17 Multi-Engine Flying

Q: In an underpowered twin, what is the role of the second engine?

A: It doubles your chance of engine failure, and it will fly you to the scene of the accident.

In normal conditions, operating a twin is not very different from operating a fast, heavy, high-powered, complex single. These issues are discussed in section 16.11. Normal multi-engine takeoff procedures are discussed in section 13.6. This chapter will be devoted to the issues that are unique to multi-engine aircraft, namely what to do if one engine quits.

17.1 Engine Out Scenarios

This section discusses some of the things you might observe when an engine fails, and what you can do in response.

First, we must deal with an important basic question: How serious is the loss of an engine? Answer: *it depends*.

- There are some situations where it is hard to notice and hardly worth noticing. On a scale of 1 to 10, this is level 1 or 2.
- There are some situations where it is extremely serious, and you must *immediately* respond in *exactly* the right way if you want to survive. On a scale of 1 to 10, this is level 10.
- Indeed there are some situations where the loss of an engine would be beyond serious: The situation would be unrecoverable, no matter how you respond. On a scale of 1 to 10, this level 11.

You must maintain proficiency so you can deal with level-10 situations. You must use good judgement so you stay away from any possibility of a level-11 situation.

17.1.1 Takeoff

If an engine fails during the takeoff roll, you have a decision to make: In some cases you must close the throttles and try to stop on the remaining runway, whereas in other cases you must try to fly away using the remaining engine. You won't have a lot of time to think about it, so you want to be 100% prepared to do the right thing instantly. This is discussed in section 17.2.1 and section 17.2.2.

17.1.2 Climb

For our next scenario, suppose you are at a reasonable altitude, at a reasonable airspeed, climbing with full power on both engines. Then one engine fails. Among other things, you will notice that the single-engine rate of climb is not half of the two-engine rate of climb. No, indeed! The reason is simple: as shown in figure 17.1, when an engine is shut down, you are not splitting the difference between two-engine performance and level flight; you are splitting the difference between two-engine performance and a zero-power descent.

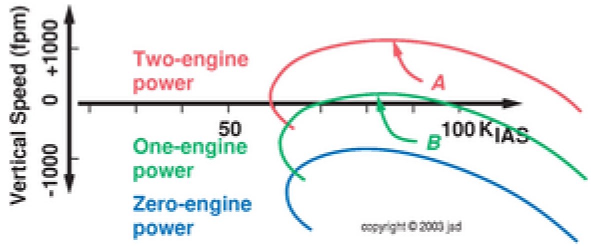

Figure 17.1: Two-, One-, and Zero-Engine Performance

The power curves in the figure are roughly representative of a Piper Apache, a well-known light twin trainer. Point *A* corresponds to the two-engine best rate of climb, 1150 fpm at 86 knots. Point *B* corresponds to the single-engine best rate of climb, 160 fpm at 82 knots. (These numbers apply to a fully-loaded aircraft at sea level in the clean configuration.) We see that the single-engine rate of climb is less than 15% of the two-engine rate of climb.

At density altitudes above 5000 feet the Apache cannot climb at all on one engine. Also, if the engines, propellers, and paint job are not quite factory-new, the performance will be even less than these book values suggest.

You must not allow yourself to think that just because airliners can climb with an engine out, your favorite light twin can climb with an engine out.

It is legal to operate a light twin with anemic or nonexistent single-engine climb performance. In such cases, engine failure at low altitude is perhaps the most critical situation that arises in general aviation with any appreciable frequency. Like a single-engine aircraft with partial power failure, you need to make a forced landing. The problem with the twin is that (because of the asymmetric thrust) if you mishandle the situation, your chance of getting into a spin is much higher than it would be in a single.

On the other hand, even if you are not climbing, you are probably not descending very fast. You can treat it as another “noisy glider” situation as mentioned in section 15.1. If you start out several thousand feet above the ground, you can probably travel dozens of miles while gradually descending. Look around and find a nice place for a forced landing.

17.1.3 Slip String

Generally, the best way to fly any airplane is to keep the airflow aligned with the fuselage. That is, we want zero slip angle, as defined in section 19.6.3. Alas, in a multi-engine airplane with asymmetric thrust, this can be rather tricky to perceive. The most direct way to get information about this angle is to use a *slip string*.

To create a slip string, tape a piece of yarn to the nose of the airplane, on the centerline, in front of the window where you can see it. Leave a foot or two of yarn dangling free, so it will align itself with the airflow.

The deflection of the string is proportional to the airplane’s slip angle, that is, the angle of the *X*-axis relative to the free-stream relative wind. There is not equality, merely proportionality, because the fuselage disturbs the local airflow pattern. The sideways component of the flow is increased more than the fore-to-aft component, so the sensitivity of the string is increased. That is, the angle of the string is larger than the actual slip angle.

Even in ordinary single-engine airplane, slip strings work surprisingly well. The straight-back component of propwash decreases the sensitivity somewhat, and the helical component of propwash biases the string slightly to one side.

The slip string is commonly referred to as a "yaw string", even though it measures the slip angle, not the yaw angle (i.e. heading) or yaw rate. The slip angle measures the angle between the fuselage and the relative wind, whereas yaw is defined relative to some fixed spatial direction. Heading and heading change (i.e. yaw rate) are easy to perceive by looking out the window, while it is not easy to perceive slip angle except by reference to a slip string. Heading can also be perceived using the directional gyro (heading indicator), and yaw rate can be perceived using the rate-of-turn gyro or by observing the rate of motion of the directional gyro.

The inclinometer ball is often referred to as the slip/skid instrument, but that is a misnomer. It measures inclination, not slip. As we shall see in section 17.1.5, it is quite possible for the airplane to be inclined but not slipping. To repeat: there is no good way to determine the slip angle without a slip string.

17.1.4 Coordination

For our next scenario, imagine that you are in level flight at cruise airspeed at a comfortable altitude. Let's also suppose that your airplane has a slip string installed. Then (surprise!) an engine fails. To simplify the discussion, let's suppose the right-hand engine[1] is the one that quits. You will immediately notice that the airplane will develop a slip angle. In this case, the airplane will yaw to the right, shown in figure 17.2. This is because one engine is producing lots of thrust, while the other is producing negative thrust, i.e. drag.[2] As always, two forces with a lever arm between them create a pure torque.

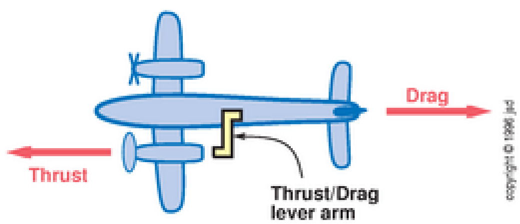

Figure 17.2: Engine Out --- Torque due to Asymmetric Thrust

This torque will produce an initial heading change. This will start out as a pure yaw; that is, it will change the direction the airplane is *pointing* without any immediate change in the direction the airplane is *going*. There will also be a tendency for the wing to drop on the side with the non-working engine: partly because of reduced propwash over the wing, and partly because if differential wingtip velocity due to the aforementioned yawing motion.

Then, after a short time (a second or so), the torques will come back into equilibrium, because of the airplane's natural yaw-wise stability (as discussed in section 8.2). That is, the uncoordinated airflow hitting the rudder will create a torque that opposes the asymmetric thrust. If you managed to keep the wings level, you will be in a boat turn to the right. The slip string will off-center to the right, indicating an asymmetric airflow, and indicating that you need to apply some[3] left rudder.

Figure 17.3: Engine Out --- No Pilot Action

At this point, you could either (i) sit there and be a spectator, as shown in figure 17.3, or (ii) press on the left rudder pedal to center the slip string, as shown in figure 17.4. From the point of view of directional control, your choice doesn't matter very much. That is, in case (i) the airflow strikes the whole airplane (including the rudder) at a nonzero angle, while in case (ii) the airflow strikes the airplane at a zero angle, with the rudder deflected. In either case, the amount of tail force produced is approximately the same. The most important difference is that the airplane will climb better in case (ii), because the airflow will be aligned with the fuselage.

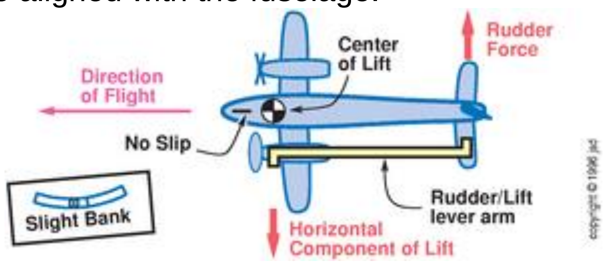

Figure 17.4: Engine Out --- Correct Compensation

Let's see what happens after the yaw-wise torques have returned to equilibrium. Let's assume for now that you are keeping the wings level. In case (i), the airplane will make a steady turn toward the dead engine. This is a boat turn, due to the uncoordinated airflow striking the fuselage, as discussed in section 8.10. This will be a genuine MV-turn (as defined in section 8.8), changing the direction of motion; the heading will follow the MV-turn in order to maintain a constant slip angle.

What is perhaps more surprising is that even in case (ii), if you keep the wings level the plane will make a MV-turn toward the dead engine (toward the right in this case). It will not turn as rapidly as in case (i), but it will turn nonetheless. The reason is that the rudder force, in addition to creating a torque, is creating an unbalanced force. This force is changing the direction of motion of the overall airplane. A possible (but non-optimal) way to stop this turn would be to apply even more pressure on the left rudder pedal, which would create a wings-level non-turning slip, as will be discussed below in conjunction with figure 17.5. For now, though, let's consider the correct strategy, which is to keep the slip string centered and apply a *bank* (to the left, in this case) to stop the turn. This is shown in figure 17.4. This uses a leftward component of lift[4] paired with the rightward rudder force. Once again we have a pair of forces with a lever arm between them, i.e. a pure torque. This lift/rudder pair cancels the thrust/drag pair discussed above.

To reiterate: when engine trouble develops, the first result of the asymmetric thrust is to make the airplane yaw toward the dead engine. The airplane changes its *heading* immediately (whereas only later does it gradually change the direction it is *going*). That is, a slip angle develops immediately. If you don't deflect the rudder, the slip angle will grow until the uncoordinated airflow striking the rudder develops enough torque to stop further yawing. This is the basic yaw-wise stability mechanism as discussed in section 8.2. The result is that the airplane does *not* spin around and around like a Frisbee --- it just develops a few degrees of slip angle and then stabilizes.

17.1.5 Perception and Initial Response

In a high-power low-airspeed situation, engine failure will be extremely noticeable. In other situations, with more airspeed and/or less power, engine failure may be harder to perceive than you might have guessed, especially if it is a gradual failure. Perceiving the initial yaw is particularly tricky during a turn --- the turn just proceeds a little faster or slower than normal. The subsequent boat turn may not be super-easy to perceive, either.

There are various ways to perceive and deal with the slip and yaw.

1. If you see a sudden yaw, or if you see a wing suddenly drop, apply opposite rudder immediately.
2. If you happen to have a slip string installed, the procedure is simple. If the string is deflected to one side, step on the rudder pedal on the opposite side, until the string is centered. The mnemonic is: “Step away from the string”.

Step away from the string.

3.

4. A less-elegant, less-accurate technique is to use the inclinometer ball. If the ball is deflected to one side, step on the rudder pedal on the same side. The mnemonic is: “Step on the ball”. Center the ball, then (to establish zero slip, as discussed below) relax the pedal force to let the ball go off center about one-third to one-half of its width.

Step on the ball.

5.

6. A commonly-used technique is to roll the wings level and then apply the rudder as needed to stop the boat turn. The advantage of this procedure is that it can be done without reference to instruments. The main disadvantage is that it doesn’t help you regain or retain control in a turn.[5] (There are rare situations where even though an engine has failed you might *want* to be turning.) Then, once you’ve got the wings level and the turn stopped, you should establish the optimal zero-slip condition, by raising the dead engine a few degrees and releasing some of the rudder pressure.[6]

At this point, you will find yourself maintaining a rudder deflection and a bank angle, both toward the side with the working engine. Use the rudder deflection (not the bank!) to identify which engine has failed. The mnemonic is: *working foot, working engine; dead foot, dead engine*. Specifically, if the right foot is *not* being used to deflect the rudder, bend your right knee. Raise that knee an inch or two, pat it a couple of times, and say “right engine has failed”. (More on this later.)

To maintain zero slip, you will need to bank the plane very slightly toward the working engine. The mnemonic is: *raise the dead*. This also implies that the inclinometer ball will be slightly displaced toward the working engine. This correct procedure (figure 17.4) requires slightly *more* aileron and slightly *less* rudder than you would need for wings-level, ball-centered, non-turning, uncoordinated flight (figure 17.5).

Having the slip string centered but the inclinometer ball not centered may seem a bit counterintuitive, so let’s examine the aerodynamics of the situation a little more closely.

The asymmetric thrust produces a yaw-wise torque, which cannot remain unopposed. The rudder is part of the solution, but remember that while the rudder is producing the desired torque it is also producing a force. We need the forces to be in balance, as well as the torques.

Suppose you try to maintain zero bank instead of raising the dead. Initially the airplane is not in equilibrium, because the rudder is producing an unbalanced force toward the dead-engine side. There are then two possibilities:

1. Suppose the sideways force remains unbalanced. This will cause the airplane to turn. This will be a wings-level, coordinated turn. I call this a pseudo boat turn. Unlike an ordinary boat turn, the airflow is coordinated along the fuselage, but unlike a regular turn, the horizontal force is not coming from the wings.
2. Suppose you push a little harder on the rudder pedal, establishing a slip toward the dead engine. The ball is centered, the wings are level, and the rate of turn is zero. (Any two of those things implies the third, regardless of engine status.) The forces are in balance because there is enough uncoordinated airflow over the fuselage to create a sideways force that balances the rudder force. The slip string is offset to the left, indicating that you are applying too much left rudder. In this situation, as shown in figure 17.5, you are using the fuselage as an airfoil. The problem is that the fuselage has a really poor lift-to-drag ratio. The sideways fuselage lift force is accompanied by a huge drag force, which steals energy from you. It would be much better to use the wings, as previously discussed in conjunction with figure 17.4.

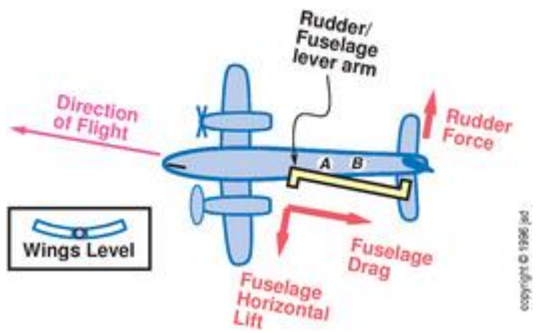

Figure 17.5: Engine Out --- Wings-Level Nonturning Slip

The proper technique is counterintuitive, because in any normal situation, proper coordination implies that the rate of turn will be proportional to the amount of bank (as in section 11.5.2) ... but an engine out, proper coordination requires a slight bank when you are not turning.

