Confidence Intervals of COVID-19 Vaccine Efficacy Rates

Frank Wang

LaGuardia Community College, CUNY, fwang@lagcc.cuny.edu

www.manaraa.com

Confidence Intervals of COVID-19 Vaccine Efficacy Rates

Abstract

This tutorial uses publicly available data from drug makers and the Food and Drug Administration to guide learners to estimate the confidence intervals of COVID-19 vaccine efficacy rates with a Bayesian framework. Under the classical approach, there is no probability associated with a parameter, and the meaning of confidence intervals can be misconstrued by inexperienced students. With Bayesian statistics, one can find the posterior probability distribution of an unknown parameter, and state the probability of vaccine efficacy rate, which makes the communication of uncertainty more flexible. We use a hypothetical example and a real baseball example to guide readers to learn the beta-binomial model before analyzing the clinical trial data. This note is designed to be accessible for lower-level college students with elementary statistics and elementary algebra skills.

Keywords

Creative Commons License

Cover Page Footnote

Frank Wang is Professor of Mathematics at LaGuardia Community College of the City University of New York. His research interests include general relativity and dynamics systems.

www.manaraa.com

Introduction

On November 18, 2020, the drug maker Pfizer issued a press release summarizing its Phase 3 study of COVID-19 vaccine (Pfizer 2020b). The announcement received a flurry of media coverage (LaFraniere et al. 2020; Zimmer 2020), focusing on the 95% vaccine efficacy rate. Pfizer's clinical trial involved 41,135 volunteers; there were 170 confirmed cases of COVID-19, with 162 observed in the placebo group and 8 in the vaccine group. Based only on these numbers reported in the New York Times (Zimmer 2020), I wrote a pedagogic note for *Numeracy* to verify Pfizer's claim of 95% vaccine efficacy rate and test its statistical significance (Wang 2021). I also used the method that students learn in a standard elementary statistics course to construct an approximate confidence interval of the efficacy rate, [92.1%, 98.5%], which was not given in the press release and news articles.

Since the completion of the manuscript of Wang (2021) in early December 2020, three leading COVID-19 vaccine developers (Pfizer-BioNTech, Moderna, and AstraZeneca-Oxford) published additional details about their clinical trials (Polack et al. 2020; Ramasamy et al. 2020; Baden et al. 2021). The US Food and Drug Administration (FDA) held meetings on December 10 and 17, 2020 to discuss Emergency Use Authorization of COVID-19 vaccines by Pfizer-BioNTech and Moderna. The FDA meeting videos and documents are open to the public, and a wealth of information is now available (FDA 2020b; FDA 2020c). Using a beta-binomial model, Pfizer provided a confidence interval of [90.3%, 97.6%] for the vaccine efficacy rate, and using the Clopper and Pearson method the confidence interval is [90.0%, 97.9%]. The two intervals are similar, and in reasonable agreement with my simplified treatment based on a one-proportion *z*-test.

To emphasize the need for confidence intervals, I often use the following example from Selvin (2004): in a particular state it was noted that more than half the women in prison for murder had killed their husbands, and less than a fifth of the men in prison for murder had killed their wives. Can we draw conclusions about male and female spousal relations? No. If only four women were convicted for murder and 660 men, the confidence interval associated with the women would be extremely wide and convey little information about the population parameter.

The COVID-19 pandemic offers a realistic example to illustrate the importance of confidence intervals. The AstraZeneca-Oxford vaccine was described to be puzzling for scientists (Callaway 2020). Among people who received a lower dose followed by a standard dose, the efficacy rate is 90%. However, for participants who received two full doses, the efficacy rate is 62%. The results may appear to be confusing, but once we know that the first trial involved only 33 COVID-19 cases, 3 in the vaccine group and 30 in the placebo group, we expect the confidence interval to be wide. Table 1 shows the raw counts of the numbers of COVID-19 cases and participants in three drug makers' clinical trials. Based on the company's

publication (Ramasamy et al. 2020), confidence interval for the initial trial is [67.4%, 97.0%], which overlaps with that of the two-dose regimen, [41.0%, 75.7%]. There is no statistically significant difference between these two reported AstraZeneca-Oxford vaccine efficacy rates, and speculation about what works better seems premature.

Table 1
Raw Data from Three Drug Makers' Publications

	Vaccine group **No. cases**/volunteers	Placebo group **No. cases**/volunteers
Pfizer all	**8**/17,411	**162**/17,511
Pfizer age ≥ 65 years	**1**/3,848	**19**/3,880
Moderna all	**11**/13,934	**185**/13,883
AstraZeneca I	**3**/1,367	**30**/1,374
AstraZeneca II	**27**/4,440	**71**/4,455
AstraZeneca combined	**30**/5,807	**101**/5,829

There are numerous ways to calculate confidence intervals (Newcombe 1998). In addition to Pfizer's methods mentioned above, Moderna employed a "stratified Cox proportional hazard model" and AstraZeneca utilized the "Poisson regression model with robust variance." The technical details about these methods can dazzle even statisticians. This note focuses on the Bayesian approach, which is not typically covered in elementary statistics courses. Some statisticians and scientists argue that the standard classical technique has undesirable features, and a Bayesian method can be more attractive (Efron 1986; Cousins 1995). Kranz (2020) analyzed the FDA meeting video and claimed that the Pfizer scientist used the Bayesian credible interval to make the presentation more understandable for advisory committee members who are not statisticians.[1] In some situations the frequentist and Bayesian confidence intervals are similar or even identical, and many textbooks have actually unknowingly used the Bayesian interpretation for the frequentist confidence intervals. For instance, the widely recommended *Chicago Guide to Writing about Numbers* (Miller 2004) contains the following example:

> Suppose our sample yields a mean math test score of 73.1 points with a standard error of 2.1 points. The 95% confidence interval is $73.1 \pm (2 \times 2.1)$, so we can be 95% sure that the average test score for the population falls between 68.9 and 77.3 points.

