
Quantum Mechanics II - The Path Integral

Nic Ford

1 Introduction

This article is part of a series on physics for mathematicians, and as the title suggests, it's a sequel to my earlier article on quantum mechanics. It's not necessary to have read that article in particular to understand this one, though; any exposure to the basic framework of quantum mechanics should be just fine, and we'll briefly review the relevant parts.

We're going to be discussing the *path integral formulation* of quantum mechanics, an alternative way to set up the theory first described by Richard Feynman in his Ph.D. thesis. The original motivation was to find a way to express quantum mechanics in terms of the Lagrangian (rather than Hamiltonian) description of classical mechanics, but it results in a very different-looking picture: rather than providing a PDE that lets you evolve the state of the system through time, it expresses the relevant quantities in terms of an integral over *all possible paths* in configuration space connecting a given pair of points.

As one might imagine, it takes quite a bit of nontrivial mathematical work to make sense of such an integral and to show that it produces the same physics as the Hamiltonian approach to quantum mechanics. As happens a lot when physics interacts with mathematics, physicists tend not to worry a whole lot about this, and it is possible to get a lot computationally in this way. We'll start out by presenting this informal picture to get an idea of why one should expect anything like a path integral to show up at all, but afterwards we'll outline a couple of the schemes for making mathematical sense of it. We'll conclude with a short discussion of the advantages of viewing quantum mechanics from this perspective.

One thing this article *doesn't* spend enough time on is explaining how to use the path integral to do concrete computations. I made this choice not because these computations are unimportant but because the topic is covered very well in many of the references. I encourage the reader to supplement this piece by at least reading up on how to use the path integral to solve the quantum harmonic oscillator.

I found many sources very helpful when preparing this article. They include:

- Several of the books I referred to in the first quantum mechanics article discuss this topic well, including:
 - *Quantum Field Theory: A Tourist Guide for Mathematicians* by Gerald Folland
 - *Quantum Mechanics for Mathematicians* by Leon Takhtajan
 - *Quantum Theory for Mathematicians* by Brian Hall

Takhtajan's discussion is by far the most thorough of these three.

- The book *Lectures on Quantum Mechanics* also has a brief but good section on the path integral.
- You can learn much more about the mathematically rigorous versions of the path integral outlined here from the following books:
 - *Functional Integration and Quantum Physics* by Barry Simon
 - *Mathematical Theory of Feynman Path Integrals: An Introduction* by Sergio Albeverio, Raphael Høegh-Krohn, and Sonia Mazzucchi
 - *Quantum Physics: A Functional Integral Point of View* by James Glimm and Arthur Jaffe.
- These notes by Matthias Blau go through the concepts at a physicist's level of rigor, and they include a lot of the explicit computations I skip here. These notes by Hitoshi Murayama are similar.

I am grateful to Harry Altman and Jordan Watkins for helpful comments on an earlier draft of this article.

2

The Informal Argument

We'll start by going through the argument as it's usually presented in physics classes, which is pretty similar to what Feynman wrote in his thesis. The goal here is to understand why one should believe something like Feynman's formula should be true. The argument will raise many mathematical questions, up to and including how one is even supposed to define the path integral itself. All of these questions are worthwhile, but we will hold off on answering any of them until the next section.

We will consider a quantum particle in \mathbb{R}^d , so that the Hilbert space is $\mathcal{H} \cong L^2(\mathbb{R}^d)$, and write Q_1, \dots, Q_d and P_1, \dots, P_d for the position and momentum observables respectively. We will assume throughout that the Hamiltonian can be expressed as a power series in the Q 's and P 's. We can express states in terms of wavefunctions $\psi : \mathbb{R}^d \rightarrow \mathbb{C}$ with $\int_{\mathbb{R}^d} |\psi|^2 = 1$, so that $(Q_j \psi)(x) = x_j \psi(x)$ and $P_j \psi = -i\hbar \partial_j \psi$. Throughout this article, we will use the notation $\langle -, - \rangle$ for inner products in \mathcal{H} and \cdot for inner products in \mathbb{R}^d .

The Q 's commute, so they have a common set of generalized eigenstates given by $\psi_y(x) := \delta(x - y)$ for $y \in \mathbb{R}^d$. In particular, this means that for any $\psi \in \mathcal{H}$, $\psi(x) = \langle \psi_x, \psi \rangle$. The same is true of the P 's; in this representation their generalized eigenstates take the form

$$\phi_p(x) = (2\pi\hbar)^{-d/2} \exp(ip \cdot x / \hbar)$$

for $p \in \mathbb{R}^d$. It will be important going forward that we may write any state ψ in terms of either of these generalized bases, that is,

$$\psi = \int \langle \psi_y, \psi \rangle \psi_y dy = \int \langle \phi_p, \psi \rangle \phi_p dp.$$

The first formula follows directly from the definition of the delta function, and the second amounts to the Fourier inversion formula.

