From Path Integrals to Fractional Quantum Statistics

Joshua Horowitz

(Dated: April 14, 2009)

Fractional quantum statistics, an exotic phenomenon in which particles confined to a plane exhibit properties between those of bosons and fermions, is developed within the context of Feynman's path-integral formulation of quantum mechanics. Starting from the geometric concept of the configuration space, the path-integral formulation is introduced and applied to the general problem of many-body quantum systems. In two dimensions, this analysis produces the anyon. Finally, some connections are made between this theoretical construct and actual physical phenomena, including the fractional quantum Hall effect.

I. INDISTINGUISHABLE PARTICLES AND CONFIGURATION SPACE

The concept of indistinguishable particles was first used by Josiah Willard Gibbs in the 19th century, decades before the advent of quantum mechanics [**?**]. Gibbs discovered that a näive calculation of the entropy of an ideal gas led to an anomalous result, in which allowing two identical gases to mix would lead to an increase in their entropy. This result made no physical sense, as such a mixing is macroscopically reversible.

Gibbs was able to resolve this paradox using indistinguishability. He noticed that the calculation of the näive entropy overcounted the actual number of microstates by counting each permutation of the particles as a different microstate. That is, if we write out a microstate as a sequence of positions and momenta for each particle, the näive entropy considered

$$(\vec{x}_1, \vec{p}_1; \vec{x}_2, \vec{p}_2; \ldots; \vec{x}_N, \vec{p}_N)$$

to be a different microstate than

$$(\vec{x}_2, \vec{p}_2; \vec{x}_1, \vec{p}_1; \ldots; \vec{x}_N, \vec{p}_N),$$

and so on, for every possible permutation of the particles' order. If the particles are indistinguishable, as we have assumed, then there is really no measurable difference between these states. Once Gibbs added a $\frac{1}{N!}$ multiplier to his microstate count, in order to compensate for the $N!$-fold overcounting, he obtained an entropy function which no longer predicted an increase in entropy during the mixing of identical gases.

Gibbs was the first physicist to propose that the indistinguishability of particles had physical significance. It is understandable that this significance had remained undeveloped for so long, as indistinguishability rarely matters in classical physics. The reason why Gibbs ran into it, while centuries of physicists before him never did, was that Gibbs was considering entropy, a quantity which relies on the global structure of the state space, rather than the local structure. This line of reasoning is worth explaining further, as it will prove to be important when we move to quantum mechanics, another situation in which the global structure becomes important.

State space is a geometric object which encodes every possible state a system can have. Each point of state space corresponds to a different complete description of the state of the system. A closely related concept is that of *configuration space*, which encodes only the positions of the objects in the system, while completely forgetting any non-positional information. We will look at configuration spaces from now on, as they are more directly used in quantum mechanics (and easier to visualize).

As an example, take a pair of distinguishable particles, A and B, and place them on a line segment of length L. We encode the configuration of the system as a pair of positions (a, b), where each of a and b is between 0 and L. These pairs can be thought of as the points of a filled-in square. This filled-in square, seen in Figure **??**(a), is thus the configuration space for our system.

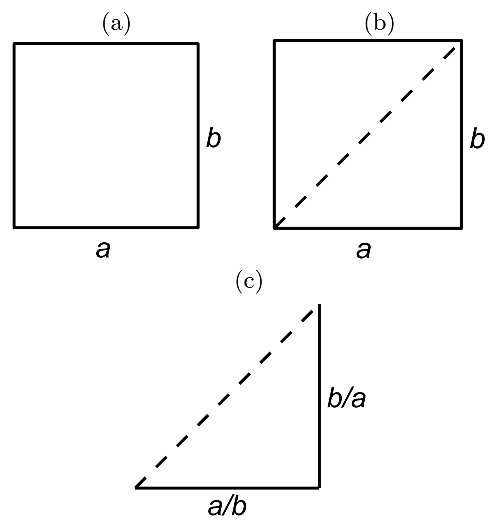

FIG. 1: Configuration spaces for (a) two distinguishable particles which are free to coincide, (b) two distinguishable particles which cannot coincide, and (c) two indistinguishable particles which cannot coincide.

This model calls for some refinements, however. For instance, note that the points on the diagonal of the square represent configurations of the form (x, x), in which particles A and B reside at the same position. We may not actually want to admit this possibility, in which case we must remove this diagonal from our phase space. This possibility is illustrated in Figure **??**(b), where the dotted line denotes a set subtracted from the phase space. Note that this cut separates our configuration space into two disconnected triangles.