While this kind of statement can be found in many standard elementary textbooks, strictly speaking such an interpretation is not what Jerzy Neyman envisioned when he first introduced the concept of confidence intervals (Neyman 1937). In the Appendix we will find the Clopper-Pearson confidence intervals and interpret them the classical way.

We will use a simplified method based on Bayesian statistics to calculate confidence intervals for data shown in Table 1 to allow students to better understand

---

[1] The phrase *credible interval* was coined to distinguish it from the classical *confidence interval*. In this note I use "confidence intervals" regardless of whether the construction is classical or Bayesian.

and communicate COVID-19 vaccine clinical trials. We refrain from cluttered notation commonly appearing in Bayesian literature, but keep the mathematics at the level of elementary statistics and elementary algebra to make this note accessible for most college students. We start with a discussion of the beta probability distribution, related to the polynomial and power functions. We then introduce Bayes's rule and the beta-binomial model, and guide the readers to understand the theory through examples. After the preparatory work, we apply the method to analyze drug makers' data to find the Bayesian confidence intervals.

The Beta Distribution

We first use the familiar normal distribution to recall probabilistic properties. A normal distribution is characterized by two parameters, the mean and the standard deviation. For the standard normal distribution, or the z-distribution, the mean is 0 and standard deviation is 1. The probability density function (pdf) is

$$f(z) = \frac{1}{\sqrt{2\,\pi}} e^{-z^2/2} \,.$$

The graph of this function is the famous bell curve, and the area under the curve between two z scores represents the probability. The middle 95% of the area under the z-distribution pdf is bounded by $z = -1.96$ and $z = 1.96$. In principle, the areas need to be found by integration, but students and practitioners use a table (in the old days) or a computer program to retrieve pre-calculated integral values for the standard normal distribution.

The normal distribution is ubiquitous in an elementary statistics course and in real applications, but it is not the only type of probability distribution. An important class of distribution is the beta distribution, characterized by two parameters α and β. We use the notation $x \sim \text{Beta}(\alpha, \beta)$ to denote that the random variable X follows a beta distribution. We use upper-case letters to refer to random variables, and lower-case letters to refer to their actual observed values. The probability density function is proportional to

$$p(x) \propto x^{\alpha-1}(1-x)^{\beta-1}, \quad 0 \le x \le 1 \,.$$

The pdf of $\text{Beta}(1, 1)$ is simply 1, a constant function. The pdf for $\text{Beta}(2, 2)$ is

$$p(x) = 6x(1-x) \,,$$

whose graph is a familiar opening downward parabola that students should know how to hand-sketch when learning elementary algebra. For other values of α and β, students can use a graphing device to plot the curves. Consider Beta$(3, 19)$, the pdf is

$$p(x) = kx^2(1-x)^{18},$$

where k is a proportionality constant to ensure that the area under the curve between 0 and 1 is unity, so that the area under the curve between two x values represents the probability.[2] The graph is shown in Figure 1. We will use this distribution to model the hospitalization rate later.
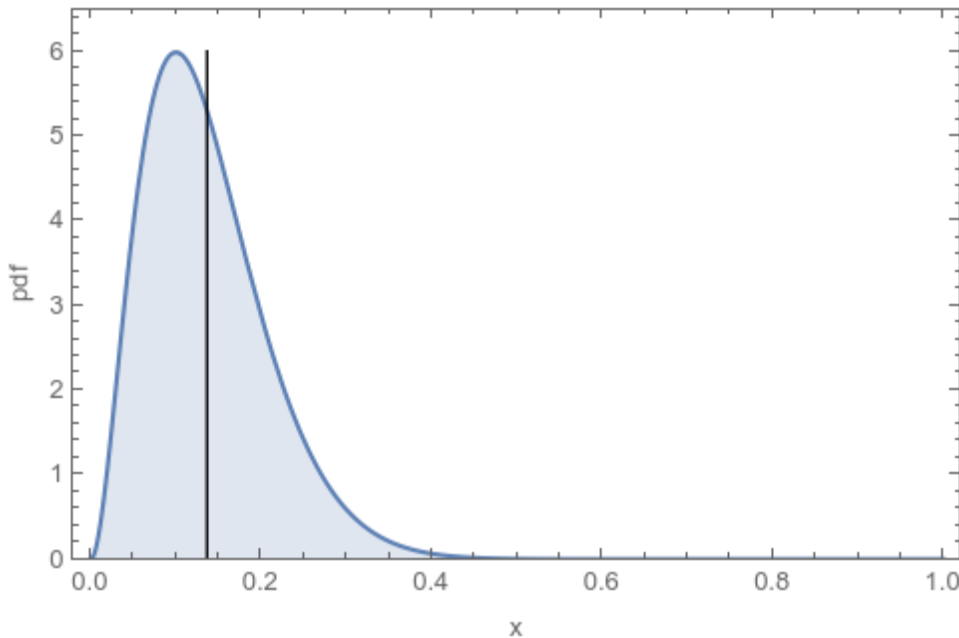

Figure 1. Probability density function of Beta$(3, 19)$, which is proportional to $x^2(1-x)^{18}$. The vertical line indicates the mean $3/22 = 0.136$, which is the horizontal coordinate of the centroid of the shape.

The mean of a continuous random variable following a certain probability distribution is the horizontal coordinate of the centroid of the area under the pdf. It is an elementary calculus problem to find the area and centroid of a shape, but we will not get into the details. For a Beta distribution, we cite the well-established results for the mean and variance (Gelman et al. 2013):

---

[2] The proportionality constant k can be expressed as the beta function which is coded in standard mathematical and statistical software.

$$E(X) = \frac{\alpha}{\alpha + \beta} \ , \quad \text{var}(X) = \frac{\alpha\,\beta}{(\alpha + \beta)^2\,(\alpha + \beta + 1)} \,.$$

For example, for Beta(3, 19), the mean is $3/(3 + 19) = 0.136$, which is indicated by the vertical line in Figure 1.

