Multiparticle spaces in quantum mechanics
Part 1: Indistinguishability

Santiago Quintero de los Ríos
for homotopico.com

October 11, 2018

Suppose you already know how to treat one-particle quantum-mechanical systems, and now you wish to go big, go statistical, or maybe go to fields. It is therefore necessary to understand systems with multiple (identical) particles, and how to represent them mathematically.

The fundamental hypothesis here is that **particles of the same kind are indistinguishable**. The question of how reasonable this hypothesis is is "a good question... for another time", so very much like The Force Awakens' writers, I will put a shade on that lamp, ignore the issue, and move on [1]. In the end, we should not say that "particle a is in state x and particle b is in state y"; at most we should say "there are two particles, one in state x and one in state y". How do we go about describing these systems mathematically, given that we already know how to describe a one-particle system? It seems cumbersome to try and work out how to mathematically write the second statement, whereas the first one seems simpler if we already know how to state that "a particle is in state x" and "a particle is in state y".

Then what we will do is write "particle a is in state x and particle b is in state y" mathematically, and *then* find out how to turn this statement into the indistinguishable case "there are two particles, one in state x and one in state y". We will see that the indistinguishable case can be thought of as a "distinguishable case" but with *permutation symmetry*.

This means that we first have to discuss a little bit what a symmetry is, and then apply this to the permutation symmetry of multiparticle states. We will obtain that there are **two types** of indistinguishability, one of which naturally takes us to Pauli's exclusion principle.

Symmetries.zip

Consider an observable A of a system S (this system does not even have to be quantum). Suppose that we have two states ψ_1, ψ_2 of this system that are *identical* if we *only* compare how they "look" with the A observable (for example, if the observable A is "distance to a point o", then all the points in a circle centered at o are identical for anyone who can only check the A observable). What we mean is that the probability distribution of A in states ψ_1 and ψ_2 are the same. Under this observable, the states ψ_1 and ψ_2 are identical. In this case we say that ψ_1 and ψ_2 are A-equivalent or A-symmetric (or symmetric with respect to A). There might be some other observable B which can distinguish between ψ_1 and ψ_2, or otherwise the states ψ_1 and ψ_2 will, by definition, be exactly the same.

In the context of quantum mechanics, we say that two states $|\psi_1\rangle, |\psi_2\rangle$ are A-symmetric if the probability distribution of the observable A is the same for both. That is, they are A-symmetric if and only if for all possible values a of A, the cumulative distributions are

$$P\left(A \leq a|\psi_1\right) = P\left(A \leq a|\psi_2\right).$$

[1] Maybe the hypothesis is justified in that otherwise the standard definition of entropy would not be extensive (Gibbs' paradox)... But maybe that is a problem with the definition of entropy, not the distinguishability (oof long word!) of particles. Again, I'm not opening that can of worms. Yet. Stay tuned!

Recall then that the probability distribution of an observable A is given by the eigenvalues of an associated Hermitian operator $\hat{A}$. Namely, if $|a_1\rangle, |a_2\rangle, \ldots$ are orthonormal eigenstates with eigenvalues $a_1, a_2, \ldots,$ then the probability distribution of A given a state ψ is given by[2]

$$P(A = a_j|\psi) = \left|\langle a_j|\psi\rangle\right|^2.$$

This implies that two states $|\psi_1\rangle, |\psi_2\rangle$ are A-symmetric if and only if for all eigenvalues a_j of $\hat{A}$,

$$\left|\langle a_j|\psi_1\rangle\right|^2 = \left|\langle a_j|\psi_2\rangle\right|^2.$$

We say that a mapping $T : \mathcal{H} \to \mathcal{H}$ is A-symmetric (or an A-symmetry) if for all $|\psi\rangle \in \mathcal{H}$, the states $|\psi\rangle$ and $|\psi'\rangle = T|\psi\rangle$ are A-symmetric. This means that for all eigenvalues a_j of A,

$$\left|\langle a_j|T|\psi\rangle\right|^2 = \left|\langle a_j|\psi\rangle\right|^2.$$

This immediately implies that the expectation value of A is the same for $|\psi\rangle$ and $|\psi'\rangle$ (as can be seen by inserting some well-placed I using the completeness relation for the $|a_j\rangle$ eigenstates).

