THE DIVERGENCE COMPLEX AND FEYNMAN DIAGRAMS

AARON MAZEL-GEE & EUGENE RABINOVICH

1. Motivation: Integrals

We wish to do integrals on a finite-dimensional, compact, connected, orientable manifold X of dimension N. Normally, we integrate an n-form μ over X, but we can replace the de Rham complex $\Omega^\bullet(X)$ with the complex of alternating multi-vector fields $V_\bullet$. This is a graded commutative algebra, just as $\Omega^\bullet$ is. Moreover, any choice of volume form μ on X gives an isomorphism $V_{n-\bullet} \cong \Omega^\bullet$ which is simply given by $Y_1 \wedge Y_2 \wedge \cdots Y_k \mapsto \mu(Y_1, Y_2, \cdots, Y_k, -, -, \cdots, -)$. We can transfer the de Rham differential d to the complex $V_\bullet$ to get a degree -1, square-zero operator Δ_μ. $(V_\bullet, \Delta_\mu)$ is not a differential-graded algebra since Δ_μ is actually a second-order differential operator: in local coordinates where $\mu = dx^1 \wedge \cdots dx^n$ and we define $\xi_i := \partial_i$, then $\Delta_\mu = \sum_i \frac{\partial^2}{\partial \xi_i \partial x_i}$. The reason that this happens is that whereas the de Rham differential has a multiplication by a form in its definition, the dual operation on the space of multivector fields is to differentiate with respect to the corresponding multivector field. Now, suppose we want to compute the zeroth homology group of the complex $V_\bullet$. By construction, this is just the n-th de Rham cohomology group of $\Omega^\bullet$. Since X is compact and connected, the n-th de Rham cohomology of X is 1-dimensional, and the map $V_0 \to \mathbb{R}$ given by $f \mapsto \int_X f\mu$ descends to an isomorphism on $H^0(V_\bullet)$. We see, therefore, that integration of functions with respect to the volume form μ corresponds to computing the zeroth homology group of the complex $V_\bullet$. In physics, we're interested in integration over an infinite-dimensional manifold. On an infinite-dimensional space, there's no such thing as a top form and the de Rham complex and $V_\bullet$ are no longer isomorphic; thinking about V_0 is our best alternative to thinking about volume forms.

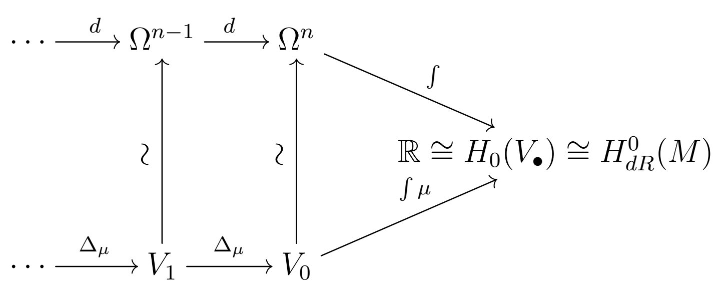

Figure 1. The de Rham and divergence complexes

Moreover, the volume form of interest in QFT looks like $e^{-S/\hbar}\mu$, where S is some function on X, and we wish to compute the expectations of functions:

$$\langle f \rangle := \frac{[f]}{[1]} = \frac{\int_X f e^{-S/\hbar}\mu}{\int_X e^{-S/\hbar}\mu}.$$

In the quantum field theory context, these are the expectation values of operators, which can be used to give the scattering probabilities of different physical processes. Now, let us assume S has a single minimum and non-degenerate Hessian; then, as $\hbar \to 0$, the measure $e^{-S/\hbar}\mu$ is supported only near the minimum of S. We can choose local coordinates such that the minimum of S is at 0 and the measure μ is just $dx^1 \wedge \cdots dx^n$. In these coordinates, we can write $s(x) = s(0) + \frac{1}{2}\sum a_{i,j}x_i x_j - b(x)$ for some symmetric, positive-definite a and $b(x)$ vanishing at least to order 3; furthermore,

$$-\hbar\Delta_{\exp(-S/\hbar)\mu} = \sum_{i=1}^{N} a_{ij}x_i\frac{\partial}{\partial\xi_j} - \sum_{i=1}^{N}\frac{\partial b(x)}{\partial x_i}\frac{\partial}{\partial\xi_i} - \hbar\sum_{i=1}^{N}\frac{\partial^2}{\partial x_i\partial\xi_i}.$$

Expanding b in a "Taylor" series around 0 (i.e. computing in a formal neighborhood of zero), and computing the homology of the complex $(V_\bullet, \Delta_{\exp(-S/\hbar)\mu})$ order by order in $\hbar$ leads us to the homological problem we consider next.

