The Maxwell Top

Eric D. Black
California Institute of Technology

April 11, 2014

The dynamics of a rigid, spinning body are some of the richest and most mathematically interesting in all of classical physics. They form the basis of our understanding of a whole host of natural phenomena ranging from precession of the equinoxes all the way down to nuclear magnetic resonance. In this lab you are going to study one kind of relatively-simple rotational motion, that of a top. Now, you have probably seen tops before and may even have played with them as a child, so you will not be surprised when I define a top as an axially-symmetric object supported at one fixed point, usually the bottom.

Figure 1: A spinning toy top.

Most tops are straightforward in their construction and have their center of mass located above their balance point, as shown in the example in Figure 1. This leads to some interesting dynamics, including precession and nutation, which you may remember from the tail end of the first term of your freshman physics lecture course, Ph1a. (See Chapter 15, "Gyroscopes" of your freshman physics textbook [1]. If you're interested, you can also read the paper that describes the development of this particular experiment [2].) In this lab you are going to study a variation on the symmetric top, usually attributed to James Maxwell, where the center of mass can be adjusted so that it falls *below* the pivot point, leading to some interesting behavior.

An object having its center of mass below its support, or pivot point, may also be familiar to you if you have ever seen "dynamic-art"

sculptures, such as the one shown in Figure 2. These appear to defy gravity by balancing upright when it looks like they should fall over, but how they work is really very simple. The center of mass is just below the pivot point (the tips of the stick-figure's toes in Figure 2), and so if the sculpture tilts, any torque exerted by gravity tends to bring the figure upright, restoring its "balance" and making this configuration a stable equilibrium. In fact, this system is technically a pendulum and will rock back and forth if tilted and then released.

The example of Figure 2 has two support points, but the idea works just as well with only one, as in Figure 3, where the little man appears to be balancing on the tip of a jackhammer. You could spin this second sculpture and make it a top, and in fact it would be a Maxwell-type top because its center of mass is below its balance point. It would exhibit all the interesting dynamics of the traditional Maxwell top, but analyzing its motion would be difficult because of its complicated shape. Its *moment-of-inertia tensor* would be a mess, and the theory of its dynamics would be more trouble than it is worth. Making a top that has *axial symmetry*, *i.e.* whose shape is symmetric about one axis, greatly simplifies things, and that is exactly what we have for you to experiment with in this lab.

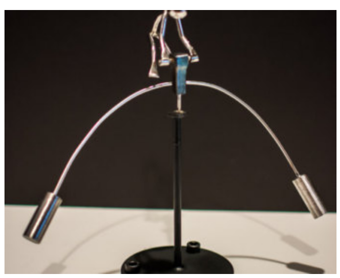

Figure 3: Another example of a dynamic sculpture, this time with a single balance point.

1 Dynamics of a spinning top

1.1 Gyroscopic motion

To understand the dynamics of the Maxwell top, first consider a plain, ordinary top, as shown in Figure 5. In this figure, the top is spinning about the z-axis with an angular rotational speed ω. We may apply a torque to it to

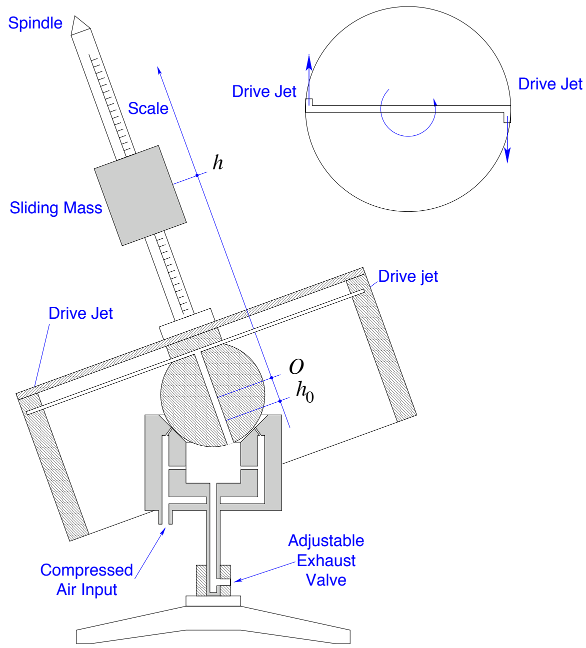

Figure 4: The top you will be using in this lab. Its pivot point is a smooth ball inside a cup, floated on an air bearing to reduce friction. The pivot point is located at the center of the ball and is labeled "O" in this drawing. Most of the mass of the top is in a heavy, steel "skirt" that hangs below the pivot point, but there is a sliding mass that you can raise or lower to move the center of mass above or below the pivot point. Drive jets tap air from the bearing, diverting it out in such a way as to spin the top up and keep it going.