Given some ψ representing the state of the system at time 0, the state at some time $t > 0$ is given by $e^{-itH/\hbar}\psi$, where H is the Hamiltonian of the system in question. Let's write $\psi_{q,t} = e^{itH/\hbar}\psi_q$ for the state which has definite position q at time t , so that $\psi_q = \psi_{q,0}$. This gives us a convenient way to express our time-evolved wavefunction:

$$\psi(q', t) = \langle \psi_{q'}, e^{-itH/\hbar}\psi \rangle = \langle \psi_{q',t}, \psi \rangle = \int \langle \psi_{q',t}, \psi_{q,0} \rangle \langle \psi_{q,0}, \psi \rangle dq = \int \langle \psi_{q',t}, \psi_{q,0} \rangle \psi(q) dq.$$

So, if we can compute $\langle \psi_{q',t}, \psi_{q,0} \rangle$ as a function of the variables q, q', t , we'll have an expression for the dynamics of any wavefunction. This function is sometimes called the **propagator**. Our goal is to find an expression for it that doesn't mention operators or the Hilbert space.

2.1 To First Order

We'll build up from the case where $q' = q + \Delta q$ and $t = \Delta t$, where the changes in position and time are small enough that we can discard everything past first order. In this case, we have

$$\begin{aligned} \langle \psi_{q+\Delta q, \Delta t}, \psi_{q,0} \rangle &= \langle \psi_{q+\Delta q}, e^{-i(\Delta t)H/\hbar} \psi_q \rangle \\ &= \langle \psi_{q+\Delta q}, (1 - i(\Delta t)H/\hbar + O(\Delta t^2)) \psi_q \rangle \\ &= \delta(\Delta q) - \frac{i(\Delta t)}{\hbar} \langle \psi_{q+\Delta q}, H \psi_q \rangle + O(\Delta t^2) \end{aligned}$$

We can eliminate all mention of operators from the second term by employing the following trick. We assumed that H is a power series in the Q 's and P 's, so use the commutation relations $[Q_j, P_k] = i\hbar\delta_{jk}$ to place, in each term, all P 's to the left of any Q 's. Write $H_{\text{cl}}(q, p)$ for the formal power series — in ordinary, commuting variables — that results from this process. (For example, if $H = Q_1P_1 + Q_3P_2^2 = P_1Q_1 + i\hbar + P_2^2Q_3$, then $H_{\text{cl}} = i\hbar + p_1q_1 + p_2^2q_3$.)

Expand $\psi_{q+\Delta q}$ in terms of generalized eigenstates for the P 's:

$$\langle \psi_{q+\Delta q}, H \psi_q \rangle = \int \langle \psi_{q+\Delta q}, \phi_p \rangle \langle \phi_p, H \psi_q \rangle dp.$$

With H in the “ P 's before Q 's” form described above we may then move each of the Q and P factors in each term of H to the corresponding side of the inner product, so that each is acting on one of its own eigenvectors. We're left with

$$\int H_{\text{cl}}(q, p) \langle \psi_{q+\Delta q}, \phi_p \rangle \langle \phi_p, \psi_q \rangle dp = (2\pi\hbar)^{-d} \int H_{\text{cl}}(q, p) \exp(ip \cdot \Delta q/\hbar) dp,$$

where the last expression comes from plugging in the formula from above for $\langle \psi_q, \phi_p \rangle = \phi_p(q)$ twice.

Another Fourier transform lets us write $\delta(\Delta q) = (2\pi\hbar)^{-d} \int e^{ip \cdot \Delta q/\hbar} dp$, so that both terms in our expression are integrals over p . Evaluating all the inner products and combining the terms gives us

$$\begin{aligned} \langle \psi_{q+\Delta q, \Delta t}, \psi_{q,0} \rangle &= (2\pi\hbar)^{-d} \int \exp\left(\frac{ip \cdot \Delta q}{\hbar}\right) \left(1 - \frac{i(\Delta t)}{\hbar} H_{\text{cl}}(q, p)\right) dp + O(\Delta t^2) \\ &= (2\pi\hbar)^{-d} \int \exp\left(\frac{i(\Delta t)}{\hbar} \left[p \cdot \frac{\Delta q}{\Delta t} - H_{\text{cl}}(q, p)\right]\right) dp + O(\Delta t^2), \end{aligned}$$

where in the second expression we have turned $1 - i(\Delta t)H_{\text{cl}}/\hbar$ back into an exponential.

2.2 The Role of the Lagrangian

There is already something interesting about this expression. When moving between Hamiltonian and Lagrangian mechanics, we turn a Hamiltonian into its corresponding Lagrangian using a Legendre transform, which (in simple cases) takes the form $L = p\dot{q} - H$. At least formally, this looks a lot like the expression we have in square brackets inside this integral!