The amount of bank for a typical airplane can be estimated using the following argument: The lift-to-drag ratio of the airplane is roughly ten-to-one. In level flight the thrust must therefore be one tenth of the lift. The lever arm between the wings and the rudder is typically about three times the lever arm between the thrust and the drag. Since the torques must cancel, the rudder force (and the horizontal component of lift) must be one third of the thrust, and therefore one thirtieth of the lift. We conclude that the horizontal component of the lift is one thirtieth of the total lift. One thirtieth of a radian is two degrees --- not exactly a huge bank. (In a four-engine airplane with one of the outboard engines failed, the bank will be larger.)

To find out exactly how much bank you need to maintain coordinated airflow over the fuselage, it helps to use a slip string. At a safe altitude, set up for single-engine flight at V_{YSE}. Apply enough rudder pressure so that the slip string indicates zero slip angle. Bank as required to maintain nonturning flight. Experiment with slightly greater or lesser rudder pressure, to see what produces the best climb performance.

You will discover that in optimal single-engine flight, the inclinometer ball is *not* centered, but the slip string is centered. The airplane is inclined, but it has zero slip angle. Make a note of how much inclination is indicated by the inclinometer ball; typically it will be off-center by one-half or one-third of its diameter. You can use this information to set up a good approximation of engine-out coordinated flight during subsequent flights when you don't have a slip string installed.

The inclinometer ball measures the inclination of the wings relative to the E-down[7] direction. The inclinometer is sometimes referred to as a slip/skid ball, but that is a misnomer[8] because the slip string (as discussed in section 17.1.3) provides your only direct information about the slip angle.

Achieving zero slip is the key to optimal climb performance.[9] The idea is to have the airflow aligned with the fuselage. Centering the inclinometer ball is not what determines performance. Practice with the slip string until you learn how much inclination is required for a given amount of asymmetric thrust.

17.1.6 Yaw Control at Reduced Speeds

So far, we have discussed engine-out climb rate (section 17.1.2) and discussed the value of maintaining coordinated flight (section 17.1.4). We now begin a discussion of airspeed. As you might imagine, this is rather important.

In the previous section, we considered the case where you started out with plenty of airspeed. In the opposite extreme case, where you start out with a very low airspeed (below V_{MC}, as defined in section 17.1.7), you must *immediately* reduce power on the working engine and start diving. If this means making a power-off landing, so be it. If your speed is only slightly below V_{MC} you might be able to use partial power, but it won't be super-easy to figure out how much you can get away with. Dive until you achieve V_{MC}, then advance the throttle on the working engine and carry out the rest of the engine-out procedure as described below.

Now let's consider the intermediate case, which is a standard training exercise: At a moderately high airspeed, one engine is shut down. You then gradually reduce speed and see what happens. Again, to simplify the discussion, let's assume the right engine has failed.

The amount of asymmetric thrust does not depend on airspeed; it depends only on the power output of the engine. In contrast, the amount of force the rudder produces depends on the airspeed squared, and on the rudder's angle of attack. Therefore as you slow down you will need progressively more rudder deflection in order to maintain zero slip. If you do it properly, the sideways force developed by the rudder will remain unchanged, and the bank angle will remain unchanged (for now).

At some point you will run out of rudder deflection. The pedal (or the rudder itself) will hit the stops. You will be unable to maintain zero slip.

Now, suppose you continue to slow down beyond this point. As a slip develops, the airflow hits the tail and rudder at an angle. This gives the tail/rudder an angle of attack over and above whatever angle of attack you created by deflecting the rudder. You are using the slip angle as a substitute for additional rudder deflection. Up to a point, this higher angle of attack allows the tail/rudder to produce a higher coefficient of sideways lift, allowing it to produce the required force in spite of the lower airspeed.

In addition to the air hitting the rudder, you now have the uncoordinated airflow hitting the fuselage. You are relying on the rudder to produce at least 100% of the torque needed to oppose the asymmetric thrust. The air hitting the fuselage makes a small unhelpful contribution to the torque budget, and (more noticeably) contributes to the sideways force budget, producing an undesirable boat turn. This boat turn is in addition to the pseudo boat turn that the rudder is producing, so you will need to increase the bank angle to maintain nonturning flight.

Obviously there is a limit to this process. If you persist in increasing the rudder's angle of attack, at some point the rudder will stall. Remember, the amount of asymmetric thrust does not depend on airspeed, whereas the *absolute maximum* amount of force the rudder can produce depends on the airspeed squared. Therefore, for any nonzero amount of asymmetric thrust, there must be some airspeed below which the rudder cannot develop enough torque. At that point there will be an uncontrollable yaw toward the dead engine. The airplane will spin like a Frisbee.

You might think you could improve the situation by releasing the rudder pedal, thinking this would reduce the rudder angle of attack. Alas, it won't work. It will just cause the airplane to establish a greater slip angle. Remember the rudder needs to produce a certain amount of force to oppose the asymmetric thrust, and the airplane's natural yaw-wise stability will adjust the tail/rudder's angle of attack, trying to create the necessary force.[10]

If the rudder stalls, it will be about as unpleasant as anything you can imagine. There will be a sudden uncontrollable yawing motion. Because of the yawing motion, the wingtip on the side with the good engine will have a higher airspeed than the wingtip on the other side. Because of the difference in airspeed (plus the difference in propwash patterns) the good-side wing will produce much more lift, so you will get an uncontrollable roll. As the inside wing drops, it will probably stall (since you were already at a low airspeed). You are now in a spin. There is no guarantee that it will be possible to recover from such a spin; multi-engine airplane certification regulations do not require spin recoveries.

On some planes (such as an Apache, a common trainer) low-speed engine-out performance is limited by the rudder, as described above. On some other planes (such as a Seneca, another common trainer) you don't need to worry about the rudder because the wings will stall first.[11] This is not much of an improvement, because a stall with asymmetric power is also rather likely to result in a spin.

To prevent such nasty things from happening, you need to maintain a safe airspeed. The manufacturer gives you some guidance in this regard, as is discussed in the following section.

17.1.7 Minimum Control Speed --- Definitions

The symbol V_{MC} denotes "minimum control airspeed". There are at least four different definitions of this term, including:

I)

FAR 23 (the certification requirements for typical general-aviation[12] airplanes) gives a very specific definition of V_{MC}, namely:

FAR 23.149 Minimum control speed.

(a) VMC is the calibrated airspeed at which, when the critical engine is suddenly made inoperative, it is possible to maintain control of the airplane with that engine still inoperative, and thereafter maintain straight flight at the same speed with an angle of bank of not more than 5 degrees. The method used to simulate critical engine failure must represent the most critical mode of powerplant failure expected in service with respect to controllability.

(b) VMC for takeoff must not exceed 1.2 VS1, where VS1 is determined at the maximum takeoff weight. VMC must be determined with the most unfavorable weight and center of gravity position and with the airplane airborne and the ground effect negligible, for the takeoff configuration(s) with--

(1) Maximum available takeoff power initially on each engine;
(2) The airplane trimmed for takeoff;
(3) Flaps in the takeoff position(s);
(4) Landing gear retracted; and
(5) All propeller controls in the recommended takeoff position throughout.

[...]

II)

FAR 1 (the "definitions" section) defines V_{MC} as "minimum control speed with the critical engine[13] inoperative". It does not specify any restrictions as to weight, configuration, altitude, et cetera.

III)

The FAA Practical Test Standards for the multi-engine rating call for demonstrating "V_{MC}" in a particular way that emphasizes losing yaw control *without* stalling the wing or rudder, even though (as discussed below) for many airplanes V_{MC} (under definition I or II) is limited by wing stall and/or rudder stall.

IV)

In common parlance, pilots apply the term V_{MC} to the airspeed where the airplane (multi-engine or otherwise) becomes uncontrollable, no matter what the reason, no matter what the configuration, and no matter whether any engine is inoperative.

Note that none of these definitions require that the airplane exhibit a positive rate of climb at V_{MC}.[14] Also note that during a V_{MC} demonstration, the pilot is not required to optimize the climb rate or to maintain zero slip --- although zero slip may be an advantage if it can be achieved.

The V_{MC} number in the Pilot's Operating Handbook is determined according to the FAR 23.149 definition. This airspeed is marked with a red radial line on the airspeed indicator, and is sometimes called the *FAR 23.149 red-radial-line* airspeed.[15]

There are various ways to lose control; whichever happens first determines where the manufacturer sets the red-radial-line:

a)

In some airplanes, under some conditions, you can maintain control, even with an engine out, right down to the point where wing stalls. This is discussed below in conjunction with figure 17.6. A stall with asymmetric thrust could be rather sudden and nasty.

b)

In others (with a smaller rudder, larger wing, and/or higher-thrust engines), there will be conditions under which the rudder will stall before the wings do, as is discussed below in conjunction with figure 17.8. A rudder stall could be very sudden and very nasty.

c)

In yet others (larger rudder area but a shorter tail-boom, so that the rudder is closer to the wings), there will be situations where neither the wing nor the rudder is stalled, but the boat-turn forces are so large that it requires more than 5 degrees of bank to counteract them and maintain nonturning flight. The airplane would be perfectly controllable if the bank were not limited to 5 degrees. Since a bank of 15 or 20 degrees is not particularly dangerous, the 5 degree limitation must be considered arbitrary. If your airplane, at a given weight and altitude, does run up against this limitation, the resulting “loss of control” is neither sudden nor nasty. The airplane will just make a gentle boat turn toward the dead engine, as is discussed below in conjunction with figure 17.7.

Possibility (c) is in some ways attractive, but you have no guarantee that this is what will happen. Rudder stall depends on slip angle, so you may be wondering why FAR 23.149 should mention a *bank* angle as opposed to a *slip* angle. Bank does not cause slip.[16] If you want to establish any connection between bank and slip, you must consider:

1. bank angle (i.e. the angle between wings and horizon)
2. slip angle (as indicated by the slip string)
3. rate of turn
4. asymmetric thrust

If any three of these are zero, the fourth is guaranteed to be zero. More generally, other things being roughly equal, given any three of these you can estimate the fourth. The problem is that other things are generally *not* equal --- depending on weight, airspeed, airplane design, et cetera, five degrees of bank could correspond to a large slip angle or perhaps no slip angle at all. So this regulation is not 100% logical.

Some people seem to assign a near-religious significance to the “5 degree bank” mentioned in FAR 23.149. However, the real significance is quite limited:

- This regulation applies to the manufacturer during certain tests. It does not apply to you in your ordinary flying. If you have a real engine failure, you are limited only by the laws of aerodynamics.
- This regulation does not even apply to you during the checkride for a multi-engine rating. In particular, the FAA Practical Test Standard says you should bank for “best performance and controllability”. Alas, that’s inconsistent; best controllability requires a lot more bank than best performance, and the PTS doesn’t tell you how to make the tradeoff.
- Five degrees is no guarantee of optimum performance. The optimal bank could be five degrees, or more, or less (usually less).
- There is nothing in FAR 23.149 or anywhere else that guarantees the airplane is well behaved at the “official” 5-degree bank angle. The maximum bank you can safely use in nonturning engine-out flight could be five degrees, or more,[17] or less.[18] In particular, there is no guarantee that by limiting yourself to 5 degrees you will always get aerodynamic warning (in the form of a nice, gentle boat turn) before you get a nasty rudder stall or wing stall. If you want to demonstrate a gentle warning, you might need to limit yourself to much less than 5 degrees.

One thing we learn from this is that you should not use bank angle or anything else as a substitute for proper airspeed control.

For that matter, airspeed control requires a little thought, too. Perhaps because FAR 23.149 uses words like “most critical” and “most unfavorable”, people commonly assume that it is always possible to control the airplane at red-radial-line airspeed, no matter what. This assumption is wrong --- dangerously wrong --- in many airplanes. For example, there are some airplanes where the certified takeoff configuration[19] calls for the flaps to be extended, and the FAR 23.149 red-radial-line is essentially equal to the stalling speed in the takeoff configuration. Then if you operate with the flaps retracted, you will lose control of the airplane at an airspeed well above red-radial-line.[20]

Specific procedures for dealing with engine failure are discussed below, in section 17.2.

17.1.8 Effect of Altitude, Weight, etc.

FAR 23 tells us that the airplane, when operated under a particular set of circumstances, can maintain directional control at red-radial-line airspeed. The question is, what happens under other circumstances?

Let's discuss an example; call this example #1. It is a a non-turbocharged airplane for which the handbook calls for flaps retracted during takeoff. Then, under standard conditions (takeoff configuration, maximum weight, etc.), the situation is shown in figure 17.6. The single-engine stall speed for the example airplane is shown by a black vertical line in the middle of the figure. The FAR 23 red-radial-line is shown as a bright red tick mark on the airspeed axis. The manufacturer had to set it a knot or two above the stalling speed, since that is what limits the low-speed handling for this airplane in this configuration.

Also, in this figure, the magenta curve shows the airspeed below which the rudder cannot develop enough force to oppose the asymmetric thrust. Thirdly, the dotted cyan curve shows the airspeed below which the boat turn forces are so large that it would require more than 5 degrees of bank to maintain nonturning flight.

Since the example airplane is not turbocharged, as altitude increases there is less thrust available on the good engine. The required rudder force declines accordingly. This is why the magenta and cyan curves trend to the left as they go up. Note that in this configuration, for this airplane, rudder performance is not a limitation --- the wing stall is the only relevant limitation.

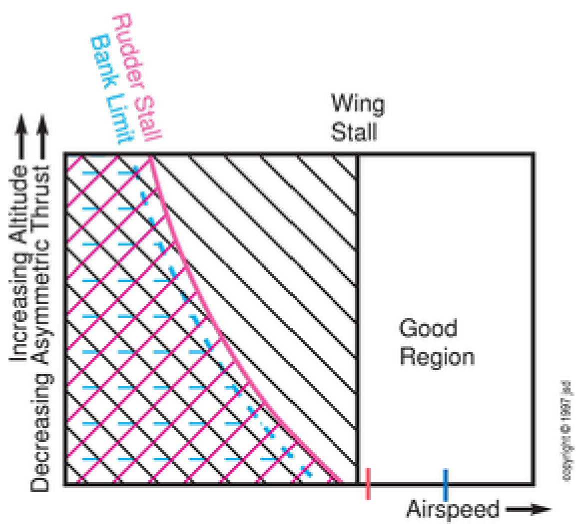

Figure 17.6: Speed & Altitude affect Directional Control (basic)

Now, suppose that several things change:

- You fly this airplane (#1) at reduced weight, about half of the legal maximum.
- You extend the flaps.[21]
- You limit yourself to 3 (not 5) degrees of bank.
- You limit yourself to less than full rudder deflection.

These new conditions have several consequences. For starters, the reduced weight will lower the stalling speed. Extending the flaps lowers the stalling speed some more. This is indicated by the black line, which moves to the left as we go from figure 17.6 to figure 17.7.

The amount of torque developed by the engine depends on altitude in the same way as before, and is unaffected by the weight, flaps, and other variations. The amount of force the rudder can produce is also unaffected. Therefore the magenta curve is the same in the two figures.

