# How a wing lifts

*“It’s easy to explain how a rocket works, but explaining how a wing works takes a rocket scientist.” - Philippe Spalart*

## Introduction

Some twenty years after completing my Aeronautics degree, I finally discovered the well-buried truth: almost nobody knows how a wing works! (Not in exact detail).

Fortunately there is one man; if ever anyone knows. It’s Doug McLean, a retired technical expert from the Boeing company. In the following I shall be drawing a great deal from him, which is downright plagiarism, however his explanations are hidden within a thick and expensive textbook that is way beyond the foreknowledge (and price-range) of the average reader except in his honest attempt to explain how a wing works to a lay audience. I hope he won’t mind me repeating his explanations, and I’ll try my best not to miss important details when I have to condense, and put his explanation in my own words.

Words in **bold** appear in the glossary at the end of this document.

## The mathematical framework

We have classical mathematics: the unsteady Navier-Stokes equations, which we’re very sure will model the flow around any type of wing at any angle of attack and calculate the ensuing lift and drag. Unfortunately, we don’t have a powerful enough computer to do this as yet, so it’s rather a question of faith.

Instead, we have to simplify the equations somewhat, which restricts us to classical (non-delta) wings, and at angles of attack below the stall. Then we get reasonable answers.

As the aerodynamicists of the early 20th century discovered, you can simplify the equations a great deal and still get sensible predictions for wing lift.

One of the key properties of air is its viscosity (‘syrupiness’) which might well be fundamental to how lift is generated. We can simplify the equations even more by restricting viscosity to the **boundary layer**. Before computers were widely available, even that was ditched and viscosity was ignored completely, and the entire flow was assumed to be inviscid. All the aeroplanes flying in the early to mid 20th century had wings developed without the inclusion of viscosity, but flew perfectly adequately.

The equations that govern the lift of a wing are moderately simple, but when they interact with the billions of air molecules flowing around a wing, we get some really subtle and complex things going on that are hard to explain in a few simple words. As we’ll see, the short explanations generally given as to how a wing work are incomplete, or often just plain wrong, generally because they’re just too brief. Only a fairly lengthy explanation will do the complex flow picture around a wing justice.

As McLean says: “We have to acknowledge that it isn’t possible, without maths, to predict the existence of (wing) lift without knowing a-priori some things about the flow that produces it. Providing iron-clad proof that lift must exist is too much to expect. What we can try to do is to explain (the lift), not predict or prove it.”

## White lies to children

Traditional lay-audience explanations of how a wing works generally follow the same culpability as the white lies we tell to our children to try to explain the world to them in short explanations that they can get a handle on (including the dreaded “because it just does!”).

With only a paragraph or two to spare in a pilot's training manual or model aircraft website, the author doesn't have the time, nor often the knowledge, to devote pages of text to the explanation, so instead offers a traditional 'quick sound bite' that promulgates the same old errors of antiquity. Given the subtlety of the phenomenon, over-simplification is a great temptation, but is a major weakness.
There are many cockeyed theories of wing lift "in circulation." (McLean's pun!)

## Lie 1: Equal transit time and the Bernoulii explanation

You're bound to have run into this old chestnut:

What you are told is that the airflow meets the front of the wing (which is called the leading edge) and is cleaved into two flows, one over the upper surface of the wing and one over the bottom surface of the wing, and that they have to meet up at the back (the trailing edge) again at the same time.

The upper surface is longer because of the aerofoil's special shape, so the flow over the top surface has to travel faster to get to the trailing edge at the same time as the flow over the bottom surface does. A faster flow has lower pressure from Bernoulli's equation which states - correctly - that a faster-moving flow produces less pressure. This lower pressure 'sucks' the wing upwards creating lift.

This explanation has many flaws:

Flaw 1: Why do the upper and lower flows *have* to join up at the same time at the trailing edge? It's never explained, so comes under the category of "it just does".

