Lebesgue Integrals and Measure Theory

This essay will be best enjoyed if you already have an understanding of the traditional integral because it involves a certain twist in that understanding - some undoing - that I hope you'll find as pleasing as I do. It's not required knowledge but without that original step forward being thoroughly entrenched, the step backwards is harder to appreciate.

Anyone who already has this knowledge is welcome to skip the first section.

Integrals

What's the point? They're vital in everything from designing aeroplanes to modelling the stock market. Broadly, it's how we sum up a whole spectrum of values. The tricky part is that there are an infinite number of values in that spectrum so we can't just add them up like we would normally. Lucky for us, a clever man called Leibniz came up with a workaround in 1684. If we plot the values on a graph so their value is just their height then the sum of all the values is just the area under that graph. Finding the integral is just finding that area - and he knew how to do that.

The modern method is to start by approximating the area under the graph with a series of rectangles, all with the same width. The total area of these rectangles (their width times their height) is roughly the same as the area we want.

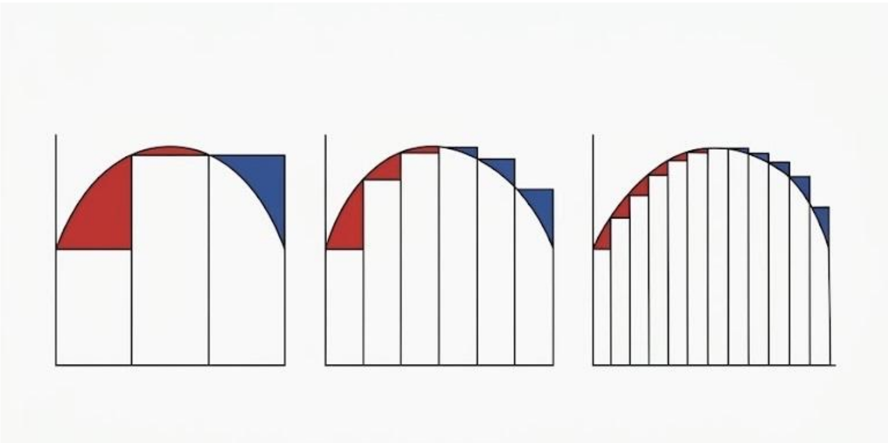

As you can see though, it's not perfect. Those blue regions and red regions are the discrepancies.

Now imagine using more and more, thinner and thinner rectangles. Intuitively, you'd guess that the blue and red areas would decrease, giving us a more accurate value for our area under the curve.

This intuition is fine if you're just getting to grips with the concept but we might want to make it more rigorous. We can think of it like this:

Halving the width of each rectangle would double the number of blue/red regions but quarter their areas on average, so half their total area. This means the thinner the rectangles, the lower the miscounted area and therefore the more accurate our value.

In fact, if you make these rectangles as thin as you want, your approximation can get arbitrarily close to the real answer.

The Problem

This approach to integration works for most 'normal' functions you come up with (e.g. $y = x^2$, $y = 2x+3$, $y = 3.5^{(5x)}$) but around the turn of the 20th century mathematicians started thinking up weirder and more wonderful ones. Dirichlet, in particular, made a real stumper.

His function took in every number between 0 and 1 and outputted 1 if it was rational (can be written as a fraction like 273/425) and 0 if it was irrational (can't, like pi).

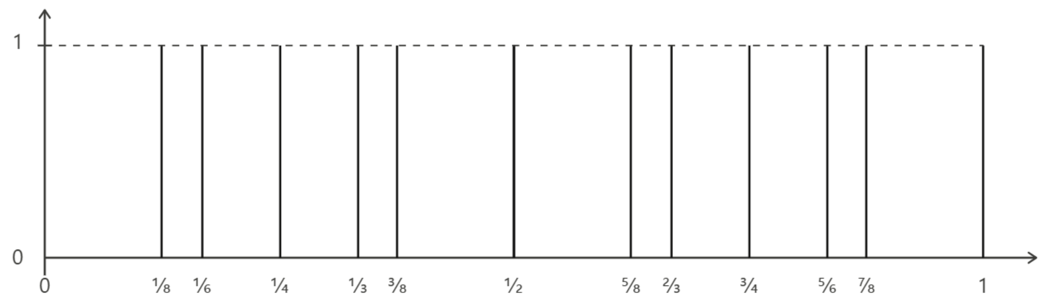

Not all rationals are shown as there are infinite

So let's try to integrate this. First, we'll have to place some rectangles on our shape and here we run into a problem. Inside the rectangles there are rational points and irrational points, so what should the height of it be – 1 or 0? Maybe we plough on anyway, hoping the problem will sort itself out as we make these rectangles thinner and thinner. After all, there were multiple values within our rectangles in the first example.

The problem is that within any region, no matter how small, there's infinitely many rationals and infinitely many irrationals. That means no matter how thin we make them, the height doesn't converge (get closer and closer to) a specific value. That's what's different here. We might have broken the integral, but in doing so we've taught ourselves something valuable about how they work.

