Quantum theory appendix 2: the meaning of probabilities

Not everyone who has thought about probability agrees with the view that it is about large ensembles of events. These dissidents argue that there are many cases where we use probabilities to refer to individual cases. For example, when we place a bet on a sporting event, we are basing our wager on a prediction of what will happen in an individual race or game. In this circumstance we talk about odds, which determine the payoff of the bet. These odds do not refer to hundreds of thousands of similar games. The bookies do not wait for hundreds of games to go by and then pay off the bettors. And how could they, for are two races or soccer games really the same?

The dissidents also point to a weakness in the interpretation of probabilities as referring to many cases, which is that there is always a chance that the predictions, being only probabilistic, do not come through. What does it mean to predict that a fair coin will turn up heads 50% of the time? If one tosses the coin 100 times then one expects roughly 50 heads, but how roughly. If you get 49 or 52 heads you will still be happy to consider the coin fair. But what if there are 37 heads out of 100 tosses? Does this mean that the coin cannot be fair? Not necessarily, for a very small fraction of the time a fair coin will come up heads 37 out of 100 times. How can one tell if the coin is fair? Do more tosses? At 100,000 tosses the chances of getting 37,000 are much less than the chance of getting 37 out of 100. But a small proportion of the time this will happen as well, so how does one know this is not that rare case?

You can see where this is going. The only way to be sure a prediction about probabilities is coming out right is to test it an infinite number of times. But this is the one thing we can't do. So it seems like there is no way in the real world where we make tests only finite numbers of times to contradict a probabilistic prediction. In that case, could probability really be about large numbers of cases?

There is a straightforward answer to this objection, which is to ask for a little modesty of purpose. We human beings live finite lives, within which we can use good advise far more than certainty. So a prediction that will be good in a vast number of cases is just fine for our purposes, even if every once in a blue moon it fails due to the fact that our experiments deal only with finite ensembles. There are many reasons why real predictions about real experiments go wrong, and the funkiness of probabilities applied to finite cases is rarely the cause of a problem, once the number of cases is large enough. And for small numbers, its a quantifiable, controllable risk, which is not the same for all the problems that can bedevil experiments in real life. So far all practical purposes, large ensembles are as good as infinite ensembles.

But isn't our purpose more elevated than this? Don't we seek a fundamental theory? Yes, but how can a probabilistic theory be fundamental? Isn't this indeed the whole point? Maybe there is a real fundamental theory out there, but its not quantum mechanics.

Notwithstanding the “for all practical purposes” answer to making sense of probabilities, a growing number of mathematicians and philosophers reject the idea that probability is about proportions of large numbers. They insist that we use probabilities in a way that are meaningful for single events. When we assert that there is an 80% chance or rain tomorrow we are making a statement about tomorrow that we act on, irrespective of what happens on similar days. We act on a forecast of 80% rain by bringing an umbrella and raincoat. We consider postponing the spring garden party. In other words, we make a bet based on our evaluation of the costs and benefits of different courses of action for tomorrow, given the forecast. To decide whether to postpone the garden party we consider the costs-mainly but not only social-of the postponement, together with the costs of not postponing if it does rain. We make a bet.

These kinds of bets are hard because there is no way of knowing today, when the decision has to be made, what the right call will have turned out to be. Either way there is a cost. If you postpone and it turns sunny tomorrow you will have irritated the many guests who were forced to change plans. If you don’t postpone and it rains, no one will have a good time.

In this context the practical import of the number in the forecast is to influence your bet. The higher the chance of rain forecast tomorrow, the more you are willing to bet the right answer is to postpone.

One school of thought about probabilities hold that all there is to probabilities of single events is their influence on betting. This school, who are called Bayseans, holds that the entire meaning of a probability for a single event is that it is a measure of how much you would be willing to bet on that event occurring.

Certainly the odds in wagers are probabilities whose meaning is cashed out directly as betting odds. The claim of the Bayseans is that all use of probabilities is as short cuts for asking how much you are willing to bet on some outcome. The only exception they will give is if there are really an infinite number of cases, so that relative frequency probabilities can have a firm meaning. Otherwise, even if there are many cases, you are betting that those finite number of cases will be typical.