Actually, it doesn't. You'd be amazed that for over a hundred years very few people actually bothered to check whether the flows *did* meet at the trailing edge at the same time. And then quite recently, the University of Cambridge's Professor Holger Babinsky decided to test it out in a wind tunnel and discovered something very surprising: they *don't* meet at the same time at the trailing edge, the upper surface flow gets there considerably earlier. See the video at www.cam.ac.uk/research/news/how-wings-really-work from which the following pictures are taken at successively increasing time intervals:

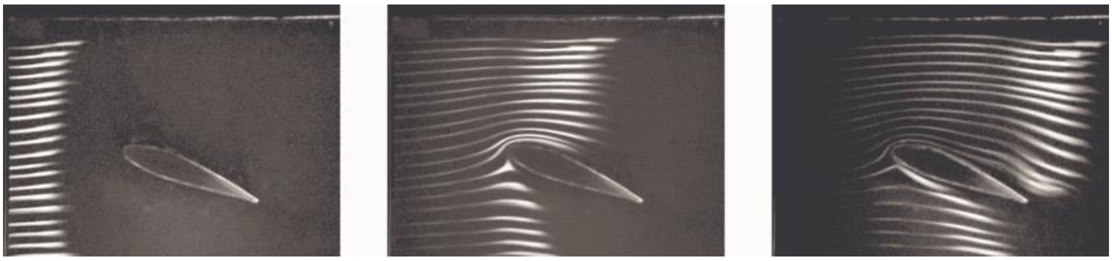

Flaw 2: This explanation is completely at odds with the fact that aeroplanes can fly upside-down when the longer surface is then at the bottom, provided the wing is at positive angle of attack (and then the upper surface flow *still* reaches the trailing edge first).

Flaw 3: Simply saying that the upper surface flow must go faster doesn't explain *how* it goes faster. It doesn't identify the physical force that accelerates it to higher velocity.

Flaw 4: Cause-and-effect
So this explanation has it that some unknown effect makes the upper flow faster, and that this then lowers its pressure due to Bernoulii's principle.

Unfortunately, fluid flows don't follow this one-way causality where A causes B which causes C. In actual fact, the causality is both ways: A causes B but B sustains A, a circular causality. Yes, the faster flow and the lower pressure do occur, and can be described by Bernoulii's principle, but one does not singularly cause the other, they both sustain each other. More on this later.

Flaw 5: Saying "the flow next to the wing goes faster" isn't quite right. Parcels of air right next to the surface get trapped in the **boundary layer** and never *reach* the trailing edge.

Flaw 6: The increase of lift caused by increasing **angle of attack** is completely ignored.

Flaw 7: If we do a simple calculation assuming equal transit time, we would find that in order to generate the required lift for a typical small airplane, the distance over the top of the wing would have to be about twice as long than under the bottom. Here's what such an aerofoil would look like, enormously fat:

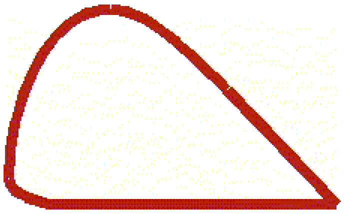

Flaw 8: The lower pressure "sucks the wing upwards".
I fell for this one. It's easy to get into very lazy thinking about what suction is, and I did.

Actually, we need to go right back to basics: What is suction? We tend to think of it in terms of sucking a drink up a straw, or the suction force of a 'vacuum' cleaner, or suction cups. If you don't think about it too hard, you can get into the mindset that a lower pressure can produce a pulling force.

But think about what pressure is: it's the ceaseless bombardment of air molecules banging into the surface of the wing. There's no way they can produce a pulling force. The pressure may be low, but it's positive. Correctly, suction just means a pressure lower than ambient.

With suction cups, it's the higher air pressure outside pushing *down* on the suction cup that holds it to the table. And when we suck on a straw, it's not the sucking force that we make, it's the higher air pressure pushing on the drink in the rest of the glass that pushes the liquid up the straw; we're not pulling the liquid up the straw.

A tornado doesn't suck the roof off a house solely because of its low pressure; it's a *difference* in pressure between the pressure inside the house and the lower pressure inside the tornado. It's actually the pressure inside the house that *pushes* the roof off.

So, low pressure can't exert a pull on the upper surface of an aerofoil. It's more correct to say that the lower surface pressure does all the pushing the wing up - and then some - and all the upper surface pressure does is *not* push back down so hard on the upper surface.