Lebesgue Integrals

The solution to this is very elegant and it comes in one fell swoop; what if we made the rectangles horizontal instead of vertical. It seems so simple and yet it took mathematicians over 60 years.

The problem arose because we couldn't find a way to split the x-axis up. There's no way to do it that doesn't leave each rectangle with an infinite number of values, jumping between 1 and 0. So what if we don't? If we use horizontal rectangles instead of vertical, we don't have to split up the x-axis at all. We split up the y-axis.

Personally, I would never have come up with this reframing myself and I thought I had a pretty good understanding of the integral. I did but it was a different kind of understanding – understanding from the inside, like only ever seeing a house from the inside and claiming to know it fully. I knew what the integral was and how it worked but not the choices that were made in its creation, the ways it could have been different. I didn't understand the wider space that it occupied, so to speak. You didn't either, most likely. It's a privileged position to have been shown all this. Lebesgue had to figure it out for himself.

I can only imagine him stroking his chin, thinking to himself 'I need to look at this from another angle… a perpendicular angle!'

Anyway- back to the maths.

For Dirichlet's Function, there are only really two rectangles, one for each possible y-value. Finding the integral is the same as finding their total area.

The first rectangle barely counts because its height is 0 and so its area is also 0. The second one though, has a height of 1 and its width is the length of all the rationals between 0 and 1.

I expect alarm bells will be going off in your head at that. How on earth do the rationals have a length? We need a new way of thinking about length entirely. Luckily, there's a field of maths exactly for this. You should hear the sound of a door screeching open now, as we momentarily leave Lebesgue Integrals behind and enter the separate but connected room of Measure Theory.

Measure Theory

The job of this field of maths is to give every set a size – to assign it some number that's meaningful, or helpful in some way. It's a broad definition for a broad field.

There are so many different ways to assign a size, called measures, because these measures are just tools, created as needed. They can't be right or wrong, just helpful or not helpful. What ties them together is something called the principle of countable additivity. It says that if a set is countable (you can order its elements), its size is the same as the total size of all its elements.

The measure that's helpful for us is Lebesgue's own that he created for his integral. Its geometric, taking sets as shapes and assigning them sizes that we might think of as length, area or volume. It's more flexible than that though. Eventually, we'll tease out of it a length for the set of rational numbers between 0 and 1 and then we can find the area of that pesky rectangle.

The Lebesgue Measure of a line segment from a to b is just b minus a which matches intuition; the length from 3 to 5 is 5 minus 3, or 2.

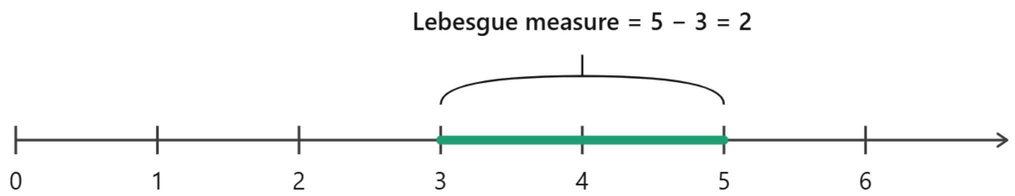

It gets a bit more interesting when we consider single points, say the point x. If we think of it as a line segment, it starts and ends at x, so its length is x minus x or 0. The size of a point is 0. We're very close to working out the length of our set now.

All that we need to do now is invoke the principle of countable additivity from before. First though, let's check it applies. Is the set of rational numbers between 0 and 1 countable? Remember countable means we can put its elements (in this case the rationals) in an order, so we can talk about the first in our set and the second and so on.

A first thought might be to order our rationals from smallest to biggest. You'd start with 0 but it's unclear where you go after that. If you propose any rational, say 1/2000, I can always find a smaller one just by adding 1 to the denominator. There is no second smallest rational.

All hope's not lost though. We know how to order the positive integers (1,2,3 etc.) and we can cleverly use this to do the same with the rationals. Every rational is, after all, a combination of two integers – its numerator and its denominator. This is our way in. We can begin by listing out every rational between 0 and 1 with denominator 1 (namely 0/1 and 1/1) and then every one with denominator 2 (1/2), ignoring repeats like 2/6 and 1/3. In this way, we can order all our rationals, meaning the set is countable! We can use the principle of countable additivity, meaning we can add up the sizes of each of them to get the size of the total set, meaning we can find the Lebesgue integral, meaning we can finally take a rest.

Conclusion

We're so almost there. We've journeyed through two different fields of maths. Now all that's left to do is piece it together.

To find the size of the set of rationals between 0 and 1 we add up the size of every rational that makes it up. As single points, this is 0 and so the total size is 0 + 0 + 0 + 0 and so on, which is 0. The length of our rectangle is 0 and so its area is 0 and the integral is 0.

This might not be the most exciting. But that's fine because answers aren't what make up mathematics, they are only the objects that populate it. Maths is the concepts, the understanding, the building in which the answers are housed. That's why the real revolutions in maths don't occur when we find out new answers, but when we see things in new ways. Lebesgue isn't remembered because he solved Dirichlet's problem. He's remembered because of the new concepts he came up with, because he laid one new rectangular brick in the building that is mathematics and measured it against his own tape measure.