In particular, if T is an A-symmetry, then for all $|\psi\rangle$ and $|\psi'\rangle = T|\psi\rangle$ we have that

$$\begin{aligned}\|\psi'\|^2 = \left|\langle\psi'|\psi'\rangle\right|^2 &= \sum_j \left|\langle a_j|\psi'\rangle\right|^2 \\ &= \sum_j \left|\langle a_j|\psi\rangle\right|^2 \\ &= |\langle\psi|\psi\rangle|^2 = \|\psi\|^2.\end{aligned}$$

And thus, T is *norm*-preserving. Unfortunately, this fact alone does not say much about T. However, if we also require T to be invertible and linear, then it must be a unitary transformation, that is, $T^\dagger T = TT^\dagger = I$, or equivalently $\langle(T\psi)|T\phi\rangle = \langle\psi|\phi\rangle$. This follows from a straight-forward application of the polarization identity

$$\langle u|v\rangle = \frac{1}{4}\left(\|u+v\|^2 - \|u-v\|^2 + i\|u-iv\|^2 - i\|u+iv\|^2\right).$$

Similarly, if we require T to be invertible and additive but *anti*linear, $T(au+v) = a^*T(u) + T(v)$, then if it preserves norms it must be antiunitary, i.e. $\langle(T\psi)|T\phi\rangle = \langle\phi|\psi\rangle$.

These results are quite similar to Wigner's theorem, which states that any invertible mapping T that preserves the *norms* of inner products and that maps *rays* into rays (i.e. if $v = \alpha u$ for some scalar α then there exists a scalar β such that $T(v) = \beta T(u)$), (and whose inverse does too) must be unitary or antiunitary.

What we have shown is that any *linear* A-symmetry must be a unitary transformation. In particular, if $\hat{U}$ is unitary A-symmetry, then the $\hat{U}$-conjugate of $\hat{A}$ satisfies

$$\begin{aligned}\left\langle\psi\middle|\hat{U}^{-1}\hat{A}\hat{U}\middle|\psi\right\rangle &= \left\langle\psi\middle|\hat{U}^\dagger\hat{A}\hat{U}\middle|\psi\right\rangle \\ &= \left\langle(\hat{U}\psi)\middle|\hat{A}\middle|(\hat{U}\psi)\right\rangle \\ &= \left\langle\psi\middle|\hat{A}\middle|\psi\right\rangle.\end{aligned}$$

This is because $\hat{U}^\dagger = \hat{U}^{-1}$, and since $\hat{U}$ is an A-symmetry, it preserves the expectation value of A. This holds for any $|\psi\rangle \in \mathcal{H}$, therefore $\hat{U}^{-1}\hat{A}\hat{U} = \hat{A}$, which means that $\hat{A}\hat{U} = \hat{U}\hat{A}$, or equivalently, $[\hat{A}, \hat{U}] = 0$.

There are different levels on which we can describe symmetry. The first one we discussed is determined *pointwise*, that is, only two states being symmetric with respect to an observable. Then, one step further, we

[2]Of course, there is the subtlety of the case where $\hat{A}$ has a continuous spectrum without eigenstates, but this can be formally extended as probability densities. I think. Don't think about it too much.

described transformations T that are also A-symmetries. Under these transformations, every pair $|\psi\rangle, T|\psi\rangle$ are A-symmetric. If this operator is to be linear, then it must also be unitary. Most often, in the standard QM references, the symmetries that are considered are unitary transformations that are symmetric with respect to the Hamiltonian (energy) observable.