Bayes's Rule and the Beta-Binomial Model

Bayes's rule can be simply stated as "the posterior is proportional to the likelihood times the prior." Formally,

$$p(\theta|y) \propto p(\theta)p(y|\theta) \,,$$

where $p(\theta|y)$ and $p(\theta)$ are the posterior and prior distributions of the parameter θ, respectively, and $p(y|\theta)$ is the likelihood. The beta-binomial model uses the binomial sampling model as the likelihood, which is

$$p(y|\theta) = C_y^n \, \theta^y \, (1 - \theta)^{n-y} \,,$$

where C_y^n is the binomial coefficient; see Wang (2021) for a discussion or any statistics textbooks for background information. For this likelihood, it is natural to use a beta distribution as the prior. As a result, the binomial model with beta prior distribution introduced in the preceding section has a posterior like this:

$$p(\theta|y) \propto \theta^{y+\alpha-1} (1 - \theta)^{n-y+\beta-1} \,.$$

The posterior is also a beta distribution, $\theta|y \sim \text{Beta}(\alpha + y, \beta + n - y)$, and the mean is

$$E(\theta|y) = \frac{\alpha + y}{\alpha + \beta + n} \,.$$

The posterior mean invariably lies between the sample proportion y/n and the prior mean $\alpha/(\alpha + \beta)$ (Gelman et al. 2013). Before getting too abstract, we use some examples from Connor (2021) published in *Numeracy* to illustrate the idea.

Hospitalization Rate

Suppose 2 out of 20 young students in a classroom need hospitalization after being infected by COVID-19. Naively, we estimate the hospitalization rate for young students to be $2/20 = 0.1$, but if 2 out 20 senior citizens who live in an assisted

living facility need hospitalization after being infected by COVID-19, it seems unreasonable to conclude that the hospitalization rate for senior citizens is also $2/20 = 0.1$, the same as students' rate. The essence of Bayesian statistics is that one needs to incorporate previously established research into the current one, by including the previous knowledge as the prior distribution. If we use $\text{Beta}(200, 100)$ as the prior for people who live in an assisted living facility (Connor 2021, Figure 3), then the posterior distribution is $\text{Beta}(200 + 2, 100 + 18)$ and the Bayes-estimated hospitalization rate is

$$\frac{200 + 2}{200 + 100 + 20} = 0.631.$$

On the other hand, if we use $\text{Beta}(40, 4000)$ for students (Connor 2021, Figure 4), then the posterior distribution is $\text{Beta}(40 + 2, 4000 + 18)$ and the Bayes-estimate rate is

$$\frac{40 + 2}{40 + 4000 + 20} = 0.010.$$

The hospitalization rate for people who live in an assisted living facility is about 63 times higher than the rate for students using the Bayesian model, which makes more sense. In further details, the likelihood for both groups is based on the observation of 2 cases out of a sample of 20, and is modeled by this binomial distribution,

$$P(y|\theta) = C_2^{20}\, \theta^2 (1 - \theta)^{18},$$

where θ is the unknown true rate. The prior $\text{Beta}(200, 100)$ for people in an assisted living facility has the following pdf:

$$P(\theta) = k\theta^{199}(1 - \theta)^{99},$$

where k is a proportionality constant. This prior distribution can be viewed as the following: previously 199 people with similar background required hospitalization, and 99 did not. The posterior is proportional to the likelihood multiplied by the prior,

$$P(\theta|y) \propto \theta^{199+2}(1 - \theta)^{99+18},$$

which is the pdf of the distribution $\text{Beta}(200 + 2, 100 + 18)$. From this analysis, the posterior mean 0.631 is just the weighted mean of the prior rate $200/(200 +$

$100) = 0.667$ and the sample proportion $2/20 = 0.1$. Similarly, the prior pdf for students is

$$P(\theta) = k\theta^{39}(1-\theta)^{3999},$$

and the posterior pdf is:

$$P(\theta|y) \propto \theta^{39+2}(1-\theta)^{3999+18}.$$

A nice feature of the Bayesian method is that we obtain a posterior probability density function for the true parameter θ as seen above, so that we can discuss the probability of the parameter. In the next section we will see how a Pfizer scientist used the distribution when communicating about the vaccine efficacy.

In Figure 2, we show the prior and posterior probability density functions for two groups. We include the likelihood function in the same plot, but the likelihood function is not a probability density function. Note that $P(y|\theta) = C_2^{20}\,\theta^2\,(1-\theta)^{18}$ is proportional to the pdf of Beta(3, 19). See also Figure 1. For people in an assisted living facility, we can see that the posterior mean 0.631 is between the prior mean 0.667 and the raw rate 0.1. This can be considered as the phenomenon "regression to the mean." For students in a classroom, we make a similar observation that the posterior mean 0.010 is between the prior mean $40/(40 + 4000) = 0.0099$ and the raw rate 0.1, although the data 2 out of 20 are dominated by the prior distribution and we can hardly distinguish the prior and posterior probability density functions from Figure 2.
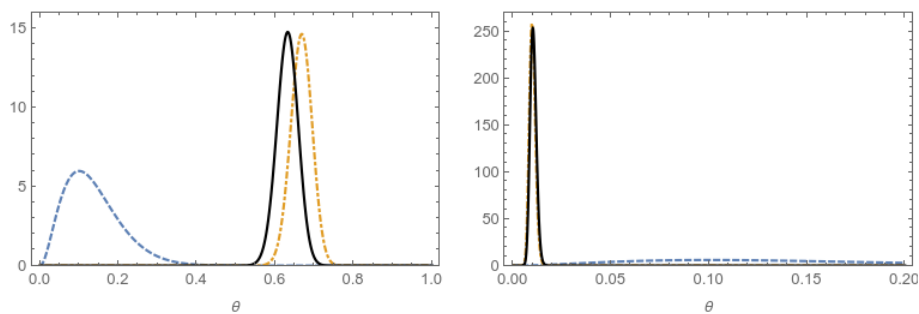

Figure 2. The dot-dashed curve is the prior, dashed curve is proportional to the likelihood, and the solid curve is the posterior, for people in an assisted living facility (left) and students in a classroom (right). The dashed curve for both plots is the same function as the curve in Figure 1 but shown in different scale.

