Flightlab Ground School
9. Rolling Dynamics

The Aircraft in Roll

The dynamics of an aircraft in roll are surprisingly complex, given the apparent simplicity of the maneuver. Of course, one person's complexity is just another person getting started. At the U.S. Navy Test Pilot School, for instance, "The classic roll mode is a heavily damped, first order, non-oscillatory mode of motion manifested in a build-up of roll rate to a steady state value for a given lateral control input."[1] Well, ok, that sounds right.

Our *Maneuvers and Flight Notes* training guide describes piloting technique during aerobatic or unusual attitude rolling maneuvers. Here the emphasis is on the general characteristics of aircraft response.

A roll starts with the creation of an asymmetric lift distribution along the wingspan. In the case of aileron roll control, deflecting an aileron down increases wing camber and coefficient of lift; raising the opposite aileron reduces camber and coefficient of lift. The resulting spanwise asymmetry produces a rolling moment.

As the aircraft begins to roll in response to the moment produced by the ailerons, the lift distribution again begins to change. The rolling motion induces an angle of attack increase on the down-going wing, and an angle of attack decrease on the up-going wing (Figure 1). This creates an opposing aerodynamic moment, called roll damping (or rolling moment due to roll rate, C_{lp}). Roll damping increases with roll rate (and varies with other factors we'll get to). When the damping moment produced by the roll rate rises to equal the opposing moment produced by the ailerons, the roll rate becomes constant.

In Figure 1 you can see that as the airplane rolls, the lift vector tilts to accommodate itself to the new direction of the relative wind, creating new vectors of thrust and drag. As a result, the rolling motion produces adverse yaw all by itself, a yawing moment that goes away when the roll stops. This yaw due to roll rate, C_{np}, is in addition to the adverse yaw created by the displaced ailerons, and increases with coefficient of lift. Depending on wing planform, at aspect ratios above 6 or so, adverse yaw due to roll rate actually becomes more significant than that due to aileron deflection.

Figure 1
Roll Damping, C_{lp}
Yaw due to Roll Rate, C_{np}

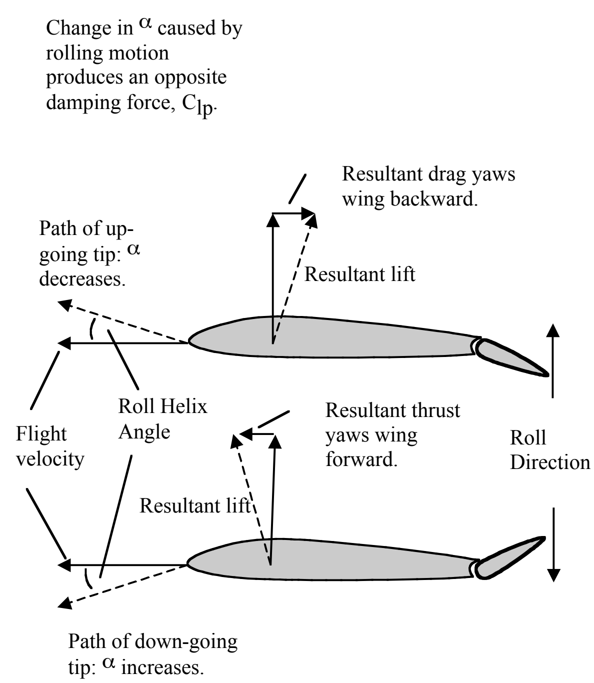

[1] *U.S. Naval Test Pilot School Flight Test Manual: Fixed Wing Stability and Control,* USNTPS-FTM-No. 103, 1997. p. 5.55.

Adverse Yaw, $C_{n_{\delta_a}}$

Aerodynamic coupling effects keep rolling from being a one-degree-of-freedom proposition. Rolling moments come with yawing moments attached, and those yawing moments affect roll behavior.

Induced drag increases when an aileron goes down, decreases when an aileron goes up. The result is usually an adverse yawing moment, opposite the direction of roll. In the absence of a sufficient counteracting yaw moment—supplied in part by the aircraft's inherent directional stability, in part by aileron design, and in the remainder by coordinated rudder—the aircraft will begin to sideslip. The velocity vector will shift from the plane of symmetry toward the roll direction if too little coordinating rudder is applied, and shift opposite the roll direction if the rudder gets too emphatic an in-turn boot. In a perfectly coordinated, ball-centered roll and turn, with adverse yaw properly countered by rudder deflection, the "instantaneous" velocity vector remains on the plane of symmetry, as Figure 2 describes.

