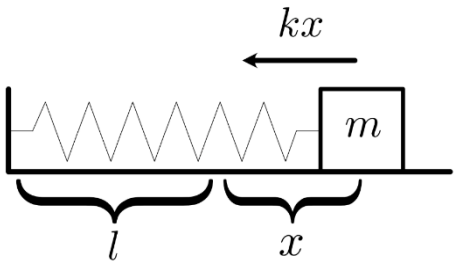
To Master Physics, First Learn the Harmonic Oscillator

Of all the systems you're likely to meet in your first physics class, far and away the most important is the **simple harmonic oscillator**—in other words, the basic setup of a block attached to a spring. No other system so thoroughly permeates virtually every corner of physics, from classical mechanics to quantum mechanics to quantum field theory. It's a setup that you are guaranteed to meet again and again throughout your physics education, and so it's a good idea to learn it well from the get-go. But why is this seemingly innocuous problem so ubiquitous in so many areas of physics?

We'll start by reviewing the basic setup, and after that we'll get into why it is that it's so important. We've got a block of mass "m" sitting on a flat table, which we'll take to be frictionless for simplicity. It's hooked up to a spring, which is itself anchored to a wall on the other end. If you give the block a little kick, or if you pull it out a way and then let it go, it will oscillate back and forth from side to side like a sine or cosine function. We call these sinusoidal oscillations **simple harmonic motion**.



Here's how we understand all this using Newton's laws. Whenever we move the block, we stretch or compress the spring, which therefore tries to pull or push the block back to its equilibrium position. The **equilibrium** is the position where the spring is happy and relaxed; it's neither compressed nor stretched, and so it exerts no force on the mass. That position will depend on the spring we're using, so let "l" denote the equilibrium length of the spring.

Now suppose the system is *not* in equilibrium; for example, suppose that we pull the block over to the right by a distance "x". Now the spring is stretched, and so it's going to try to pull the block back over to the left toward equilibrium. How hard it pulls will be proportional to how far we've moved the block, as well as how *stiff* the spring is. The stiffness is measured by another parameter called the **spring constant** "k", which determines how big a force the spring exerts when we've stretched it by a distance "x":

$$F = -kx.$$

Then the $F = ma$ equation for the block is

$$ma = -kx.$$

If you've learned some calculus, then you know that the acceleration $a(t)$ is nothing but the second derivative of the position $x(t)$ with respect to time. And so we can alternatively write this equation as

$$\frac{\mathrm{d}^2x}{\mathrm{d}t^2} = -\Omega^2 x.$$

I've also introduced another symbol here, capital omega $\Omega = \sqrt{k/m}$, because this combination is so important, as we'll see in a second.

So to solve for the trajectory of the block, we just have to ask ourselves what kind of function $x(t)$, when we take its derivative twice in a row, gives us back the same function times this negative number $-\Omega^2$. The answer is a sine or cosine! Recall that the derivatives,

$$\frac{\mathrm{d}}{\mathrm{d}t}\sin(\Omega t) = \Omega\cos(\Omega t) \qquad \text{and} \qquad \frac{\mathrm{d}}{\mathrm{d}t}\cos(\Omega t) = -\Omega\sin(\Omega t).$$

Those factors of Ω on the right come from the chain rule: the derivative of $\sin(\Omega t)$ with respect to its argument is $\cos(\Omega t)$, and then we multiply by the derivative of Ωt with respect to t, which contributes the factor of Ω. Now taking the second derivative, we get

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\sin(\Omega t) = -\Omega^2 \sin(\Omega t)$$

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\cos(\Omega t) = -\Omega^2 \cos(\Omega t),$$

which is exactly what we wanted: the second derivative is $-\Omega^2$ times the function that we started with. Then the general solution to the $F = ma$ equation is

$$x(t) = x_0 \cos(\Omega t) + \frac{v_0}{\Omega}\sin(\Omega t),$$

where x_0 and v_0 are the initial position and velocity of the block—how far did we pull the mass out at $t = 0$ and/or how fast did we kick it?

So we indeed find that the block sinusoidally oscillates back and forth around its equilibrium position in simple harmonic motion. The rate at which it oscillates is determined by $\Omega = \sqrt{k/m}$, which we therefore call the **natural frequency** of the system. Ω had better have dimensions of one over time since it's supposed to be a frequency, so let's check:

$$\sqrt{\frac{k}{m}} \sim \sqrt{\frac{\mathrm{N/m}}{\mathrm{kg}}} = \sqrt{\frac{1}{\mathrm{s}^2}} = \frac{1}{\mathrm{s}},$$

Ω tells you how fast the block will oscillate once you set it moving. Notice that you don't get to pick this oscillation speed! It's fixed by the stiffness of the spring and the mass of the block, and it will always oscillate at this same rate. (Unless you do something like kick the block so hard that it breaks the spring or crashes into the wall.)