Baseball Batting Average

This real-life example is from Irizarry and Love (2015). José Iglesias is a professional baseball player. In April 2013, he made 9 hits out of 20 times at bat. The raw rate, called batting average, is $9/20 = 0.450$. It is strikingly high, as no one has finished a season with an average of 0.400 since Ted Williams did it in 1941. Irizarry and Love calculated the batting averages for all players with more than 500 at bats during the previous three seasons, and found the mean to be 0.275 and standard deviation 0.027. They used a normal distribution to compute the posterior distribution, but here we will use the beta-binomial model. From the records of the previous three seasons, we set the following equations to solve for α and β,

$$\frac{\alpha}{\alpha + \beta} = 0.275 \, , \quad \frac{\alpha\,\beta}{(\alpha + \beta)^2\,(\alpha + \beta + 1)} = 0.027^2.$$

One can employ a computer algebra system such as Maple or Mathematica to solve this system of two equations with two unknowns to obtain $\alpha = 74.935$ and $\beta = 197.556$. With the prior distribution Beta$(74.935, 197.556)$, the posterior distribution is Beta$(74.935 + 9, 197.556 + 11)$, and the Bayes-estimated rate for Iglesias is

$$\frac{74.935 + 9}{74.935 + 197.556 + 20} = 0.287.$$

The posterior distribution

$$P(\theta|y) \propto \theta^{74.935-1+9}(1 - \theta)^{197.556-1+11}$$

is shown in Figure 3, along with the prior and likelihood. The posterior pdf allows us to construct a Bayesian 95% confidence interval, denoted by $[L, U]$. The lower bound is the value such that the area under the pdf between 0 and L is 0.025, and the upper bound is the value such that the area under the pdf between U and 1 is 0.025. This is an elementary area problem that students learn in calculus, but many computer programs exist to perform numerical integration and locate the end points. Using the R command qbeta(c(0.025, 0.975), 74.935+9, 197.556+11), we obtain the 95% confidence interval [0.237, 0.340]. We conclude that based on José Iglesias's performance in April 2013, there is a 95% probability that his true batting average is between 0.237 and 0.340. We can talk about the probability of the parameter (batting average), which we cannot do under the classical framework.
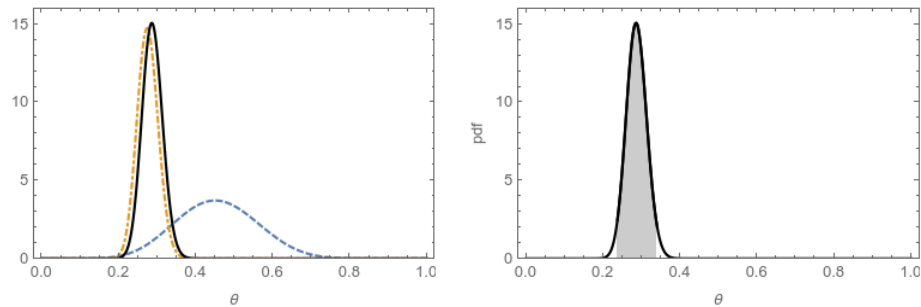
Figure 3. Left: The dot-dashed curve is the prior, dashed curve is proportional to the likelihood, and the solid curve is the posterior. Right: The shaded area under the posterior pdf between the end points of the 95% confidence interval, [0.237, 0.340], is 0.95.

The Bayes-estimated rate 0.287 turned out to be a more accurate prediction than the raw rate 0.45. From May to September, Iglesias had 97 hits out of 330, or a batting average of 0.293.[3] The Bayesian prediction that he would not be as good the remainder of the season is another example of regression to the mean.

Vaccine Efficacy Rates

According to the FDA issued guidance, the vaccine efficacy rate should be at least 50% to be considered a success, although the lower bound of the confidence interval can be as low as 30% (FDA 2020a), another reminder that confidence intervals are crucial. To ensure public trust, Pfizer, Moderna, and AstraZeneca agreed to make their full study protocols publicly available. Pfizer used the FDA recommendations to construct a minimally informative prior beta distribution for the Bayesian confidence interval (Pfizer 2020a), as we will describe below.

From Pfizer's protocol, we learn that the vaccine efficacy is defined as

$$\text{VE} = 1 - \text{IRR},$$

where IRR is the ratio of COVID-19 illness rate in the vaccine group to the illness rate in the placebo group. The parameter θ is defined as

$$\theta = \frac{1 - \text{VE}}{2 - \text{VE}}.$$

From the above definition, we can find a formula for VE in terms of θ:

---

[3] Reviewer 1 used a uniformly distributed prior ($\alpha = 1$, $\beta = 1$) to find the 95% confidence interval [0.257, 0.660], which also encompasses the true batting average. The empirical approach using other players' data, however, provides a much more precise prediction.

$$\text{VE} = \frac{1 - 2\theta}{1 - \theta} \, .$$

In Wang (2021), I used the *New York Times* description to express the efficacy rate as follows:

$$\text{VE} = 1 - \frac{P(\text{COVID+}|\text{vaccine})}{P(\text{COVID+}|\text{placebo})} \, ,$$

where $P(\text{COVID+}|\text{vaccine})$ and $P(\text{COVID+}|\text{placebo})$ are the number of cases in vaccine and placebo groups, respectively. Pfizer's illness rate is measured in units of per 1000 person-years. Because we lack detailed temporal information of the clinical trial results, we make an approximation of equal surveillance time for the vaccine and control groups to simplify the calculation. With such an approximation and some algebraic operations, we obtain

$$\theta = \frac{P(\text{COVID+}|\text{vaccine})}{P(\text{COVID+}|\text{vaccine}) + P(\text{COVID+}|\text{placebo})} \, .$$

