Welcome to the first lecture of this class.

You may be used to having a first lecture devoted to general comments and motivating examples.

This one will be different.

We will dive into the heart of the subject right away.

In fact, today we will accomplish a lot.

By the end of this lecture, you will know about all of the elements of a probabilistic model.

A probabilistic model is a quantitative description of a situation, a phenomenon, or an experiment whose outcome is uncertain.

Putting together such a model involves two key steps.

First, we need to describe the possible outcomes of the experiment.

This is done by specifying a so-called sample space.

And then, we specify a probability law, which assigns probabilities to outcomes or to collections of outcomes.

The probability law tells us, for example, whether one outcome is much more likely than some other outcome.

Probabilities have to satisfy certain basic properties in order to be meaningful.

These are the axioms of probability theory.

For example probabilities cannot be negative.

Interestingly, there will be very few axioms, but they are powerful, and we will see that they have lots of consequences.

We will see that they imply many other properties that were not part of the axioms.

We will then go through a couple of very simple examples involving models with either discrete or continuous outcomes.

As you will be seeing many times in this class, discrete models are conceptually much easier.

Continuous models involve some more sophisticated concepts, and we will point out some of the subtle issues that arise.

And finally, we will talk a little bit about the big picture, about the role of probability theory, and its relation with the real world.