What fooled me were the technical diagrams in my Uni course textbooks using arrows to show the pressure difference around the aerofoil compared to ambient. The arrows over the majority of the upper surface point upwards so seem to suggest that they're pulling away from the upper surface as if there's a 'suction force'.

As an aside, many believe that the greater reduction in pressure on the upper surface compared to the slight increase in pressure on the lower surface is a necessary characteristic of lift. But it's not actually, it's just a side effect of an aerofoil of finite thickness.

### Lie 2: The momentum-based explanation

Now it's unfair to call this one a lie, because it's broadly correct, but it's too short on detail, there are large chunks missing from the explanation.

The air behind the wing is deflected downwards. This is true, we can see it with smoke, and it's called the downwash. To acquire downwards momentum, the air must have a downwards force exerted on it by the aerofoil.

Thus, by Newton's third law (forces always occur in pairs), the aerofoil must have an equal upward force exerted on it by the air. This explanation avoids Lie 1's one-way causality fallacy: the mutuality of the force exchange between the aerofoil and the air is explicitly acknowledged.

So far, so good, and you can crudely picture how the lower surface flow is kind of 'wedged' downwards by the angle of attack causing a downwards-sloping lower surface.

Flaw 1: This is too crude a mental image: the flow turning has to be caused by a pressure force, but pressure isn't mentioned.

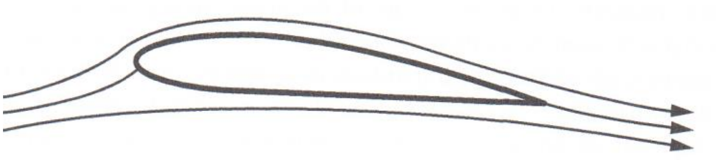

[In a thin vertical region directly ahead of the wing there's very visible upwash. BUT this doesn't reduce the wing lift for an interesting reason: way above and below the dramatic and visible upwash, there's a great deal of downwash which cancels it out ahead of the wing. However, this downwash is spread-out thinly (is of low magnitude) over a great vertical distance: it's very weak but there's a large vertical swath of it – it's large in summation - but it's easily missed in photographs of the airflow around a wing.

In a wind-tunnel, this near-invisible downwash cancelling-out the upwash simply isn't there because way above and below the front of the wing, the airflow impacts the floor and ceiling of the tunnel. These impacts generate pressure forces that push down through the flow and onto the wing to cancel-out the effect of the upwash.]

Then it's also emphasized that the flow pattern over the upper surface also contributes strongly to the overall downwash: the upper surface flow goes downwards too at the trailing edge.

Flaw 2: Yes it does, but the reason given as to *why* the upper surface flow also goes down is usually wrong.

In many textbooks, it's stated that the reason the upper surface flow remains attached to, and follows the curvature of, the upper surface is to do with the Coandă effect.

Applying the term Coandă effect is inaccurate and therefore confusing. What is happening over the upper surface appears Coandă-like, but it *isn't* the Coandă effect. The Coandă effect is erroneously seen as implying that viscosity plays a role in the ability of a flow to follow a curved surface. It is asserted that differences in speed in adjacent layers within the boundary layer cause shear forces which cause the flow of the fluid to want to bend in the direction of the slower layers. Actually, there's no physical basis for any direct relationship between shear forces and the tendency of the flow to follow a curved path.

The Coandă effect properly works on a jet of fluid exiting from a nozzle at high pressure. What is happening is that there is a significant entrainment of the air around the jet into the jet. If the jet impinges on a curved surface, the airflow can only be entrained into the side of the jet away from the surface, but not near the surface because there's a solid boundary in the way. This asymmetry bends the flow onto the surface. So the Coandă effect is really an effect of entrainment. This doesn't happen on an aerofoil.

In fact, the flow turning to follow a convex surface doesn't require viscosity at all. When a flow turns to follow a curved path - which it naturally will - the pressure field naturally adjusts itself so as to provide the force needed to accelerate the fluid towards the centre of curvature.

