

In this lecture we complete our discussion of multiple continuous random variables.

In the first half, we talk about the conditional distribution of one random variable, given the value of another.

We will see that the mechanics are essentially the same as in the discrete case.

Here, we will actually face some subtle issues, because we will be conditioning on any event that has 0 probability.

Nevertheless, all formulas will still have the form that one should expect.

And in particular, we will see natural versions of the total probability and total expectation theorems.

We will also define independence of continuous random variables, a concept that has the same intuitive content as in the discrete case.

That is, when we have independent random variables, the values of some of them do not cause any revision of our beliefs about the remaining ones.

Then, in the second half of the lecture, we will focus on the Bayes rule.

This will be the methodological foundation for when, later in this course, we dive into the subject of inference.

The Bayes rule allows us to revise our beliefs about a random variable.

That is, replace an original probability distribution by a conditional one, after we observe the value of some other random variable.

Depending on whether the random variables involved are discrete or continuous, we will get four different versions of the Bayes rule.

They all have the same form, with small differences.

And we will see how to apply them through some examples.