The rudder deflection necessary to handle adverse yaw depends on the ratio of yaw moment to roll moment the ailerons produce. While the ratio is basically a function of the aileron system design, it increases with coefficient of lift, C_L. This means that as airspeed goes down, the need for rudder coordination becomes greater. The nature of induced drag rise at high angles of attack is the major reason, since induced drag increases as the square of the coefficient of lift. As the drag curve becomes steeper, a given aileron deflection produces a greater difference in induced drag across the span, and the yaw/roll ratio increases. Differential or Frise ailerons, initially designed to reduce aileron *forces*, also reduce adverse yaw by increasing the drag of the up-going aileron.

Another factor is the reduction in directional stability caused by the disrupted fuselage wake at angles of attack approaching stall. Because energy is removed from the free stream, more rudder deflection is needed as weathercock stability goes down in the tired-out air.

Configuration is also important. Partial-span flaps cause an aircraft to fly at a more nose-down angle for a given overall coefficient of lift. As a result, the aileron portion of the wing experiences a relative washout (leading edge down) and generates a lower local coefficient of lift than when the flaps are up. That lower local coefficient translates into less adverse yaw. Flaps also reduce dihedral effect, so the sideslip that does occur has less effect on roll.

Spoilers

Spoilers are generally thought to produce proverse, roll-direction yaw, but they can cause adverse yaw. Spoilers increase profile drag. They also decrease induced drag, since they kill lift. When the increase in profile drag predominates, as it does at high speed, spoilers can generate proverse yaw. At low speeds, when induced drag is more important, they can generate adverse yaw, since the induced drag on the wing going down, the one with the deflected spoiler killing lift, is less than on the wing going up, where the spoiler remains tucked away.

Spoilers are useful in situations when aeroelastic aileron reversal could have been a problem (B-52, and just about all of the swept-wing, high aspect ratio transports that followed), or when it's necessary to extend the wing area available for flaps. They have an advantage over ailerons of producing powerful rolling moments at high angles of attack, but the disadvantage of lesser moments at low angles. The classic problem with spoilers is a possible nonlinear response as their location moves forward on the wing. Small deflections may generate no roll, or even a temporary reversed roll response. (As the spoiler first rises, the tripped airflow can reattach to the wing. This results in an effective increase in camber and therefore in lift. Spoiler movement has to be nonlinear with control wheel or stick movement—so that the spoilers can quickly pop up high enough to defeat any tendency for the airflow to reattach.) Many designs use spoilers and ailerons in combination, with the ailerons providing both rolling moment and control feel, and a possible way of overcoming nonlinear spoiler response.

If you're stuck in coach, the most entertaining window seat on a Boeing is just back of the trailing edge, where you can watch the slot-lip spoilers being used for bank control when the flaps are extended. When the spoilers rise, the slots above the flaps open up. The change in pressure pattern reduces the lift gained from flap deployment and causes the aircraft to roll. Deployed symmetrically, the spoilers provide aerodynamic braking.

A pilot who looks up to find himself flying inverted in a spoiler-equipped aircraft has a quandary, since spoilers become more effective at higher coefficients of lift (higher α). Does that mean the pilot should pull while inverted, to increase roll response, at the risk of altitude loss and airspeed gain from the resulting nose-low attitude? It's hard to find someone with a satisfactory answer. Aggressive use of the rudder might be warranted to help roll the aircraft using dihedral effect and roll due to yaw rate.

Turn Coordination

Usually, we roll in order to turn. Steep turns deserve to be regarded as unusual attitudes, not just because of the high bank angles but also the high-gain response those angles require when you're trying to be perfect. The concentration level goes up when you fly a steep (45-degree plus), coordinated turn, while holding altitude.

We've defined coordination during roll in terms of keeping the velocity vector on the plane of symmetry. While the ailerons are deflected, that means using rudder to correct for adverse yaw and yaw due to roll rate—companion phenomena whose relative magnitudes can be difficult to figure out, but then you don't have to figure: just push the rudder to center the ball.

Once the bank is established, the ailerons move to the position necessary to maintain the bank angle. Rudder into the turn is often needed to counteract yaw damping, C_{nr}, caused by the fuselage and tail resisting the yaw rate and by the outside wing moving faster than the inside wing and producing more drag. That might be the predicament shown in the center aircraft at the bottom of Figure 3. An over-banking tendency requires aileron deflection against the turn—possibly causing proverse yaw (since the inside aileron is down), which of course modifies the rudder requirement.

The gyroscopic precession of the propeller creates a force parallel to the vertical turn axis. Therefore, precession causes the nose to move perpendicularly to the horizon, regardless of bank angle. For a clockwise propeller as seen from the rear, that means nose down turning right, nose up turning left. At high bank angles, when the aircraft's y-axis approaches alignment with the turn axis, more rudder deflection may be needed to counter precession. Step on the ball.

Banked turns are combinations of yaw and pitch (Figure 2). Coordination (keeping the velocity vector on the plane of symmetry) means establishing both the yaw rate and the pitch rate appropriate for the bank angle.

