

Hamiltonian and Lagrangian Mechanics

Nic Ford

1 Introduction

This article is the first in a series I plan to write about physics for a mathematically trained audience. We're going to start by talking about classical mechanics, the stuff that your first physics class was probably about if you've ever taken one. The formulation of classical physics usually presented in introductory physics classes is called Newtonian mechanics; it talks about things like masses and forces and Newton's laws of motion. Newtonian mechanics is easy to teach and to work with without much machinery, but it has some features that can make it difficult to analyze mathematically. Physical systems and their interactions are described in terms of coordinates with velocity and force vectors all over the place, and it can be difficult to know how to deal with things like symmetries and constraints.

There are other, equivalent ways of describing classical mechanics, sometimes collectively called "analytical mechanics," which are much easier to describe in a coordinate-free way. At the cost of a bit more abstraction, the analytical formulations have two big advantages: they make it easier to set up and solve some very complicated mechanics problems and, probably more importantly for our purposes, they make the relationship between classical mechanics and its generalizations most clear. The two most prominent such formulations are called Hamiltonian and Lagrangian mechanics, and they're what we're going to discuss in this article. They are, as we'll see, two different ways of saying the same thing, but they highlight different enough aspects of the situation that they're worth talking about separately.

This article assumes some mathematical background beyond what's usually used to present these ideas in a physics class that covers them. In particular, the reader is expected to be familiar with the basics of the theory of smooth manifolds to the level of someone who's finished a one-semester class on the subject. I will assume you remember a little bit about physics, but that you have never seen the Hamiltonian and Lagrangian frameworks discussed here.

I am very grateful to Jeff Hicks and Jake Levinson for their many helpful comments on earlier drafts of this article. Some of the examples and a couple ideas about the presentation are adapted from Gerald Folland's *Quantum Field Theory: A Tourist Guide for Mathematicians* and Michael Spivak's *Physics for Mathematicians: Mechanics I*, both of which I recommend.

2 Hamiltonian Mechanics

2.1 The Newtonian Setup

We'll start by briefly describing, in coordinates, the sort of Newtonian mechanics problem we're eventually going to be describing in a coordinate-free way. The prototypical example to keep in mind is that of a collection of N particles moving in \mathbb{R}^3 where particle i has mass m_i . We'll write the position of particle i as \mathbf{q}_i , with the boldface there to remind you that it's an element of \mathbb{R}^3 and not an \mathbb{R} -valued coordinate on \mathbb{R}^{3N} . We'll write $\mathbf{p}_i = m_i(d\mathbf{q}_i/dt)$ for the **momentum** of particle i .

In the Newtonian setup, we describe physics in terms of **forces**; we imagine that there is some vector \mathbf{F}_i we can compute for each particle for all time which tells us how that particle is accelerating, or equivalently, how its momentum is changing. Specifically, the relationship is given by "Newton's second law":

$$\mathbf{F}_i = \frac{d\mathbf{p}_i}{dt} = m_i \frac{d^2\mathbf{q}_i}{dt^2}.$$

In general one could imagine these forces depending on any data whatsoever about the physical system, but we're going to be most interested in the case of **conservative** forces. This is the case where there is a function V on \mathbb{R}^{3N} called a **potential** for which the force on particle i is given by

$$\mathbf{F}_i = -\frac{\partial V}{\partial \mathbf{q}_i}.$$

So the force is given by the gradient of a function which depends only on positions, not on momenta. This condition is equivalent to saying that the integral of the force vector field around a closed loop — a quantity called the **work** done by the force while traveling around the loop — is always zero.

The name "conservative" comes from the fact that, suitably interpreted, this last condition is what we mean by saying that energy is conserved. There are many physical phenomena that are often modeled as nonconservative forces; friction is probably the most familiar example. But an overwhelming amount of physical evidence points toward the belief that the fundamental laws of physics do conserve energy, and that physical models of things like friction are merely "neglecting" the energy that leaks into forms like heat and sound that are more difficult to model. It is possible, but somewhat painful, to set up Hamiltonian mechanics in a way that allows for things like friction, but we're going to focus on the conservative case in this article.

Throughout this short description we've already done things that make it difficult to keep track of what needs to be done with all these quantities when we change coordinates. The force on a particle is given by a gradient, and each momentum coordinate is "attached" to both a mass and a particular spatial coordinate. It will often be convenient to switch to a coordinate system that does not isolate each particle so neatly in its own triple of coordinates, or even one that mixes what we are now calling position and momentum coordinates. This all cries out for a description that describes the physical system in terms of points on a manifold, to which we can assign coordinates only once we know how all the mathematical objects involved are defined intrinsically.

2.2 Configuration Space and Phase Space

We'll start our quest for a coordinate-free description of mechanics by fixing a smooth manifold Q which we'll call **configuration space**. You should think of a point in Q as corresponding the "position" of each component of a physical system at some fixed time. Some examples worth keeping in mind are:

1. A particle moving in \mathbb{R}^3 . In this case, Q is just \mathbb{R}^3 .
2. N particles moving in \mathbb{R}^3 . We specify the configuration of this system by specifying the position of each particle, which we can do using a point in \mathbb{R}^{3N} .
3. Two particles connected by a rigid rod of length ℓ . We could describe the configuration of this system using a point in $\{(a, b) \in \mathbb{R}^3 \times \mathbb{R}^3 : \|a - b\| = \ell\}$.
4. A rigid body moving through space. We could describe its configuration using a point in $\mathbb{R}^3 \times SO(3)$, specifying the location of the object's center of mass and its orientation.

In particular, specifying a point $q \in Q$ gives you an instantaneous snapshot of the system, but it doesn't tell you anything about it's changing. Even if you have a complete description of the physics, this doesn't provide enough information to predict how the system will evolve in the future. (Imagine a ball rolling on a table; if you just know its position and not its velocity you don't know where it's about to move.)

