as operator-valued functions on this lattice — honest functions now, not distributions — which can introduce some ϵ 's into the formulas we derived in the last two sections. If we need an infrared cutoff, we can additionally restrict to a large box with side length L , where L is an integer multiple of ϵ .

Back in the first article, we in fact needed to restrict our fields to a finite-volume box in exactly this way to make sense of the *free* theory. We ended up having to subtract a quantity from the Hamiltonian which was finite when the cutoff was present but which blew up when it was removed. This is a particularly simple example of renormalization using an infrared cutoff, one which only required messing with the constant term in the Lagrangian.

One attractive feature of lattice regularization from a conceptual perspective, especially when you use both the ultraviolet and infrared cutoffs, is that you end up with a quantum system with finitely many degrees of freedom. This means, at least theoretically, everything takes place in an ordinary Hilbert space, exactly like in ordinary quantum mechanics, with none of the ugly infinities that have plagued our story so far. Because of this, even when using another regularization scheme, I sometimes find it helpful to picture lattice regularization whenever I get confused about the physical meaning of some mathematical step in the renormalization story.

Unfortunately, just like in the case of a hard cutoff, it's not really feasible to use lattice regularization to do much computationally, and it has the additional unpleasant feature of breaking the Lorentz symmetry (or, after Wick-rotating, the $SO(4)$ symmetry). However, a variant of this idea where time is also discretized ends up being very useful for doing *numerical* computations. It's an especially important tool for theories like quantum chromodynamics where the perturbation-theoretic perspective we're emphasizing in this series ends up not producing very many usable results.

3.5 Dimensional Regularization

Dimensional regularization is the scheme we're actually going to use for our renormalization computations. The idea is to take advantage of the fact that, after Wick-rotating, all our integrals are spherically symmetric, and this ends up meaning that we can generalize the expression to a number of dimensions d other than 4. The resulting expression in fact ends up being an *meromorphic* function of d , and this function of d serves as the regularized value of the diagram.

It shouldn't be especially obvious that this procedure is well-defined. This is not the place to go into a ton of detail: it would increase the length of the article substantially, and it's already explained well in other sources. I found the discussion in Section 7.3 of Folland's book to be nicely written and amenable to a mathematician's sensibilities, and there's also a careful exposition in Section 4.1 of *Renormalization: An Introduction to Renormalization, the Renormalization Group, and the Operator-Product Expansion* by John Collins. To get a deep understanding of what's going on, you'll probably need to consult one of these other sources, but this shouldn't be necessary to follow the rest of this article.

But we can maybe make it a bit clearer in the context of an example, so let's look at the "tadpole" diagram we examined earlier when discussing the hard cutoff. After Wick-rotating, we wrote the value of the diagram in the form

$$\frac{-i\lambda}{2} \int \frac{d^4 k^E}{(2\pi)^4} \frac{1}{|k^E|^2 + m^2}.$$

Now, let's consider the same expression, but as an integral over \mathbb{R}^d for arbitrary d :

$$\frac{-i\lambda}{2} \int \frac{d^d k^E}{(2\pi)^d} \frac{1}{|k^E|^2 + m^2}.$$

Regardless of the value of d , this integral is spherically symmetric, so we can switch to spherical coordinates and reduce it to an integral over just the radial coordinate. We get

$$\frac{-i\lambda}{2} \frac{A_d}{(2\pi)^d} \int_0^\infty \frac{r^{d-1}}{r^2 + m^2},$$

where $A_d = 2\pi^{d/2}/\Gamma(d/2)$ is the area of the unit $(d-1)$ -sphere in \mathbb{R}^d .

This last expression gives us the meromorphic function of d mentioned earlier. Performing the integral gives the value

$$\frac{-i\lambda}{2} \frac{2\pi^{d/2}}{(2\pi)^d \Gamma(d/2)} \frac{m^{d-2}}{2} \Gamma(d/2) \Gamma(1-d/2) = \frac{-i\lambda}{2} \frac{m^{d-2}}{(4\pi)^{d/2}} \Gamma(1-d/2).$$

The fact that our original four-dimensional integral diverged is reflected in the fact that this function has a *pole* at $d = 4$ arising from the factor of $\Gamma(1-d/2)$: near $d = 4$, it looks like

$$\frac{-i\lambda}{2} \frac{m^2}{16\pi^2} \left(\frac{2}{d-4} + \log \frac{m^2}{4\pi} + \gamma - 1 + O(d-4) \right),$$

where γ is the Euler–Mascheroni constant. As we’ll see when we dig into explicit computations in the next section, the goal is to find a way for this pole to cancel with a pole from another diagram when we rewrite everything in terms of our more physically meaningful parameters.