As bank angle increases, pitch rate becomes increasingly sensitive. Pitch rate controls load factor, and ***for a constant-altitude turn, the required load factor goes up exponentially with bank angle.*** Getting the pitch rate/load factor right at high bank angles is difficult because the load requirements change so rapidly with even small changes in bank angle. Because the load factor goes up exponentially, so do the stick forces, at least in aircraft with reversible elevator controls.

Figure 2
Relative Yaw and Pitch Rates. Constant-Altitude Turn

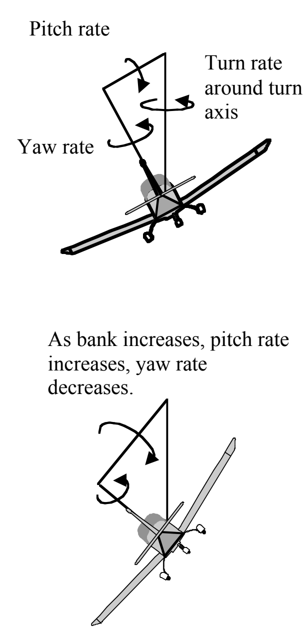

Figure 3
Coordinated Yaw Rate

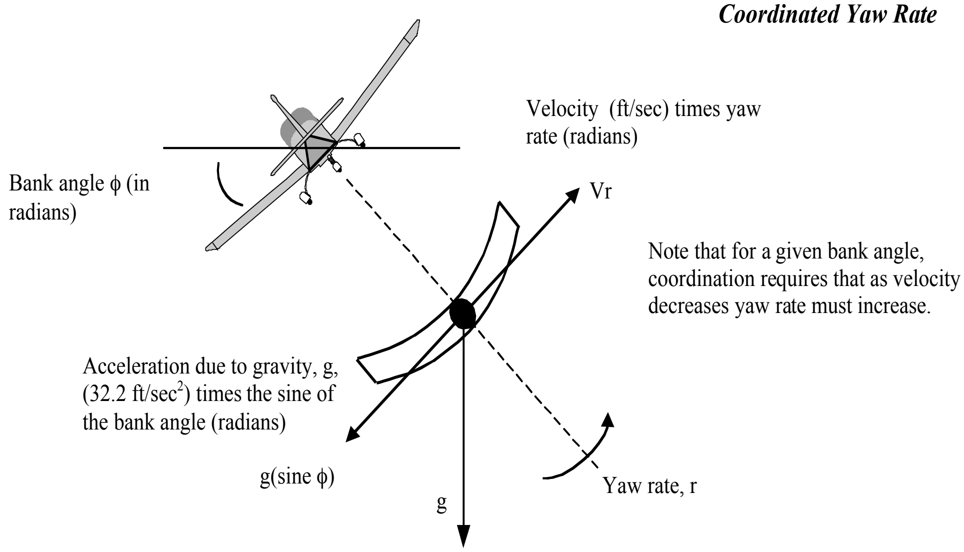

Ball stays centered when the acceleration toward the outside of the turn due to yaw rate and velocity equals the acceleration to the inside due to bank angle and gravity.

$$32.2\,(\text{sine}\ \phi) = Vr$$

"Stepping on the ball" controls yaw rate.

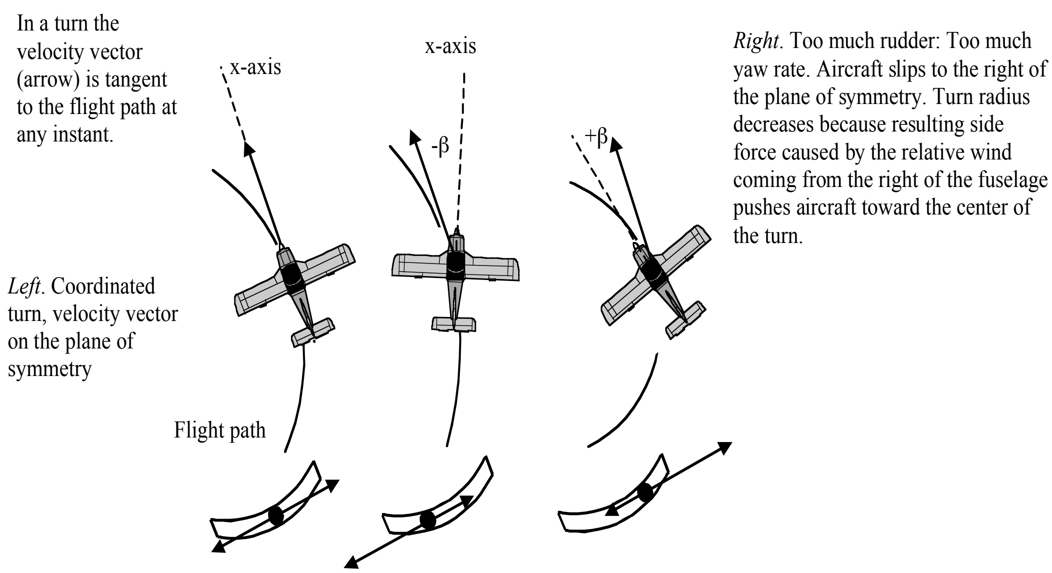

Center. Too little rudder: Too little yaw rate. Aircraft slips to the left of the plane of symmetry. Turn radius increases because resulting side force caused by the relative wind coming from the left of the fuselage pushes aircraft away from the center of the turn.