There's one more thing we need to review before we step back and discuss why this simple system is so prevalent, and that's the potential energy of the spring. As the block slides back and forth, it's constantly speeding up or slowing down, and so its kinetic energy $\frac{1}{2}mv^2$ is changing. Its *total* energy, though, kinetic plus potential, is a constant, where the potential energy is given by $U = \frac{1}{2}kx^2$. We can check that by taking the rate of change of the total energy:

$$\frac{\mathrm{d}}{\mathrm{d}t}\{\frac{1}{2}m(\frac{\mathrm{d}x}{\mathrm{d}t})^2 + \frac{1}{2}kx(t)^2\} = m\frac{\mathrm{d}x}{\mathrm{d}t}\frac{\mathrm{d}^2x}{\mathrm{d}t^2} + kx\frac{\mathrm{d}x}{\mathrm{d}t}.$$

We therefore get,

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \frac{\mathrm{d}x}{\mathrm{d}t}(m\frac{\mathrm{d}^2x}{\mathrm{d}t^2} + kx),$$

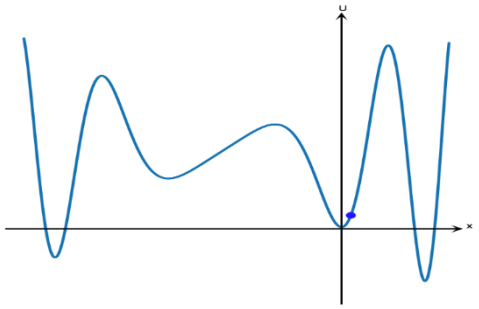
but the thing in parentheses vanishes by $F = ma$! Therefore, the total energy is indeed a constant. Notice that the force $F = -kx$ is related to the potential energy $U = \frac{1}{2}kx^2$ by

$$F = -\frac{\mathrm{d}U}{\mathrm{d}x}$$

In other words, the force is equal to minus the slope of the potential energy curve. That's the general relationship between force and potential energy. Notice that it means that in equilibrium, i.e. where the force is equal to zero, the slope of the potential must vanish.

In fact, you can think about the potential energy curve as if it were a frictionless hill with a particle sliding along it—in the case of the harmonic oscillator, the hill is a parabola $U \propto x^2$.. The equilibrium point is at the bottom of this well; if you set a particle at rest there at $x = 0$, it will happily sit there forever. If you set the particle down away from the equilibrium, however, it will rock back and forth about the bottom of the well, in the same way that the block oscillates back and forth around its equilibrium position.

Thinking about the potential energy curve as a hill is a very powerful way of quickly developing your intuition for how a particle in a general potential will behave, without ever doing the hard work of trying to solve the $F = ma$ equation.



Which brings us to the reason why the harmonic oscillator is so incredibly important and far-reaching. Consider a particle moving in *any* potential $U(x)$, like the random example above. We would like to solve once again the $F = ma$ equation,

$$m\frac{\mathrm{d}^2x}{\mathrm{d}t^2} = -\frac{\mathrm{d}U}{\mathrm{d}x},$$

to determine the trajectory $x(t)$. But outside of a handful of potential energy functions described in textbook examples, finding a simple solution to this equation for a complicated potential like this will be a hopeless task.

So how can we make progress? The first thing you should do when someone hands you a potential energy curve is to identify its stable equilibrium points. An equilibrium is again a point where the force on the particle vanishes, and so the slope of the potential will be zero there. A *stable* equilibrium point is one that's at the bottom of a well, as opposed to the top of a hill.

We can of course write down the *exact* solution to the $F = ma$ equation for a particle that's set down at rest at an equilibrium point x_{eq}: the force vanishes, and so the particle will just happily sit there forever!

$$x(t) = x_{\mathrm{eq}}.$$

On the other hand, if you give the particle at a stable equilibrium a little tap, it will oscillate back and forth around the bottom of the potential, just like the block on a spring oscillates around the bottom of its parabolic well.

You'll see this behavior all the time in your daily life if you pay attention. When you rock gently in a chair you're oscillating around a stable equilibrium. When you drop an ice cube in a glass of water it bobs up and down around its equilibrium height. Your coat hanging on the hook is swaying slightly around its equilibrium axis. And the basic reason is that when you're nearby a stable equilibrium point of almost any potential, the bottom of the hill looks just like the parabola of the simple harmonic oscillator.