change this rotational speed by applying a force F to the edge of the top as shown. (The air jets will do this in our top.) The torque on the top τ will then be

$$\tau = rF$$

where r is the radius between the center of mass and the point where the force F is applied. You are probably familiar with Newton's law in the case of linear motion, $F = ma$. For rotating objects the same law applies, except that it is written

$$\tau = I\frac{d\omega}{dt}$$

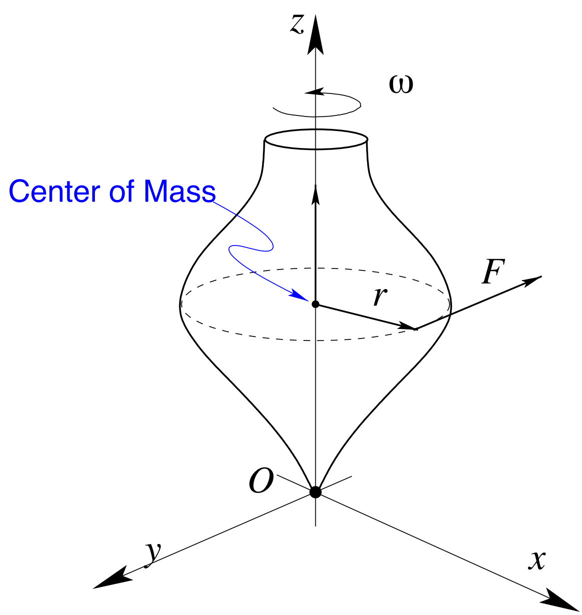

Figure 5: Symmetric top subject to an external force $\vec{F}$ which produces a torque $\vec{\tau} = rF\hat{z}$.

where torque takes the place of force, I, the moment of inertia, takes the place of mass, and $d\omega/dt$ is the angular acceleration. As long as the top is *rigid, i.e.* not changing shape, we can rewrite the second equation as

$$\tau = \frac{d(I\omega)}{dt} = \frac{dL}{dt}$$

where $L = I\omega$ is the angular momentum.

As long as the torque is applied in such a way as to increase (or decrease) its rotational speed around the $\hat{z}$-axis, this is just a one-dimensional equation and offers no surprises. If, however, we try and *tilt* the rotational axis through an angle ϕ, as shown in Figure 6, the situation gets a little more complicated, and a lot more interesting.

When we introduce motion in three dimensions, it becomes helpful to write the necessary equations in vector form. In vector terms, the expression for torque and Newton's law can be expressed as

$$\vec{\tau} = \vec{r} \times \vec{F}$$

($i.e.$ torque is the cross product of the moment arm with the applied force) and

$$\vec{\tau} = \frac{d\vec{L}}{dt}$$

($i.e.$ torque is equal to the change in angular momentum with time).

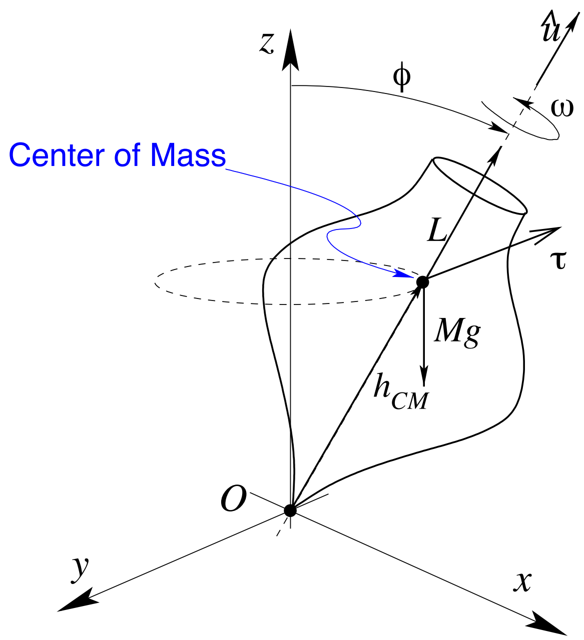

Figure 6: Precession of the top. In this sketch $\vec{\tau}$ is parallel to the y axis and is pointing to the negative direction of the y axis. The vector $\vec{h_{CM}}$ is in the plane Oxz