Figure 4
Roll Helix Angle

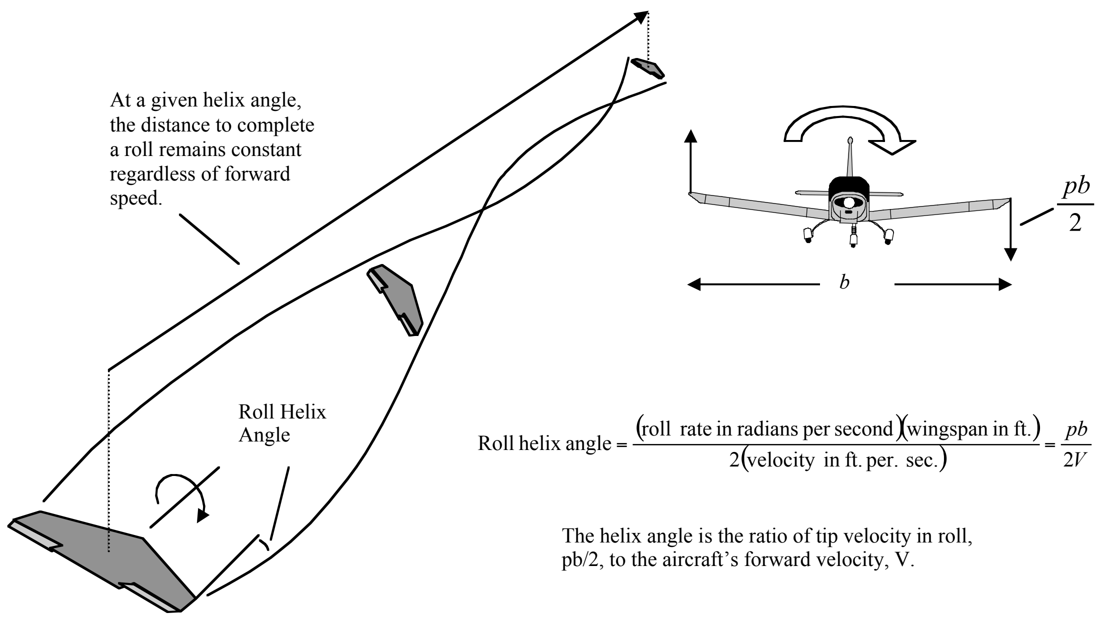

$$\text{Roll helix angle} = \frac{(\text{roll rate in radians per second})(\text{wingspan in ft.})}{2(\text{velocity in ft. per. sec.})} = \frac{pb}{2V}$$

The helix angle is the ratio of tip velocity in roll, pb/2, to the aircraft's forward velocity, V.

Roll Helix Angle

As an airplane rolls, its wingtips follow a helical path through the sky, like the shape of a stretched spring. You see the helix at airshows when the pilot fires up the tip smoke. The angle between the resultant flight path of the wingtip and flight path of the aircraft is called the *roll helix angle* (Figure 4). The roll helix angle increases with increasing aileron displacement. Maximum attainable helix angle depends on aircraft design and mission. But on a given aircraft a given aileron deflection always builds to a given roll helix angle.

For any aileron deflection (roll helix angle), coordinated roll rate, p, varies directly with true airspeed. "Coordinated" means there's no sideslip affecting the rate. For the statement to be true, there must also be no aero-elastic effects (wing bending caused by aileron deflection at high speeds).

As the aircraft's forward velocity increases, it will maintain the helix angle—by virtue of rolling faster. As a result, for a given helix angle an aircraft will complete a roll in the same forward distance traveled, regardless of airspeed. Slippage aside, the helix angle is like the pitch of a screw of given length. Regardless of rpm, it takes the same number of turns to fasten it down.

When you know the helix angle, wingspan, and aircraft velocity, you can solve the formula in Figure 4 for roll rate. The helix angle provides a way of establishing minimum acceptable roll rates for different aircraft (about 0.09 minimum for fighters; 0.07 for transports). But it doesn't account for roll inertia, so it's only a partial description of rolling character and is no longer used as a certification or acceptance measure.

Timed bank changes are used instead. The FAA's requirements are listed in FAR Parts 23.157 and 25.147 (e).