How much I am willing to bet on rain tomorrow is however not a property of anything in the atmosphere or weather system. How much I am willing to bet on a coin coming up heads three times in a row is not a property of the coin. This definition of probability as betting odds makes probabilities out to be descriptions only of our belief or knowledge about a system.

Could the probabilities in quantum mechanics be betting probabilities of this kind? In the last few years a number of physicists have proposed precisely this. According to them, the probabilities for different outcomes of an experiment on an electron is not a property of the electron at all, it is a property of our knowledge and belief about how the electron will behave in experiments we might impose on it. But then the same must be true for the wave-as the probability is the square of the height of the wave. The wave associated with an electron is from this point of view something that has no objective meaning-it is a way of symbolizing our knowledge about the electron and its interaction with various devices we use to measure it.

Chris Fuchs, who is one of the most eloquent advocates of this application of the Bayesian notion of probabilities to quantum theory, puts a pragmatic twist on this move. Pragmatic in the sense of the American pragmatist philosophers, William James, Dewey and Charles Sanders Pierce. To Fuchs, the wave in the quantum description of an electron is a tool we observers use to keep track of and reason about the odds we would be willing to bet on the outcomes of the different measurements we might make on the electron.

This viewpoint has become commonplace in a recently emergent field of research which is called quantum information theory. Originally this area was motivated by problems that arose in attempts to design and build a quantum computer. A quantum computer is one in which properties of quantum systems are used to accelerate the power of a computing or communications device. This is a fascinating direction which may yet pay off hugely for technology, but this is not our concern here. What is interesting is that a new way of talking about quantum mechanics arose from seeing a quantum system through the lens of computer science and information theory.

A computer manipulates and stores bits of information, stored in a memory. A bit is an observable that can be on or off. A quantum computer would treat this as two possibilities that can be added or superposed like a wave. As the addition of waves is happening in an abstract space of two possibilites-on and off- rather than in real space, there is a more general notion of a wave, which is a quantum state. A quantum state is a way of assigning numbers and phases to different possible results of an observation in a way that can be added together. This more general notion of a quantum state includes the electron waves that we were just discussing, but it includes much else.

One of the things that can be described with this general notion of a quantum state is a quantum version of a bit of information. This can be on, or off, or a continuum of possibilities in between. The quantum state assigns a complex number to each distinct possibility, which is called its amplitude. The square of an amplitude for something to be observed is the probability that it will be seen. A quantum computer is then a machine that can manipulate and store these general quantum bits. These are called qbits.

The pragmatist or Baysean point of view arose when quantum computer scientists tried to make sense of the notion of probabilities that arose in their descriptions of quantum computers. It turns out that a quantum computer, being quantum, is partly stochastic, ie the answer it can give to a particular computation is not a single answer, but a probability distribution of answers. To make sense of this for a single computation, they needed to employ a notion of probability that can be applied to single events. They only one available is the Baysean notion of probabilities as bets.

This can certainly seem confusing. Its not if we remember Neils Bohr's advice that we keep a clear distinction between what is out there in nature and what is part of our knowledge or description of nature. The quantum computer is certainly doing something real and objective out in the world. But the description we have of the quantum computer is an aspect of our knowledge we have of the world. When we employ a quantum state to make a probabilistic prediction for the outcome of a quantum computation we are expressing our betting odds that that will be the outcome.

This point of view is indeed an extension and a deepening of Bohr's way of understanding quantum mechanics. Like the stochastic approach it denies the possibility that a quantum wave or quantum state is a complete description of an individual system or an individual measurement. They differ in what they assert the quantum description is about. The stochastic approach asserts that the quantum state and probabilities apply to a large ensemble of similarly prepared systems. The Bayesian approach asserts the quantum state is about a single system and not an ensemble, but it is nothing real, it is just a description of our knowledge of that individual system. In other words, quantum mechanics is a fancy way of placing bets on the outcomes of experiments we set up and perform.

Both of these approaches leave open the possibility that sometime in the future we may be able to discover a more complete description of what is going on when an individual electron or atom is interacting with an individual experiment. But they agree that quantum mechanics is not that objective description of an individual system. And they agree that if once we discover that deeper level of description, quantum mechanics may come to be understood as an approximation to it.