We can go one step further and consider symmetries with respect to *all* observables, which we call **absolute symmetries**[3]. Any two states that are related by these symmetries are *absolutely indistinguishable*. What this means is that if T is an absolute symmetry and $|\psi\rangle = T|\phi\rangle$, then $|\psi\rangle$ and $|\phi\rangle$ have the exact same distributions for *all observables*. This means that they are indistinguishable![4]

Of course, if our system allows for these kinds of absolute symmetries, then *our state space is defined with redundancies*, since we *should* define the states and the observables in such a way that two states are equal if and only if they have the same distribution for all observables. In truth, absolute symmetries arise from being sloppy in the definition of states.

We will now attempt to characterize absolute symmetries. So let T be an absolute symmetry. Consider any state $|\psi\rangle$, and complete it to an orthonormal basis $|\psi\rangle = |\psi_1\rangle, |\psi_2\rangle, |\psi_3\rangle, \ldots$. Since T is an absolute symmetry, for all Hermitian $\hat{A}$, we have that $\left\langle (T\psi) \middle| \hat{A} \middle| (T\psi) \right\rangle = \left\langle \psi \middle| \hat{A} \middle| \psi \right\rangle$. In particular, we consider the projection operators with respect to the elements of the basis,

$$P_{\psi_j} = |\psi_j\rangle\langle\psi_j| .$$

Then we have that

$$\begin{aligned} \left\langle (T\psi_i) \middle| P_{\psi_j} \middle| (T\psi_i) \right\rangle &= \left\langle (T\psi_i) \middle| \psi_j \right\rangle \left\langle \psi_j \middle| T\psi_i \right\rangle \\ &= \left| \left\langle \psi_j \middle| T\psi_i \right\rangle \right|^2 . \end{aligned}$$

On the other hand,

$$\begin{aligned} \left\langle (T\psi_i) \middle| P_{\psi_j} \middle| (T\psi_i) \right\rangle &= \left\langle \psi_i \middle| P_{\psi_j} \middle| \psi_i \right\rangle \\ &= \left\langle \psi_i \middle| \psi_j \right\rangle \left\langle \psi_j \middle| \psi_i \right\rangle \\ &= \left| \left\langle \psi_i \middle| \psi_j \right\rangle \right|^2 \\ &= \delta_{ij} . \end{aligned}$$

This means that

$$\left| \left\langle \psi_j \middle| T\psi_i \right\rangle \right|^2 = \delta_{ij} ,$$

so that, say, $\left\langle \psi_j \middle| T\psi_j \right\rangle = e^{i\theta_j}$. From here we obtain that

$$T|\psi_j\rangle = \sum_k \left\langle \psi_k \middle| T\psi_j \right\rangle |\psi_k\rangle = e^{i\theta_j} |\psi_j\rangle .$$

Now we don't know that T is linear or antilinear, but it certainly does preserve rays. For any given state $|\varphi\rangle$ we repeat the previous process and obtain that there exists a unitary scalar α_φ such that $T|\varphi\rangle = \alpha_\varphi|\varphi\rangle$. Then if $|\psi\rangle = \beta|\phi\rangle$, we have that

$$T|\psi\rangle = \alpha_\psi|\psi\rangle = \alpha_\psi\beta|\phi\rangle = \alpha_\psi\beta\alpha_\phi^{-1}T|\phi\rangle .$$

Then T maps rays into rays, and it preserves the norms of inner products,

$$|\langle (T\psi)|(T\phi)\rangle| = |\alpha_\psi^*\alpha_\phi \langle\psi|\phi\rangle| = |\langle\psi|\phi\rangle|,$$

and thus it satisfies the hypotheses of Wigner's theorem, so T must be unitary or antiunitary. Suppose that it is unitary, so it preserves inner products. Then for any $|\psi\rangle, |\phi\rangle$, it follows that

$$\langle\psi|\phi\rangle = \langle (T\psi)|T\phi\rangle = \alpha_\psi^*\alpha_\phi \langle\psi|\phi\rangle .$$

[3] Don't quote me on this, I just made this up.

[4] If this is not convincing, then ask yourself "how you I tell two such states apart?" There is no observable that can distinguish between them.

