8 Yaw-Wise Torque Budget

In aircraft (unlike cars, bikes, or small boats) you have separate control over which way it is *pointing* relative to which way it is *going*.

8.1 Overview

This chapter discusses the yaw-wise motion of the airplane, which has to do with which way the airplane is *pointing*.[1]

Normally you want the aircraft to be *pointing* the same direction as it is *going* through the air. That is, you want the slip angle to be small. There are several reasons for this:

- Precision: If your objective is to turn to the left, it doesn't make sense to let the maneuver begin with a big inadvertent yaw to the right.
- Efficiency: Slipping creates unnecessary drag.
- Comfort: Passengers really hate being sloshed from side to side. Maybe it doesn't bother you, but it will bother your passengers. Also note that in many small aircraft, passengers are at a mechanical disadvantage because they are seated farther from the pivot point (the center of mass) than the pilots are. That means any given yaw *angle* produces more sideways *displacement* at the passengers' location.
- Safety: Whereas if you stall in coordinated flight the nose will just drop straight ahead, if you manage to stall in sufficiently uncoordinated flight, you will get a spin (see chapter 18) or a snap roll (see chapter 11), which is much harder to recover from.

Maintaining zero slip angle while maneuvering requires coordinated use of the ailerons and rudder, so pilots speak of "zero slip angle" and "good coordination" almost interchangeably. (Situations that call for an intentional slip are discussed in section 11.2.)

This chapter considers, one by one, the various phenomena that affect the airplane's yaw-wise motion. There are surprisingly many such phenomena, including the helical propwash, yaw-wise inertia, adverse yaw, P-factor, and gyroscopic precession --- plus the stability and damping created by the vertical fin and rudder.

8.2 Yaw Stability

For ordinary objects such as cars, bicycles, and small boats, there is only one "steering" control, and it serves to control both the direction you are pointing and the direction you are going. But the two ideas are not necessarily linked.

As an extreme example of non-linkage, take a Frisbee and draw on it the picture of an airplane. When you throw the Frisbee, the picture of the airplane will be yawing like crazy, turning around and around and around. You have no control over which way it is pointing. You do, however, have some modest control over which way the Frisbee as a whole is going: if there is a nonzero bank angle, the Frisbee will follow a curved flight-path.

Far and away the most powerful technique[2] for changing the direction an airplane is going is to put it into a bank, so that the horizontal component of lift forces a change in flight-path, as mentioned in section 4.3 and especially section 6.2. This is not yaw; bank by itself will not change the direction you are pointing.

The vertical fin and rudder are responsible for controlling the yaw angle, which is the main topic of this chapter.

An airplane has *partial* linkage between the direction you are going and the direction you are pointing. That is:

- It has less linkage than a car. For example, in preparation for a crosswind landing, you may well choose to have the airplane going through the air in one direction while pointing in another direction.
- It has more linkage than a Frisbee (which utterly lacks a rudder). Most of the time, you want the airplane's slip angle to be small.

Figure 8.1 shows a situation where the airplane's heading has been disturbed out of its usual alignment with the airflow. There are lots of ways this could happen, including a gust of wind, a momentary uncoordinated deflection of the controls, or whatever.

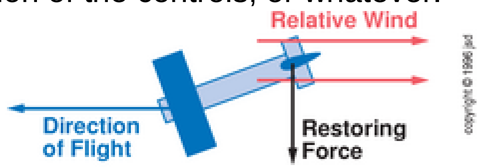

Figure 8.1: Response to a Yaw Angle

In this situation, the relative wind is striking the vertical fin and rudder at an angle. Like any other airfoil, the fin/rudder produces lift in proportion to its angle of attack, so it will produce a force (and therefore a torque) that tends to re-align the airplane with the wind. We say that the airplane has lots of yaw-wise stability.

The colloquial name for yaw-wise stability is "*weathervaning tendency*". That is, the airplane tends to align itself with the relative wind, just as a weathervane does. Section 8.11 discusses weathervaning during taxi.

8.3 Yaw Damping

In most airplanes, pure yawing motions are reasonably well damped. The process is analogous to the process that produces damping of pure vertical motions and pure rolling motions (see chapter 5). When the tail is swinging to the right with an appreciable velocity, it sees a relative wind coming from ahead *and to the right*. The resulting angle of attack produces a leftward force that damps the rightward motion.

A leftward force in proportion to a rightward velocity is exactly what constitutes damping.

In some airplanes, there is a lightly-damped Dutch roll mode, involving the yaw axis along with others, as discussed in section 10.6.1.

8.4 Helical Propwash

One of the very first things that people find out about when they start learning to fly is that it takes right[3] rudder (sometimes a lot of right rudder) to keep the airplane going straight at the beginning of the takeoff roll. The physics of the situation is portrayed[4] in figure 8.2.

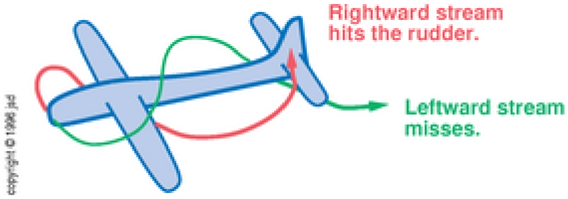

Figure 8.2: Helical Propwash

It would be nice if the propeller would just take the air and throw it straight backwards, but it doesn't. The propeller airfoil necessarily has some drag, so it drags the air in the direction of rotation to some extent. Therefore the slipstream follows a helical (corkscrew-like) trajectory, rotating as it flows back over the aircraft.

The next thing to notice is that on practically all aircraft, the vertical fin and rudder stick up, not down, projecting well above the centerline of the slipstream. That means the helical propwash will strike the left side of the tail, knocking it to the right, which makes the nose go to the left, which means you need right rudder to compensate.

You don't notice the effect of the helical propwash in cruise, because the aircraft designers have anticipated the situation. The vertical fin and rudder have been installed at a slight angle, so they are aligned with the actual airflow, not with the axis of the aircraft.

In a high-airspeed, low-power situation (such as a power-off descent) the built-in compensation is more than you need, so you need to apply explicit left rudder (or dial in left-rudder trim) to undo the compensation and get the tail lined up with the actual airflow.

Conversely, in a high-power, low-airspeed situation (such as initial takeoff roll, or slow flight) the helix is extra-tightly wound, so you have to apply explicit right rudder.

Helical propwash sometimes contributes to left/right asymmetry in multi-engine aircraft, as discussed in section 17.1.12.

8.5 P-Factor

The term *P-factor* is defined to mean "asymmetric disk loading". It is an extremely significant effect for helicopters. When the helicopter is in forward flight, the blade on one side has a much higher airspeed than the other. If you tried to fly the blades at constant angle of attack, the advancing blade would produce quite a bit more lift than the retreating blade.

8.5.1 Blade Speed

For airplanes, the same effect can occur, although it is usually small. For the effect to occur at all, you need to have an angle between the propeller axis and the relative wind. To be specific, imagine that the aircraft is in a nose-high attitude, but its direction of motion is horizontal (i.e. the relative wind is horizontal). Then the downgoing blade will be going down and a little bit forward, while the upgoing blade will be going up and a little bit backward. The downgoing blade will effectively have a slightly higher airspeed. Since this blade is on the right-hand side of the airplane (once again assuming a typical American engine) it will tend to torque the airplane around to the left and you'll need right rudder to compensate.

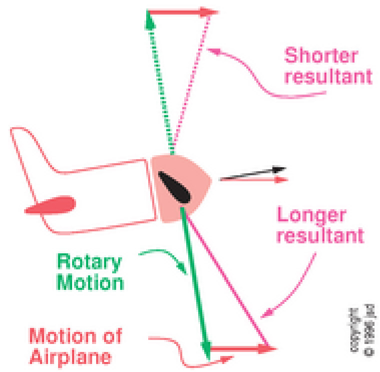

Figure 8.3: P-Factor

The situation is depicted in figure 8.3. The airplane is in level flight, with a 10 degree nose-up attitude. The motion of the blade through the air is shown in magenta. It consists of the rotational motion (shown in green) plus the forward motion of the whole airplane (shown in red). The motion of the downgoing blade is shown with solid lines, while the motion of the upgoing blade is shown with dotted lines. You can see that the speed of the downgoing blade is larger than the speed of the upgoing blade.

This is the main contribution to P-factor: the advancing blade sees more relative wind, while the retreating blade sees less relative wind.

8.5.2 Blade Angle

There is a widespread misconception that P-factor arises because the angle of the right (downgoing) propeller blade is larger than the angle of the left (upgoing) propeller blade. Many books erroneously call attention the angle of the blade relative to the ground. The blade doesn't care about the ground; the only thing that matters is the angle of attack, i.e. the angle between the blade and its own motion through the air.

The correct analysis is shown in figure 8.4. As a point of reference, the left panel shows level pitch attitude in normal level flight, where no P-factor occurs. Meanwhile, the right panel shows the airplane in a 10 degree nose-up attitude (still in level flight). Since we want to discuss angle of attack, I have attached a "reference line" pointer to each of the blades, just like the reference line used in section 2.2. The angle of attack of the propeller blade is just the angle between the reference line and the blade's motion through the air.

You can also think of the blade's angle of attack as the angle between the reference and the blade's relative wind. Relative wind and direction of motion are the same concept, just reversed 180 degrees. Be careful though, because there are various different relative winds, including the instantaneous wind relative to the moving blade and the average wind relative to the overall airplane.

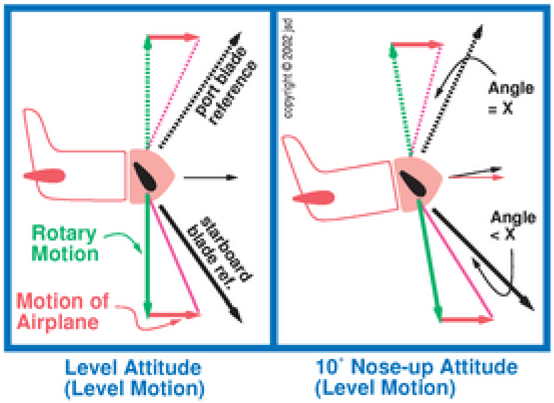

Figure 8.4: P-Factor and Angle of Attack

When the propeller disk is inclined to the direction of flight (so that P-factor really is occurring) the downgoing blade has a slightly greater angle of attack (compared to the upgoing blade) as shown in figure 8.4. This occurs because the vector representing the airplane's motion has "better leverage" when it meets the resultant, because the resultant is shorter and because it is more nearly perpendicular to the airplane's motion.

This angle-of-attack effect is of course zero when propeller axis is aligned with the direction of flight.[5] The effect is never very large, because

1. At low speeds, the airplane's forward velocity (as represented by the horizontal red arrow in figure 8.4 is so small that it can't have much effect on anything.
2. At high speeds, the airplane has a low angle of attack, so the angle between the propeller disk and relative wind is necessarily small (except for helicopters, tilt-rotors, and such).
3. At very high speeds, when you are going fast enough to over-run the geometric pitch of the propeller (so that the resultant coincides with the reference line in figure 8.4), you might think that a small difference in angle of attack would be a 100% effect. I suppose that's true, but in this case the total thrust is practically zero, and 100% of nothing is nothing.

This angle-of-attack effect is in addition to (and usually smaller than) the airspeed effect discussed previously. Both are small compared to the helical propwash effect.

Remember, we don't care whether the downgoing blade makes a bigger angle *to the vertical* than does the upgoing blade. The blade doesn't care which way is up --- all it cares about is where the relative wind is coming from. Imagine a tailwheel-type airplane stationary in the run-up area on a windless day. You can incline the propeller disk as much as you want relative to vertical, but there will be no P-factor unless there is *wind* blowing through the propeller disk at an angle.

8.5.3 Initial Takeoff Roll

There are quite a lot of myths surrounding P-factor. For some reason, P-factor gets blamed for the fact that typical aircraft require right rudder on initial takeoff roll. This is impossible for several reasons.