2. Statement of the General Problem

Motivated by the example of finite-dimensional integration, we consider the complex $V_\bullet := \mathbb{R}[\![x_1, \cdots, x_N, \xi_1, \cdots, \xi_N, \hbar]\!]$ with the ξ's in degree 1 and all other variables in degree 0. Given a symmetric, positive-definite 2-tensor ($N \times N$ matrix) a_{ij} and a $b(x) \in \mathbb{R}[\![x_1, \cdots, x_N]\!]$ consisting of only cubic and higher terms, we can define the following homological degree -1, square-zero operator

$$Q = \sum_{i,j=1}^{N} a_{ij}x_i\frac{\partial}{\partial\xi_j} - \sum_{i=1}^{N}\frac{\partial b(x)}{\partial x_i}\frac{\partial}{\partial\xi_i} - \hbar\sum_{i=1}^{N}\frac{\partial^2}{\partial x_i\partial\xi_i}.$$

$V_\bullet$ represents the structure of a formal neighborhood of M, with the a_{ij} and $b(x)$ representing the power series expansion of S, and is a graded-commutative algebra. Q is not a derivation, so $(V_\bullet, Q)$ is not a dga. (As an aside, we can compute the failure of Q to be a derivation; this is called the *Schouten-Nijenhuis bracket* and gives a Poisson bracket on $V_\bullet$). Since everything in the motivating case above was $\hbar$-linear, we should expect that the homology of this complex should be $\mathbb{R}[\![\hbar]\!]$ in degree 0, since it was $\mathbb{R}$ in the motivating case. We elevate this to an ansatz. Now, Q cannot create power series that don't have at least a linear term; in particular, this means that $1 \in V_0$ is not exact. Using this and our ansatz, we can define $\langle f\rangle \in \mathbb{R}[\![\hbar]\!]$ for any $f \in V_0$ by $[f] = \langle f\rangle[1]$. $\langle f\rangle$ is the homological analog of computing the expectation of f via integration. We wish to manage to the combinatorics of computing $\langle f\rangle$ in some easily enumerable way. This is where Feynman diagrams come in. But first, let's do a simple example.

3. Wick's Lemma

Set $b = 0$ and consider the problem of finding $[x_{i_1}x_{i_2}\cdots x_{i_n}]$. To this end, we compute a specific boundary element of V_0:

$$Q\left(\sum_{k=1}\xi_k(a^{-1})_{ki_1}x_{i_2}x_{i_3}\cdots x_{i_n}\right) = \sum_{i,j} x_i(a^{-1})_{ji_1}a_{ij}x_{i_2}x_{i_3}\cdots x_{i_n} - \hbar\sum_{\alpha=1}^{n}(a^{-1})_{i_1i_\alpha}x_{i_2}x_{i_3}\cdots\hat{x}_{i_\alpha}\cdots x_{i_n}$$
$$= x_{i_1}x_{i_2}\cdots x_{i_n} - \hbar\sum_{\alpha=1}^{n}(a^{-1})_{i_1i_\alpha}x_{i_2}x_{i_3}\cdots\hat{x}_{i_\alpha}\cdots x_{i_n},$$

where the hat indicates removing that variable from the product. So, we see that the homology class of $x_{i_1} \cdots x_{i_n}$ is the same as the homology class of $\hbar$ times all ways of "contracting" x_{i_1} with another x_{i_a} and replacing the product with $(a^{-1})_{i_1 i_a}$. When $n = 1$, this formula actually implies that $\langle x_i \rangle = 0$. By induction, the expectation value of any product of an odd number of x_i's is zero. On the other hand, for a product of an even number of x_i's, we can use induction to show that

$$\langle x_{i_1} x_{i_2} \cdots x_{i_{2n}} \rangle = \hbar^n \sum_{\text{pairings } \{1\pi(1)\}\{2\pi(2)\}\cdots\{2n\pi(2n)\}} (a^{-1})_{i_1 i_{\pi(1)}} (a^{-1})_{i_2 i_{\pi(2)}} \cdots (a^{-1})_{i_n i_{\pi(n)}}$$

We can represent the above equation as an equality between two sums of diagrams:

(1)

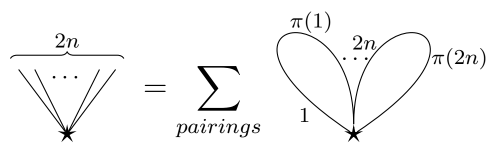

On the left-hand side, we have a diagram with one marked vertex and $2n$ half-edges emanating from that vertex. The external half-edges represent the x_i's. The RHS is a sum of diagrams, where each diagram is a possible way of pairing off the $2n$ half-edges of the LHS diagram and making a loop from those two half-edges. Each loop contributes a factor of $\hbar$ and each internal edge contributes an a^{-1} factor. In this case, we don't really need the diagrammatic approach, but when $b \neq 0$, we'll see how the diagrams can greatly simplify our work.

4. Feynman Diagrams

Consider now the case $N = 1$, $b = x^3/6$. Let's try to apply our recursive approach here. We compute:

$$Q(x^n \xi) = ax^{n+1} - \frac{1}{2}x^{n+2} - n\hbar x^{n-1}.$$

Rearranging, we notice that

$$[x^{n+1}] = \frac{1}{2a}[x^{n+2}] + \frac{n\hbar}{a}[x^{n-1}]. \tag{2}$$

Let's try to use this relation to find $[x]$ (so we start with $n = 0$):

$$[x] = \frac{1}{2a}[x^2] = \frac{1}{2a}\left(\frac{1}{2a}[x^3] + \frac{\hbar}{a}[1]\right) = \left(\frac{1}{2a}\left(\frac{1}{2a}[x^4] + \frac{2\hbar}{a}[x]\right) + \frac{\hbar}{a}[1]\right).$$

It's easy to see that this will get very complicated very fast; we already have a second appearance of $[x]$ on the RHS of the above equation. We can bring the term involving $[x]$ to the LHS, but this doesn't help very much. We might try to interpret this as an equation for $[x^{n+2}]$ in terms of strictly lower powers of x, but for completely arbitrary b, we can have infinitely many terms replacing $[x^{n+2}]$. Thus, we need to understand the combinatorics of the process much better to get any further. To this end, let us drop the assumption $N = 1, b = x^3/6$ and instead work with completely general b and N. First write

$$b(x) = \sum_{n=1}^{\infty} \sum_{\vec{\imath} \in \{1,2,\cdots,N\}^n} \frac{b_{\vec{\imath}}^{(m)}}{m!} x_{i_1} \cdots x_{i_n},$$

where $b^{(m)}$ is the completely symmetric tensor of m-th order partial derivatives of b at zero. Then, a very simple modification of the arguments we've seen already shows that

$$
\begin{aligned}
(3)\quad Q\left(\sum_{i,\vec{\imath}} f_{i,\vec{\imath}}\, x_{\vec{\imath}}\, (a^{-1})_{i,j}\, \xi_j\right) \\
= \sum_{i,\vec{\imath}} f_{i,\vec{\imath}}\, x_i x_{i_1} \cdots x_{i_n} - \sum_{m=2}^{\infty} \sum_{i,\vec{\imath},j,\vec{\jmath}} \frac{1}{m!}\, b^{(m+1)}_{j,\vec{\jmath}}\, x_{j_1} \cdots x_{j_m}\ (a^{-1})_{i,j}\, f_{i,\vec{\imath}}\, x_{i_1} \cdots x_{i_n} \\
- \hbar \sum_{i,\vec{\jmath}} \sum_{k=1}^{n} f_{i,\vec{\jmath}}\, (a^{-1})_{i,j_k} x_{j_1} \ldots \hat{x}_{j_k} \ldots x_{j_n}
\end{aligned}
$$

for any n-tensor f.

This is kind of a mess. To clean up the mess, we define a combinatorial/topological object that'll help us handle the recursion relations much better. This object is called a **Feynman diagram** Γ, which consists of the following data:

- A finite set $V = V_{int} \coprod \{\star\}$ of vertices, of which we distinguish one by the star.
- A finite set HE of half-edges of the diagram.
- A map $i : HE \to V$ assigning to each half-edge the vertex on which it's incident.
- An involution $\sigma : HE \to HE$. We will call the fixed points of the involution **external half-edges** and the non-fixed points the **internal half-edges**. The size two orbits of σ are called **internal edges**.
- An ordering of $i^{-1}(\star)$.