So if we want to describe the state of a physical system in a way that allows us to do physics, the state needs to carry some additional information. Different formulations of analytical mechanics do this in different ways, and unfortunately the version used by the Hamiltonian formulation is one of the more opaque choices: the state of a physical system is given by specifying a point in the cotangent bundle of Q , which to match with physicists' conventions we will call **phase space**. We'll usually use the coordinates (q, p) to refer to a point in phase space. (So q is a point in Q and p is a cotangent vector at q .) When the system is in the state (q, p) , we'll call p the **momentum**.

The first time I encountered this setup I was confused by the fact that momentum is represented by a cotangent vector rather than a tangent vector — after all, the velocity of a particle is definitely a tangent vector, and momentum is supposed to be a multiple of it.

It will be easier to talk about this once we have the finished picture in front of us, but we can say a bit right now. While velocities should inarguably be tangent vectors — a velocity is literally the time derivative along the path that a particle is following — it's actually not clear that this extends to momenta. When we use the word "momentum" we will mean something more general than "mass times velocity"; the two will coincide for Newtonian mechanics in rectangular coordinates but they can be different in general. For example, if we have a particle of mass m moving in \mathbb{R}^2 and use polar coordinates, the momentum corresponding to the θ coordinate turns out to be the angular momentum $xp_y - yp_x$, which is *not* $m(d\theta/dt) = (xp_y - yp_x)/(x^2 + y^2)$.

Of course I have not yet said what it means for one expression or another to be the "right" generalization of momentum to a given coordinate system, but the point is that the relationship between momentum coordinates and derivatives of the corresponding position coordinates depends on the physical meaning of those coordinates; it's not something you can extract just by looking at configuration space. (Indeed, this is true even in rectangular coordinates: the relationship depends on the mass of the particle, which is a physical quantity.) We'll return to this question later.

2.3 Symplectic Geometry

To properly describe Hamiltonian mechanics, we'll need some basic facts about symplectic geometry, which we'll briefly go over now in case they aren't familiar.

A **symplectic manifold** is a smooth manifold M together with a choice of a nondegenerate closed 2-form ω on M . (That is, the antisymmetric bilinear form ω defines on each tangent space is nondegenerate and $d\omega = 0$. We'll see soon how each of these two conditions is relevant.) A diffeomorphism between two symplectic manifolds that preserves the symplectic form is called a **symplectomorphism**.

The main thing that will turn out to make Hamiltonian mechanics go is the fact that the cotangent bundle of a manifold naturally has the structure of a symplectic manifold. The cotangent bundle of any manifold comes with a canonical symplectic form which can be described pretty simply. We start by defining the **tautological 1-form** on T^*Q . Given a tangent vector v at a point $(q, p) \in T^*Q$, we'll write $\theta(v) = p(\pi_*(v))$, where $\pi : T^*Q \rightarrow Q$ is the projection map. Given a local coordinate system q_1, \dots, q_n on a chart on Q , we also get coordinates p_1, \dots, p_n on each cotangent space. I encourage you to check that in these coordinates

$$\theta = \sum_{i=1}^n p_i dq_i.$$

We then define $\omega = d\theta$, so that in these same coordinates

$$\omega = \sum_{i=1}^n dp_i \wedge dq_i,$$

which is clearly nondegenerate.

One very striking difference between Riemannian and symplectic geometry is that in a neighborhood of any point on any symplectic manifold (even if it's not a cotangent bundle) there is a coordinate system $q_1, \dots, q_n, p_1, \dots, p_n$ for which $\omega = \sum_i dp_i \wedge dq_i$. This result is called "Darboux's theorem" and the q 's and p 's are said to provide **canonical coordinates**. This means that, very unlike on a Riemannian manifold, a symplectic manifold has no local geometry, so there's no symplectic analogue of anything like curvature.

Even though phase space will end up being the only symplectic manifold we'll use to do physics, it's actually cleaner to describe the required machinery in more generality, so for now M will be an arbitrary symplectic manifold. We'll return to the case of phase space soon.

Since ω puts a nondegenerate bilinear form on each tangent space, it gives an isomorphism between the tangent and cotangent spaces at each point of M , and therefore an isomorphism between vector fields and 1-forms. We will especially be interested in this isomorphism in the case where the 1-form is df for some function f . In this case, we'll write X_f for the unique vector field for which $\omega(Y, X_f) = df(Y)$ for all Y . (There is an arbitrary sign choice to make here — I could have said that $\omega(X_f, Y) = df(Y)$. As always happens with such things, this decision seems to about evenly split the authors of books on this subject. Hamilton's equations, discussed below, do have an arrangement of signs that everyone agrees on, and I've made choices in this section that are consistent with that.)

This vector field is sometimes called the **symplectic gradient** of f ; if we had a Riemannian metric instead of ω here then this construction would of course give the usual gradient. It's worth emphasizing, though, that while a Riemannian gradient of f (usually) gives a direction in which f is increasing, the symplectic gradient gives a direction in which f is constant, since $X_f(f) = df(X_f) = \omega(X_f, X_f) = 0$.

Given any vector field X at all on a smooth manifold, the existence and uniqueness of solutions of ODE's lets us define a **flow**, that is, a one-parameter group of diffeomorphisms $\phi^t : M \rightarrow M$ for which

$$\left. \frac{d}{dt} \right|_{t=0} \phi^t(a) = X|_a$$

for any point $a \in M$, where the notation $X|_a$ means the tangent vector we get by restricting X to a . (In general the flow might only be defined for t in some neighborhood of 0, but this will always be enough for our purposes.) The flow is used to construct the **Lie derivative** of a tensor field with respect to a vector field. In order to take any sort of derivative of a tensor field on a manifold it's necessary to be able to compare values of the tensor field at different points, and the flow gives us a way to do this. We define

$$\mathcal{L}_X(T) = \left. \frac{d}{dt} \right|_{t=0} (\phi^t)^*(T).$$

A vector field that arises as a symplectic gradient — that is, as X_f for some f — is called a **Hamiltonian vector field** and the corresponding flow is called a **Hamiltonian flow**. Note that since the definition of X_f depends on ω , in order for the Hamiltonian flow corresponding to a function to make sense, it's necessary for ω to be preserved by the flow. Otherwise after running time forward using the flow our vector field won't be X_f for the same f anymore!