For diagrams that are more complicated than this one, it can take a bit more trickery to write the integral in a form that we can apply this scheme to, but it’s always possible. The most important such trick, which we won’t go into, involves what are called *Feynman parameters*; if you’re interested in learning how that works, I encourage you to check out Collins or Folland. In this article, I will restrict myself to just telling you what the final expressions end up being.

While dimensional regularization ends up being pretty nice to use computationally, that niceness comes at the cost of being quite a bit harder to interpret physically: it’s far from clear what to make of the values of our integrals at non-integer values of d , and it’s hard to tell a story in which taking d to 4 is somehow like lifting an ultraviolet cutoff. While this lack of concreteness is definitely a count against it, it is so much nicer computationally than the alternatives — especially in more complicated theories where it becomes desirable to have a regularization scheme that respects the theory’s symmetries — that the tradeoff ends up being worth it.

4 Perurbative Renormalization

With our choice of regularization scheme in hand, let’s start renormalizing ϕ^4 theory. I think it’s best at this point to actually get our hands dirty with some integrals; while we’ll have a lot to say on a more abstract level soon, those theoretical considerations will almost certainly make more sense after you’ve seen an example of the type of computation they pertain to.

To illustrate the process, we’ll compute scattering amplitudes for two incoming and two outgoing particles up to order λ^2 . This means looking at the 4-point function $\tilde{G}^{(4)}$, but since we also need to relate the parameters in the Lagrangian to μ and Z , we’ll also need to examine $\tilde{G}^{(2)}$.

As we discussed earlier, we want to express all our scattering amplitudes in terms of physically measurable parameters, rather than the unmeasurable parameters m and λ that appear in the Lagrangian, so we need to pick some parameters to use for this. The physical particle mass μ is a natural replacement for m , but we need a replacement for λ as well.

Since λ ends up attached to vertices in the Feynman diagram of degree 4, we have that, to first order, $i\mathcal{M}(p_1, p_2; p_3, p_4) = -iZ^2\lambda \bmod \lambda^2$. This suggests a way to pick a physically measurable quantity λ_r to use instead of λ : we can use the value of a four-particle scattering amplitude for some specific choice of incoming and outgoing momenta. A popular choice is to take $q = (\mu, 0, 0, 0)$ and declare that

$$-iZ^2\lambda_r = i\mathcal{M}(q, q; q, q).$$

Physically, you can think of the right-hand side as the limit of the scattering amplitude as the velocities of the incoming and outgoing particles approach 0. We'll call λ_r the **renormalized coupling constant**.

We have three equations that constrain the values of m , λ , and Z : the equation defining λ_r that we just discussed, and the two equations relating μ and Z to the self-energy $\Sigma(p^2)$ that we discussed earlier. Since there are only finitely many Feynman diagrams that contribute to each term in the power series for our scattering amplitudes, it would hypothetically be possible to use these three constraints to explicitly solve for m , λ , and Z in terms of μ and λ_r , and then plug the resulting formal power series into whatever scattering amplitude we're interested in computing. The hope would then be that, as we remove whatever regularization scheme we imposed at the beginning of this process, the resulting functions of μ and λ_r converge to something nice and finite.

4.1 Rewriting the Lagrangian

While this method would work, there is a trick (sometimes called **renormalized perturbation theory**) that makes it considerably nicer to perform the computation. It involves rewriting the Lagrangian in a form that involves μ and λ_r directly, so that the values of the Feynman diagrams can more easily be expressed in terms of the variables we ultimately care about. Let's see how this works.

First, recall that the value of Z depended on our choice of interpolating field ϕ : we had $Z^{1/2} = \langle p|\phi(0)|\Omega\rangle$, where $|p\rangle$ was a one-particle state. So if we define $\phi_r = Z^{-1/2}\phi$ and use this as our interpolating field instead, then its Z will just be 1, which will mean we don't have to keep track of the Z 's in our formulas for scattering amplitudes anymore.

This of course comes at the cost of *introducing* Z 's into the Lagrangian: we get

$$\begin{aligned}\mathcal{L} &= \frac{1}{2}((\partial\phi)^2 - m^2\phi^2) - \frac{\lambda}{4!}\phi^4 \\ &= \frac{Z}{2}((\partial\phi_r)^2 - m^2\phi_r^2) - \frac{Z^2\lambda}{4!}\phi_r^4 \\ &= \frac{1}{2}((\partial\phi_r)^2 - Zm^2\phi_r^2) - \frac{Z^2\lambda}{4!}\phi_r^4 + \frac{\delta Z}{2}(\partial\phi_r)^2,\end{aligned}$$

where $\delta Z = Z - 1$. We'll do a somewhat similar transformation to get μ and λ_r to show up as coefficients in the Lagrangian as well: if we set $\delta m^2 = Zm^2 - \mu^2$ and $\delta\lambda = Z^2\lambda - \lambda_r$, then we can write

$$\mathcal{L} = \frac{1}{2}((\partial\phi_r)^2 - \mu^2\phi_r^2) - \frac{\lambda_r}{4!}\phi_r^4 + \frac{\delta Z}{2}(\partial\phi_r)^2 - \frac{\delta m^2}{2}\phi_r^2 - \frac{\delta\lambda}{4!}\phi_r^4.$$


The final three terms are called **counterterms**, and the role they serve in the computation will become more apparent when we dig into an explicit example in just a moment. It is worth emphasizing right now — especially because some more old-fashioned sources can be confusing

on this point — that this is the *same* Lagrangian that we started with, just expressed in terms of a different set of quantities than the ones we started with. We are not adding new terms to the Lagrangian, just splitting apart the terms that were already there.