The main difference is that in the classical version, p and \dot{q} are related by one of the Hamilton-Jacobi equations, $dq/dt = \partial H/\partial p$, whereas in this integral, p and $\Delta q/\Delta t$ have nothing to do with each other. The analogy with Lagrangian mechanics can be made to work, though, in the (very common) case where the Hamiltonian has the form $H = P^2/2m + V(Q)$ for some $m > 0$ and some function V . In this case, we have

$$(2\pi\hbar)^{-d} \int \exp\left(\frac{i(\Delta t)}{\hbar} \left[p \cdot \frac{\Delta q}{\Delta t} - \frac{p^2}{2m} - V(q) \right]\right) dp,$$

which is a Gaussian integral.

We'll take a brief detour to discuss how to compute integrals like this one. Consider first the integral $\int_{\mathbb{R}^d} e^{-Q(x)/2}$, where $Q(x) = x \cdot Ax + b \cdot x + c$, where A is a positive definite symmetric matrix, $b \in \mathbb{R}^d$, and $c \in \mathbb{R}$. Complete the square to write $Q(x)$ in the form $(x - x_0) \cdot A(x - x_0) + c'$; we then have that x_0 is the minimum of Q and $c' = Q(x_0)$. So

$$\int_{\mathbb{R}^d} e^{-Q(x)/2} = e^{-Q(x_0)/2} \int_{\mathbb{R}^d} e^{-(x-x_0) \cdot A(x-x_0)/2} = e^{-Q(x_0)/2} \int_{\mathbb{R}^d} e^{-x \cdot Ax/2} = e^{-Q(x_0)/2} \sqrt{\frac{(2\pi)^d}{\det A}};$$

the last equality follows by performing an orthogonal change of coordinates to diagonalize A and then rescaling each coordinate of x by the square root of the corresponding eigenvalue of A .

This equation in fact still holds if A is complex-valued with positive definite real part, with the caveat that we have to be careful about picking the right square root of the now possibly complex $\det A$. (We won't belabor this question here.) If the real part of A is just positive semidefinite (as in our present case, where A is a purely imaginary scalar) then the integral is no longer absolutely convergent. The strategy we'll go with to deal with this is to replace A with $A + \epsilon I$ and then letting ϵ go to 0 from the right.

At any rate, this argument shows that in order to compute our integral, we need to find the value of p at which $p \cdot (\Delta q/\Delta t) - H$ is stationary and plug it into the exponent in the integral. In other words, we want the value of p for which $\Delta q/\Delta t = \partial H/\partial p$. But this exactly means that $p \cdot (\Delta q/\Delta t) - H$ is the Lagrangian!

When all the factors are accounted for, we therefore have

$$\begin{aligned} \langle \psi_{q+\Delta q, \Delta t}, \psi_{q, 0} \rangle &= \left(\frac{2\pi i \hbar \Delta t}{m} \right)^{-d/2} \exp\left(\frac{i\Delta t}{\hbar} \left[\frac{1}{2} m \left(\frac{\Delta q}{\Delta t} \right)^2 - V(q) \right]\right) + O(\Delta t^2) \\ &= \left(\frac{2\pi i \hbar \Delta t}{m} \right)^{-d/2} \exp\left[\frac{i\Delta t}{\hbar} L\left(q, \frac{\Delta q}{\Delta t}\right)\right] + O(\Delta t^2). \end{aligned}$$

(Here $\sqrt{i} = e^{i\pi/4}$.) The argument here would also work if the Hamiltonian contained a term of the form $f(Q)P$ for an arbitrary function f , since we would still be left with a Gaussian integral to compute, but when it matters we'll restrict our attention to the case where this term is absent. This formula will be our starting point for the general case.

2.3 In General

We can turn this first-order formula into a method for computing $\langle \psi_{q',t}, \psi_{q,0} \rangle$ even when t is large. The trick is simply to divide the interval $[0, t]$ into small pieces and use the first-order formula on each one: just as we inserted a (generalized) basis of momentum eigenstates into the middle of the expression before, we may do the same with position eigenstates as a way of cutting the time interval into two pieces. That is, for $0 < t_1 < t$, we have

$$\langle \psi_{q',t}, \psi_{q,0} \rangle = \int \langle \psi_{q',t}, \psi_{q_1,t_1} \rangle \langle \psi_{q_1,t_1}, \psi_{q,0} \rangle dq_1.$$

Divide the whole interval into N pieces, writing $t_j = jt/N$ for $j = 0, \dots, N$. We get:

$$\langle \psi_{q',t}, \psi_{q,0} \rangle = \int \langle \psi_{q',t}, \psi_{q_{N-1},t_{N-1}} \rangle \cdots \langle \psi_{q_1,t_1}, \psi_{q,0} \rangle dq_1 \cdots dq_{N-1}.$$

