Let us now go through an example.

Suppose that we have an unknown random variable Theta that has a uniform distribution between 4 and 10.

We observe some other random variable X that's related to Theta according to the following model.

This is the conditional distribution of X given Theta.

For any given value of theta, X is going to take values between theta minus 1 and theta plus 1.

And the conditional distribution is uniform on that range.

One way of thinking about this particular observation model is that what we observe is the true value of Theta plus some noise term.

And this noise term is uniform on the range from minus 1 to plus 1.

So given a value of Theta, we may observe anything, because of noise, that's up to one lower or one higher than the true value of Theta.

And if we take this description, actually, this random variable U has this distribution no matter what Theta is.

And therefore, U is independent of Theta.

But in any case, this particular interpretation will not matter.

Let us see how do we proceed.

In Bayesian estimation, the first step is always to put our hands on the posterior distribution of Theta.

And to find the posterior, we can start by first finding the joint.

So let us look at the x theta plane.

That's where the joint distribution is going to live.

And our first step will be to locate those values of X and Theta that are possible, given our description.

From this model here, we know that theta minus 1 is going to be less than or equal to x.

And x is going to be less than or equal to theta plus 1.

And we translate this into two inequalities, namely that theta is less than or equal to x plus 1, and from here, that theta is larger than or equal to x minus 1.

So these are the constraints that we have on the possible values of x and theta.

So here we plot the line where theta is equal to x plus one.

And here we plot the line on which theta is equal to x minus 1.

And these two inequalities that we've got here tell us that we need to live somewhere in between those two lines.

In addition, we have the fact that theta lives between 4 and 10.

And that places these two limits as well.

So to summarize, this shape here is the set off all possible x's and thetas.

Outside this shape, the joint PDF is going to be zero.

What is it going to be inside here?

Well, because the prior is uniform, that is, it is constant, and the model is also uniform, to obtain the joint we multiply these two.

And since they are constant, we obtain a joint that's also constant.

So the joint PDF is equal to a constant over that set.

We can easily calculate the area of this set.

It is 12.

So the joint is equal to 1 over 12 inside this set.

And of course, it's 0 elsewhere.

So we have a uniform joint PDF.

Now, let us look at the posterior.

If I tell you that X takes on this specific value, this means that we now live in this universe.

And it means that all of those thetas are possible.

The posterior distribution is going to be a distribution that tells us the probabilities of these different thetas.

What kind of distribution is it?

Well, we know that the conditional is just a section out of the joint but keeps, otherwise, the same shape.

Since the joint is constant, it's uniform over that set, it means that the posterior, or the conditional, is also constant over that set.

So the conclusion is that the posterior distribution of Theta is a uniform distribution on this set.

Given this knowledge, what is the conditional expectation?

The conditional expectation of a uniform is just the midpoint of that uniform.

And so this is going to be our estimate of Theta, the conditional expectation of Theta, given the observation that we have obtained.

And then a similar argument applies no matter what x we have obtained.

For any given x, our estimate is going to be the midpoint of the corresponding interval.

So what kind of shape do we get by doing this, by joining the mid-points?

It's going to be a straight line over this region.

It's also going to be a straight line over this region except that, because of the change in shape, it's going to be a straight line with a different slope.

And similarly, in this region, it's also going to be a straight line with a different slope.

So what have we plotted here?

For any given value of X, we have plotted the corresponding conditional expectation of Theta given that value of X. And as a function of x, this gives us a certain curve.

And this blue curve that we have calculated is a particular function of x.

And we can think of this function g as being our estimator.

So the way we're going to be processing the data will be that whenever we obtain an x, we apply this particular

function g.

And we come up with an estimate.