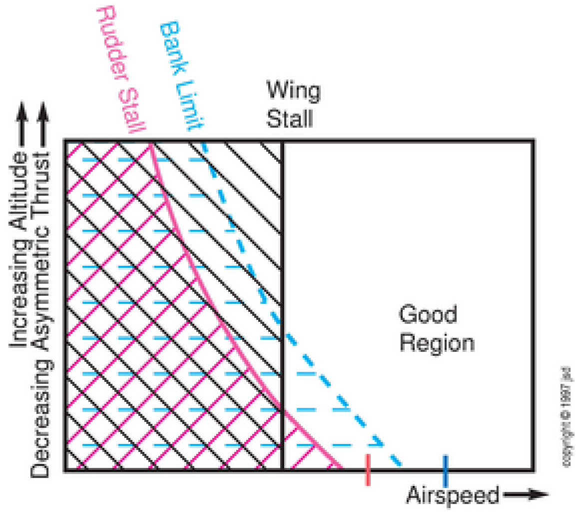

Figure 17.7: More Flaps, Less Weight, etc.

At the reduced weight, less lift is needed for supporting the weight of the airplane. As always, the horizontal component of lift, at any particular bank angle, is proportional to the weight of the airplane. Therefore, at any particular bank angle, you have less ability to oppose a boat turn. This is one reason why the cyan dotted curve moves to the right as we go from figure 17.6 to figure 17.7.

Limiting yourself to less than full rudder deflection does *not* reduce the amount of torque that must be produced in order to oppose the asymmetric thrust; it just means that the airplane will establish a slip to create the necessary force. (If there were an unbalanced torque, the airplane would not only turn, it would accelerate in the yaw-wise direction, rotating faster and faster.)

In this slipping condition, the fuselage produces a boat turn on top of whatever pseudo boat turn the rudder is producing, so you will need more bank to oppose the turn, and you will run up against the bank limitation sooner. This is the second reason that the dotted cyan curve (the bank limit) moves to the right.

And of course, if you limit yourself to a smaller bank, you will run up against the bank limit sooner. This is the third reason that the cyan curve moves to the right as we go from figure 17.6 to figure 17.7.

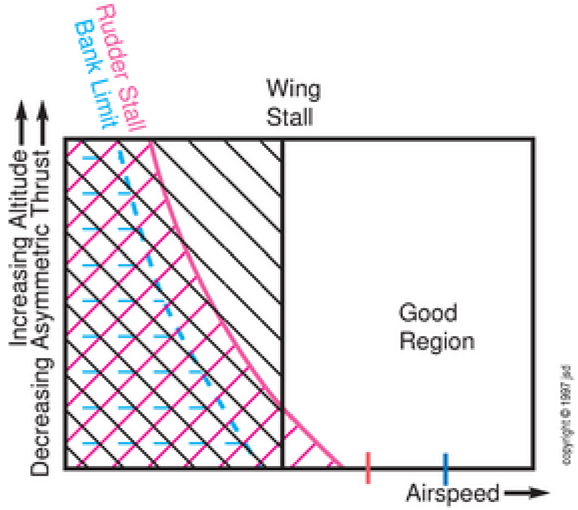

Figure 17.8: More Bank

Conversely, if you allow yourself a large bank (15 or 20 degrees) you can push the dotted cyan curve very far to the left, as indicated in figure 17.8.

Now let's consider what happens in different airplanes. For example #2, let's consider an airplane that has somewhat smaller wings. To compensate, the manufacturer specifies that flaps are to be extended in the certified takeoff configuration. The result is that the certified performance of the new plane is identical to the performance of example airplane #1, as shown in figure 17.6. The interesting wrinkle is this: if you fly the new airplane with flaps *retracted*, the performance is as shown in figure 17.9. Note the higher wing stall speed. The airplane will become uncontrollable at an airspeed well above the FAR 23.149 red-radial-line.

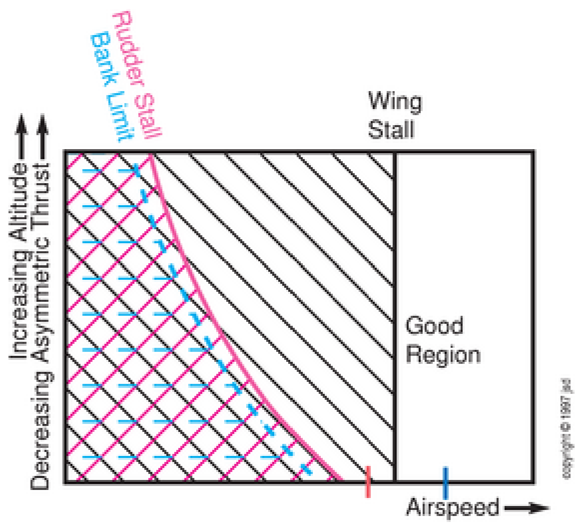

Figure 17.9: Certified Flaps Not Used

As a pilot, it is not important for you to memorize the details of what's going on in these figures. The point of all this is to convince you it's complicated, and highly dependent on circumstances that you don't have much control over. The one thing that's worth remembering is that you're OK down to redline airspeed *with departure flaps extended*. You "might" be OK down to lower speeds, but you don't generally know how much lower, and there's no safe or easy way to find out.

For example #3, let's take an airplane where the wing has a very low stalling speed. For such a plane, figure 17.6 never applies; figure 17.7 (or figure 17.8, depending on bank angle) applies even at max weight with the flaps retracted.

Let’s summarize what we know so far, in a form that is perhaps more directly useful when you are actually in the cockpit.

- Maintain a safe airspeed. This speed should be above the wing’s stalling speed *in the current configuration* and above the FAR 23.149 red-radial-line, whichever is higher. Leave yourself a reasonable margin of safety.
- The best procedure is to establish zero slip (or minimum slip, if full rudder deflection isn’t enough to establish zero slip).

- Then bank to establish the desired rate of turn (usually zero). The amount of bank increases as the weight decreases; use whatever bank angle does the job. Remember, though, that maintaining a safe airspeed is more important than getting exactly the right slip angle or bank angle.
- In your multi-engine training, you were probably given the chance to demonstrate “loss of directional control” under conditions where the “loss” resulted in a gentle boat turn toward the dead engine. You absolutely must not assume that the airplane will always behave this way. In other circumstances, you might get a sudden rudder stall or wing stall, either of which could result in a spin.
- If you want to demonstrate the gentle boat turn, you can *arrange* that it occurs before any of the nastier alternatives, as suggested in figure 17.7. You just have to put sufficiently strict limits on the bank and rudder deflection. Reduced weight helps, too. Turbocharging makes it easier to perform the demonstration at a safe altitude.
- If you go exploring speeds below the red-radial-line, things get dicey. If you allow yourself unlimited[22] bank, there is no doubt that you can maintain directional control right down to the point where the wing stalls and/or the rudder stalls. You can get a good estimate[23] of the wing’s stalling speed, but I can’t think of a safe way for you to find out whether or not the rudder will stall before the wings.[24] Please do *not* try to find this out experimentally!

For more information on engine-out procedures, see section 17.2.

17.1.9 Effect of Center of Mass

We know that we have to pay careful attention to the location of the airplane’s center of mass, since it has a big effect on the angle of attack stability; see for example section 6.1.1.

This leads us to wonder what effect center-of-mass position has on V_{MC}. There are two possible answers:

1. CM location has no effect whatsoever if you use the unwise wings-level technique depicted in figure 17.5.
2. CM location does matter if you use the recommended procedure depicted in figure 17.4. As the CM (or, more precisely, the center of lift) moves aft, V_{MC} increases.

In both cases, you need to create a torque to oppose the asymmetric thrust. You create it using a pair of forces with a lever arm between them. One force comes from the rudder.

In case (1), the rudder force is paired with a horizontal force due to air hitting the side of the fuselage. This fuselage horizontal lift depends on the shape of the airplane, but does not at all depend on the CM location.

There is a deep theorem of physics that says that for any two axes parallel to each other, the torque around one is the same as the torque around the other (provided there are no overall unbalanced forces on the system). In the zero-bank case, it means that V_{MC} can't depend on center of mass location (unless the airplane is actually turning, i.e. being accelerated sideways).

To understand the basis of this theorem, refer again to figure 17.5. Let's pick two pivot points *A* and *B* somewhere along the rudder/wing lever arm, as shown in the figure. (You can, if you wish, imagine them to be two possible locations of the center of mass; the CM is no better or worse than any other pivot point.)

When we calculate the total torque around each pivot point:

- The lever arm from *A* to the rudder is long, but the lever arm from *A* to the other horizontal force is short.
- The lever arm from *B* to the rudder is short, but the lever arm from *B* to the other horizontal force is long.

The total torque around *A* is exactly the same as the total torque around *B*. The total torque is the only thing that affects V_{MC}, and that is the same no matter what pivot point is used.

In case (2), the story is slightly different. The rudder force is paired with the horizontal component of lift from the wings, tail, et cetera. This component arises because you are in a slight bank, as illustrated in figure 17.4. The location of this force depends *indirectly* on the CM location, according to the following chain of reasoning:

a)

The large vertical component of lift must be located very close to the center of mass, to oppose the force of gravity; otherwise the airplane would be out of equilibrium in pitch.

b)

The small horizontal component of lift is located at the same place as the large vertical component.

Here's another way of saying the same thing: the location of the lift vector depends directly on the shape of the airplane, but you have to adjust the shape of the airplane in order to keep the center of lift located very close to the center of mass. Note that we are not talking about the lift of the wings alone, but the lift of the entire airplane including the tail. In the particular example illustrated in figure 17.4, the center of mass is located rather far forward. The tail has been adjusted to produce a negative amount of lift in order to maintain equilibrium in pitch. The horizontal component of lift depends directly on this contribution from the tail, which in turn depends on CM location.

As the center of lift moves aft, the lever arm between it and the rudder gets shorter. This means you need more rudder deflection and more bank to oppose any given amount of asymmetric thrust.

17.1.10 Effect of Drag (e.g. Landing Gear)

To reiterate: in engine-out flight you have two problems: impaired rate of climb, and asymmetric thrust which can lead to uncontrollable yaw if you're not careful.

You may be thinking that it is possible to counteract the asymmetric thrust using asymmetric drag. Technically, that's true, but as we shall see, it isn't particularly practical.

An unrealistically good type of asymmetric drag is shown in figure 17.10. A source of additional drag (a small parachute) is attached far out on the wing (on the working-engine side). Because it has a long lever arm, a modest amount of drag force will create a significant amount of yaw-wise torque. This will help you maintain directional control. Of course, the drag will exacerbate your rate-of-climb problems.

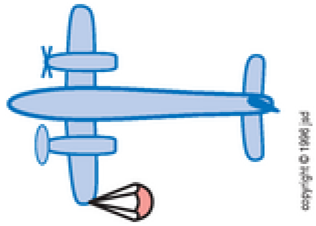

Figure 17.10: Asymmetric Drag (Useful, but Unrealistic)

If the parachute is attached at a different point, the results will be different. If it is attached near the working engine, as shown in figure 17.11, its contribution to the yaw budget will be exactly the same as if you had throttled back the working engine; added drag is the same as reduced thrust. The effect on climb performance is also the same as if you had throttled back. Obviously, using the throttle is more convenient and practical than adding asymmetric drag.

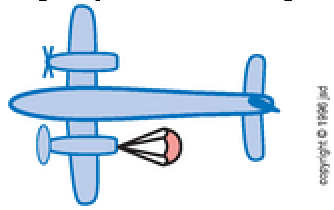

Figure 17.11: Asymmetric Drag (Useless)

Now we can do a more detailed analysis of how the landing gear contributes to the yaw-wise stability and equilibrium. Let's take the gear-up situation as a starting point, and see what *differences* arise when you put the gear down.

With the gear up, the forces are in equilibrium: thrust balances drag. With the gear down, there is extra drag. Eventually equilibrium will be restored somehow. Let's assume[25] the airplane just slows down, so that the extra drag of the gear will be balanced by reduced drag on the rest of the airplane.

So we have two new forces: a rearward contribution from the gear, and a forward contribution from the reduced drag on the rest of the airplane.

First, let's see what happens when the slip angle is zero. In that case the two new forces are oriented right along the line between them. This contributes nothing to the yaw-wise torque budget, because the forces have no component perpendicular to the lever-arm between them.

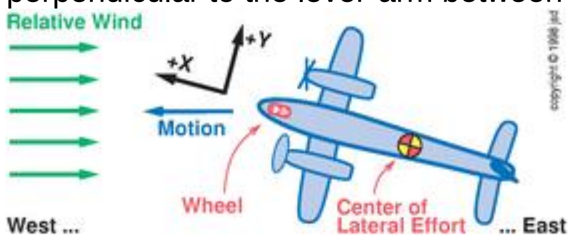

Figure 17.12: Extended Landing Gear (Worse than Useless)

Next, let's see what happens when (as shown in figure 17.12) a slip angle has developed. Once again, the new force on the wheel will be mostly a drag force, rearward in the direction of the relative wind. The other new force (the reduced drag on the rest of the airplane) will act in the opposite direction, centered at a place called the center of lateral effort.

Now we have a pair of forces with a component perpendicular to the lever arm. This will create a yaw-wise torque. The torque will grow in proportion to the slip angle. On most airplanes the nose wheel is far ahead of the center of lateral effort, so this will make a *negative* contribution to the yaw-wise stability.

As a final refinement, we consider the fact that when the wheel meets the air at an angle (as shown in figure 17.12), it acts a little bit like an airfoil and produces a force *perpendicular* to the relative wind, i.e. a sideways lift force. This force grows in proportion to the slip angle and makes another negative contribution to the yaw-wise stability.

To summarize this subsection:

- Extending the landing gear always creates drag, which impairs the rate of climb.
- To the extent that the landing gear creates *symmetric* drag, it contributes nothing to yaw-wise equilibrium.
- The landing gear typically makes a negative contribution to yaw-wise stability.[26]
- Usually the only contribution that is even theoretically helpful comes from the asymmetry of having one landing gear in the propwash of the working engine. However, this is not a practical advantage since you could achieve better rate of climb, the same equilibrium, and better stability by keeping the landing gear retracted and slightly reducing power on the working engine.

Of course, during the descent and landing phases, there are some obvious advantages to extending the landing gear.

There is a more-or-less endless list of other contributions to the yaw budget, but they are usually small and unimportant, especially if you maintain a steady speed, maintain zero slip angle, and keep the airplane balanced left/right.

Here are a couple of small items; you can probably think of others.

- With an engine out, if you don't crossfeed properly you can wind up with an unbalanced distribution of fuel. Then you will get weird yawing moments whenever you accelerate or decelerate.
- Normally when you make a power change, the forces is out of equilibrium until the new drag (at the new airspeed) comes in to balance with the new power setting. In an engine-out situation, this causes an out-of-equilibrium torque as well.

17.1.11 Roll Control

Whenever one or more engines are producing power, propeller drag will cause a rolling moment, as discussed in section 9.5. You will need to deflect the ailerons to the right to compensate.