Recall that we extracted our Feynman rules in the previous article by dividing the Lagrangian (after converting it to a Hamiltonian) into a “free part” and an “interacting part.” Now that we have written the Lagrangian in this new form, we will perform this split differently, taking $\frac{1}{2}((\partial\phi_r)^2 - \mu^2\phi_r^2)$ as the free part and the other four terms as the interacting part. This will naturally result in a different set of Feynman rules. If you are interested, it’s a nice exercise to trace through the original derivation of the Feynman rules and see how it changes; here I’ll simply state what the new rules are.

The new terms correspond to new types of vertices that can be present in a diagram. In addition to our familiar degree-4 vertex, which contributes a factor of $-i\lambda_r$ whenever it appears, we get a second type degree-4 vertex carrying a factor of $-i\delta\lambda$. There will also be a new degree-2 vertex corresponding to the δZ and δm^2 terms. (We could introduce two separate types of vertices here, but it’s equivalent — and simpler — to consolidate them.) This vertex will contribute a factor of $-i(\delta m^2 - p^2\delta Z)$, where p is the momentum of either of the two edges attached to it, which have to be equal by momentum conservation. The factor of p^2 arises from the fact that the corresponding term in the Lagrangian contains a derivative of the field operator.

We’ll draw the new vertices like this:



Here are a couple examples of diagrams with these new types of vertices, which might make the rules a bit clearer:



4.2 Renormalizing to First Order

Now that we’ve rewritten our Lagrangian in this new form, we need a way to nail down the values of $\delta\lambda$, δm^2 , and δZ . Like we mentioned earlier, we have three equations at hand to determine the values of these three parameters: the defining equation of λ_r we just discussed, and the two

equations relating the self-energy $\Sigma(p^2)$ to μ and Z from our earlier discussion of 1PI diagrams, which were $\Sigma(\mu^2) = \mu^2 - m^2$ and $1 - \Sigma'(\mu^2) = 1/Z$.

After the reparameterization we just performed, our Z is 1 and the coefficient on the mass term in our Lagrangian is equal to μ^2 , and these two facts combine to make these equations look particularly nice. I encourage you to verify that they end up in the form

$$\begin{aligned} i\mathcal{M}(q, q; q, q) &= -i\lambda_r \\ \Sigma(\mu^2) &= 0 \\ \Sigma'(\mu^2) &= 0. \end{aligned}$$

These are called **renormalization conditions**. At any given order in perturbation theory, there are only finitely many diagrams contributing to the left-hand sides of each of these equations. As long as we can explicitly compute a regularized value for every diagram that shows up in this way, we can use our renormalization conditions to solve for $\delta\lambda$, δm^2 , and δZ to any order in perturbation theory, and we can then use these values to compute whatever scattering amplitudes we're interested in. This procedure is called **perturbative renormalization**.

As a warm-up, let's start by working everything out to first order. The only proper diagrams that could contribute terms of order λ_r^1 to the four-point function are these two:

The renormalization condition specifying λ_r tells us that, when we set all the momenta to q , the sum of these diagrams should be $-i\lambda_r$, i.e.,

$$-i\lambda_r - i\delta\lambda = -i\lambda_r \text{ mod } \lambda_r^2.$$

This just means that $\delta\lambda$ is zero to first order in λ_r , which shouldn't be that surprising: we essentially defined λ_r so that this would happen.

Something slightly more interesting happens with the two-point function. The two 1PI diagrams that contribute to the self-energy are:

(We could have included another diagram like the one on top but with a degree-4 counterterm vertex on the bottom, but that diagram would contain a factor of $\delta\lambda$, which we just learned is zero to first order!)