In the case of the tilted top shown in Figure 6, gravity pulls down on the center of mass of the top, which would pull a non-spinning top downward and simply increase the tilt angle ϕ as the top falls over. If the top is spinning, however, it has an angular-momentum vector $\vec{L}$ that points along the top's axis $\hat{u}$. The torque, and thus the change in the angular-momentum vector, is perpendicular to the axis $\hat{u}$, which leads the top to move "sideways" in a circle around the z-axis, and this motion is called *precession* or sometimes *gyroscopic motion*.

1.2 Quantitative prediction of precession

To be *science*, any predictions you make have to be not just testable but *falsifiable, i.e.* able to be proven wrong. To be called physics, both your theory and your observations have to be quantitative, *i.e.* in the form of numbers.

So, let's make a quantitative prediction about precession that we can test in the lab. Let's calculate what we think the precession rate will be. First, we evaluate the torque that produces it. The magnitude of the torque $\vec{\tau}$ due to gravity $M\vec{g}$ is

$$\tau = \left|\vec{h}_{CM} \times M\vec{g}\right| = h_{CM} M g \sin\phi, \tag{1}$$

where ϕ is the angle of the top's axis from the vertical, and $\vec{h}_{CM}$ is the vector pointing to the top center of mass. Because $\vec{g}$ is always vertical, $\vec{\tau}$ must always lie in the horizontal plane, and therefore the change in angular momentum $d\vec{L}$ must also be in the horizontal plane.

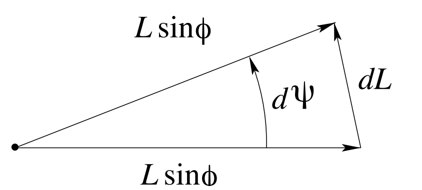

Figure 7: Projection and variation of the angular momentum in the horizontal plane.

To find the precession rate, we start with the projection of $\vec{L}$ in the horizontal plane. Remember that the magnitude of $\vec{L}$ does not change, only its direction, and that is true for the projection in the horizontal plane as well (see Figure 7)

$$dL = L\sin\phi d\psi,$$

where $d\psi$ is the change in the direction of this projection.

This gives us

$$L\sin\phi\frac{d\psi}{dt} = \tau,$$

which we can combine with Equation 1 to get

$$\frac{d\psi}{dt} \equiv \Omega = \frac{h_{CM} M g}{L}.$$

Substituting $L = I\omega$ gives us

$$\boxed{\Omega = \frac{Mg}{I\omega} h_{CM}} \tag{2}$$

This is a clear and testable (falsifiable) prediction. The precession frequency Ω is something we can measure, and all of the quantities on the right side of Equation 2 are things we can measure as well. Note that Ω does not depend on the angle ϕ. This is another kind of prediction known as a *null hypothesis, i.e.* that the outcome of an experiment does *not* depend on a particular input. More on this later.

2 Experimental setup

As mentioned earlier, the Maxwell top you will be using in this lab is illustrated in Figure 4. The base of the top is a spherical ball, which rests inside a cup. The ball is part of the top, the cup is part of the base, and the two are machined to tight tolerances to have the same radii of curvature and thus fit closely together, allowing the top to pivot or spin on the base. The ball and cup (or ball and socket, for you biology or mechanical-engineering students) must have some lubrication between them to work properly, and in our case that lubrication is provided by compressed air. Now is a good time to emphasize, *under no circumstances, ever, put the ball inside the cup without lubrication, i.e. air flowing!* If you do, the ball and cup will scratch one another, ruining both the top and its base. Also, these tops are individually fitted to their bases. *Never mix tops and bases of different numbers!* The radii of curvatures are matched, and putting a ball in the wrong-sized cup could also scratch and ruin both. Each base and each top is individually numbered so that you don't mix them up.