Figure 5
Time to Reach
Roll Rate

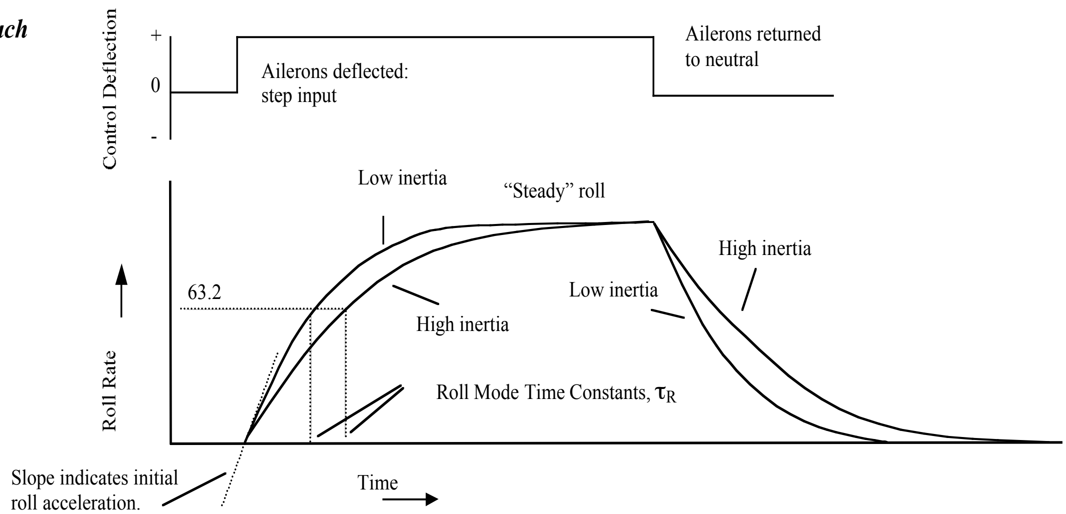

Roll Acceleration, Time Constant

Roll acceleration (and thus how quickly the aircraft reaches the final roll rate for a given aileron deflection) depends on how much rolling moment the ailerons produce versus the moment of inertia about the roll axis the ailerons have to overcome. We're assuming the necessary rudder coordination, and thus no inadvertent moments due to sideslip:

$$\frac{\text{Rolling moment due to aileron deflection}}{\text{Moment of inertia about the roll axis}} = \text{Rolling acceleration}$$

The moment of inertia about the roll axis, I_{xx}, depends on how mass is distributed in the aircraft. If there's lots of fuel in the wings, some engines hanging out there, and maybe full tip tanks, roll inertia will be higher than if most of the mass were confined to the aircraft's fuselage.

Figure 5 shows the effect of inertia on roll acceleration and deceleration. Span, wing planform, and the rolling moment due to the aileron "step" deflection are the same for both curves. Only the moment of inertia, I_{xx}, is different. Notice the difference in initial slope between the two curves, and how maximum roll acceleration (not rate) happens at the initial control input, before damping forces have a chance to build. Notice the difference in the time required for reaching the steady rate in the two inertia cases, and for the roll to stop when the ailerons are neutralized.

The roll mode time constant, τ_R, is the time it takes to achieve 63.2 percent of the final roll rate, in seconds, following step input. If you could maintain the initial acceleration (i.e. no damping to slow things down) the airplane would reach the target roll rate in only 63.2 percent of the time actually required. In theory, the rolling aircraft "remembers" this constant. Because of damping, it takes the same amount of time to achieve the next 63.2 percent of the final roll rate, then the next 63.2 percent after that, and so on. At least theoretically, the asymptotic curve keeps trying but never flattens out. After five time constants the roll rate will be at 99.5% of the final value. Close enough.

The greater the moment of inertia versus the aileron authority, the longer the time constant will be. Two aircraft may have the same maximum roll rate, but the one with the greater time constant will take longer getting to it *and longer to stop rolling when the ailerons are returned to neutral*. Therefore inertia characteristics, and not just maximum roll rates, need to be taken into account when comparing the rolling performance of different aircraft. In a given aircraft, roll rate can magically increase with increasing change of angle. Roll rate measured from a 45-degree bank to an opposite 45-degree bank may be lower than if measured from, say, 60 degrees to 60 degrees. That's because the time spent accelerating is a smaller fraction of the total.

It's difficult for a pilot to measure the time constant without special equipment, because there's no easy visual reference on roll initiation, and it's usually less than a second (and imperceptible for a fast roller). Since the time constant is the same to stop a roll as it is to start one, and the visual references are clearer, you can try to measure the time to stop instead. (It helps if there's minimal freeplay in the control system.) Begin with a 45-degree bank angle and make a coordinated roll to upright using an immediate, step aileron deflection (the amount of deflection doesn't matter, since the time constant is the same, see Figure 6). Hold that aileron input until the wings become level with the horizon, then instantly neutralize the controls and watch for any additional roll. 63.2 percent of the time it takes for any remaining roll to subside is the time constant due to roll inertia. You might think that the more aileron you use and the faster you roll to upright the greater your overshoot past level. But it doesn't work that way, because the faster you roll the greater the aerodynamic damping available to stop things when you neutralize the controls.