The November 18, 2020 press release from Pfizer (Pfizer 2020b) and the *New York Times* article (Zimmer 2020) did not provide the number of volunteers in each group, so I resorted to Bayes's rule to relate the inverse probabilities:

$$\frac{P(\text{vaccine}|\text{COVID+})}{P(\text{placebo}|\text{COVID+})} = \frac{P(\text{vaccine})}{P(\text{placebo})} \times \frac{P(\text{COVID+}|\text{vaccine})}{P(\text{COVID+}|\text{placebo})} \, .$$

See Wang (2021) for further discussion. With the approximation $P(\text{vaccine}) = P(\text{placebo})$,[4] we can express θ as

$$\theta = \frac{P(\text{vaccine}|\text{COVID+})}{P(\text{vaccine}|\text{COVID+}) + P(\text{placebo}|\text{COVID+})} \, .$$

This definition allows us to estimate the raw rate for Pfizer to be $\theta = 8/170$, based on the observed 8 people in the vaccine group among 170 COVID-positive volunteers. We also have the vaccine efficacy rate as follows:

---

[4] To analyze the clinical trial of the Sputnik V vaccine in Russia, which involved 14,964 volunteers in the vaccine group and 4,902 people in the placebo group (Logunov et al. 2021), we have developed a more general formulism and will present it in a future publication.

$$\text{VE} = \frac{1 - 2 \times \dfrac{8}{170}}{1 - \dfrac{8}{170}} = \frac{154}{162} = 1 - \frac{8}{162} = 0.951.$$

I show extra steps in arithmetic so that the reader can relate to my earlier treatment (Wang 2021).

To find the Bayesian confidence interval, Pfizer constructed a minimally informative prior. An uninformative prior is a constant function, or $\text{Beta}(1, 1)$. According to Pfizer's protocol (Pfizer 2020a), they set the prior VE to have a mean of 30%, the minimal requirement from the FDA guidance, and at this rate $\theta = (1 - 3/10)/(2 - 3/10) = 7/17 = 0.4118$. Recall that for a beta distribution, the mean is $\alpha/(\alpha + \beta)$. Let us keep $\beta = 1$, and we can solve for α:

$$\frac{\alpha}{\alpha + 1} = \frac{7}{17}, \quad \alpha = \frac{7}{10}.$$

This formula is the basis for Pfizer's prior, $\text{Beta}(0.700102, 1)$. The extra digits in Pfizer α is due to rounding error and seems superfluous.

We use the binomial sampling model for the likelihood. From the observed 8 cases in the vaccine group and 162 cases in the placebo group,

$$P(y|\theta) = C_8^{170} \, \theta^8 (1 - \theta)^{162} \,.$$

The pdf for $\text{Beta}(7/10, 1)$ is

$$P(\theta) = k\theta^{7/10-1}(1 - \theta)^{1-1} = k\theta^{-3/10} \,,$$

where k is a proportional constant. The posterior pdf is

$$P(\theta|y) \propto \theta^{7/10-1+8}(1 - \theta)^{162} \,,$$

which is the pdf of $\text{Beta}(7/10 + 8, 1 + 162)$. See Figure 4 for the prior and posterior probability density functions, together with the likelihood. The Bayes-estimated rate is

$$\theta = \frac{\dfrac{7}{10} + 8}{\dfrac{7}{10} + 1 + 170} = 0.0501, \quad \text{VE} = \frac{1 - 2 \times 0.0501}{1 - 0.0501} = 0.947.$$
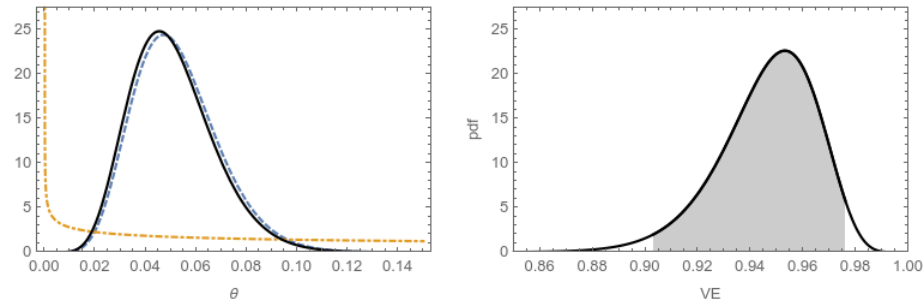
Figure 4. Left: The dot-dashed curve is the prior, dashed curve is proportional to the likelihood, and the solid curve is the posterior. Right: The shaded area under the posterior pdf of VE between the end points of the 95% confidence interval, [0.903, 0.976], is 0.95.

Because the prior is minimally informative, the posterior rate is not too different from the raw rate. Unlike examples in the preceding section, here the data dominate the prior distribution. From the posterior distribution, we can use the R command qbeta(c(0.025, 0.975), 7/10+8, 1+162) to solve the area problem to find the 95% confidence interval for θ to be [0.0232, 0.0880], and the corresponding VE confidence interval [90.3%, 97.6%]. This is identical to the confidence interval in the briefing document that Pfizer submitted to the FDA for the December 10, 2020 meeting. Pfizer scientist Dr. William Gruber said the following during the Advisory Committee meeting (FDA 2020b).

> There's 95 percent probability that efficacy falls in the intervals shown; meaning, that over 97.5 percent likelihood that the efficacy is greater than 90 percent. Likewise, the probability that vaccine efficacy is at least greater than 30 percent greatly exceeds FDA COVID-19 vaccine guidance.

Dr. Gruber's testimony is an example that one can talk about the probability of the true parameter, in this case vaccine efficacy rate, under the Bayesian framework.