2.1 Safety

Now is an excellent time to cover safety, not only yours, but that of the hardware as well. For your safety,

1. Wear covered shoes, in case you drop a top on your toes. You should be doing this any time you are in a lab, anyway.

2. Remove jewelry, especially rings on your fingers or necklaces that dangle and might get caught in the spindle, before working with a spinning top. If you have long hair tie it back. These are the same precautions you would take before working with a lathe.

3. Handle the top gently. Don't grab it while it's spinning, and if you choose to put your hands lightly on it to help it spin down, be careful of friction. Your hands can heat up quickly doing that, and any rings you might have forgotten to take off can get damaged or snagged, which will make you unhappy, even if you are lucky enough not to get injured.

Your own safety is your first priority, but you should also give some consideration to the hardware.You'll heal, after all, but the tops won't. They are not just older than you are, they are older than I am! Let's keep them in good shape for future generations, like those who came before us did for us. With that in mind,

1. Always run a decent amount of air pressure while you are working with these tops, at least 30 psi any time the top is in the base, whether it is spinning or not.

2. Do not switch tops with different bases. Each top has its own base and only works properly with that one. You can damage a top, a base, or both by mixing them.

3. Make sure the ball and socket are both clean and free of grit or any foreign objects before starting. The cups are highly polished, and even a small scratch from a piece of sand or grit can ruin the performance of the air bearing.

4. Clean all tape off the top once you are finished with that part of the experiment. The tops are also carefully balanced, and even a small amount of tape can set up a vibration in the top at high speed.

3 First laboratory week

All but one of the parameters that go into our prediction (Equation 2) are fairly easy to measure, the one exception being I, the moment of inertia of the top. This requires some cleverness, and we'll spend the first session of this lab learning two clever ways to measure it that are actually examples of widely-used experimental techniques. It is often a good idea to measure something you want to know two different ways. It serves as a "reality check" and helps identify any errors or mistakes you might make in a single measurement. Even the best experimentalists sometimes make mistakes. The intelligent scientist is humble enough to admit that he or she is fallible and builds such checks and safeguards into their experimental design.

3.1 The obvious method: acceleration in response to a known, constant torque

The first method you are going to use is to measure the moment of inertia is, as I have said, the obvious one. Apply a known, constant torque, and measure the angular acceleration of the top. This method relies on the torque being constant, *i.e.* not changing as the speed of the top changes, and at first glance you might think you have to take my word for it that this is the case. However, as we will see, there is a way for you to check this yourself. But first, let's measure the torque with the top at rest.

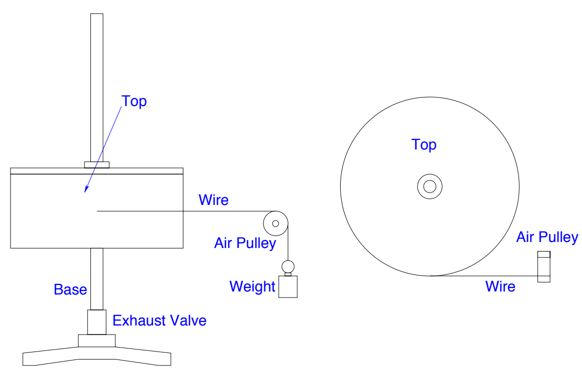

Figure 8: Direct measurement of the torque applied by the air jets by balancing it out with a hanging weight of known mass.

As shown in Figure 8, hang a known weight from a wire draped over a pulley and attached to the edge of the top. Use tape to attach the wire (or string). There are air-bearing pulleys in the lab for this that have very little friction, and your TA can show you how to set them up. Adjust the air flow using the exhaust valve at the base, as shown in Figure 8, until the torque from the jets just cancels the torque from the weight, and the top holds its position. Now, calculate the torque that the jets apply from the weight of the hanging mass and the radius of the top.

Once you have the torque, remove the weight and the tape, but don't change the valve setting. (That would alter the torque and make your previous measurement meaningless.) Now, you need to measure the acceleration that this torque produces. You might be tempted to just haul off and measure the number of revolutions vs. time, but there is a better way to do it.

