The Bounce of the Superball
Transcript

Date: Tuesday, 7 December 2010 - 12:00AM

Location: Museum of London

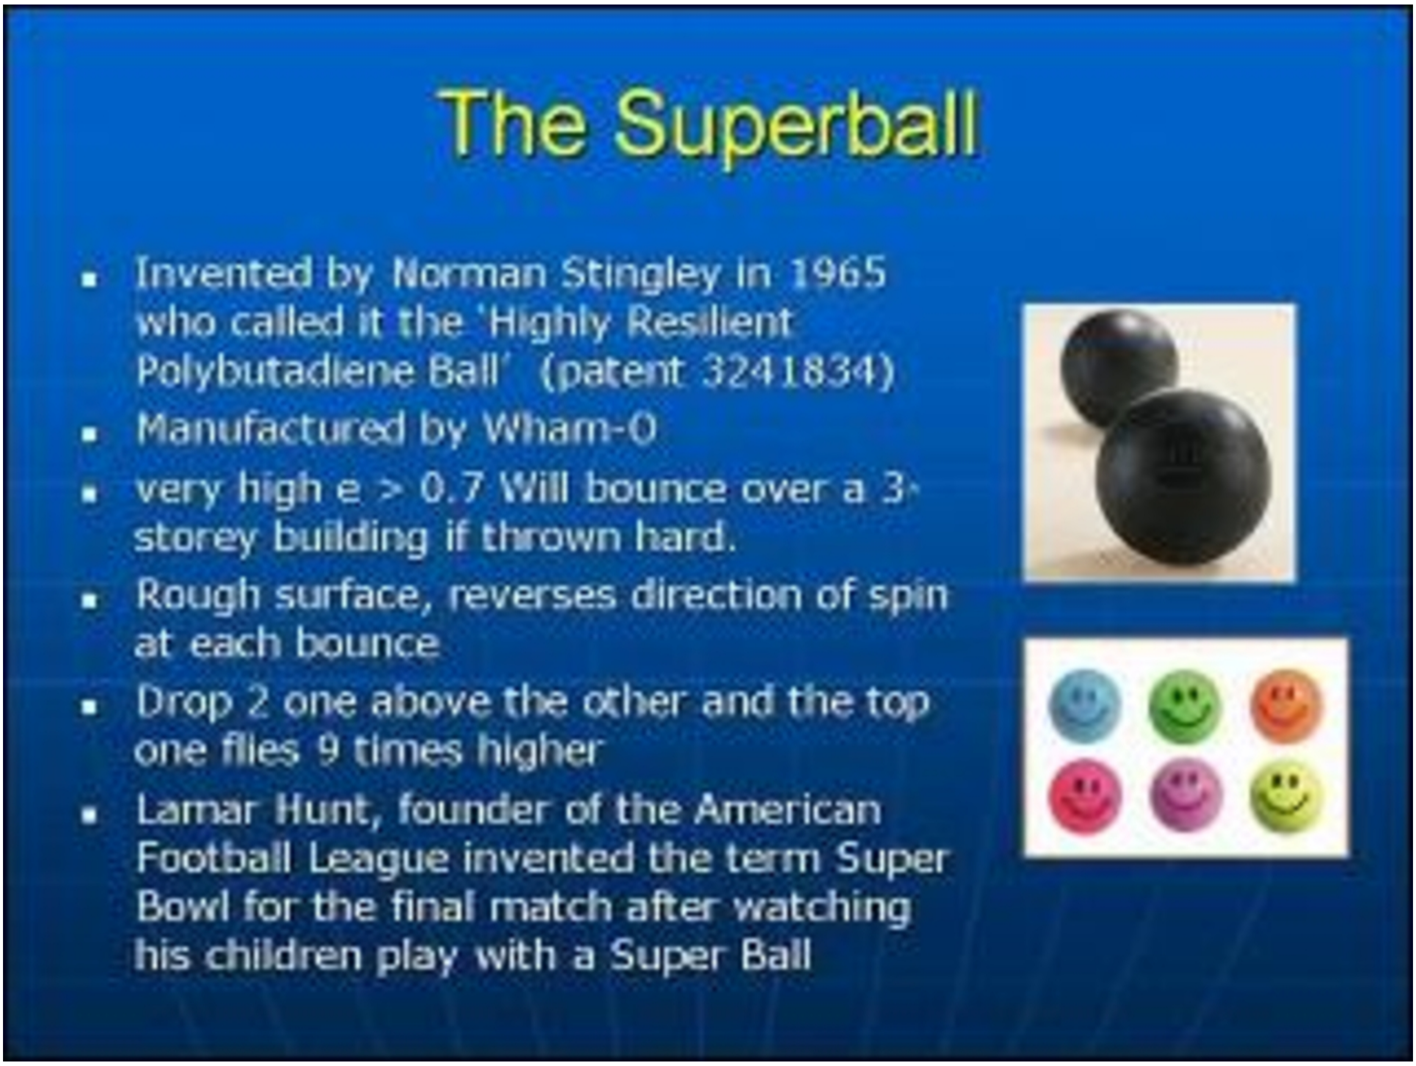

The Bounce of the Superball

John D Barrow FRS

Gresham Professor of Geometry

Welcome to today's lecture. I shall talk today about a mixture of problems that involve balls being thrown, balls being struck, balls bouncing on the grounds - problems of projectile motion and problems of impacts. I am going to use practical applications to various sporting events and other types of real world situation to illustrate something that is unexpected and unusual about each of these situations.

