Geometric Astronomy

BY

DYLAN HOLMES

December 2021

Geometric astronomy

Dylan Holmes

December 7, 2021

Abstract

Why do planets have seasons? How, quantitatively, does the amount of planetary tilt affect how much daylight you get? What shape does a sundial's shadow trace throughout the day? Given a time and place on Earth, what direction is the sun? What is the zodiac like on other planets? How do you define reasonable coordinate systems and calendars starting from scratch on the surface of a planet?

I answer questions like these from first principles, in a quantiative, purely geometric way. There's no physics at all—no gravity, no atmosphere—and only very simple geometry—spheres moving in uniform circular orbits. These simplified answers provide intuition to explain what we see in the heavens and what we might see on other worlds.

Contents

Note 1: Welcome! I intend for this document to be a useful reference, where each separate section is self-contained enough that you can read it and learn something new without having to read everything prior or brush up on notation. On the other hand, you can also read it straight through, cumulatively.

Note 2: Who should read this? I think the geometric intuition provided here can be useful to anyone, and I use a lot of pictures to get the point across. If you do want to follow the detailed calculations, it will help to know about vectors and trigonometry—that is, about sine and cosine, how to multiply matrices, and what a cross product is.

Note 3: This is an incomplete α version of this document. Much useful material is here (around fifty pages worth, including diagrams), but a lot of other interesting material has yet to be assembled. If you find it useful and/or you would like to see more, please let me know.

0.1 Quick reference

Rodrigues's rotation formula. Suppose you revolve a vector $\vec{\mathbf{v}}$ around the axis $\vec{\mathbf{k}}$ by an angle of θ, according to the right hand rule. Rodrigues's rotation formula tells you how to compute the end result as a linear combination of the vectors $\vec{\mathbf{v}}$, $\vec{\mathbf{k}} \times \vec{\mathbf{v}}$, and $\vec{\mathbf{k}}$:

$$\vec{\mathbf{v}}_{new} = \vec{\mathbf{v}} \cos\theta + (\vec{\mathbf{k}} \times \vec{\mathbf{v}}) \sin\theta + \vec{\mathbf{k}}(\vec{\mathbf{k}} \cdot \vec{\mathbf{v}})(1 - \cos\theta)$$

Trigonometric identities If you define the trigonometric functions $\langle \cos\theta, \sin\theta \rangle$ to be the coordinates of the point $\langle 1, 0 \rangle$ after it has been rotated counterclockwise by angle θ, and you remember the formula for a rotation matrix $\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$, you can easily derive many trigonometric identities.

For example, because you can rotate by angle $\alpha + \beta$ by rotating first by α, then by β, you find by matrix multiplication that:

$$\cos(\alpha \pm \beta) = \cos\alpha \cos\beta \mp \sin\alpha \sin\beta$$

$$\sin(\alpha \pm \beta) = \sin\alpha \cos\beta \pm \sin\beta \cos\alpha$$

The subsolar point. The subsolar point is the point where the line between the sun and a planet intersects the planet's surface. The noonday sun is directly overhead at that latitude, which I denote as β.

On un-tilted planets, the subsolar point is always at the equator. On tilted planets, the subsolar point moves throughout the year: if u is the year angle (ranging from $u = 0$ to 2π over a complete orbit and starting with $u = 0$ in northern springtime) and q is the obliquity (planetary tilt), then:

$$\sin\beta = \sin u \sin q$$

In particular, the subsolar point ranges between latitudes $\pm q$ over the year, reaching the extreme endpoints during summer and winter.

Local compass directions. If you're at latitude α on a spinning planet with spin axis $\vec{\mathbf{p}}$, you can define local coordinates using cross products. Letting $\vec{\mathbf{z}}$ denote your local up direction (i.e., from the center of the planet to your position):

$(\vec{\mathbf{z}} \times \vec{\mathbf{p}}) \times \vec{\mathbf{z}}$	North, the counterclockwise pole.
$\vec{\mathbf{p}} \times \vec{\mathbf{z}}$	East, the rising sun.
$(\vec{\mathbf{p}} \times \vec{\mathbf{z}}) \times \vec{\mathbf{z}}$	South, the clockwise pole.
$\vec{\mathbf{z}} \times \vec{\mathbf{p}}$	West, the setting sun.

These compass directions can be defined everywhere except at the two poles, where $\vec{\mathbf{z}} \times \vec{\mathbf{p}} = \vec{\mathbf{0}}$.

Spin axis and local noon. At each latitude α, the position of the spin axis $\vec{\mathbf{p}}$ and the noonday sun $\vec{\boldsymbol{\eta}}_\odot$ can be represented in local coordinates:

$$\vec{\mathbf{p}} = \overrightarrow{\mathbf{north}} \cos\alpha + \vec{\mathbf{z}} \sin\alpha$$

$$\vec{\boldsymbol{\eta}}_\odot = \vec{\mathbf{z}} \cos\omega + \overrightarrow{\mathbf{north}} \sin\omega$$

You can see that both vectors must always lie in the north-zenith plane. Here, ω (the zenith angle) completely specifies the position of the noonday sun by describing the angle that the noonday sun makes with the $\vec{\mathbf{z}}$ axis.

The sun's daily circle in local coordinates. On a spinning planet where days are much shorter than years, the position of the sun is given by:

$$\vec{\mathbf{N}}(\delta) = \vec{\boldsymbol{\eta}}_\odot \sin(\delta) + \overrightarrow{\mathbf{east}} \cos(\beta)\cos(\delta) + \vec{\mathbf{p}} \sin(\beta)(1 - \sin\delta)$$

Here, $\vec{\mathbf{N}}$ is the position of the sun over time, $\vec{\boldsymbol{\eta}}_\odot$ is the position of the sun at noon, and $\vec{\mathbf{p}}$ is the (counterclockwise) spin axis of the planet. As for the angles, δ is the day angle, ranging from 0 to 2π during one planetary rotation such that $\delta = \frac{\pi}{2}$ at noon.

Alternatively, expressed in local north-west-up coordinates, the sun's position becomes a function of your viewing latitude α, the latitude of the subsolar point[1] β , and the day angle δ:

$$\vec{\mathbf{N}}(\delta) = \begin{bmatrix} -\sin\alpha\cos\beta\sin\delta \;+\; \cos\alpha\sin\beta \\ -\cos\beta\cos\delta \\ \cos\alpha\cos\beta\sin\delta \;+\; \sin\alpha\sin\beta \end{bmatrix}$$

The equinox effect A tilted planet has two annual *equinoxes*, during which the planet behaves as if it has no tilt. In particular, you can observe the following characteristic effects:

1. Day and night are equal in length everywhere on the planet.
2. The line between the planet and the sun passes directly through the equator.

[1]The latitude β of the subsolar point is equivalent to the time of year u, because $\sin\beta = \sin u \sin q$.

3. The sun rises due east and sets due west everywhere. (Except at the two poles, where the sun revolves around the horizon.)
4. Ignoring physical climate phenomena, the planet has a uniform season everwhere—no winter or summer extremes.
5. Everywhere on the planet, the shadow of a vertical sundial stick traces out a straight line from west to east during the day, instead of the usual curve. (The two poles are an exception.)
6. Everywhere on the planet sees the sun rise and set during the day. The sun doesn't stay above the horizon all day or below the horizon all day anywhere.
7. Unlike every other time of year, there is no place where the sun simply kisses the horizon exactly once without rising or setting. Hence there is no place to see a sundial cast a parabolic shadow.

The shape of sundial shadows A sundial is a vertical stick placed on a level plane. Over the course of the day, the tip of the stick traces out a certain shape. This shape is always a conic section (hyperbola, parabola, ellipse), though it may be degenerate.

The type of shape depends on the number of times the sun crosses the horizon in a single revolution. Most of the time it crosses twice—at sunrise and sunset—and the shadow forms a hyperbola.

At the extreme polar regions, it hovers above (or below) the horizon throughout the entire revolution, crossing the horizon zero times—and the shadow forms an ellipse.

There is an intermediate case, occuring throughout the polar circle, where on a certain special day the sun spends all day above (or all day below) the horizon but gently touches down on the horizon exactly once. In this case, the shadow forms a parabola.

A parabolic shadow occurs at whichever latitude α is complementary with the subsolar latitude β, provided neither is at the equator. Closer to the equator and the shadows are all hyperbolic. Closer to the pole, and the shadows are all elliptical.

0.2 What is this document?

My living room is located in the northern hemisphere, facing south. Late one afternoon, the sun shone through the leaves of a big tree in my yard, splashing a hazy golden pattern on the wall. "That's interesting," I thought. "I bet if I stuck a dark spot onto the window, its shadow would trace out a sundial pattern on the wall. I wonder if I could predict exactly where the shadow would go, just from first principles. I wonder if I could extrapolate, from the shadow pattern on one day, what the pattern would be throughout the rest of the year. And how exactly are latitudes, time of year, and shadow connected, anyways—what can you predict from what?"

These were practical questions, and I had a lot of theoretical gaps to fill in. I decided what I needed first was a very rough, intuitive explanation of how the sun moves in the sky and how it looks from different spots on Earth. I didn't mind neglecting essentially all of the physical details, like atmosphere and even the eccentricity of our orbit, as long as I could derive how the trajectory of the sun in the sky follows from the fact that we're orbiting it on a spinning planet with a tilted axis. I wanted an explanation that was general enough to predict what you'd see on other worlds.

A major obstacle in my project was that most of the explanations I could find were more like *glossaries* for the kind of explanation I was looking for: they would define terms (the ecliptic plane, mean solar time) and angle measurements (zenith angle, obliquity), introduce an inscrutable diagram, and call it a day. But if these sundial shadows were ever going to end up on my wall, I needed a coordinate system—not just angle measurements, but the sines and cosines that put them into space. And how did the astronomers even construct their coordinate systems in the first place? What could you use to anchor yourself in this starry universe?

My work to fill this gap became Section 1.6, on coordinate systems, our year-long orbit, and switching between "this planet orbits the sun" and "the sun orbits this planet"[2]. In Section 2, I considered *daily* motion: the spinning planet and what it means to have a tilted spin axis. In Section 3, I found a formula for the position of the sun in local coordinates, and in Section 4 classified all the sorts of sundial shadows you can get.

0.3 Self-study homework problems

1. Explain how, in a circular orbit, the right hand rule allows you to choose an unambiguous $\vec{\mathbf{z}}$ axis. What about an unambiguous choice of x or y axis?
2. If a translation T sends the axes of coordinate system A onto the axes of coordinate system B, what transformation sends the coordinates of a point represented in system A to its coordinates in

[2]Both views are reasonable from a relativistic standpoint. As far as I can tell, the only thing the geocentrists got wrong was believing that celestial bodies can *only* orbit the Earth.

system B? (Consider the origin, for example.) Same question for a rotation R. What about a translation followed by a rotation?