(For ease of reading, I'm using only one integral sign at the front to stand for all of the q_j integrals.) Writing $\Delta t = t/N$, $q_0 = q$, and $q_N = q'$, our first-order formula then implies that our propagator equals:

$$\lim_{N \rightarrow \infty} \left[(2\pi\hbar)^{-dN} \int \exp \left(\frac{i\Delta t}{\hbar} \left[\sum_{j=0}^{N-1} p_j \cdot \frac{q_{j+1} - q_j}{\Delta t} - H_{\text{cl}}(q_j, p_j) \right] \right) dq_1 \cdots dq_{N-1} dp_1 \cdots dp_{N-1} \right],$$

or, in the case where the Hamiltonian is quadratic in the P 's and we can perform the p integrals,

$$\lim_{N \rightarrow \infty} \left(\left(\frac{2\pi i\hbar\Delta t}{m} \right)^{-dN/2} \int \exp \left[\frac{i\Delta t}{\hbar} \sum_{j=0}^{N-1} L \left(q_j, \frac{q_{j+1} - q_j}{\Delta t} \right) \right] dq_1 \cdots dq_{N-1} \right).$$

Recall that in Lagrangian mechanics, the path $\gamma(t)$ that a particle follows has the property that its *action*

$$S[\gamma] = \int_0^t L(\gamma(u), \dot{\gamma}(u)) du$$

is stationary. Intriguingly, the summation appearing inside this second expression can be expressed as an action: it is (up to a constant factor) the action of a path that visits each q_j at time t_j and travels with constant velocity in between.

As N grows, the set of paths appearing in this way grows more refined, approximating any particular continuous path arbitrarily closely in (say) the L^∞ metric. This inspired Feynman to, somewhat fancifully, think of this expression as a limit of ‘‘Riemann sums’’ of an integral over the space of *all* paths connecting q to q' . This is written

$$\begin{aligned} \langle \psi_{q',t}, \psi_{q,0} \rangle &= C \int_{C(q',q;t)} \exp \left[\frac{i}{\hbar} \int_0^t L(\gamma(u), \dot{\gamma}(u)) du \right] \mathcal{D}\gamma \\ &= C \int_{C(q',q;t)} \exp \left(\frac{i}{\hbar} S[\gamma] \right) \mathcal{D}\gamma. \end{aligned}$$

Here $C(q', q; t)$ is the set of all continuous paths $\gamma : [0, t] \rightarrow \mathbb{R}^d$ with $\gamma(0) = q$ and $\gamma(t) = q'$. This expression has a lot of mathematical problems if one tries to take it too literally. $\mathcal{D}\gamma$ refers to the (entirely fictional) translation-invariant measure on the space of all such paths. C is supposed to be a constant depending only on m and t , the analogue of the $(2\pi i\hbar\Delta t/m)^{-dN/2}$ appearing in

the earlier expression, which of course doesn't converge to anything as $N \rightarrow \infty$. Note that the Δt outside the sum in the earlier expression has become the du inside the integral.

Plugging this into our original expression for the time-evolved wavefunction produces the formula

$$\psi(q', t) = C \int_{C(q'; t)} \exp\left(\frac{i}{\hbar} S[\gamma]\right) \psi(\gamma(0), 0) \mathcal{D}\gamma,$$

where $C(q'; t)$ is the set of all continuous paths γ with $\gamma(t) = q'$, but with the left endpoint unconstrained.

3 Rigorous Path Integrals

Despite seeming like total mathematical nonsense, this formalism is very powerful. It's possible to get meaningful answers out of it by just treating it formally and computing until you've eliminated all mention of spaces of paths and are left with just a number or a function. This is essentially what physicists do, and it works well enough for that purpose. But the ideas that gave rise to the path integral are compelling enough to make one wonder if there is any way at all to make mathematical sense of it. And it is in fact possible to make more progress on this than you might think!

We will focus entirely on the very common case where $H = H_0 + V(Q)$, where $H_0 = P^2/2m$ and V is a function. (Depending on the framework, there will be some conditions on V .) In this case, everything we've done up through concluding

$$\langle \psi_{q', t}, \psi_{q, 0} \rangle = \lim_{N \rightarrow \infty} \left(\left(\frac{2\pi i \hbar \Delta t}{m} \right)^{-dN/2} \int \exp \left[\frac{i \Delta t}{\hbar} \sum_{j=0}^{N-1} L \left(q_j, \frac{q_{j+1} - q_j}{\Delta t} \right) \right] dq_1 \cdots dq_{N-1} \right)$$

can be made rigorous following a more careful version of essentially the same computation.