We can expand any potential $U(x)$ in powers of x around a stable equilibrium point using a Taylor series, like you might have learned about in your math classes. First, let's slide our origin over so that it's on top of whatever equilibrium point we're interested in—that will make the formula simpler:

$$U(x) = U(0) + U'(0)x + \frac{1}{2!}U''(0)x^2 + \frac{1}{3!}U'''(0)x^3 + \cdots.$$

Here's what's going on in this formula. First of all, if you're sitting right at $x = 0$, then when you evaluate the potential, you're obviously going to get $U(0)$.. If you take a tiny step away from there, so that x is a small but non-zero number, the potential is still going to be close to $U(0)$, but now it will be shifted slightly away by the slope $U'(0)$ times x. Then

$$U(x) \approx U(0) + U'(0)x$$

is a very good approximation to our function near $x = 0$.. We use this approximation all the time in physics. But it's actually not that useful to us in this particular case, because remember that we've put our origin at a stable equilibrium point, where $U(0) = 0$ and $U'(0) = 0$. To get a better approximation, we have to include the terms with higher powers of x in the Taylor series. By including more and more of these terms, we get a better and better approximation to our function, as you can see by dragging the slider below. And when we add up *all* the infinite number of terms, we reproduce the exact function!

But we don't need an infinite number of terms to get a good approximation to our potential near the stable equilibrium point. We just need the first non-zero one—the quadratic term:

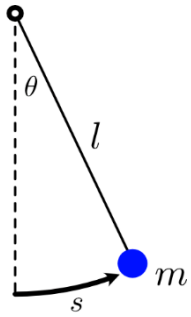
$$U(x) \approx \frac{1}{2} U''(0) x^2.$$

That's the potential energy of a spring, with spring constant $k = U''(0)$! *Then almost any potential energy function, however complicated, reduces to the simple harmonic oscillator potential in the neighborhood of its stable equilibrium points*. (The "almost" exception is for the rare case when $U''(0)$ *also* happens to vanish, in which case we need to include the next term in the series to get a better approximation.)

And this is why the simple harmonic oscillator is so prevalent. Systems tend to settle into stable equilibrium, and small disturbances make them oscillate around it. So the first thing we should do with any potential energy function is find its stable equilibrium points, and then ask what happens when we perturb slightly away from them. A particle that's released there will oscillate around the equilibrium in simple harmonic motion, with natural frequency.

$$\Omega = \sqrt{\frac{U''(0)}{m}}$$

To understand the physics farther away from the equilibrium points is usually much harder. The bigger x gets, the more important the higher power corrections in the Taylor series like x^3 and x^4 become. But we can often get approximate solutions by treating them as perturbations of our harmonic oscillator solution.



You've seen all this in action before if you've studied the simple pendulum, though you may or may not have seen it presented in this language. If a pendulum is inclined at an angle $\theta = s/l$, where s is the arc length coordinate that the particle has traced out, then the mass will be at a height $h = l - l\cos\left(\frac{s}{l}\right)$ above the bottom. (I'll write the coordinate as "s" instead of "x" so we don't get it confused with the xy coordinates of the point.) Then the potential energy function $U = mgh$ is,

$$U(s) = mgl\left(1 - \cos\left(\frac{s}{l}\right)\right).$$

The stable equilibrium is of course at $s = 0$, where the pendulum hangs straight down at rest. And we can check by taking the derivative:

$$U'(s) = mg\sin\left(\frac{s}{l}\right) \Rightarrow U'(0) = 0.$$

$$U''(s) = \frac{mg}{l}\cos\left(\frac{s}{l}\right) \Rightarrow U''(0) = \frac{mg}{l}.$$

Note that it's positive—that's how we know we're at a stable equilibrium, as opposed to the *top* of the arc, where the slope of the potential also vanishes but the second derivative is negative. That's an *unstable* equilibrium because if you take a tiny step away from it the pendulum will swing far away.

Our Taylor expansion around equilibrium is then,

$$U(s) \approx \frac{1}{2}\frac{mg}{l} s^2,$$

describing a simple harmonic oscillator with natural frequency

$$\Omega = \sqrt{\frac{g}{l}},$$

which is the familiar frequency of a pendulum! Taylor expanding around the equilibrium in this context is often called the **small angle approximation**, because it amounts to expanding the $F = ma$ equation for the pendulum around $\theta = 0$. In that neighborhood, the system reduces to a simple harmonic oscillator, and so the pendulum sinusoidally oscillates around its equilibrium position.