So we'd like to characterize the X for which $\mathcal{L}_X \omega = 0$. To do this we invoke *Cartan's magic formula*, which says that $\mathcal{L}_X = \iota_X \circ d + d \circ \iota_X$. (Here ι_X is the **interior product** with X , which is the map from d -forms to $(d-1)$ -forms defined by $\iota_X \alpha(Y_1, \dots, Y_{d-1}) = \alpha(X, Y_1, \dots, Y_{d-1})$.) This is where we use the fact that ω is closed: we see that

$$\mathcal{L}_X \omega = \iota_X(d\omega) + d(\iota_X \omega) = d(\iota_X \omega).$$

If X corresponds to α under the isomorphism between vector fields and 1-forms given by ω , then $\iota_X \omega = -\alpha$ by definition, so we see that flowing along X_α preserves ω if and only if α is closed. In particular, since X_f corresponds to df , *all Hamiltonian flows preserve ω* .

It will be important for us to analyze how functions change along Hamiltonian flows; we will, in fact, basically be translating all the physical questions this framework can address into what values functions take along a Hamiltonian flow. That is, if X_f is a Hamiltonian vector field, a is a point in M , and g is a function on M , we'd like to compute dg/dt along the flow of X_f through a . By definition, this is just $X_f(g)$, so by the definition of X_g ,

$$\left. \frac{dg}{dt} \right|_a = X_f(g) = dg(X_f) = \omega(X_g, X_f).$$

This fact will turn out to be important enough to warrant a definition: we'll write $\{g, f\} = \omega(X_g, X_f)$ and call it the **Poisson bracket** of g and f . As we just saw, the Poisson bracket measures how g changes along the Hamiltonian flow corresponding to f . In particular, $\{g, f\} = 0$ if and only if f 's Hamiltonian flow preserves g . Note also that the Poisson bracket is antisymmetric (because ω is), which means that f 's Hamiltonian flow preserves g if and only if g 's Hamiltonian flow preserves f . (The Poisson bracket in fact turns out to put a Lie algebra structure on $C^\infty(M)$ — that is, it also satisfies the Jacobi identity — but we won't need this fact here.)

So, to summarize:

- Phase space, being the cotangent bundle of configuration space, has a natural symplectic structure. In coordinates, the symplectic form is given by $\omega = \sum dp_i \wedge dq_i$.

- On any symplectic manifold, we can associate to each function f a vector field X_f , and vector fields arising in this way are called Hamiltonian vector fields. Flowing along a Hamiltonian vector field always preserves the symplectic form.
- This construction lets us define the Poisson bracket $\{g, f\} = \omega(X_g, X_f)$, which measures both how f changes when flowing along X_g and how g changes when flowing along X_f .
- Since flowing along a vector field X preserves ω if and only if the corresponding 1-form α is closed, we can reverse this entire process if M is simply connected. In that case, $\alpha = df$ for some f , so $X = X_f$, and f is uniquely determined up to adding a constant. So if M is simply connected (or if not, then in an open neighborhood of any point), *there is a one-to-one correspondence between vector fields whose flow preserves ω and smooth functions on M modulo constants.*

2.4 Phase Space and Hamiltonians

We're now ready to see how this machinery can allow us to do physics. We fix a manifold Q called configuration space, and we write $P = T^*Q$ for its cotangent bundle, which we'll call phase space. The basic assumption of Hamiltonian mechanics is that the way we "run time forward" in our physical system is by following the Hamiltonian flow corresponding to a distinguished function H , which we'll call the **Hamiltonian**. That is, if our system is in state (q, p) at time t_0 and ϕ^t is the flow along X_H , then our system is in state $\phi^t(q, p)$ at time $t + t_0$.

Suppose we are using local coordinates $q_1, \dots, q_n, p_1, \dots, p_n$ in which the symplectic form can be written as $\omega = \sum_i dp_i \wedge dq_i$. Given two vector fields

$$X = \sum (a_i \partial_{q_i} + b_i \partial_{p_i}), \quad X' = \sum (a'_i \partial_{q_i} + b'_i \partial_{p_i}),$$

we get that $\omega(X, X') = \sum (b_i a'_i - a_i b'_i)$. I encourage the reader to verify that this means that for a function f ,

$$X_f = \sum_i \left(\frac{\partial f}{\partial p_i} \partial_{q_i} - \frac{\partial f}{\partial q_i} \partial_{p_i} \right),$$

and that the Poisson bracket is given by

$$\{f, g\} = \sum_i \left(\frac{\partial f}{\partial q_i} \frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i} \frac{\partial g}{\partial q_i} \right).$$

If we've chosen a Hamiltonian H , then the value of a function f evolves through time according to solutions of the differential equation $df/dt = \{f, H\}$. Plugging in q_i and p_i for f , we get **Hamilton's equations**:

$$\frac{dq_i}{dt} = \frac{\partial H}{\partial p_i}, \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}.$$

As we saw in the last section, Hamiltonian flows always preserve their corresponding function, so the Hamiltonian itself ought to measure some scalar quantity that doesn't change as time moves forward. In classical physics there's really only one such quantity to choose: the value Hamiltonian at a point in P ought to be physically interpreted as the total **energy** of the system when it is in that state.

In particular, consider the case where $\{f, H\} = 0$. This happens exactly when H 's Hamiltonian flow preserves f , that is, f is conserved by the laws of physics. But it is also equivalent

to the claim that f 's Hamiltonian flow preserves H , that is, flowing along X_f preserves H . This phenomenon gives us the Hamiltonian mechanics version of a result called *Noether's theorem*: going between X_f and f gives us a one-to-one correspondence between Hamiltonian vector fields whose flow preserves H (that is, vector fields whose flow preserves both H and ω) and scalar functions which are conserved by the laws of physics.