Then $\alpha_\psi^* \alpha_\phi = 1$. However, since $|\alpha_\psi| = 1$, then $\alpha_\psi^{-1} = \alpha_\psi^*$, and this implies that $\alpha_\psi^{-1}\alpha_\phi = 1$, or rather, $\alpha_\psi = \alpha_\phi$. In conclusion, if T is a unitary absolute symmetry, then is is multiplication by a unit scalar:

$$T\left|\psi\right\rangle = \alpha\left|\psi\right\rangle \quad \forall\left|\psi\right\rangle.$$

And we *know* that this is a redundancy in our description of the states! In truth, we should work in the *projective* Hilbert space $P\mathcal{H}$, in which we assume that vectors which are related by multiplication by any unitary scalar represent *one and the same state*.

Describing multiparticle states

Suppose that we have two systems S_1, S_2, described by Hilbert spaces $\mathcal{H}_1, \mathcal{H}_2$, respectively. The mathematical setting for describing the joint system is the [completion] $\mathcal{H}$ of the **tensor product**[5] $\mathcal{H}_1 \otimes \mathcal{H}_2$. The inner product in $\mathcal{H}$ is term-wise as

$$\langle u_1 \otimes u_2 | v_1 \otimes v_2 \rangle := \langle u_1 | v_1 \rangle \langle u_2 | v_2 \rangle.$$

There are a few different notations for product elements of $\mathcal{H}$, but we will mostly stick to $|u_1 u_2\rangle := |u_1\rangle \otimes |u_2\rangle$. Note that not every state (actually *very few* states) in $\mathcal{H}$ can be written as a tensor product of two states in $\mathcal{H}_1$ and $\mathcal{H}_2$. A general state is a *linear* combination of states of the form $|u_1, u_2\rangle$, with $u_i \in \mathcal{H}_i$, and if $\{|\varphi_i\rangle\}, \{|\psi_j\rangle\}$ are bases for $\mathcal{H}_1$ and $\mathcal{H}_2$, then $\{|\varphi_i \psi_j\rangle\}$ is a basis for $\mathcal{H}$.

Any observable $\hat{A}$ over the space $\mathcal{H}_1$ can be "extended" to an observable over $\mathcal{H}$, by identifying it with $\hat{A} \otimes \mathrm{id}_{\mathcal{H}_2}$. Similarly for observables over $\mathcal{H}_2$.

I will not go into all the nuances and interesting details of these kind of states (e.g. entangled states and such), since what we are interested in is representing indistinguishability.

Indistinguishability defined

For the time-being, we will only work with 2-particle spaces, and then we will go nuts and go to general k-particle spaces. Consider the case where $\mathcal{H}_1 = \mathcal{H}_2 := \mathcal{H}$, in which case the total Hilbert space is $\mathcal{H} \otimes \mathcal{H}$. By definition, the states $|\psi_1\psi_2\rangle$ and $|\psi_2\psi_1\rangle$ are considered *different*. However, we want them to *represent the same state*. There are two ways to go about this:

1. Scrap everything and rethink a suitable Hilbert space that clearly represents indistinguishable particles without any redundancy, or

2. force any two states that represent the same configuration (e.g. as above) to be *absolutely symmetric*.

If you're all the way down here, you'll be glad to read that we will not choose, I repeat, not choose option 1. Even more, working with option 2 will actually *suggest* which Hilbert space to use that would work for option 1. It's a win-win.

Then we choose option 2. Define a transformation T[6] that acts on product terms as

$$T\left|\psi_1\psi_2\right\rangle = \left|\psi_2\psi_1\right\rangle,$$

and extend it by linearity. This is a *permutation* transformation. Since we require T to be a (unitary) absolute symmetry, then by the result of the previous section T must behave as multiplication by some unit scalar α, that is, $T = \alpha I$. However, we also have that permuting twice should do nothing (if you switch 1 and 2,

[5] The pragmatic justification for this is that this is the simplest non-trivial way to make a Hilbert space out of $\mathcal{H}_1$ and $\mathcal{H}_2$, but I am sure that this can be justified using the spectra of observables of the joint system and a little bit of probability theory. Stay tuned!