Now suppose that we have two indistinguishable particles (which, like before, are not allowed to coincide). It no longer makes sense to speak of particle A or particle B, as though the particles were labeled with little flags, but we can still describe the configuration using a pair of positions (a, b). The important thing to realize is that the configuration described by (a, b) is the exact same configuration as that described by (b, a). Our old picture is now misleading, as it has two separate points in the configuration space representing the same configuration. We need to "identify" the points (a, b) and (b, a). This identification means that we fold one of the triangles over onto the other, and glue them together. We end up with the picture in Figure **??**(c). This triangle is the true configuration space for a pair of non-coinciding indistinguishable particles.

Here, we can start to see the local versus global distinction in action. Comparing the spaces of Figures **??**(b) and **??**(c), one finds that they are locally identical. That is, an ant living on one of these configuration space would not be able to tell which one it was. However, the spaces are undeniably different; the space in Figure **??**(b) is disconnected while that of Figure **??**(c) is connected. This property of connectivity is a global property. Properties such as this are not usually essential to the classical behavior of a system, as classical systems always evolve according to local rules. However, as Gibbs discovered, the global structure of a state/configuration space becomes important when doing statistical mechanics, as statistical mechanics looks at all the possible states of a system at once. And, as we will see, global structure plays a surprisingly important role in quantum physics.

II. INDISTINGUISHABLE PARTICLES IN QUANTUM MECHANICS

The configuration space becomes absolutely essential in the study of quantum mechanics of indistinguishable particles. The state of a many-body quantum system is described using a many-body wavefunction: a function which gives a complex amplitude for every possible collection of positions the many bodies can have. This is usually denoted $\psi(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)$, implying that the wavefunction is a function of ordered N-tuples of positions. However, we know that there is a hidden constraint. If the particles are indistinguishable, we must have

$$\psi(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N) = \psi(\vec{x}_2, \vec{x}_1, \ldots, \vec{x}_N),$$

and so on, for every reordering of the positions. That is, inputs to the wavefunction ψ which represent the exact same configuration must give the exact same output.

However, we can reformulate this symmetry constraint in a more fundamental and natural way. Since $(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)$ and $(\vec{x}_2, \vec{x}_1, \ldots, \vec{x}_N)$ represent the same configuration of the system, they are really just different ways of labeling a single point in configuration space. Why not say that the wavefunction is defined on the configuration space, rather than on the set of tuples like $(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)$, which are really just arbitrary labels for configurations? Then we no longer need to impose a symmetry constraint on the wavefunction. The indistinguishability of the particles is accounted for by the particular geometric structure that indistinguishability gives to the configuration space.

It is worth noting that this is a somewhat atypical way of handling indistinguishability in quantum systems. In this treatment, the wavefunction is always *symmetric*. The reader may be more familiar with the following argument, which leads to either symmetric or *antisymmetric* wavefunctions:

> If our particles are indistinguishable, then the interchange of a pair of particles in our wavefunction cannot introduce any physically measurable change. However, the value of a wavefunction is not itself physical, only its squared norm, which gives a physically measurable probability. Therefore, we need only have
>
> $$|\psi(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)|^2 = |\psi(\vec{x}_2, \vec{x}_1, \ldots, \vec{x}_N)|^2,$$
>
> which implies
>
> $$\psi(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N) = e^{i\theta}\psi(\vec{x}_2, \vec{x}_1, \ldots, \vec{x}_N),$$
>
> where $e^{i\theta}$ is some complex phase. But performing this transposition a second time brings us back to where we started, so we must have $(e^{i\theta})^2 = 1$. Therefore, our phase can be either $+1$ or -1. These possibilities correspond to wavefunctions which are either symmetric or antisymmetric. In the first case, we have bosons, and in the later case, fermions.

This argument leads to some reasonable results, but it is fundamentally nonsensical. The wavefunction $\psi(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)$ is supposed to represent the amplitude for a configuration in which the particles are at the positions $\{\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N\}$. If our particles are indistinguishable, this is the *exact same* configuration as that represented by the input $(\vec{x}_2, \vec{x}_1, \ldots, \vec{x}_N)$. What could it possibly mean to say that the value of the wavefunction depends on what order you present the positions in? It is far less mysteriously formal to say that indistinguishable particles are always represented by symmetric wavefunctions [**?**].