The total rotation is expected to scale with the square of time.

$$\theta(t) = \frac{1}{2}\alpha t^2$$

Consider the time τ it takes the top to go around one full revolution, starting from rest.

$$\theta(\tau) = \frac{1}{2}\alpha\tau^2 = 2\pi$$

After twice this time (2τ), because of the square, the top will have gone around *four* times. After 3τ, it will have gone around nine times, and so on. In general, after $N\tau$, the top will have made N^2 full revolutions.

$$\theta(N\tau) = \frac{1}{2}\alpha(N\tau)^2 = N^2 2\pi$$

It's a lot easier to count revolutions and hit a stopwatch button after every square (1, 4, 9, 16, ...) than it is to hit the button after each individual revolution, and you'll get data up to a higher speed.

LABORATORY EXERCISE:

1. Do this. Turn the air on to at least 30 psi, set the top on its base, and suspend a known weight (2 grams is a good one) from a string attached to the edge of the top, as shown in Figure 8. Adjust the exhaust valve until you get as close to equilibrium as possible. Calculate the weight of the suspended mass, measure the radius of the top, and from that calculate the torque applied. Estimate the uncertainties in your measurements and your calculation.

2. Remove the weight, string, and any leftover tape, and get ready to measure the angular acceleration. If you have a stopwatch setting on

your watch or app on your phone, you can use that. If not, ask your TA, and they can provide you with one. There is a mark on the top itself and a pointer on a stand nearby. Set the pointer up so that its point is close to the edge of the top, and start with the top at rest and the mark and the pointer aligned. Start your watch when you release the top. Hit the "lap" button on the stopwatch on the first, fourth, ninth, sixteenth, etc. times the mark passes the pointer, for as many rotations as you can. Record your data in your lab notebook.

3. Plot your results in terms of the square root of the number of revolutions, $\sqrt{N}$ vs. time, $N\tau$, and fit a straight line to your data. Calculate the angular acceleration α.

4. Plot your residuals. Are your results consistent with the assumption of a linear relationship between $\sqrt{N}$ and $N\tau$, *i.e.* that the torque is constant for all speeds of the top?

5. Calculate the moment of inertia of the top I from the relation

$$\tau = I\alpha$$

6. If you have gotten far enough in *Taylor* to know how, estimate the uncertainty in your value for I.

3.2 The not-so-obvious method: the loaded torsional pendulum

The moment of inertia affects the frequency of a torsional pendulum in exactly the same way that mass affects the frequency of a linear mass on a spring. Where the natural frequency of a mass on a spring is given by

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{k_\ell}{m}},$$

(where m is the mass, and k_ℓ is the linear spring constant), the frequency of a torsional pendulum is

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{k}{I}},$$

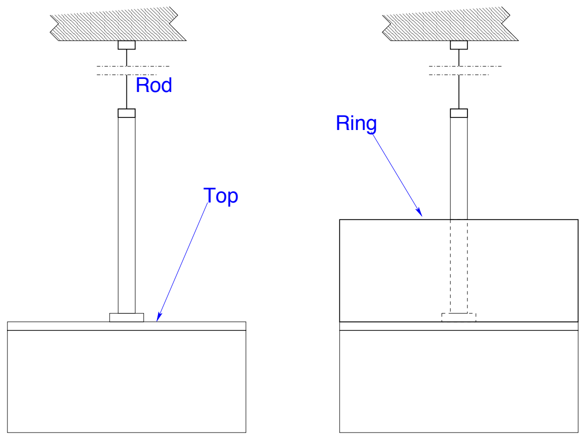

Figure 9: Unloaded (left) and loaded torsional-pendulum setups for measuring the moment of inertia of the top.

where k is, in this case, the rotational spring constant, *i.e.* the restoring torque per unit of angular displacement. If we were to attach our top to a torsional pendulum with a known k, we could then measure its natural frequency and then calculate I. (This is the way astronauts used to measure their body weight during extended missions on the international space station and, before that, skylab. Being weightless, stepping on a bathroom scale doesn't help you find your mass in such an environment. Instead, they would sit in a chair in a contraption that bounced them back and forth with springs of known restoring force. The natural frequency told them their mass, and thus whether they had gained or lost weight over the course of the mission.)