Losing an engine will cause additional roll-wise problems on top of all your other problems. That's because the working engine creates more propwash over its wing, producing more lift on that side. You need to deflect the ailerons toward the working engine to compensate. Many airplanes have *aileron trim* to help you deal with this.

17.1.12 Critical Engine

On a typical twin, you will notice that the left engine causes more yaw trouble than the right engine does. There are several reasons for this, including helical propwash, twisted lift, and possibly P-factor.

First: Helical propwash was discussed in section 8.4 in connection with single-engine airplanes. The multi-engine story is partly the same and partly different. To be specific, let's consider a plane where the engines rotate clockwise as seen from behind.

Typically, in normal flight, most of the propwash misses the vertical tail, as shown in figure 17.13. However, because it spreads out on its way from the engine to the tail area, some fraction of the propwash does manage to hit the tail. The effect may be large or small, depending on the size and shape of the airplane.

You need to apply right rudder to compensate, just like in single-engine planes. If your airplane requires right rudder during the initial takeoff roll, it must be due to propwash; it can't be due to twisted lift (because there is no lift yet), and it can't be due to P-factor (because the prop disks are not inclined).

With one engine out, as long as you are able to maintain zero slip, the effect will be roughly half as large, because only one engine's propwash is acting on the vertical tail, as shown in figure 17.14.

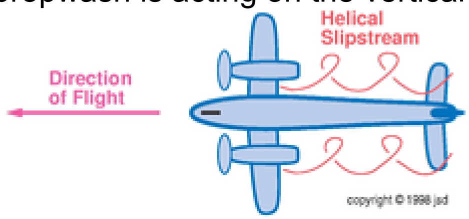

Figure 17.13: Normal: Helical Propwash Effect May Be Small or Large

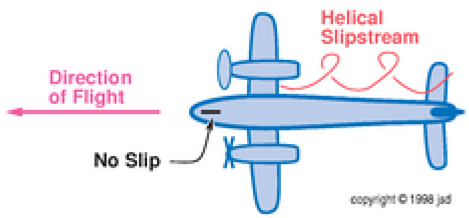

Figure 17.14: Engine Out: Helical Propwash Has Half As Much Effect

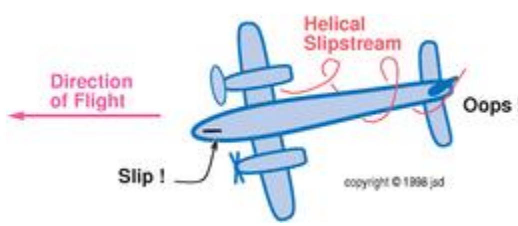

Figure 17.15: Slip: Helical Propwash Has More Effect

If you don't apply enough rudder to maintain zero slip, more of the tail will move into the propwash, as shown in figure 17.15. (At low airspeeds, you could easily have a situation where you *can't* apply enough rudder to prevent this.) Since the vertical tail sticks up, not down, the propwash from the right engine will be rotating in such a way as to reduce rudder effectiveness.[27] If possible, you should apply additional right rudder to compensate.

It is a bit ironic that propwash affects the yaw-wise torque budget *more* when you already have a big slip angle. Normally you don't allow that to happen unless you are forced to, so this effect is usually only noticeable at low airspeeds --- such as a V_{MC} demonstration, or a crosswind takeoff (especially a crosswind from the left).

In a plane with *four* propellers, the tail will be much more affected by the propwash from the inboard engines than from the outboard engines. By using the engines one at a time, and in various pairings, you can shed a lot of light on the effects discussed in this section.

Secondly: As mentioned in section 17.1.11, propeller drag creates a rolling moment and requires right aileron no matter which engine is running. This aileron deflection will produce a certain amount of twisted lift, even though the magnitude of the lift vector is the same on both sides, as discussed in section 8.8.4. You will need to apply right rudder to compensate. This will be most noticeable in high-power low-airspeed situations.

Thirdly: P-factor (asymmetric disk loading) makes a small contribution to the yaw-wise torque budget. I measured this in a light twin, as discussed in section 8.5.4, using both engines. The effect was small, but could be observed if you looked closely. With only one engine, the effect would be half as large.

I also calculated from theory that when the airspeed decreases from cruise to V_{MC}, the corresponding increase in angle of attack causes the center of effort of the propeller disk to move to the right by about one inch. That's not zero, but it's not very much, either.

Most of the effects that people blame on P-factor are really mainly due to a combination of adverse yaw and helical propwash.

To summarize: Some yaw contributions are unbiased, requiring rudder deflection depending on which engine is out, according to the simple rule: working foot, working engine. These include the asymmetric thrust (as diagrammed in figure 17.4) and the increased lift over the working engine's wing (as mentioned in section 17.1.11).

Some other contributions are biased to the right, requiring right rudder no matter which engine is out. These include helical propwash acting on the tail, propeller drag acting via twisted lift, and P-factor. These are what make one engine more critical than the other.

Terminology: The engine you most regret losing is called the *critical engine*. In a twin where both engines rotate clockwise, that will be the left engine. With the left engine out, you will run out of rudder authority sooner, because the biased contributions add to the unbiased contributions. (If the right engine were out, the biased contributions would work in your favor, reducing the amount of left rudder required.)

Some twins have counter-rotating propellers. (That is, one engine rotates clockwise while the other rotates counterclockwise.) In that case both engines cause equally much yaw trouble, and either (or neither) can be considered the critical engine.

17.2 Engine Out Procedures

Engine failure is an emergency. You might want to review the general discussion of emergencies in chapter 15.

Make sure you know the emergency checklist for your airplane. Not all airplanes are the same. The following discussion applies to a "generic" airplane, and serves to illustrate some important concepts, but should not be taken as a substitute for airplane-specific knowledge.

17.2.1 Basic Takeoff Considerations

During takeoff, it is important to be able to detect any problems promptly. Early in the takeoff roll, you should glance at the gauges (RPM, manifold pressure, fuel flow, and EGT) to make sure the readings are normal --- and that both engines are the same. Make sure the airplane "feels" like it is pulling straight, i.e. no unusual steering effort is required to keep it going straight.

If anything funny happens while there is adequate runway remaining ahead of you, close both throttles immediately and stop straight ahead. In a high-powered airplane, such as an airliner, there will be a point where it is not possible to stop on the runway but it is possible to continue accelerating then fly away safely on one engine. See section 17.2.2.

A light twin taking off on the same runway will use a smaller fraction of the runway for a normal takeoff, but will have worse single-engine performance. As a consequence, there will typically be a time even after liftoff when it is better to close the throttles and re-land on the remaining runway. Indeed, even if the remaining runway is not quite enough, you might want to land on it: Suppose that because of density altitude or whatever, your aircraft has poor single-engine climb performance. You will sustain vastly less damage if you land and slide off the end of the runway at low speed, rather than making an unsuccessful attempt to climb out on one engine.

In many light twins, the climb performance is OK with the landing gear retracted but very poor with it extended. Therefore a common rule is the following: when there is no more useful runway ahead, retract the gear. If an engine fails before that point, you know you are committed to landing; if it fails after that point, you know you are committed to climbing.

Some other twins have a very different problem: when the gear is *partially* retracted it is markedly draggier than either the fully-retracted or fully-extended position. In such aircraft, if the gear is down you have to leave it down, unless/until you have plenty of altitude.

17.2.2 Balanced Field Length; Takeoff Decision Speed

Sometimes you need to make a more sophisticated stop-versus-go decision. This requires a bit more pre-flight planning. The result will be expressed in terms of a *takeoff decision speed*, denoted V_1. During the takeoff roll, note the point where the airspeed crosses V_1. If you lose an engine before that point, stop. If you lose an engine after that point, continue the takeoff.

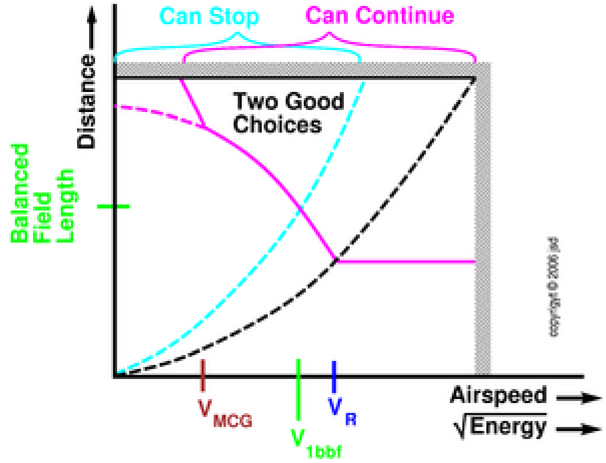

Figure 17.16: Accelerate-Stop, Accelerate-Go, or Both

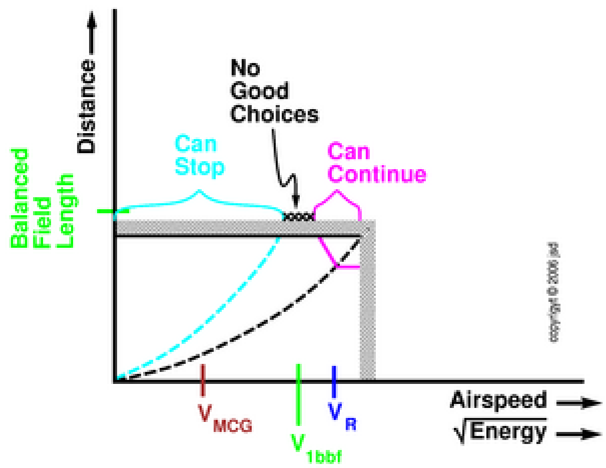

Figure 17.17: Accelerate-Stop, Accelerate-Go, or Neither

To understand V_1, refer to figure 17.16 and figure 17.17. In each figure, magenta curve is the *accelerate-go* distance, i.e. the runway length required to accelerate up to a given speed, lose an engine, and then go ahead with the takeoff. The required distance is plotted as a function of the speed at which the engine failure occurs. In figure 17.16, the magenta curve has three parts, each controlled by a different limiting factor:

- The rightmost part of the curve, where it runs horizontally, represents the runway usage in the normal case with all engines running. You attain flying speed well before reaching the end of the runway.
- If the engine failure occurs during the takeoff roll, the runway requirement is increased, because the acceleration will be impaired. This is represented by the gently sloping section in the middle of the curve.
- If the engine failure occurs at a very low airspeed, you might have uncontrollable yaw problems. Continuing the takeoff under such circumstances is not recommended. V_{MCG} denotes the minimum control speed for ground operations.

The dashed black curve in each figure is the distance you have actually travelled down the runway, as a function of speed, assuming both engines are working normally.

The dashed cyan curve in each figure is the *accelerate-stop* distance, i.e. the runway length required to accelerate up to a given speed and then stop straight ahead. It is plotted as a function of the speed at which the decision is made. You can see that if the decision is made at a low speed, very little runway is used.

The cyan curve continues well to the right of V_R, representing the case where you actually have become airborne, but decide to close the throttles and re-land straight ahead. If you have plenty of runway, this might be a very sensible thing to do.

Note: In this case (landing straight ahead, just after liftoff) we need to relabel the horizontal axis in figure 17.16. In all cases, what matters is the amount of mechanical energy you must get rid of in order to stop the plane. Before liftoff, speed is the only contribution to the energy, so we can label the axis either as speed or as square root of energy; they are equivalent. After liftoff, √energy remains correct but airspeed does not.

The horizontal black line near the top of each figure represents the actual runway length. You are allowed to use less distance than this, but you can't use more.

If you lose an engine at a "sufficiently" low airspeed, you can stop. To ascertain how low is "sufficiently" low, look at the place where the dashed cyan curve (representing accelerate-stop) crosses the horizontal black line (representing runway length).

If you lose an engine at a "sufficiently" high airspeed, you can continue the takeoff and fly away. To ascertain how high is "sufficiently" high, look at the place where the magenta curve (representing accelerate-go) crosses the horizontal black line (representing runway length).

It is useful to compare figure 17.16 with figure 17.17, which represents the same aircraft operating on a shorter runway.

In figure 17.16, midway through the takeoff roll, there is a region where you can either abort the takeoff and stop on the runway, or continue the takeoff and fly away on one engine.

In figure 17.17, there is a nasty region where you can *neither* stop on the remaining runway, nor fly away on one engine. See below for more discussion of this.

Of particular interest is the point where the accelerate-go curve crosses the accelerate-stop curve in figure 17.16. This determines the *balanced-field length*, i.e. minimum runway length to guarantee that at every point during the takeoff there will be at least one good option.

If the actual runway is barely longer than the balanced-field length, there is only one value of V_1 that makes sense, namely $V_1 = V_{1bbf}$. In general, V_1 represents a decision speed, and the suffix "bbf" refers to "barely balanced field". On such a runway you need to pay close attention to the airspeed as it crosses V_1. You mustn't try to continue the takeoff when you should be stopping, and you mustn't try to stop when you should be continuing.

If there is lots of extra runway, you have some freedom to choose a V_1 value higher or lower than V_{1bbf}. The lower limit is where the accelerate-go curve (magenta) crosses the runway-length line (black). The upper limit is where the accelerate-stop curve (cyan) crosses the runway-length curve. Given such a choice, most people choose a relatively high value, since accelerate-stop is usually preferable to accelerate-go. It is better to be on the ground wishing you were in the air, instead of being in the air wishing you were on the ground. Whatever you do, don't choose an extremely-high or extremely-low value; you should distribute some of your safety margin to the accelerate-stop maneuver and some to the accelerate-go maneuver.

The details of the curves in figure 17.16 will depend on factors such as braking conditions, wind, temperature, weight, type of airplane, et cetera. If you have four powerful engines and ice on the runway, V_{1bbf} could be quite low compared to V_R. On the other hand, if you have two smallish engines and good braking, V_{1bbf} will be practically equal to V_R; most light piston twins fall into this category. In the extreme case where you can't climb with an engine out, the concept of balanced field loses its meaning.

- Some Pilot's Operating Handbooks will give you charts and tables for V_{1bbf} for various conditions. They call it V_1, on the unjustifiable assumption that you will choose $V_1 = V_{1bbf}$ even if you have extra runway.
- Sometimes they don't mention V_1 at all; they just suggest a value for V_R and tabulate the accelerate-stop and accelerate-go distances for various conditions. Check both tables; whichever number is larger tells you how much runway you want to have when choosing $V_1 = V_R$.
- You can gain additional insight as to the shape of the accelerate-stop curve by looking at other tables: add the "takeoff ground roll" distance to the "landing ground roll" distance. Similarly add the "takeoff over obstacle" distance to the "landing over obstacle" distance.
- Obviously you should adjust the black line (runway length) according to the runway you will be using.

Airliners are not allowed to take off unless the available runway exceeds the balanced field length. In contrast, in general aviation, you may use a shorter runway if you want. In that case, there will be a period during the middle of the takeoff roll where you can neither stop nor continue safely on one engine. In such a case you must shut down the good engine and apply the brakes. It is much better to hit the trees at the end of the runway when you are "almost" stopped then to hit them when you are "almost" at full flying speed. This seems obvious on paper, but when you are in the cockpit it takes a lot of willpower to actually shut down the good engine. Think about this. Promise yourself that you will do it right.