[6] This notation is not standard. It is usually called p or P or something like that or maybe something entirely different because nobody wants confusion with the momentum operator.

then switch them again, you get 1 and 2), so $T^2 = I$. But we also have that $T^2 = \alpha^2 I$, so this implies that $\alpha^2 = 1$. Then $\alpha = \pm 1$, that is, $T = \pm I$.

But wait!

This cannot possibly be true for *all* elements of $\mathcal{H} \otimes \mathcal{H}$! As a counterexample, consider *literally any pair of linearly independent elements* $|\psi_1\rangle, |\psi_2\rangle \in \mathcal{H}$. **By definition**, the tensor products $|\psi_1 \psi_2\rangle$ and $|\psi_2 \psi_1\rangle$ are not only *different*, but also *linearly independent*. This means that

$$T |\psi_1 \psi_2\rangle = |\psi_2 \psi_1\rangle \overset{!!!}{\neq} \pm |\psi_1 \psi_2\rangle .$$

So the relation $T = \pm I$ cannot be true for *all* vectors! Did we arrive at a contradiction? Is it *impossible* to represent the state space of two identical particles with permutation symmetry? Well... no. We just showed a counterexample where $T |\psi_1 \psi_2\rangle \neq \pm |\psi_1 \psi_2\rangle$, but there *are* some vectors in $\mathcal{H} \otimes \mathcal{H}$ where permutation *is* a symmetry. For example, consider the vector

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|\psi_1 \psi_2\rangle + |\psi_2 \psi_1\rangle \right) .$$

If we apply T, then we get

$$\begin{aligned} T |\psi\rangle &= \frac{1}{\sqrt{2}} \left(T |\psi_1 \psi_2\rangle + T |\psi_2 \psi_1\rangle \right) \\ &= \frac{1}{\sqrt{2}} \left(|\psi_2 \psi_1\rangle + |\psi_1 \psi_2\rangle \right) \\ &= \frac{1}{\sqrt{2}} \left(|\psi_1 \psi_2\rangle + |\psi_2 \psi_1\rangle \right) \\ &= |\psi\rangle . \end{aligned}$$

Similarly, consider the vector

$$|\psi'\rangle = \frac{1}{\sqrt{2}} \left(|\psi_1 \psi_2\rangle - |\psi_2 \psi_1\rangle \right) .$$

Upon acting with T,

$$\begin{aligned} T |\psi'\rangle &= \frac{1}{\sqrt{2}} \left(T |\psi_1 \psi_2\rangle - T |\psi_2 \psi_1\rangle \right) \\ &= \frac{1}{\sqrt{2}} \left(|\psi_2 \psi_1\rangle - |\psi_1 \psi_2\rangle \right) \\ &= -\frac{1}{\sqrt{2}} \left(|\psi_1 \psi_2\rangle - |\psi_2 \psi_1\rangle \right) \\ &= - |\psi'\rangle . \end{aligned}$$

Then for *some* states, it is true that $T |\psi\rangle = \pm |\psi\rangle$.

What do we do now? We *restrict* our state space to one where there is a permutation symmetry, and we simply *ignore* all the other vectors that do not have this symmetry. So we define the **symmetric product** of $\mathcal{H}$ as the set of all vectors $|\psi\rangle \in \mathcal{H} \otimes \mathcal{H}$ for which $T |\psi\rangle = |\psi\rangle$:

$$\mathcal{S}^2 \mathcal{H} := \{ |\psi\rangle \in \mathcal{H} \otimes \mathcal{H} \ : \ T |\psi\rangle = |\psi\rangle \} .$$