3. Sketch what a translation followed by a counterclockwise rotation looks like. Choose some numbers and compute a 4×4 matrix representing a translation followed by a counterclockwise rotation. Does the result agree with your picture? Ensure that you've multiplied the transformations in the correct order, and ensure that the rotation really is *counterclockwise* by applying it to an easy test vector like $[1, 0, 0, 1]$.
4. A planet orbits a sun in a circular orbit of radius A. Write a formula for the planet's coordinates, letting the center of the circle be the origin. Which sign convention makes the orbit clockwise versus counterclockwise? Why doesn't it make a physical difference in this simple case?
5. Using a coordinate transform, express the yearly motion of the sun as seen from the planet. (Assume, for now, that the planet does not spin.)
6. Considering the idealized symmetry of the situation—circular orbit, no spin, no tilt—what could you do to determine your position in a circular orbit? Explain how the *zodiac* constellations help.
7. Explain this statement: "A planet's spin axis $\vec{\mathbf{p}}$ basically doesn't swivel, and as a result its tilt *direction* is a matter of convention. Only the tilt angle is fundamental."
8. Using a coordinate transform, express the yearly motion of the sun as seen from a tilted planet with obliquity q. Use a centripetal coordinate system. When rotating the spin axis, choose the rotation so that at the start of the year the sun begins at height $z = 0$ moving upward. What season is this? How would you rotate the spin axis differently so that the year starts in a different season?
9. Draw a picture of the sun's annual motion as seen from the tilted planet's surface. Which positions correspond to which seasons?
10. On Earth, the sun's annual path passes through the twelve zodiac constellations throughout the year. In astrology, the sun's constellation at the time of your birth is your 'sun sign'. Do the stars look the same on Earth and Mars? The zodiac constellations might be different on Earth and on Mars— why?
11. Suppose a planet with obliquity q has a circular orbit of radius A around its sun. At any given moment, the line between the planet and the sun forms the hypotenuse of a right triangle where the legs are parallel and perpendicular to the planet's spin axis.

Draw this in a diagram. What angles and side lengths do you know? (Remember your formula for the sun's annual motion.) What happens to this diagram in the limit as A becomes infinitely large? Assuming the sun is far away, derive the formula for the sun's latitude β in terms of the time of year u and the obliquity q: $\sin\beta = \sin u \sin q$. (The $\sin u$ term may be different depending on which season starts the new year. Explain the possibilities.)

1 Coordinate systems in space

1.1 Orientation using the distant stars

If only the cosmos were overlaid with a convenient grid. Then we could read off numbers to tell exactly where we are, what direction we're headed, and how quickly we're going there. Unfortunately, no such grid has yet been found—somehow, we must find our own way to chart where we are.

If we were in an empty universe, with no stars, no moon, no surrounding debris to disrupt the blackness, there would be no way to get your bearings. It would be similarly difficult if we were in a thick fog, or a storm of churning debris. Wayfinding depends on having some clear constant to measure against—measurements of position, orientation, and speed are all done relative to some identifiable waypoints.

Fortunately, in *our* part of the cosmos, the blackness of space is mostly empty and we have a clear view full of scattered stars. These stars are so staggeringly far away from us that—like a distant mountain seen through a car window—they do not move, no matter how much we move. Over all the distances that our planet and solar system move, over the timescales that we care to measure, the stars stay in place. They appear as if they are bright pinpoints of light poked through an enormous black orb that surrounds us. It's quite a convenient backdrop for getting our bearing in the universe.

Imagine being suspended in this starry void. The stars are scattered in an irregular pattern, so you can easily identify which part of the cosmic sphere you're facing at any given moment: oh, there's a familiar cluster of stars. Because space is densely packed with stars (instead of there being just a handful of stars), we can reasonably assume that you can recognize whichever direction you're facing. So, the distant stars give you a kind of *absolute orientation.* By checking against the fixed stars, you can always tell which direction you're facing and how you're rotating.

If you are alone in this starry void, you still don't have a practical sense of *position* or *movement*—remember, the stars are so far away that even solar-system sized jumps don't change how they look. To measure these properties, we need other—closer—objects.

1.2 Viewing the sun from orbit

The Earth moves around the sun in an approximately circular orbit that takes one year to complete. Suppose *you* were suspended in space, moving in such an orbit around the sun. What would you see, and what could you learn about the motion?

Once again, in this simple geometric setup, if you were suspended in an idealized black void, you would not be able to tell that you were moving: a circular path is perfectly symmetric, and so you would have no cues to tell you where you were on that circle or how fast you were moving. You would simply see a bright star suspended, context-less, in the void.

Fortunately, the starry backdrop provides a cue. As shown in Figure 1, when you move around the sun, you change your point of view—you change which part of the starry backdrop appears behind the sun.

Of course, you can only tell *relative* position and motion in space: you'll see the same thing whether *you're* moving in a circle around the sun, or the sun is moving in that same circle around you. There is, in an important sense, no discernable difference[3].

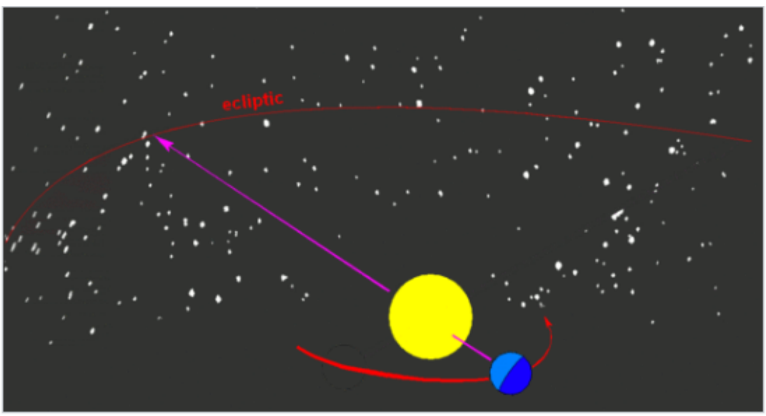

Figure 1: Throughout the year, the line of sight from the Earth to the sun changes, making the sun appear to move against the starry backdrop. When the planet's motion repeats, the sun's transit across the starry backdrop repeats; the sun visits the same part of the starry backdrop at the same time each year. This figure is adapted from Wikipedia's *Ecliptic* article.

The line of sight from the Earth to the sun changes throughout the

[3]But suppose you want to fire a laser through the center of the sun. It's sunrise—the exact moment the center of the sun appears at the horizon. If you intend to fire the laser now, what angle should your laser be pointed at so the laser will pierce the center of the sun? (Neglect atmospheric refraction. The sun is approximately eight minutes away at light speed. The earth completes one full rotation in twenty-four hours.) How does the answer depend on whether the earth goes around the sun, or conversely?

year, going in a complete circle. In practice, you will see the sun appear against a different part of the unchanging starry backdrop (the fixed stars) throughout the year[4]. Because the orbital motion is cyclical, the sun's transit across the starry backdrop is also cyclical; the sun visits the same parts of the backdrop at the same time each year.

If you get your bearings by identifying constellations in the night sky, then you can describe what part of the sky the sun is in by saying what constellation it's passing in front of. This is what the twelve constellations of the zodiac are: twelve named waypoints that the sun visits throughout each year.

1.3 All orbital motion is planar

In idealized two-body problems, the only orbital shapes are conic sections: ellipses and circles, parabolas, and hyperbolas. These are all 2D shapes, confined to a plane.

As a result, in any idealized orbital problem, you can always find a convenient coordinate system in which the planet's orbital motion is exclusively in the x-y plane and has no z component.

Note that, suprisingly, all the planets in our solar system happen to share approximately the same orbital plane. That is, they all move in orbits that occupy roughly the same plane.

1.4 Centripetal coordinate systems

Any circular motion (such as a planetary orbit) gives you an unambiguous way to define a coordinate system.

- The first axis, $\vec{\mathbf{r}}$, is in the radial direction, pointing from the planet to the star.
- The second axis, $\vec{\mathbf{t}}$, is in the tangential direction, pointing in the direction the planet is moving. In circular motion, this is always perpendicular to the radial direction
- The third axis, $\vec{\mathbf{z}}$, is in the axial direction—the central axis around which the planet moves in a circular path. It is defined by the right hand rule as the cross product of $\vec{\mathbf{r}} \times \vec{\mathbf{t}}$ (which is incidentally in the same direction as the planet's orbital angular momentum as defined in physics).

Of course, the position and velocity of the planet are changing all of the time and so these potential axes are changing in time. But if

[4] On a planet, it's hard to get your bearings in daylight, so suppose you mark the sun's position at noon, then return at night to see what stars are in that position.

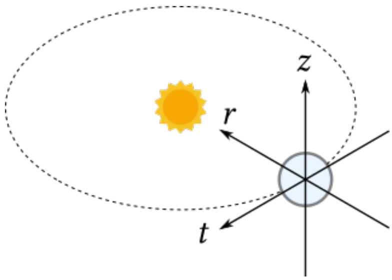

you freeze any moment in time, the axes form an orthogonal coordinate system in that moment that you can use.

Note that in this coordinate system, the planet's orbit remains within a 2D plane, the r–t plane.

1.5 Example: Planetary impact

Suppose a planet is orbiting its star when a stray meteor crashes into it, giving it a kick in some particular direction. What happens next?

Does it wobble in its orbit? Corkscrew away? To see what happens next, imagine the moment just after the planet has been kicked. If you freeze time in that moment, you can identify two vectors: the planet's velocity, and the vector between the planet and the sun. From that moment onward, the planet's motion will be in the plane defined by those two vectors.

The orbital motion is once again planar, though you might have to find a new coordinate system where that plane is level.

1.6 Transformations in homogeneous coordinates

You can represent translations and rotations as matrices if you use *homogeneous coordinates*. In homogeneous coordinates, you represent a 3D point $\langle x, y, z \rangle$ as $\langle x, y, z, 1 \rangle$. That is, you represent three dimensional points as four dimensional points with a constant 1 in the fourth slot.

In homogeneous coordinates, a counterclockwise rotation by angle θ around the z axis is achieved using the matrix:

$$\mathbf{R}_z(\theta) \equiv \begin{bmatrix} \cos\theta & +\sin\theta & 0 & 0 \\ -\sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

And a 3D translation in the amount $\langle +\Delta x, +\Delta y, \Delta z\rangle$ is achieved using the matrix:

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & +\Delta x \\ 0 & 1 & 0 & +\Delta y \\ 0 & 0 & 1 & +\Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

1.7 Distant stars and points at infinity

By convention, you can multiply the coordinates of a homogeneous point by any amount $\alpha \neq 0$ without changing its meaning. For example, the homogeneous point $\langle 1, 2, 3, 1\rangle$ and the homogeneous point $\langle 5, 10, 15, 5\rangle$ are different ways of expressing the same point. You can always normalize a homogeneous point by dividing all the components by its fourth component; this will put it back into standard form with its fourth component equal to 1.