My first and simplest example of throwing things - simple because the thrown object weighs a lot, about 16lbs – is the shot-put. Its weight means that it is not affected by air resistance or other cunning aerodynamic features, like a golf ball would be.

Let us briefly recap our knowledge. In a previous lecture, when I talked about curves - about the expected trajectory of a projectile like this and what you should you do to make it go as far as possible. Returning to this topic, we shall see that there are a couple of surprises, that the real world situation is not what you would expect from the simple textbook idealisation.

A textbook idealisation is the following: if you throw something, from the ground here, any mass, and you launch it at an angle with the ground (theta θ), with a speed v, then it will follow a parabolic path, something first noticed by Galileo long ago, and the maximum or the range that it will have before it comes back to the ground will be this distance here. It depends on the square of the launch speed, and it depends on the acceleration due to gravity, and it depends on the sine of twice the launch angle. A sine of any angle can never be bigger than one, so the largest range would come when the sine of two theta is equal to one, and that is when two theta is equal to 90 degrees. So, when theta is 45 degrees, you will get the maximum range in this situation.

Putting the shot is not as simple as that because you do not launch from ground level. If you were coaching a team to put the shot as far as possible, you might naïvely think it best to launch at 45 degrees. However, if you thought twice about it, you might realise that that might not be the best thing to do.

So, in the case of a shot putter who is some height (h) – if it is one of the gigantic male shot putters, h could be about two metres tall or six foot six in the case of someone like Geoff Capes, – then the maximum range becomes slightly different. It is determined by this height at which you launch things, and so the formula for the range is the old formula when you launch from the ground plus a correction which takes into account the fact that you might be launching some height (h) above the ground.

When h is zero, this just stays as the old formula. The point, of course, is that the angle needed to give the maximum range is no longer 45 degrees. The angle depends on how tall you are, what the value of h is, and also what the speed is at which you launch the shot.

So, we shall use some simple numbers: acceleration due to gravity 9.8m per square second, height of our shot putter 2 metres, and when you throw the shot, you launch it at 14m per second. In this case, you are looking at about 43.5 degrees rather than 45.

If you look at the formula as a whole, you can see there are three factors that you might work on if you wanted to be better at throwing the shot. You could perhaps become taller or encourage taller athletes into this event, or you could move to another planet where the acceleration due to gravity is smaller. As we shall see in a moment, this is a real factor you could take into account on Earth, because the acceleration due to gravity is different at different places on the Earth. But if you were to shot-put on the Moon, g would be six times smaller, and you would throw six times farther, all other things being equal. However, if you were to change g by 1%, if you were to change h by 1%, these would be the differences that you would make: you would add about 20cm to your 20m or so shot-put by changing g by 1%; if you increased your height by 1%, you would just add 2cm to the distance you threw. But the really big change comes by becoming better and stronger and more dynamic and being able to throw the shot with a greater launch speed, because that comes in as the square of the launch speed and you get a 40cm improvement for a 1% increase in launch speed. So it is clear what you should do, as a shot put coach: your team has to become dynamic, much stronger, much better at launching things at high velocity.

Here is a picture of the results from that formula. If you had different launch speeds and you were 2m tall, here are the optimal projection angles which are the maxima of these curves. As you can see, they are quite close to 45 degrees – a little bit less. If you are a world-class shot putter, you will be looking at 42 degrees; if you were 50 feet, a 15 meter shot-putter, 41; if you were a school first competitor, 39.

That is the first factor to take into account. You will see many articles in education journals, like the American Journal of Physics, which discover and re-discover this little detail and use it as an example of slightly unusual projectile motion.

But if you know anything about athletics, you know that real top class shot putters do not launch their shot projectiles anywhere near that optimum angle of 42 degrees, or 43 degrees. It is nowhere near that, certainly nowhere near 45 degrees. It is closer to about 37 or 38. So what is going on here?

There is another constraint in this problem. So far we have been thinking that you could change v and you could change the launch angle at will, and that the two are not correlated in any way. But if you try to lift a heavy weight, at different angles to your body, you will soon learn that there is a great difference between holding 100lb weight out like this to pushing it above your head. So the angle at which you apply a force is correlated with your strength, and therefore with the speed with which you can launch something. You cannot launch something with equal speed at any angle.