- Nearly everybody these days learns to fly in nose-wheel type aircraft, which means the propeller disk is vertical during the initial the takeoff roll. Since there is no angle between the relative wind and the propeller axis, P-factor obviously cannot occur.
- Now let's suppose, just for sake of argument, that you are flying a taildragger, in which the propeller disk is actually non-vertical during the initial takeoff roll. Common experience is that

the most right rudder is required at the very beginning of the takeoff, before much forward speed has been achieved. The FAA **Airplane Flying Handbook** (reference 15) says this is because P-factor is worst at low airspeeds. But wait a minute --- real P-factor is proportional to airspeed. In the initial moments of the takeoff roll, there is no relative wind, so there can't possibly be any P-factor. Of course, if you are taking off into a headwind, there could be a little bit of P-factor --- but does that mean if you take off with a slight tailwind there will be a negative amount of P-factor, requiring left rudder? Don't bet on it.

The real reason that you need right rudder on initial takeoff roll is because of the helical propwash, as discussed in section 8.4. P-factor exists in some circumstances, but it cannot possibly explain the behavior we observe during initial takeoff roll.

8.5.4 Observing P-Factor

It is not easy to observe P-factor. It is usually swamped by other effects such as helical propwash (section 8.4) and twisted lift (section 8.8.4).

An important preliminary experiment is to observe what happens during the takeoff roll in a multi-engine aircraft. (To be specific, let's consider the case where both engines rotate clockwise as seen from the rear.) In some airplanes, the propwash hits the tail, and you must apply right rudder to compensate, just like in single-engine planes.

In other airplanes, most of the helical propwash misses the vertical tail in normal flight. This causes no problems and no compensation is required. See section 17.1.12 for details on this. This is the perfect way to illustrate that there is no P-factor when the propeller disk is not inclined.

If you really want to observe nonzero P-factor, you can proceed as follows: Take a twin-engine (or four-engine) aircraft with non-counter-rotating propellers. Attach a slip string. Establish coordinated cruising flight, with the same amount of power on both sides. Confirm that the ball and string are centered. Pull up into a nonturning climb at very low airspeed (i.e. very high angle of attack), maintaining cruise power. Maintain coordinated flight as indicated by the slip string. Observe the rolling tendency due to propeller drag. Shift weight (e.g. fuel) from left to right to get rid of the rolling tendency, so you can fly straight *without* deflecting the ailerons, i.e. without incurring any twisted lift.

You will observe the inclinometer ball will be slightly off-center. This can be attributed to P-factor. To be explicit:

- Each engine is producing the same *amount* of thrust.
- But the result is asymmetric thrust, because the *location* of the thrust vectors is shifted (because of asymmetric disk loading).
- We have gone to some trouble to minimize other contributions, such as twisted lift. Also note that even if there were some helical propwash hitting the tail, that wouldn't be a good explanation for the off-center inclinometer. You should be able to compensate for helical propwash using the rudder alone, without banking.
- To summarize: we have symmetric power settings but asymmetric thrust. A detailed analysis of the effects of asymmetric thrust can be found in section 17.1.5.

The effect of P-factor is not very large. You can easily compensate using a little bit of right rudder and right bank. Indeed, in typical situations you can just ignore it entirely.

You can never use rudder deflection as an indication of P-factor, because any situation that exhibits

P-factor will also exhibit a large amount of helical propwash.

The single-engine version of the previous experiment goes like this: Put a slip string on each wing, far enough out on the wing that it is not unduly disturbed by the propwash, yet close enough in that you can see it. In a high-wing aircraft, you'll have to put the string on the bottom of the wing. Put strings on both sides in symmetric locations, so you can tell for sure what string position corresponds to symmetric airflow. Then confirm that in normal nonturning cruising flight, you have symmetric airflow (as indicated by the strings) and zero inclination (as indicated by the inclinometer ball). Finally, set up a situation in which the largest possible P-factor occurs: flaps retracted, minimum airspeed, and full power.

Once again, the indication of P-factor in this situation would be to have the ball be off-center when the strings were centered. I have tried this experiment, but the P-factor was too small to observe.

Here's another possible experiment. Take your favorite aerobatic airplane and paint the starboard rudder pedal green and the port rudder pedal red, just so we can keep straight which is which. Now go to a safe altitude and set up for inverted slow flight. In this high-power, low-speed situation, do you need to push the port (red) pedal or the starboard (green) pedal? If P-factor is more important, the answer will be port, because that is now the downgoing, advancing blade. If helical propwash is more important, the answer is starboard, because the relationship between the propeller, rudder, and rudder pedals is unchanged by the inversion.

8.6 Gyroscopic Precession

A spinning object will respond to a torque in one direction with a motion in another direction. This remarkable and counterintuitive phenomenon --- *gyroscopic precession* --- is discussed in more detail in section section 19.9.

Gyroscopic precession is often quite noticeable at the point where a taildragger raises the tail, early in the takeoff roll.[6] If the airplane were an ordinary non-spinning object, you could raise the tail using the flippers alone. The flippers do not actually dictate the *motion* of the fuselage; they just produce a *force* and a pitch-wise torque. For a gyroscope, this pitch-wise torque produces a yaw-wise motion. If you try to raise the tail of a real airplane using flippers alone, it will yaw to the left because of precession.

To get a gyroscope to actually start *moving* in the pitch-wise direction, you need to apply a torque in the yaw-wise direction. This is what the vertical fin and rudder are for. See section 19.9.

Of course, an airplane has some plain old mass in addition to its gyroscopic properties. In order to lift this ordinary mass you need to use the flippers. Therefore, the tail-raising maneuver requires both flippers and rudder --- flippers to change the pitch of the ordinary mass, and rudder to change the pitch of the gyroscope.

8.7 Canted Engine

Often the engine is mounted in such a way that direction of the thrust vector is a little to one side of the axis of the airplane. This is done in order to compensate for various nonidealities such as helical propwash. It contributes to the yaw-wise torque budget in the obvious way.

8.8 Rudder Usage During Rolls

Turning the airplane properly requires coordinated use of ailerons and rudder. Getting it exactly right is a bit tricky.

Remember that in an airplane, the direction you are moving is not necessarily the same as the direction you are pointing. There are several crucial things that happen during a turn:

1)
You use the wings to change the direction you are going, i.e. to re-orient your momentum vector. I call this the *MV-turn*.

2a)
You use the rudder to change your heading (i.e. to overcome yaw-wise inertia, i.e. to provide yaw-wise acceleration).

2b)
You use the rudder to overcome steady adverse yaw due to twisted lift, as discussed below.

2c)
You use the rudder to overcome transitory adverse yaw due to differential drag.

Item 1 is relatively straightforward: you put the airplane into a bank. The horizontal component of lift will change the direction of motion. Note that MV is a bit of a pun; it might stand for “momentum vector” or “mass times velocity” ... or both.

Item 2a is important because if the airplane didn’t have any vertical tail, banking would cause it to just *slip* off in the new direction without changing its heading. It is much nicer to yaw the plane to align its axis with the new direction of motion, so you apply the rudder, thereby creating a yaw rate that matches the MV-turn rate.[7]

Now we come to item 2b. We must consider adverse yaw. As discussed in section 8.8.4, during a steady roll, the aerodynamic forces produced by the two wings are equal in magnitude, but one force vector is twisted slightly forward while the other one is twisted slightly rearward. This causes a yawing moment in exactly the wrong direction: if you are rolling to the right it tries to make the airplane yaw to the left. To compensate you must deflect the rudder whenever the ailerons are deflected.

Finally, we come to item 2c. Suppose you are flying an airplane where there is a lot of mass out on the wings. Whenever you are starting or ending a roll maneuver, you need to accelerate one wing upward and the other wing downward. As discussed in section 8.8.3, this briefly requires extra lift on one wing and reduced lift on the other wing. This unequal lift produces unequal induced drag. This drag causes additional adverse yaw.

For any given rate of roll, you need to use lots more rudder at low airspeeds, for reasons discussed in section 8.8.6.

Procedures for maintaining coordination during turns are summarized in section 8.8.7; the intervening sections describe in a little more detail what is the problem we are trying to solve.

8.8.1 Analysis of a Roll

To make the discussion more concrete, let’s consider a roll starting from straight-and-level flight and rolling to the right. As we can see from figure 8.5, there are multiple timescales in the problem.

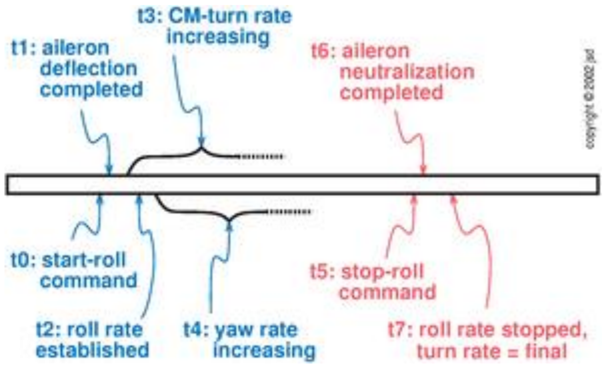

Figure 8.5: Timeline for Roll Maneuver

- [t0, t1] the time it takes you to move the ailerons;
- [t1, t2] the time it takes for the roll rate to reach the value commanded by this aileron deflection;
- [t2, t3] the time it takes the yaw rate to reach the value corresponding to the rate of MV-turn;
- [t3, t4] the time you hold the ailerons deflected.

Let's analyze what happens if you move the ailerons fairly abruptly. Although generally I recommend flying with a smooth, gentle touch, (1) there will be times when you want to roll the airplane on short notice, so let's learn how to do it; and (2) the abrupt case makes it easier to understand what is going on.[8]

In some airplanes, such as a Piper Cub, the roll rate will reach its final very quickly (within a small fraction of a second), because the airplane has very little roll-wise inertia. Practically all the mass (pilot, passenger, fuel, and engine) is arranged in a straight line right on top of the roll axis, so they don't contribute much moment of inertia. In other airplanes, such as a Cessna 310, the roll rate responds much more slowly, because lots of mass (engines and tip tanks) is situated far from the roll axis.

Before the roll rate is established (i.e. during the time [t1, t2]) the plane will experience transitory adverse yaw due to differential induced drag. The nose will swing a little toward the outside of the turn. The effect is usually rather small, since

1. these differential drag forces are typically small (during slow flight) or very, very small (during cruise), compared to the differential lift forces that cause the roll, and
2. these forces must act against the yaw-wise inertia, which is at least as large as the roll-wise inertia.

The rest of the discussion applies no matter how slowly or abruptly you moved the ailerons.

After the time t2, a steady roll rate exists. Even though the ailerons are deflected, there is no difference in lift from one wing to the other, for reasons discussed in section 8.8.4. Since there is no difference in lift, there will be no difference in induced drag, hence no transitory adverse yaw.

However, one wingtip is diving, so its force vector is twisted slightly forward. The other wingtip is rising, so its force vector is twisted slightly rearward. Even though each force has practically the same *magnitude* as it would in non-rolling flight, the twist means there is a slight component of force in just the right direction to produce a steady adverse yawing moment.

In addition, because the airplane is rolling, a bank is developing. This bank causes an MV-turn; that is, the airplane is changing its direction of motion. In order to keep it *pointing* in the same direction it is *moving*, you need to deflect the rudder during the roll, as discussed in section 8.8.5.

At time t6, the ailerons are neutralized, but the rolling motion has not yet stopped. (Again, there is a delay due to roll-wise inertia.) At this point there are several things going on:

1. There is a difference in lift between the two wings, as needed to damp out the roll. This creates a negative amount of transitory adverse yaw. This requires a left-rudder contribution to compensate.
2. However, the airplane is still rolling, and a still-increasing rate of yaw is needed to coordinate with the still-increasing rate of MV-turn. This requires a right-rudder contribution.
3. Similarly, because the airplane is still rolling, the twisted lift requires a right rudder contribution.

In practical situations, the first item (transitory adverse yaw) is usually smaller than the other two. During the interval [t5, t7] the roll rate is decreasing, so you need less and less rudder deflection.

Analogous statements would apply if you started from a left turn and used right aileron and right rudder to roll out of the turn. Similarly, it is easy to do a similar analysis for rolling into a left turn and/or rolling out of a right turn.

8.8.2 Designers' Tricks

Imagine an airplane without a vertical fin. It would behave be more like a Frisbee than a boat --- if you gave it a yaw rate, inertia would make it just keep on yawing until some torque acted to stop it. Even if it were not yawing, there would be no reason to expect the yaw angle (i.e. heading) to be anywhere close to the desired value.

In a real airplane, of course, the vertical fin and rudder supply the forces required to keep the yaw angle and yaw rate under control. An overview of how you use the rudder during turns can be found in section 8.8.