Something special happens with points that have a zero in the fourth component $\langle x, y, z, 0\rangle$. They behave like they are infinitely far away.

A distant object, such as a mountain, won't seem to move very much as you change position. A very distant object, such as the stars in the sky (or, to a lesser extent, the moon) won't seem to move at all even if you travel by car.

Points at infinity capture this phenomenon in an idealized way. Translations have no effect at all (try applying a translation matrix to a point at infinity, for example), the way movement doesn't change the appearance of a distant object. Rotations still do have an effect—points at infinity have a particular direction; you can look toward them or away from them for example.

2 Planets that spin

2.1 Axial tilt (obliquity)

Some planets spin. I use the term *spin axis* to refer to the imaginary axle around which it spins. By analogy, the *orbital axis* is the imaginary axle around which the planet orbits; for an ideal circular orbit, this is the line that perpendicularly pierces the center of the circle.

Some planets, like our Earth, have axial tilt. *Axial tilt* means that the spin axis and the orbital axis are not parallel. (If a planet only orbited a star and did not spin, there would be no such thing as axial tilt.)

The two major questions I wanted to know about axial tilt are these:

- In what coordinate system is the axis tilted? What are the angles defined relative to?
- How does the tilt change throughout the year?

You can define tilt in terms of the cross product A planet has an axial tilt whenever its spin axis $\vec{\mathbf{p}}$ and its orbital axis $\vec{\mathbf{z}}$ aren't parallel but instead point in different directions. In that case, the cross product $\vec{\mathbf{p}} \times \vec{\mathbf{z}}$ is nonzero. You can describe $\vec{\mathbf{p}}$ as a *rotated version* of $\vec{\mathbf{z}}$. You have taken $\vec{\mathbf{z}}$ and rotated it by some angle q around some axis.

In fact, the axis has a convenient mathematical definition—it's just the cross product of those vectors $\vec{\mathbf{p}} \times \vec{\mathbf{z}}$.

That angle q is called the *obliquity* of the spin axis. As I'll describe below, you can calculate it using the usual dot product rule:

$$\vec{\mathbf{p}} \cdot \vec{\mathbf{z}} = |\vec{\mathbf{p}}|\ |\vec{\mathbf{z}}| \cos q$$

$$q \equiv \cos^{-1}\left[\frac{\vec{\mathbf{p}} \cdot \vec{\mathbf{z}}}{|\vec{\mathbf{p}}|\ |\vec{\mathbf{z}}|}\right]$$

This is just a pure geometric fact:

> **Cross product yields the unique rotational axis**. If $\vec{a}$ and $\vec{b}$ are two vectors with the same length but different directions, you can describe $\vec{b}$ uniquely as a rotated version of $\vec{a}$.
>
> A rotation is defined by the axis of rotation and the amount (angle) of rotation. Here, the axis is given by the cross product $\vec{a} \times \vec{b}$ and the angle is given by the dot product $\vec{a} \cdot \vec{b} = |\vec{a}|\ |\vec{b}|\ \cos\theta$.

The spin axis basically doesn't swivel Surprisingly, Earth's spin axis (basically) points resolutely in the same absolute direction at all times. That is, in an idealized sun-centered coordinate system where the Earth moves in a perfect circle around the sun, the spin axis never changes direction (Figure 2).

Because the Earth's spin axis never changes direction, both the obliquity angle q and the rotational axis $\vec{\mathbf{p}} \times \vec{\mathbf{z}}$ are constant throughout all time. They're a fundamental property of the orbit itself.

2.2 Seasonal coordinate system

If a planet moves in a circular orbit around its star, the orbital axis $\vec{\mathbf{z}}$ is always perpendicular to the radial vector $\vec{\mathbf{r}}$ between the planet and the star. In fact, over the course of an entire orbit, the $\vec{\mathbf{r}}$ vector sweeps through *every* vector perpendicular to $\vec{\mathbf{z}}$.

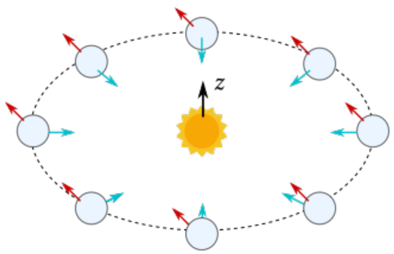

Figure 2: Over the course of a full orbit, the radial axis (turquoise) sweeps through all possible vectors perpendicular to the orbital axis (black).

If the planet is tilted, then its spin axis $\vec{\mathbf{p}}$ and orbital axis $\vec{\mathbf{z}}$ are not parallel. Their cross product $\vec{\mathbf{z}} \times \vec{\mathbf{p}}$ is therefore a nonzero vector that is perpendicular to $\vec{\mathbf{z}}$.

You can therefore find a place in the orbit where $\vec{\mathbf{r}} = \vec{\mathbf{z}} \times \vec{\mathbf{p}}$. This is called an *equinox*. (The other equinox occurs when $\vec{\mathbf{r}} = -\vec{\mathbf{z}} \times \vec{\mathbf{p}}$.)

As we will see, the equinox is a special moment in the year for a tilted planet. It is useful to use this equinox-based centripetal coordinate system:

$$\vec{\mathbf{z}} = \vec{\mathbf{z}}$$
$$\vec{\mathbf{r}} = \vec{\mathbf{z}} \times \vec{\mathbf{p}}$$
$$\vec{\mathbf{t}} = (\vec{\mathbf{z}} \times \vec{\mathbf{p}}) \times \vec{\mathbf{z}}$$

and choose the year angle so that $u = 0$ occurs at the equinox[5].

2.3 Amount of sunlight as a function of obliquity

Suppose you have a planet with obliquity (axial tilt angle) q, and suppose the year is much longer than the length of the day[6].

[5] Many cultures do choose a calendar system that starts at the equinox, celebrating new year's on the first day of spring.

[6] This is an idealized case, where the ratio of day length to year length goes to zero.

At the poles, you'll experience half a year of uninterrupted sunlight[7]. Other places will receive continuous sunlight for less than half the year.

Linear relation between latitude and continuous sunlight. Quantitatively, the poles (which are at $\pm 90°$ latitude) receive half a year of continuous sunlight. The *polar circles*, which are defined as $\pm(90 - q)°$ latitude and therefore depend on the amount of obliquity, receive a single day of continuous sunlight. (As a fraction of a long, long year, this is just zero.) The amount of continuous sunlight varies linearly between these limits, and is identically zero between the two polar circles (between $(90 - q)°$ and $-(90 - q)°$ latitudes).

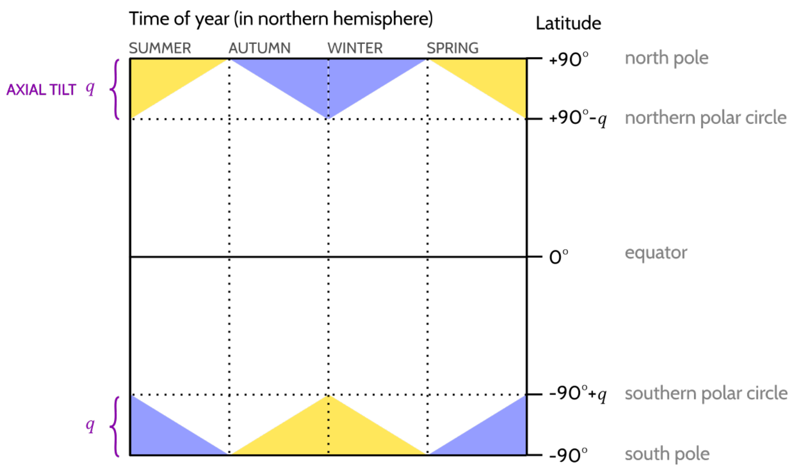

Figure 3: On a tilted planet (with obliquity q), the polar regions receive continuous sunlight for an entire day or more. In the idealized case where years are much longer than days, the amount of continuous sunlight ranges linearly from 0 to half a year as you travel from the polar circles (within q latitude of either pole) to the poles themselves. Here, yellow regions receive at least a day of continuous sunlight; blue regions receive at least a day of continuous darkness; and white regions receive less than a day of either, although of course the amount of sunlight varies seasonally everywhere.

[7] I mean "sunlight" geometrically: a place gets sunlight whenever the sun is physically above the horizon there. This definition neglects physical effects such as atmospheric refraction, which make the sky look bright even when the sun is below the horizon.

A picture (Figure 3) is clearer than a formula, but if you want a formula relating the fraction of the year f that a particular latitude ℓ has continuous sunlight, it's:

$$f(\ell) = \max\left[0, \frac{1}{2} - \frac{90° - |\ell|}{2q}\right]$$

2.4 Spin direction depends on perspective

Every spinning planet has two poles. It spins clockwise around one pole and counterclockwise from another. Therefore, there's no real answer to the question "Does the planet spin counterclockwise or clockwise" unless you have a viewing direction in mind.

Earth spins counterclockwise around its north pole and clockwise around its south pole.

2.5 Hemisphere bias in clockwise, north, south

Surprisingly, the etymologies of the terms clockwise, north, and south are based on the behavior of the sun in Earth's northern hemisphere:

- In the northern hemisphere, the shadow of a sundial traces out a clockwise path throughout the day. Mechanical clocks were designed to imitate this pattern, which is why that direction is called "clockwise".
- The term South is the derived from the term *sunward*. In the northern hemisphere, the sun travels across the *southern* half of the sky.
- Similarly, North is the *nether* side of the sky, where the sun doesn't shine[8]

As a result, if our linguistic heritage were developed in the southern hemisphere, clockwise, north, and south would have the opposite meanings—the shadows of sundials move *counterclockwise* in the southern hemisphere, the sunward side is toward the *north*, and the nether side is toward the *south*.

It makes sense that people living in the northern hemisphere developed these terms based on their particular experiences. However, when everyone on Earth uses this terminology, then we are all taking an implicitly northern perspective.

[8] Incidentally, this is why moss, which prefers low light, has a slight bias for growing on the north sides of trees in the northern hemisphere.

Cosmically, it would be nice to have planetary-scale terminology that matches the local conditions everywhere and doesn't privilege a specific hemisphere. Especially for other planets in other star systems: terms like north, south, clockwise, and counterclockwise—not to mention summer and winter—are hemisphere-dependent.