Here is a picture of various tests, under control conditions, of what the launch projectile speed would be for different launch angles. You can see the tendency that, as you increase the angle from zero to 90, the speed with which you can launch it steadily decreases.

So if we fold this constraint into our previous formula here, so that v and theta are now constrained by that relation, what we discover is that the optimal angle for launch is much less than our 43 to 45. It is indeed down somewhere between 34 and 38 degrees, depending on how good you are, how fast you actually launch it, and the other detailed factors of the individual shot putter.

Here is a collection of the optimal ranges against projection angles, and you should be interested in the peak. You can see that if you are a very good, world-class shot putter – the world record is about 23m – your optimal angle is up here and, as you get weaker, your optimal angle slips down. This is the second surprise about this simple problem: the initial conditions, the things that you think are independent variables, are not. There is an additional constraint in the problem.

We mentioned shot putting on another planet. That is a bit extreme. What about shot putting on this planet? We saw that the range depended on v squared over g. This is a very important factor. If we were to work out high jumping performance or long jumping performance – long jumping is another type of projectile motion – the maximum distance you can achieve is always proportional to this combination, v squared over g. So as g gets smaller, the strength of gravity is weaker, less constraining; you go further, you go higher.

Here is a schematic picture of the Earth. As you are aware, the Earth rotates on its axis once every day, which means that if you were to hold a spring balance and walk over the surface of the Earth, that spring balance would record the net attractive force of gravity of the balance mass towards the centre of the Earth.

The key factor is that this measured acceleration due to gravity changes as you go from the Poles to the Equator; it gets smaller and then increases again. It changes because of two factors. The first, the smaller, is that the Earth is not completely spherical. It is slightly fatter at the Poles and, as a result, there are some places, closer to the Poles, which are further from the centre of the Earth than places near the Poles, so they will feel a very slightly weaker force of gravity. But the dominant effect is created by the rotation of the Earth. So if you are located at a point here, then you are going round in a circle, with this radius, every day, and that produces on you a reactive force, which points outwards horizontally in this picture, and therefore opposes the force of gravity pulling you towards the centre. That force is zero at the Poles because you are going round in a circle of zero radius, and it is a maximum at the Equator because you are rotating in a circle of maximum radius. So that centrifugal force opposes the attractive force of gravity towards the centre of the Earth, and so, as you walk from the Pole to the Equator, the effective value of the acceleration due to gravity decreases.

Here is that effect as a formula. The net acceleration due to gravity looks like this - the acceleration due to gravity because of the attraction of the mass that you are using - and here is the centrifugal effect caused by the rotation. It depends on the radius of your latitude and your angular velocity, one revolution per second.

I have put an m here for mass, so this is the weight that you would measure at a radius r, and this is the angular velocity. If you have a mass, and you weigh it on a balance at the Poles, and you weigh it at the Equator, it will weigh less at the Equator than it weighs at the Poles. So, if you are a weightlifter for example, you would want to get close to the Equator and high up. A venue like Mexico City is really excellent. A 200kg mass in Mexico City would weigh 200.8kg in Helsinki. So this is a significant difference when it comes to setting records or such considerations.

Similarly, the high jump or long jump that you could achieve is also proportional to one over the acceleration due to gravity. All things being equal, besides g, a 2m high jump in Helsinki would give you 2.05m in Mexico City, and an 8m long jump turns into 8.20m in Mexico City. So these are not inconsiderable differences.

We want to move on to consider something else about projectile motion now, and to add more realism to the problem. So far, we have been looking at the motion of projectiles in the way that you would in school physics or mathematics, assuming that there is no air resistance and that everything is happening in a perfect vacuum.

This is a fairly good approximation for shot putting, which is why I chose it. The shot is so heavy, it hardly creates any significant air resistance, but if you move a body through a medium, whether it is water or air, then there is a force of drag created by the medium you are moving through, and that force is proportional to the cross-sectional area that you are presenting to the medium as you move through it. It is proportional to the square of the speed with which you move, and it is proportional to the density of the medium that you are moving through. So you cannot parachute jump in a vacuum.

Here is a simple picture. Here is the effect that I just mentioned. If we were to have a sphere moving through air, the drag that occurs because of the air's proportional to the density of air, so if you go from one place in the world to another, where the density changes a little bit, this drag will change. It is a very insignificant effect. But it depends on the area that is presented, so πr^2, and the square of the velocity. That is what really counts – how fast you are moving through the medium.

Here is a picture of projectile motion, with a launch angle of 60 degrees, and these are times in seconds through the trajectory. What you see at the top is the idealised situation for a launch speed of 45m per second, which is motion in a vacuum, and the trajectory is a perfect parabola. So this is the sort of thing that we were looking at just now, with the shot put thrown from ground level initially.

If you add air resistance into the problem, there is a big change. First of all, you will notice that the trajectory is no longer a parabola, it is no longer symmetrical between the first half and the second half, it does not go anywhere near as far, and it does not remain in the air for anywhere near as long. You are looking at a trajectory here that goes slightly under 100m, whereas in a vacuum, it would be 177m. So, including air resistance in projectile motion has a major effect.