17.2.3 Procedure: Low Altitude

Once you are airborne and assured of single-engine climb performance, the following checklist applies to our generic airplane at low altitudes: three things, five things, four things. Specifically:

- Three things: airspeed, ball and needle.
- Five things: mixtures, propellers, throttles, gear, flaps.
- Four things: identify, verify, feather, secure.

Now let's spell each item out in more detail, for the case where your initial speed is above V_{MC}:

- Three things: airspeed, ball, and needle. That is, pitch to maintain best-climb speed. Then the easiest thing to do is apply enough rudder pressure to center the inclinometer ball. "Step on the ball". This involves more rudder pressure than you need to establish zero slip, but at this stage of the game you are in a hurry and centering the ball is a rough-and-ready approximation. With the ball centered, nonturning flight will require a slight bank toward the working engine. (Wings-level non-turning flight is really overdoing it. It involves a slip toward the dead engine, which puts an unnecessary burden on the rudder and degrades climb performance.)
- Five things: mixture controls forward (rich), propeller controls forward (fine pitch), throttles forward (maximum allowable power), landing gear retracted, flaps retracted.

- At this point you might want to check airspeed, ball, and needle again. The airplane has probably slowed down quite a bit, so you may need to make a pitch adjustment and retrim.
- Four things: identify, verify, feather, secure. Let's suppose the right engine has failed.
  - Identify: raise your right knee (dead foot, dead engine) and say aloud, "the right engine has failed".
  - Verify: retard the right throttle. There should be no change in the situation. If you retard the wrong throttle you will notice immediately; push it forward again, go back to step 1 (airspeed, ball, and needle), and try again.
  - Feather: grab the correct propeller control, pull it back a little ways and listen to make sure you've got the right one, then pull it all the way into the feather position. (If this is a simulated emergency, just pull it back half an inch or so and tell your instructor that you are simulating the feather.)
  - Secure: when the engine has stopped spinning, shut off its mixture control, its fuel supply, its boost pumps, its alternator, its magnetos, et cetera. Close its cowl flaps and open the cowl flaps on the working engine.

- Finally, check airspeed, ball, and needle again. Make sure you are trimmed for best-climb speed. Establish zero slip by applying somewhat less rudder pressure than would be necessary to center the ball. In zero-slip flight the ball will be off-center by one-third to one-half of its diameter. Use the rudder trim to hold this arrangement. In this condition, nonturning flight will require banking a few degrees toward the good engine: "raise the dead".

Here are the same items again, for the where you have a fair bit of initial altitude, but your initial speed is below V_{MC}:[28]

- To avoid losing directional control, reduce power on the good engine. Don't worry about which is which; retard *both* throttles. If this means making a power-off landing, so be it. Otherwise, if you have sufficient altitude, dive to trade altitude for speed.
- Five things: After you have gotten back to V_{MC}, advance both throttles, both propeller controls, and both mixture controls. Retract the gear (if appropriate) and retract the flaps.
- Three things: airspeed, ball, and needle. As always, "step on the ball" to get the airflow approximately coordinated. You may be unable to establish zero slip, even with full rudder deflection, in which case you should apply full rudder and let the airplane establish whatever slip is necessary to oppose the asymmetric thrust. To establish nonturning flight, the wings should be almost horizontal, with the dead engine raised slightly.
- Continue speeding up to best-climb speed. This will probably require cashing in additional altitude.
- Four things: identify, verify, feather, secure --- the same as before.

Reading about these things is good, but not sufficient. You really should to up with an instructor and practice these things. Practice until the right actions become routine. Review it at least once every six months.

17.2.4 Procedure: Higher Altitude

Finally, here is the procedure for the case where you have a reasonable airspeed and a reasonable altitude, say 1000 feet AGL or more. You should *not* be in any big hurry to feather the offending engine. If the problem is minor, restarting will be a lot easier if the engine is not feathered. The checklist should be:

- Three things: airspeed, ball, and needle.
- Five things: mixtures, propellers, throttles, gear, flaps.
- Four things: identify, verify, debug, think.

Take a systematic approach to debugging. Start somewhere on the panel and then check everything you come to, systematically.

It doesn't hurt to be logical, but remember that in an actual emergency, you will be much less logical than you normally are. Unless it is obvious what the problem is, check everything, in order. Don't just check the things that come to mind. Systematic habits are more likely to stay with you.

After you've checked everything once, then try applying logic. What was the last thing you fiddled with before the failure? Did you just shut off the fuel boost pumps? Maybe you should switch them back on; look at the fuel pressure... or did you miss the boost pumps and turn off the magnetos instead? Did you just switch from the inboard to the outboard tanks? Maybe you should switch back, or switch to crossfeed.

Remember that you may be unable to climb or even maintain altitude on one engine. See section 17.1.2 for a discussion of this.

17.2.5 Airspeed Management

The airspeed that gives the best single-engine rate of climb is referred to as V_{YSE}. The value of V_{YSE} for standard conditions (max weight, sea level, etc.) is marked on the airspeed indicator by a blue radial line, and is commonly called blueline airspeed.

If an engine fails, you should (except in certain special situations) maintain a speed at or above V_{YSE}. Maintain thine airspeed lest the ground arise and smite thee.

One exception to the foregoing rule: If you need altitude to avoid an obstacle, you'll be better off at V_{XSE} (best *angle* of climb) as opposed to V_{YSE} (best *rate* of climb). In typical trainers, the single-engine performance is so anemic that V_{XSE} will be only slightly slower than V_{YSE}, for reasons illustrated in figure 7.8. Indeed, if you are above the single-engine absolute ceiling, the climb rate is negative and V_{XSE} is slightly *faster* than V_{YSE}.

Another exception: The optimal airspeed on final approach is typically less than V_{YSE}. You're not climbing, so you don't need to worry about climb performance. With the good engine at idle, you can go as slow as you want. (But then you've got big problems if you need to go around, as discussed in section 17.2.6.)

Yet another exception: Suppose your airplane has enough single-engine climb performance that the minimum level-flight speed V_{ZSE} is significantly slower than V_{YSE}. (See figure 7.7.) Further suppose you lose an engine at night at low altitude over a dark forest, at a very low airspeed. You don't want to dive all the way to V_{YSE}, because that could take you into the trees. A more modest dive will produce a speed above V_{ZSE}. Thereafter you can speed up in level flight, or climb at constant airspeed. In this scenario you don't need *best* rate of climb as long as you have *some* rate of climb. Related issues are discussed in section 7.5.3 and section 13.3.

Another relevant airspeed is the minimum control airspeed, V_{MC}. As discussed in section 17.1.7, you could get into big trouble if the airspeed gets too much below V_{MC}. At any speed above V_{MC} you should apply full power on the good engine and speed up to best-climb airspeed. Don't be shy about diving to get to best-climb speed; otherwise, if you start at a low airspeed, the airplane might not be able to climb or speed up at all.

At speeds below V_{MC}, you will be forced to use *less* than full power on the good engine, to keep the yaw from getting out of hand while you speed up to V_{MC}. Losing an engine at an airspeed below V_{MC} is such a nasty situation that most people don't practice during training. To recover, you have to partially close the throttle on the good engine, which takes a lot of willpower. You don't have much time to think. Then you have to dive, cashing in quite a lot of altitude to get the needed airspeed. The usual procedure calls for speeding up to V_{MC} plus a few knots, to give yourself a little margin, before returning the good engine to full power.

17.2.6 Engine-Out Go-Arounds

The first thing to be said about engine-out go-arounds is that you should make every possible effort to make sure that you do not ever need to perform one. The most common reason for a go-around is that you are about to land long and run off the end of the runway. Therefore, if at all possible, fly to somewhere that has a really long runway before attempting any engine-out landing.

The second thing to be said is that for typical airplanes there is a certain height above the ground --- often a surprisingly great height --- below which an engine-out go-around is simply not possible. The reason for this is simple: the typical approach speed is quite slow --- not only below V_{YSE} (best-climb speed) but near or even below V_{ZSE} (zero-climb speed, as defined in figure 7.7). If you try to climb out at low airspeed, the rate of climb may well be negative. In order to speed up from approach speed to any reasonable climb speed, you will need to cash in quite a lot of altitude. You will also consume time (and altitude) while you retract the landing gear, et cetera. In a Seneca, the decision to go around must be made above 400 feet AGL; below that altitude, you *are* going to touch down. If the runway is obstructed, land on the taxiway, or the infield, or whatever. If you have enough runway to touch down but not stop, consider doing a touch and go (which works better if you leave the gear down). Also consider landing anyway, with the expectation of going off the end at low speed; this is vastly preferable to hitting obstructions at high speed during an unsuccessful go-around.

17.2.7 Low-Speed Engine-Out Demonstrations

There are several key ideas I want my students to know about low-speed engine-out performance, including:

#0)

If you are ever in a position where you can descend to a safe landing *without* using high power on the good engine, by all means do so. This is not a game where you get extra points for climbing when you don't have to.

The rest of the discussion assumes you need the maximum achievable power from the good engine.

#1)

Starting from moderate speeds, as you slow down you will need more and more rudder to maintain coordinated flight. This is the coordinated regime. The amount of bank needed to maintain nonturning flight is basically constant.

#2)

There comes a point where you run out of rudder authority and cannot maintain coordinated flight. As the speed decreases further, the slip angle automatically increases, and more boat turn gets added to the pseudo boat turn. This is the uncoordinated regime. The bank angle must increase as airspeed decreases if you want to maintain nonturning flight.

#3)

You can maintain control down to V_{MC} (i.e. FAR 23.149 red-radial-line) in the takeoff configuration.[29]

#4)

If you persist in engine-out flight down to sufficiently low airspeed, at some point the wings and/or rudder will stall and you will be very sorry.

#5)

If you are below V_{MC}, you should reduce power on the good engine, dive to regain V_{MC}, and then re-open the throttle on the good engine.

#6)

If you are below V_{YSE}, you should dive to regain V_{YSE}, obstructions permitting.

#7)

To clear distant obstacles, you want to dive to achieve V_{XSE} as soon as possible. To clear nearby obstacles, you don't want to dive below their altitude, obviously. For a combination of obstacles, you face some tricky tradeoffs. The best solution is to make sure you never get into a low-altitude low-airspeed situation.

The aircraft manufacturer is supposed to specify a minimum safe speed for intentional engine cuts, denoted V_{SSE}, which is typically quite a bit higher than V_{MC}.

To demonstrate these key ideas, you should start in the takeoff configuration at a speed at or above V_{SSE}. Then cut one engine, and gradually reduce airspeed. This will demonstrate idea #1 immediately. If there is a chance you will reach V_{MC} before you have a chance to demonstrate idea #2, it is a good idea to artificially limit the available rudder deflection, perhaps by blocking the pedal with the toe of your other shoe. We do not wish to demonstrate idea #4. After demonstrating flight slightly above V_{MC} (idea #3), return to V_{YSE} (idea #6) and then resume normal flight.

To demonstrate a portion of idea #5, we use a separate maneuver. Starting with both engines at idle, perform a power-off stall. Recover to V_{MC}, then using only one engine, recover to V_{YSE}.

The FAA commercial pilot multi-engine practical test standard ("PTS") contains a task called "ENGINE INOPERATIVE --- LOSS OF DIRECTIONAL CONTROL DEMONSTRATION". The requirements are a bit confusing. For one thing, the PTS speaks of banking "for best performance and controllability" but doesn't say how to trade off performance versus controllability. Best climb performance typically requires less bank than best ultra-low-speed controllability.

Among many examiners, the traditions concerning this task are as follows:

1. Start at a safe altitude and safe airspeed. The PTS calls for V_{YSE} plus ten knots.
2. The PTS calls for flaps set for takeoff. However, there are some planes (such as a Seneca) where the certified takeoff checklist calls for zero flaps, and where the red-radial-line is essentially right at the stalling speed. In such planes you can lower the stalling speed by extending the flaps, which will make the demonstration safer and easier. Most examiners are happy to permit this. Call it "short field" takeoff configuration if you like. In other planes (such as an Apache) where the stalling speed is already well below the red-radial-line, don't bother extending the flaps.
3. Reduce power on one engine to idle. Do not actually stop or feather the engine.
4. Depress the rudder to establish zero slip. This gives best performance.
5. Bank to establish nonturning flight. This will be a very shallow bank.
6. Block the rudder motion to ensure that you run out of rudder before the airspeed gets close to the edge of the envelope. Stay away from the red-radial-line and wing-stall speed, whichever is higher. The PTS calls for staying 20 knots above the wing-stall speed.[30]
7. Gradually decrease airspeed.

8. After you run out of rudder deflection, the unwritten rule is that you should not increase the bank. That means the airplane will start to turn. The turn is your signal that it is time to begin the recovery phase. You are *not* being asked to demonstrate key idea #2.

Before the checkride, you should discuss these unwritten rules with your examiner, to make sure you are both singing the same tune.

The airspeed limit is needed to ensure safety. The artificial limits on rudder deflection and bank are needed so that you can demonstrate a nice gentle boat turn, by pretending to run out of control authority; otherwise the airplane would be controllable at all safe airspeeds and there would be nothing to demonstrate.

Note that in everyday (non-checkride) flight, if you run out of rudder authority at a speed above red-radial-line (and if you are sure you want to be flying so slowly) you would just smoothly enter the uncoordinated regime and increase the bank.

You should *not* do demonstrations the way FAR 23.149 seems to suggest:

- Do not suddenly shut down one engine at a low airspeed. Shut it down at or above V_{SSE} and then slow down. (Alternatively, I suppose it would be safe to fly below V_{SSE} and *gradually* reduce power on one engine, but I can't think of a reason why you would want to.)
- Do not explore low-speed performance at low altitude. The practical test standard calls for a minimum of 3000 feet AGL.

Full-blown FAR 23 V_{MC} determinations should be left to professional test pilots. For that matter, not even test pilots dare to experiment with loss of control at low altitude. They are not crazy; they experiment at a series of safe altitudes and then extrapolate.

1
We won't discuss aircraft that have centerline thrust, e.g. the Cessna 337 Mixmaster.

2
At this initial moment, other contributions to the yaw-wise torque budget are negligible -- assuming the airplane was in more-or-less coordinated flight before the engine failed.

3
Remember, this section is assuming cruise airspeed. At lower airspeeds, it may be critically important to *immediately* apply *full* rudder.

4
This means total lift, including the contributions of the wings, horizontal tail, et cetera. The center of lift will be located quite close to the center of mass, for reasons discussed in section 6.1.3.

5
If you forget to roll the wings level before using the rudder to stop the heading change, you could easily find yourself stepping on the wrong rudder. For instance, if the left engine fails during a turn to the right, you might be tempted to stop the turn by stepping on the left (wrong!) rudder.

6
To know how much bank and/or how much inclinometer ball deflection corresponds to zero slip, you can (a) recall from your training flights what configuration corresponds to best performance, (b) recall from flights with a slip string what configuration corresponds to zero

slip, or (c) let the inclinometer ball go off-center by half its width, which is usually "close enough" to the right answer.