Where, then, do bosons and fermions come from? We will see that they arise quite naturally from the path-integral formulation of quantum mechanics. This alternative perspective will also put us in the position to look a little further and see the possibility of a third, fascinating possibility: the anyon.

III. PATH INTEGRALS AND MANY-BODY SYSTEMS

In 1948, Richard Feynman presented a new formulation of quantum mechanics, called the *path integral formulation*, which reproduced the standard results (such as the Schrödinger equation), but which was far more general and opened up whole new directions in the study of relativistic quantum mechanics and quantum field theory. It would take a paper the length of this one to adequately cover this formulation in a precise or comprehensive way, but the fundamental concept can be summarized very quickly.

Feynman's formulation gives a way to calculate the probability that a quantum system in the configuration C_i at time t_i will be measured to be in the configuration C_f at time t_f. Or rather, it gives the *quantum amplitude* for such a transition, so that the entire time-evolution operator $U(t_i, t_f)$ can be reconstructed. The formulation requires knowledge of a Lagrangian for the system [**?**]. So if we have a Lagrangian for a classical system, we can use path-integrals to quantize it.

The fundamental object used in Feynman's formulation is the *history*. This is a description of a way in which the system might have evolved from a configuration C_i at time t_i to a configuration C_f at time t_f. These histories are not required to be physically realizable in any sense; they are allowed to be extremely fanciful stories. Particles can abruptly change direction mid-flight, objects can venture into regions of enormous potential without any hesitation, and the system's energy can fluctuate arbitrarily. The most that is required is that the system evolve smoothly enough during the history that we can compute its Lagrangian. Our configuration-space model gives us an easy way to visualize such histories. They are no more than paths in configuration space connecting C_i to C_f. Figure **??** shows two such paths.

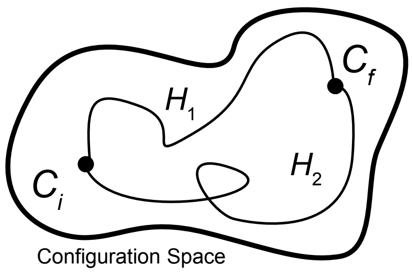

FIG. 2: Two histories, H_1 and H_2, each connecting the configuration C_i to C_f.

Feynman's method proceeds by assigning to each such history H an phase

$$A(H) = e^{iS(H)/\hbar}, \tag{1}$$

where $S(H)$ is the classical action of the history, computed by integrating the Lagrangian along the history. Once we know what the amplitude of each history is, we can find the amplitude of the transition by summing up the amplitudes for all possible histories [**?**].

In most situations, the path integral formulation presented above successfully takes a Langrangian and turns it into a description of how a quantum system will evolve over time. But there is a hole in this mechanism. To understand it, we will have to understand the structure of *history space*. In this space, each point represents an entire history connecting C_i to C_f. Small movements through this space correspond to small variations of whole histories. The Lagrangian action assigns a real number to every point in history space, and the classical principle of stationary action says that any classically realizable history is a stationary point for this action function.

But being a stationary point is a *local* property which depends only on the derivatives at the point of concern. Any change in the action which does not affect these derivatives is a non-physical change, at least as far as classical physics is concerned.

The most basic thing you can do to change a function while keeping its derivatives is to add a constant. It is well known that adding a constant to an action does not change any of the resulting classical physics. In our path-integral formulation, this added constant would result in a constant phase shift in our amplitudes. This itself is not a problem; global phase shifts are as non-physical as constant action shifts. But there is a more subtle possibility.

If the space of histories splits apart into *disconnected pieces*, one can add a constant to the action not across the entire space, but just inside a single piece. This change will not affect any derivatives, because no derivatives depend on more than one disconnected piece of the history space. Classically, the action is only used to compare histories which lie close to one another, related by a small variation. If two histories are not connected by any succession of small variations, the values of their actions need not have any relationship [**?**].

However, in the path integral formulation, we are adding up amplitudes from *all* histories, connected or not. If we use our the "gauge-invariance" of the classical Lagrangian to manipulate the value of the action on these disconnected components, we will get varying results from our path integrals. Therefore, when a space of histories is disconnected, the path integral formulation as described above does not suffice to give well-defined results.

