A confident dialogue

"I found that the estimated odds ratio is 5.0 and the 95% confidence interval is 3.4 to 7.3"
"Great. What does that mean?"

"There is 95% probability that the parameter (the true odds ratio) is between 3.4 and 7.3"
"Nope."

"Why not?"
"The unknown parameter is a number, say, 6.5 or 9.1, and in frequentist statistics a number does not form a probability distribution (it's a property of a random variable). Either the parameter resides in the interval [3.4, 7.3], or it does not. Therefore, the probability that the interval [3.4, 7.3] contains the parameter is either 1 (say, if the parameter were 6.5) or 0 (say, if it were 9.1). That probability is never 95% or anything other than 1 or 0."

"So where is this 95% coming from?"
"The explanation is a little long and a little abstract. Imagine you were able to replicate your study (i.e., the process that generated your estimate) an infinite number of times. Then, you would have obtained an infinite list of odds ratio estimates besides your estimate of 5.0. Right?"

"Right. I can imagine such a list of estimates that originate in one estimator."
"Now, imagine a parallel list of 95% confidence intervals that you would have computed for each estimate on that list, besides your interval [3.4, 7.3]."

"Okay. I can imagine such a list of confidence intervals."
"Then 95% of these intervals (which you can't really compute) will contain the parameter."

"Isn't that equivalent to claiming that there is 95% probability that *my* interval, [3.4, 7.3], contains the parameter (the true odds ratio)?"
"No. It's equivalent to saying that there is 95% probability that an interval generated by this estimator *will* contain the parameter. In frequentist statistics, "probability" is about the future, not about your already realized interval. Again, either the parameter is located in the interval [3.4, 7.3], or it is not. No probability may be attached (other than 0 or 1)."

"Still a little foggy for me."
"Perhaps an analogy would help. Suppose you bought this morning a lottery ticket where the probability of winning happens to be very high: 95% of the tickets are expected to win. The drawing is scheduled for tonight at 8:00PM. What is the probability of winning tonight (given that a ticket was purchased)?"

"95%, of course."
"Right. Now, it is 9:00PM and the drawing was done. The winning lottery tickets were determined, but you haven't checked the results yet. You hold the ticket you bought in your pocket, not knowing whether it was one of the lucky tickets."

“Okay.”
“Tell me, what is the probability that your ticket won?”

“95%?”
“No. The drawing was done. No frequentist chance is involved anymore. Either you won or you did not win. The probability of 95% is concerned with what *will* happen to all tickets, not with what *has* already happened to one particular ticket.”

“But I still feel that there is 95% probability that I won.”
“Ah… Now you begin to sound like a Bayesian. Have you ever heard about Bayesianism, a school of statistics and philosophy of science, which many reject (including me)?”

“No, I have not. I do remember, however, something called Bayes’ theorem.”
“Yes, that’s the foundation of Bayesianism. Now back to your feelings. You are actually saying that along with your ticket you are holding a strong belief in favor of the possibility that you won. You want to say that you *feel* 95% certain about having won. Feelings, however, are not the basis of the computation of a 95% confidence interval in frequentist statistics. If you want the Bayesian version, it is called a 95% credible interval, but it’s not the same computation. Bayesians allow for probability distributions of parameters, where the probabilities in questions are degrees of belief (often translated into betting behavior: how you would spread the bets on possible values of the parameter).”

“So what is the frequentist interpretation of the 95% confidence interval I have computed?”
“I am sorry to tell you that it has no interpretation whatsoever. Nothing. Nada (Spanish). Gornisht mit gornisht (Yiddish).”

“I see. I need to think about a 95% confidence interval as a random variable. This random variable forms a probability distribution that I can’t really depict. All that I have computed is one realized value of that random variable [3.4, 7.3], about which I cannot say anything probabilistically.”
“Exactly.”

“And what about other interpretations of a 95% confidence interval I have heard?”
“What interpretations?”

“My 95% confidence interval [3.4, 7.3] is a plausible range of values of the parameter.”
“Ah? I can add any value to your plausible range of values.”

“How?”
“Easily. You claim that anything outside [3.4, 7.3] is not a plausible value. For example the value 9.1 is not plausible (or less plausible). Right?”

“Right.”
“I will make it as plausible as 7.1 in a second. Is there any reason why I cannot compute a 97% confidence interval, or a 98% confidence interval, or a 99.9% confidence interval, which are wider? Can’t I increase the so-called plausible range to my liking simply by increasing the confidence level?”

“I see. And what about precision? May I say that the 95% confidence interval tells me how precise my estimate is?”
“Well, if precision means “distance from the unknown parameter”, it is a false idea. There is no way to tell the distance between an estimate and the unknown parameter.”

“What else could ‘precision’ mean?”
“We may speak about precision of an estimator, as estimated by the confidence interval. But the confidence interval is a function of the confidence level, too. (I can shrink or extend your claim of precision by changing the confidence level, as before.) The only way the confidence interval could give a notion of precision is by fixing the confidence level. Then, the width of the confidence interval is a function of the standard error, which itself is a measure of precision (a measure of spread). Using the confidence interval as an estimate of precision is nothing more than redundant mathematics on the standard error.”

“One last question, if I may.”
“Sure.”

“Where is this taught?”
“Right here (and in basic statistics courses, I hope.)”