Let's see how to recover Newtonian mechanics. In mechanics problems, energy is usually given as a sum of two terms, one representing **kinetic energy**, written T , and one representing **potential energy**, written V . In our Newtonian example from above, the kinetic energy is the usual

$$T = \sum_i \frac{1}{2} m_i \left| \frac{d\mathbf{q}_i}{dt} \right|^2 = \sum_i \frac{|\mathbf{p}_i|^2}{2m_i},$$

and the potential energy is simply our potential function V . So our Hamiltonian all together is:

$$H(q, p) = T(p) + V(q) = \sum_i \frac{|\mathbf{p}_i|^2}{2m_i} + V(q),$$

and then, combining the three coordinates for each particle into a single vector, Hamilton's equations give us

$$\begin{aligned} \frac{d\mathbf{q}_i}{dt} &= \frac{\partial H}{\partial \mathbf{p}_i} = \frac{\mathbf{p}_i}{m_i} \\ \frac{d\mathbf{p}_i}{dt} &= -\frac{\partial H}{\partial \mathbf{q}_i} = -\frac{\partial V}{\partial \mathbf{q}_i}. \end{aligned}$$

Note that Hamilton's first equation exactly tells you how to compute the velocity of a particle once you know its momentum, which does something to address the concern we had earlier. Importantly, we see that this relationship depends on Hamiltonian; asking which velocity corresponds to a given momentum is meaningless until you've specified the laws of physics.

For our mechanical Hamiltonian, since the kinetic energy term is a homogeneous quadratic function of the momentum, we think of it as corresponding to a Riemannian metric on configuration space. In order to get agreement between the two ways of translating between velocity and momentum — using the inner product or going through Hamilton's first equation — we need to include the masses of the particles in the metric, so that in our case for two tangent vectors v, v' we have

$$\langle v, v' \rangle_T = \sum_i m_i \langle \mathbf{v}_i, \mathbf{v}'_i \rangle$$

where $\langle \cdot, \cdot \rangle$ is the usual inner product on \mathbb{R}^3 . This induces a metric on the cotangent space given by

$$\langle p, p' \rangle_T = \sum_i \frac{\langle \mathbf{p}_i, \mathbf{p}'_i \rangle}{m_i},$$

so following this convention the Hamiltonian would be written

$$H(q, p) = \frac{1}{2} \langle p, p \rangle_T + V(q).$$

2.5 Examples

2.5.1 The Harmonic Oscillator

First, let's consider a harmonic oscillator. This is a physical system with one degree of freedom q in which the potential energy has the form $\frac{1}{2}kq^2$ for some k . (The factor of $\frac{1}{2}$ is of course purely for convenience.) This is a decent model for, for example, a mass attached to a light, frictionless spring.

If the mass of this particle is m , then our Hamiltonian is $H = p^2/2m + kq^2/2$, and Hamilton's equations are

$$\frac{dq}{dt} = \frac{p}{m} \quad \frac{dp}{dt} = -kq.$$

This is of course a very easy pair of differential equations to solve: you get, writing $\alpha = \sqrt{k/m}$, that $q = A \sin(\alpha(t - t_0))$ and $p = A\alpha \cos(\alpha(t - t_0))$ for some A and t_0 .

So far this analysis is basically identical to what we would have gotten using regular Newtonian mechanics. Still, even though we just found a solution, we can get some practice with this machinery by performing a change of coordinates that makes the solution even easier. These solutions lie on the ellipse $(\alpha q)^2 + p^2 = A^2$, which suggests that we ought to rescale q and p and switch to polar coordinates.

So let's first try setting $r = \sqrt{kq^2 + p^2/m}$ and $\theta = \arctan(\sqrt{km}q/p)$; these are the polar coordinates corresponding to $\tilde{p} = p/\sqrt{m}$ and $\tilde{q} = \sqrt{k}q$. Sadly, this doesn't quite do what we want: these aren't canonical coordinates, that is, the symplectic form isn't $dr \wedge d\theta$. Indeed,

$$\omega = dp \wedge dq = \sqrt{\frac{m}{k}} d\tilde{p} \wedge d\tilde{q} = \sqrt{\frac{m}{k}} r dr \wedge d\theta.$$

It would be possible to work out the form of the Poisson bracket in these coordinates and see what equations we get, but it's even easier to just find coordinates that are canonical and use those. We can do this by replacing r with $s = \frac{1}{2}\sqrt{m/kr^2} = r^2/2\alpha$. We then have $\omega = ds \wedge d\theta$ and $H = \alpha s$, and so Hamilton's equations are

$$\frac{d\theta}{dt} = \frac{\partial H}{\partial s} = \alpha \quad \frac{ds}{dt} = -\frac{\partial H}{\partial \theta} = 0.$$

This analysis makes it obvious that s is a conserved quantity — that's literally what the second equation says. This is equivalent to saying that $\{s, H\} = 0$, which we could have checked in the original coordinates if we wanted. In this case this is all kind of silly, since s is just a constant multiple of H ; the next example will feature a less silly version of this phenomenon.

2.5.2 The Two-Body Problem

Consider two particles, with masses m_1 and m_2 , moving under the influence of a conservative force that depends only on their relative positions, that is, on the difference $\mathbf{q}_1 - \mathbf{q}_2$. (You might imagine for example two celestial bodies moving under the influence of gravity.) So our configuration space is $\mathbb{R}^3 \times \mathbb{R}^3$, and our Hamiltonian is

$$H = \frac{|\mathbf{p}_1|^2}{2m_1} + \frac{|\mathbf{p}_2|^2}{2m_2} + V(\mathbf{q}_1 - \mathbf{q}_2)$$

for some function V .