From Pfizer's document, we find that the 65-and-older subgroup has one case in the vaccine group and 19 in the placebo group (see Table 1). We expect a wider confidence interval for this age group. Similar to the above procedure, the posterior distribution is Beta$(7/10 + 1, 1 + 19)$. Below, we show how to use R to find the 95% confidence interval for θ, [0.00765, 0.219], and the corresponding end points of the confidence interval for VE, [71.9%, 99.2%].[5] We can say that there is a 95% probability that the vaccine efficacy rate is between 71.9% and 99.2% for people of age 65 and over, based on the data.

---

[5] Reviewer 2 noticed that the Bayesian confidence interval [71.9%, 99.2%] is slightly different from the Clopper-Pearson confidence interval [66.7%, 99.9%] in Pfizer's document submitted to the FDA. In general, the Bayesian confidence interval is narrower than the classical one. Additionally, our simplified treatment did not take the minor difference in surveillance time into account.

```
> (theta <- qbeta(c(0.025, 0.975), 7/10+1, 1+19))
[1] 0.007646717 0.219527315
> (VE <- (1-2*theta)/(1-theta))
[1] 0.9922944 0.7187252
> curve(dbeta(x, 7/10+1, 1+19))                    #posterior
> curve(dbeta(x, 7/10, 1), lty = 4, add = TRUE)    #prior
> curve(dbeta(x, 2, 20), lty = 2,add = TRUE)       #likelihood
```

The R output of the graph is shown in Figure 5. We again point out that the likelihood, in this case $P(y|\theta) = C_1^{20}\, \theta^1\, (1-\theta)^{19}$, is not a probability density function. This function is proportional to the pdf of Beta$(2, 20)$, and we use it to show the relative magnitude of the likelihood.
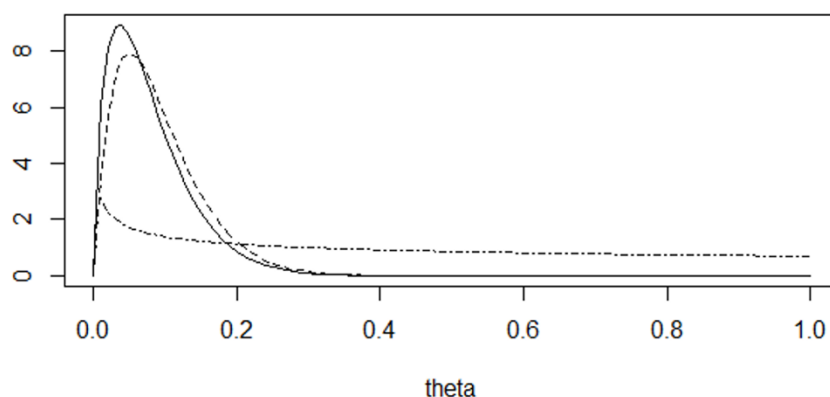

Figure 5. In the R graphic output, the dot-dashed curve is the prior, dashed curve is proportional to the likelihood, and the solid curve is the posterior.

Table 2 summarizes the confidence intervals of the data in Table 1 using the minimally informative prior Beta$(7/10, 1)$. They are in good agreement with the published ones, which were based on classical or nonparametric methods (except for the group Pfizer all). Although the values are similar, the Bayesian and classical interpretations of confidence intervals are very different. See above for a Bayesian interpretation, and the Appendix for the classical interpretation.

Table 2
Comparison of VE Confidence Intervals %

	Based on this work	Published values
Pfizer all	[90.3, 97.6]	[90.3, 97.6]
Pfizer age $\geq$ 65 years	[71.9, 99.2]	[66.7, 99.9]
Moderna all	[89.4, 96.8]	[89.3, 96.8]
AstraZeneca I	[70.7, 97.0]	[67.4, 97.0]
AstraZeneca II	[41.5, 75.8]	[40.0, 76.5]
AstraZeneca combined	[55.8, 80.4]	[54.8, 80.6]

As mentioned in the Introduction, AstraZeneca-Oxford reported a vaccine efficacy rate of 90.0% for volunteers who received a lower amount then the full

amount in the second dose, and 62.1% for volunteers who received two full doses. The company combined the data, and claimed an overall 70.4% efficacy rate. Our calculated confidence interval for the combined data is narrower than the confidence intervals of individual trials, as we would expect it.

Although the AstraZeneca-Oxford vaccine's efficacy rate is less impressive than that of Pfizer and Moderna, it can still make a significant impact on public health. Imagine two otherwise identical communities, but one is vaccinated and the other not. Even if the efficacy rate is only 70%, for every 100 people who become infected by COVID-19 in the unvaccinated community, there will be on average 30 sick people in the vaccinated community. Furthermore, the Bayesian model gives the probability of every possible vaccine rate, and public health professionals may use decision theory to allocate medical resources.

Concluding Remarks

Under the Bayesian framework, one estimates a parameter based on the observed data and obtains a posterior probability density function for the parameter. This distribution allows one to talk about the probability of the parameter, which is often more natural when communicating uncertainty. Specifically, we can state that there is a 95% probability that Pfizer's vaccine efficacy rate is between 90.3% and 97.6%, based on the clinical trial data. We have presented a simplified method to analyze vaccine data based on Bayesian statistics. The posterior probability distribution in a beta-binomial model is an elementary polynomial or power function that students with basic algebra skills are familiar with. The explicit expression for the probability of the unknown parameter can be graphed, and students can use it to communicate confidence intervals more flexibly. With the minimally informative prior, we found that the Bayesian confidence intervals are similar to classical ones reported in medical literature. One can recast a classical confidence interval into a Bayesian one using our method, and speak about the probability of vaccine efficacy rates like what the Pfizer scientist did during the FDA meeting.