We call this symmetric product the **boson state space**[7] (of two particles). Similarly, we can also define the **antisymmetric product** of $\mathcal{H} \otimes \mathcal{H}$ as

$$\Lambda^2 \mathcal{H} := \{ |\psi\rangle \in \mathcal{H} \otimes \mathcal{H} \ : \ T |\psi\rangle = - |\psi\rangle \} ,$$

[7] "Why "boson"? What does this have to do with spin? Aaaah!" Well, again. A good question... for another time.

and we call this the **fermion state space** in two particles. These *fermionic* states have the particularity that two particles cannot be in the same state! That is, the vectors $|\psi\,\psi\rangle$ (which represent two particles in the same state) are *not* in the fermion state space. This is in stark contrast with the bosonic case, where $|\psi\,\psi\rangle \in \mathcal{S}^2\mathcal{H}$.

The beauty of this is that we are *not* violating any of the axioms of quantum mechanics by restricting ourselves to this space. The boson and fermion state spaces *are* Hilbert spaces (one must simply check that they are closed subspaces of $\mathcal{H} \otimes \mathcal{H}$), and in fact, we can even find a nice, juicy basis for them, given that we already have one for $\mathcal{H}$. That comes next week, pinky promise.

There are (spoiler alert) more than two particles in the universe

This astonishing discovery means that we're going to need a bigger boat. What if we have k particles? The procedure is almost exactly the same as in the previous section, but we have to be a bit more careful. In this case, the ambient space "with redundancies" will be the k-th tensor product of $\mathcal{H}$, which we write as $\otimes^k\mathcal{H}$.

But now permutations get *a lot* more complicated since we can exchange any number of particles. In general a k-permutation is a rearrangement of the numbers $\{1,\ldots,k\}$, which is to say a bijective[8] function $\sigma : \{1,\ldots,k\} \to \{1,\ldots,k\}$. The set of all k-permutations is denoted $\mathfrak{S}_k$.

Our requirement of *symmetry under permutations* can be rewritten as follows: For any permutation $\sigma \in \mathfrak{S}_k$, define a transformation T_σ on product vectors as

$$T_\sigma\,|\psi_1\cdots\psi_k\rangle = \left|\psi_{\sigma(1)}\cdots\psi_{\sigma(k)}\right\rangle,$$

and extend it to the entire space $\otimes^k\mathcal{H}$ by linearity. As an example, let $k = 5$, and consider the permutation

$$\sigma : (1,2,3,4,5) \mapsto (2,4,5,1,3).$$

Then the action of T_σ is

$$T_\sigma\,|\psi_1\psi_2\psi_3\psi_4\psi_5\rangle = \left|\psi_{\sigma(1)}\psi_{\sigma(2)}\psi_{\sigma(3)}\psi_{\sigma(4)}\psi_{\sigma(5)}\right\rangle = |\psi_2\psi_4\psi_5\psi_1\psi_3\rangle$$

(and extended by linearity). Note that $\psi_1,\ldots,\psi_k$ do not have to be different a priori for this to make sense.

Also note that $\sigma \mapsto T_\sigma$ is a linear action, or what we call a **representation** of $\mathfrak{S}_k$ on $\otimes^k\mathcal{H}$. This is because $T_\sigma \circ T_{\sigma'} = T_{(\sigma\circ\sigma')}$ for any pair of permutations $\sigma, \sigma' \in \mathfrak{S}_k$.

Our requirement of *permutation invariance* is that T_σ is an *absolute (unitary) symmetry* **for all permutations** $\sigma \in \mathfrak{S}_k$. This means that for each $\sigma \in \mathfrak{S}_k$, there is a unit scalar α_σ such that $T_\sigma = \alpha_\sigma I$. In particular, let's choose a *transposition*, which simply switches two numbers, $\tau : (1,\ldots,i,\ldots,j,\ldots,k) \mapsto (1,\ldots,j,\ldots,i,\ldots,k)$. We write $(i\ \ j)$ for the transposition that switches i and j. Then we're back in known waters, since transpositions are their own inverses! This means that $T^2_{(i\ j)} = I$, which, again, means that $\alpha_{(i\ j)} = \pm 1$, as in the two-particle case. However, this does not mean that it is 1 for *all* transpositions or -1 for *all* transpositions; it might happen that $\alpha_\tau = 1$ for some transpositions but $\alpha_\tau = -1$ for others. Just because this is getting too long, I will leave as an appendix[9] the proof that, indeed, it should be the same for all transpositions, i.e. $\alpha_{(i\ j)} = 1$ for all transpositions or $\alpha_{(i\ j)} = -1$ for all transpositions.