The diagram on top is the “tadpole” that we looked at in the last section. Our other two renormalization conditions say that both the sum of these two diagrams and the derivative of that sum should vanish after setting $p^2 = \mu^2$. This works out to

$$0 = \Sigma(\mu^2) = -i \left(\frac{\lambda_r}{2} \frac{\mu^{d-2}}{(4\pi)^{d/2}} \Gamma(1 - d/2) + \delta m^2 + \mu^2 \delta Z \right) \text{ mod } \lambda_r^2$$

$$0 = \Sigma'(\mu^2) = i \delta Z \text{ mod } \lambda_r^2.$$


From this, we conclude that δZ vanishes to first order, but that

$$\delta m^2 = \frac{-\lambda_r \mu^{d-2} \Gamma(1 - d/2)}{2(4\pi)^{d/2}} \text{ mod } \lambda_r^2.$$

The only interesting thing about this expression for us will be the fact that it has a pole at $d = 4$ coming from the gamma function, which exactly cancels the corresponding pole in the tadpole. (In fact, because the value of the tadpole happens not to depend on p^2 at all, it cancels the *entire value*, but that won't generalize past this simple example, as we'll see momentarily.) Because this pole ends up cancelling, the *sum* of the values of our two diagrams can be finite when we send d to 4 even though each diagram *individually* isn't.

4.3 Renormalizing to Second Order

To second order, the proper diagrams contributing to the four-point function are:



We won't go through the computation of the dimensionally regularized values of the last three diagrams here; if you're interested, you can find this computation in Section 7.4 of Folland or Section 10.2 of Peskin and Schroeder. Instead, we'll just state the result. To do this, it will be convenient to introduce the **Mandelstam variables**

$$s = (p_1 + p_2)^2, \quad t = (p_1 - p_3)^2, \quad u = (p_1 - p_4)^2.$$

The value of the middle diagram turns out to be $-i(-i\lambda_r)^2 V(s)$, where

$$V(a) = \frac{\Gamma(2 - d/2)}{2(4\pi)^{d/2}} \int_0^1 \frac{dx}{(\mu^2 - ax(1-x))^{2-d/2}},$$

and the other two are the same except with t and u in place of s .

Our renormalization condition then tells us that

$$-i\lambda_r = -i(\lambda_r + \delta\lambda + (-i\lambda_r)^2(V(s) + V(t) + V(u))) \text{ mod } \lambda_r^3$$

when we plug in $q = (\mu, 0, 0, 0)$ for all four momenta. This makes the Mandelstam variables equal to $4\mu^2$, 0, and 0 respectively, so we learn that

$$\delta\lambda = \lambda_r^2(V(4\mu^2) + 2V(0)) \text{ mod } \lambda_r^3.$$

For our present purposes, the most important feature of $V(a)$ is its behavior as d approaches 4. If we define

$$W(a) = \int_0^1 dx \log(\mu^2 - ax(1-x)),$$

then near $d = 4$ we have

$$V(a) = \frac{1}{32\pi^2} \left(\frac{-2}{d-4} - \gamma + \log(4\pi) - W(a) \right) + O(d-4).$$

(The γ here is again the Euler–Mascheroni constant.)


This implies that our three new diagrams all take values that diverge as d approaches 4 due to that $-2/(d-4)$ term. But, because this divergent term doesn't depend on the external momenta, it will cancel when we add on the contribution from the counterterm (that is, from the second diagram). All together, we get that up to second order

$$\begin{aligned} i\mathcal{M}(p_1, p_2; p_3, p_4) &= -i(\lambda_r + \delta\lambda - \lambda_r^2(V(s) + V(t) + V(u))) \\ &= -i \left(\lambda_r + \frac{\lambda_r^2}{32\pi^2} (W(s) + W(t) + W(u) - W(4\mu^2) - 2W(0)) \right) + O(d-4). \end{aligned}$$

At this point, everything is nice and finite, so we are free to let d go to 4, which of course eliminates the $O(d-4)$ stuff at the end.


This means that we have done the thing we set out to do at the very beginning: we've shown that, if we write our scattering amplitude as a function of λ_r and μ and hold these two quantities constant as we let d approach 4, then (at least to second order) the scattering amplitudes converge to something finite. Although we discussed this briefly at the outset, it's worth emphasizing here that in this setup our “unrenormalized” parameters λ , m , and Z definitely *do* depend on d , and in fact they will blow up as d goes to 4. This definitely has some implications for how we interpret the Lagrangian we started with, and we'll have a lot more to say about this in the final section of this article.

We won't go into as much detail about the other two counterterms, δm^2 and δZ . If you were interested in computing them to second order, the relevant 1PI diagrams would be:



As we saw in the first-order computation, our renormalization conditions would force δm^2 to swallow up any divergences in the two-point function that are constant as functions of p^2 , and δZ to swallow up any divergences that are proportional to p^2 ; unlike in the first-order case, here there actually is such a divergence for δZ to take care of, arising from the final diagram in the list. If you're interested in digging into this computation in detail, you can look at Sections 4.4 and 4.5 of Pierre Ramond's *Field Theory: A Modern Primer*.