While East and West have a definite geometric meaning on any spinning planet that orbits a star—the planet sees the star rise in the east and set in the west—North and South and summer and winter have no universal meaning.

2.6 Red and black terminology

Here is my proposal for planet-wide terminology:

- Any spinning planet has two poles: the *red pole*, around which it spins counterclockwise per the right hand rule, and the *black pole* around which it spins clockwise.
- The poles divide the planet into a *red hemisphere* and a *black hemisphere.*
- Similarly, there are *red and black polar circles*, and *red and black tropics.*
- The two equinoxes are the *redward equinox* and the *blackward equinox*, which occur as the sun crosses the equator into one of the two hemispheres.
- And we can refer to right handed spin as *redwise* and left handed spin as *blackwise.*

Incidentally, this language matches the etymological origins of the Red and Black Seas on Earth.

While you could argue that north, south, clockwise, and counterclockwise are already harmless, having been stripped of their hemisphere-specific origins, I believe as a matter of taste and universality that it is nice to have fresh neutral terminology.

I use a mixture of traditional and new terminology in this document.

3 The position of the sun as seen from the planet's surface

3.1 Local coordinates: the direction of sunrise and sunset

A spinning planet has a spin axis vector $\vec{\mathbf{p}}$, defined as the pole around which it spins *counterclockwise.* (On Earth, this is the north pole.)

At each point on a planet, there's a zenith vector $\vec{\mathbf{z}}$ which points straight overhead. The zenith vector $\vec{\mathbf{z}}$ depends where you are on the surface; it points radially outward from the planet's center to your position. The spin axis $\vec{\mathbf{p}}$ is location-independent.

You can compute compass directions from these two vectors: $\vec{\mathbf{z}} \times \vec{\mathbf{p}}$ points West toward the setting sun, while $\vec{\mathbf{p}} \times \vec{\mathbf{z}}$ points East toward the rising sun[9]. (Figure 4)

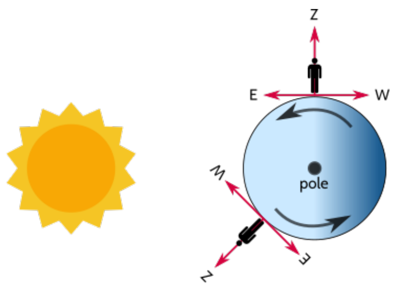

Figure 4: The direction of sunrise (East) can be defined in local coordinates: Take a top-down view of a spinning planet. Every spot on the planet has a locally-defined zenith vector $\vec{\mathbf{z}}$ which points straight up. • As you can see, if the planet spins counterclockwise around its spin axis $\vec{\mathbf{p}}$, it will make the sun appear in the direction $\vec{\mathbf{p}} \times \vec{\mathbf{z}}$.

3.2 North, East, South, West in local coordinates

If you're on a spinning planet whose spin axis is $\vec{\mathbf{p}}$ and your local "up" direction is $\vec{\mathbf{z}}$, you can define the four compass directions as follows:

$(\vec{\mathbf{z}} \times \vec{\mathbf{p}}) \times \vec{\mathbf{z}}$	North, the counterclockwise pole.
$\vec{\mathbf{p}} \times \vec{\mathbf{z}}$	East, the rising sun.
$(\vec{\mathbf{p}} \times \vec{\mathbf{z}}) \times \vec{\mathbf{z}}$	South, the clockwise pole.
$\vec{\mathbf{z}} \times \vec{\mathbf{p}}$	West, the setting sun.

3.3 The height of the sun at noon depends on latitude

I used to believe that the sun passed directly overhead at noon. It turns out that, instead, the noonday sun will be more in the nothern or southern half of the sky, depending on your latitude α.

[9]This definition does not depend on your planet's definition of North. It just requires that $\vec{\mathbf{p}}$ is whichever axis the planet rotates counterclockwise around.

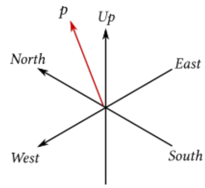

Figure 5: Local compass directions. Using the right-hand rule, you can compute the four compass directions as cross products of $\vec{\mathbf{p}}$ (the spin axis) and $\vec{\mathbf{z}}$ (the local upward direction). This process works everywhere except the poles, where the two vectors are parallel. $\vec{\mathbf{z}} \times \vec{\mathbf{p}} = \vec{\mathbf{p}} \times \vec{\mathbf{z}} = \vec{\mathbf{0}}$.

> **Changes in latitude produce equal changes in zenith angle**. If you change your latitude by an amount α toward the north, you will see the noonday sun move by an amount α toward the south.

The situation is easiest to understand on a planet with no axial tilt (obliquity). On a planet with no axial tilt, the noonday sun is directly overhead at the equator. If you move northward by an amount α, the sun will move southward in the sky by that same amount. So, at the poles where the latitude is $\pm 90°$, the sun simply sits at the horizon year-round. In general, on a planet with no tilt, the sun's zenith angle from "straight overhead" is exactly the opposite of your latitude.

If a planet has axial tilt, the sun will be directly overhead at various latitudes throughout the year, not just at the equator. Still, the equal-changes rule still applies: if the sun is directly overhead at latitude β at some point during the year, and you are at latitude α, then the sun's declination is at angle $\beta - \alpha$.

> **Noonday zenith angle equals the difference in latitude between viewer and sun**. If you are at latitude α and the sun is directly above latitude β, then the noonday sun will be at zenith angle $\beta - \alpha$ on the north side of straight overhead.
>
> (On a planet with no tilt, $\beta = 0$.)

3.4 On a tilted planet, the latitude of the sun oscillates

If you draw a line from the star to the center of a planet, it will pierce the surface of the planet at some point. This is called the *subsolar point*.

On a planet with no axial tilt, the subsolar point is always on the equator. On a planet with obliquity q, the subsolar point oscillates back and forth between latitude $+q$ and latitude $-q$ throughout the year. The latitudes $\pm q$ are called the *northern and southern tropics*; they're the furthest points from the equator where the sun can ever be directly overhead.

Quantitatively, if u is the year angle (a measure that goes from 0 to 2π over the course of one full orbit) and β is the latitude of the subsolar point, and q is the obliquity, then (as we will see in the next section):

$$\sin(\beta) = \sin(u)\sin(q)$$

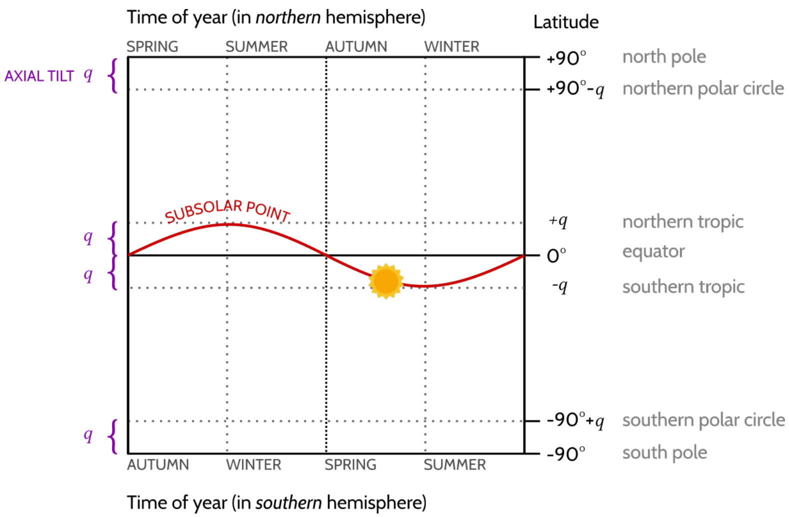

Figure 6: On a tilted planet with obliquity q, the sun's latitude—that is, the point where a line from the sun to the planet crosses the planet's surface—wanders annually between latitudes $\pm q$. The motion is periodic but not quite sinusoidal.

This is sort of a strange compound periodic motion. Note that if $q = 0$ (no axial tilt), we get $\sin(\beta) \equiv 0$ and the sun is always directly above the equator.

Otherwise, if we have some axial tilt $q \neq 0$, then the four quarters of the year have special meaning:

When $u = 0$ or $u = \pi$, we have that $\beta = 0$ and the subsolar point crosses the equator.

When $u = \pi/2$ or $u = 3\pi/2$, we have that $\beta = \pm q$, its extremal value.

These four important points correspond to the four seasons.

3.5 Exact and approximate formulas for the sun's latitude

The subsolar point, as we've defined it, is the point where the line between the planet and the sun crosses the planet's surface. Let's compute the latitude of the subsolar point as it changes throughout the year.

On a planet with axial tilt (obliquity) q, the sun appears to orbit the planet in a yearly circle; the sun's orbital plane is tilted by angle q with respect to the planet's spin axis.

Without axial tilt, the position of the sun over time has the coordinates for circular motion:

$$\langle A\cos u, A\sin u, 0\rangle$$

where A is the distance between the sun and the planet, u is the year angle (ranging from 0 to 2π over the course of the year), and the axes are chosen so that the z direction coincides with the planet's spin axis.

To account for the axial tilt, we now have to rotate the sun's trajectory by angle q. This breaks the symmetry of the sun's orbit—previously, all points on the year angle u were equivalent. After rotation, the sun's position will rise and fall past the equator. This symmetry-breaking produces the four seasons.

Because the situation is symmetric, we can rotate around any axis in the x-y plane. Our choice of axis just determines what seasonal time of year $u = 0$ represents. Let's rotate around the x axis so that $u = 0$ will correspond to spring in the $+z$ hemisphere.

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos q & \sin q \\ 0 & -\sin q & \cos q \end{bmatrix} \begin{bmatrix} A\cos u \\ A\sin u \\ 0 \end{bmatrix} = \begin{bmatrix} A\cos u \\ A\sin u\cos q \\ -A\sin u\sin q \end{bmatrix}$$

The position of the sun in its rotated position forms a right triangle whose height is $-A\sin u\sin q$, whose hypotenuse has length A, and whose internal angle β is the latitude of the subsolar point. By trigonometry, we obtain the relationship that defines the subsolar latitude:

$$\sin\beta = \sin u\sin q$$

Now the size of the latitude $|\beta|$ is at most the size of the obliquity q. If the obliquity is small, then the small-angle approximation says that $\sin(x) \approx x$, so we can approximate:

$$\beta \approx q\cdot\sin u \qquad \text{when } q \text{ is small}$$

Under this approximation, we see that the subsolar point basically oscillates sinusoidally between latitudes $\pm q$ (the tropics) over the course of the year.

3.6 Decomposing vectors into compass directions

Let $\vec{\mathbf{p}}$ be the spin axis of your planet. Pick a spot on the surface of a planet and let $\vec{\mathbf{z}}$ be the local zenith direction ("straight up").