The idea that an aircraft that accelerates quickly into a roll can stop just as quickly takes aerobatic students by surprise. In performing 360-degree rolls, most will start out leading the recovery by too much and stopping short of wings-level. (That's normal, but students who over-roll and stop past wings level are obviously behind the aircraft.)

At a given altitude, roll mode time constant varies inversely with true airspeed (TAS). That's true to experience—the faster you go, the quicker you can accelerate into a roll. But at a constant TAS, the time constant increases with altitude. The air is less dense, and for a given TAS there's less dynamic pressure available to overcome inertia. There's also less damping as altitude increases, so it takes longer for the roll rate to settle.

Airplanes with noticeable time constants (as caused by high roll inertia and/or low roll damping, and limited aileron control power) require that pilots learn to "shape" their control inputs, first using large initial deflections for maximum acceleration and then reducing the deflection once the desired rate is achieved. Then they have to check the airplane's roll motion with opposite, anticipatory aileron inputs when capturing a bank angle or returning to level

Figure 6
Time to Reach
Roll Rate

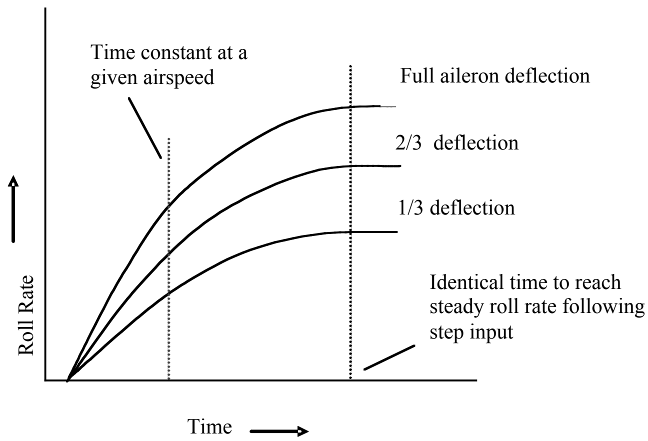

flight. The wheel or stick becomes an acceleration controller. Pilots can adapt, but the workload increases. Because of its huge x-axis roll inertia, the B-52 has this kind of response.

On the other hand, aircraft with short time constants tend to feel quick and responsive. Because they accelerate quickly, the stick becomes essentially a rate controller. Within bounds, that's what pilots prefer.

The FARs don't specify time constant requirements, but, for the military, MIL-STD-1797A 4.5.1.1 sets the maximum τ_R between 1 and 1.4 seconds, depending on aircraft mission and phase of flight.

In an aerobatic aircraft, because of the rapid roll acceleration, you can sometimes get the impression that the aircraft's steady roll rate is faster than it really is. We feel acceleration much more profoundly than rate, especially as passive recipients. Studies have shown that pilots tend to estimate ultimate roll rate based on initial acceleration. One of the malicious joys of instructing from the backseat in a true high-performance tandem aerobatic trainer is watching your student's head snap sideways when you demonstrate maximum acceleration point rolls. A given roll rate can seem much faster (and not such great fun) when you're the pilot-not-flying, because you're not performing the initiating control inputs that prepare the rest of your body for the ride.

Figure 7
Dihedral Effect, Rudder use, Roll rate

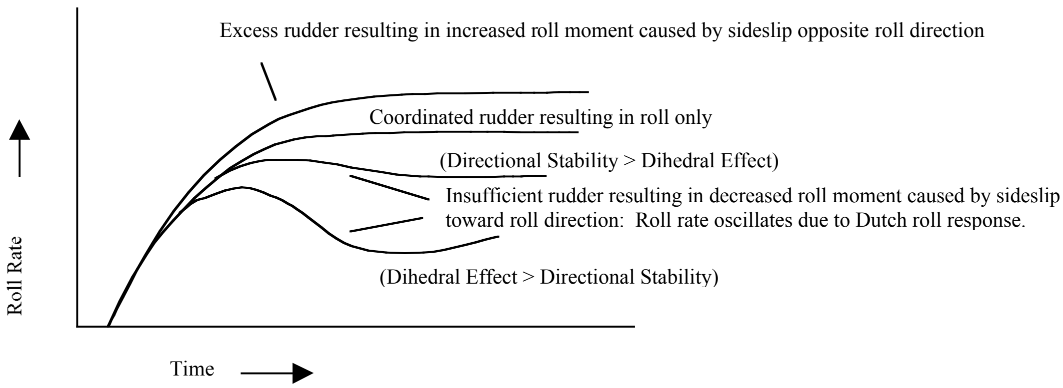

Roll Rate

Figure 6 shows a time history, from the initial acceleration to the point where acceleration stops and a constant rate is achieved, for rolls using three different aileron instantaneous step deflections at the same airspeed. One thing it reveals is that at a given airspeed it doesn't take an airplane any longer to reach its highest, full-deflection roll rate than it does to reach lesser rates using lower deflections.