This is exactly the situation we run into with many-particle systems. Suppose we have two particles in a three-dimensional space. Denote by C the configuration in which the particles are located at $\vec{x}_1$ and $\vec{x}_2$. We want to know the amplitude associated to the transition from C at time t_i to C at time t_f. So we look at histories: ways of moving the two particles around so that at time t_i they are positioned at $\vec{x}_1$ and $\vec{x}_2$, and at time t_f they are located at the same two points. One such history will have the two particles sitting where they are for the full time interval without moving. Many more will have them moving around in some complicated pattern, before they each return back to their respective starting positions. But there will also be a whole class of histories which switch the two particles' positions, so that the particle which started at $\vec{x}_1$ ends at $\vec{x}_2$ and visa versa. And though you can get from any non-switching history to any other non-switching history through continuous deformation, no amount of deformation will turn a non-switching history into a switching history, while keeping itself fixed to C at times t_i and t_f. We can see this using a space-time diagram, like the ones in Figure **??**.

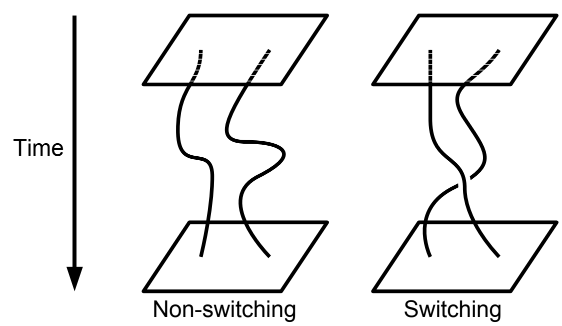

FIG. 3: Space-time diagrams of two histories taking configuration C to configuration C, one which does not switch the particles and one which does.

So our space of histories breaks up into disconnected components, between which there can be no continuous transformation. The classical action alone cannot guide us in comparing the phases of these separate classes. There must be some other factor involved. We will update our amplitude formula to include this unknown component-dependent factor:

$$A(H) = \rho([H])e^{iS(H)/\hbar},$$

where $[H]$ is the component H lies in, and ρ is the function which associates a phase to that component.

IV. THE ORIGIN OF BOSONS AND FERMIONS

Let us take the example addressed above of two particles in three-dimensional space, moving from the configuration C at time t_i to the same configuration C at time t_f. Figure **??** showed two disconnected histories between these configurations, one which left each particle where it started, and the other which swapped their positions. It is not completely obvious, but this non-switching/switching classification serves to classify the mutually deformable classes of histories. That is, any two non-switching histories can be continuously deformed into each other, and any two switching histories can be continuously deformed into each other. This is true even if the particles are not permitted to move through one another. There is just so much room in three dimensions that the only thing keeping one history from being deformed into another is the fundamental obstacle of which particle goes on to which end position.

So we have two components, which we can call [non-switch] and [switch] We want to know what their special phases $\rho([\text{non-switch}])$ and $\rho([\text{switch}])$ are. Based on what has been said so far, we are completely lost. These phases seem to be completely arbitrary. But there is an essential aspect of the path-integral formulation which has not been mentioned yet. This is the principle of concatenation. If we have a history H_1 connecting C_1 to C_2 and a second history H_2 connecting C_2 to C_3, we can glue them together to form a history $H_1 \circ H_2$ connecting C_1 to C_3. In order for the path-integral formulation to give time-evolution operators which make any sense, we need the amplitudes of H_1 and H_2 to multiply to the amplitude of $H_1 \circ H_2$. That is, we have the rule

$$A(H_1 \circ H_2) = A(H_1)A(H_2).$$

Figure **??** shows the action of concatenation in terms of space-time diagrams.

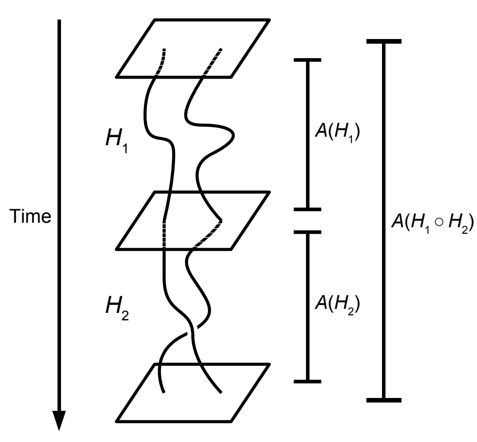

FIG. 4: Concatenation of two histories corresponds to stacking their space-time diagrams in time. The amplitudes corresponding to the histories should multiply: $A(H_1 \circ H_2) = A(H_1)A(H_2)$.