Aircraft manufacturers know about how turns are affected by twisted lift and yaw-wise inertia. They generally try to provide the needed yaw-wise torque automatically, using various tricks. One trick is to interconnect the rudder and ailerons with a spring. That means you automatically get a certain amount of rudder deflection in proportion to the aileron deflection. They choose the proportionality factor so that you can more or less fly "with your feet on the floor" at cruise airspeeds. Of course, vastly more rudder is needed at lower airspeeds; fortunately you can easily overpower the interconnect spring by pushing on the controls in the obvious way.

Here's another trick, which you may have noticed on many airplanes: when one aileron goes down a little, the other one goes up a lot. (This is called *differential aileron deflection*.) The designers were trying to arrange for the upward-deflected aileron to generate a lot of parasite drag. If they do it just right, the drag force is just enough to overcome twisted lift and yaw-wise inertia during a steady roll. The so-called *Frise aileron* uses a similar trick. It has lip on the bottom, well ahead of the hinge. The lip sticks down into the airstream when the main part of the aileron is deflected up. Again, the purpose of the lip is to generate drag on the wing with the upward-deflected aileron.

In addition to overcoming yaw-wise inertia (during a steady roll), the designers also want to overcome transitory adverse yaw (when ailerons have been deflected but the roll hasn't yet started). Fortunately, transitory adverse yaw is rather small, and by adjusting the amount of differential deflection, and the amount of the Frise effect, pretty good cancellation can be achieved.

The bad news is that this compensation only works at one airspeed. The designers arrange it so you can fly with your feet on the floor during cruise. This is a mixed blessing, because it can lull you into complacency. At lower airspeeds, where it is most important, you still need to use lots of rudder to keep things coordinated. Don't forget!

8.8.3 Transitory Adverse Yaw

Suppose you wish to roll into a right turn. You will deflect the ailerons to the right, as shown in figure 8.6. During the brief time after the ailerons are deflected and before the steady roll is established, this will increase the lift created by the left wing, and decrease the lift created by the right wing. Unfortunately, there is no way to produce lift without producing drag, so the left wing will be dragged backwards while the right wing lunges forward. This is the exact opposite of what we wanted; the airplane yaws to the left even though we wanted it to turn to the right. Being a good pilot, you have anticipated this, so you apply right rudder as well as right aileron, to make sure the nose swings the right way.

Figure 8.6: Transitory Adverse Yaw

Even if you don't get the footwork exactly right, the nose will eventually swing around and point more-or-less the right way, because of the airplane's inherent yaw stability (as discussed in section 8.2).

Once a steady roll rate is established (no roll-wise acceleration), the two wings are producing the same amount of lift, so this type of adverse yaw will no longer exist.[9]

Now let's consider what happens if you wish to roll out of the turn. The airplane is banked to the right and already turning to the right. You will deflect the ailerons to the left. This will cause extra drag on the right wing, and reduced drag on the left wing. The airplane will yaw to the right, continuing and exaggerating the turn that you were trying to stop. Anticipating this, you apply left rudder along with the left aileron

8.8.4 Steady Adverse Yaw -- Twisted Lift

Now let's consider what happens during a steady turn. As illustrated in figure 8.7, the airplane as a whole is moving forward, but the left wingtip is moving forward and *up* while the right wingtip is moving forward and *down* (because of the rolling motion).

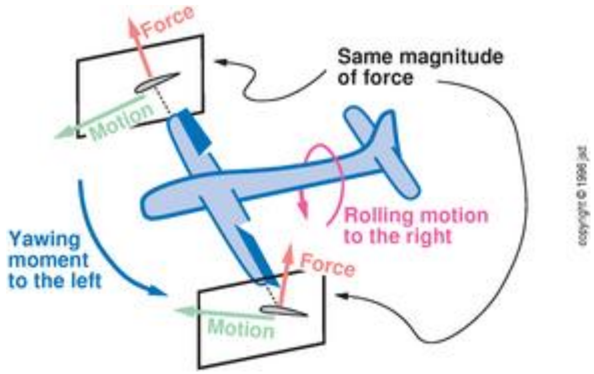

Figure 8.7: Steady Adverse Yaw -- Twisted Lift

Let's see what the *local* angle of attack is at the wingtip. We use the trusty formula

angle of attack + angle of climb = pitch + incidence (8.1)

In the figure, the right wingtip has a negative angle of climb, since it is going forward and down. But the deflected aileron gives it a lower incidence, effectively twisting that section of airfoil nose-down. By the same token, the left wingtip has a positive angle of climb (due to the rolling motion) and an increased incidence (due to the aileron).

In a steady roll, the incidences just cancel the climb angles, so that the left wing and the right wing end up flying at the same angle of attack. If they didn't cancel, you wouldn't have a *steady* roll.

The cancellation means there is no roll-wise torque, but the yaw direction is a different story. As you can see in the figure, the force vector for the downgoing wing is twisted forward, while the force vector for the upgoing wing is twisted rearward. This pair of fore-and-aft force components creates a yaw-wise torque. You need to deflect the rudder to compensate.

Some people try to argue that these force-components should be called "drag" forces since they are directed fore and aft, in the same direction as the overall relative wind. However, it is much better to think of them as components of the local lift, since the twisted lift remains perpendicular to the *local* relative wind. The strongest argument is this: a drag force should dissipate energy in proportion to force times airspeed, but it is clear that the twisted lift forces do not dissipate energy.[10]

8.8.5 Yaw-Wise Inertia

(In this section, we will assume that you are flying at such a low airspeed that the designers' tricks discussed in section 8.8.2 are not sufficient to produce automatically coordinated turns.)

Whenever the airplane is in a bank, it will make a MV-turn. A pure MV-turn, however, is not what you want. A pure MV-turn means that even though the airplane is *moving* in a new direction, the *heading* hasn't changed. The airplane has a nonzero slip angle. The uncoordinated airflow acting on the tail will eventually set up a yawing motion that matches the MV-turn rate, converting it from a pure MV-turn to a more-or-less[11] coordinated turn. If the yaw-wise damping is weak, as it usually is, the nose will slosh back and forth several times as it tries to catch up with the MV-turn.

At any particular MV-turn rate, once the yaw rate is established, no further yaw-wise torque is required. Like a toy top, once the airplane starts rotating in the yaw-wise directon it will be happy to continue rotating.[12] The only time you need a yaw-wise torque is when the yaw rate is changing.

So, we see that during a steady roll,

- The ailerons are deflected.
- The bank angle is increasing and, correspondingly, the rate of MV-turn is increasing.
- To match the MV-turn rate, the yaw rate must increase.
- To increase the yaw rate, the rudder should be deflected.

Conclusion: the rudder should be deflected when the ailerons are deflected.

8.8.6 Amount of Rudder Required

As we have seen, there are actually three different reasons why you need to apply the rudder during roll maneuvers: twisted lift, differential induced drag, and yaw-wise inertia. The amount of rudder deflection you need depends on the shape of your airplane, and also depends on airspeed.

* Twisted Lift

Example 1: Consider an airplane with long wings and with most of the mass concentrated near the middle of the airplane. A typical glider is an excellent example, but almost any ordinary-shaped airplane will do. In this case there will be very little roll-wise inertia, and accordingly very little transitory adverse yaw. There will also be rather little yaw-wise inertia. Therefore in such a plane, the dominant effect will be steady adverse yaw due to twisted lift.

Example 2: Suppose you are flying along in *any* airplane on a sunny summer day. You encounter a situation where your right wing is in an updraft, while your left wing is in a downdraft. You deflect the ailerons in order to maintain zero bank, zero roll rate, and constant heading. This combination of non-horizontal relative wind and deflected ailerons creates twisted lift, the same as shown in figure 8.7 (except that the roll rate is zero in this case). Therefore this is a perfect example of steady adverse yaw, and you must deflect the rudder to compensate. (This could not be explained by differential drag or yaw-wise inertia. This is pure twisted lift.)

The yawing moment due to twisted lift is essentially independent of airspeed. It just depends on the deflection-angle of the ailerons. Meanwhile, though, the force produced by the rudder is proportional to airspeed squared. Therefore you need lots more rudder deflection (per unit aileron deflection) when the airspeed is low.

* Differential Induced Drag

Example 3: Consider an aircraft where there is a lot of mass located far away from the roll axis. A twin with heavy engines mounted way out on the wings, plus tip-tanks full of fuel, is a good example. Such a plane will have lots of roll-wise inertia, and therefore lots of transitory adverse yaw. You will still have to worry about yaw-wise inertia and twisted lift, but *in addition* to those effects you will need to apply *extra* rudder deflection when ailerons are first deflected, before the steady roll develops.

The amount of rudder required depends dramatically on airspeed. In addition to the rudder-force issue discussed above, the amount of transitory yawing moment itself increases when the airspeed decreases. The key to understanding this is to realize that whereas the coefficient of lift is more or less proportional to the angle of attack (for moderate angles of attack), the coefficient of induced drag is more or less proportional to the *square* of the angle of attack.

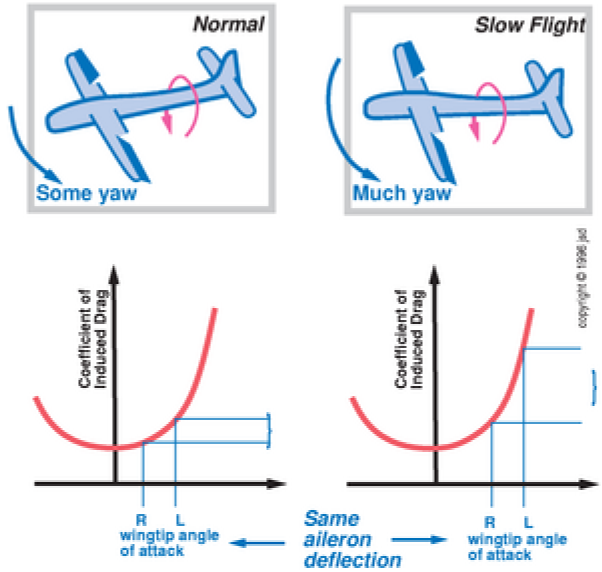

Figure 8.8: Slow Flight Means More Transitory Adverse Yaw

The left side of figure 8.8 shows the same situation as in figure 8.6, along with the coefficient of drag curve. On this curve I have indicated the different angles of attack for the two wingtips, and the correspondingly different amounts of drag. We see that the coefficient of drag curve is relatively flat on the bottom, so at relatively small angles of attack (high airspeeds), a difference in angle of attack doesn't cause too much difference in drag.

In contrast, the right side of figure 8.8 shows the same aircraft in *slow flight*. Both wings are operating at a higher angle of attack. Because the coefficient of drag curve is steeper in this regime, the same difference in angle of attack (i.e. the same aileron deflection) creates more difference in drag (i.e. more transitory adverse yaw).

* Yaw-Wise Inertia

Example 4: Consider a long, thin, single-engine biplane carrying lots of cargo. Since it has a rather short wingspan, there will be rather little twisted lift, i.e. rather little steady adverse yaw. Similarly, since all the mass is close to the roll axis, there will be very little roll-wise inertia, i.e. very little transitory adverse yaw. There will, however, be lots of yaw-wise inertia.

Example 5: Let's return to the case where your right wing is in an updraft, while your left wing is in a downdraft. This time, however, you don't deflect the ailerons; you just accept the resulting roll rate. During the steady roll, you will need to deflect the rudder to supply the yaw-wise angular momentum to match the ever-increasing MV-turn rate. (This rudder requirement could not be explained by twisted lift or differential drag. This is pure yaw-wise inertia. Also note that no designers' tricks could maintain coordination in this situation, since the ailerons are not deflected.)

Once again, the amount of rudder required increases markedly at low airspeeds. There are three main contributions; the first two essentially cancel each other:

- The roll rate depends on the deflection-angle of the ailerons, times airspeed. That means at low airspeeds the roll rate is less, which reduces the amount of rudder required.
- The amount of turn that results from a given bank angle increases at low airspeeds, as discussed in section 16.5. This increases the amount of rudder required.

- Finally, as always, rudder effectiveness depends on airspeed squared, increasing the amount of rudder deflection required at low airspeeds.