We define the **valence** of a vertex v to be the size of $i^{-1}(v)$ and denote it by $val(v)$. We require as part of the definition that $val(v)$ to be at least 3 for all non-starred vertices v. We can also make the following definitions:

- The **in degree** of a Feynman diagram Γ is $val(\star)$. We denote this by $d_{in}(\Gamma)$.
- The **out degree** of a Γ is the number of external half-edges. We denote this by $d_{out}(\Gamma)$.
- An **embedding** of Feynman diagrams $\Gamma' \hookrightarrow \Gamma$ is a pair of injective maps $f_1 : HE' \to HE$ and $f_2 : V' \to V$ on vertices such that $i \circ f_1 = f_2 \circ i'$ (the maps preserve the incidence relations), elements of the same orbit under σ' get taken to elements of the same orbit under σ, and $f_2(\star') = \star$; we also require that $d_{in}(\Gamma') = d_{in}(\Gamma)$ and that f_1 preserve the ordering on the half-edges incident on $\star$.
- An **isomorphism** $\Gamma' \to \Gamma$ of Feynman diagrams is a bijection on half-edges and vertices that preserves the incidence relations and pairings of half-edges, as well as the ordering of the half-edges at the marked vertex. An **automorphism** is defined similarly.
- A Feynman diagram is **closed** if all of the σ orbits are of size 2, i.e. there are no external half-edges.
- The **Betti number** of a Feynman diagram is its first Betti number, which can be computed as the number of orbits of σ minus the number of unmarked vertices in Γ.
- A **corolla** is a Feynman diagram with such that $V = \{\star\}$ and all half-edges are external. A corolla is depicted on the LHS of equation 1.

We note that this combinatorial data is equivalent to the specification of a 1-dimensional CW complex with some extra structure. We call a Feynman diagram connected if it's connected as a CW complex. We assume that all Feynman diagrams are connected from now on. Figure **??** gives an example of a Feynman diagram.

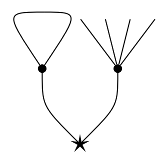

(A) A Feynman diagram with $d_{in} = 2$, $d_{out} = 4$, and Betti number 1.

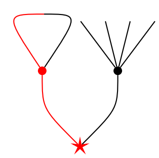

(B) An embedding of Feynman diagrams. Note that external half-edges can be identified with internal half-edges.

Figure 2. A Feynman diagram and an embedding of Feynman diagrams.

We now re-express equation 3 in terms of Feynman diagrams. Let $\widetilde{FD}$ denote the set of Feynman diagrams with orderings of the external half-edges. The ordering of the external half-edges will prove to be useful in what follows. Given a power series $f \in \mathbb{R}[\![x_1, \cdots, x_N]\!]$, we can define a function $ev_f : \widetilde{FD} \to V_0$ as follows. First, we define the function for a homogeneous polynomial f and extend by linearity to arbitrary f. So, let $f = \sum_{\vec{\imath} \in \{1,\cdots,N\}^{n+1}} f_{\vec{\imath}}\, x_{i_1} \cdots x_{i_{n+1}}$. We define $ev_f(\Gamma)$ to be zero unless $d_{in}(\Gamma) = n+1$. If $d_{in}(\Gamma) = n+1$, we define $ev_f(\Gamma)$ by first labeling the half-edges of Γ by numbers in $\{1, \cdots, N\}$. To a Feynman diagram labeled thus, we form the product of factors:

- For the marked vertex, a factor $f_{\vec{\imath}}$, where $\vec{\imath}$ is the vector of labels on the half-edges incident to the starred vertex, read off in the ordering given to those half-edges.
- For each internal edge with labels i, j on the half-edges, a factor a^{-1}_{ij}.
- For each external half-edge with label i, a factor x_i.
- For each internal vertex with valence m, a factor $b^{(m)}_{\vec{\imath}}$, where $\vec{\imath}$ is the vector of labels formed by reading the labels on the half-edges incident on the vertex in any order ($b^{(m)}$ is a symmetric tensor).

Then, $ev_f(\Gamma)$ is the sum over all possible labellings of Γ in this way. Note that if f is a degree n homogeneous polynomial, then $ev_f(\cdot)\hbar^{\beta(\cdot)}$ evaluated on the n-fold corolla gives f again.