In order to continue, we now need an important result of group theory, which we will not prove, which states that every permutation can be written as a (non-unique) product of transpositions. Even though this product is not unique, what *is* unique is the *parity* of the number of transpositions needed to express a permutation. That is, a permutation that can be written as an *even* number of transpositions can *only* be written as an even number of transpositions, and similarly, a permutation that can be written as an *odd* number of transpositions can *only* be written as an odd number of transpositions.

[8] Recall that this means that for each i there is one *and only one* number j such that $\sigma(j) = i$. This implies that when you write down $\sigma(1),\ldots,\sigma(k)$, you obtain exactly $1,\ldots,k$ but in disorder.

[9] that is, a post in the future

The **sign** of a permutation σ, denoted $\mathrm{sgn}(\sigma)$, is defined as 1 if it can be written with an even number of transpositions, and -1 if it can be written with an odd number of transpositions. With this definition, we are nearly done. Suppose that $\alpha_{(i\ j)} = 1$ for all transpositions. Then we have, for any permutation σ, that

$$T_\sigma = T_{\tau_1 \cdots \tau_m} = T_{\tau_1} \cdots T_{\tau_m} = 1 \cdots 1 I = I,$$

where $\tau_1, \dots \tau_m$ are transpositions that compose σ. This is what we call the **trivial representation** of $\mathfrak{S}_k$. In this case, we require the canonical action of permutations to be the trivial representation:

$$T_\sigma = I.$$

However, we fall into the same trap as in the previous section. By *definition*, this is simply not true for products of linearly independent vectors, so we must restrict our scope to a subspace of $\otimes^k \mathcal{H}$ where there *is* this kind of permutation invariance. We define the k**-th symmetric product** of $\mathcal{H}$ as

$$\mathcal{S}^k \mathcal{H} := \left\{ |\psi\rangle \in \otimes^k \mathcal{H} \ : \ T_\sigma |\psi\rangle = |\psi\rangle \text{ for all } \sigma \in \mathfrak{S}_k \right\}.$$

And we call it the **boson state space** of k particles. An example of such a state is

$$|\psi\rangle = \frac{1}{\sqrt{3}} \left(|\psi_1 \psi_2 \psi_2\rangle + |\psi_2 \psi_1 \psi_2\rangle + |\psi_2 \psi_2 \psi_1\rangle \right) \in \mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H}.$$

But now we have the other choice, which is that $T_{(i\ j)} = -1$ for all transpositions. In this case, the symmetry is not as simple. If σ is a permutation, then we decompose it into transpositions, $\sigma = \tau_1 \cdots \tau_m$. The action of T_σ is, then

$$T_\sigma = T_{\tau_1 \cdots \tau_m} = T_{\tau_1} \cdots T_{\tau_m} = (-1)^m I = \mathrm{sgn}(\sigma) I.$$

This follows since $(-1)^m$ is precisely 1 if m is even and -1 if m is odd. Then our other choice for permutation symmetry is the requirement that

$$T_\sigma = \mathrm{sgn}(\sigma) I,$$

which we call the **alternating representation** of $\mathfrak{S}_k$. But again! The same trap as before! This relationship does not hold for many vectors in $\otimes^k \mathcal{H}$, so we restrict to a subspace where this relationship does hold, which is the k**-th alternating** (or exterior) **product** of $\mathcal{H}$:

$$\Lambda^k \mathcal{H} := \left\{ |\psi\rangle \in \otimes^k \mathcal{H} \ : \ T_\sigma |\psi\rangle = \mathrm{sgn}(\sigma) |\psi\rangle \text{ for all } \sigma \in \mathfrak{S}_k \right\}.$$

And this space is called the **fermion state space** of k particles. Again, it can be seen that there are no vectors in which there are two particles in the same state (for if there were, a permutation of such two particles would need to yield a -1... which it doesn't). This means that *two fermionic particles cannot be in the same state*, something that is called *Pauli's exclusion principle*.