The key mathematical ingredient is the *Lie-Kato-Trotter formula*, which says that, if A and B are self-adjoint operators on a Hilbert space \mathcal{H} and $A + B$ is essentially self-adjoint on the intersection of their domains, then for any $\psi \in \mathcal{H}$, we have

$$e^{i(A+B)\psi} = \lim_{N \rightarrow \infty} (e^{iA/N} e^{iB/N})^N \psi.$$

Applying this to $\frac{1}{\hbar} H_0$ and $\frac{1}{\hbar} V$ produces the formula with the integral over both the p and q variables, and the p integrals can then be performed to produce this expression just as we did above.

It's entirely reasonable to be satisfied with this as a definition of the path integral: you can treat the integral over paths as just a formal expression, a suggestive way of writing this limit. But it's actually possible to make sense of the integral in a way that actually "takes place" on the space of continuous paths, and there are some interesting things to learn by doing so. A lot of approaches have been developed since Feynman's original proposal; we'll sketch two of them here.

3.1 The Wiener Measure and Analytic Continuation

One strategy involves taking the Schrödinger equation for our Hamiltonian

$$\frac{d\psi}{dt} = \frac{-i}{\hbar} H\psi = \frac{-i}{\hbar} \left(\frac{-\hbar^2}{2m} \nabla^2 \psi + V\psi \right)$$

and replacing the real number t with a purely imaginary parameter, writing $t = -i\tau$ with τ real, which has the effect of turning $e^{-itH/\hbar}$ into $e^{-\tau H/\hbar}$ and the Schrödinger equation into

$$\frac{d\psi}{d\tau} = \frac{\hbar}{2m} \nabla^2 \psi - \frac{1}{\hbar} V\psi.$$

(This move is sometimes called performing a **Wick rotation** in the time variable.) It's useful to think of this new equation as a version of the heat equation, just with the potential term added on.

If we went through the same argument as before, but with the goal of computing $e^{-\tau H/\hbar}$, the resulting path integral expression would be

$$\langle \psi_{q'}, e^{-\tau H/\hbar} \psi_q \rangle = C \int \exp \left[\frac{-1}{\hbar} \int_0^\tau \left(\frac{1}{2} m |\dot{\gamma}(u)|^2 + V(\gamma(u)) \right) du \right] \mathcal{D}\gamma,$$

where the integral again takes place over the space of all continuous paths γ with $\gamma(0) = q$ and $\gamma(\tau) = q'$. We mentioned before that this expression is problematic for a couple reasons: there is no measure on the space of continuous paths that can actually play the role of $\mathcal{D}\gamma$ here, and there is no number which can play the role of C . In this heat-equation-like setting, though, it turns out that there *is* a measure which can play the role of

$$C \exp \left[\frac{-1}{\hbar} \int_0^\tau \frac{1}{2} m |\dot{\gamma}(u)|^2 du \right] \mathcal{D}\gamma.$$

Note the formal similarity between this expression and the portion of the path integral expression we are attempting to replace. The normalizing factor out front serves the role of C , the product of the exponentials inside the integral serves the role of the middle factor, and dq_1, \dots, dq_{n-1} serves the role of $\mathcal{D}\gamma$. None of these three factors has any meaning individually in the limit as n goes to infinity — even $|\dot{\gamma}(u)|^2$ is meaningless, since the set of paths differentiable at even one point has Wiener measure zero! — but their product does.

The measure that does the job is called the **conditional Wiener measure**. It is a probability measure $\mu_{q',q}^\tau$ on the set $C(q', q; \tau)$ of continuous paths on $[0, \tau]$ starting at q and ending at q' , and it is uniquely characterized by the property that, for any partition $0 = \tau_0 < \tau_1 < \dots < \tau_{n-1} < \tau_n = \tau$ and any Borel sets $A_1, \dots, A_{n-1} \subseteq \mathbb{R}^d$,

$$\mu_{q',q}^\tau(\{\gamma : \gamma(\tau_i) \in A_i \forall i\}) = \int_{A_1} \dots \int_{A_{n-1}} p(q_n, q_{n-1}; \tau_n - \tau_{n-1}) \dots p(q_1, q_0; \tau_1 - \tau_0) dq_1 \dots dq_{n-1},$$

where $q_n = q', q_0 = q$, and

$$p(q_f, q_i; \Delta\tau) = (2\pi\hbar\Delta\tau/m)^{-d/2} \exp \left(-\frac{m(q_f - q_i)^2}{2\hbar\Delta\tau} \right).$$

(The “conditional” in the name refers to the fact that we are constraining the value of the right endpoint; the unconditional Wiener measure — obtained by removing the factor involving q' — constrains only the left.)

The Wiener measure was originally developed to describe Brownian motion, and the measure of the set we just described can be thought of as the probability that a particle undergoing Brownian motion visits each q_i at the corresponding time τ_i , conditioned on the fact that it starts at q and ends at q' . The function $p(q_f, q_i; \Delta\tau)$ gives the probability for a particle starting at q_i to end up at q_f after time $\Delta\tau$. Notice that the Wick rotation has the effect of making the

integrand in the path integral positive, since the exp is now applied to a real rather than purely imaginary quantity, which enables us to interpret the resulting measure probabilistically.