During World War II, lots of this country's leading mathematicians were involved in a rather mundane and tedious business of calculating range tables for artillery shooting and shell motion in order to calculate, in great detail, what these trajectories would be expected to be like in different conditions of wind and so forth.

Notice also that the heights are not consistent – 53 here, 76 there. So real world projectiles do not quite behave like the ones in the textbooks.

Here are two more pictures that show you the effects. This is a situation where you are launching at quite a low speed, about 10m per second, about as fast as a top class sprinter could run. What you see is that the dotted trajectory here has got a launch angle just a bit bigger than 45 degrees, and yet it has got a longer range than the 45 degree trajectory in vacuum.

Over here is an example of what happens when you launch at very high speed, but with a very low angle of launch, and then you see a very pronounced difference in what happens to the motion of the projectile. It rises pretty steadily, almost just proportional to distance here, and then undergoes a very sudden change and drop. So the trajectories start to appear very different.

Here is a collection of golf ball trajectories, taken from film from a golf range, and you can see this type of low trajectory going proportional to distance, and then a much more sudden drop.

The person who first started to think about calculating these trajectories was Peter Tait back in 1890. Tait was a famous mathematician, who created knot theory and many other interesting areas of mathematics. He was the youngest ever senior wrangler, the top student in mathematics in Cambridge, when he was about twenty, and he was a schoolmate of James Clark Maxwell, who I think was a year above him. All through his life, he and Maxwell were great friends and great competitors, first at school, and then at university, and then they collaborated later in life. He also wrote some famous books and papers with Lord Kelvin; "Course of Physics", by Kelvin and Tait, was a canonical text.

Tait was also very interested in golf, and his son was a champion amateur golfer. He was the first person to consider the realistic motion of a golf ball, taking into account the air resistance that would be encountered, but also taking into account the fact that the golf ball is a sphere which will rotate, with some back-spin perhaps, and this creates another aerodynamic factor to affect its motion that we will look at in a moment. What he showed was that the trajectory is not a parabola, in general. Approximately, when the launch angle is small, so you can approximate theta by sine theta, the distance in the x direction, after time t, varies just as the logarithm of one plus some constant times the time, and in the y direction, here's the angle theta, here is this factor, so this is x, time minus gt squared. So if we had no air resistance, we would just have this factor and this factor here.

This is the first picture that he drew in his "Papers in Nature", back in the 1890s, trying to calculate the optimal way to drive a golf ball and, for the first time, to study the motion of a projectile that had finite size, encountered resistance and span as it moved.

What he included was the fact that if you have a ball of this sort, and it is moving through a medium, it encounters drag, it has a weight, a force acting downwards, but it also receives some lift. It does so because the spin ensures that air passes over the top of the ball faster than it passes by at the bottom, and so the pressure of the air at the top is lower than at the bottom, resulting in an upward lift force that tends to push the ball upwards. This is the third factor - in addition to gravity and the drag - that Tait included in that first description.

In practice, golf balls are of course dimpled, and that dimpling is not just decorative or for fun. The dimples both decrease the drag on the ball as it moves through air and help increase the lift. They do that by changing the nature of the air flow that is very close to the ball, that is in some sense almost stuck to the ball for a while by frictional effects. This is what mathematicians call the boundary layer.

At the top, where the speed is faster, the boundary layer becomes turbulent, and the flow becomes disordered. Down by the bottom, it remains low speed and much more orderly - what mathematicians call lamina or smooth.

The turbulent boundary layer clings to the ball for much later before it breaks away, and it is this effect that increases the lift and reduces the drag. So a dimpled golf ball behaves very differently to a smooth sphere. With a smooth sphere, you have a breakaway of that boundary layer much earlier: much more drag and less lift. With the dimpled surface at high velocity, the boundary layer hangs on for much longer before it breaks away. At the bottom, where the speed is lower, you do not get the turbulent flow and get a different effect.

So the dimpling of golf balls is a rather significant factor, and you could imagine that there exists something of a golf ball crystallography. There are dozens, maybe hundreds, of different designs that seem to exist for the dimpling of golf balls, and it would take another lecture to talk about them. I think Ian Stewart once wrote an article called “Golf Ball Crystallography”, in which he studied the geometrical patterns of the dimples that are put on golf balls, as people try to optimise the turbulent creativity round the surface of the ball.

Here are two that have icosahedral symmetry. If you are into this sort of geometry and you like golf balls, here is a subject for you.