The centrifugal force generated when the flow follows a curved path is countered by a pressure gradient at right angles to the local flow direction. This normal pressure gradient and the flow curvature have a circular relationship in which they cause and support each other simultaneously.

So which is more correct (or less wrong), the Bernoulli approach or the momentum approach? In a way, they're both right in part: they are not contradictory.

The Bernoulii explanation is based on the near-surface flow and the surface pressure, whereas the momentum explanation is based on a manifestation of lift that extends way out into the far-field. A more satisfying explanation ties the workable bits of both explanations together.

## A fuller explanation

### Aerofoil shape and angle of attack

How much lift force a wing produces depends on the shape of its **aerofoil**, the **angle of attack** at which it approaches the oncoming airflow, and the airflow speed and air density.

Starting with simple shapes, a curved-plate or 'cambered' aerofoil works better than a flat-plate aerofoil. Adding thickness to produce a streamlined aerofoil which is rounded at the **leading edge** will then produce more lift with less drag.

### The aerofoil reference frame

It's easier to visualise the airflow in a reference frame that moves with the aerofoil - so that the aerofoil appears stationary - than one that is fixed to the ground. Either frame is equally valid as far as lift is concerned.

### Lift involves action and reaction (Newton's third law)

This is the momentum explanation described earlier, but we now have to explain *how* the moving airflow pushes upwards as the reaction force called lift.

### Lift is felt as a pressure difference on the aerofoil surfaces

A fluid above a temperature of absolute zero exerts a pressure. When the aerofoil isn't moving, this pressure is the same all over the wing surface, at just the ambient air pressure.

However, when the aerofoil is flying, the average pressure on the lower surface is usually higher than ambient (unless the aerofoil is very thick), and the average pressure on the upper surface is always lower than ambient. This pressure difference causes lift.

This pressure difference is surprisingly weak - *much* less than the ambient pressure - causing roughly only 0.2 Newtons of lift per square centimetre of wing area, which is why acres of wing area are required to lift a heavy aeroplane.

### Lift involves force and acceleration (Newton's second law)

We need to explain how the airflow maintains this pressure difference by examining the forces exerted on the air, and the resulting accelerations *of* the air, not just at the surface of the aerofoil, but in an extended region around the aerofoil. This requires considering that:

The air flows as if it were a continuous material that deforms to follow the contours of the aerofoil. The air consists of a huge number of free molecules that move randomly in all directions and collide frequently with their neighbours.

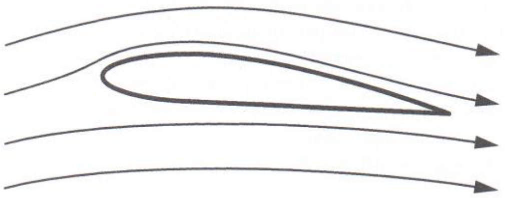

Because of this continuous interaction between molecules, the air flows as if it were a continuous material which deforms and changes course to flow around the aerofoil, and fills all of the space around the aerofoil and touches all of its surface.

Because of this deformation, changes in flow direction are gradual, and the speed and direction of the flow vary over a wide region around the aerofoil. An extended pattern of variations such as this is referred to as a flowfield or velocity field.

The aerofoil's solid surface forces the flow to follow the direction of the aerofoil contour, with the result that the speed and direction of the flow are affected over a wide area.

To produce the downward airflow - and hence lift - the rear of both surfaces of the aerofoil must have a predominantly downward slope, so the aerofoil must have either camber and/or a positive angle of attack. Realise that whilst it seems obvious that the air ‘smacking’ into the underside of the aerofoil gets ‘wedged’ downwards, in reality, it is turned by a region of pressure that dwindles slowly with distance way below the aerofoil.

The aerofoil affects the air pressure over a wide region: a pressure field. A diffuse ‘cloud’ of low pressure always forms above the aerofoil, and a diffuse ‘cloud’ of high pressure usually forms below. Where these clouds touch the aerofoil they create the pressure difference that causes the lift on the aerofoil.

These pressure differences are generally largest at the aerofoil surface then die away gradually in all directions: above, below, ahead, and behind, the aerofoil.