We can already see another case of Noether's theorem here. The fact that V depends only on $\mathbf{q}_1 - \mathbf{q}_2$ means that if we translate both particles by the same vector and leave their momenta fixed, H is unchanged. For concreteness let's consider translating in the positive x direction; this corresponds to flowing along the vector field $\partial_{x_1} + \partial_{x_2}$ (writing x_i for the x component of \mathbf{q}_i). These translations also self-evidently preserve the symplectic form, and so our vector field must be Hamiltonian. And indeed, it's X_f for $f = (p_1)_x + (p_2)_x$, the x component of the total momentum of the system. You could also check directly that the Poisson bracket $\{f, H\}$ is zero. So we see that a Hamiltonian that is preserved by translations in some direction corresponds to physics that preserve the component of total momentum in that direction.

There is a common change of coordinates that makes this system a bit easier to analyze: write

$$\mathbf{Q} = \frac{m_1 \mathbf{q}_1 + m_2 \mathbf{q}_2}{m_1 + m_2} \quad \mathbf{q} = \mathbf{q}_1 - \mathbf{q}_2.$$

Now, given *any* diffeomorphism f from a manifold Q to itself, we can lift it to a diffeomorphism on the cotangent bundle by setting $f^\sharp(q, p) = (f(q), (f^{-1})^*(p))$. We call f^\sharp the **cotangent lift** of f . It turns out that a diffeomorphism on a cotangent bundle has the form of a cotangent lift if and only if it preserves the canonical 1-form θ . To compute f^\sharp in coordinates, first note that $(f^{-1})^*(p)(v) = p((f^{-1})_*(v)) = p((f_*)^{-1}(v))$ by definition, so the matrix for $(f^{-1})^*$ is the transpose of the inverse of the Jacobian of f .

So in particular, the cotangent lift gives us a natural way to turn any diffeomorphism on configuration space into a symplectomorphism on phase space. Once can check that doing this for our change of coordinates here gives us the momentum coordinates

$$\mathbf{P} = \mathbf{p}_1 + \mathbf{p}_2 \quad \mathbf{p} = \frac{m_2 \mathbf{p}_1 - m_1 \mathbf{p}_2}{m_1 + m_2},$$

and our Hamiltonian becomes

$$H = \frac{|\mathbf{P}|^2}{2M} + \frac{|\mathbf{p}|^2}{2m} + V(\mathbf{q}),$$

where $M = m_1 + m_2$ and $m = m_1 m_2 / (m_1 + m_2)$. (The reader is encouraged to verify these computations; it's good practice!)

The point of this change of coordinates was to "decouple" the two parts of the Hamiltonian. The coordinate \mathbf{Q} is called the **center of mass** of the system; what we've shown is that our original system is equivalent to one with a free particle of mass M moving with the center of mass and a particle of mass m moving under the influence of the potential V .

2.5.3 Central Potentials

As one more example of the relationship between symmetries and conservation laws, let's consider a particle moving in a potential that depends only on the distance of that particle from the origin. That is,

$$H = \frac{|\mathbf{p}|^2}{2m} + V(|\mathbf{q}|).$$

This is a good model for a planet moving around the sun under the influence of Newtonian gravity; in this case we'll have $V(r) = -GMm/r$, where M is the mass of the sun and G is the gravitational constant.

But no matter what V is, the fact that it depends only on the length of \mathbf{q} means that the physics is preserved by any rotation about the origin. It's worth being precise about what we

mean by this: rotation about the origin is a diffeomorphism on configuration space, and to extend it to a symplectomorphism on phase space we need to take its cotangent lift.

If R_θ is the rotation by θ around the z axis, then

$$(R_\theta)^\sharp(x, y, z, p_x, p_y, p_z) = (\cos \theta x - \sin \theta y, \sin \theta x + \cos \theta y, z, \quad (1)$$

$$\cos \theta p_x - \sin \theta p_y, \sin \theta p_x + \cos \theta p_y, p_z). \quad (2)$$

(The transpose of the inverse of R_θ is just R_θ itself, since R_θ is orthogonal.) This is the map that has to preserve the Hamiltonian if our analysis is to go through, which means it's important that H depends only on $|\mathbf{p}|$ and $|\mathbf{q}|$.

To get the vector field whose flow produces this symmetry, we take the derivative of this with respect to θ at $\theta = 0$. We get

$$X = -y\partial_x + x\partial_y - p_y\partial_{p_x} + p_x\partial_{p_y},$$

which is X_{L_z} where $L_z = xp_y - yp_x$.

We call L_z the **angular momentum** of our particle about the z axis, and this analysis shows that any physics arising from a Hamiltonian which is symmetric under rotations about the z axis.

2.5.4 A Note

It is easy to construct Hamiltonians which aren't invariant under rotations or translations. Indeed, the one from the last example isn't preserved by translations, and correspondingly we shouldn't expect momentum to be conserved by, say, Newtonian gravity. Nonetheless, it's believed by most physicists that the fundamental laws that the universe runs on, whatever they are, do have these two symmetries — the results of a physical experiment don't depend on where you do it or which way you were facing — and that therefore conservation of linear and angular momentum hold in general.

If you are presented with a Hamiltonian that doesn't have this symmetry, like the one in the last example, the assumption is that there's some part of the physics that you're neglecting, and that if you included it the symmetry would appear again. For example, if we imagine the last example to be about a planet moving around the sun, we are neglecting the influence of the planet's gravity on the sun, and if we included it we would be in the situation from the previous example about the two-body problem.

There is another symmetry that classical physics obeys: it also shouldn't matter *when* an experiment is performed. Under our formalism, time translation comes from flowing along the vector field given by H itself, so this symmetry corresponds to the conservation of energy. This example is a bit different from the others, though, because the relationship is true by definition! This is an artifact of the way we set up the Hamiltonian formalism: it picks out time translation as "special," as the flow that corresponds to the Hamiltonian, and specifying the Hamiltonian is the way we specify the laws of physics.