It was crucial to perform the Wick rotation to make this work. Heuristically, since almost all paths fail to be differentiable, you can think of the “divergent” integral $\exp(-\int_0^\tau |\dot{\gamma}|^2)$ as compensating for the “infinte” factor $\mathcal{D}\gamma$. If we tried to do this for the original path integral, we would find ourselves with an extra factor of i inside the exponential, which would ruin this delicate balance. There is no analogue of the Wiener measure for the original, un-rotated path integral.

If we interpret the path integral in terms of the Wiener measure, then it does in fact solve our Wick-rotated Schrödinger equation. Since the Wiener measure is meant to stand for all the factors in the path integral except the one containing the potential, we should integrate it against that last factor to get the propagator. That is, we have the **Feynman–Kac formula**, which says that if V is continuous and bounded below, then

$$\langle \psi_{q'}, e^{-\tau H/\hbar} \psi_q \rangle = \int_{C(q', q; \tau)} \exp\left(-\frac{1}{\hbar} \int_0^\tau V(\gamma(u)) du\right) d\mu_{q', q}^\tau.$$

Now, we are left with the question of how this formula helps us solve the *actual* Schrödinger equation. One tempting strategy try to extend the definition of the Wiener measure to cover complex values of τ ; we already said this idea is no good when τ is purely imaginary (that is, when t is real) but maybe we could consider some sort of limit as τ approaches the imaginary axis. Sadly, this can't work: the construction of the Wiener measure fails as soon as τ has *any* nonzero imaginary part at all.

Instead, the strategy is to consider an *analytic continuation* of the Feynman–Kac formula. There are a few ways of doing this that work, including allowing τ , m , or \hbar to become complex. In all cases, as soon as the parameter leaves the real axis, the function in question is no longer itself expressible as an integral over the space of paths. Nonetheless, by performing this analytic continuation and taking the appropriate limit as whichever parameter approaches the imaginary axis, we can indeed recover a solution to the Schrödinger equation. There are details on these approaches in Glimm and Jaffe's book, and a short discussion with good references in the final chapter of the book by Albeverio et al.

3.2 Fresnel Integrals

One indication that it might be challenging to make mathematical sense of the path integral comes from the fact that the integrand is a complex number of absolute value 1, so if anything resembling convergence is going to happen, it must rely on rapidly oscillating values cancelling each other out rather than on the integrand decaying rapidly enough. This is true even in one dimension: the integral $\int_{-\infty}^{\infty} \exp(ip^2) dp$ converges only conditionally, since $\exp(ip^2)$ is obviously not Lebesgue integrable.

The approach taken in the book by Albeverio et al. mentioned in the introduction revolves around assigning a consistent meaning to integrals of this type. They start by considering integrals on \mathbb{R}^d of the form

$$\int e^{-\frac{1}{2}i|x|^2} \phi(x) dx.$$

If ϕ has rapid enough decay — for example, if it belongs to the Schwarz space — then this integral is well-defined, and we may take a Fourier transform to show that it equals

$$(2\pi i)^{-d/2} \int e^{\frac{1}{2}i|x|^2} \hat{\phi}(x) dx,$$

where $\hat{\phi}(x) = \int \exp(i(x \cdot y))\phi(y)dy$ is the Fourier transform of ϕ . Note the similarity between this integral and the one that appears inside the limit in our application of the Lie–Kato–Trotter formula.

The first integral actually still makes sense if we replace $\phi(x)dx$ with any complex measure $d\mu(x)$ with finite total variation. So, if we have a function f which is the Fourier transform of such a measure — that is, $f(x) = \int \exp(i(x \cdot y))d\mu(y)$ — we can *define*

$$(2\pi i)^{-d/2} \int e^{\frac{1}{2}i|x|^2} f(x)dx$$

by setting it equal to

$$\int e^{-\frac{1}{2}i|x|^2} d\mu(x),$$

since, when μ has finite total variation, this integral is well-defined even if the first one is not. For example, $f(x) = 1$ is in the class of functions that arise in this way.

The authors call the functions f that arise in this way **Fresnel integrable functions** and they call the integral being defined the **Fresnel integral** of f , after the classical Fresnel integrals that take a similar form. The Fresnel integrable functions include the Schwarz functions, but also include other functions. For such functions — including, for example, $f(x) = 1$ — the function $e^{\frac{1}{2}i|x|^2} f(x)$ is not Lebesgue integrable in general, so the only meaning its integral has in this context is given by the integral against $d\mu$.