So we should look at how the non-switching and switching histories concatenate with one another. We obtain the following "multiplication table":

$\circ$	[non-switch]	[switch]
[non-switch]	[non-switch]	[switch]
[switch]	[switch]	[non-switch]

We have to find phases $\rho([H])$ which agree with this multiplication table. The relationship $[\text{non-switch}] \circ [\text{non-switch}] = [\text{non-switch}]$ gives

$$\rho([\text{non-switch}])^2 = \rho([\text{non-switch}]),$$

which implies $\rho([\text{non-switch}]) = 1$. Once we have this, $[\text{switch}] \circ [\text{switch}] = [\text{non-switch}]$ gives us $\rho([\text{switch}])^2 = \rho([\text{non-switch}]) = 1$. So $\rho([\text{switch}])$, the phase associated with switching the two particles, can have the values $+1$ or -1. This should sound familiar. We have rederived the possible phase-factors for swapping two particles, except now they correspond to the amplitudes of histories in a Feynman path integral, rather than the values of a many-particle wavefunction. We will call particles with $\rho([\text{switch}]) = +1$ bosons, and those with $\rho([\text{switch}]) = -1$ fermions.

In general, if we have N particles, the disconnected classes of histories will correspond to the possible ways to permute these N particles. These permutations, along with their method of composition, yield a structure called the *symmetric group*, which mathematicians denote by S_N. The question then arises as to how to associate a complex phase to every permutation in S_N in a coherent way, so that concatenation of permutations corresponds to multiplication of phases. This is a well-studied problem, which can be mathematically stated as the question of finding the one-dimensional representations of S_N. (The "one-dimensional" here refers to the fact that we are associating complex numbers to permutations, rather than higher-order matrices.) The answer to this question is that there are precisely two such representations:

1. The trivial representation, which associates the phase $+1$ to every permutation.

2. The alternating representation, which associates the phase $+1$ to any permutation built out of an even number of transpositions, and -1 to any permutation built out of an odd number of transpositions.

This means that our binary classification for the pair of particles generalizes to any number of particles. In three dimensions, we have derived the general theory of bosons and fermions.

V. ANYONS

We have seen that in three dimensions, the space of histories breaks up into pieces corresponding to different ways to permute the particles. The difference between bosons and fermions is the way they associate phases to these permutations. Perhaps, then, if we can get the space of histories to break up in a different way, we can find new possibilities for many-body statistics. This is exactly what happens when we confine particles to a two-dimensional space.

Let us take a pair of indistinguishable particles in two dimensions which are not allowed to coincide. There are two basic ways we can swap the positions of these particles: counterclockwise and clockwise. This gives rise to two different histories, which we call H_1 and H_2. They are illustrated in Figure **??** using space-time diagrams.

If we were in three dimensions, we could deform these histories into one another continuously. To see how to do this, note the diagrams on the top of Figure **??**. The arrows in these diagrams depict how the two particles move in order to take each other's place. In three dimensions, we can connect the first diagram to the second by a continuous sequence of such diagrams, by rotating the diagram around the axis connecting the two particles. This is an instance of the general fact mentioned earlier, that in three dimensions any two histories which switch a pair of particles can be connected by a continuous transition.

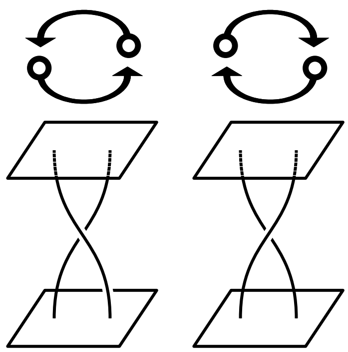

FIG. 5: Space-time diagrams of two different ways to swap a pair of particles in the plane.

But in two dimensions, a rotation of the diagram into the third dimension is not possible. Is there some other way we can connect these histories continuously? The space-time diagrams in Figure **??** suggest that the answer is *no*. This is because a deformation of the history on the left into the history on the right corresponds to a continuous deformation of the strand configuration on the left into the strand configuration on the right. It is easy to see that such a deformation would involve making the two strands pass through each other. This is not permitted, since it means that the particles coincide at some moment in time.

So if our particles live in two dimensions, the connectivity of histories to one another is determined by more than just the way they permute the particles. There is a richer structure which relates to the way the particles wind around one another during the trajectory of the history.