These sorts of effects - the smooth flow and the turbulent flow round the edge of a ball – can be seen happening with cricket balls. Throwing a cricket ball is a more complicated motion because of its seam, but if you shine one side of the ball and leave the other side alone, or if you are the old England captain and have lots of grit and sandpaper in your pocket, allowing you to rough up one side of the ball while nobody is looking, then you end up with a ball that is very smooth on one side and very rough on the other. Consequently, the air flow is rather smooth and lamina round the smooth side, but turbulent and disordered around the rough side. In the case of a cricket ball, that will cause a lateral force, and the ball will swing - it will move in the air. So, if you want to win by means foul or fair, bear that in mind. But at the moment, we do not seem to need to do that sort of thing…

The next thing I want to talk about, just very briefly, is an odd topic. It is catching a moving or flying ball, what the Americans call a fly-ball. I always wondered what that was. I always thought it was something that you put in the pantry or something to catch insects, but it seems not. A fly-ball is a ball that just comes at you in the field, out of the blue.

The interesting problem posed, a long time ago, by Mr Chapman, was what do you do… what should you do… what does the brain do in order to be in the right place at the right time to catch a ball that comes towards you? It is one of these things that you instinctively do, but people who study what the brain is doing, how it calculates unconsciously what to do, would love to discover the simplest algorithm to programme a robot to catch a cricket ball coming towards it, so that it is in the right place, at the right time, when the ball just comes to the hand.

This is a complicated problem, and to make it even tractable, you would want to make a couple of simplifications. Firstly, imagine the ball is coming straight towards you as a fielder, so you are not going to have to move sideways at all. You either move forwards or you move backwards, or you stay where you are, if you are lucky and it is coming straight towards you. Of course, in practice, this is the hardest situation. If you are a fielder and the ball is coming slightly laterally to you, you have more information because you can see the arc of the ball and you actually do better at catching it, whereas if it is just coming straight towards you, it is a little harder – there is less information.

What do you do? There is a curious consideration here. If you look at the equations of the projectile motion - we shall ignore the effects of air resistance, so there is none of this drag effect and there is no lift effect, only the simple parabolic motion - then you can see that there are two things that could happen. It could be that you need to walk forward, at a steady speed, so the ball is going to come to rest there – you are going to walk forwards so you are in the right place; or the ball is going to go over your head, so you want to walk backwards to be in the right place; or you might just stand still.

These are pictures of the tangent of the angle that the ball is making with you. If you look at the ball and measure this angle (theta) that it is making as it comes towards you, then the thing to do is to move so that the rate of change of that angle, with time, is a constant. So, if this angle changes, increases with time, then the ball is just going to go over your head and you will not catch it. If it decreases with time, this is the trajectory of the ball that is just going to hit the ground in front of you and you will not catch it either. But the ball that is going to end up in the same place as you are, at the same time, is the ball where the rate of change of tan theta is a constant, and you will then end up in the right place, at the right time. So the formula to give your robot for catching the ball is to look where the ball is and move such that the rate of change of the tangent of the angle between the ball and you is a constant, and you will catch it. This is a strange but rather simple instruction that gets you in the right place, at the right time.

Unfortunately, if you add air resistance to the problem, you will remember that the trajectory is no longer a parabola. What happens is that, even if you are in the right place, at the right time, this tangent of the angle will increase with time, so this simple rule no longer works. What the optimum strategy should be is a challenging and interesting question.

Peter Brancazio, an American sports scientist, introduced an interesting consideration. He felt that somehow the circuits in the brain that are involved with balance - normally associated with hearing - might play a more significant role here than those involved with vision, because they have a different structure, they work faster, and they are more basic. He does not mean that you listen for the ball or take your cue from the crack of the ball against the bat, because exactly the same considerations hold if someone was just throwing the ball silently.

However, there is a little experiment that you can do. If you put your finger in front of your head like this, and you move your finger from side to side, you will notice that you cannot focus on the finger. But if you hold your finger still and move your head from side to side, your finger remains completely in focus. Try it.

In the second case, where the finger remains in focus, you are using circuits associated with hearing and balance; you can see that somehow there is much more going on – there is the possibility for processing information in a different way. That remains an unsolved, but rather interesting, problem.

Approaching the subject of impacts and bouncing balls, let us now look at a classic problem involving an impact between two moving objects. These objects are usually taken to be billiard or snooker balls. When I was at the University of California in Berkeley in the 1970s, an old Hungarian Professor in the department told us once that he often used snooker as an example when describing mechanics and motion. He said that he had never seen a game of snooker. He had never even touched a snooker cue. He knew a lot about the game, but his knowledge came entirely from books on mechanics.

Suppose that we have two masses, M and m. This one starts off with speed V and this one with speed v, and after hitting one another, this one has speed U and this has speed u. In the collision, momentum is conserved; the sum of the MV beforehand is the same as the sum afterwards, and there is a relation between the relative velocities in this direction. Beforehand, it is V minus v, while afterwards it is U minus u, and that is reversed, so U minus u would be minus e times V minus v - we switch the sign. E is usually called the coefficient of restitution, and it tells you how bouncy the mass is, how much is lost in the collision. If it is perfectly elastic, then e is one, and there is no loss of this component of the velocity. If one of these masses was a lump of putty that just stuck to its surface, e would be zero and there would be no rebound at all.