Like with momentum, it is possible to "break" the time translation symmetry (and therefore energy conservation) by using a Hamiltonian that depends explicitly on time. This is useful when the forces acting on the particles or the constraints of the physical system change over time. (An example of the latter that's often trotted out in physics classes is a bead attached to a spinning circle of wire.) I've chosen not to consider the case of time-dependent Hamiltonians or Lagrangians in this article for simplicity, but the theory does continue to work just fine in that setting.

3 Lagrangian Mechanics

Recall that Hamilton's equations are given by

$$\frac{dq_i}{dt} = \frac{\partial H}{\partial p_i}, \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}.$$

As I mentioned briefly in that section, the first equation can give us a sort of answer to the question we had earlier about the relationship between momentum and velocity: it supplies, for every point $(q, p) \in T^*Q$ a tangent vector $v \in T_qQ$, and it tells us to interpret that tangent vector as a velocity. Provided that this procedure is invertible, which it will be in all of the cases we care about, we can think of it as giving us a “change of coordinates” from the cotangent bundle to the tangent bundle. This will turn out to give us another formulation of mechanics, called **Lagrangian mechanics**, which, while formally equivalent to everything we've done so far, sheds light on different aspects of mechanics than the Hamiltonian picture.

3.1 The Legendre Transform

We'll take this assignment of tangent vectors to points in phase space as our starting point. If the tangent vector v comes from (q, p) in this way, our goal will be to rewrite Hamilton's equations in a way that depends on q and v rather than on q and p . Put another way, the Hamiltonian gave us a way to turn paths in the cotangent bundle into paths in the tangent bundle, and we'd like to see what restrictions Hamilton's equations impose directly on these new paths.

Hamilton's first equation was, in some sense, “used up already” in the definition of v . By switching coordinates to v and interpreting v as a velocity, this equation tells us simply that $v = dq/dt$, that is, at every point $(q, v) = \gamma(t)$ along our path, we should have $v = \gamma'(t)$. This means that we might as well just talk about paths in configuration space rather than its tangent bundle; we can identify such a path with its lift to the tangent bundle and do away with one of our two equations.

So it remains to translate the second equation into something to do with v . Note that the q coordinate has very little to do with our goal here: we're trying to turn a statement about the cotangent bundle into a statement about the tangent bundle, and all of the action is happening in the fibers of these two bundles. So it will be cleaner to ignore q for now and examine how our coordinate change procedure behaves on a general vector space.

Suppose we have a smooth function H on a vector space U . (U will end up being the cotangent space at a point of configuration space.) Then for each point $p \in U$, dH gives a linear map from the tangent space T_pU to \mathbb{R} . But since U is a vector space, there is a canonical identification of each of its tangent spaces with U itself, so we can in fact think of dH as giving us a way of assigning, to each $p \in U$, a linear map from U to \mathbb{R} , that is, an element of U^* .

So H gives us a map $W_H : U \rightarrow U^*$. (It's important to emphasize that W_H has no reason to be linear, so this is not any sort of inner product.) Geometrically, W_H takes p to the element of U^* corresponding to the linear part of the linear approximation to H at p .

Suppose now that W_H is invertible. (This will happen, for example, if H is a constant plus a nondegenerate quadratic form in p , as is the case for our mechanical Hamiltonians from the last section.) Then it turns out that W_H^{-1} arises in the same way as W_H : there is a function L so that $W_L = W_H^{-1}$. We call L the **Legendre transform** of H .

In fact, L can be computed explicitly: one can show that when W_H is invertible, $W_L = W_H^{-1}$ if and only if

$$L(W_H(p)) + H(p) = \langle W_H(p), p \rangle$$

up to an additive constant, where $\langle \cdot, \cdot \rangle$ is the pairing between U^* and U . (Usually we let the constant be zero.) In particular, this makes it clear that the Legendre transform is an involution: H is also the Legendre transform of L .

To see all this, first note that by definition, for any $p \in U$, we have $\langle W_H(p), p' \rangle = dH_p(p')$, where on the right hand side we think of p' as living in the tangent space at p . So, taking the derivative of both sides and pairing with an arbitrary p' , we get that they are equal if and only if

$$\langle (DW_H)_p(p'), W_L(W_H(p)) \rangle = \langle (DW_H)_p(p'), p \rangle,$$

where $(DW_H)_p$ is the total derivative of W_H at p . Since W_H is invertible, the left side of this pairing can be anything, so this is true if and only if $W_L(W_H(p)) = p$. The reader is encouraged to fill in the missing steps of this argument; it's a good exercise in following all the relevant definitions.

So what did this giant mess of symbols get us? To any function H on U we've associated a new function L on U^* so that their "coordinate-change functions" are inverses of each other. In coordinates p_1, \dots, p_n on U and v_1, \dots, v_n on U^* , this means that

$$v_i = \frac{\partial H}{\partial p_i}, \quad p_i = \frac{\partial L}{\partial v_i},$$

and

$$L = \langle v, p \rangle - H.$$

In the case we're interested in, when we're doing this in every fiber of the cotangent bundle and H is the Hamiltonian of some physical system, we call L the **Lagrangian** of that same system.

What does this look like for the mechanical Hamiltonians we were working with before? There we had $H = T + V = \sum_i \frac{|\mathbf{p}_i|^2}{2m_i} + V(q)$. So we get $\mathbf{v}_i = \partial H / \partial \mathbf{p}_i = \mathbf{p}_i / m_i$ and $\langle v, p \rangle = \sum_i \frac{|\mathbf{p}_i|^2}{m_i} = 2T$. This means the Lagrangian turns out to be $L = 2T - (T + V) = T - V$.

It's worth stressing again that despite the fact that we've arrived at an expression for L that looks very similar to H , there is an additional important difference between the two aside from the fact that the sign on V has flipped: H is a function on the cotangent bundle and L is a function on the tangent bundle! The relationship between momenta and velocities — that is, between p and v — depends entirely on the physics being modeled, so unless you've picked a Hamiltonian or a Lagrangian this relationship remains unspecified. Only when this relationship has been established does it even make sense to write something like $L + H = 2T$; if we were being more careful we would actually write something like $L(q, W_H(p)) + H(q, p) = 2T(p)$.