Figure 7 shows the effects of sideslip on roll rate. Roll moments caused by sideslip are a function of effective dihedral and of both sideslip angle and angle of attack. (See Flightlab Ground School, Lateral/Directional Stability) The two dashed lines in Figure 7 show an uncoordinated, aileron-only bank, with adverse yaw allowed to do its worst. In a laterally stable aircraft, sideslip produced by adverse yaw reduces roll rate because of the opposing rolling moment from dihedral effect. Rolling in one direction while yawing in another can also set off Dutch roll oscillation, seen in the figure as an oscillation in roll rate over time. Note that the worse case, when rolling without coordinating rudder, comes when dihedral effect is high and directional stability is low.

For geometrically similar airplanes, roll rate varies inversely with span (cutting the span in half gives twice the rate). The reason is that roll damping varies directly with span. At any given roll rate, the longer the span the faster the wingtip moves, and therefore the greater the damping caused by the larger roll-induced change in angle of attack. This explains the short wingspans of aircraft designed to roll fast.

Roll damping is also an inverse function of true airspeed. Because TAS increases with altitude, roll damping decreases as you climb—as does directional and longitudinal damping. (See Damping versus Altitude in "Longitudinal Maneuvering Stability.")

At high angles of attack approaching C_{Lmax}, roll damping begins to decrease as the wing's C_L curve begins to level out (Figure 8). Turbulence or yaw rate causing a wing to drop can then force the angle of attack of the tip section past the point of stall. Now the coefficient of lift, instead of rising as it normally does as the wing descends, starts falling down the post-stall side of the lift curve, and damping disappears. The transformation of damping into autorotation is the essence of spin departure. Quickly reducing the angle of attack usually restores roll damping, and lateral control, before a real spin can get underway.

Muscle Versus Roll Rate

Aileron systems are designed primarily in terms of the lateral control required at speeds near stall—a function of aileron size. At high speeds, roll rate is a function of the available aileron deflection.

As mentioned earlier, ***for a given aileron deflection (thus roll helix angle), coordinated roll rate, p, varies directly with true airspeed***. Roll rate also depends on how big a gorilla is driving. In an aircraft without boosted or powered flight controls, ***the control force felt by the pilot increases as the square of the true airspeed***. As a result, aileron forces go up faster than roll rates, and ultimately the force required for maximum deflection can exceed the pilot's muscle power.

For example, a Spitfire had a maximum roll rate of around 105 degrees per second at about 175 knots EAS. A clipped-wing Spitfire made it to about 150 degrees per second at the same speed. A P-51B Mustang's roll rate peaked at only about 90 degrees per second at around 260 knots EAS. When these airplanes went slower, maximum-performance roll rates decreased due to the slower TAS. When they went faster, roll rates decreased because the pilot couldn't fully deflect the controls, as Figure 9 illustrates. The roll rate of the Japanese Zero went down drastically at high speed because of aileron reversal caused by wing twisting (see below).

By way of comparison, if a contemporary, high-performance aircraft designed for top aerobatic competition rolls less than 360 degrees per second, at maximum sustained level flight speed, it's considered a slug.