A simple situation is where one of the balls is stationary. If you were hitting a golf ball with a club, this would be stationary. So, initially v would equal zero; if you then solve these two equations together, you can arrive at two simple formulae which tell you what the speed would be of this ball as it is hit away, and this would be the final speed of the object that is hitting it.

Let us use some numbers. In the case of this golf ball, e is about 0.7, so it is fairly bouncy. It has got a small mass - about 0.46kg - and a typical golf club head would have a mass of 0.2kg. In a moment, we shall use calculations to demonstrate why this is a good mixture.

If you put these figures into these formulae, you get some real numbers. When hitting the ball, the club head will move at about 15m per second, which will result in the ball moving at about 34m per second. As you can see, these formulae work quite well to describe something as simple as hitting a golf ball.

Here is an interesting question: if we looked at these equations and knew the mass of a golf ball, could we work out what should be the best mass for the club head? How should we design the club head? Or, similarly, if you were given the mass of a club head, what would be the best mass for the golf ball?

This is similar to the shot putting problem in that there is a hidden constraint here. As you increase the club head mass, it is going to have an effect on how hard or how fast you can swing it. So a very high club head mass might give you a big speed for the ball, but it will be much harder to swing it. Consequently, you have to experiment in order to test that hidden constraint. A series of experiments was carried out which looked at the logarithm of the speed of the club head as you swing it against its mass. You can see from this graph, as you would expect, as the mass of the club head goes up, the speed you can achieve goes down. This slope is about 5.3. This means that the speed is inversely proportional to about the fifth power or the fifth root of the mass. That is our hidden constraint.

Here is our formula for the final speed of the ball, in terms of the mass of the club head, the speed of the club head, the coefficient of restitution – which is about 0.7 - and the sum of the masses. We put that in here, and see that the final speed depends on the mass of the club head and the mass of the ball. So, what we want to know is: when is this a maximum? Which ratio of M over m is a maximum?

That is a simple calculus problem. We just work out DU by DM, the differential of U with respect to M, and set it equal to zero. When we do that, we find this is the condition for the speed to be a maximum. The ratio of the two masses, the mass of the club head over the mass of the ball is n minus one, where n is this number, which is 5.3.

Plenty of golfing ‘trial and error’ required for this. I do not play golf, but I see people using little lead strips to change the mass of their club head, small adjustments in situ. If you have a ball with a mass (m) of 0.046, a fairly standard ball, you could predict that the optimal mass of the golf club head to make the launch speed of the ball as great as possible is 0.2kg, and that is about right.

Whenever you have an object with a heavy bit on the end for hitting things, whether it is a tribal war club or a hammer or a golf club, this type of analysis applies. You can work out the optimum ratio of the mass of the club head to the thing that you are hitting if you want to hit it as effectively as possible.

The overall efficiency in this case – that is, the amount of kinetic energy, half MV squared, in the club head that gets transferred to the ball - is not very great. It is approximately 43 percent. Of course, Tiger Woods, when not busy doing other things, surpasses most of the figures that we have been looking at. For various reasons, he can achieve tremendous drive lengths. He can achieve enormous velocity when hitting the ball, probably more than most of his competitors, because he has enormous flexibility. The back-swing goes far further than for you or for other professional golfers.

We shall now move onto something else related to impacts, something known as the centre of percussion. It has a fairly familiar application - in cricket - and a not so familiar one, on snooker and pool tables.

What is a centre of percussion? It is really providing an answer to the question: if you hit a ball with a bat or a racquet, where is the best place to hit it?

In the case of a cricket bat, there are a lot of complicated considerations involving normal modes of vibration of the bat, but there is one factor that is especially dominant and interesting in this case (tennis racquets are more complicated still).

Imagine that you have something like a bat (a rod in this diagram, suspended at the top). You can carry out an experiment like this by just suspending a cricket bat on a hook, and then hitting it in some way, or throwing a ball at it. There are two things that happen to the bat: its centre of gravity moves as a whole, so the bat moves as a whole – we are assuming there is no friction; also, there will be a rotational effect, meaning that the bat will be displaced from the vertical. So there are two aspects of the motion: the rotation tends to move the top backwards, but the overall motion - what we call translational motion - tends to move the top forwards. Both those effects create a force on the handle, in opposite directions. What we would like to know is whether there is a way of hitting the bat so that these forces cancel out, so that there is no net force on your hands at the handle?

If you play cricket you will know that if you hit the ball wrong, you get a horrible of reaction on your hands gripping the bat; even if you hit the ball almost right, you still feel a reverberation. But is there an ideal circumstance that means, if you hit the ball right, you will not feel that reaction?

There is, and it is called the centre of percussion. So these two effects - the translational force of acceleration and the movement of the bat - cancel each other out if the force is applied at this spot which, in this case, is about seven-tenths of the distance down the bat.

