The bad jelly bean counter example

Note: this is an *optional* thought-exercise. I'm happy to discuss it with anyone who is confused by it, but do not stress out if you do not find it helpful. The details of this will **not** be on an exam. I just find it a helpful way to think about why a the confidence level of a confidence interval is **not** the probability that the true value of the parameter is contained in the confidence interval. The example is taken from: M. Jésus Bayarri and James O. Berger, James O. 2004."The interplay of Bayesian and frequentist analysis." *Statistical Science*. **19(1)**. 58–80. http://projecteuclid.org/DPubS?service=UI&version=1.0&verb=Display&handle=euclid.ss/1089808273

Imagine a jelly bean counter. If θ is the true number of jelly beans in your jar, then 50% of the time it will report a count of $\theta + 1$ and 50% of the time it will report a count of $\theta - 1$. In other words, it is always off by one; half the time it is too high and half the time it is too low.

You are only mildly curious about the number of jelly beans that you have, so you decide to use the counter twice. So $N = 2$ for your tiny study.

Before you do your measurements a statistician points out that your expected outcomes can be described in an expected relative frequency table (we'll cover the probability calculations involved in chapter 5):

	Low on count 1	High on count 1	*row sums*
Low on count 2	25%	25%	50%
High on count 2	25%	25%	50%
column sums	50%	50%	100%

This (excessively clever) statistician points out that if you construct a funky confidence interval that contains only one point: 1 + the minimum count you observe, then this will actually be a 75% confidence interval. This is the true because if the jelly bean counter is reports a number that is too low on the first or the second counting, then the minimum value + 1 will be exactly right. The possible measurement outcomes in which this confidence interval contains the true value are shown in red below. Note that they add up to 75% of experimental outcomes:

	Low on count 1	High on count 1
Low on count 2	25%	25%
High on count 2	25%	25%

Possbile scenario 1

Suppose, you conduct your counts with the faulty jelly bean counter and the machine reports $Y_1 = 432$ and $Y_2 = 434$.

Using the statistician's funky confidence interval, you can say that you are "75% confident that the true number of jelly beans is in the interval from 433 to 433."

But what is the probability that the true count is 433? It is 100% because there is no other number of jelly beans that could have resulted in those two counts (assuming that we are correct about how the counter misbehaves)

Possbile scenario 2

Suppose that you do the counts and you get $Y_1 = 501$ and $Y_2 = 501$.

The funky CI lets us say that we are "75% confident that the true number of jelly beans is in the interval from 502 to 502."

However doesn't it seem like the true count of 500 and true count of 502 are equally good explanations? So shouldn't the probability of the true count being 502 be equal to 50%? We'll discuss this next week[1].

Take home message

The confidence level of a confidence interval is **not** the probability that the true value of the parameter is contained in the confidence interval!

[1] it is actually a bit tricky