Our four-particle scattering amplitude computation didn't actually require finding the values of δm^2 and δZ because there don't happen to be any proper four-particle diagrams containing the degree-2 counterterm vertex that contribute terms up to second order. Such diagrams would start to show up at third order, though. Here's an example:



Notice, though, that the values of δm^2 and δZ are completely determined by looking at the two-point function. So, once that computation has been done, the value of *this* diagram is nailed down. That means we just have to hope that its value cancels with some other divergent values from some other diagrams to leave us something that converges as d goes to 4. It's not especially obvious that this should always work out! That's the question we'll turn to next.

5 Does This Always Work?

Now that we have that first computation under our belt, it's worth saying some words about why this procedure worked and to what extent we should expect it to generalize. Unlike a lot of what we've done in this series up to this point, this corner of the quantum field theory story is one in which it's actually possible to prove some rigorous theorems. We won't go through any of those proofs here, but my hope is that this quick overview will give you a sense of what can be established and provide a foundation if you choose to explore the literature more thoroughly.

5.1 Power Counting and Weinberg's Theorem


After we rewrote our Lagrangian in terms of our more physical variables μ and λ_r , we ended up with three counterterms: δm^2 , δZ , and $\delta\lambda$. We saw that each of these counterterms ended up with a value which diverges as d approaches 4, but, because the renormalization conditions forced the sum of *all* the relevant diagrams to be finite, this divergent value had to cancel the divergences coming from the diagrams without counterterm vertices.

Specifically, because the degree-2 counterterm vertex contributes $-i(\delta m^2 + p^2 \delta Z)$ to its diagram, we should expect δm^2 to take care of divergent values from diagrams with two external legs that are constant as functions of p^2 , and δZ to take care of those that take the form of a constant times p^2 , where p is the external momentum. And, since the degree-4 counterterm vertex contributes $-i\delta\lambda$, we should expect $\delta\lambda$ to take care of divergent values arising from diagrams with *four* external legs that are constant as functions of the external momenta.

In light of this, a good first question to ask is if we can tell whether the value of a diagram diverges just from looking at the diagram — that way, we can try to determine whether all the divergences that arise are of the type that can be cancelled by our counterterms.

Suppose we have a diagram with v internal vertices, e internal edges, and $l = 1 - v + e$ loops. (For the moment, we'll restrict to the case where the diagram has no counterterm vertices.) The value of the diagram will be an integral over l four-dimensional variables of a rational function of degree $-2e$, since each edge contributes a propagator of degree -2 . In order for the integral to converge, therefore, we at least need $4l - 2e < 0$. (The exception to this is if $l = 0$, in which case there is no integral, so the value of the diagram is always finite.)

This procedure is called **power counting**, and the quantity $4l - 2e$ is called the **superficial degree of divergence** of the diagram, often written as D . The preceding discussion implies that if $D \geq 0$, then the integral definitely diverges. The converse is not true, of course: the integral $\iint (1 + x^2)^{-12} dx dy$ obviously diverges, since the integrand doesn't depend on y , even though the degree of the integrand is much less than 2. A similar phenomenon can happen when a Feynman diagram contains a divergent *subdiagram*, as in this example:



Here we have $e = 5$ and $l = 2$, which gives $D = 4l - 2e = -2$, but the value of this diagram clearly diverges, since the loop on the top will contribute a factor of $\int \frac{d^4 k}{(2\pi)^4} \frac{i}{k^2 - m^2 + i\epsilon}$, which has no chance of converging.

In this example, while the diagram as a whole has a negative superficial degree of divergence, it contains a subdiagram which doesn't. (Formally, the "subdiagrams" are the diagrams that arise as connected components after deleting some number of edges from the original diagram.) Happily, it turns out that this is the *only* way our naive criterion can fail: a result called **Weinberg's Theorem** states that the integral arising from a Feynman diagram converges if and only if the superficial degrees of divergence of it and all its connected subdiagrams are negative. Talagrand has a nice discussion of Weinberg's Theorem in Section 15.2 if you're interested in learning more.

For ϕ^4 theory, we can turn this into a fairly simple criterion. Since each external leg is connected to one internal vertex, edge internal edge is connected to two, and each internal vertex has degree 4, we have $4v = 2e + n$, where n is the number of external legs. We can therefore compute that


$$D = 4l - 2e = 4(1 - v + e) - 2e = 4 - n.$$

The only superficially divergent diagrams are therefore the ones with 2 or 4 external legs. (It's a simple graph theory exercise to show that it's impossible for a diagram to have 1 or 3 external legs.)

5.2 Some Words about Divergent Subdiagrams

This is a hopeful conclusion: the only superficially divergent diagrams in our theory are exactly the ones that our counterterms are equipped to handle. (If we had found a superficially divergent diagram with *six* external legs, for example, we might have been in trouble: we don't have a counterterm vertex with degree 6, so it's hard to see how that divergence could ever get cancelled.) But we should say a bit about the other case that Weinberg's Theorem leaves us with: why should we expect our procedure to handle superficially convergent diagrams with divergent subdiagrams?