We can also define two types of operations, Δ^{μ}_m and Δ^{κ}_k, on $\mathbb{Q}^{\widetilde{FD}}$. Given $\Gamma \in \widetilde{FD}$ with $n+1$ external half-edges, $\Delta^{\mu}_m(\Gamma)$ is the graph gotten from Γ by attaching the $n+1$-th external half-edge of Γ to some half-edge of a vertex of valence $m+1$. The m new external half-edges created in this way are given some ordering $n+1, \cdots, n+m$. On the other hand, Δ_k^κ takes the k-th and $n+1$-th external half-edges and pairs them off into a loop (notice that we need $1 \leq k \leq n$ for this to make sense). Figure **??** illustrates these operations. Now, consider the set $\mathbb{Q}^{\widetilde{FD}}$ of all $\mathbb{Q}$-formal sums of elements of $\widetilde{FD}$.

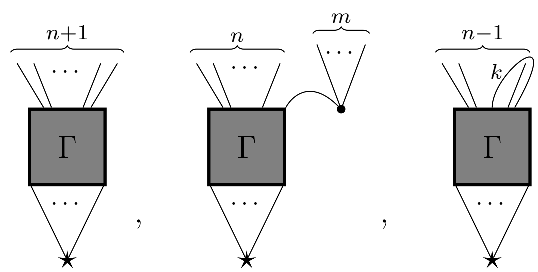

FIGURE 3. A Feynman diagram Γ, $\Delta_m^\mu(\Gamma)$, and $\Delta_k^\kappa(\Gamma)$. Whatever is inside the gray box is the same across all three diagrams.

We define the function $E : \mathbb{Q}^{\widetilde{FD}} \to \mathbb{Q}^{\widetilde{FD}}$ by

$$\Gamma \mapsto \sum_{m=2}^{\infty} \frac{1}{m!} \Delta_m^\mu(\Gamma) + \sum_{k=1}^{n} \Delta_k^\kappa(\Gamma)$$

when Γ has $n+1 > 0$ external half-edges and $\Gamma \mapsto \Gamma$ when Γ is closed. We extend to all of $\mathbb{Q}^{\widetilde{FD}}$ by linearity. Then, equation 3 tells us that the following diagram commutes, since the Betti number remains unchanged under a Δ^μ operation and goes up by one under a Δ^κ operation:

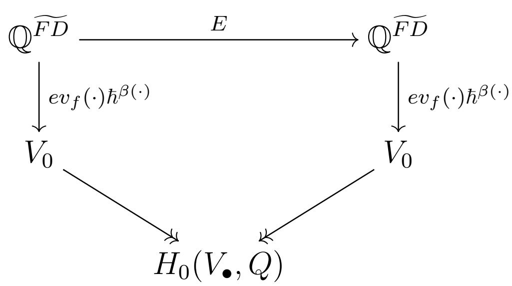

Strictly speaking, $ev_f(\cdot)\hbar^{\beta(\cdot)}$ is only a function on Feynman diagrams, but we can extend it to $\mathbb{Q}^{\widetilde{FD}}$ by linearity.

So, if we start with an $n+1$-fold corolla Ψ, which computes the homology class of degree $n+1$ homogeneous polynomials, we can apply the operation E without changing the homology class corresponding to the new diagrams created by E, i.e.

$$[ev_f(\Psi)\hbar^{\beta(\Psi)}] = [ev_f(E^k\Psi)\hbar^{\beta(E^k\Psi)}]$$

for all $k \geq 0$, and if we start with a degree $n+1$ homogeneous polynomial f, we have

$$[f] = [ev_f(\Psi)\hbar^{\beta(\Psi)}] = [ev_f(E^k\Psi)\hbar^{\beta(E^k\Psi)}]. \tag{4}$$

$E^k\Psi$ is a sum over all possible diagrams obtained from Ψ by a sequence of k Δ operations (of type Δ^μ or Δ^κ), with each diagram obtained thus weighted by a factor of $\left(\prod_{v \in V_{int}} (val(v) - 1)!\right)^{-1}$, since every time E creates an internal vertex of valence of $m+1$

by a Δ^μ_m operation, it's with a factor of $1/m!$. In the power series topology, the $k \to \infty$ limit of $ev_f(E^k\Psi)\hbar^{\beta(E^k\Psi)}$ is a sum over only closed diagrams with $d_{in} = n+1$, since large numbers of external vertices correspond to very high-degree polynomials and these give a small contribution in the power series topology. $ev_f(\cdot)\hbar^{\beta(\cdot)}$ is a polynomial in $\hbar$ for closed diagrams, so to compute $\langle f \rangle$ for any degree $n+1$ polynomial we simply need to figure out which isomorphism classes of closed diagrams contribute to the above sum and with what weights. In fact, our main result is

Theorem 4.1. *If f is a homogeneous degree $n+1$ polynomial, then*

$$\langle f \rangle = \sum_\Gamma \frac{ev_f(\Gamma)\hbar^{\beta(\Gamma)}}{|Aut(\Gamma)|},$$

where the sum is over representatives of isomorphism classes of closed (and connected) Feynman diagrams Γ with $d_{in}(\Gamma) = n+1$.