7

See section 19.4 for a discussion of E-down and related concepts.

8

The confusion is understandable, since asymmetric thrust is about the only way you can maintain an inclination without being in a slip of some kind.

9

... or cruise performance, for that matter --- engine-out or otherwise.

10

In fact, an undeflected rudder produces a less stall-resistant shape, which will probably stall at a higher airspeed.

11

If the red-radial-line (as defined in the following section) is down near the bottom of the green arc, it is a good guess that wing stall is what limits the airplane's low-speed controllability. Conversely, if the red-radial-line is much higher than the bottom of the green arc, you can guess that rudder stall is what limits the low-speed controllability (unless the red-radial-line is artificially high because of the arbitrary 5 degree bank limit in FAR 23.149).

12

A very similar regulation, FAR 25.149, applies to transport-category aircraft (e.g. airliners).

13

The notion of "critical engine" is discussed in section 17.1.12.

14

If the airplane weights more than 6000 pounds, FAR 23.66 requires the airplane to be able to climb with one engine inoperative, at an airspeed "equal to that achieved at 50 feet" after takeoff. Even this does not require climb at V_{MC}.

15

The similar term *red arc* frequently refers to the region at the high end of the airspeed indicator. There are of course other red lines and red arcs (on tachometers, oil-temperature gauges, etc.), but they are not relevant to the present discussion.

16

... except perhaps in a minor way that is not relevant here. See section 11.5.6 for a general discussion of slip angle versus bank angle.

17

There are many airplanes that are quite nicely behaved even under conditions that require more than 5 degrees of bank in order to maintain non-turning engine-out flight.

18

That's right: there is no guarantee that 5 degrees is safe. It is commonly assumed that "the manufacturer must have tested a 5 degree bank because that is the maximum allowed". But in fact, best control might be achieved at 2 or 3 degrees, and there is no reason to assume that the manufacturer ever tried more than that. Remember, in non-slipping flight the bank required is quite modest (and independent of airspeed), and there is not much that the manufacturer can achieve by slipping that could not be better achieved by more rudder deflection.

19

This means the takeoff configuration as specified in the Pilot's Operating Handbook or Airplane Flight Manual. Remember that these documents are legally part of the airplane. You can't have a certified V_{MC} without a certified takeoff checklist.

20

In this example, we are assuming that wing stall (not rudder stall) is what limits low-speed handling.

21

This option is available to us because, for this airplane, the certified takeoff configuration did not call for flaps to be extended.

22

We are talking about rather modest bank angles, perhaps 15 or 20 degrees, that would not get you into trouble in other circumstances.

23

It may be a good idea to check the wing's stalling speed by performing a stall with both engines at zero thrust. The zero-engine stalling speed won't be quite the same as the one-engine stalling speed, but it should be a useful estimate.

24

I have done calculations that indicate that for certain light trainers, the wings will almost always stall before the rudder, but you absolutely should not assume that this is true for all airplanes in all circumstances.

25

At the end of this section, other possibilities will be considered.

26

One could imagine designing an airplane with the landing gear so far aft that they were behind the lateral center of effort, in which case they would increase yaw-wise stability.

27

The propwash from the left engine is actually helpful.

28

It's a little hard to see how this situation could arise in the course of normal flying. However, (a) such a situation is sometimes created as part of a training exercise, and (b) it could arise if the pilot mishandles an engine-out situation, squandering the initial airspeed.

29

In other configurations, you can maintain control down to red-radial-line or V_S, whichever is higher.

30

In airplanes where V_{MC} is at or near V_S, V_S + 20 may seem like a generous or even excessive margin of safety. In other airplanes, however, V_S + 20 is not nearly enough. You need to be careful, since there are plenty of airplanes where V_S + 20 is near (or even below) V_{MC}. A better criterion might be to stay above V_S + 10, and above red-radial-line + 10, whichever is higher.

18 Stalls and Spins

Caution: Cape does not enable user to fly.
--- warning label on Superman
costume sold at Walmart

Spins are tricky. After reading several aerodynamics texts and hundreds of pages of NASA spin-tunnel research reports, I find it striking how much remains unknown about what happens in a spin.

18.1 Stalls: Causes and Effects

Here's a basic yet important fact: if you don't stall the airplane, it won't spin. Therefore, let's begin by reviewing stalls.

As discussed in section 5.3, the stall occurs at the critical angle of attack, which is defined to be the point where a further increase in angle of attack does not produce a further increase in coefficient of lift.

Nothing magical happens at the critical angle of attack. Lift does not go to zero; indeed the coefficient of lift is at its maximum there. Vertical damping goes smoothly through zero as the airplane goes through the critical angle of attack, and roll damping goes through zero shortly thereafter. An airplane flying 0.1 degree beyond the critical angle of attack will behave itself only very slightly worse than it would 0.1 degree below.

If we go far beyond the critical angle of attack (the "deeply stalled" regime) the coefficient of lift is greatly reduced, and the coefficient of drag is greatly increased. The airplane will descend rapidly, perhaps at thousands of feet per minute. Remember, though: the wing is *still supporting the weight of the airplane.* If it were not, then there would be an unbalanced vertical force, and by Newton's law the airplane would be not only descending but *accelerating* downward. If the wings were really producing zero force (for instance, if you snapped the wings off the airplane) the fuselage would accelerate downward until it reached a vertical velocity (several hundred knots) such that weight was balanced by fuselage drag.

18.2 Stalling Part vs. All of the Wing

We can arbitrarily divide the wing into sections. Each section contributes something to the total lift. It is highly desirable (as discussed in section 5.4.3) to have the coefficient of lift for sections near the wing-root reach its maximum early, and start decreasing, while the coefficient of lift for sections near the tips continues increasing[1] (as a function of angle of attack).

Therefore it makes perfect sense to say that the sections near the roots are stalled while the sections near the tips are not stalled. If only a small region near the root is stalled, the wing *as a whole* will still have an increasing coefficient of lift --- and will therefore not be stalled.

We see that the wing will continue to produce lots of lift well beyond the point where part of it is stalling. This is the extreme slow-flight regime --- you can fly around all day with half of each wing stalled (although it takes a bit of skill and might overheat the engine).

18.3 Boundary Layers

There is a very simple rule in aerodynamics that says the velocity of the fluid right next to the wing (or any other surface) is *zero.* This is called the *no-slip boundary condition*. Next to the surface there is a thin layer, called the *boundary layer*, in which the velocity increases from zero to its full value.

18.3.1 Separated versus Attached Flow

The wing works best when there is *attached flow*, which means that air flowing near the wing follows the contour of the wing. The opposite of attached flow is *separated flow*.

For attached flow, as we move through the boundary layer from the wing surface out to the full-speed flow, there is practically no pressure change. Sometimes it helps to think about attached flow in the following way: Imagine removing the boundary layer and replacing it with a layer of hard putty that redefines the shape of the wing. Then imagine "lubricating" the new wing so that the air slides freely past it; the no-slip condition no longer applies. Bernoulli's principle can be used to calculate the pressure on the surface of the putty (whereas it could not be applied inside the boundary layer). Forces are transmitted from the air, through the putty, into the wing and the rest of the airplane. The putty-covered wing may not be the most desirable shape, but it won't necessarily be terrible.

For separated flow, the putty model does not work. Suppose I want to pick up a piece of lint from the floor using a high-powered vacuum cleaner. If I keep the hose 3 feet away from the floor, it will never work; I could have absolute zero pressure at the mouth of the hose, but the low pressure region would be "separated" from the floor and the lint. If I move the hose closer to the floor, eventually it will develop low pressure near the floor. This is part of the problem with separated flow: there is low pressure somewhere, but not where you need it. Separation can have multiple evil effects:

- Separation means the air doesn't follow the contour of the wing. This is somewhat like having a really thick boundary layer. The wing can't force the air into the optimal flow pattern, so not as much low pressure is produced.
- Whatever low pressure is produced isn't all attached to the wing surface. This is a new problem that an attached flow would not have, no matter how thick the boundary layer.
- On a non-streamlined object such as a golf ball, there is a lot of drag (specifically: form drag, as discussed in section 4.4) because separation disrupts a desirable high-pressure area behind the ball.

18.3.2 Laminar versus Turbulent Flow

In the simplest case, there is *laminar flow*, in which every small parcel of air has a definite velocity, and the velocity varies smoothly from place to place. The other possibility is called *turbulent flow*, in which:

- at any given point the velocity fluctuates as a function of time, and
- at any given time the velocity changes rapidly as we move from point to point, even for nearby points.

The closer we look, the more fluctuations we see.

Attached turbulent flow produces a lot of mixing. Some bits of air move up, down, left, right, faster, and slower relative to the average rearward flow.

For separated laminar flow, there will be some reverse flow (noseward, opposing the overall rearward flow) but the pattern in space will be much smoother than it would be for turbulent flow, and it will not fluctuate in time.

You can tell whether a situation is likely to be turbulent if you know the *Reynolds number*. You don't need to know the details, but roughly speaking small objects moving slowly through viscous fluids (like honey) have low Reynolds numbers, while large objects moving quickly through thin fluids (like air) have high Reynolds numbers. Any system with a Reynolds number less than about 10 is expected to have laminar flow everywhere. If you drop your FAA "Pilot Proficiency Award" wings into a jar of honey, they will settle to the bottom very slowly. The flow will be laminar everywhere, since the Reynolds is slightly less than 1. There will be no separation, no turbulence, and no form drag --- just lots of skin-friction drag.

Systems with Reynolds numbers greater than 10 or so are expected to create at least some turbulence. Airplanes operate at Reynolds numbers in the millions. The wing will have a laminar boundary layer near the leading edge, but as the air moves back over the wing, at some point the boundary layer will become turbulent. This is called the *transition to turbulence* or simply the *boundary layer transition*. Also at some point (before or after the transition to turbulence) the airflow will become separated. The designers try to keep the region of separation rather small and near the trailing edge. In order to make a wing develop a lot of lift without stalling, it helps to minimize the amount of separation.

18.3.3 Boundary Layer Control

One scheme[2] for controlling separation involves the use of vortex generators. Those are the little blades you see on the top of some wings, sticking up into the airstream at funny angles. Each blade works like the moldboard of a plow, reaching out into the high-velocity airstream and turning the layers over --- plowing energy into the inner layers.

Re-energizing the boundary layer allows the wing to fly at higher angles of attack (and therefore higher coefficients of lift) without stalling. This improves your ability to operate out of short and/or obstructed fields.

The vorticity created by these little vortex generators should not be confused with the bound vortex, the big vortex associated with the circulation that supports the weight of the airplane. As discussed in section 3.12, to create lift there must be vortex lines running *along* the span, associated with air circulating *around* the wing. Vortex generators can't provide that; their vortex lines run chordwise, not spanwise.[3]

Boundary-layer turbulence (whether created by vortex generators or otherwise) also helps prevent separation, once again by stirring additional energy into the inner sublayers of the boundary layer.

On a golf ball, 99% of the drag is form drag, and only 1% is skin-friction drag. The dimples in the golf ball provoke turbulence, adding energy to the boundary layer. This allows the flow to stay attached longer, maintaining the high-pressure region behind the ball, thereby decreasing the amount of form drag. The turbulence of course increases the amount of skin-friction drag, but it is worth it.[4]

Bernoulli's principle does not apply inside the boundary layer, separated or otherwise. As discussed in section 3.4, Bernoulli's principle applies in situations where pressure (potential energy) and airspeed squared (kinetic energy) add up to a constant. This is not the case in the boundary layer, because friction there converts a significant amount of the energy into heat.

Do vortex generators play the same role as dimples on a golf ball? Not exactly. Unlike a golf ball, a wing is supposed to produce lift. Also unlike a golf ball, a wing is highly streamlined; consequently, its form drag is not predominant over skin-friction drag. Vortex generators are typically used to improve lift at high angles of attack (by fending off loss of lift due to separation). They might improve performance at high speeds, i.e. low angle of attack, by decreasing form drag at the expense of skin-friction drag -- but probably not by much.

If you want ultra-low drag, and don't care about short-field performance, you want a wing with as much laminar flow as possible. Designing a “laminar flow wing” is exquisitely difficult, especially in the real world where the laminar flow could be disturbed by rain, ice, mud, and splattered bugs on the leading edge.

There is always *some* separation on every airfoil section. The separation grows as the angle of attack increases. If there is too much separation, it cuts into the wing's ability to produce lift. If there were no separation, the wing could continue producing lift up to a very high angle of attack (thereby achieving a fantastically high coefficient of lift).

Having lots of separation is the dominant cause (but not the definition) of stalling.[5] Remember: the stall occurs at the critical angle of attack, i.e. the point where max coefficient of lift is attained.

18.3.4 Recap: Turbulence and Boundary Layers

A full discussion of turbulence and/or separated flow is beyond the scope of this book; indeed, trying to *really* understand and control these phenomena is a topic of current research. There is nothing simple about it. But there are a few things we can say.

- The opposite of *separated flow* is *attached flow*.
- The opposite of *turbulent flow* is *laminar flow*.
- Separated flow need not be very turbulent, nor vice versa.
- Laminar flow need not be attached, nor vice versa.
- Turbulence doesn't cause separation (and may help prevent it).

For more information, see e.g. reference 10.

18.4 Coanda Effect, etc.

The name *Coanda effect* is properly applied to any situation where a thin, high-speed jet of fluid meets a solid surface and follows the surface around a curve. Depending on the situation, one or more of several different physical processes might be involved in making the jet follow the surface.

As a pilot, you absolutely do *not* need to know about the Coanda effect or what causes it. Indeed, many professional aerodynamicists get along just fine without really understanding such things. The main purpose of this section is to dispel the notion that a normal wing produces lift "because" of some type of Coanda effect.

Using the Coanda effect to explain the operation of a normal wing makes about as much sense as using bowling to explain walking. To be sure, bowling and walking use some of the same muscle groups, and both at some level depend on Newton's laws, but if you don't already know how to walk you won't learn much by considering the additional complexity of the bowling situation. Key elements of the bowling scenario are not present during ordinary walking.

18.4.1 Tissue-paper Demonstration

You can demonstrate one type of Coanda effect for yourself using a piece of paper. Limp paper, such as tissue paper, works better than stiff paper. Drape the paper over your fingers, and then blow horizontally, as shown in the following figures.[6]

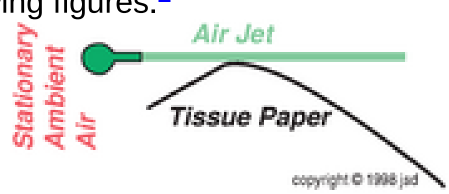

Figure 18.1: Tissue Paper; No Coanda Effect

If the jet passes just above the paper, as shown in figure 18.1, nothing very interesting happens. The jet just keeps on going. The paper is undisturbed.