In the computation we did earlier, we saw examples where a divergence was cancelled by a diagram consisting of just a counterterm vertex connected directly to the external legs. In fact, essentially the same thing happens whenever a diagram contains a divergent subdiagram. For example, consider the following two diagrams:



We showed in our first-order computation that the divergence in the “tadpole” diagram wound up being cancelled to first order in λ_r by the δm^2 counterterm. You can use this fact along with the Feynman rules to show that the same cancellation happens to *fourth* order in λ_r between the two diagrams pictured here. (You shouldn’t have to evaluate either integral to do this; just write out the integrals corresponding to each one and add them, grouping together as many terms as possible. There’s really no substitute for working out this computation out yourself, so I strongly encourage you to take some time and do it!)

The same thing will in fact happen in general: if the divergence in some diagram is cancelled by a counterterm vertex, then the same cancellation will occur whenever that diagram appears as a subdiagram of a superficially convergent diagram. Because of this, what we’d want to show is that, as long as the shapes of the superficially divergent diagrams match the counterterm vertices as they do here, all the divergences arising from all divergent Feynman diagrams end up getting absorbed by the counterterms after we impose our renormalization conditions.

As you might imagine, the combinatorics involved in proving this get quite intricate. For example, what if there are two divergent subdiagrams which partially overlap, or if one divergent subdiagram contains another? While we won’t prove it here, it’s perhaps comforting to know that there is actually a theorem to the effect that this all works out.

It’s called the **BPHZ Theorem**, and its proof works by giving a procedure for splitting up the sum of the values of the diagrams, both with and without counterterms, into a sum of integrals, each of which separately converges as the cutoff is removed. The procedure requires splitting up each counterterm as a sum in which each term corresponds to a nested sequence of divergent subdiagrams. If you’re interested in the details of how all this works, the best source I’ve encountered by far is Part IV of Michel Talagrand’s *What Is a Quantum Field Theory*.

(There is one other question here worth touching on: we may now know that the superficially divergent diagrams all have the right number of external legs to be cancelled by our counterterms, but how do we know that they have the right dependence on the external momenta? This issue is also taken care of by the BPHZ story, but the following loose heuristic might be helpful. Because of the form the propagator takes, if you take the derivative of a Feynman integral with respect to one of the external momenta, the superficial degree of divergence will decrease by 1. I encourage you to turn this into a quick argument that we should expect diagrams with four external legs to contribute divergences that are at worst constant in the external momenta, and that we should expect diagrams with two external legs to contribute divergences that are at worst proportional to p^2 .)

5.3 Renormalizability


The upshot of the preceding discussion is that it is possible in ϕ^4 theory, at each order in perturbation theory, to use our renormalization conditions to assign values to the counterterms

that make all the divergences in all the scattering amplitudes go away. Because of this, we say that ϕ^4 theory is “renormalizable.” It is possible in other quantum field theories for things to not work out so nicely.

Suppose, for example, that we had instead chosen to look at ϕ^6 theory, the theory whose Lagrangian is

$$\mathcal{L} = \frac{1}{2}((\partial\phi)^2 - m^2\phi^2) - \frac{\lambda}{6!}\phi^6.$$

If we tried to repeat our power counting analysis here, we would run into a big problem: there are superficially divergent diagrams with *arbitrarily many* external legs. For example, all of these diagrams have superficial degree of divergence 4:



Applying the procedure we used earlier would give us a degree-6 counterterm vertex which could cancel the divergence from the first of these, but it would be no help for the rest. We could imagine modifying the Lagrangian to add terms with higher powers of ϕ to take care of them, but we’d need to end up with a counterterm of the form $\delta g_n \phi^{2n}$ for every $n \geq 4$, and each counterterm would need to have its value nailed down by a corresponding renormalization condition. In other words, extracting any predictions from our theory would require knowing the values of infinitely many parameters!

All in all, for any given theory, this story can end in one of three ways:

- If there are diagrams with arbitrarily high superficial degrees of divergence, the theory is said to be **nonrenormalizable**. This is the case for the ϕ^6 theory we just discussed, and it means that it’s not possible to remove all the divergences using only finitely many renormalization conditions.
- If there are infinitely many superficially divergent diagrams but the possible degrees of divergence are bounded above, the theory is said to be **renormalizable**. In this case, as the name suggests, it is possible to use only finitely many renormalization conditions to assign finite values to all scattering amplitudes to every degree in perturbation theory. (In the case of ϕ^4 theory, we needed three.)
- An even better case is if there are only finitely many superficially divergent diagrams. In this case the theory is called **superrenormalizable**.

As we saw during our analysis of ϕ^4 theory, you can always determine whether a given theory is renormalizable by looking at its Lagrangian and doing a little graph theory. As a quick exercise, I encourage you to show that if you modify our theory by adding terms to the Lagrangian which are polynomials in ϕ , then it will only be renormalizable if every term you add has degree at most 4. For more details, I recommend the discussion in Section 7.2 of Folland’s book.