8.8.7 Summary: Coordinated Turning Procedures

A proper turn consist of two ingredients: a MV-turn and a heading change. In an idealized “basic” airplane, you would use the ailerons to bank the airplane and lift the MV around the corner, and you would use the rudder to change the heading and combat adverse yaw. In a typical modern airplane *at cruise airspeeds,* deflecting the ailerons alone creates a fair approximation of the proper torques in both directions (roll-wise and yaw-wise). In all airplanes *at low airspeeds,* proper rudder usage is vitally important.

The basic rule is simple:

- if you are rolling to the right, you must apply right rudder;
- if you are rolling to the left, you must apply left rudder.

The *amount* of rudder will depend inversely on the airspeed.

Another version of the rule substitutes the word “aileron” for “roll”:

- right aileron requires right rudder;
- left aileron requires left rudder;

In a steady roll, the two versions are more or less equivalent; at the beginning and end of a roll (when the roll rate does not match the aileron deflection) the truth lies somewhere in between. Split the difference.

These rules you to anticipate the need for rudder deflection. As discussed in section 11.5, you have many ways of knowing when you’ve got it right:

1. The acid test involves looking out the window. You should perceive that the rate of heading change is proportional to the amount of bank.
2. You can also look to the side and perceive that the wings flap straight up and down, not slicing fore and aft as you roll.
3. You can see that the inclinometer ball remains almost centered.
4. Yet more information comes from the seat of your pants.

By the way: If you think about it for a moment, you can see that in inverted flight (negative angle of attack) you will have a *negative* amount of adverse yaw --- if you deflect the yoke to the left you will need to push on the right rudder pedal, and vice versa --- just the opposite of what you would do in noninverted flight. When you are actually in the plane, hanging upside down, this is not as confusing as it seems on paper. A little thought and a little practice will make it fairly self-evident which wing you should lower to make a MV-turn and which rudder pedal you should push to change the heading.

As mentioned at the beginning of this chapter, there are lots of reasons why you should use the rudder properly during turns. Alas, the learning process is complicated by the fact that in many cases the airplane will “cover up” small mistakes for you. In particular, whenever the airplane is in a slip, the vertical fin will automatically try to return the plane to zero slip angle. This is the yaw-wise stability discussed in section 8.2. The plane will (under most conditions) eventually establish an approximately correct rate of heading change anyway. The goal of correct rudder usage is to establish the correct yaw-wise motion *without* a slip developing even temporarily.

The dependence of adverse yaw on airspeed can lead to trouble. Pilots spend almost all of their time buzzing around at cruise airspeeds, where ignoring the rudder is OK or nearly so. Sometimes this leads to complacency. The problem arises on approach and/or departure, where airspeeds are much lower. Proper coordination becomes more challenging, exactly at the place where it is most important (since the margins for error are also smaller). If you mishandle the ailerons at low speed and low altitude, you could well cause a spin or a snap roll, with no chance for recovery.

Section 11.5 describes a few useful tricks for perceiving exactly how much rudder is needed to achieve perfect coordination.

8.9 Long-Tail Slip

Now let's see what happens while the airplane is in an established turn. In particular, let's consider an airplane with a fairly long fuselage, flying in a fairly tight turn. As shown in figure 8.9, there is no way that the airflow can be lined up with the front part of the fuselage and the back part of the fuselage at the same time. The fuselage is straight, and the path through the air is curved. You can't have a straight line be tangent to a circle at two different points. You have to choose.

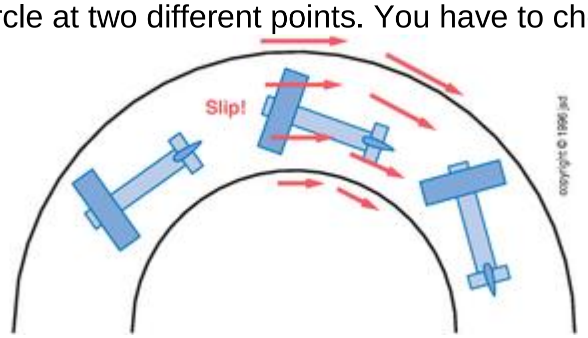

Figure 8.9: Airplane in a Tight Turn --- Rudder Neutral

If left to its own devices, the airplane will choose to have the vertical fin and rudder lined up with the airflow. The fin/rudder combination is, after all, an airfoil. Airfoils are good at producing tremendous forces if the wind hits them at an angle of attack. Besides, the tail is way back there where it has a lot of leverage.

Because of the air hitting the sides of the fuselage, and other effects, the fin/rudder might not completely determine the slip angle, but it will be the main determining factor. For sure, the airflow at the front of the fuselage --- and over the wing --- will have a significant slip component.

This will occur whenever the airplane is in a turn (unless you explicitly deflect the rudder to compensate). I call this the *long-tail slip* effect. This slip sounds like a bad thing, but in fact it can be put to good use; without it there would be no roll-wise stability, for reasons discussed in section 9.3. Remember: an inadvertent turn will be a slipping turn.

You can see from the geometry of the situation that the amount of long-tail slip is proportional to the length of the airplane and inversely proportional to the turning radius. The latter depends on the *square* of the airspeed, as well as the bank angle.

In a stubby, fast aircraft like a V-tailed Bonanza in a 15 degree bank at 165 knots, the long-tail slip effect will be small fraction of a degree --- hardly noticeable. On the other hand, in a long, slow glider, maneuvering to stay in a thermal using a 45 degree bank at 50 knots, the effect will be fifty or a hundred times greater! You will need several degrees of rudder deflection. You may need to push the rudder pedal all the way to the floor just to keep the air flowing straight over the wings. (Even if you decide to accept a little slip over the wings in order to reduce the crossflow over the fuselage and stabilizer, you will still want inside rudder, and lots of it.)

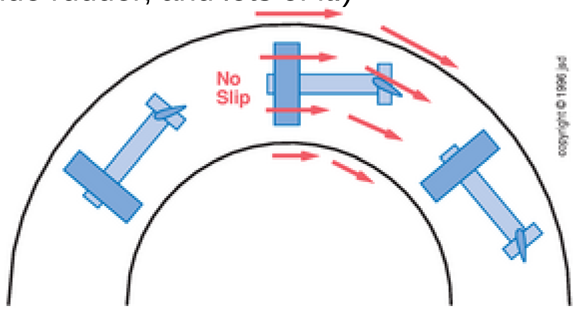

Figure 8.10: Airplane in a Tight Turn --- Rudder Deflected

I emphasize that even though you are holding inside rudder (bottom rudder) during the turn, this is definitely not a skidding turn (unless you get carried away and use *too much* inside rudder). This rudder usage is completely unrelated to the uncoordinated "boat turn" discussed in section 8.10.

We would like the airflow to be aligned perpendicular to the wings and parallel to the fuselage everywhere, but in a tight turn this is not possible. We have to compromise and "split the difference". The lowest-drag arrangement is to have at the nose a slight crossflow from inside the turn, and at the tail a slight crossflow from outside the turn.

The best way to check the alignment is with a *slip string* --- a piece of yarn exposed to the airflow where the pilot can see it. Non-experts commonly call this a yaw string, but this is a misnomer. In fact the string measures slip angle, not yaw angle. This is discussed in more detail in section 17.1.3.

If (as is usually the case) you don't have a slip string, you can try to infer the alignment by looking at the inclinometer ball. Remember, however, that the inclinometer ball and the slip string actually measure quite different things. The distinction is noticeable whenever the rudder is deflected, particularly in a twin with an inoperative engine, as discussed in section 17.1.4.

8.10 Boat Turn

My friend Larry has a sailboat. It doesn't have ailerons. You steer it with the rudder.[13] This changes the direction the boat is pointing. As shown in figure 8.11, this causes the water to flow crosswise past the hull, creating a sideways force that eventually changes the direction the boat is going.

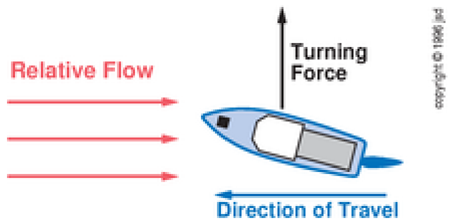

Figure 8.11: Boat Making a Boat Turn

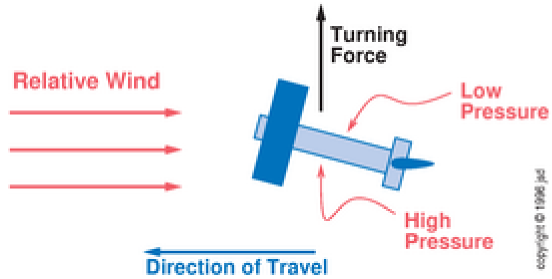

Figure 8.12: Airplane Making a Boat Turn

All the same words can be applied to an airplane. Keeping the wings level, you press the right rudder pedal. This causes the airplane to yaw to starboard. As shown in figure 8.12, air will then hit the fuselage on the port side, creating a sideways force[14] that will gradually shove the airplane around in a right-hand turn. (There will also be a lot of drag, but that is not our concern at the moment.) The force of the wind on the rudder (needed to yaw the plane) is smaller than, and in the opposite direction to, the resulting force of the wind on the fuselage.

In powered flight, the horizontal component of thrust will make an additional contribution to the boat turn. Remember, a turn results from a net force that is not aligned with the way you are going. This includes engine thrust, whenever there is a nonzero slip angle.

To reiterate: the airplane will turn to the right if you hold the right rudder pedal down --- even if the wings are not banked. Of course, turning the airplane properly (using the wings) is ten times more effective and more efficient than a boat turn

8.11 Weathervaning During Taxi

When the airplane is on the ground, it feels the force of the ground *and* the force of the wind.

Since the tail is far, far behind the wheels, a crosswind will create a yaw-wise torque. It will tend to blow the tail downwind, forcing the nose to turn upwind, just like a weathervane.

Now, suppose you are moving (as opposed to parked). The weathervaning tendency causes the nose to turn into the wind. The wheels are still on the ground, making lots of friction, so the airplane will roll in the direction determined by the wheels, i.e. the direction it is heading. Therefore the airplane will travel toward the *upwind* side of the runway. This may seem ironic or even paradoxical, but it's true --- the crosswind causes the airplane to move upwind.[15] You have to deflect the rudder to downwind to compensate.

8.12 Asymmetric Thrust

In a multi-engine airplane, if the engine on one side has failed, or for any reason is developing less thrust than its counterpart on the other side, this will produce a torque (possibly a very large torque) in the yaw-wise direction. This is discussed in section 17.1.4.

8.13 Yaw-Wise Torque Budget --- Summary

We have finally come to the end of this section, having covered the most important causes and effects of yaw-wise torques and motions. There are quite a number of such processes:

- The helical propwash effect is important, especially in high-power / low-airspeed situations.

- Gyroscopic precession means that deflecting the flippers will cause a yawing motion (and deflecting the rudder will cause a pitching motion).
- Adverse yaw means that deflecting the ailerons will cause a yawing moment.
- The long-tail slip effect means that an inadvertent turn will be a slipping turn. This effect is very significant in gliders. It is much less noticeable in typical powered aircraft, but it has important implications for roll stability, as discussed in section 9.3.
- P-factor exists in principle but is usually insignificant.
- Actual motion in the yaw-wise direction will create a yawing moment that tends to damp the motion.
- Yawing the airplane changes the direction it is *pointing* which does not automatically change the direction it is *going*; the “boat turn” effect exists but is feeble and inefficient.
- The pilot can deflect the rudder to oppose the unwanted yawing effects and create the desired ones.

Some of these ideas will be revisited when we discuss “Dutch roll” in section 10.6.1.

Perceiving coordination and maintaining coordinated flight is important. Further discussion of this topic appears in chapter 11, along with a discussion of how and why to perform intentional slips.

1
For a more-precise definition of what we mean by yaw, see figure 19.8 in section 19.6.1. For a definition of terms such as *yaw angle*, *heading*, and *slip angle*, please refer to section 19.6.3. The terminology and general principles of forces and moments are discussed in section 19.7. Roll-wise and pitch-wise motion are discussed in chapter 9 and chapter 6.

2
... but not the only technique, as discussed in section 8.10.