Recall that Hamilton's first equation now just tells that the tangent vector we pick at every point of our path should be the time derivative of the path at that point, so we are just left with translating the second into a statement about Lagrangians. That equation was

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}.$$

Now, since we performed our Legendre transform just in the fibers of the cotangent and tangent bundles, nothing interesting happened to derivatives with respect to q coordinates, so the fact that $L = \langle v, p \rangle - H$ means that $\partial L / \partial q_i = -\partial H / \partial q_i$. This, combined with the fact that $p_i = \partial L / \partial v_i$ gives us **Lagrange's equation**:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial v_i} \right) = \frac{\partial L}{\partial q_i}.$$

It is more common for authors to write \dot{q}_i where I've written v_i here, using the usual physicists' convention of dots to indicate time derivatives. The reason I didn't do this was to avoid a common confusion: when you write expressions like $\partial L/\partial \dot{q}_i$ it's tempting to assume that one is supposed to compute \dot{q}_i from q_i or something. But the Lagrangian is a function on the tangent bundle, not just on configuration space, and $\partial L/\partial v_i$ is just a derivative with respect to one of the coordinates on the tangent bundle. Given a smooth path in configuration space there is a natural way to lift it and obtain a path in the tangent bundle, and the physical assumption we are making is that the paths that happen physically are exactly the ones whose lifts satisfy Lagrange's equation.

Note that we end up with half as many equations as before — we have one for every position coordinate, rather than one for every position or momentum coordinate. But because the Lagrangian depends on both q and v and we fix v_i to be the time derivative of q_i , these end up being second-order differential equations rather than the first-order Hamilton's equations, so we still need the same amount of information in our initial conditions to solve them as before.

3.2 An Example

Let's look at a concrete example of an at least somewhat nontrivial mechanics problem and see how to solve it using the Lagrangian formalism. This problem still would be feasible to tackle using the techniques from a Newtonian mechanics class, but the Lagrangian approach makes it quite straightforward.

Consider a two-dimensional world with a mass hanging from a very light spring. The end of the spring without the mass is fixed in place and the other end is free to swing around. We'll pick coordinates r, θ for our configuration space, where r is the current length of the spring and θ is the angle the spring makes with the vertical. We'll write the corresponding time derivatives as v_r and v_θ . These coordinates have the nice property that their time derivatives are always perpendicular, so the speed of the particle is $\sqrt{v_r^2 + (rv_\theta)^2}$. Therefore, the kinetic energy is simply $T = \frac{1}{2}mv_r^2 + \frac{1}{2}mr^2v_\theta^2$.

There are two contributions to the potential energy: gravity and the restoring force from the spring. Springs are well modeled by potentials of the form $V_{\text{spring}} = \frac{1}{2}k(r - \ell)^2$ for some constant k , where ℓ is the "natural length" of the spring. We encountered a potential of this form when we discussed the harmonic oscillator. Gravity (at least in cases like this where we can neglect the varying distances from the center of the earth) produces a constant acceleration in all falling bodies, and I encourage you to check that this is the same as asserting that $V_{\text{gravity}} = mgh$ where g is that constant acceleration and h is the height of a particle above an arbitrary reference height.

Putting this all together, we get

$$L = \frac{1}{2}m(v_r^2 + r^2v_\theta^2) - \frac{1}{2}k(r - \ell)^2 + mgr \cos \theta.$$

We can plug this into Lagrange's equation and get, simplifying a bit and using physicists' dot notation for derivatives,

$$m\ddot{r} = m\dot{\theta}^2 + mg \cos \theta - k(r - \ell)$$

and

$$mr\ddot{\theta} + 2m\dot{r}\dot{\theta} = -mgr \sin \theta.$$

3.3 The Calculus of Variations

There is another, very different way to obtain the Lagrange's equation involving a technique called the **calculus of variations**. The calculus of variations is a sort of infinite-dimensional calculus performed on spaces of functions rather than finite-dimensional vector spaces. It provides tools for doing things like finding local minima or maxima of some functional on such a space of functions. I'll sketch the part of this story that produces Lagrangian mechanics.

Suppose we are given an arbitrary function $L : \mathbb{R} \times TQ \rightarrow \mathbb{R}$, where we think of the first \mathbb{R} as representing a time coordinate. Given any path $\gamma : [0, 1] \rightarrow Q$, we consider the following quantity, called the **action**:

$$S(\gamma) = \int_0^1 L(t, \gamma(t), \gamma'(t)) dt.$$

(Here again $t \mapsto (\gamma(t), \gamma'(t))$ is the natural lift of γ to a path in the tangent bundle of Q ; we are abusing notation slightly writing the components the way we are here.) We want to find, out of all paths γ with a fixed starting and ending point, the ones that locally minimize or maximize S .

We'll answer this by considering smooth homotopies $h : (-\epsilon, \epsilon) \times [0, 1] \rightarrow Q$ for which $h(0, t) = \gamma(t)$ and which leave the endpoints fixed. As usual, we'll write $h_u(t) = h(u, t)$. If γ minimizes the action, then it ought to in particular be a local minimum within any such homotopy. That is, 0 should be a critical point of the function $u \mapsto S(h_u)$.

For ease of notation from here on out, I'll write $\tilde{\gamma}(t) = (t, \gamma(t), \gamma'(t))$ for the argument to S , so $S(\gamma) = \int_0^1 L(\tilde{\gamma}(t)) dt$.