Figure 9
Roll Rates versus Airspeed, Muscle-Powered Reversible Controls

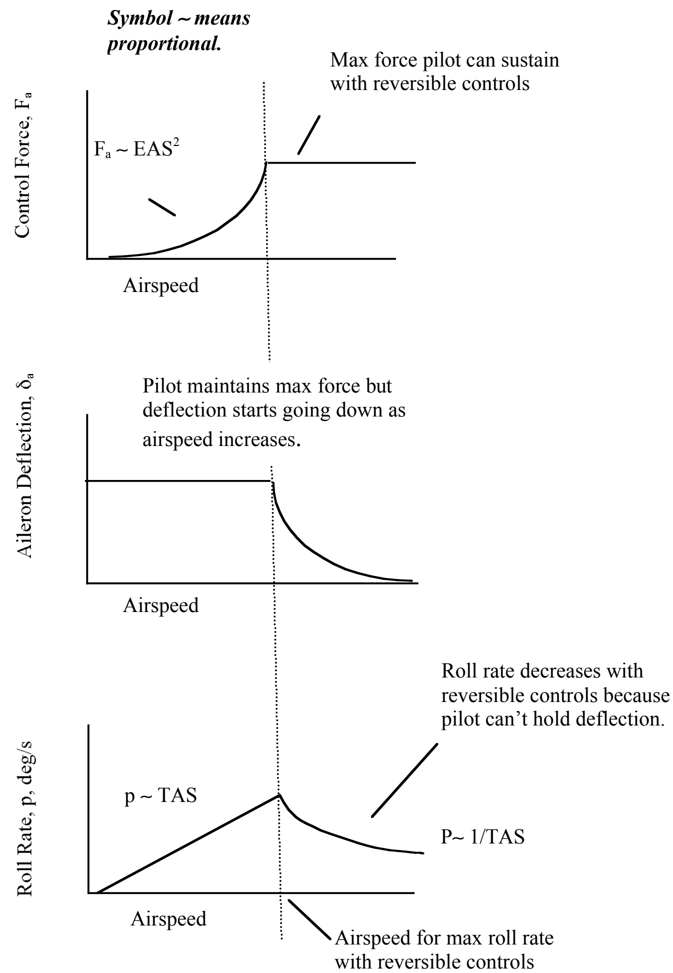

Figure 8
Loss of Roll Damping past Stall

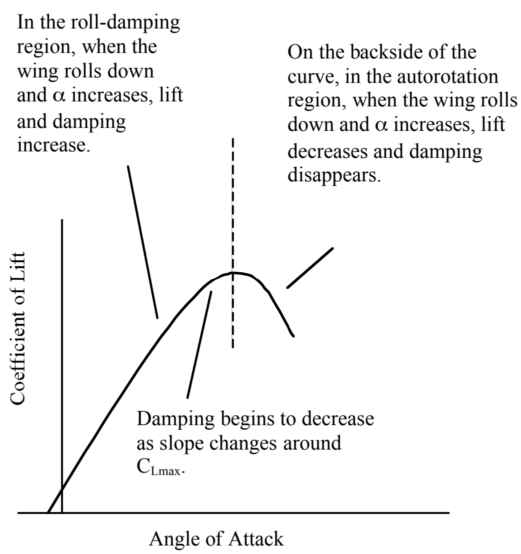

In testing for certification under FAR Part 23.157, aircraft are required to roll from a thirty-degree banked turn to a thirty-degree bank in the opposite direction within time limits specified for aircraft weight and configuration. Under FAR Part 25.147 (e), "Lateral control must be enough at any speed up to *V*FC/*M*FC to provide a peak roll rate necessary for safety, without excessive control forces or travel."

Figure 10
Aileron Reversal

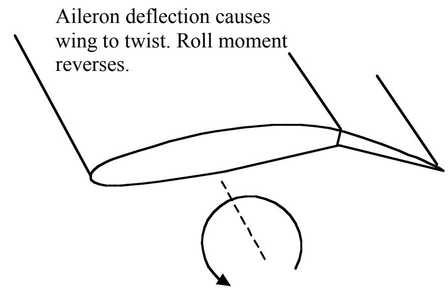

Aeroelastic Aileron Reversal

At high speeds, aeroelastic deformation also puts a cap on roll rates. The down aileron produces a twisting moment on the wing, which forces the leading edge to deflect downward, reducing the angle of attack (Figure 10). This reduces lift and consequently rolling moment. Roll rate then starts going down and at a certain speed, V_R, when the decrease in lift due to twisting equals the increase in lift due to aileron deflection, the ailerons will no longer create a normal rolling moment. Beyond this speed "aileron reversal" occurs. A down-going aileron then produces a down-going wing.

One of the cures for aileron reversal, not surprisingly, is to increase the torsional stiffness of the wing (at the expense of added weight). On swept wings it's necessary to increase the bending stiffness because the geometry of a swept wing causes it to twist as it bends. Moving the ailerons inboard or extending their span inboard also helps raise V_R on a swept wing. Spoilers are another option, as mentioned.

Comparing Roll Performance

You first need to specify flight regime in order to make useful roll performance comparisons between our aerobatic training aircraft and larger transports—for example the roll rates at approach speed and configuration versus cruise. Because of yaw/roll coupling, you also need to consider sideslip and yaw rate contributions based on aircraft dynamics and pilot technique. The question is complicated by all of the derivatives that have to be plugged in.

The people who create simulation algorithms have the derivatives plugged in. Comparisons of roll rates (and control forces) made between your aircraft and our trainers based on the performance of your aircraft's *simulator* should be useful—especially if the simulation occurs within the boundaries of the angle of attack, α, and sideslip, β, envelopes supported by your aircraft manufacturer's flight-test data. Rolls through 360 degrees can be modeled reliably, as long as they happen within α/β boundaries, even if the subject aircraft has never been tested in full rolls itself. Combined high angles of attack and high sideslip angles may not be well supported, however, because they're usually not flight-tested for aircraft not receiving spin certification, and their nonlinear effects make reliable modeling difficult. During the unusual-attitude portion of your simulator training (if indeed there is such), ask to observe roll rates at different airspeeds, configurations, and altitudes—and especially with different contributions from the rudder. Then, assuming you passed the check and won't be called a troublemaker, go ahead and ask where the test data stops and the extrapolation begins.