If you rotate the zenith vector $\vec{\mathbf{z}}$ by any angle θ in the north-south direction, the result can be expressed as a linear combination of the local zenith direction $\vec{\mathbf{z}}$ and the local redward direction $\overrightarrow{\mathbf{north}}$:

$$\vec{\mathbf{z}}_{new} = \vec{\mathbf{z}}\cos\theta + \overrightarrow{\mathbf{north}}\,\sin\theta$$

(Note that you can define $\overrightarrow{\mathbf{north}}$ as a cross-product combination of the zenith and spin axis vectors: $\overrightarrow{\mathbf{north}} \equiv (\vec{\mathbf{z}} \times \vec{\mathbf{p}}) \times \vec{\mathbf{z}}$.)

In particular, the spin axis $\vec{\mathbf{p}}$ can be expressed in terms of these local coordinates: if α is your current latitude, then by elementary geometry[10], $\vec{\mathbf{p}}$ is just a version of $\vec{\mathbf{z}}$ that has been rotated northward by the complementary angle $\pi - \alpha$.

$$\vec{\mathbf{p}} = \overrightarrow{\mathbf{north}}\,\cos\alpha + \vec{\mathbf{z}}\sin\alpha$$

Similarly, no matter where you are on the planet, the noonday sun lies somewhere in the north-zenith plane. The angle that the sun deviates from directly overhead is called its *zenith angle* ω. You can express the position of the noonday sun in terms of these local coordinates as:

$$\vec{\boldsymbol{\eta}}_{\odot} = \vec{\mathbf{z}}\cos\omega + \overrightarrow{\mathbf{north}}\,\sin\omega$$

> **Azimuthal angle of the sun**. The position of the sun at noon can be specified with a single number, the azimuthal angle ω. This is because the noonday sun must always lie in the north-zenith plane; the only question is at what angle.
>
> Given the azimuthal angle, the position of the noonday sun is $\vec{\boldsymbol{\eta}}_{\odot} = \vec{\mathbf{z}}\cos\omega + \overrightarrow{\mathbf{north}}\,\sin\omega$.
>
> Note that ω is equal to the difference in latitudes between the viewer and the sun: $\omega = \beta - \alpha$.

[10] In the limit as the distance between the planet and star grows much greater than the radius of the planet

3.7 Noon height determines the time of year

The position $\vec{\boldsymbol{\eta}}_\odot$ of the sun at noon is determined by the latitude and time of year. This is the principle behind a *calendar sundial*, which tells you the time of year given that you know the height of the sun at noon.

The postion of the sun at noon is always in the north-up plane; the exact orientation is specified by the azimuthal angle ω. By geometric reasoning, this angle is equal to the difference in latitude between the viewer and the sun: $\omega = \beta - \alpha$.

The subsolar point's latitude β changes with the year angle u and the obliquity of the planet q via $\sin\beta = \sin u \sin q$.

Hence the time of year (u) can be determined if you know your latitude (α), the azimuthal angle of the sun at noon (ω), and the obliquity of the planet q:

$$\sin u = \frac{\sin(\omega + \alpha)}{\sin q}$$

3.8 Rodrigues's rotation formula

Suppose you revolve a vector $\vec{\mathbf{v}}$ around the axis $\vec{\mathbf{k}}$ by an angle of θ, according to the right hand rule. Rodrigues's rotation formula tells you how to compute the end result as a linear combination of the vectors $\vec{\mathbf{v}}$, $\vec{\mathbf{k}} \times \vec{\mathbf{v}}$, and $\vec{\mathbf{k}}$:

$$\vec{\mathbf{v}}_{new} = \vec{\mathbf{v}}\cos\theta + (\vec{\mathbf{k}} \times \vec{\mathbf{v}})\sin\theta + \vec{\mathbf{k}}(\vec{\mathbf{k}} \cdot \vec{\mathbf{v}})(1 - \cos\theta)$$

3.9 The sun's daily motion as observed from the surface

Throughout the course of the day, the sun seems to revolve around the polar axis $\vec{\mathbf{p}}$. You can use Rodrigues's formula to compute the coordinates of the sun as seen from a particular position on the planet[11].

In local coordinates, the pole and noonday sun are at:

$$\begin{aligned} \vec{\mathbf{p}} &= \overrightarrow{\mathbf{north}}\cos\alpha + \vec{\mathbf{z}}\sin\alpha \\ \vec{\boldsymbol{\eta}}_\odot &= \vec{\mathbf{z}}\cos(\beta - \alpha) + \overrightarrow{\mathbf{north}}\sin(\beta - \alpha) \end{aligned}$$

where α is your latitude and β is the latitude where the sun is directly overhead.

[11]Note that in the analysis that follows, I'm assuming that—as on Earth—years are much longer than days. This means that we can consider β (the sun's latitude, i.e. the latitude of the subsolar point) to be constant over the course of one day. The actual motion is a little more complicated.

We can use Rodrigues's formula to compute how the sun revolves around the pole. If we let θ be the sun's angle away from noon (so $\theta = 0$ corresponds to noon), the sun's position through the day is:

$$\vec{\mathbf{N}}(\theta) = \vec{\boldsymbol{\eta}}_\odot \cos\theta + (\vec{\mathbf{p}} \times \vec{\boldsymbol{\eta}}_\odot)\sin\theta + \vec{\mathbf{p}}(\vec{\mathbf{p}} \cdot \vec{\boldsymbol{\eta}}_\odot)(1 - \cos\theta)$$

Now, it turns out that the quantities $\vec{\mathbf{p}} \times \vec{\boldsymbol{\eta}}_\odot$ and $\vec{\mathbf{p}} \cdot \vec{\boldsymbol{\eta}}_\odot$ have a simple form[12]:

$$\vec{\mathbf{p}} \times \vec{\boldsymbol{\eta}}_\odot = \overrightarrow{\mathbf{east}} \cos\beta$$
$$\vec{\mathbf{p}} \cdot \vec{\boldsymbol{\eta}}_\odot = \sin\beta$$

Note that these expressions don't depend on α, the viewer's latitude: the pole and sun have a fixed angle between them, so while changing your position on the planet changes their local coordinates, it doesn't change their relative angle measurements.

So the position of the sun throughout the day is:

$$\vec{\mathbf{N}}(\theta) = \vec{\boldsymbol{\eta}}_\odot \cos(\theta) + \overrightarrow{\mathbf{east}} \cos(\beta)\sin(\theta) + \vec{\mathbf{p}}\sin(\beta)(1 - \cos\theta)$$

Now, θ is an unusual way for measuring the sun's daily progress. It has the wrong handedness (θ gets more negative throughout the day) and a strange origin ($\theta = 0$ at midday). Instead, we can use the angle measure $\delta \equiv \frac{\pi}{2} - \theta$.

$$\vec{\mathbf{N}}(\delta) = \vec{\boldsymbol{\eta}}_\odot \sin(\delta) + \overrightarrow{\mathbf{east}} \cos(\beta)\cos(\delta) + \vec{\mathbf{p}}\sin(\beta)(1 - \sin\delta)$$

3.10 The sun's daily motion in local (north, west, up) coordinates

We've obtained a formula for the position of the sun $\vec{\boldsymbol{\eta}}_\odot(\theta)$ as seen from the planet's surface. That formula was a function of the spin axis orientation ($\vec{\mathbf{p}}$), subsolar latitude (β), and time of day (θ) (where $\theta = 0$ is noon).

By picking the right-handed coordinate system

$$\overrightarrow{\mathbf{north}} = [1, 0, 0]$$
$$\overrightarrow{\mathbf{west}} = [0, 1, 0]$$
$$\overrightarrow{\mathbf{up}} = [0, 0, 1],$$

[12] To obtain this form, I used the fact that the cross product distributes across vector sums, the fact that cross product of a vector with itself is zero, and the sum identities for sine and cosine.

we can express that formula for the sun's position in *local* (north, west, up) coordinates, as seen from a particular viewing latitude α:

$$\vec{\mathbf{N}}(\delta) = \begin{bmatrix} -\sin\alpha\cos\beta\sin\delta \ + \ \cos\alpha\sin\beta \\ -\cos\beta\cos\delta \\ \cos\alpha\cos\beta\sin\delta \ + \ \sin\alpha\sin\beta \end{bmatrix}$$

Figure 7: The position of the sun $\vec{\mathbf{N}}$ during the course of the day, as a function of the viewer's latitude α, the sun's latitude β, and the time of day δ (an angle ranging from 0 to 2π over the course of one revolution). The three coordinates represent local 'north', 'west', and 'up' relative to the viewer.

Recall that the planet spins around its axis, which makes the sun seem to orbit the planet in the opposite direction. Accordingly, the sign on δ—and hence the sign on the westward term—has been chosen so that the planet's spin around its axis and the sun's daily trip around the planet go in opposite directions. The sun correctly rises in the east and sets in the west.

3.11 Time of sunrise and sunset

The height of the sun above the horizon, as we have found, is:

$$N_z(\delta) = \cos\alpha\cos\beta\sin\delta \ + \ \sin\alpha\sin\beta$$

where α is your latitude, β is the subsolar latitude, and δ is the day angle (from 0 to 2π during one day, reaching $\pi/2$ at noon). Sunrise and sunset occur whenever this height reaches zero; the corresponding value of δ tells you the time.

If the coefficient $\cos\alpha\cos\beta$ is equal to zero (i.e., if either you or the sun are at the poles), then the sun spends all day at the horizon in perpetual twilight. Otherwise, we can rearrange terms to find:

$$\sin\delta = \frac{\sin\alpha\sin\beta}{\cos\alpha\cos\beta}$$

which has two solutions $\delta_1 + \delta_2 = \pi$, since the $\sin$ function achieves the same height at supplementary angles.

3.12 The center, radius, and orientation of the sun's daily motion

The sun's journey throughout the day is produced by revolving it around the spin axis $\vec{\mathbf{p}}$. For this reason, the sun's daily path is always some kind of circle. In this section, we'll compute the center, radius, and orientation of that circle in order to get intuition.