3
All the examples in this section assume a typical American engine that rotates clockwise as seen from behind.

4
The figure exaggerates the curvature of the stream lines.

5
Interestingly, it goes to zero again when the axis is perpendicular to the direction of flight, as in a helicopter.

6
but if you pay attention you can notice it in many other situations

7
Yaw-wise acceleration (which may be a somewhat unfamiliar subject) is discussed in more detail in section 8.8.1.

8
This analysis ignores the overbanking tendency and various other small effects.

9
Although in a steady turn you may need some rudder deflection because of the long-tail slip effect, as discussed in section 8.9, and in a steady roll you will need some rudder deflection because of twisted lift and roll-wise inertia, as discussed in the following sections.

10
The real drag vector gets twisted, too, but the consequences are too small to worry about.

11
It won’t be exactly coordinated because of the long-tail slip effect, as discussed in section 8.9.

12

But you will generally need some rudder deflection to compensate for the long-tail slip effect.

13

Boat lovers' note: there are some ocean liners that do use roll-control devices rather like ailerons, although they are primarily for passengers' comfort, not for steering. Also, to be sure, there are some boats that can be steered by banking them. On my sailboard, for instance, you have to bank it the wrong way (i.e. to the outside of the turn) by shifting your weight. On some light racing yachts you can steer them pretty well just by shifting the weight of the crew around. Many speedboats bank into the turns. But we're getting off the subject. The point is that Larry's boat (like lots of others) leans to leeward whether you are turning left, turning right, or going straight. The reason it doesn't tilt any more than it does is because it has tons of lead in the keel. You can't bank it by shifting your weight, and it wouldn't turn much if you did. You steer it with the rudder.

14

This force is classified as a lift force, since it is perpendicular to the relative wind --- even though it is produced by the fuselage (not the wings), and even though it is horizontal. See the official definitions in chapter 4.

15

In those rare cases where there is inadequate friction on the wheels (such as a seaplane, or an airplane taxiing on a slick icy surface) it is quite possible for the wind to blow the airplane downwind. This of course has nothing to do with torque; it's just a plain force.

9 Roll-Wise Torque Budget

Many non-pilots think pilots must have super-human fast reflexes. But in fact, good pilots are known for their smoothness, not their quickness.

This chapter considers the various forces that could impart a rolling moment to the airplane.[1]

9.1 Dihedral

Back in section 8.2, we discussed how an uncoordinated relative wind[2] will affect the yaw-wise motion; now let's see how it will affect the roll-wise motion.

The first thing that people think of in this connection is *dihedral*. The word comes from the Greek word for "two planes" and just means that the two wings are not coplanar, as shown in figure 9.1.

Figure 9.1: Definition of Dihedral

In the presence of dihedral, any uncoordinated (side-to-side) airflow will hit the bottom of one wing and the top of the other wing, as shown in figure 9.2. This means one wing will be forced up and the other forced down. If you work out all the angles between the total relative wind and the wings, you find that indeed the angle of attack is increased on the upwind wing and reduced on the downwind wing. The difference in lift produces a rolling moment. Any process whereby uncoordinated airflow produces a rolling moment it is called a *slip-roll coupling*; dihedral is a good example of this.

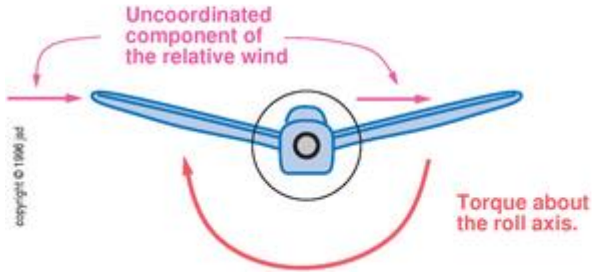

Figure 9.2: Dihedral --- Slip Produces Rolling Moment

The rolling moment (i.e. the roll-wise torque) will be proportional to the dihedral angle, and proportional to the amount of slip.

9.2 Other Forms of Slip-Roll Coupling

Dihedral is only one of several reasons why an airplane might have a slip-roll coupling. A high-wing airplane has a certain amount of slip-roll coupling because of *interference effects*. That is, when the airplane is in a slip, the fuselage interferes with the airflow over the wing. As shown in figure 9.3 and figure 9.4, the stream lines have to bend a little in order to flow around the fuselage. This creates an updraft at the root of the upwind wing, and a downdraft at the root of the downwind wing. This creates a rolling moment that tends to raise the upwind wing.

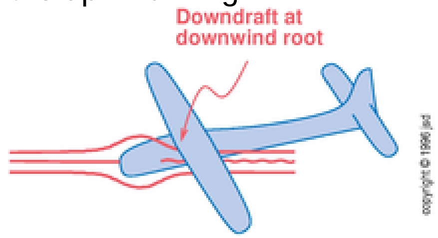

Figure 9.3: Redirection --- High-Wing Airplane in a Slip (1)

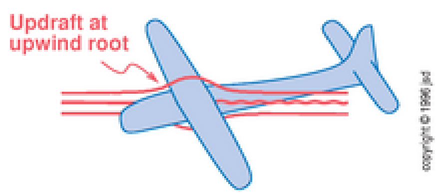

Figure 9.4: Redirection --- High-Wing Airplane in a Slip (2)

As shown in figure 9.5 and figure 9.6, on a low-wing aircraft the effect is reversed. There is an updraft at the root of the downwind wing, and a downdraft at the root of the upwind wing. This contributes a *negative* amount of slip-roll coupling.

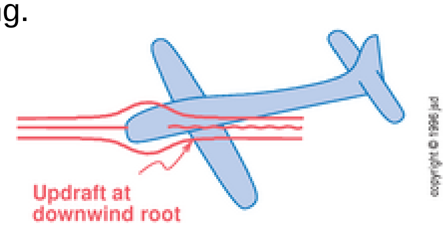

Figure 9.5: Redirection --- Low-Wing Airplane in a Slip (1)

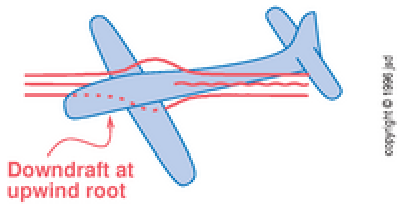

Figure 9.6: Redirection --- Low-Wing Airplane in a Slip (2)

A related interference effect is shown in figure 9.7. A fuselage moving sideways through the air is a very non-streamlined object. Downstream of such an object you expect to find a large, messy wake. The air in the wake is less capable of producing lift when it flows over the wing.

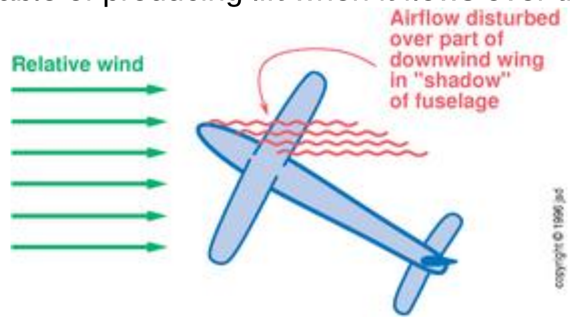

Figure 9.7: Turbulence --- High-Wing Airplane in a Slip

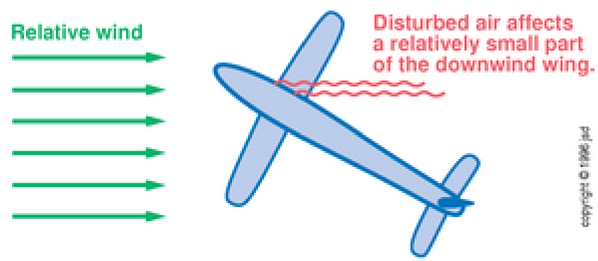

Figure 9.8: Turbulence --- Low-Wing Airplane in a Slip

Because the air that the wing really cares about is coming from ahead and *below*,[3] this type of interference is more pronounced in a high-wing airplane --- the fuselage is in a stronger position to disturb the relevant airflow. This is can be seen by comparing figure 9.7 with figure 9.8.

The magnitude of the effect of the wake is very difficult to predict. It will depend not only on the general shape of the fuselage, but also on the details of the surface finish.[4] It will also depend very nonlinearly on the airspeed and slip angle.

In general, interference effects mean that (to achieve an adequate amount of slip-roll coupling) low-wing airplanes typically need more dihedral than high-wing airplanes. You can check this by looking at typical airplanes at your local airport.

A third effect is illustrated in figure 9.9. If you put a swept-wing airplane into a slip, the more-forward wing produces more lift. That wing (the left wing in the figure) presents effectively more span to the airstream. It is a common mistake to think that the increased span explains the increased lift. The mistake is to overlook the fact that when it presents more span, it necessarily presents less chord. Lift, other things being equal, is proportional to wing area, and it is well known that area is not changed by a rotation. The correct explanation has more to do with the direction of airflow. Air flowing spanwise along an airfoil doesn't produce lift. The key idea is that the chordwise component of the airflow is bigger for the left wing.

Figure 9.9: Swept Wing Airplane in a Slip

A fourth effect is shown in figure 9.10. In practically all aircraft, the rudder sticks up above the roll axis. When the aircraft is in a slip, the rudder produces a substantial force. This force times this lever arm produces a roll-wise torque.

Anything else that sticks up above the roll axis and produces sideways drag or sideways lift contributes the same way. This includes the wings of a high-wing airplane, although the effect is small since spanwise flow along a wing doesn't create much force --- just a little bit of sideways drag. This is another reason why high-wing airplanes can get by with less dihedral (for the same amount of slip-roll coupling).

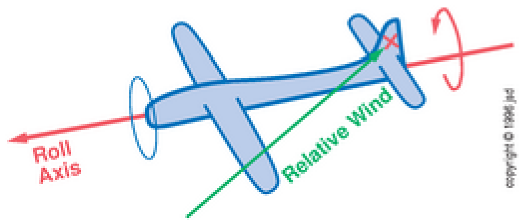

Figure 9.10: Tall Rudder in a Slip

All four effects just mentioned are in the same direction, and can be combined: You can have a high-wing, swept-wing airplane with lots of dihedral and a really high tail --- in which case you would probably have more slip-roll coupling than you need.[5]

The propwash contributes a negative amount of slip-roll coupling when the engine is producing power. If you yaw the nose to the right, the uncoordinated component of the wind will blow more of the propwash to the right wing. The extra lift on the right wing will roll you to the left.

Slip-roll coupling is the reason why you can make a relatively normal turn with the rudder (inelegant though it is). If you gently press on the right rudder, you will cause a skid that will eventually produce a bank to the right. Of course the skid itself will also cause a boat turn to the right. If you hold a constant rudder deflection, the boat-turn force will only be proportional to the rudder deflection, whereas the bank (and the associated non-boat turn) will keep getting larger and larger because of the slip-roll coupling.

9.3 Roll-Wise Stability

We are now all set to understand how the airplane responds if, for some reason, one wing goes a bit lower than the other.

The airplane will start to turn. If the turn were perfectly coordinated, the airplane would be happy to keep turning around and around and around. Fortunately, as we recall from our discussion of the long-tail slip effect (section 8.9), "an inadvertent turn will be a slipping turn". This tiny amount of slip, acting through the slip-roll coupling, will tend to roll the airplane back to wings-level straight-ahead flight.

This process gives the airplane a slight amount of *roll-wise stability*.

Airplane designers always make sure the airplane has a certain amount of slip-roll coupling, for exactly this reason.

The roll-wise stability is rather weak, because the two necessary ingredients are individually weak: The slip-roll coupling is usually moderately weak, and the long-tail slip effect is so weak that (except for glider pilots) most pilots never notice it unless it is pointed out.

Common experience indicates that roll-wise stability is indeed rather weak. If you are cruising along in turbulent air and take your hands off the controls for a couple of moments, you do not expect the nose to pitch up or down 30 degrees, and you do not expect it to yaw left or right 30 degrees, but you would not be at all surprised to have a 30 degree bank develop.

Even in the best of conditions, the stability generated by the long-tail slip with slip-roll coupling can only overcome a small amount of uncommanded bank. For larger bank angles, the overbanking tendency (section 9.4) takes over and creates roll-wise instability.

* Dihedral in the Absence of Slip

Before going on, let’s take another look at what happens in a coordinated turn. Sometimes it is argued that when the airplane is in a bank, the lowered wing has a bigger footprint (a bigger projection on the ground) than the raised wing, as shown in the left part of figure 9.11.