Acknowledgements

References

Baden, L. R., H. M. El Sahly, B. Essink, K. Kotloff, S. Frey, R. Novak, D. Diemert, S. A. Spector, N. Rouphael, C. B. Creech, J. McGettigan, S. Khetan, N. Segall, J. Solis, A. Brosz, C. Fierro, H. Schwartz, K. Neuzil, L. Corey, P. Gilbert, H. Janes, D. Follmann, M. Marovich, J. Mascola, L. Polakowski, J. Ledgerwood, B. S. Graham, H. Bennett, R. Pajon, C. Knightly, B. Leav, W. Deng, H. Zhou, S. Han, M. Ivarsson, J. Miller, and T. Zaks. 2021. "Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine." *New England Journal of Medicine,* 384: 403–416. https://doi.org/10.1056/NEJMoa2035389

Callaway, E. 2020. "Oxford COVID Vaccine Results Puzzle Scientists." *Nature,* 588: 16–18. https://doi.org/10.1038/d41586-020-03326-w

Connor, C. 2021. "Computing for Numeracy: How Safe Is Your COVID-19 Social Bubble?" *Numeracy,* 14(1): Article 12. https://doi.org/10.5038/1936-4660.14.1.1382

Cousins, R. D. 1995. "Why Isn't Every Physicist a Bayesian?" *American Journal of Physics,* 63(5): 398–410. https://doi.org/10.1119/1.17901

Efron, B. 1986. "Why Isn't Everyone a Bayesian?" *American Statistician,* 40: 1–11. https://doi.org/10.2307/2683111

Food and Drug Administration. 2020a. "Development and Licensure of Vaccines to Prevent COVID-19: Guidance for Industry." https://www.fda.gov/media/139638/download

Food and Drug Administration. 2020b. "Vaccines and Related Biological Products Advisory Committee December 10, 2020 Meeting Announcement." https://www.fda.gov/advisory-committees/advisory-committee-calendar/vaccines-and-related-biological-products-advisory-committee-december-10-2020-meeting-announcement

Food and Drug Administration. 2020c. "Vaccines and Related Biological Products Advisory Committee December 17, 2020 Meeting Announcement." https://www.fda.gov/advisory-committees/advisory-committee-calendar/vaccines-and-related-biological-products-advisory-committee-december-17-2020-meeting-announcement

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. 2013. *Bayesian Data Analysis*, 3rd edition. Boca Raton, FL: CRC Press. https://doi.org/10.1201/b16018

Irizarry, R. A., and M. I. Love. 2015. *Data Analysis for the Life Sciences*. https://doi.org/10.1201/9781315367002

Kranz, S. 2020. "A Quiz about a 95% CI Interpretation in the FDA COVID Vaccine Meeting." https://skranz.github.io/r/2020/12/11/CovidVaccineFDAHearingCI.html

LaFraniere, S., K. Thomas, N. Weiland, D. Gelles, S. G. Stolberg, and D. Grady. 2020. "Politics, Science and the Remarkable Race for a Viable Vaccine." *New York Times*, November 22.

Logunov, Denis Y., Inna V. Dolzhikova, Dmitry V. Shcheblyakov, Amir I. Tukhvatulin, Olga V. Zubkova, Alina S. Dzharullaeva, Anna V. Kovyrshina, Nadezhda L. Lubenets, Daria M. Grousova, Alina S. Erokhova, Andrei G. Botikov, Fatima M. Izhaeva, Olga Popova, Tatiana A. Ozharovskaya, Ilias B. Esmagambetov, Irina A. Favorskaya, Denis I. Zrelkin, Daria V. Voronina, Dmitry N. Shcherbinin, Alexander S. Semikhin, Yana V. Simakova, Elizaveta A. Tokarskaya, Daria A. Egorova, Maksim M. Shmarov, Natalia A. Nikitenko, Vladimir A. Gushchin, Elena A. Smolyarchuk, Sergey K. Zyryanov, Sergei V. Borisevich, Boris S. Naroditsky, Alexander L. Gintsburg. 2021. "Safety and Efficacy of an rAd26 and rAd5 Vector-based Heterologous Prime-boost COVID-19 Vaccine: An Interim Analysis of a Randomised Controlled Phase 3 Trial in Russia." *Lancet,* 397: 671–681. https://doi.org/10.1016/S0140-6736(21)00234-8

Miller, J. E. 2004. *The Chicago Guide to Writing about Numbers*. Chicago: University of Chicago Press. https://doi.org/10.7208/chicago/9780226526324.001.0001

Newcombe, R. G. 1998. "Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods." *Statistics in Medicine,* 17: 857–872. https://doi.org/10.1002/(SICI)1097-0258(19980430)17:8<857::AID-SIM777>3.0.CO;2-E

Neyman, J. 1937. "Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability." *Philosophical Transaction of the Royal Society of London,* 236(767): 333–380. https://doi.org/10.1098/rsta.1937.0005

Pfizer. 2020a. "PF-07302048 (BNT162 RNA-based COVID-19 Vaccines) Protocol C4591001." https://pfe-pfizercom-d8-prod.s3.amazonaws.com/2020-09/C4591001_Clinical_Protocol.pdf

Pfizer. 2020b. "Pfizer and BioNTech Conclude Phase 3 Study of COVID-19 Vaccine Candidate, Meeting All Primary Efficacy Endpoints." https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-conclude-phase-3-study-covid-19-vaccine

Polack, Fernando P., Stephen J. Thomas, Nicholas Kitchin, Judith Absalon, Alejandra Gurtman, Stephen Lockhart, John L. Perez, Gonzalo Pérez Marc, Edson D. Moreira, Cristiano Zerbini, Ruth Bailey, Kena A. Swanson, Satrajit Roychoudhury, Kenneth Koury, Ping Li, Warren V. Kalina, David Cooper,

Robert W. Frenck, Jr., Laura L. Hammitt, Özlem Türeci, Haylene Nell, Axel Schaefer, Serhat Ünal, Dina B. Tresnan, Susan Mather, Philip R. Dormitzer, Uğur Şahin, Kathrin U. Jansen, and William C. Gruber. 2020. "Safety and Efficacy of the BNT162b2 mRNA COVID-19 Vaccine." *New England Journal of Medicine,* 383: 2603–2615. https://doi.org/10.1056/NEJMoa2034577