First, we'll use our earlier formula for the position of the sun $\vec{\mathbf{N}}(\delta)$. To obtain three points on the circle, we'll plug in three different convenient times of day: dawn ($\delta_1 = 0$), dusk ($\delta_2 = \pi$), and noon ($\delta_3 = \pi/2$). At those times, the sun will be at, respectively:

$$\vec{\mathbf{N}}_1 = \begin{bmatrix} \cos\alpha\sin\beta \\ -\cos\beta \\ \sin\alpha\sin\beta \end{bmatrix} \quad \vec{\mathbf{N}}_2 = \begin{bmatrix} \cos\alpha\sin\beta \\ +\cos\beta \\ \sin\alpha\sin\beta \end{bmatrix} \quad \vec{\mathbf{N}}_3 = \begin{bmatrix} \cos\alpha\sin\beta - \sin\alpha\cos\beta \\ 0 \\ \cos\alpha\cos\beta + \sin\alpha\sin\beta \end{bmatrix}$$

The sun is at opposite ends of the circle at dawn ($\vec{\mathbf{N}}_1$) and dusk ($\vec{\mathbf{N}}_2$); therefore, the average of those two positions will be the center of the circle:

$$\mathsf{center}(\alpha, \beta) \equiv \sin\beta \cdot \begin{bmatrix} \cos\alpha \\ 0 \\ \sin\alpha \end{bmatrix}$$

Note that, therefore, the subsolar latitude β controls how far the center is displaced from the origin, while the viewer latitude α controls in what direction. The center can be displaced anywhere in the north-up plane—no east-west displacement allowed.

Also note that when α and β have the same sign—i.e., when the viewer and the subsolar point are in the same hemisphere—the center is displaced *above* the horizon, creating the long days of summer. When they have opposite signs, the center is displaced below the horizon, creating the long nights of winter.

> **Seasonally varying daylight**. Quantitatively, the center of the sun's daily circle is translated above the horizon by the amount $\sin\alpha\sin\beta$. So, when the viewer (α) and subsolar point (β) are in the same hemisphere, you get the long days of summer. When they differ, the displacement is negative—long winter nights.

As for the radius of the circle, we can simply take the distance between dawn ($\vec{\mathbf{N}}_1$) and dusk ($\vec{\mathbf{N}}_2$) and divide by two. We find:

$$\mathsf{radius}(\alpha, \beta) \equiv \cos\beta$$

Hence the radius of the circle is controlled by the subsolar latitude. This makes sense: the daily motion of the sun is a revolution around the polar axis, so the closer the sun gets to either pole, the smaller the radius will be.

Finally, we can compute the orientation of the circle. One convenient way to express orientation is to find a normal vector perpendicular to the circle. To do so, we'll take the cross product of two radii: $\vec{\mathbf{N_1}} - \mathsf{center}$ and $\vec{\mathbf{N_3}} - \mathsf{center}$. We find:

$$\mathsf{normal}(\alpha, \beta) \equiv \cos^2 \beta \cdot \begin{bmatrix} \cos \alpha \\ 0 \\ \sin \alpha \end{bmatrix}$$

Surprisingly, the normal vector is proportional to the center vector. The circle is always aligned with the direction in which it is displaced. Simply put, *the sun's daily circle always points toward the origin.* The circle is viewer-centered.

To recap:

Measuring the sun's circle. The sun's daily path is a circle of revolution with the following properties:

$$\begin{aligned} \mathsf{center}(\alpha, \beta) &\equiv \sin \beta \cdot \begin{bmatrix} \cos \alpha \\ 0 \\ \sin \alpha \end{bmatrix} \\ \mathsf{radius}(\alpha, \beta) &\equiv \cos \beta \\ \mathsf{normal}(\alpha, \beta) &\equiv \cos^2 \beta \cdot \begin{bmatrix} \cos \alpha \\ 0 \\ \sin \alpha \end{bmatrix} \end{aligned}$$

Here, α is the viewer's latitude and β is the latitude of the subsolar point. The coordinate axes are viewer-local, representing the north, west, and upward directions respectively.

3.13 Diagrams of the sun's daily circle

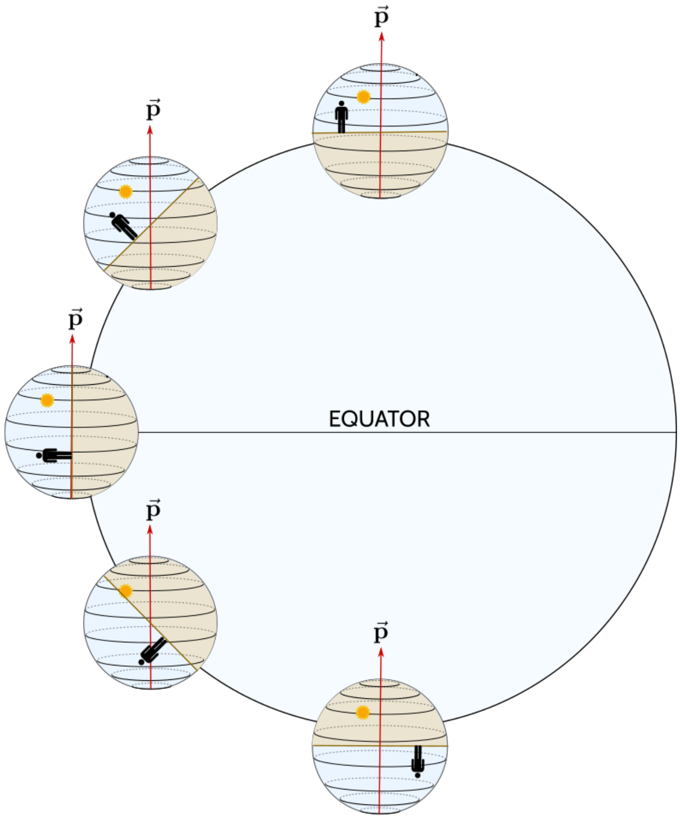

Figure 8: The sun moves in a particular circle perpendicular to the spin axis $\vec{\mathbf{p}}$; the position of the circle is the same for all viewers, and is determined by the sun's latitude. However, from different viewing latitudes, the *angle* between the spin axis and the ground plane changes, causing a different amount of daylight per revolution (It is dark whenever the sun is below the local horizon). • You can tell that in these pictures, it is summer in the upper hemisphere—the sun spends more than half the day above the horizon there, and the pole spends the entire revolution above the horizon. At the equator, light and dark are equally balanced. In the lower hemisphere, there is more dark than light, and the pole is in darkness all day long. • Note that all viewers shown here are at the same longitude and hence the same time of day.

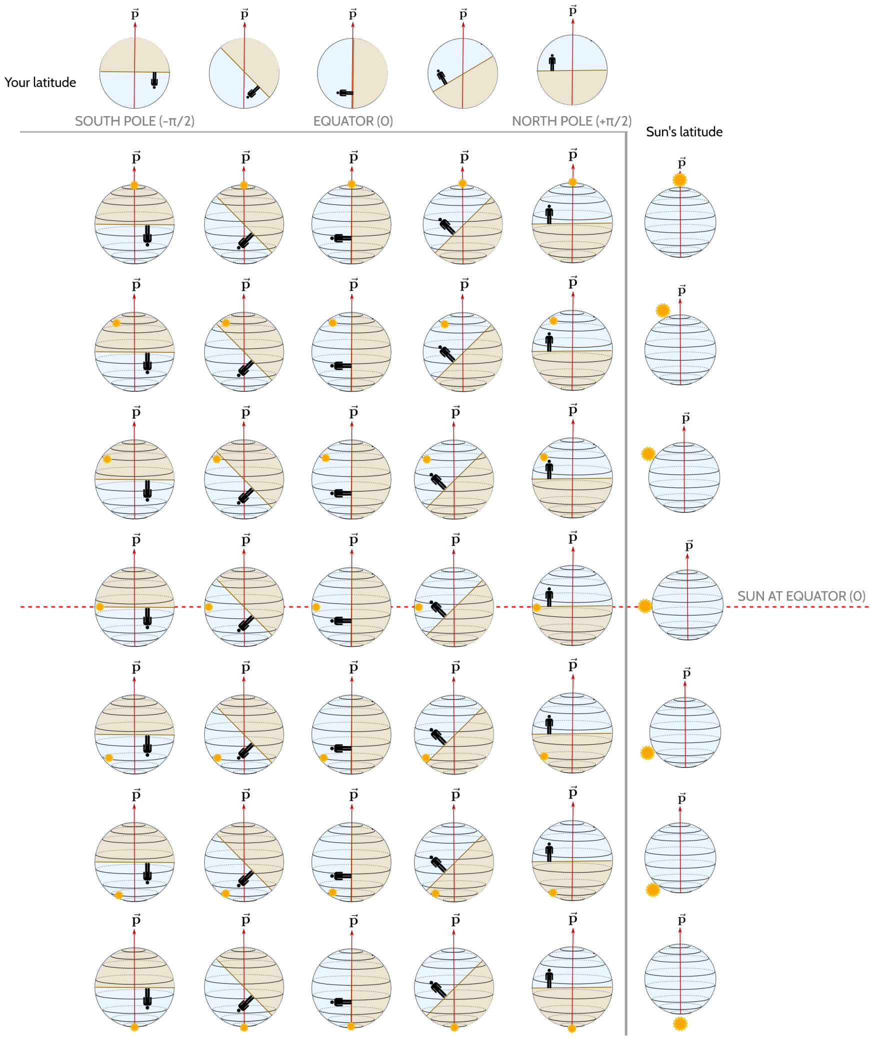

Figure 9: Here, you see how the daily path of the sun changes with the viewer's latitude (α, horizontal axis) and subsolar point (β, vertical axis). On a planet with obliquity q, the sun oscillates between latitudes $\pm q$ throughout the year and never goes beyond them. Note that rows correspond to the same moment in time on different parts of the globe. It is dark when the sun is below the horizon.

3.14 The term that represents seasonally longer days

We've seen that the height of the sun above the horizon is:

$$N_z(\delta) = \cos\alpha\cos\beta\sin\delta \ + \ \sin\alpha\sin\beta$$

[matching signs for β and α mean that you're in the same hemisphere as the sun]

The zeroes of N_z, corresponding to times of sunrise and sunset, occur whenever δ satisfies the equation

$$\cos\alpha\cos\beta\sin\delta = -\sin\alpha\sin\beta$$

Observations:

- A *midnight sun* occurs when the sun is above the horizon for a full day. A midnight sun occurs whenever the sun is away from the equator ($\beta \neq 0$) and the viewer is β degrees away from the pole.

  Then the sunrise-sunset equation becomes:

  $$\sin\beta\cos\beta\sin\delta = -\sin\beta\cos\beta$$

  If you fix $\beta \neq 0$ (sun away from the equator) and put the viewer at $\alpha = \pi/2 - \beta$, the equation becomes:

  $$\sin\beta\cos\beta\sin\delta = -\sin\beta\cos(\beta)$$

3.15 The sun and the viewer locations are symmetric

Note how α, the viewer location, and β, the subsolar location, are on equal footing in the daylight equation.

The excess daylight term vanishes (days and nights have the same length) both when you're at the equator ($\alpha = 0$) in which case the sun's position is irrelevant, and when the sun is at the equator ($\beta = 0$) in which case your viewing position is irrelevant. There's an edge case when you're at the equator and the sun's at one of the poles (or vice versa), in which case the sun is constantly at the horizon.