We shall now look at a real cricket bat, with a view to calculating that distance. If we have got a long rod, like a bat, with a length of 2r, then the place to hit it so that you get no reaction is at a distance h plus its moment of inertia divided by its mass times r away from this pivot. So, for a uniform rod, which is what a cricket bat really looks like, this moment of inertia (that is, its tendency not to be moved) is one-third MR squared, and so h is 4r over 3, which is two-thirds of the total distance down here, 2r. So this is where you should hit your cricket ball, about two-thirds of the way down the cricket bat from the handle. Of course, this is not a perfect rod. There are little details here. The moment of inertia is a little bit different, the shape of the bat at the back is a little bit different, and you could do this calculation in much greater detail, but this is the basic idea. Here is the perfect spot.

We can now apply this idea a little more carefully in a less familiar situation. If we were to hit a billiard or a snooker ball with a cue, the same idea applies. There is an ideal place to hit the ball with the cue such that, when you hit it, the combination of the spin and the translational motion cancel out, down here, and the ball does not slip at all – it just rolls.

If you watch people playing crown bowls on grass or, nowadays, indoors, the skill in that sport is to deliver the ball by changing the angle at which you throw it so that you automatically create this type of situation; the ball rolls right from the word go, it does not skid and then do something unpredictable.

In this case, suppose you hit the ball right through the centre, there will not be any sliding at the bottom. The ball will just move as a whole. If you hit it above the centre, at any old point, it is going to slide and it is going to rotate, to spin back. As it rotates, these points down here move in that direction. We want to know whether there is a place that you can hit it, above the centre, such that the translational velocity in this direction cancels out the rotation in the opposite direction at the bottom, causing the ball to roll straight away and not slip.

Once again, this amounts to calculating the centre of percussion. The sliding speed that you get from hitting it is calculated by dividing the force (f) by the mass, times the duration of the force. The rotation speed in the opposite direction is the force times the radius, times the duration, times the distance between h and the radius, divided by the moment of inertia. This gives us the 'h = r + I/Mr' condition that we just looked at, but now we are using a sphere instead of a rod. For a sphere, the moment of inertia is two-fifths Mr squared, which leads us to predict that the location of the ideal spot is seven-tenths times the ball's diameter, 2r. That is the place to hit a snooker ball, or a pool ball, if you just want it to roll truly and not slide.

A little while ago I thought that, by looking at this, I would be able to calculate for myself the ideal height of a cushion on a snooker and billiard table. My thinking was that if you made the table correctly then the height of the nub from the cushion, the bit that protrudes, ought to be such that whenever a ball hits it, it would roll true off it and not skid. The logic should surely be to make the height of the cushion equal to the centre of percussion, equal to this seven-tenths of the ball diameter.

I discovered that there were rules about the construction of snooker and billiards tables that specify quite precisely the height of that cushion, and I was surprised to find that my intuition was wrong - it was 0.635 plus or minus 0.10 times the ball's diameter. I asked David Alciatore, an applied mathematician in Arizona and something of a world expert on the mechanics of billiards and snooker, and he told me, "You are right with this calculation, this is where it ought to be, but it is made just a little bit smaller so that, after the ball bounces off the cushion, it is not pushed to roll so quickly and so hard into the guttering of the table, in order to stop the wear of the table very close to the cushion." So this tiny difference aims to reduce the wear on the table near the cushion – very unsatisfactory, but that is how it is!

You now know how to hit the ball, as well as the reason why the cushion is at about that height. To give you a good approximation, you should hit the ball at the same height of the nose of the cushion above the table. Any money you win from this tip, please send ten percent to me!

We have talked enough about these sorts of impacts. Finally, we shall look at bouncing balls and, ultimately, the superball.

Here is a time-lapse photograph of a simple bouncing ball. As it hits the ground, the coefficient of restitution (e) is less than one, so when it bounces back up its speed is not as great as its initial downward. As a result, it does not reach as great a height, and each successive bounce is lower and lower. The trajectories in this example are almost parabolic. This is the simple motion of a ball, with no spin.

Once you add some spin to the ball, much more complicated things happen when the bounce takes place. As this example illustrates, when a ball impacts with back-spin or top-spin, you have to take into account the change of the spin as well as the change in the speed.

We are about to turn to this rather strange object called the superball, which you probably remember from your childhood. I shall risk throwing one here. They are much tamer than they used to be when I was young, but the history of these objects is amusing. They were only invented in 1965, so I can remember having one very soon after they appeared in this country. They began life with the rather unappealing name, "Highly Resilient Polybutadiene Ball". This is its US patent number. They were manufactured by a company in America called Wham-O, which places it very definitely in the 1960s!