Note that the factor of $(2\pi i)^{-d/2}$ is a part of the formal expression of the Fresnel integral we are now defining; to emphasize this, the authors write

$$\widetilde{\int} e^{\frac{1}{2}i|x|^2} f(x)dx := \int e^{-\frac{1}{2}i|x|^2} d\mu(x).$$

The tilde is supposed to be taken to “include” the factor of $(2\pi i)^{-d/2}$. This factor serves the role of the quantity we’ve been calling C .

We can further extend the definition of the Fresnel integral — tilde and all — to an infinite-dimensional separable real Hilbert space \mathcal{P} . That is, if μ is a complex measure of bounded total variation on \mathcal{P} and f is its Fourier transform, we can define $\widetilde{\int} e^{\frac{1}{2}i|x|^2} f(x)dx$ via the exact same formula, but since \mathcal{P} is infinite-dimensional there is no longer an analogue of the expression with the $(2\pi i)^{-d/2}$.

Making sense of the Feynman path integral in this formulation is then a matter of applying this machinery when \mathcal{P} is an appropriate Hilbert space of continuous paths in some \mathbb{R}^d . (This is *not* the Hilbert space that quantum mechanics takes place in!) Write \mathcal{P} for the space of differentiable paths $\gamma : [0, t] \rightarrow \mathbb{R}^d$ such that $\gamma(0) = 0$ and the time derivative $\dot{\gamma}$ is in L^2 . In order to produce the path integral formula we’re after, we take the inner product on \mathcal{P} to be given by

$$\langle \gamma_0, \gamma_1 \rangle = \frac{m}{\hbar} \int_0^t \dot{\gamma}_0(u) \cdot \dot{\gamma}_1(u) du.$$

Suppose our initial state is given by the wavefunction $\psi_0(x)$. Then, defining

$$f(\gamma) = \exp \left[\frac{-i}{\hbar} \int_0^t V(\gamma(u)) du \right] \psi_0(0),$$

the authors show that the path integral solves the Schrödinger equation, in the sense that

$$\begin{aligned}\psi(x, t) &= \widetilde{\int} e^{\frac{i}{\hbar} \int_0^t |\dot{\gamma}|^2} f(\gamma) d\gamma \\ &= \widetilde{\int} \exp\left(\frac{i}{\hbar} \int_0^t \frac{1}{2} m |\dot{\gamma}(u)|^2 du\right) \cdot \exp\left(\frac{-i}{\hbar} \int_0^t V(\gamma(u)) du\right) \psi_0(0) d\gamma.\end{aligned}$$

As in the Wiener measure approach, we are able to give meaning to the path integral in sort of an indirect way. The Fresnel integral is not literally an integral of a function against some measure on the space of paths, just a functional defined on that space in a way that gives it many of the properties we would expect from such an integral. In order for this to be useful, we need to know that the f just defined is in fact Fresnel integrable for nice choices of the potential V . The paper “Functions in the Fresnel Class” by Chang, Johnson, and Skoug is a good reference for this.

4 What Is It Good For?

The path integral describes exactly the same physics as the Hamiltonian formalism that’s usually used for nonrelativistic quantum mechanics, and anything that can be done in one picture can also be done in the other. In particular, the set of systems that can be solved analytically in either picture is pretty small: basically just free particles and particles in one of a small list of potentials, most notably the quadratic potential that gives rise to the harmonic oscillator. Because it has been written up well in so many other places, we won’t go over the path integral computation for the quantum harmonic oscillator; if you’re interested you can find it, for example, in the notes by Blau linked in the introduction.

4.1 The Classical Limit

Still, as often happens when one is given two different ways of looking at the same situation, there are some things that are easier to see from the path integral point of view. One is the relationship between quantum and classical physics: classical mechanics arises, in some sense, as the $\hbar \rightarrow 0$ limit of quantum mechanics, and the path integral formalism offers a particularly nice way to see how this happens.

This picture starts from the stationary phase approximation for finite-dimensional integrals of the form $\int_{\mathbb{R}^n} g(x) e^{if(x)/\hbar} dx$. If \hbar is very small, we should usually expect that, in any bounded region in \mathbb{R}^n , the exponential in the integrand is oscillating so quickly that the integral will be very close to zero. This is the right picture *except* in the neighborhood of a critical point of f , where the integral will instead resemble a Gaussian.

This can be formalized when f and g are sufficiently nice. (In particular we require all the critical points of f to be nondegenerate.) The result is a formula of the form

$$\int_{\mathbb{R}^n} g(x) e^{if(x)/\hbar} dx = \sum_{c \in C(f)} g(c) e^{if(c)/\hbar} K_c + o(\hbar^{n/2}),$$

where $C(f)$ is the set of critical points of f and K_c depends only on the second derivatives of f at c . With some more work, it is possible to extract a power series in \hbar for the error term.