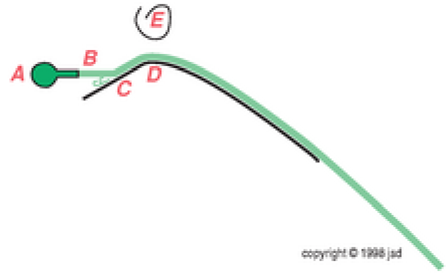

Figure 18.2: Tissue Paper; Coanda Effect

On the other hand, if the jet actually hits the paper as shown at point *C* in figure 18.2, the downstream part of the paper will rise up. This is because the air follows the curved surface; as it does so, it creates enough low pressure to lift the weight of the paper.

The air in your lungs, at point *A*, is at a pressure somewhat above atmospheric. At point *B*, after emerging from the nozzle, the air in the jet is at atmospheric pressure.

As discussed in section 3.3, the fact that the fluid follows a curved path proves that there is a force on it. This force must be due to a pressure difference. In this case, the pressure on the lower edge of the jet (where it follows the curve of the tissue paper near point *D*) is less than atmospheric, while the pressure on the upper edge of the jet (near point *E*) remains more-or-less atmospheric. This pressure difference pulls down on the jet, making it curve. By the same token it also pulls up on the paper, creating lift.

People who only half-understand Bernoulli's principle will be surprised to hear that the jet leaves the nozzle at high speed at atmospheric pressure. It's true, though. In particular, the crude statement that "high velocity means low pressure" is an oversimplification that cannot be used in this situation. The correct basis of Bernoulli's principle is that *for a particular parcel of air* the mechanical energy density (pressure plus kinetic energy per unit volume) remains more-or-less constant. If you want to compare two *different* parcels of air, you'd better make sure that they started out with the same mechanical energy. In this case, the air in the jet leaves the nozzle with a higher mechanical energy than the ambient air. Your lung-muscles are the source of the extra energy.

When this high-velocity, atmospheric-pressure air smacks into the paper at point *C*, it actually creates above-atmospheric pressure there. Indeed, we can use the streamline-curvature argument again: if the air turns a sharp corner, there must be a very large pressure difference.

In order to make this sharp turn, the air needs something to push against. A good bit of the required momentum comes from the air that splatters backward, as suggested by the squiggles in the figure, just below and upstream of the point of contact. This process is extremely messy. It is much more complicated than anything that happens near a wing in normal flight. To visualize this splatter, blow a jet of air onto a dusty surface.[7] Even if you blow at a very low angle, some of the dust particles blow away in the direction *opposite* to the main flow.

18.4.2 Blowing the Boundary Layer

Since we saw in section 18.3 that de-energizing the boundary layer is bad, you might think adding energy to the boundary layer should be good... and indeed it is. One way of doing so uses vortex generators, as discussed in section 18.3. Figure 18.3 shows an even more direct approach.

- We use a pump to create a supply of air at very high pressure.

- The air comes out a nozzle. The result is a jet of high-velocity air at the same pressure as the local air.[8]
- The jet shoots out of a slot in the top of the wing, adding energy to the boundary layer at a place where this could be very helpful.

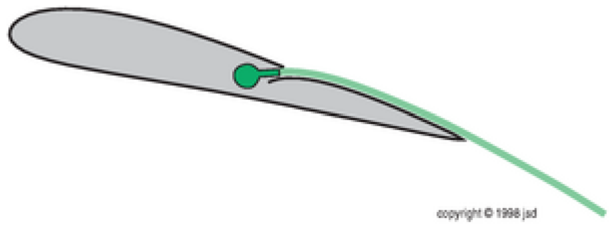

Figure 18.3: Blowing the Boundary Layer

Once again, the Coanda effect cannot explain how the wing works; you have to understand how the wing works before you consider the added complexity of the blower.

In this case we expect one spectacular added complexity, namely *curvature-enhanced turbulent mixing*. This phenomenon will not be discussed in this book, except to say that it does not occur near a normal wing, while it is likely to be quite significant in the situation shown in figure 18.3.

Curving flows with lots of shear can be put to a number of other fascinating uses, but a discussion is beyond the scope of this book. See reference 22.

18.4.3 Teaspoon Demonstration

Another example uses a jet of water following a curved surface. You can easily perform the following experiment: let a thin stream of water come out of the kitchen faucet. Then touch the left side of the stream with the convex back side of a spoon. The stream will not be pushed to the right, but instead will follow the curve of the spoon and be pulled to the left. The stream can be deflected by quite a large amount. In accordance with Newton's third law of motion, the spoon will be pulled to the right.

I don't understand everything I know about this situation, but it is safe to say the following:

1. This jet of water-in-air differs in fundamental ways from the air-in-air situation described above.
2. This effect has practically nothing to do with the way a normal wing produces lift.

To convince yourself of these facts, it helps to have a higher velocity and/or a larger diameter than you can conveniently get from a kitchen faucet. A garden hose will give you a bigger diameter, and if you add a nozzle you can get a higher velocity. You can easily observe:

- The amount of lift[9] you can produce is pathetically small, compared to the dynamic pressure and area of the water jet.
- The lift-to-drag ratio is terrible. Indeed this makes it very hard to measure the lift; if you get the angle slightly wrong you will inadvertently measure a drag component instead.
- The water spreads out when it hits the surface, making a thin coating over a wide area of the surface. This is in marked contrast to what happens in the air-in-air jet, as you can demonstrate by placing thin strips of tissue paper side by side. You can easily blow on one strip and lift it without disturbing its neighbors.

- Some of the spreading layer flows backwards, ahead of the point of contact of the jet, corresponding to a negative amount of upwash. This is grossly different from what happens near a real wing.
- The effect does not depend on curvature-enhanced turbulent mixing with the ambient air. This is quite unlike what happens in a real airplane with boundary-layer blowing.

It appears that surface tension plays two very important roles:

1. At the water/air interface it prevents mixing of the air and water.
2. At the water/wing interface it plays a dominant role in making the water stick to the surface.

In both respects this is quite unlike the air-in-air jet, where the air/wing surface tension has no effect and there is no such thing as air/air surface tension.

To convince yourself of this: Take a thin sheet of plastic. Get it wet on both sides, and drape it over a cylinder. You will not be able to lift it off the cylinder using a tangential water jet. The surface tension holding the wet plastic to the cylinder is just as strong as the tension between the plastic and the jet. In contrast, when the same piece of plastic has air on both sides, you can easily lift it off the cylinder using an air jet.

18.4.4 Fallacious Model of Lift Production

You may have heard stories saying that the Coanda effect explains how a wing works. Alas, these are just fairy tales. They are worse than useless.

1. For starters, these fairy tales often claim that blowing on tissue paper (as described just above) proves that “high velocity means low pressure” which is absolutely not what is being demonstrated. The high-velocity air coming out of your mouth is at atmospheric pressure. If you blow across the top of a *flat* piece of paper, it will not rise, no matter what you do. There is no low pressure in the jet (unless and until it gets pulled around a corner). Therefore the Coanda stories give a wrong explanation of normal wings and basic aerodynamics. And by the way, such stories cannot even begin to explain the operation of flat wings --- yet we have seen in section 3.10.1 that a barn door doesn’t behave very differently from other airfoils.
2. The Coanda-like notion of airflow following a curved surface cannot possibly explain why there is upwash in front of the wing. In figure 18.2 there must be a stagnation point on the upper surface of the paper near point *C*. This is completely different from the situation near a normal wing, where the stagnation line must be somewhere below the leading edge of the wing. Upwash is important, since it contributes to lift while creating a *negative* amount of induced drag. A further consequence, by the way, is that these Coanda-like stories be reconciled with the known behavior of stall-warning devices, as discussed in section 3.5.
3. As mentioned above, the distribution of velocities necessary to create curvature-enhanced turbulent mixing is produced by a high-speed jet but is not produced by a normal wing.
4. Sometimes the fairy tales say that the jet “sticks” to the surface because of viscosity. This implies that if the viscosity of the fluid changes, the amount of lift an airfoil produces should change in proportion. In fact, though, the amount of lift produced by a real wing is independent of viscosity over a wide range. Also, many of the processes responsible for the real Coanda effect require the production of turbulence, so they only work if the viscosity is sufficiently *low*.[10]
5. In the real Coanda effect, we know where the high-velocity air comes from. It comes from a nozzle. Upstream of the nozzle is a pump (or a rocket engine, or some other device) to supply the necessary energy. The jet makes high-velocity air above the wing, not below, because

that's where we aim the nozzle. An ordinary wing is completely different. It is wonderfully effective at creating high-velocity air above itself, without nozzles, without pumps, and without the huge energy[11] budget that pumps etc. would require.

6. The fairy tales generally neglect the fact that the wing speeds up the air in its vicinity, and just assume that the relative wind meets the wing at the free-stream velocity and follows the curve in a Coanda-like way. As a consequence, they miscalculate the pressure gradients by a factor of ten or so.
7. Finally, in the real Coanda effect we know how big the jet is. Its initial size is determined by the nozzle. The amount of mixing depends on the speed of the jet, the speed of the ambient air, the curvature of the surface, and other known quantities.

   In contrast, (a) the typical fairy tales imply that the entire flow pattern of a normal wing can be explained by mentioning the magic words "Coanda effect", yet (b) they cannot explain how thick a chunk of air is deflected by the wing. One inch? Six inches? A chord-length? A span-length? Some amount proportional to the viscosity of the air? It would be very hard to calculate how much ... nonsensical things are often rather hard to calculate.
8. Even when there is a real Coanda effect, as in figure 18.3, it is just a small part of what is happening in the vicinity of the wing.

18.4.5 Fact versus Fallacy

Don't let anybody tell you that squirting a spoon or blowing on tissue paper is a good model of how a wing works.

If you want to "get the feel" of lift production, the obvious methods are the best. These include holding a model airfoil downstream of a household fan, or sticking it out of a car window. (You can think of this as the first step toward a home-made wind tunnel.) Among other things, you will quickly discover that precise control of angle of attack is important.

18.5 Spin Entry

Case 1: In normal flight, rolling motions are very heavily damped, as discussed in section 5.4. Even though the static stability of the bank *angle* is small or even negative, you cannot get a large roll *rate* without a large roll-inducing torque; when you take away the torque the roll rate goes away.

Case 2: Near the critical angle of attack, the roll damping goes away. Suppose you start the aircraft rolling to the right. The roll rate will just continue all by itself. The right wing will be stalled (beyond max lift angle of attack) and the left wing will be unstalled (below max lift angle of attack).

Case 3: At a sufficiently high initial angle of attack (somewhat greater than the critical angle of attack), the roll will not just continue but accelerate, all by itself. This is an example of the *departure*[12] that constitutes the beginning of a snap roll or spin. The resulting undamped rolling motion is called *autorotation*.

At a high enough angle of attack, the ailerons lose effectiveness, and at some point they start working in reverse.[13] Figure 18.4 shows how this reversal occurs. Suppose you deflect the ailerons to the left. This raises the angle of attack at the right wingtip and lowers it at the left wingtip. Normally, this would increase the lift on the right wing (and lower it on the left), creating a rolling moment toward the left. Near the critical angle of attack, though (as seen in the left panel of the figure), raising or lowering the angle of attack has about the same effect on the coefficient of lift, so no rolling moment is produced (for now, at least).

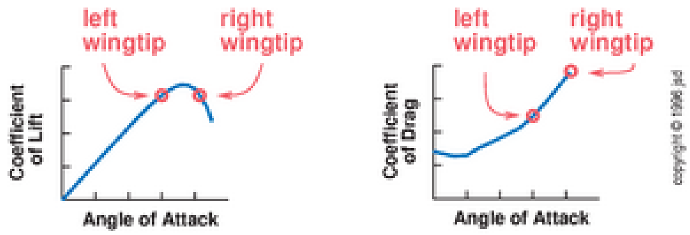

Figure 18.4: Lift and Drag at Departure

We see that at this angle of attack, anything that creates a rolling moment will cause the aircraft to roll like crazy, and indeed to keep accelerating in the roll-wise direction. There will be no natural roll damping, and you will be unable to oppose the roll with the ailerons.

There are two main ways of provoking a spin at this point:

- Suppose the airplane is in a steady slip to the left. That is, you are steadily pushing on the right rudder pedal. Then the slip/roll coupling (as discussed in section 9.1 and section 9.2) will cause it to spin to the right.
- Suppose the airplane is not in much of a slip, but you suddenly cause it to yaw to the right. The left wingtip will temporarily be moving faster, and the right wingtip will temporarily be moving slower. This difference in airspeeds will create a difference in lift, causing a spin to the right. The initial yawing motion could come from a sudden application of rudder, or from adverse yaw, or whatever. Note that in the right panel of figure 18.4, the aileron deflection has a tremendous effect on the drag. This means that ailerons deflected to the left cause a yaw to the right, which in turn provokes a roll to the right, which is the opposite of what ailerons normally do.

18.6 Types of Spin

18.6.1 Spin Modes

The word “spin” can be used in several different ways, which we will discuss below. The spin family tree includes:

- “departure”, i.e. onset of undamped rolling;
- incipient spin --- i.e. one that has just gotten started; or
- well-developed spin, which could be
  - a steep spin, or
  - a flat spin.

Figure 18.5 shows an airplane in a steady spin. You can see that the direction of flight has two components: a vertical component (down, parallel to the spin axis) and a horizontal component (forward and around).

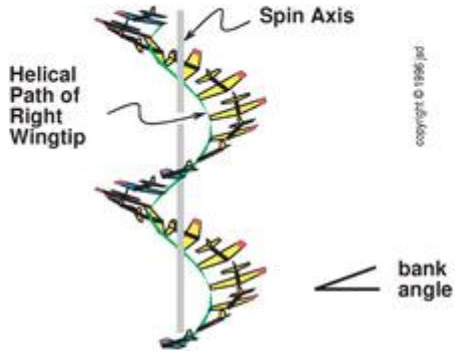

[Figure 18.5](#): Airplane in a Steady Spin

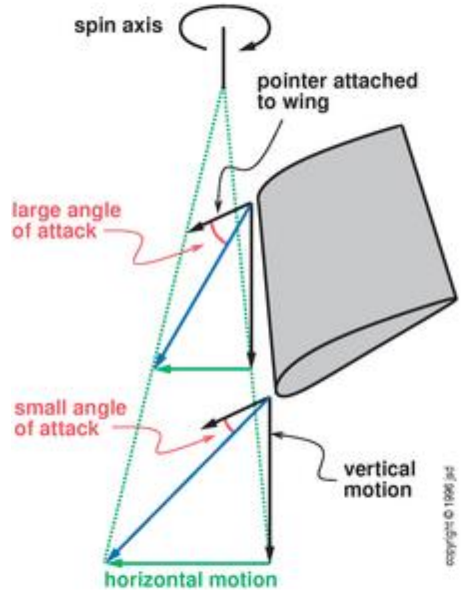

[Figure 18.6](#): Steep Spin --- Geometry

[Figure 18.6](#) is a close-up of a wing in a steep spin. We have welded a pointer to each wingtip, indicating the direction from which the relative wind would come if the wing were producing zero lift; we call this the Zero-Lift Direction (ZLD). (For a symmetric airfoil, the ZLD would be aligned with the chord line of the wing.) Remember that the angle between the direction of flight and the ZLD pointer is the angle of attack.