3.16 The sideways planet and the time-space exchange

On a planet with the largest possible obliquity, the spin axis is actually orthogonal to its orbital axis. (During some parts of the year, in other words, the spin axis sometimes points directly at the sun.) On this kind of *sideways planet*, some marvelous things happen.

First, the subsolar latitude wanders from pole to pole over the course of the year. In general, the subsolar latitude ranges between the tropics at $\pm q$; here, that range encompasses the entire planet.

Specifically, as we've seen, the subsolar latitude β generally obeys the equation: $\sin(\beta) = \sin u \sin q$,[13] which on a sideways planet ($q = \pi/2$) becomes

$$\beta = u$$

so the subsolar latitude corresponds directly to the time of year: at the north pole during northern summer, the south pole during southern summer, and changing linearly in between.

Second, the position of the sun in the sky becomes, with the substitution $\beta = u$:

$$\vec{\mathbf{N}}(\delta) = \begin{bmatrix} -\sin\alpha\cos u\sin\delta \ + \ \cos\alpha\sin u \\ -\cos u\cos\delta \\ \cos\alpha\cos u\sin\delta \ + \ \sin\alpha\sin u \end{bmatrix}$$

Note that the α and u terms can be exchanged, keeping the value of N_z. This means that you can directly interconvert viewer position and time of year.

3.17 Near the equinoxes, the axis tilt effects vanish.

4 Sundials

4.1 Sundials and shadows

A sundial is a device for telling time based on the position of the sun. When a vertical ($\vec{\mathbf{z}}$-aligned) stick (called a *gnomon*) is placed into the ground, the tip of the stick traces out a shadow based on the time of day. Specifically, the tip of the gnomon casts a shadow at the point where a ray passing between the sun and the tip of the gnomon meets the ground plane.

If we choose a coordinate system where the tip of the gnomon is at $\langle 0, 0, 0\rangle$, then by a similar-triangle argument, a light source at $\langle X, Y, Z\rangle$ will cast a shadow onto the ground at $\langle -fX/Z, -fY/Z, -f\rangle$, where f is the height of the gnomon.

For simplicity, let's assume $f = 1$. Let's also ignore the Z component, just considering the 2D shadows being drawn on the ground plane.

[13] Here, u is the year angle, which ranges from 0 to 2π over the course of the year, and q is the obliquity—the angle of spin axis tilt.

Then, in general, we have what I call the sundial principle:

> **Sundial principle.** In a standard sundial setup, a light source at point $\langle X, Y, Z \rangle$ will cast a shadow onto the ground at:
> $$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} -X/Z \\ -Y/Z \end{bmatrix}$$

(This is also, incidentally, the principle for how an idealized pinhole camera works. The pinhole and the tip of the gnomon play analogous roles, as do the ground plane and the back of the camera.)

4.2 Shadows at infinity and phantom shadows

Any light source above the tip of the gnomon $Z = 0$ will cast a proper shadow onto the ground. But the tip of the gnomon at $Z = 0$ represents the limit for shadow-making: a light source at $Z = 0$ will emit a ray of light that passes parallel to the gorund plane and never intersects it. Call this a *shadow at infinity*; as a light source approaches $Z = 0$ from above, its shadow goes further and further away.

Something even stranger happens when the light source is below $Z = 0$. A light source below the tip of the gnomon will still send a ray through the tip of the gnomon. The emitted ray is heading *away* from the ground plane and so will never physically intersect it. But if you imagine rewinding the ray's path *backwards* behind the light source, you'll eventually find a spot that meets the ground plane. Because this places the light source *between* the tip of the gnomon and this imaginary intersection point, it doesn't physically exist. But it has some interesting mathematical properties, so let's call it a *phantom shadow.* (Figure 10).

4.3 A formula for the sundial's shadow position

If we combine our sundial equation with our formula for the sun's position $\vec{\mathbf{N}}(\delta)$ throughout the day (δ varies from 0 to 2π, reaching $\pi/2$ at noon) at a particular viewing latitude (α) and sun's latitude (β), we find that the gnomon's daily shadow follows an arcing curve

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \dfrac{+\sin\alpha\cos\beta\sin\delta - \cos\alpha\sin\beta}{\cos\alpha\cos\beta\sin\delta + \sin\alpha\sin\beta} \\ \dfrac{+\cos\beta\cos\delta}{\cos\alpha\cos\beta\sin\delta + \sin\alpha\sin\beta} \end{bmatrix}$$

4.4 Hyperbola, parabola, circle

Consider the shape that the tip of the gnomon traces out over the course of the day. You can prove that this shape will generally follow one of

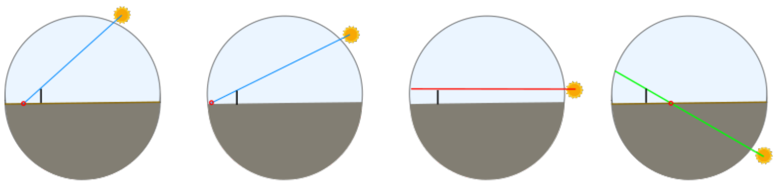

Figure 10: The height of the light source determines what kind of shadow is cast. Real shadows (first two images) occur when the light source is above the gnomon. The lower the light source, the further away the shadow—as with sunrise and sunset. • Exactly at the level of the gnomon (third image), it casts no real shadow, just a shadow "at infinity" in some particular direction. Below the gnomon (fourth image), it casts a "phantom shadow", an imaginary shadow *between* the light source and the gnomon. Phantom shadows are useful for understanding conic sections with two separate branches—one branch is a real shadow cast during the day, and the other is a phantom shadow cast at night.

three possible curves: an ellipse, a parabola, or a hyperbola. These are *conic sections*, the quadratic curves you can draw in two dimensions.

The type of shape is determined by the number of times the sun crosses the horizon when the planet completes a single spin. If you imagine the sun's daily path as a circle located somewhere in the sky, you can see that it can cross the horizon twice, once, or zero times.

In summer near the polar regions, there are times where the sun never sets—it remains in the sky for the entire day. The shadow of the sundial therefore traces out a closed shape, an ellipse.

In the most familiar case, the sun will cross the horizon twice a day—once when it rises, once when it sets—and the shadow will be a hyperbola.

If the conditions are exactly, perfectly right, you get an intermediate case—a parabola. The parabolic shadow occurs when the bottom of the sun's daily circle *just* kisses the horizon, without passing below it. In this way, it manages to touch the horizon just once, a simultaneous dusk and dawn.

4.5 The equinoctal line and the infinite circle

Sundial shadows are, in general, ellipses, parabolas, or hyperbolas. However, in special circumstances, you can get shadows that are special or degenerate versions of these.

On a planet with no tilt, or on a tilted planet during the equinoxes, the sun rises due east, passes directly overhead, and sets due west. As such, the sundial's shadow traces out a very special hyperbola—a straight line[14].

The poles are an exception. On a planet with no tilt, or a tilted planet during the equinoxes, the sun sits exactly at the horizon, in perpetual twilight, without rising or setting. Its long, twilit shadows cast a kind of *circle at infinity*; the gnomon's shadow sweeps out an infinitely wide complete circle.

4.6 The kissing condition for parabolic shadows

A parabolic shadow is a unique intermediate case between the elliptic shadows (which occur in the polar regions of tilted planets) and hyperbolic shadows (which occur everywhere else). A parabolic shadow occurs when the sun's daily circular path touches the horizon exactly once during a full spin. We can show that this happens when the viewer's latitude (α) and the latitude of the subsolar point (β) are complementary.

> **Parabolic kissing condition.** A sundial casts a parabolic shadow when the sun touches the horizon exactly once during the day.
>
> This happens only when the viewer and the subsolar point are in the same hemisphere (neither can be at the equator) and their latitudes are exactly complementary—the viewer is as far from the pole as the subsolar point is from the equator.
>
> Mathematically, the latitudes of the viewer (α) and the subsolar point (β) must satisfy
>
> $$\cos(\alpha + \beta) = 0, \qquad \alpha, \beta \neq 0$$

Observations Because the subsolar point can't be at the equator, parabolic shadows only occur on tilted planets.

Because the latitude of the subsolar point is always changing, the latitude where you can see a parabolic shadow is always changing. You never get an exact parabolic shadow in the same place two days in a row.

To derive the kissing condition, remember that the sun's daily path is a

[14] The mathematical solution is actually a *pair* of straight lines on top of each other—the gnomon casts a real shadow from east to west during the day, then a phantom shadow from west to east at night

circle. In order for the circle to touch the horizon exactly once, it must be resting with its lowest point directly on the horizon[15].

Recall the formula for the height of the sun above the horizon and set it equal to zero:

$$N_z(\delta) = \cos\alpha\cos\beta\sin\delta \; + \; \sin\alpha\sin\beta = 0$$

This is an expression involving the viewer's latitude (α), the latitude of the subsolar point (β), and the time of day (δ, an angle ranging from 0 to 2π over the course of one day and equal to $\pi/2$ at noon).

Right away, note that we must prevent the case where the coefficient $\cos\alpha\cos\beta$ is zero, because then the sun's height won't depend on time of day—it'll be at a constant height, which means it can't cross the horizon at exactly one value of δ. So we won't allow the viewer or the subsolar point to be at the poles $\alpha, \beta \neq \pm\pi/2$.

With that in mind, the sun's unique lowest point occurs at midnight, $\delta = -\pi/2$, in which case its height is

$$N_z(-\frac{\pi}{2}) = -\cos\alpha\cos\beta + \sin\alpha\sin\beta$$

By a trigonometric identity, this is:

$$N_z(-\frac{\pi}{2}) = -\cos(\alpha + \beta)$$

which is zero just when $\alpha + \beta = \frac{\pi}{2}$ or $\alpha + \beta = -\frac{\pi}{2}$. Combined with our prior restriction that neither α nor β may be at the poles, this is the kissing condition for parabolic shadows.

4.7 The parabolic point teleports between the poles

On a tilted planet, there is always a unique spot where you can see a parabolic sundial shadow—except during the equinoxes, when there is no such spot. (Figure 11.)

4.8 The handedness of circular motion depends on the reference point

TODO

As previously discussed, clockwise motion is always defined relative to a particular viewing plane. Circular motion in a plane can be considered clockwise or counterclockwise, for example, based on if you're

[15] Or with its *highest* point directly on the horizon and the rest of it directly below—but of course if the sun is below the horizon all day, it won't cast a real parabolic shadow, only a phantom one.