This suggests an analogous picture for the path integral

$$\langle \psi_{q',t}, \psi_{q,0} \rangle = C \int_{C(q',q;t)} \exp\left(\frac{i}{\hbar} S[\gamma]\right) \mathcal{D}\gamma.$$

When \hbar is very small, the integral should be dominated by contributions from the paths γ which are critical points of the action S . But these are exactly the solutions to the classical equations of motion!

While this argument means that the classical path is the dominant contribution to the path integral for small \hbar , that doesn't directly imply that particles follow these classical paths in this same limit. But it's not too hard to give a very rough heuristic argument for this as well.

Suppose the initial state is given by some function $\psi(q, 0)$, and for any q' and t , suppose there is a unique classical path $\gamma_{q',q;t}$ which visits q and time 0 and q' at time t . The stationary phase argument tells us that

$$\psi(q', t) = \mathcal{A} \int \exp\left(\frac{i}{\hbar} S[\gamma_{q',q;t}]\right) \psi(q, 0) dq.$$

Here \mathcal{A} is some constant that we will not worry about for the rest of this sketchy argument.

Now, suppose that the initial state is sharply peaked around some initial position q_0 and momentum p_0 . (The latter condition means that the Fourier transform of $\psi(-, 0)$ is peaked.) This means we can replace this q integral with an integral over some small $[q_0 - \Delta q, q_0 + \Delta q]$, over which we will replace S with its linear approximation. But (I encourage you to check) $-\partial S / \partial q[\gamma_{q',q;t}]$ is the *momentum* of the particle at $t = 0$ when moving along the classical path. Calling this momentum $p_{\text{cl}}(q', q)$, we conclude that

$$\psi(q', t) \approx \mathcal{A} \int_{q_0 - \Delta q}^{q_0 + \Delta q} \exp\left(-\frac{i}{\hbar} q \cdot p_{\text{cl}}(q', q)\right) \psi(q, 0) dq,$$

where we have absorbed the constant part of the linear approximation into \mathcal{A} .

Now, to make use of that fact that the momentum is sharply peaked around p_0 , we perform a Fourier transform, giving us:

$$\begin{aligned} \psi(q', t) &\approx \mathcal{A} \int_{q_0 - \Delta q}^{q_0 + \Delta q} \int \exp\left(-\frac{i}{\hbar} q \cdot p_{\text{cl}}(q', q)\right) \exp\left(\frac{i}{\hbar} q \cdot p\right) \hat{\psi}(p, 0) dp dq \\ &= \mathcal{A} \int_{q_0 - \Delta q}^{q_0 + \Delta q} \int \exp\left[\frac{i}{\hbar} q \cdot (p - p_{\text{cl}}(q', q))\right] \hat{\psi}(p, 0) dp dq. \end{aligned}$$

The integrand will be close to zero if p isn't close to p_0 . But also, if $p - p_{\text{cl}}(q', q)$ is at all big — on the order of $\hbar/\Delta q$ or bigger — then the q integral will be the integral of a rapidly oscillating exponential, which will also make it very close to zero. We conclude that the only way for $\psi(q', t)$ to have any appreciable amplitude is if p_0 is close to $p_{\text{cl}}(q', q)$, that is, if the classical motion with initial position q_0 and momentum p_0 ends up at q' .

While this is a somewhat satisfying picture, the real value added by the path integral approach in the low \hbar is not really primarily in verifying that quantum mechanics reproduces classical mechanics. More interesting — though beyond the scope of this article — are the higher-order terms in the stationary phase approximation, which give corrections to the classical picture when \hbar is small but not quite negligible, the so-called **semiclassical approximation**.

4.2 Symmetry

One large advantage to thinking in terms of path integrals — and the reason I’m including it in this series of articles — comes from relativistic quantum field theory.

We’ll have much more to say about this later in the series when we start exploring quantum field theory properly, but put simply, because the path integral formalism is built out of the Lagrangian rather than the Hamiltonian, it is easier to express it in a way that transparently accounts for the symmetries of a system right from the start. For this reason, a lot of the computations one needs to do in quantum field theory are a lot easier to write down in the path integral language.

There’s not actually anything especially quantum about this observation. Hamiltonian mechanics, in both its classical and quantum forms, involves building some sort of time-translation operator out of the Hamiltonian of the theory, and so in particular it requires one to pick a single forward time direction. This requires breaking the Lorentz symmetry of special relativity, which mixes time and space coordinates. Lagrangian mechanics is easier to work with in a way that doesn’t require such a choice. Something similar is true for many other symmetries and constraints a theory might have, which we’ll also hopefully see later on.

We’re going to leave this discussion at this very general point for now. We haven’t focused on performing concrete calculations with the path integral here, primarily because they’re written up well in many other places. When we do eventually come around to our discussion of quantum field theory, though, the path integral will be an indispensable conceptual tool.