So we need

$$\left. \frac{d}{du} \right|_{u=0} \int_0^1 L\left(t, h(u, t), \frac{\partial h}{\partial t}(u, t)\right) dt = 0,$$

and we can turn this into a condition which doesn't mention h with a bit of computation. We can pull the derivative inside the integral sign and use the chain rule to get that the left side is

$$\int_0^1 \sum_{i=1}^n \left[\frac{\partial h_i}{\partial u}(0, t) \frac{\partial L}{\partial q_i}(\tilde{\gamma}(t)) + \frac{\partial^2 h_i}{\partial u \partial t}(0, t) \frac{\partial L}{\partial v_i}(\tilde{\gamma}(t)) \right] dt$$

for any coordinate system $(q_1, \dots, q_n, v_1, \dots, v_n)$ on the tangent bundle; here h_i is the q_i component of h . Then, using integration by parts, we can transform the second term a bit more, turning it into:

$$\sum_{i=1}^n \left[\left(\frac{\partial h_i}{\partial u}(0, t) \frac{\partial L}{\partial v_i}(\tilde{\gamma}(t)) \right) \Big|_{t=0}^1 - \int_0^1 \frac{\partial h_i}{\partial u}(0, t) \frac{d}{dt} \left(\frac{\partial L}{\partial v_i}(\tilde{\gamma}(t)) \right) dt \right].$$

The fact that the endpoints are fixed means that the boundary term on the left is zero: at $t = 0$ and $t = 1$, $h_u(t)$ doesn't vary as I change u , so $\partial h / \partial u = 0$. So we're left with just the second term which we combine with the missing term above to get that we want

$$\int_0^1 \sum_{i=1}^n \frac{\partial h_i}{\partial u}(0, t) \left[\frac{\partial L}{\partial q_i}(\tilde{\gamma}(t)) - \frac{d}{dt} \left(\frac{\partial L}{\partial v_i}(\tilde{\gamma}(t)) \right) \right] dt = 0.$$

Finally, we note that, since $\partial h / \partial u$ could be any vector at all at each t , the only way that integral is zero for every h is if, for each i , the part in square brackets is identically zero as a function of t . But that is exactly the same as saying that γ must satisfy Lagrange's equation!

So we get a description of the laws of physics that looks much more “global” than the one from symplectic geometry above: a physical system gets from one point in configuration space to another by following a path that is a critical point of the action functional S . The original description told us how the position and momentum (or velocity) evolve through time if we know them at one moment, whereas this new description puts a restriction on the entire path at once. We’ll have more to say about this in the next section.

4 A Little Philosophy

The first time I encountered Noether’s theorem gave me a particular feeling about physics that I’ll do my best to explain here. There is a sense in which, for example, the fact that the laws of physics are invariant under rotation is “obvious” and fact that they conserve angular momentum isn’t. After all, an expression like $L_z = xp_y - yp_x$ looks a little arbitrary. Why should it be a fundamental of nature that that’s conserved rather than, say, $xp_y + yp_x$?

Now that we’ve built up the machinery that connects rotation and angular momentum, it’s very tempting to say that we’ve “reduced” the fact that angular momentum is conserved to “merely” the fact that physics is rotationally symmetric. This is, of course, not true: we have reduced it to the fact that physics is rotationally symmetric *and* the fact that physics can be described using Hamiltonian mechanics. (There is of course a Lagrangian version of Noether’s theorem too, and in fact this is the form in which it was originally stated.) Indeed, trying to prove Noether’s theorem from Hamiltonian mechanics gets things a bit backwards: the fact that something like Noether’s theorem is true is baked into the fact that all the flows we consider are along Hamiltonian vector fields.

This is emphatically a physical and not a mathematical assertion — it is certainly possible to imagine a universe where physics doesn’t work this way, but this does not seem to be the universe we live in. We can’t hope to prove that conservation of angular momentum is “the only way it could logically have been” without making some assumptions about the nature of the laws of physics. What we can do, and what in some sense is the main goal of all of physics, is to try to find a way to state those laws with as few moving parts as possible, and I think that at least is something that this whole discussion manages to do. In either math or physics, any time you can take an arbitrary-looking algebraic expression and show that it falls out naturally from less arbitrary-looking mathematical objects you’ve made some progress toward making the world more intelligible.

This is partly why I took the particular path through the material that I did. It’s actually much more common in mechanics texts to do the Lagrangian version first and derive the Hamiltonian picture from it using a Legendre transform. I always had some trouble making this approach fit in my head without seeming like magic. Until you’ve seen how it fits into the story it’s unclear why the Lagrangian has anything to do with anything, and especially unclear why one would want to minimize the action. Even if, as I keep saying, it’s necessary to make up some physical assumption at some point to make this all go, I feel like the assumptions made here feel at least a little less made up; the reader is of course free to disagree.

As I sort of said when discussing it above, it isn’t quite as “magical” as it might first appear — the claim is not that the universe searches over all possible paths for the global minimum of the action, but rather that the one that occurs is a critical point of the action under any perturbation. Sometimes a path with this property is called “stationary.” So the rule is not quite as global as it might be — it is global in the sense of pertaining to the whole path at once rather than a particular point along it, but not in the sense of pertaining to all possible paths, even those far

away from the one under consideration.

Still, it was very striking to people in the early 19th century that Newtonian mechanics could be recast in such a strange way, and I think this reaction makes a lot of sense. The fact that the action-principle version is mathematically equivalent to the more local description shouldn't detract from this; even in pure mathematics one can often learn a great deal from expressing the same fact in two different ways. A lot of this work happened around the end of the 18th century, and, in keeping with the Enlightenment-era times, many authors took the action principle as evidence of the benevolent guiding hand of Nature making sure the particles don't waste too much of their precious action on frivolous non-stationary paths or something.

This is, of course, not the way most modern physicists talk, but there is a different sense in which we now understand the story told in this article to be offering us a hint about something beyond Newtonian mechanics. Basically all the theories of modern physics, including quantum mechanics, quantum field theory, general relativity, and more speculative extensions to them like string theory, are most naturally expressed not in terms of forces and equal and opposite reactions but in terms of Lagrangians and Hamiltonians.