The downwash, the changes in airflow speed, and the clouds of low and high pressure are all necessary for the production of lift. They support each other in a reciprocal cause-and-effect relationship, and none would exist without the others.

We can imagine tiny ‘parcels’ of air moving around the wing, and investigate how they move. Newton's second law describes how a pressure difference imposes a net force on a parcel: it must cause a change in the speed or direction (or both) of the parcel's motion.

The pressure difference causes a change in the parcel's motion, but the existence of the pressure difference *depends* on the parcel's motion. The relationship is thus reciprocal: a fluid parcel changes speed or direction in response to a pressure difference, and its resistance to changing speed or direction (its inertia) sustains the pressure difference.

There are both horizontal and vertical examples of this reciprocity:

In the pressure ‘clouds’ both above and below the aerofoil, the pressure is higher above than below: there is a vertical pressure gradient. This causes a downward force on fluid parcels, which is resisted by the downward acceleration of the flow. This pressure gradient is sustained because the air parcel has inertia and therefore resists having its path deflected from a straight line.

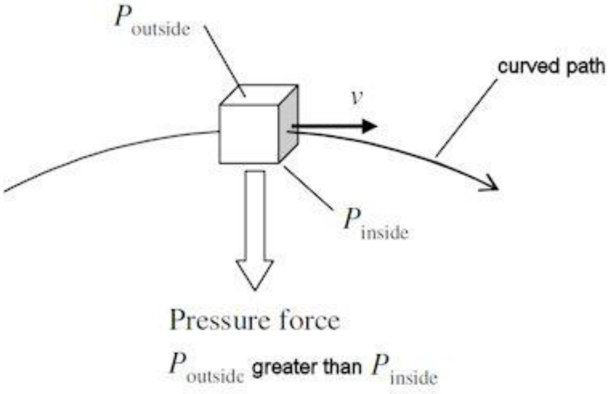

Similarly, parcels of air encountering the cloud of low pressure over the upper surface of the wing are speeded up by the low pressure as they enter the cloud from upstream, then slow down as they leave the cloud. Conversely, parcels entering the higher pressure cloud below the wing are slowed down within the cloud.

Outside of the viscous **boundary layer** and viscous **wake**, these changes in parcel speed with changes in pressure follow Bernoulli's principle (high speed equals low pressure and low speed equals high pressure).

Once again, the relationship is reciprocal: the differences in pressure in the horizontal direction cause changes in flow speed, and the fluid's resistance to acceleration, because of its inertia, sustains the pressure differences.

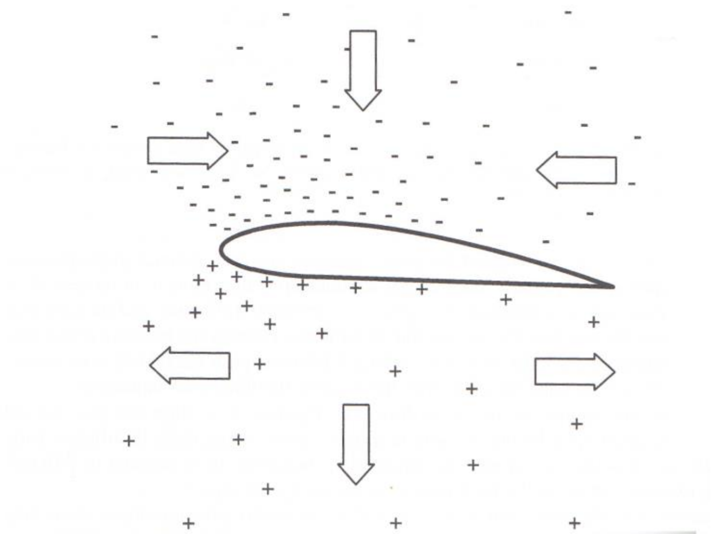

Pressure differences exert net forces on fluid parcels in different parts of the flow field

You might be surprised that air mass (inertia) is so important, given the tenuous nature of air, however even a modest volume of air (one cubic metre) has considerable mass (1.225 kilograms). The wings of even a light aeroplane cause significant changes in speed and direction to thousands of kilos of air per second.