These spaces will be the setting for doing quantum mechanics of many particles. They seem pretty similar, but in reality, they are extremely different.

Our objective for next week will be writing down an orthonormal basis for these spaces, given a basis for $\mathcal{H}$. Then we will consider states spaces where there can be arbitrarily many particles, and hopefully find a way into field theory.

The takeaway

The assumption of indistinguishability implies that the Hilbert space that represents multiparticle states (particles of the same kind) must not distinguish between "particle a is in state x" and "particle b is in state y". Coming up with such a Hilbert space from scratch is not trivial, so what we do instead is construct a Hilbert space whose elements represent multiparticle states, but also in which the particles are distinguishable. If one particle is represented by the Hilbert space $\mathcal{H}$, then such a "distinguishable" multiparticle space is $\otimes^k \mathcal{H}$. Then, in order to pass to indistinguishability, we considered *subspaces* of $\otimes^k \mathcal{H}$ in which *permutation is an absolute symmetry*. There are only two of those, namely the *bosonic space* which is the symmetric product $\mathcal{S}^k \mathcal{H}$, and the *fermionic space* which is the antisymmetric product $\Lambda^k \mathcal{H}$. In particular, the fermionic space does *not* have any elements in which two particles are in the same state. This is Pauli's Exclusion Principle.

References

- Folland, G. B. (2008). *Quantum Field Theory: A Tourist Guide for Mathematicians*, Chapters 3 and 4.
- Sakurai, J. J. (1995). *Modern Quantum Mechanics, Revised Edition*, Chapters 1, 4, and 6.
- Beck, M. (2012). *Quantum Mechanics: Theory and Experiment*, Chapter 8.
- Weinberg, S. (2002). *The Quantum Theory of Fields I: Foundations*, Chapter 2, Appendix A.

Appendix: T_σ is the same for all σ

The quick way to do this, which requires a bit of machinery from algebra, is as follows (a more down-to-earth proof, I think, will follow soon). Note that the representation $T : \mathfrak{S}_k \to \{I, -I\} \cong \mathbb{Z}_2$ is actually a one-dimensional (since it maps to multiples of the identity), and therefore $\mathrm{im}(T)$ is abelian. This means that T must factorize through the abelianization $\mathfrak{S}_k^{ab} = \mathfrak{S}_k/[\mathfrak{S}_k, \mathfrak{S}_k]$, where $[\mathfrak{S}_k, \mathfrak{S}_k]$ is generated by the commutators $ghg^{-1}h^{-1}$. However, the map $\mathrm{sgn} : \mathfrak{S}_k \to \{1, -1\} \cong \mathbb{Z}_2$ is surjective whose kernel is precisely[10] $[\mathfrak{S}_k, \mathfrak{S}_k]$, so $\mathfrak{S}_k^{ab} \cong \mathbb{Z}_2$. This means that $T = \tilde{T} \circ \pi$, where $\tilde{T} : \mathbb{Z}_2 \to \mathbb{Z}_2$ is a homomorphism and $\pi = \mathrm{sgn}$ is the projection to the abelianization. There are only two possible endomorphisms of $\mathbb{Z}_2$, namely the trivial one and id. So if $\tilde{T}$ is trivial, then T is trivial too. If $\tilde{T} = \mathrm{id}$, then $T = \mathrm{sgn}$.

[10] One of the inclusions is trivial to prove, the other requires noting that the product of two transpositions is a product of 3-cycles, and that 3-cycles are always commutators. This is not hard, but not obvious!