Its coefficient of restitution is very high – bigger than 0.7. If I went outside and threw this very hard, it would bounce over the building. It is quite different to tennis balls or other types of balls that we are used to. It is produced in an unusual way, baked in a particular way with unusual things at the centre, highly compressed, but the key to its behaviour is its roughness. This is an example of the motion of a rough ball, and it behaves quite differently to a billiard ball. It never slips, so the motion is always of no slip when it hits the ground, and it will have back-spin and can behave in a quite unusual way. I cannot do it here, but imagine that I had a table like this, going in that direction, and I threw this ball under the table. If it was a tennis ball, it would just bounce and keep going in that direction, but this ball, after two bounces, would come back. If I throw it at the right speed, I can throw it down there and it will come straight back at me. It does not behave like a tennis or billiard ball. Each time it bounces, the direction of its spin reverses.

There is another game that you can play if you have more than one superball, but I am not going to do it as it is rather dramatic. I have done a similar demonstration before with a ping-pong ball and a basketball. If you drop two balls, one on top of the other, then the top ball will bounce nine times higher than it would if dropped on its own. If I do that here, it will hit the ceiling, but you can try this afterwards; try it with any two balls and you will have a surprising outcome.

Here is another curious fact that I discovered while preparing this talk, something that I did not know until last week. You will be familiar with the so-called Super Bowl, American football's version of the cup final. Lamar Hunt, founder the American Football League, created the name 'Super Bowl' after watching his children spend an afternoon playing with a superball, very soon after it came on the market. That is the origin of the American Super Bowl – it is just a deviation from superball.

Here are some pictures showing what goes on with this type of ball. If you were to throw it against the floor at a fairly sharp angle like this, it will bounce up. However, from here it will not do what a tennis ball would do, which is just to keep on doing that (see diagram), but it will come back, and after one or two bounces, it will essentially have returned to where it began, except for losing a small amount of energy. Here is another picture. It is possible, if you throw the ball with the right composition at the right angle, to make it come back on exactly the same trajectory that it went in on.

When we throw one of these balls, it has a bounce here and comes out. We have to take into account the rotation that the ball has, so there is a back-spin here and, in this case, there will be a forward spin.

It is an interesting mental exercise to ask yourself what would happen next here. The overall motion is symmetric in time (in the sense that Newton's Laws of Motion do not distinguish the future from the past), and so what has happened is that the velocity in this sense, with rotation in that sense, has changed the sine of this velocity and this has become a top spin. It is going forwards, so after it bounces, it is just the time reverse of this, and you have got to end up with this situation coming out, just the time reverse of what you had before. Therefore, after two bounces, the ball returns to looking how it did before.

So, what happens in these collisions? What is the simple way of looking at it? You can study them in detail. You have to conserve the total energy of these collisions, which has two components: the ordinary kinetic energy, m is the mass of the ball, v is its speed, half mv squared; but there is also the rotational energy of the ball, and that rotational energy looks like a half times its moment of inertia, two-fifths mr squared, times the angle of velocity squared. This is conserved at each bounce, assuming that there is no loss of energy in heat and sound (which we shall just ignore). We shall also assume that the coefficient of restitution is one. It is certainly not very far off.

At this point, there is an angle of momentum, which is Iω for the ball, and a back rotation, MR times the linear speed (v). This quantity, the total angle of momentum of the situation before and afterwards, is conserved about this contact point. Those two simple conserved quantities must be considered alongside the notion that the vertical velocity that comes out is minus e times the velocity that goes in, coefficient of restitution – V(out) = -eV(in); this could be different for the vertical component and for the horizontal component. The key point to note is that there is no slip at this point, that this ball is rough. That is what we mean by rough. These are the collisions of a maximally rough ball that does not slide at the point where it hits the wall. At that point, the velocity is reversed in this way.

Here are a couple of pictures of situations with various velocities. When coming in with no rotation at 20 degrees and at low speed, you go out with a top-spin. Here is what happens with a tennis ball. It is fairly similar, but you go out with a smaller spin because the coefficient of restitution is smaller. However, if you throw it at high speed and you have more back-spin, this is what the superball does; it just comes back out in the same direction, with a top-spin, equal and opposite to the starting back-spin. The tennis ball on the other hand, with the back-spin, comes out in the other direction, with just a slight top-spin, so it is quite different.

Here is what would happen if you tried to play superball snooker. In ordinary snooker, if you hit the ball at this angle around this table, the angle of reflection is equal to the angle of incidence, meaning that you would come back to the same point if you hit the ball hard enough and under perfect conditions. On this same shaped table, if you hit the superball, you go up there, you come back in that direction, and you go there and you come back to the same point - but you are going through a very different sequence of trajectories. So playing snooker with the superball is quite a challenge!

I hope that I have given you a new intuition about what happens when you throw things, when you try to catch things, when you try to hit things, when you try to get out of the way of things; an idea that the motion of trajectiles and the motion of impacts does not just stop with the simplest things that you learn at school, with parabolic motion or simple conservation of momentum. There are other simple things in the world that we come across every day which behave quite differently, and, in many respects, are understood, but in some respects, are still not completely understood.