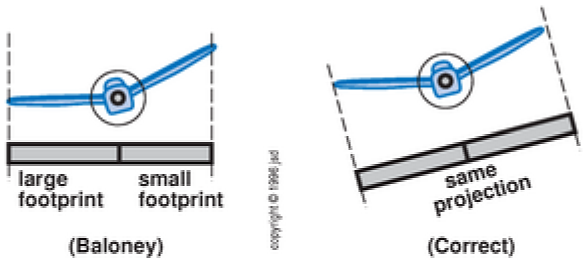

Figure 9.11: Dihedral Has No Effect in the Absence of Slip

A similar argument was used back in section 9.2 to explain why swept wings produce a slip-roll coupling. There is one slight difference: the swept-wing effect is real (because it involves the direction of the air) whereas the supposed effect of dihedral in a coordinated bank is completely imaginary. The wing doesn’t know or care where the ground is. It cares only where the air is coming from. In a coordinated turn, the air is coming from straight ahead, so dihedral has no effect.

Other myths about dihedral involve the angle of the lift vectors of the two wings. The correct answer is the same: in the absence of slip, dihedral has no effect. As long as the air is coming from straight ahead, the lift vectors are symmetrically disposed, as shown in figure 9.12.

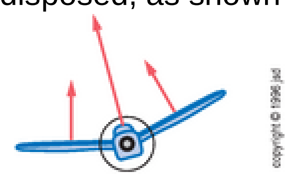

Figure 9.12: Symmetric Lift Vectors

In a coordinated turn, the aircraft is happy to continue turning forever; it will definitely not have any tendency to return to wings-level flight. Indeed, it will have the opposite tendency, called the overbanking tendency, which we now discuss.

9.4 Differential Wingtip Speed; Overbanking

Figure 9.13 shows the aircraft in a coordinated turn. The outside wingtip follows a path of length 2 π *R* (big *R*) while the inside wingtip has the proverbial "inside track" --- its path is only 2 π *r* (little *r*). Since the outside wingtip travels farther in the same amount of time, it must be moving faster.

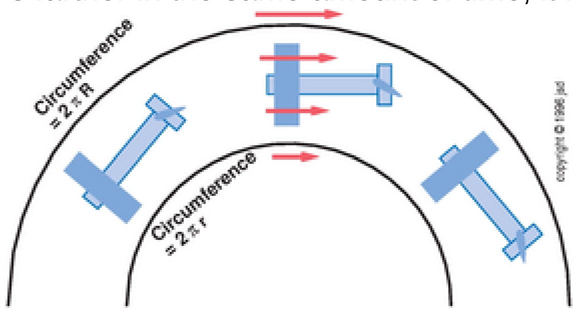

Figure 9.13: Overbanking Tendency

The same fact is depicted a second time in the figure --- the relative wind is depicted to be stronger on the outside wingtip. Since the lift generated by an airfoil depends on the square of the airspeed, the outside wing would produce more lift (other things being equal). This means that the aircraft in a turn (especially a properly coordinated turn) will tend to bank into the turn more and more. The tighter the turn, the more pronounced this *overbanking tendency* becomes. The next thing you know, you are in a spiral dive (as discussed in section 6.2).

In order to combat this tendency, you need to deflect the ailerons against the turn.

The strength of this effect depends on the ratio of the wingspan to the radius of turn. If you have stubby wings, high airspeed, and shallow bank angle, you'll never notice the effect. On the other hand, in a glider you might have long wings, low airspeeds and steep turns --- in which case you might need quite a bit of outside aileron deflection just to maintain a steady bank angle.

It is interesting to combine this with what we learned about long-tail slip effect (section 8.9) --- in the slow, steeply banked turn in the glider, you would be holding a substantial amount of inside rudder (to prevent the long-tail slip) and a substantial amount of outside aileron (to counteract the overbanking tendency). If you are not expecting this, it will appear very strange. You are holding completely crossed controls, yet the turn is perfectly coordinated. You can confirm this by observing that the slip string is perfectly centered.

You don't want to have to figure this out while sitting in the glider, trying to make a steep turn. Sometimes it pays to read the book before you go flying.

9.5 Rolling Moment due to Propeller Drag

The engine makes a contribution to the roll-wise torque budget.

As we remarked earlier, the propeller does not throw the air straight back; there is some rotational drag on the propeller blades. According to Newton's law of action and reaction, you can see that if the prop throws the air down on the right, it tends to make the airplane roll to the left.

To put it more crudely: take a model airplane (where the propeller rotates to the right) and hold it by the propeller. If you start the engine, the airplane will rotate to the left.

As shown in figure 9.14, some of the rotating air hits the top of the right wing and the bottom of the left wing.[6] This tends to reduce the amount of roll --- but it can never reduce it to zero or cause a roll to the right. Similarly, any air intercepted and “straightened out” by the tail reduces the rolling moment somewhat. Using Newton’s law again, we see that if any air escapes while still rotating down to the right, the airplane will roll to the left.

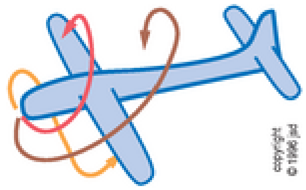

Figure 9.14: Rolling Moment Produced by Propeller

The only way to restore equilibrium is to take a corresponding amount of air and throw it down on the left. Airplane designers have long since learned about this propeller drag rolling moment, and they take steps to compensate for it. For instance, they set the left wing at a slightly higher angle of incidence than the right wing. This is called, unsurprisingly, *asymmetric incidence*. It is especially useful to apply this trick to the part of the wing that flies in the propwash, so that the effect increases as engine power increases. On a Piper Cherokee, the roll-wise trim is easily adjustable on the ground --- in the flap extension mechanism for each flap there is a turnbuckle that allows the flap to be raised or lowered until the roll-wise trim is just right.

If the roll-wise trim is just right in cruise, it will be nowhere near right during a soft-field takeoff. In that case, the propeller drag will be worse because of the high power, and the fancy rigging of the airfoils will be less effective because of the low airspeed. The result: you will have to deflect the yoke to the right, using the ailerons to counter the prop drag rolling moment.

9.6 Engine Inertia

Newton’s second law asserts that force equals mass times acceleration. There is a rotational version of this law, asserting that the rotational force (i.e. torque) equals the rotational inertia times the rotational acceleration. That means whenever the engine RPMs are increasing or decreasing, a torque is produced.

There is also a rotational version of Newton’s third law, asserting that if you impart a clockwise rotational momentum to one thing, you must impart a counter-clockwise rotational momentum to something else.

Consider an airplane which has the engine aligned in the usual way, but where the propeller-drag effects (discussed in section 9.5 are negligible. The easiest way to arrange this is to have a single engine driving two counter-rotating propellers. The Wright brothers used this trick in their first airplane.

While (and only while) the engine speed is changing, the airplane will tend to roll. It will roll to the left if the engine is speeding up, and it will roll to the right if the engine is slowing down.

In steady flight in this airplane, the the engine’s rotational inertia has no effect. The fact that the engine / dual propeller system is producing power does not imply that it is producing any net torque.

To clarify the distinction, compare the two situations shown in figure 9.15 and figure 9.16. We have an ordinary single-engine airplane. We have removed the propeller and bolted a huge brake drum onto the propeller shaft. In figure 9.15, the brake shoes are attached to the floor of the hangar. When we run the engine, the brake will produce a huge torque that will make the airplane want to roll to its left. This is completely analogous to the propeller drag effect discussed in section 9.5. There is an instrument called a *prony brake* that measures the torque-producing capability of an engine in precisely this way.

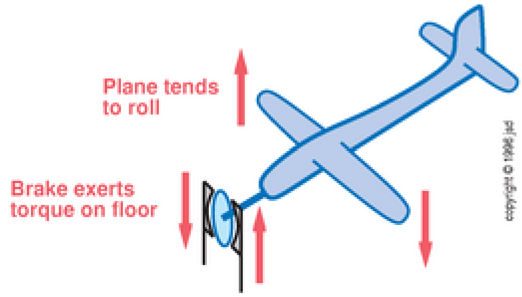

Figure 9.15: Prony Brake Attached to the Floor

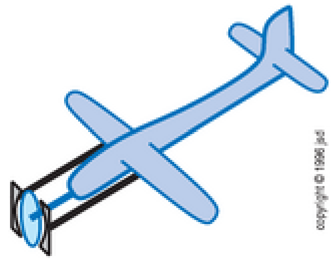

Figure 9.16: Prony Brake Attached to the Airplane Itself

In figure 9.16, the brake shoes are attached not to the floor, but to the airplane itself. Even if the engine is producing torque (straining against its engine mounts) all the torques flow in a closed circuit and cancel. The airplane as a whole exhibits *no* rolling tendency.

Newton's law is quite explicit about this: if you want to give the airplane some left-rolling momentum, you have to give *something else* some right-rolling momentum. This "something else" could be the air (as in figure 9.14) or perhaps the hangar floor (as in figure 9.15).

The angular motion of the internal engine parts can only affect the rolling moment if you *change* their rotational speed.

Engine rotational inertia should not be confused with propeller drag. In a direct-drive propeller installation, the propeller-drag torque does act on the fuselage via the engine mounts, but that is a coincidence, not a law of physics. In a gear-drive installation, most of the propeller-drag torque acts on the fuselage via the gearbox.

Finally, consider the case where the engine rotates one way and the propeller rotates the other way (which is easy to arrange using a gearbox). In a steady slow-flight situation, I guarantee you will need to deflect ailerons to compensate for propeller drag; engine inertia *per se* will have no effect on pilot technique.

9.7 Climbing and Descending Turns

In a level turn both wingtips are moving horizontally. In a climbing turn, both wingtips will be climbing, but they will not make equal angles to the horizon. This is because the climb angle depends on the ratio of the vertical speed to the forward speed. As a result of the different climb angles, we get different angles of attack for the two wingtips. The geometry of the situation is shown in figure 18.6 (in the chapter on spins). Another way to think about this is to recognize that it involves rotating in a non-horizontal plane, as discussed in section 19.6.4.

Let's do an example, as shown in table 9.1.

		Airspeed (K_{TAS})	Vertical speed (fpm)	Angle of climb	Angle of attack
Climbing turn	inside wingtip	99.46	500	2.844°	4.485°
	outside wingtip	100.54	500	2.814°	4.515°
	difference	2.2%			0.7%
Descending turn	inside wingtip	99.46	-500	-2.844°	4.515°
	outside wingtip	100.54	-500	-2.814°	4.485°
	difference	2.2%			-0.7%

Table 9.1: Climbing and Descending Turns

The example involves an airplane with a 35-foot wingspan turning at standard rate (3 degrees per second) at 100 K_{TAS} while climbing or descending at 500 fpm. We can calculate the resulting angle of attack at the wingtips.

We see that the change in angle of attack typically has less effect than the change in airspeed. In a climbing turn, the angle effect contributes to the overbanking tendency, while in an ordinary descending turn, it somewhat reduces it.

In a spin (which has a higher vertical speed, lower airspeed, and vastly higher rate of turn) the angle effect is extremely significant, as discussed in section 18.6.1. Far from reducing the rolling moment, increasing the angle of attack on the inside wing (which is stalled) only makes the situation worse.

9.8 Roll-Wise Torque Budget --- Summary

There are several effects that can give rise to a rolling moment. The most important ones are:

- A slip tends to cause a rolling moment, for several reasons, including: dihedral produces slip-roll coupling; the fuselage shadowing one wing (especially on a high-wing airplane) produces slip-roll coupling; swept wings produce slip-roll coupling, and a tall rudder that sticks up above the roll axis produces slip-roll coupling.
- This slip-roll coupling combines with the long-tail slip effect (discussed in section 8.9) to give the airplane a small amount of roll-wise stability.
- At medium or large bank angles the overbanking tendency creates roll-wise instability; the bank angle will tend to get larger and larger. This produces a spiral dive.
- There exists some medium-small bank angle where the two just-mentioned effects cancel. (That is, the overbanking tendency cancels the stability due to slip-roll coupling plus long-tail slip.) At this bank angle, the airplane will happily continue turning, at constant bank angle, without any help from the pilot.

- The airplane tends to roll left in high-power / low-airspeed situations, because of propeller drag.
- If you suddenly change the speed of rotation of the engine, the rest of the airplane will be subjected a brief rolling impulse. (Similarly, if you change the *orientation* of the plane of rotation, gyroscopic precession will cause yawing and/or pitching moments.)
- If the speed and direction of rotational motion is unchanging, engine torque will have no noticeable effects. (Engine torques will of course exist, but they will be part of "closed circuits" of torque within the fuselage, so they will not affect the handling of the airplane.)