6 Taking Stock

We started this whole discussion with the goal of extracting usable numbers from ϕ^4 theory, and with the perturbative renormalization procedure in hand we have essentially accomplished this. There are a few issues it's worth reflecting on at this stage.

6.1 The Meaning of the Perturbation Series

We can use our perturbative renormalization procedure to compute scattering amplitudes as functions of μ and λ_r to any desired degree in perturbation theory. I find it most straightforward to think of the output of this procedure as a formal power series in λ_r ; because our theory is renormalizable, the BPHZ theorem guarantees that each coefficient in this power series converges to something finite as the cutoff is removed.

Our theory is a toy example that doesn't really correspond to anything physical, but if it did, we could use this to compare its results to experiment. We'd have to measure the mass μ of our particle and perform a low-momentum scattering experiment to determine the value of λ_r ; with those quantities in hand we'd be able to predict the value of any other scattering amplitude to any desired degree in perturbation theory and compare it with a scattering experiment. Indeed, computations very much like this have been performed for quantum electrodynamics, resulting in some of the most precise agreements between theory and experiment anywhere in science.

In the last article, we talked a bit about the question of whether the perturbation series converges. The conclusion was that, while as far as I know there aren't any theorems in this area, no one really expects that it does. Nothing about the renormalization story we just finished telling changes that conclusion. (Indeed, that conclusion was really *only* about the renormalized version of the perturbation series, since that's the one with finite coefficients.)

This doesn't actually have much of a practical effect on the way we use the series to predict scattering amplitudes: in practice, one just stops computing when the number of Feynman diagrams starts to become unmanageable and uses that as the prediction, and this works quite well. But of course it does have a rather large *theoretical* effect! In what sense can we think of our series as an approximation of anything if it doesn't converge for any nonzero value of λ_r ? Does it have any meaning at all?

The fact that results can be extracted from perturbation theory can be made to line up so well with experiment should be taken as evidence that there is *something* going on behind it. And it is possible for the first few terms of a power series to provide a good approximation for a function even if the series doesn't actually converge at any nonzero value of its parameter, for example if the series is an asymptotic expansion of the function, and that may very well be the case for the perturbation series.

But it's hard to imagine proving such a theorem as things stand now, because, for interacting theories in four dimensions, we don't have a rigorous description of the function that the series would be an approximation *of*! As we've mentioned repeatedly throughout this series, no one has managed to build a complete mathematical theory which spits out a well-defined function corresponding to the scattering amplitudes we're trying to compute.

(There *has* been some success building a rigorous nonperturbative model of ϕ^4 theory in two or three spacetime dimensions. I know very little about how this goes, but the standard reference is the book *Quantum Physics: A Functional Integral Point of View* by James Glimm and Arthur Jaffe if you are interested in learning more.)

There is a program called *perturbative algebraic quantum field theory*, which is composed of honest, rigorous mathematics, but as the name suggests its output is essentially the same formal power series that we managed to build through our more physically oriented, hand-wavy methods. As far as I can tell, doing better than this is seen within the field as an extremely difficult open problem. As Urs Schreiber put it in the introduction to his “Introduction to Perturbative Quantum Field Theory” (which is a decent reference if you’re interested in this side of the story) building a rigorous, non-perturbative model of something like Yang–Mills theory “might well be a 10^4 year problem”.

6.2 More on Renormalization Schemes

It’s also interesting to reflect on the role that the parameters in the Lagrangian played in our renormalization story. We started with a Lagrangian that contained two free parameters, m and λ . Our perturbative renormalization procedure started by rewriting the Lagrangian in terms of the more physically meaningful μ and λ_r . This required introducing counterterms, which are related to the original parameters via the equations $\delta Z = Z - 1$, $\delta m^2 = Zm^2 - \mu^2$ and $\delta\lambda = Z^2\lambda - \lambda_r$.

At each finite value of the cutoff, we can use our renormalization conditions to write δZ , δm^2 , and $\delta\lambda$ in terms of μ and λ_r . But this required assigning them values which blow up as the cutoff is taken to infinity; indeed, the fact that the divergences in our diagram integrals could be “absorbed” by the counterterms in this way is the entire reason that the procedure actually works.

This, in turn, means that Z , m , and λ are *themselves* assigned values that blow up as the cutoff is taken to infinity. In particular, while the way we wrote our Lagrangian made it look like we were describing a family of quantum field theories parametrized by m and λ , we can see now that this is a misleading description of the situation: except for the free theories (where $\lambda = 0$) *none* of the theories in the family we’re looking at correspond to finite values of m and λ .