Measuring the rotational restoring force k and then the natural frequency of a torsional pendulum is one way to find I, but it's not very precise. As it turns out, while it's certainly possible to measure k, it is fairly difficult to do so with a high degree of precision. Your uncertainty in k will be much larger than the uncertainty in the frequency, and so the former will dominate in the uncertainty in your answer I. Fortunately, there is a way to bypass the k measurement entirely and find I with much better precision, and that is to introduce a known additional moment of inertia I_1 to the system. This extra "loading" will reduce the natural frequency of the torsional pendulum to

$$f_1 = \frac{1}{2\pi}\sqrt{\frac{k}{I + I_1}}.$$

This, together with the unloaded natural frequency, gives us two equations with two unknowns, I and k, which is something we can solve.

A little algebra will quickly show you that the moment of inertia of the top is then,

$$I = \frac{I_1}{\left(\frac{T_1}{T_0}\right)^2 - 1}, \tag{3}$$

where T_0 is the period of the unloaded top, and T_1 is the period of the top plus the extra moment of inertia I_1.

I'll leave it up to you to solve for the torsional spring constant k, if you want to.

LABORATORY EXERCISE:

1. Above each Maxwell-Top station is a unistrut angle bracket for hanging your torsional pendulum, and on it there should be a metal collar with the same diameter as the top. Your TA can provide you with the rod part of the torsional pendulum. Before you do anything else, you will need to determine the moment of inertia of the ring. The formula for the moment of inertia of such a ring is

$$I_1 = \frac{1}{2} M \left(R^2 + r^2\right),$$

where R is the outer radius, r is the inner radius, and M is the total mass. Take down your ring, make the necessary measurements, and calculate its moment of inertia. If you have gotten far enough in *Taylor* to know how to do it, estimate the uncertainty of your answer, based on your best estimates of the uncertainties of your measurements of R, r, and M.

2. Put the ring back on its bracket (you'll see why in a minute), and then hang your top from the bracket using your torsional rod. Measure the period of this torsional pendulum with just the top suspended. The best way to do this is to time multiple oscillations (ten, twenty, a hundred - however many you have the patience for) and then divide to find a single period. How does the uncertainty in your period measurement depend on the number of oscillations you count?

3. Now lower the ring onto the top, and repeat your period measurement.

4. Calculate the moment of inertia of your top, and compare it with your result from the direct measurement you made with the hanging weight. Are your two values consistent with each other?

4 Second laboratory week

4.1 Dependence on center-of-mass position h_{CM}

Recall Equation 2, where we predicted how the precession frequency would depend on various parameters.

$$\Omega = \frac{Mg}{I\omega} h_{CM}$$

You have probably noticed by now that the top has two parts: the main top, and a sliding mass that can be moved up and down the spindle to adjust the overall center of mass. With the sliding mass installed, the prediction of Equation 2 becomes,

$$\Omega = \frac{(M+m)g}{(I+I_s)\omega} h_{CM} \approx \frac{(M+m)g}{I\omega} h_{CM},$$

where m is the mass of the sliding mass, and I_s is its moment of inertia, which we will neglect from here on out since $I_s \ll I$. The height of the center of mass becomes

$$h_{CM} = \frac{h_0 M + hm}{M+m},$$

where h_0 is the center-of-mass height of the top without the sliding mass installed. This is something you can calculate when you design the top, and for ours it is expected to be 6.03 ± 0.01mm by design. Putting this into our expression for Ω yields, after some algebra,

$$\Omega = \frac{Mg}{I\omega} h_0 + \frac{mg}{I\omega} h \tag{4}$$

Now we have a prediction in the form of a linear equation, where Ω is the outcome of the experiment (the dependent variable), and h is the parameter you vary (the independent variable). You may recall from the second week of class how to fit data to a theory of this form, and now you're going to do it with real data, rather than the canned stuff I gave you then.

LABORATORY EXERCISE:

1. Set the sliding mass at its lowest position on the spindle of the top.

2. Turn on the air to at least 30 psi, set the top on its base, and tilt the spindle. Verify that the spindle tries to return to the vertical, indicating that the center of mass is below the pivot point.

3. Set the sliding mass at its highest position, and again tilt the spindle. This time the top should try to tip over. (Don't let it.) This confirms that you can, in fact, move the center of mass to where it is above the pivot point.

4. Adjust the sliding mass until you think the center of mass of the system is right at the pivot point. Record this position in your lab notebook.

5. Move the sliding mass back to the bottom of the spindle, set the spindle to an angle of about 30° from the vertical, and spin it up.