In particular, as the analysis in terms of the space-time diagrams suggests, we can specify the classes of continuously connected histories by looking at how many times the two particles circle around each other in some given direction, say, counterclockwise. To be more precise: The disconnected piece of history space a given history falls into is determined by the number of counterclockwise half-twists in its space-time diagram. Figure **??** shows representative histories in the classes corresponding to -1, 0, 1, and 2 half-twists. In general, for every integer n we have a class $[n]$ containing the histories in which the particles wind around each other $n/2$ times.

The statistics of our pair of plane-bound particles will be determined by the rule which assigns complex phases to these classes. As before, we are not free to give whatever phases we want to the classes. We are bound by the way the histories in the classes concatenate. It is easy to see that if a

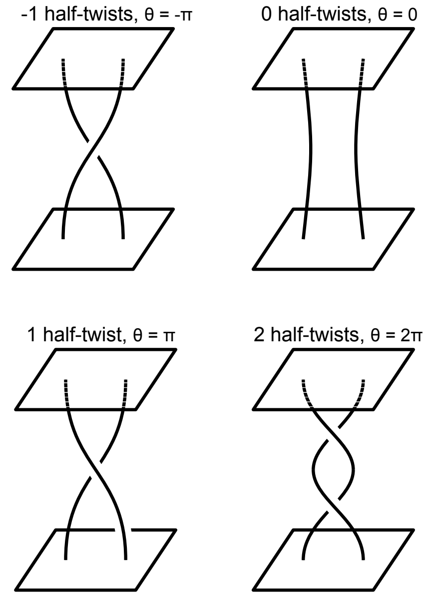

FIG. 6: Ways in which two particles can spin around one another, specifying histories which are disconnected from one another.

history which involves m half-twists is followed by a second history which involves n half-twists, the concatenation will have $m+n$ half-twists. Symbolically, $[m]\circ[n]=[m+n]$. In order for our phase assignments to be coherent with this law of composition, we must have $\rho([m])\rho([n])=\rho([m+n])$.

It turns out this rule is almost all we need to constrain our phase assignments. First, $\rho([0])\rho([0])=\rho([0])$ implies that $\rho([0])=1$. Just like [non-switch] was given the trivial phase 1 in the three-dimensional two-particle case, the most straightforward history here, [0], must be given the phase 1. Next, we can write the phase for any positive n in terms of the phase for 1:

$$\rho([n])=\rho([1+\cdots+1])=\rho([1])\cdots\rho([1])=\rho([1])^n.$$

Since $\rho([1])\rho([-1])=\rho([0])=1$ gives $\rho([-1])=\rho([1])^{-1}$, we can extend this rule to negative n. Altogether, we have

$$\rho([n])=\rho([1])^n,\text{ for all } n. \tag{2}$$

Therefore, the rule for assigning phases is entirely determined by which phase we assign to the single half-twist.

What restrictions must we place on $\rho([1])$? It must be a phase, of course, but are there any others? The class [1] is similar to the class [switch] that we considered when we were looking at three-dimensional space. Back then, we found that the relation

$$[\text{switch}]\circ[\text{switch}]=[\text{non-switch}]$$

led to the condition $\rho([\text{switch}])^2=1$, which restricted $\rho([\text{switch})$ to the values ±1. But now, in two-dimensional space, we have no such relations, since $[1]\circ[1]=[2]$, $[1]\circ[1]\circ[1]=[3]$, and so on. As long as never wrap back to [0], we won't be able to find any relations which restrict the possible values of $\rho([1])$. We are led to the conclusion that $\rho([1])$ is in fact arbitrary. We can say that

$$\rho([1])=e^{i\theta},$$

where θ is an arbitrary real constant. Then, using (**??**), we will have the general rule

$$\rho([n]) = e^{i\theta n}.$$

When $\theta = 0$, we have $\rho([n]) = 1$ for all n, so swapping particles adds no extra phase to any histories. These are bosons. When $\theta = \pi$, we have $\rho([n]) = (-1)^n$, so that a history which has n odd and which thus swaps the particles gets a phase of -1. These are fermions. But we now have the power to choose any θ from 0 to 2π. Bosons and fermions have become just two particular possibilities out of a veritable infinitude, as portrayed in Figure **??**. Since we are free to choose *any* phase, we call our particles *anyons*.

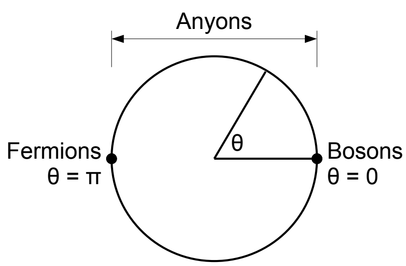

FIG. 7: The circle of possible half-twist phases, which connects bosons to fermions with a continuous range of anyons.