Proof. It will suffice to show the following three things:

1. (1) Every isomorphism class of closed Feynman diagrams appears in this sum.
2. (2) We have only finitely many contributions to the sum at each order in $\hbar$.
3. (3) Each isomorphism class appears with weight $|Aut(\Gamma)|^{-1}$.

The proof of (1) will follow from the methods used in the proof of (3). To see (2), note that for a closed Feynman diagram, all σ-orbits are size 2 and therefore the Betti number is $\frac{1}{2}|HE| - |V_{int}|$. But the marked vertex contributes $n+1$ half-edges, and each internal vertex contributes at least 3 half-edges (all internal vertices are at least trivalent), so

$$\beta(\Gamma) \geq \frac{n+1}{2} + \frac{|V_{int}|}{2}.$$

Thus, the Betti number gives a bound on the number of internal vertices, so there are finitely many isomorphism classes of graphs with any given Betti number. This proves (2).

Now, we only need to prove (3) (and (1) along the way). First note that Γ appears in the sum only if Γ (or, strictly speaking, a Feynman diagram isomorphic to Γ) can be constructed from the $n+1$-fold corolla by a sequence of Δ^κ and Δ^μ operations. For any Γ, let $\kappa(\Gamma)$ be the number of different ways Γ can be constructed in this way; we call this the **construction number** of Γ. Then, by the discussion after equation 4, each isomorphism class of closed Feynman diagrams appears with weight

$$\kappa(\Gamma)\left(\prod_{v \in V_{int}(\Gamma)} (val(v)-1)!\right)^{-1}.$$

If we show that this factor is $|Aut(\Gamma)|^{-1}$, then we're done, because we will have in particular shown that $\kappa(\Gamma) \neq 0$. Now, a separate cyclic ordering of the half-edges incident on each internal vertex of Γ gives rise to a construction of a diagram isomorphic to Γ by the following iterative method. Given a Feynman diagram Γ' with an ordering of the external half-edges and an embedding $\Phi : \Gamma' \hookrightarrow \Gamma$, we can create a new Feynman diagram $\nu(\Gamma)$ with the same attributes by looking at the highest ordered external half-edge v of Γ' and asking whether $i(\sigma(\Phi(v)))$ is a vertex in Γ that is identified with a vertex of Γ'. If so, perform a Δ^κ operation and connect v with the half-edge corresponding to $\sigma(\Phi(v))$. Otherwise, perform a Δ^μ_m operation with $m = val(i(\sigma(\Phi(v)))) - 1$, using the cyclic ordering of the half-edges given by Γ to order the newly created external half-edges. Now, we see that starting with $\Gamma' = \Psi$, the $n+1$-fold corolla, this method gives rise to a diagram isomorphic to Γ after some (finite) number of iterations, because such a construction goes through constructs each vertex of Γ exactly once and pairs off the half-edges of Γ appropriately. Figure **??** illustrates an example. We call this newly constructed Feynman diagram Γ'' and note that we have an isomorphism $\Gamma'' \to \Gamma$ handed to us by the iterative method used to produce Γ''. Moreover, any construction of Γ gives a cyclic ordering on the half-edges incident on the vertices of Γ'' induced from the total ordering that the half-edges received when their vertex was constructed by a Δ^μ operation; we can pull back this cyclic ordering to Γ, and it's easy to see that this cyclic ordering will construct the same Γ''. In other words, we have the following map

$$\begin{array}{c} \Omega \\ \big\downarrow\!\!\!\!\downarrow \\ M, \end{array} \tag{5}$$

where Ω is the set of cyclic orderings of half-edges incident on each of the vertices of Γ and M is the set of constructions of Γ.