6. Using either a strobotach or a timing light connected to a function generator, measure the rotational frequency of the top ω. Don't forget the necessary conversion factors between angular frequency ω, ordinary frequency f, and RPMs (revolutions per minute)! Adjust the bleed valve on the top's base to get smooth, relatively-constant rotation.

7. Measure the period of precession. You can use one of the stand-mounted pointers to designate a beginning and ending point for your precession measurement. Your rotational frequency may not be as stable as you like, and if this is the case you can measure it before and after your precession measurement and take the average.

8. Repeat this process for at least two more evenly-spaced positions of the sliding mass, for a minimum of three total data points. More data points are certainly better, but only do them if you have the time. Record your results in your lab notebook.

9. Do a linear fit to your data. Do your fit parameters agree with your prediction from Equation 4? Does the line pass through zero where you expected it to, based on where you put the sliding mass to get the center of mass right at the pivot point?

10. The design value for the center-of-mass height was quoted in the original paper as 6.03 ± 0.01mm [2], but the author of that paper did not

make it clear whether that was the value for h_0 or h_{CM} with the sliding mass set at the bottom of the spindle. Based on your fit parameters, which do you think it is?

11. If we had not neglected the moment of inertia of the sliding mass I_s, would it have made a difference in your fit? Why?

4.2 Dependence on ϕ, null experiments

In Equation 2 we made a prediction about the precession rate as a function of the height of the center of mass h_{CM} and the parameters M, g, I, and ω. Notably absent from this list was ϕ, the angle the spindle makes from the vertical. This absence is not a simple omission. It is a prediction, and that prediction is that the precession rate will be the same regardless of the value of ϕ. This is a bold claim, when you really think about it, and not something that you could ever fully test. How, for example, could you tell that there was absolutely *no* effect of ϕ on Ω, as opposed to just a very small one that is too low for your instruments to detect? The answer is that you can't. All you can do is set *upper limits, i.e.* if there is an effect, it is smaller than the resolution of your instrument.

Setting upper limits on quantities that may be zero forms a special class of experimental technique all its own, known as *null experiments.* The hypothesis that a particular variable has no effect on the outcome of an experiment is known as a *null hypothesis.* There has been a great deal written about null experiments, and a lot of very smart people have put a great deal of thought into their development and interpretation. Some of the most famous and significant null experiments were what we now call the Eötvös experiments, which tested the *equivalence principle, i.e.* that gravitational and inertial mass are the same thing. Newton was the first to do such an experiment, but Eötvös ran with the idea in the nineteenth century. The equivalence principle is a cornerstone of general relativity, and the Eötvös experiments provide direct experimental justification for it. Most modern tests of string theory and searches for dark matter express their results in terms of upper limits, invariably because they have failed to observe a definitive signal.

In this part of the experiment I want you to pretend that there is some variation of string theory that predicts a ϕ dependence with some unknown parameter λ.

$$\Omega(\phi) = \frac{Mgh_{CM}}{I\omega}(1 + \lambda\phi) \tag{5}$$

This is not as much of a stretch as it might at first seem. That's the problem with string theory. There are a lot of variations, and they predict practically anything you can think of.

LABORATORY EXERCISE:

1. Run the sliding mass all the way down to the bottom of the spindle, set the air pressure to at least 30 psi, and set the top in its base.

2. Measure the period of precession like you did before for three different spindle angles. You can choose any angles you want, but they should be as widely separated as you can get and still have the top work properly, *e.g.* 10°, 30°, and 45°.

3. Estimate your uncertainty in Ω, based on the uncertainty in your period measurements.

4. Plot Ω vs. ϕ, and fit a line representing Equation 5 to it.

5. From the slope of your line, calculate a value for λ and its associated uncertainty. Is this consistent with our previous expectation that $\lambda = 0$? If so, what limits can you place on λ? Remember that λ could also be negative, so in this case you may conceivably place both an upper and lower limit on its value.

References

[1] Steven C. Frautschi, *et al.*, *The Mechanical Universe: Mechanics and Heat, Advanced Edition*, Cambridge University Press; Advanced Ed., (2008).

[2] H. V. Neher, *Air-Bearing Maxell Top*, Am. J. Phys. **30**, No. 7, 503-506 (1962).