Some of these ideas will be revisited when we discuss Dutch roll in section 10.6.1.

1
: For a discussion of the terminology and general principles of forces and moments, you can refer to section 19.7.

2
: i.e. an airflow pattern that is flowing left-to-right or right-to-left over the fuselage.

3
: See the discussion of upwash in section 3.1, including figure 3.2.

4
: See the discussion of dimples on golf balls in section 18.3 and in reference 10.

5
: Excessive slip-roll coupling will cause the airplane to suffer from Dutch roll, as discussed in section 10.6.1.

6
: The figure greatly exaggerates how tightly the flow pattern is wound.

10 Equilibrium, Stability, and Damping

Three of the most useless things in aviation are:

- The airspace above you.
- The fuel not on board.
- The runway not in front of the wheels.

Several parts of this book make use of the concepts of equilibrium, stability, and damping. This section defines the concepts a little more precisely and clarifies the relationships between them.

10.1 Equilibrium

The word *equilibrium* is quite ancient. The word has the same stem as the name of the constellation "Libra" --- the scale. The type of scale in question is the two-pan balance shown in figure 10.1, which has been in use for at least 7000 years. The compound word "equilibrium" translates literally as "equal balance" and means just that: everything in balance, i.e. no unbalanced forces.

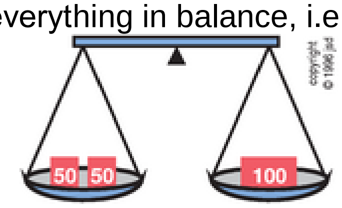

Figure 10.1: Equilibrium --- Forces in Balance

The *wheel* is more modern than the balance; its known use goes back “only” about 5500 years. It provides some more sophisticated illustrations of equilibrium and related concepts.

As indicated in figure 10.2, there are three ways to have the wheel be in equilibrium: [1] position the weight at the bottom, [2] remove the weight entirely (or put it at dead center, where the axle is) or [3] position the weight at the top.

If we attach the weight to any other point, system will be out of equilibrium. If we then let go, it will immediately start rotating.

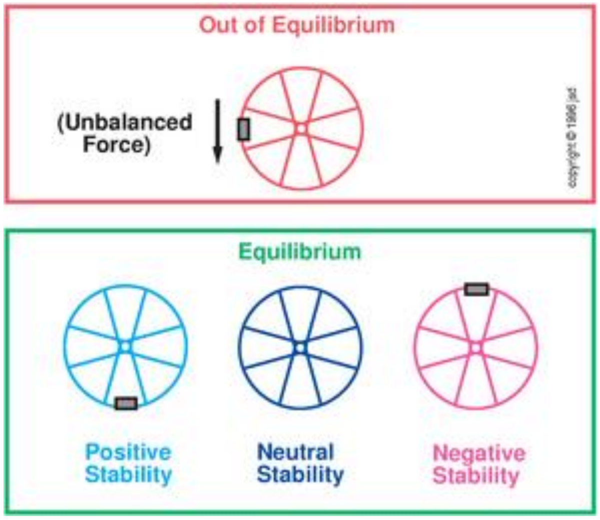

Figure 10.2: Equilibrium and Stability

10.2 Stability

Stability has to do with how the system responds if we move it a little ways from its equilibrium position. There are three possibilities:

- *Positive stability* means that if the system is displaced a little ways from its equilibrium position, it will generate a force tending to push it back towards equilibrium. The wheel with the weight positioned at the bottom is an example of positive stability.
- *Neutral stability* (also called *zero stability*) means that if the system was in equilibrium and you displace it slightly, it remains in equilibrium. No force is generated. The perfectly balanced wheel is an example of this.
- *Negative stability* means that if the system is displaced a little ways from its equilibrium position, it will generate a force that tends to push it farther from equilibrium. The wheel with the weight at the top is an example of negative stability.

It usually doesn’t make much sense to talk about stability except for systems that are in equilibrium or nearly so.

For a multi-dimensional system, we get to ask about the stability of each “mode”, i.e. each possible direction of motion. For example, consider an egg resting on a horizontal table. An ideal egg has zero stability against motion in one direction: it is free to roll around its axis of symmetry. On the other hand, it has positive stability against motion in the end-over-end direction; if you rock the egg slightly by pushing its nose down, it will tend to return to its original state.

10.3 Damping

A system exhibits *damping* if motion of the system produces a force that opposes the motion.

A bicycle wheel provides a good demonstration of a system with very little damping. Assuming the bearings are good and the wheel is not touching anything, when you spin the wheel it will keep going for more than a minute. Air friction produces very small forces that eventually cause the wheel to slow down.

A bicycle wheel that is rubbing against something is much more heavily damped. When it is in motion, rubbing friction can create large forces that oppose the motion and bring the motion to a stop.

A dynamical system can exhibit negative amounts of damping, but this is harder to demonstrate with a simple system. Negative damping tends to make the motion increase, which means that energy is being added to the system from somewhere; therefore simple friction can never produce negative damping.

Nose wheel shimmy of an airplane is a good example of what happens if a system has a negative amount of damping. If the aircraft is moving along the ground at high speed, the nosewheel will eventually hit a pebble or something. The nosewheel is then no longer aligned with the direction of travel. By the usual “castering” principle, this causes a force that tends to return the wheel to its proper position (that is, the wheel exhibits positive stability). Unfortunately, in many cases there is too much stability, and too much inertia in the castering mechanism. The result is that the wheel tends to overshoot its equilibrium position and continue to the other side, going out of alignment in the opposite direction by an even greater amount. The result is an oscillation that quickly grows to large amplitude.

Note the relationship of stability and damping: when the wheel is being forced back toward alignment, the force is toward the equilibrium position (positive stability) but is in the same direction as the motion (negative damping).

To eliminate the shimmy problem, a hydraulic “shimmy damper” is installed on the nose wheel. Figure 10.3 is cutaway drawing showing how a hydraulic damper works. It consists of an oil-filled cylinder, plus a pushrod attached to a disk inside the cylinder. When the pushrod moves from side to side, oil is forced to flow through the small holes in the disk. This creates a force proportional to the velocity of motion --- i.e. damping.

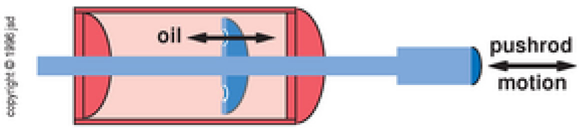

Figure 10.3: Hydraulic Damper

Sometimes the fluid leaks out of the damper, and even more commonly the linkages connecting the damper to the wheel become worn and loose. This makes the damper ineffective, whereupon the you get a vivid demonstration of negative damping. A preflight check of the damper and linkages is easy and worthwhile.

Also... as discussed in chapter 5, the airplane’s rolling motion and pure vertical motion are normally very heavily damped, but this damping goes to zero and becomes negative at the stall.

10.4 Relationship of Stability and Damping

To reiterate: stability refers to a force that arises depending on the *position* of the system; damping refers to a force that arises depending on the *velocity*.

In old-fashioned terminology, what we call "stability" was sometimes referred to as "static stability", and what we call "damping" was sometimes referred to as "dynamic stability". What's worse, occasionally *both* terms were shortened to the single word, "stability", which was unnecessarily confusing.

Also, modern usage prefers "damping" not "dampening" --- if you start talking about a "dampener" people will think you want to moisten the system.

Stability can be positive, zero, or negative; damping can also be positive, zero, or negative. A dynamical system can display any combination of these two properties --- nine possibilities in all, as shown in figure 10.4. In the top row, the bicycle wheel is dipped in molasses, which provides damping. In the middle row, there is no damping. In the bottom row, you can imagine there is some hypothetical "anti-molasses" that provides negative damping.

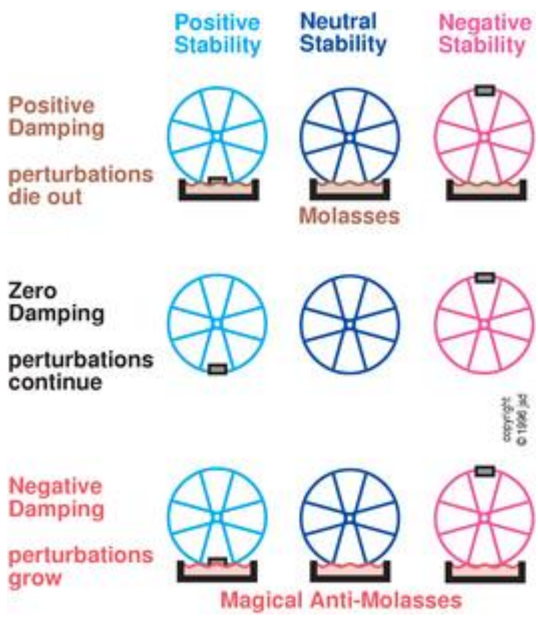

Figure 10.4: Stability and Damping --- Possible Combinations

10.5 Oleo-Pneumatic Struts

A great example of a device that provides a force that depends on position *and* a force that depends on velocity is the *oleo-pneumatic strut*, which is widely used on landing gear as a combination spring and shock absorber. It consists of a piston in a cylinder filled with both oil ("oleo") and air ("pneuma"), as shown in figure 10.5. If the piston is moved up into the cylinder, the air at the top of the cylinder is compressed. (The hydraulic oil is essentially incompressible.) This "air spring" creates a force that depends on the position. As the piston moves, the oil in the hollow part of the piston is forced to flow through the holes in the disk, creating a force that depends on the speed of motion, using the same principle as the damper discussed previously.

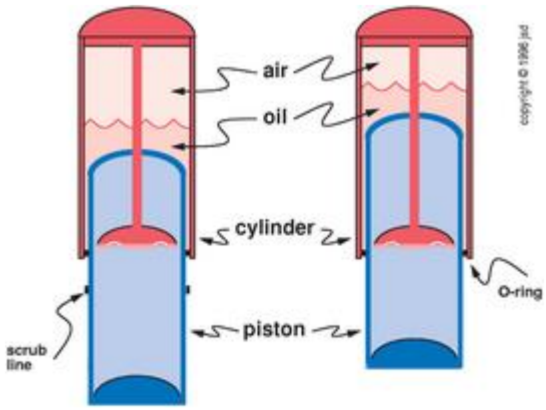

Figure 10.5: Oleo-Pneumatic Strut

It is important that the strut contain the right amount of air *and* the right amount of oil. Problems can arise more easily than you might think.

Suppose that over time, some of the oil leaks out of the strut on your airplane.[1] Your friend, Murgatroyd Fudpucker, borrows the plane and notices during preflight that one of the struts is low --- that is, not enough of the piston is protruding from the cylinder. Murgatroyd gets out a bicycle pump and adds air to the strut. The strut now sits at the correct height. During future preflight checks, a passive glance the strut will give you the impression that things are OK ... but they are not really OK.

The problem is that *oil* has been replaced with *air*. Since air is a thousand times more compressible than hydraulic oil, the amount of force it takes to make the strut "bottom out" has been greatly reduced. If you or Murgatroyd makes even a slightly hard landing, the piston will smash against the end of the cylinder, metal to metal. This has roughly the same effect on the airframe as hitting it with a sledgehammer. Repairs could be very, very expensive.

Therefore, if there is any chance that the airplane has been mis-serviced since the last time you flew it, you should check not only the height of the struts, but also their springiness. To check a main-gear strut, lift up the wing a few inches and then let it drop. Similarly, to check the nose strut, lift up the nose (perhaps by pushing down on the tail) a little ways and then let it drop. If any strut compresses more than it should (e.g. if it comes anywhere close to bottoming out), do not fly the airplane until the strut has been properly serviced with air *and oil*.

There is a thin coating of oil on exposed part of the piston, which collects dust. When the piston is shoved into the cylinder, the O-ring will scrub the dirt down the piston and cause it to collect in a ring called the *scrub line*. Observing the scrub line can tell you how close the strut has come to bottoming out recently.