What they do correspond to is finite values of μ and λ_r . I like to think of μ and λ_r as a coordinate system on a (hypothetical, not actually rigorously defined) two-dimensional space of quantum field theories. From this perspective, the formal power series in λ_r that we ended up with can be thought of as a description of what this space looks like in an infinitesimal neighborhood of the point with coordinates $(\mu, 0)$, which corresponds to a free theory of particles with mass μ . A rigorous description of what things look like further away from this point would require building the object we just finished saying no one knows how to build.

The meaning of μ and λ_r was set by our renormalization conditions, and while the conditions we chose were easy to interpret physically, there’s no mathematical reason we couldn’t use other ones. In the terminology introduced earlier, this would mean choosing a different “renormalization scheme,” and I like to think of this as a change of coordinates on this space of theories we’ve been imagining. We could, for instance, have defined λ_r using some other value of the four-point function, even one where the input momenta aren’t physically possible for a particle of mass μ . (The procedure in the proof of the BPHZ theorem actually corresponds to a renormalization scheme where we set the values of the two- and four-point functions at input momenta which are all zero.)

There are a lot of renormalization schemes in use, and the one that I believe is the most popular for doing actual computations is quite different from the one we described here. It’s called the “modified minimal subtraction” scheme, basically always abbreviated $\overline{\text{MS}}$, and it works by specifying the functional form of the counterterms rather than directly specifying the values of our renormalized parameters.

As usual for computational issues like this, I recommend the discussion in Peskin and Schroeder, where they introduce $\overline{\text{MS}}$ in Section 11.4. Near the end of that section, they also discuss one of the more important reasons one might want to change renormalization schemes: the so-called “large logarithm problem” that arises when you use a renormalization scheme where the parameters are set at low energies to compute scattering amplitudes at high energies.

6.3 The Role of Renormalizability

Regardless of the renormalization scheme you pick, our original Lagrangian with its poorly-behaved parameters m and λ is really only a meaningful object after imposing a cutoff. You can, if you like, picture an entire family of “cut-off Lagrangians,” one for each value of the cutoff, each of which does have a well-defined (although cutoff-dependent) value of m and λ ; from this perspective, our original Lagrangian is just a constraint on the form these can take. If our Lagrangian is *renormalizable*, that means it’s possible to pick a renormalization scheme and finitely many coordinates such that, when you write the scattering amplitudes in terms of these coordinates, the coefficients of the resulting perturbation series converge to something finite as we raise the cutoff to infinity.

For the first few decades of quantum field theory’s existence, it was seen as more or less mandatory that the theories used to describe the fundamental interactions of physics be renormalizable. This was quite helpful in narrowing down the form of the Lagrangian, since renormalizability is a fairly strong constraint — we caught a glimpse of this in our brief discussion of ϕ^6 theory above. Theorists were in fact able to find perturbatively renormalizable descriptions of all the currently known fundamental forces other than gravity, resulting in the now-famous Standard Model of particle physics.

(Renormalizability is, by the way, the main obstacle to constructing a sensible quantum theory of gravity: the Lagrangian that gives rise to general relativity is not renormalizable, hence the search for some other theory that might not suffer from this problem.)

More recently, though, a new perspective has arisen which partially rehabilitates the non-renormalizable theories; this is the “Wilsonian” perspective mentioned in the introduction. We’ll have more to say about this later in this series, but we can say a bit about it now.

Nonrenormalizability is only a problem if you actually want to take the cutoff all the way to infinity. As we’ll see when we explore this later on, leaving the cutoff large but finite is fine if you are only interested in energies that are much smaller than the scale of the cutoff, and it turns out that nonrenormalizable interaction terms in the Lagrangian make contributions to scattering amplitudes that shrink with the ratio between the particles’ energies and the cutoff.

In light of this, while nonrenormalizable theories are unsuitable as models of “fundamental” physics that we expect to hold for arbitrarily high energies, they can be perfectly fine **effective field theories**, that is, theories whose validity is constrained to some smaller range of energies. In fact, I believe the dominant perspective in the field is that the Standard Model itself is probably an effective field theory, that is, that there is some other, yet-to-be-discovered model which presumably includes gravity and to which the Standard Model is a low-energy approximation.

In fact, because general relativity models gravity as arising from the curvature of the metric on spacetime, one might suspect that any quantum theory that incorporates gravity will involve some change to the structure of spacetime itself, and may not even be a quantum field theory at all. This is, I think, part of what might be behind some physicists’ cavalier attitude toward some of the foundational issues we’ve talked about. If you’re willing to impose an ultraviolet cutoff, everything is much more mathematically straightforward. By trying to take the ultraviolet limit in a way that leaves the geometry of spacetime intact, we may be making some sort of conceptual

mistake from the perspective of this hypothetical theory of quantum gravity. In other words, we may be having so much trouble making mathematical sense of this limit because it's just the wrong limit to take.

We'll have to leave the Wilsonian story here for now. It will be much easier to explain it in detail once we've been acquainted with the functional integral perspective on quantum field theory. That's where we'll turn next.