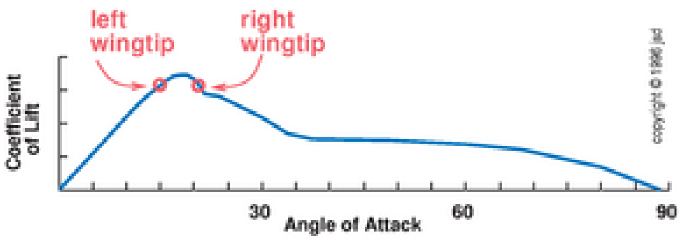

[Figure 18.7](#): Steep Spin --- Coefficient of Lift

In this situation, both wingtips have the same vertical speed, but they have significantly different horizontal speeds --- because of the rotation. Consequently they have different directions of flight, as shown in the figure. This in turn means that the two wingtips have significantly different angles of attack, as shown in [figure 18.7](#). The two wings are producing equal amounts of lift, even though one is in the stalled regime and one in the unstalled regime.

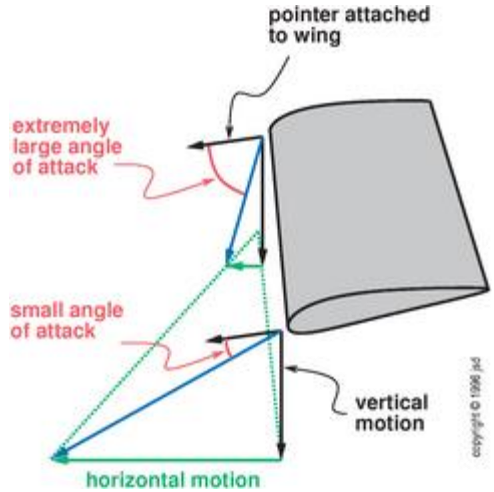

Figure 18.8: Flat Spin --- Geometry

Figure 18.8 shows another spin mode. This time the rotation rate is higher than previously. The spin axis is very close to the right wingtip. The outside wing is still unstalled, while the inside wing is very, very deeply stalled, as shown in figure 18.9.

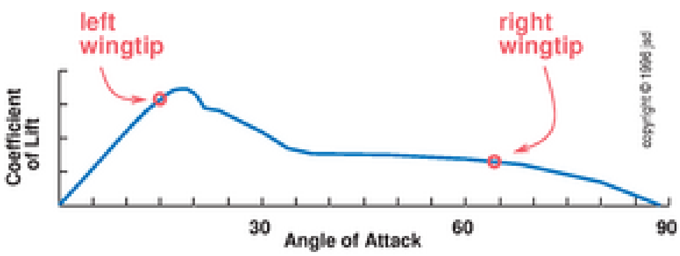

Figure 18.9: Flat Spin --- Coefficient of Lift

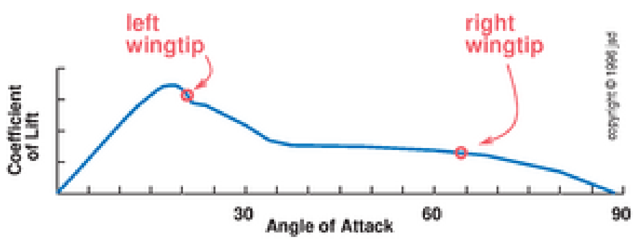

Figure 18.10: Doubly Stalled Flat Spin --- Coefficient of Lift

Figure 18.10 shows yet another possible spin mode. In this case, the outside wing is stalled, while the inside wing is, of course, much more deeply stalled. Whether this spin mode, or the one shown in figure 18.9 (or both or neither) is stable depends on dozens of details (aircraft shape, weight distribution, et cetera).

There is a common misconception that in a spin, one wing is stalled and the other wing is always unstalled. This is true sometimes but not always, especially not for flat spins.

It would be better to define “spin” as follows:

> In a spin, at least one wing is stalled, and the two wings are operating at very different angles of attack.

18.6.2 Samaras, Flat Spins, and Centrifugal Force

A samara is a winged seed. Maples are a particularly well known and interesting example.

Maple samaras have only one wing, with the seed all the way at one end. Its mode of flight is analogous to an airplane in a flat spin. In an airplane, the inside wing is deeply stalled, while in the samara the inside wing is missing entirely.

In a non-spinning airplane, if one wing were producing more lift than the other, that wing would rise. So the question is, why is a flat spin stable? Why doesn't the outside wing continue to roll to ever-higher bank angles? The secret is *centrifugal force*.[14] Suppose you hold a broomstick by one end while you spin around and around; the broomstick will be centrifuged outward and toward the horizontal.

In an airplane spinning about a vertical axis, the high (outside) wing will be centrifuged outward and downward (toward the horizontal), while the low (inside) wing will be centrifuged outward and upward (again toward the horizontal). In a steady flat spin, these centrifugal forces cancel the rolling moment that results from one wing producing a lot more lift than the other. This is the only example I can imagine where an airplane is in a steady regime of flight but one wing is producing more lift than the other.

As discussed in reference 19, an aircraft with a lot of mass in the wings will have a stronger centrifugal force than one with all the mass near the centerline of the fuselage. In particular, an aircraft with one pilot and lots of fuel in the wing tanks could have completely different spin characteristics than the same aircraft with two pilots and less fuel aboard.

18.6.3 NASA Spin Studies

In the 1970s, NASA conduced a series of experiments on the spin behavior of general-aviation aircraft; see reference 21 and reference 20 and other papers cited therein. They noted that there was "considerable confusion" surrounding the definition of steep versus flat spin modes, and offered the classification scheme shown in table 18.1.

spin mode	Steep	Mod'ly Steep	Mod'ly Flat	Flat
angle of attack	20 to 30	30 to 45	45 to 65	65 to 90
nose attitude	extreme nose-down		less nose-down	
rate of descent	very rapid		less rapid	
rate of roll	extreme		moderate	
rate of yaw	moderate		extreme	
wingtip-to-wingtip difference in angle of attack	modest		large	
nose-to-tail difference in slip	large		large	

Table 18.1: Spin Mode Classification

The angle of attack that appears in this table is measured in the aircraft's plane of symmetry; the actual angle of attack at other positions along the span will depend on position.

The NASA tests demonstrated that general aviation aircraft not approved for intentional spins commonly had unrecoverable flat spin modes.

18.6.4 Effects of Changes in Orientation of Spin

In all cases NASA studied, the flat spin had a *faster* rate of rotation (and a slower rate of descent) than the steep spin. Meanwhile, reference 6 reports experiments in which the flatter pitch attitudes were associated with the *slower* rates of rotation. This is not a contradiction, because the latter dealt with an unsteady spin (with frequent changes in pitch attitude), rather than a fully stabilized flat spin. A *sudden change* to a flatter pitch attitude will cause a *temporary* reduction in spin rate, for the following reason.

In any system where angular momentum is not changing, the system will spin faster when the mass is more concentrated near the axis of rotation (i.e. lesser moment of inertia). The general concept is discussed in section 19.8. By the same token, if the mass of a spinning object is redistributed farther from the axis, the rotation will slow down.

When the spinning airplane pitches up into a flatter attitude, whatever mass is in the nose and tail will move farther from the axis of rotation. Angular momentum doesn't change in the short run, so the rotation will slow down in the short run.

In the longer run --- in a steady flat spin --- the aerodynamics of the spin will pump more angular momentum into the system, and the rotation rate will increase quite a lot. The rotation rate of the established flat spin is typically twice that of the steep spin.

Recovering from an established flat spin requires forcing the nose down. This brings the mass in the nose and tail closer to the axis of rotation. Once again using the principle of conservation of angular momentum, you can see that the rotation rate will increase (at least in the short run) as you do so --- which can be disconcerting.

18.7 Recovering from a Spin

If you find yourself in an unusual turning, descending situation, the first thing to do is decide whether you are in a spiral dive or in a spin. In a spiral dive, the airspeed will be high and increasing; in a spin the airspeed will be low. You should be able to *hear* the difference. Also, the rate of rotation in a spiral is much less; the high speed means the airplane has lots of momentum and can't turn on a dime. In a spin, the aircraft will be turning a couple hundred degrees per second.

To get out of a spin,[15] follow the spin-recovery procedures given in the Pilot's Operating Handbook for your airplane. The literature is full of home-brew spin recovery procedures that probably work most of the time in most airplanes, but if you want a procedure that works for sure, follow the handbook for your airplane.

For typical airplanes, the spin recovery procedure contains the following items:

- Retard the throttle to idle
- Retract the flaps
- Neutralize the ailerons
- Apply full rudder in the direction opposing the spin
- Briskly move the yoke to select zero angle of attack.

Now let's discuss each of these items in a little more detail.

Retarding the throttle is a moderately good idea for a couple of reasons. For one thing (especially if you have a fixed-pitch prop) it keeps the engine from overspeeding during the later stages of the spin recovery. More importantly, gyroscopic precession of the rotating engine and propeller can hold the nose up, flattening the spin and interfering with the recovery (depending on the direction of spin).

Propwash might increase the effectiveness of the horizontal tail and therefore *assist* in the spin recovery, but (especially in a flat spin) the propwash could be blown somewhere else by the abnormal airflow --- so you may not be able to count on this.

Retracting the flaps is a moderately good idea because you might exceed the “max flaps-extended speed” if you mishandle the later stages of the spin recovery and you don’t want to damage the flaps.

Retracting the flaps may help with the spin recovery itself. Recall from section 5.4.3 that the flaps effectively increase the washout of the wings. Washout ensures that the airplane will stall before it runs out of roll damping. (This produces a nice straight-ahead stall.) In the spin, though, when you have lost all vertical damping and roll damping, the washout doesn’t help. The early stages of spin recovery are not like the early stages of stall entry.

Neutralizing the ailerons is usually a good idea for the simple reason that it is hard to think of anything better to do with them. Deflecting the ailerons effectively increases the angle of attack of one wingtip and decreases the angle of attack of the other wingtip. In a spin, the part of the wing where the ailerons are may (or may not) be in the stalled regime --- so deflecting the ailerons to the left may (or may not) produce a paradoxical rolling moment to the right.

Depressing the rudder to oppose the spin is obviously a good thing to do.

Finally, you want to move the yoke to select zero angle of attack. In typical trainers, this means shoving the yoke all the way forward, but in other aircraft, especially aerobatic aircraft, all the way forward might select a large *negative* angle of attack. Shoving the yoke all the way forward in such a plane would likely convert the spin to an inverted spin --- hardly an improvement. This is just one example of why you want to know and follow the spin recovery procedure for your specific airplane.

The relative significance of the rudder compared to the flippers in breaking the spin depends radically on the design of the airplane, the loading of the airplane, and on the spin mode, as discussed in reference 19.

In normal non-spinning flight, you should apply smooth pressures to the controls. Spin recovery is the exception: it calls for brisk, mechanical *motions* of the controls, almost without regard to the pressures involved.

If you get into a spin in instrument conditions, you should rely primarily on the airspeed indicator and the rate-of-turn gyro. The inclinometer ball cannot be trusted; it is likely to be centrifuged away from the center of the airplane --- giving an indication that depends on where the instrument is installed, telling you nothing about the direction of spin. The artificial horizon (attitude indicator) cannot be trusted since it may have tumbled. The rate-of-turn gyro is more trustworthy, since it cannot possibly suffer from gimbal lock, since has no gimbals, since it is a rate gyro (not a free gyro).

Recovery from a so-called incipient spin (one that has just gotten started) is easier than from a well-developed spin. Normal-category single-engine[16] certification requirements say that an airplane must be able to recover from a one-turn spin (or a 3-second spin, whichever takes longer) in not more than one additional turn. If you let the spin go on for several turns, you might progress from a steep spin to a flat spin. Recovery could take a lot longer --- if it is possible at all.

If you load the airplane beyond the aft limit of the weight and balance envelope, even the incipient spin may be unrecoverable; see section 6.1.1. Imperfect repairs to the wing, or slack in the control cables, could also impede spin recovery.

Finally, the spin is yet another reason why it is *NOT SAFE* to think of the yoke as simply the up/down control.[17] In a spin you have a low airspeed and a high rate of descent. If you think of the yoke as the up/down control, you will be tempted to pull back on the yoke, which is exactly the wrong thing to do. On the other hand, if you think of the yoke as (primarily) the fast/slow control, you will realize that you need to push forward on the yoke, to solve the airspeed problem.

18.8 Don't Mess With Spins

It is quite impressive how well a samara works. A maple seed descends very slowly, riding the wind much better than a parachute of similar size and weight ever could. Flat spins can be extremely stable; a wing by itself *loves* to spin. That's why spins (and flat spins in particular) are so dangerous: it takes a lot of rudder force to persuade a wing to stop spinning.

Spins are extremely complex. Even designers and top-notch test pilots are routinely surprised by the behavior of spinning airplanes. Spin-test airplanes are equipped with cannon-powered spin-recovery parachutes on the airframe, and quick-release doors in view of the distinct possibility that the pilot will have to bail out. Tests are conducted at high altitude over absolutely unpopulated areas. Therefore please don't experiment with spinning a plane except exactly as approved by the manufacturer --- one unrecoverable spin mode can ruin your whole day.

1
: This happens naturally on a rectangular wing; it can be enhanced by washout and other designers' tricks.

2
: An even more direct method of adding energy to the boundary layer uses a jet of high-velocity air, as discussed in section 18.4.2.

3
: Of course, the vortex generators contribute *indirectly* to maintaining the health of the big bound vortex, since they help maintain attachment, thereby allowing the wing to create lots of circulation.

4
: See reference 10 for a nice discussion of golf balls, cricket balls, and boundary layers in general.

5
: ... just as having lots of water is the cause, but not the definition, of drowning --- you can get very wet without drowning.

6
: You can blow directly from your lips, but it's better to use a flexible straw or a thin piece of tubing, so that you can get a better view of what's happening. If you put a nozzle at the end of the tube, the jet will keep its shape better.

7
: Finely-ground black pepper is a convenient source of suitable dust.

8
: This won't be exactly atmospheric, since the local pressure has been affected by the wing.

9

Remember, lift is the force perpendicular to the flow and perpendicular to the surface.

10

Indeed, as long as the viscosity is not *exactly* zero, the smaller the viscosity, the greater the turbulence.

11

Of course a *small* amount energy is required, because of skin friction and induced drag, but this is very small, out of all proportion to the energy that the air parcel transforms internally, from kinetic energy to pressure and back again.

12

This refers to “departure from normal flight”. It has nothing to do with takeoff or with a “departure stall” which merely refers to a stall in the takeoff configuration.

13

Under present-day certification rules, the ailerons are required to work normally up to at least stalling angle of attack. However, some older airplanes were built under older rules. These planes, including many aerobatic aircraft, have much less washout, and therefore lose aileron effectiveness earlier. All planes lose effectiveness eventually. For simplicity, this section ignores washout.

14

See section 19.4 for a discussion of the nature of centrifugal fields.

15

Recovery from a spiral dive is discussed in section 6.2.4.

16

Multi-engine aircraft are not required to be recoverable from any sort of spin, incipient or otherwise.

17

This point is discussed in chapter 7.