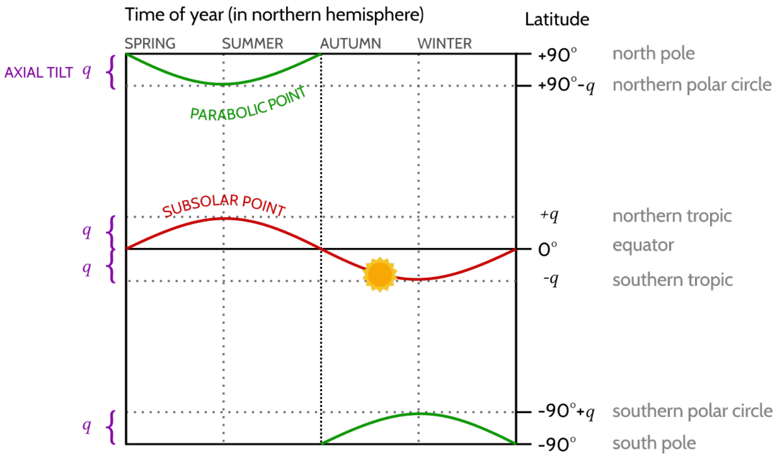

Figure 11: On a tilted planet (obliquity q), there's always a unique spot where you can see a parablic shadow. It moves throughout the year (green curve). At the equinoxes, the parabolic point disappears—and reappears at the opposite pole. At all times, the latitudes of the sun (red) and the parabolic point (green) are complementary angles.

looking at it from one side or the other. Similarly, a compass rose drawn on the ground has the terms North, East, South, West occur in clockwise order, whereas if that same compass rose is lifted up above the viewer and projected onto the sky, those directions appear in counterclockwise order. The handedness of the motion is always defined relative to a particular 'facing' direction, $\vec{\mathbf{z}}$, even though this direction might be implicit.

Given a viewing direction $\vec{\mathbf{z}}$, angular momentum defines the *intrinsic handedness* of any motion: motion is counterclockwise if its angular momentum has a positive $\vec{\mathbf{z}}$ component and clockwise if it has a negative $\vec{\mathbf{z}}$ component.

We can also define an *extrinsic* kind of handedness: suppose an object is moving through space along a path $\vec{\mathbf{x}}(t)$. If I pick any arbitrary pivot point $\vec{\mathbf{y}}$, I can ask whether the object is moving clockwise or counterclockwise around that pivot. Of course, if the path $\vec{\mathbf{x}}(t)$ is especially wobbly, the answer may change in time. But we can determine the answer by computing the angular momentum around the center point

$$\vec{\mathbf{L}}_0 \equiv (\vec{\mathbf{x}}(t) - \vec{\mathbf{y}}) \times \vec{\mathbf{x}}'(t)$$

and determining whether its $\vec{\mathbf{z}}$ component is positive or negative.

The interesting thing is that motion may be clockwise around its own "intrinsic" center, and counterclockwise around another.

For example, take clockwise motion in a circular loop. If the loop doesn't contain the origin, then as you can confirm, its handedness relative to the origin pivot is clockwise half the time and counterclockwise half the time.

For another extreme example, parabolic motion may be intrinsically counterclockwise for all time, but extrinsically clockwise with respect to the origin for all time. Check out the parabola $\vec{\mathbf{x}}(t) = \langle t, t^2 + 1 \rangle$ for example.

This is just to clear up a point about sundial shadows. In shorthand, you can say that sundial shadows generally move clockwise when you're sufficiently far north and counterclockwise when you're sufficiently far south. To be really explicit, this statement is about what we've called the *external* handedness of the motion, relative to the *origin of the sundial* where the gnomon sits.

In fact the shadow itself moves *intrinsically counterclockwise* in the northern hemisphere: from south west to due south to south east. But on a real sundial, usually a circular clock dial has been drawn around the origin. The shadow's motion, which is intrinsically counterclockwise around its center of motion is extrinsically clockwise around the gnomon origin where the time is read off.

4.9 The handedness of a sundial

On an untilted planet (as well as on the equinoxes of a tilted planet), the sun rises due west and sets due east, casting sundial shadows that move in a straight line. In all other situations, the sundial shadow moves in a curved arc, in either a clockwise or counterclockwise direction. The handedness depends on whether the viewer is north or south of the sun's latitude.

In general the rule is:

> **Sundial handedness**. When a sundial's latitude is north of the sun, the sundial's shadow moves clockwise (with respect to its gnomon!) throughout the day. When a sundial's latitude is south of the sun, the shadow moves counterclockwise. Hence sundial shadows move clockwise in the northern hemisphere and counterclockwise in the southern—except within the tropics (latitude $\pm q$) where handedness depends on time of year.

To see this, note that if the sundial is north of the subsolar latitude, then it will also cast shadows to the north. Hence the gnomon's shadow will move from the northwest (at sunrise) to due north (at noon) to northeast (at sunset). It moves from west to east on the northern side of the gnomon, hence it moves clockwise around the gnomon.

An analogous argument applies when the sundial is south of the subsolar latitude, in which case the sundial's shadow moves counterclockwise around the gnomon. And of course, the same handedness argument applies even in the polar regions where the sun may not set.

Beyond the northern tropic, the subsolar point is always to the south year-round; hence north of the tropics, sundial shadows always move clockwise. Similarly for the southern tropic and counterclockwise motion.

Within the tropics, the sun is sometimes to the north and sometimes to the south; hence, the shadow changes from counterclockwise to clockwise and back depending on time of year, with perfectly straight line motion (neither clockwise nor counterclockwise) when the sun is directly overhead.

4.10 Sundials on comets

When we derived the daily motion of the sun, we assumed that the planet's days are much shorter than its years; that is, that the planet completes a full spin around its axis much faster than it completes an full orbit around its star.

Although not all planets work this way, at least our dear Earth does. And it's a convenient assumption because it allows us to assume that the planet has basically no orbital movement over the course of the day—it's essentially *spinning in place.* Thus the daily motion of the sun follows circular paths because we, the viewers, are pirouetting next to it.

Interestingly, then, these sundial results don't depend on the shape of the orbit. They apply even on other astronomical bodies. For example, while planets are generally on closed, elliptical orbits, comets may be on open orbits—parabolas or hyperbolas—that swoop just once around the sun before flying back into the infinite dark.

But if you find a spinning comet out in space and put a sundial on it, you can expect to see the same results we've derived here: assuming the comet is spinning pretty quickly relative to its orbital movement, you'll see the sun make an approximately circular daily path, and the sundial will trace out a hyperbola, parabola, or ellipse based on how many times the sun's path crosses the horizon[16].

You could even keep a sundial on a spinning rocketship that's using its engines to follow a zig-zagging path of its own. As long as the

[16]One obstacle you might encounter, if we pause our geometric daydreaming to think of physics, is that the sun's light gets dimmer the further away you get. Eventually a real sundial's shadow would become too faint to see!

rocketship's heading doesn't change much, the sundial will exhibit the usual range of behaviors.

4.11 Sundials on walls

On an ordinary sundial, the gnomon is in the $\vec{\mathbf{z}}$ direction, pointing straight up. However, you could mount a gnomon on a tilted plane pointing in some other direction $\vec{\mathbf{n}}$.

Under these conditions, we want to know: how will the sun's shadow appear on the plane at various latitudes? What are the conditions for hyperbolic, parabolic, and elliptic shadows? The example of a sundial mounted on a west-facing wall ($\vec{\mathbf{n}} = \overrightarrow{\mathbf{west}}$) will provide concrete context.

For a west-facing sundial, the usual north-west-up coordinate system rotates into the north-down-west coordinate system. In the old coordinate system, the motion of the sun is:

$$\vec{\mathbf{N}}_{\vec{\mathbf{z}}}(\delta) = \begin{bmatrix} -\sin\alpha\cos\beta\sin\delta \;+\; \cos\alpha\sin\beta \\ -\cos\beta\cos\delta \\ \cos\alpha\cos\beta\sin\delta \;+\; \sin\alpha\sin\beta \end{bmatrix}$$

From this new coordinate system, the daily motion of the sun becomes:

$$\vec{\mathbf{N}}_{\overrightarrow{\mathbf{west}}}(\delta) = \begin{bmatrix} -\sin\alpha\cos\beta\sin\delta \;+\; \cos\alpha\sin\beta \\ -\cos\alpha\cos\beta\sin\delta \;-\; \sin\alpha\sin\beta \\ -\cos\beta\cos\delta \end{bmatrix}$$

What are the qualitative behaviors of this sundial? Note that the height of the sun above the sundial is $-\cos\beta\cos\delta$. It is only visible above the sundial in the afternoon ($\delta \geq \pi/2$). Furthermore, because the sun's motion becomes level with the gnomon twice a day at noon and midnight ($\delta = \pm\pi/2$), the shadows of a west-facing gnomon are—unusually—all hyperbolas, regardless of latitude and time of year.

In general, for any gnomon orientation $\vec{\mathbf{n}}$, there is a unique rotation that sends $\vec{\mathbf{z}}$ onto $\vec{\mathbf{n}}$, and the north-west-up coordinate system onto the tilted gnomon's coordinate system, by rotating around $\vec{\mathbf{z}} \times \vec{\mathbf{n}}$. The exception is $\vec{\mathbf{n}} = -\vec{\mathbf{z}}$, in which case there are many possible rotations; we will choose the one that sends north-west-up onto north-east-down.

Let us find the daily movement of the sun in a tilted gnomon's coordinate system. Any gnomon orientation $\vec{\mathbf{n}}$ can be expressed in north-west-up coordinates $\vec{\mathbf{n}} = [n_x, n_y, n_z]$, in which case

5 Daylight on exotic orbits

5.1 Exotic daylight on comets

While planets are generally on closed, elliptical orbits, other astronomical bodies may be on open orbits—parabolas or hyperbolas—that swoop just once around the sun before flying back into the infinite dark.

On an open orbit, a year is no longer a cycle. There are just two epochs in eternity: approaching and receding.

How does the obliquity affect things?

How will the sun appear against the fixed stars? Will the sun approach some signs asymptotically—are there vanishing point constellations?

(What if the orbit's an ellipse?)

5.2 Sundials on comets

While planets are generally on closed, elliptical orbits, other astronomical bodies may be on open orbits—parabolas or hyperbolas—that swoop just once around the sun before flying back into the infinite dark.

It's interesting to consider what a sundial would see on an open orbit.

6 Physics and other complications

6.1 The atmosphere interacts with daylight

The atmosphere lights up before the sun arrives and after it departs, meaning that 'civil' sunrise and sunset don't occur exactly when the sun disappears behind the horizon.

As a compounding effect, the atmosphere bends light, warping the apparent position of the sun.

6.2 The precession of the equinoxes

6.3 Orbits are eccentric, not just circular

6.4 Orbital speed varies with position

6.5 Seasonal lag: On wet planets, temperature change lags seasonal change