## The roles of camber, a sharp trailing edge, and a rounded leading edge

Both the aerofoil shape and the angle of attack affect the downward-turning of the air around the aerofoil.

A cambered aerofoil aligns the aerofoil surfaces better with the desired curved flow and enhances the downwash. As a result, cambered shapes produce more lift at a given angle of attack than flat plates.

The trailing edge plays an important role in controlling the downwash: if the flows from both the upper and lower surfaces didn’t leave the aerofoil smoothly from the sharp trailing edge (which they do) then the flow from one side would have to negotiate the very sharp corner around the trailing edge to get to the other side.

But viscosity prevents flow from going around a sharp edge in this way. So, with help from viscosity, the trailing edge has the effect of directing the flow - as it leaves the aerofoil - to flow in the direction dictated by the shape of the rear of the aerofoil. This tendency of the trailing edge to steer the flow downwards is a major reason why angle of attack and camber are so effective in influencing lift.

In contrast, a rounded leading edge is needed to accommodate - with a minimum of flow disruption - the wildly differing directions that the upwash approaches the aerofoil at different angles of attack.

## Glossary

### Aerofoil (airfoil) geometries:

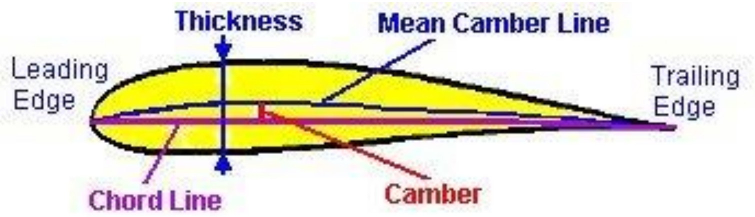

**Angle of attack: $\alpha$**
This is usually referred to as 'alpha', and corresponds to the angle between the incoming airflow direction (usually the **Freestream** direction) and some vehicle or fin datum such as the wing chord line.

**Angle of incidence (referred to as 'rigging angle' in the U.K.):**
The fore-aft tilt of the wings or tailplane with respect to the aircraft's fuselage centreline.

**Bernoulii's principle:**
Is just a statement of the Law of Conservation of Energy couched in aerodynamic terms and is expressed in the equation: $P + \frac{1}{2}\rho V^2$ = constant,

or: $\Delta P = -\frac{1}{2}\rho \Delta V^2$ where $P$ is pressure, $\rho$ is density, and $V$ is flow velocity.

**Boundary layer:**
The surface of an aerofoil is mountainous on the atomic scale, which dynamically 'snags' the airflow right at the surface, bringing it to a dead stop on the surface. The boundary layer is the thin (a few millimetres) region where the air slows progressively down from the flow speed down to zero at the surface. Within the boundary layer, viscosity and turbulence are prominent.

**Freestream:**
The undisturbed airflow ahead of the vehicle.

**Lift (equation):**
Lift is a force generated by aircraft at right-angles to their flightpath.
The equation used to calculate lift is simply the *lift coefficient, Cl*, times **dynamic pressure**, times some reference area '*S*', i.e: $L = \frac{1}{2}\rho V^2 S\, Cl$ ($\rho$ = atmospheric density.)
For aeroplanes, this reference area '*S*' is the total wing area.

**Skin:**
The outer covering or surface of the vehicle.

**Wake:**
The disturbed flow region behind the wing strongly exhibiting the effects of viscosity.

## References

Ref. 1: ‘Understanding aerodynamics – arguing from the real physics’, Doug McLean, Wiley, 2014

Ref. 2: ‘The design of the aeroplane’, Darrol Stinton, Blackwell science, 1995

Ref. 3: ‘Introduction to Flight – seventh edition’, John D. Anderson Jnr, McGraw Hill, 2012

Ref. 4: ‘Aerodynamics for engineering students – third edition’, Houghton and Carruthers, Arnold, 1988

Ref. 5: ‘Application of the leading-edge-suction analogy of vortex lift to the drag due to lift of sharp-edge delta wings’, Edward C. Polhamus, NASA Technical Note D-4739
Downloadable at:
http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19680022518_1968022518.pdf

Ref. 6: ‘Stick and rudder’, Wolfgang Langewiesche