This treatment for two particles remains essentially the same for larger numbers of particles, except that the collection of disconnected pieces of history space grow more complex. Instead of having two particles which just twist around each other a certain number of times, we have a complex interweaving of a larger number of particles. One possibility for three particles is shown in Figure **??**.

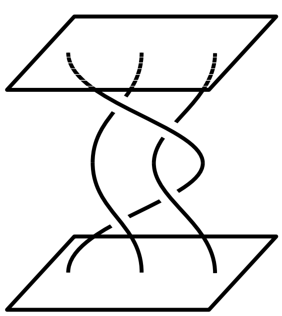

FIG. 8: A history involving three particles in two dimensions creates a triple braid in space-time.

These interweavings in space-time are, quite reasonably, called *braids*. To classify the possible statistics of many particles in two dimensions, we must concern ourselves with how these braids concatenate with one another. That is, we consider the *braid group*, which encodes the possible braids along with their method of composition. Just as quantum statistics in three dimensions came down to understanding the one-dimensional representations of the symmetry group S_N (there were only two), quantum statistics in two dimensions reduces to understanding the one-dimensional representations of the braid group B_N.

This is a complex task. Somewhat surprisingly, the final answer is much the same as in the two-particle case. A representation (that is, a rule for assigning phases to pieces of history space) is specified by a single complex phase which determines the phase given to a half-twist of any pair of particles. So our analysis extends more generally.

VI. ANGULAR MOMENTUM AND THE SPIN-STATISTICS CONNECTION

It is well known that there is a profound connection between angular momentum and quantum statistics. In particular, the spin angular momentum of a particle in three dimensions is quantized; it can only take on values of the form $\hbar\sqrt{s(s+1)}$ where s is a (non-negative) integer or half-integer. Particles with integer s are always bosons, and those with half-integer s are always fermions. This connection comes from quantum field theory, so we will not discuss it in too much depth here. However, our strange new result about particles living in two dimensions raises an important question: What strange feature of two-dimensional angular momentum is responsible for the unexpected quantum statistics found on the plane? The answer is simple. Spin in two dimensions is not quantized.

Spin in three dimensions is quantized essentially because of the structure of the group of rotations in three-dimensional space. This group is non-abelian: it contains members which do not commute with one another. This is familiar fact: the orientation of a book after two $\pi/2$-rotations about different axes depends on the order the rotations are performed in. Since the angular momentum operators generate the rotations, we obtain nontrivial commutation relations between the angular momentum operators. These commutators are the root cause of the quantization of angular momentum.

But in two dimensions, rotations become much simpler. A rotation in the plane can be characterized by a single angle θ, and any two such rotations commute with one another. So we do not have quantization of angular momentum in two dimensions, and the spin of a two-dimensional particle is free to vary continuously.

This spectrum corresponds exactly to the continuous spectrum of possible statistics we have for two-dimensional particles. Varying our anyons' spin number s from 0, for bosons, to $\frac{1}{2}$, for fermions, and up to 1, for bosons again, will cause our statistics-determining phase $e^{i\theta}$ to vary from $e^{i0} = 1$, to $e^{i\pi} = -1$, and then back to $e^{i2\pi} = 1$, going around the unit circle. To actually prove this correspondence would require results from quantum field theory beyond the scope of this paper, but we can at least glimpse the parallels. Angular momentum gives us a new way to understand the importance of dimension.

VII. PHYSICAL IMPLEMENTATION: THE AHARANOV-BOHM EFFECT

Up to this point, the possible existence of anyons has been justified in an entirely theoretical way, starting from the foundations of quantum theory. The question has remained as to whether they actually exist. Maybe anyons are just too strange to come about in our universe. For instance, as one anyon circles around another, the two anyons pick up a phase from this circling, no matter how far away they are from one another. This weird non-locality invites skepticism. Furthermore, winding one way introduces a different phase than winding the other way, even though our typical assumption of parity symmetry would forbid this.