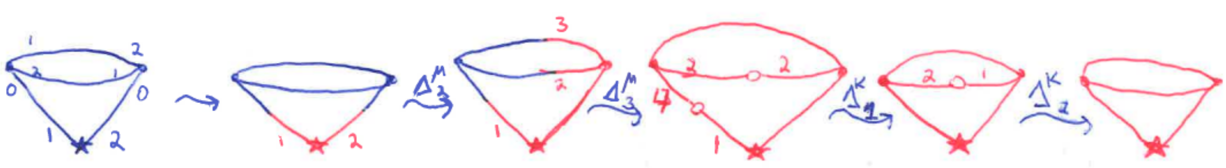

FIGURE 4. A cyclic ordering gives rise to a construction of Γ, depicted here as a sequence of embeddings into Γ. In the first picture, the numbers indicate the cyclic ordering on the half-edges incident at each vertex (except at the starred vertex, where we simply have an ordering). In the subsequent diagrams, we've only numbered the external half-edges of the embedded diagram.

The size of Ω is precisely $\prod_{v \in V_{int}(\Gamma)} (val(v)-1)!$ and the size of M is $\kappa(\Gamma)$. The automorphisms of Γ act on Ω in the natural way (we push the ordering forward via the isomorphism); if we can show that this action is fiberwise, and free and transitive on the fibers, then we'll have shown that

$$\kappa(\Gamma) = \frac{\prod_{v \in V_{int}(\Gamma)} (val(v)-1)!}{|Aut(\Gamma)|},$$

then we'll be done. (In more geometric terms, $\Omega \to M$ is a principal $Aut(\Gamma)$ bundle.) To see that this action is fiberwise, note that an automorphism is a way of permuting the half-edges and vertices of a Feynman diagram without changing the information of which half-edges are connected to each other and which half-edges are incident on a given vertex. Another way to say this is that if we name the vertices and half-edges, an automorphism might move vertex A to the spot where vertex B was, but it won't change the fact that half-edge 1 of vertex A is attached to half-edge 4 of vertex C. Moreover, the way that $Aut(\Gamma)$ acts on Ω guarantees that (e.g.) if half-edge 3 of vertex A is succeeded by half-edge 1 and preceeded by half-edge 2 in a particular cyclic ordering, then under an automorphism this will still be true

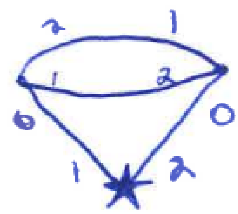

FIGURE 5. The automorphism exchanging the two top edges of Γ gives rise to a cyclic ordering distinct from the one in figure ??. It does, however, give rise to the same construction.

in the new cyclic ordering. Thus, since this is the only information we needed to uniquely determine an element of M, $Aut(\Gamma)$ acts fiberwise (see figure ??).

To see that the action is transitive on the fibers, let p, q be two different cyclic orderings of the half-edges incident on the vertices of Γ that give rise to the same construction of Γ. Note that both of these constructions will give two (possibly different) isomorphisms $\Gamma'' \to \Gamma$, with Γ'' being the same in both cases since the construction was the same. Composing one isomorphism with the inverse of another gives an automorphism of Γ; we need to check that it takes the cyclic ordering p to the cyclic ordering q. But this is a consequence of an argument we made above: the fact that the two isomorphisms $\Gamma'' \to \Gamma$ arose from the same sequence of Δ operations means that the cyclic ordering induced on the half-edges incident at each vertex of Γ'' by the sequence of Δ operations pulls back under the two isomorphisms to p and q. Therefore, the automorphism $\Gamma \to \Gamma$ takes $p \mapsto q$. Finally, we show that the $Aut(\Gamma)$ action is free on the fibers. Suppose $\phi \in Aut(\Gamma)$ fixes a cyclic ordering. Then, unless $\phi = e$, there must be some vertex $v \in \Gamma$ such that ϕ fixes v but not all half-edges incident on v; this is because if ϕ fixed all half-edges incident on all vertices that it fixes, then since ϕ has to fix the starred vertex and all half-edges incident on it by the definition of automorphism, it would have to fix all vertices connected by an edge to the starred vertex, and it would have to fix all half-edges incident on those vertices and fix the vertices attached to those half-edges and so on, and so ϕ would therefore be the identity. Since Γ is connected, this shows that ϕ must fix at least one vertex v but not the half-edges incident on it. Thus, ϕ must change the cyclic ordering at v; but this is a contradiction, so we must have $\phi = e$ and therefore the $Aut(\Gamma)$ action on Ω is free. This completes the proof. □