Please do not get the impression from the foregoing discussion that "air is bad" and "oil is good". I discovered an airplane recently where nose strut contained no air at all, but contained several inches too much oil instead. Once again, a passive, non-skeptical preflight check would not have caught the problem, because the struts were sitting at the normal height. Fortunately, I checked the springiness. There was no springiness, since trying to compress a solid column of hydraulic oil is about like trying to compress cast iron.

To reiterate: you should make sure that the struts contain the right amount of air and the right amount of oil. Servicing a strut isn't very tricky; it just has to be done right.

10.6 Oscillations

Whenever a system has positive stability but not enough damping, you can expect to see oscillations.

10.6.1 Analysis of Dutch Roll

As remarked in section 9.3, the airplane has only a small amount of stability in the roll-wise direction. You may be wondering why designers don't fix this problem by increasing the slip-roll coupling. The answer is that they are worried about Dutch roll.

Dutch roll is a messy combination of rolling, slipping, and yawing.[2] As we shall see, this combined motion is less damped than the pure rolling, slipping, or yawing motions would be.

A moderate amount of Dutch roll never killed anybody, but it does tend to provoke nausea, especially in passengers.

The Dutch-roll oscillations typically have such a short period (a couple of seconds) that it is a challenge for the pilot to overcome them by working the controls. A spiral dive, on the other hand, develops much more slowly. Therefore if it comes down to a compromise between roll-wise stability and Dutch-roll damping, designers generally increase the damping at the expense of the stability.

To understand where Dutch roll comes from, and how to fight it, gives us an opportunity to combine and apply most of the things we have learned about equilibrium, stability, and damping.

The rolling and yawing motions associated with Dutch roll are shown in figure 10.6; we will discuss the slipping component in a moment.

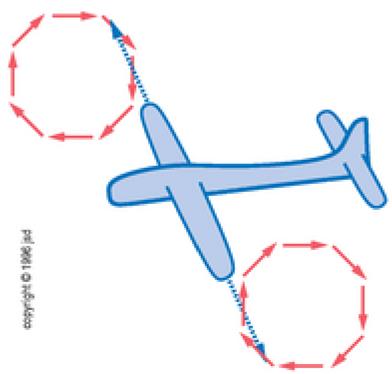

Figure 10.6: Dutch Roll

The wingtip yaws forward, then rolls up, then yaws backward, then rolls downward, then repeats. The opposite wingtip does the same thing, 180 degrees out of phase. Imagine pedaling a bicycle backwards.

To analyze the damping of the Dutch roll system, we must remember that energy is force times distance; by the same token power (energy flow) is force times velocity. The component of the force in the direction of the velocity is the only thing that matters; the component in the perpendicular direction doesn't count.

We begin by using figure 10.7 to analyze the forces that affect the rolling motion. The velocity and position of the wingtip is shown in red; net changes in the lift vector are shown in blue.

At point *A* in the figure, the wing is going upward. That means it has less angle of attack than normal (and in particular, less angle of attack than the opposite wingtip). The reduced lift corresponds to a net force opposite to the velocity, and therefore energy is being removed from the system. At point *C*, a similar analysis applies. The wingtip is descending, creating more angle of attack and more lift than normal. This corresponds to a net force which is once again opposite to the velocity, removing energy from the system. This is the same roll damping mechanism as discussed in section 5.4.

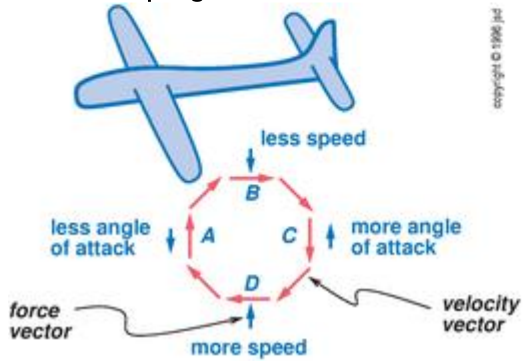

Figure 10.7: Dutch Roll --- Roll Forces

At point *B*, the wingtip has less velocity than normal, and less lift, while at point *D* the wingtip has more velocity and produces more lift. There is no effect on the damping, because the forces are perpendicular to the velocity.

We continue by using figure 10.8 to analyze the forces that affect the yawing motion. At point *B* in the figure, the vertical fin/rudder is wagging to the right. This changes the rudder angle of attack, opposing the motion. This is the same yaw damping mechanism discussed in section 8.3. Also, at this point, the port wingtip has less drag than the other, because it is moving backwards. Both of these effects take energy out of the system, providing damping. The same processes produce damping at point *D* also.

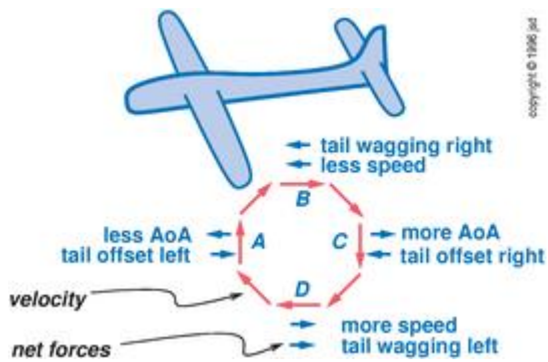

Figure 10.8: Dutch Roll --- Yaw Forces

At point *A*, there is a little less induced drag on the port wingtip because it is flying at reduced angle of attack. This has no effect on the damping, because the force is perpendicular to the velocity.

Also at point *A*, there is a yawing force because the airplane's heading is not aligned with its direction of travel; the tail is too far to the left. This provides yaw-wise stability but does nothing for the yaw damping, because the force is perpendicular to the velocity.

The analysis of point *C* is analogous to point *A*.

If the yawing and rolling motions were the whole story, Dutch roll would be no problem. According to the analysis so far, there is lots of positive damping. The Dutch roll would quickly die out.

Unfortunately, nature is not so kind, as we discover when we take the sideways motion of the aircraft into account. Refer to figure 10.9.

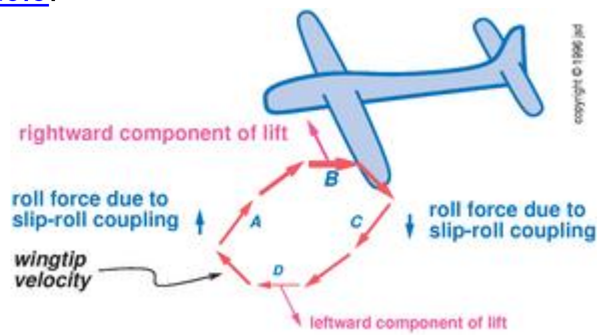

Figure 10.9: Dutch Roll --- Slip Causes Problems

At point *B* in the figure, the left wingtip is at the highest point in the cycle. The airplane is banked to the right. The wings' lift vector is inclined to the right, so there is a rightward component of lift. In fact, during the whole half-cycle from point *A* to point *C* there is at least some rightward force. Since the airplane has lots of inertia and not much damping[3] with respect to pure sideways motion, the rightward velocity just increases and increases during the whole half-cycle. The maximum rightward velocity is achieved near point *C*.

During the next half-cycle (from *C* via *D* to *A*) the airplane is banked to the left. The leftward force reduces the previously-acquired rightward velocity to zero, and then builds up a leftward velocity. The sideways velocity is zero at point *D*, and the maximum leftward velocity is achieved near point *A*.

Note that like any other lightly-damped oscillator (such as a pendulum, for instance a playground swing set) the maximum rightward *force* occurs when the plane is at is maximum leftward *position*.

The final ingredient is the slip-roll coupling.[4] A certain amount of slip-roll coupling is highly desirable because it is a necessary part of the process that produces roll-wise stability (section 9.3).

The bad news is that the slip-roll coupling contributes a negative amount of damping to the Dutch roll mode. The rightward velocity is maximum at point *C*, producing a leftward-rolling moment. The force is in the same direction as the roll velocity, so it adds energy to the Dutch roll.

Analogously, the leftward velocity is maximal at point *A*, producing a rightward-rolling moment. This, too, is in the same direction as the roll velocity, contributing negative damping.

So slip-roll coupling presents designers with a dilemma: it increases roll-wise stability, but decreases (Dutch) roll damping.

The simplest way a designer can resolve this dilemma is to notice that roll-wise stability depends on both slip-roll coupling and the long-tail slip effect. Therefore if you have a problem with Dutch roll, decrease the slip-roll coupling and increase the long-tail slip effect, for instance by making the tail boom longer and reducing the rudder area. As a rule of thumb, you can tell just by looking at a short-coupled airplane that it will have a problem with underdamped Dutch roll.

The other (all too common) design choice is to sacrifice stability. Most airplanes wind up with very, very little roll-wise stability. Consequently spiral dives are a constant threat.

10.6.2 How to Fight Oscillations

Since this book is intended for pilots, not designers, we should discuss how the pilot should use the controls in order to oppose obnoxious oscillations.

First, bit of simple advice: in an airplane that is susceptible to Dutch roll, be extra careful to avoid uncoordinated usage of ailerons and rudder since that would unnecessarily put energy into the Dutch roll mode.

Once Dutch roll gets started (due to turbulence, or klutzy control-usage, or whatever), it may be hard to stop. In some airplanes you may be able to improve the situation as follows: If the rudder pedals are moving because of the sideways force that the Dutch roll puts on the rudder, then you should rest your feet firmly on the pedals to prevent them from moving. This will increase the stability and (more importantly) the *damping* in the yaw-wise direction.

If that doesn't suffice, you can try to fight the oscillations by direct intervention. This requires some skill and lots of attention.

You should *not* think about correcting the *position* of the wing. If you deflect the ailerons to the right at point *D*, the wings will return to level (point *A*) sooner, but you will be applying a force in the same general direction as the velocity, increasing the velocity and the energy of the Dutch roll mode.

As we have seen, the airplane has plenty of stability and not enough damping, so what we need is a force that depends on the velocity, not the position. Therefore the ailerons need to be deflected to the left when the left wing has its maximum upward velocity, near point *A*. You should apply the deflection before point *A* and remove it after point *A*. Similarly, you should apply right aileron (smoothly) a little before point *C* and neutralize them (gradually) after point *C*.

A similar analysis applies to rudder usage. Don't try to correct the position. Instead, you need to apply right rudder at the point where the nose is swinging to the left with the maximum velocity (point *B*); by the same token you need to apply left rudder when the nose is swinging to the right with the maximum velocity (point *D*).

The same logic applies to phugoid oscillation (section 6.1.12), and to pilot-induced pitch oscillation associated with a botched landing. That is: when the nose is high, you should not push on the yoke to correct the nose-position; you should anticipate that the position will very soon over-correct all by itself. So, if the nose is high and dropping (or about to drop), you need a judicious pull on the yoke to prevent the pitch attitude from overshooting.

The general principle for stopping an oscillation is that your actions should increase the *damping*. (In contrast, if you try to increase the stability, e.g. by pushing when the nose is high and pulling when the nose is low, you will just make the oscillations faster, and probably bigger.)

As a consequence, you should react to the velocity, not the position. If the nose is moving with a high velocity to the left, apply right rudder. If the nose is rising rapidly, push on the yoke.

Act to increase the damping, not the stability.

Speaking of oscillations in general:

- Almost every airplane on earth has a lightly-damped phugoid mode. This is relatively easy to deal with, because the oscillations are reasonably slow. You can just look out the window, notice the pitch excursion, and deal with it.
- In contrast, a lightly-damped Dutch roll mode is relatively rare. Such a mode is relatively obnoxious, because the timescales can be comparable to human reaction times.
- There can be all sorts of pilot-induced oscillations. This includes pitch oscillations as well as the minute-by-minute heading oscillations associated with overcorrecting for navigational errors.
- Et cetera.

With a little thought, you can see that all these oscillations have important features in common.

1
: On a retractable-gear airplane, you can lose all the oil, even the oil inside the hollow piston, more easily than on a fixed-gear airplane.

2
: The constant-heading slip exercise discussed in section 16.7 is sometimes mistakenly called Dutch roll, but it's not the same

3
: In a system with lots of damping and not much inertia, like a spoon in molasses, the velocity tends to be proportional to the applied force. In the other extreme (lots of inertia, little damping) we can apply Newton's second law without worrying about frictional forces --- therefore the acceleration is proportional to the force and the velocity accumulates as long as the force is applied.

4
: That is, a slip produces a rolling moment --- by means of e.g. dihedral, sweepback, tall rudder, and/or shadow effects, as discussed in section 9.2.