But physicists have indeed found phenomena which behave this way. In particular, the acquisition of a phase through a circling action is reminiscent of the Aharonov-Bohm effect. Consider an ideal solenoid, which has a magnetic flux of Φ in its interior, without creating any magnetic field outside of itself [**?**]. Aharonov and Bohm showed that when a particle with charge q moves in a circle around this solenoid, it will acquire a phase of

$$\Delta\phi = \frac{q\Phi}{\hbar}. \tag{3}$$

This result is truly surprising! The charged particle has somehow felt the effect of the tube of flux, even though it has only moved through regions without any electromagnetic fields. There are various resolutions to this paradox, the most straightforward of which is the realization that there is still a magnetic vector potential in the exterior of the solenoid, even though there is no magnetic field, and this vector potential somehow has the ability to carry the influence of the flux tube.

For our purposes, we are interested in the fact that the phase in (**??**) is quite arbitrary. Continuously varying the flux Φ will allow us to obtain any value for $\Delta\phi$. So the Aharonov-Bohm effect has the potential to give us a way to make anyons.

Of course, the effect as described above is quite asymmetric. In the Aharanov-Bohm effect, we speak of an electric charge circling a tube of magnetic flux, and these are two very different sorts of objects. We need to have a system of indistinguishable particles in order to have true anyonic statistics. The simplest solution to this is to form a new composite particle which bundles together a charge and a tube of flux. We quite plainly call these particles "charge-flux-tube composites" [**?**]. As two of these speculative particles spin around one another, each of their charges will interact with the others' flux tube, giving the whole system the phase shift we expect of a system of anyons.

So, as long as we have a two-dimensional system with quasiparticles which act like both charges and fluxes, we will have fractional quantum statistics. It has been proposed that such a configuration arises in the fractional quantum hall effect. The two-dimensional system here is a interface between two semiconductors. Exitations of electrons trapped at this heterojunction act as anyons, and it turns out that this gives them the fractional charge which is responsible for the fractional quantum Hall effect.

As a final connection between the theory of anyons and the "real world", I would like to mention a possible *application* of fractional quantum statistics: topological quantum computation. A quantum computer is a theoretical device which exploits strange quantum phenomena, such as such as superposition, entanglement, and interference, in order to perform computations. There is good evidence to suggest that a computer capable of manipulating quantum states could be much more powerful than a classical computer [**?**].

However, there are many obstacles standing in the way of physically implementing a general quantum computer. Chief among them is the problem of decoherence. Quantum states are generally very delicate, and small perturbations from an outside environment can completely ruin the superposed structures which are necessary for non-trivial quantum computation.

Anyonic statistics may provide a solution to this problem. If we have anyons with internal structure (such as spin), we can store information in this structure. Then we can design a system in which circling the anyons around one another in a certain braid pattern performs some useful computation. The benefit to this system is that the exact positions of the anyons do not matter; all that matters is the braid they create as they twirl. Small amounts of noise will jiggle the anyons, but these jiggles will only have computational effects if they make one anyon go completely around another. This possibility can be reduced just by stationing the anyons far enough away from each other. Therefore, an anyonic quantum computer could be extraordinarily robust. Fractional quantum statistics may very well lead to a revolution in computation.

[] A. Khare, *Fractional Statistics and Quantum Theory* (World Scientific Publishing Company, 2005)

[] F. Wilczek, *Fractional Statistics and Anyon Superconductivity* (World Scientific Publishing Company, 1990)

[] J. M. Leinaas and J. Myrheim, "On the Theory of Identical Particles", *Il Nuovo Cimento* **37** (1977) 1-23

[] A. Lerda, *Anyons: Quantum Mechanics of Particles with Fractional Statistics* (Springer-Verlag, 1992)

[] X.-G. Wen, *Introduction to Quantum Many-Boson Theory* (unpublished notes: `http://dao.mit.edu/8.08/chintr-bsn.pdf`)

[] If you haven't seen a Lagrangian before, it suffices to say here that it is a function which assigns a real number to every state of a system, and which in certain situations can be taken to be $T - V$, the kinetic energy minus the potential energy.

[] Here, "summing up" actually refers to integrating over the highly infinitely-dimensional space of paths, but for our purposes that is a technical detail.

[] One would be right to object that this is impossible: the field lines at the top of the solenoid must join up with those at the bottom through a path outside of the solenoid, creating fringe effects. However, we are free to make the solenoid as long as we want, which will move these returning field lines as far away from our domain of consideration as we want. Fringe effects can be made arbitrarily small, and they are certainly not involved in causing the Aharonov-Bohm effect.

[] Charge-flux-tube composites are also referred to as "composite fermions".

[] For example, Peter Shor found a quantum algorithm which can factor integers into prime factors in a very efficient way, something which is believed to be impossible with non-quantum computers.