Ramasamy, Maheshi N. Angela M. Minassian, Katie J. Ewer, Amy L. Flaxman, Pedro M. Folegatti, Daniel R. Owens, Merryn Voysey, Parvinder K. Aley, Brian Angus, Gavin Babbage, Sandra Belij-Rammerstorfer, Lisa Berry, Sagida Bibi, Mustapha Bittaye, Katrina Cathie, Harry Chappell, Sue Charlton, Paola Cicconi, Elizabeth A. Clutterbuck, Rachel Colin-Jones, Christina Dold, Katherine R. W. Emary, Sofiya Fedosyuk, Michelle Fuskova, Diane Gbesemete, Catherine Green, Bassam Hallis, Mimi M. Hou, Daniel Jenkin, Carina C. D. Joe, Elizabeth J. Kelly, Simon Kerridge, Alison M. Lawrie, Alice Lelliott, May N. Lwin, Rebecca Makinson, Natalie G. Marchevsky, Yama Mujadidi, Alasdair P. S. Munro, Mihaela Pacurar, Emma Plested, Jade Rand, Thomas Rawlinson, Sarah Rhead, Hannah Robinson, Adam J. Ritchie, Amy L. Ross-Russell, Stephen Saich, Nisha Singh, Catherine C. Smith, Matthew D. Snape, Rinn Song, Richard Tarrant, Yrene Themistocleous, Kelly M. Thomas, Tonya L. Villafana, Sarah C. Warren, Marion E. E. Watson, Alexander D. Douglas, Adrian V. S. Hill, Teresa Lambe, Sarah C. Gilbert, Saul N. Faust, Andrew J. Pollard. 2020. "Safety and Immunogenicity of ChAdOx1 nCoV-19 Vaccine Administered in a Prime-boost Regimen in Young and Old Adults (VOV002): A Single-blind, Randomized, Controlled, Phase 2/3 Trial." *Lancet,* 396: 1979–1993. https://doi.org/10.1016/S0140-6736(20)32466-1

Selvin, S. 2004. *Biostatistics: How It Works*. Upper Saddle River, NJ: Pearson.

Wang, F. 2021. "Using COVID-19 Vaccine Efficacy Data to Teach One-Sample Hypothesis Testing." *Numeracy,* 14(1): Article 7. https://doi.org/10.5038/1936-4660.14.1.1383

Zimmer, C. 2020. "What Does It Mean If 2 Companies Report 95% Efficacy Rates?" *New York Times*, November 21.

Appendix

In Pfizer's briefing document submitted to the FDA for the December 10, 2020 meeting, two methods are mentioned in the footnotes, the beta-binomial model and the Clopper and Pearson method. Among 170 volunteers who became infected by COVID-19, 8 were from the vaccine group, and 162 from the placebo group. Below is the R code that reproduces Pfizer's reported Clopper-Pearson confidence interval for θ, [0.0205, 0.0906]. Then we use θ to compute the confidence interval for VE, [90.0%, 97.9%].

```
> testall <- binom.test(8, 8+162)
> (theta <- testall$conf.int)
[1] 0.02053273 0.09061668
attr(,"conf.level")
[1] 0.95
> (VE <- (1-2*theta)/(1-theta))
[1] 0.9790368 0.9003537
```

The Clopper-Pearson interval is commonly referred to as the "exact confidence interval." Let us decipher the meaning of R's output. For the binomial distribution, the probability mass function is

$$p(y|\theta) = C_y^n \, \theta^y \, (1-\theta)^{n-y}$$

(see the main text). Let the 95% confidence interval of θ be $[L, U]$. To find U for the observation $y = 8$ out of $n = 170$, we need to solve the following equation for θ. See Figure A.1.

$$\sum_{k=0}^{8} C_k^{170} \theta^k (1-\theta)^{170-k} = 0.025.$$

There are algorithms to solve this equation, and below we use R to verify that when $\theta = 0.0906$ the above equation is satisfied.

```
> plot(0:30, dbinom(0:30, 170, 0.09061668), type = "h")
> sum(dbinom(0:8, 170, 0.09061668))
[1] 0.025
```
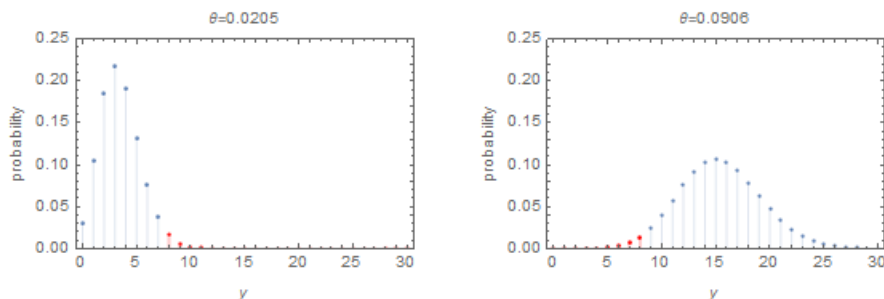
Figure A.1. Classical solutions for the construction of the 95% confidence interval.

Similarly, to find L we solve the following equation:

$$\sum_{k=8}^{170} C_k^{170} \theta^k (1-\theta)^{170-k} = 0.025.$$

Below we verify that when $\theta = 0.0205$, the equation is satisfied.

```
plot(0:30, dbinom(0:30, 170, 0.02053273), type = "h")
> sum(dbinom(8:170, 170, 0.02053273))
[1] 0.02500001
```

If we repeat the clinical trial, we may get a different y, which corresponds to a different confidence interval. The classical construction guarantees that in the limit of many repeated trials, 95% of the confidence intervals contain the unknown true value θ. The classical confidence interval reflects the relative frequency with which the statement "θ is in the interval $[L, U]$" is a true statement (Cousins 1995). Robert D. Cousins claimed that many people do not think about the classical confidence intervals this way, and wrote a paper titled "Why Isn't Every Physicist a Bayesian?" (Cousin 1995), in the spirit of Brad Efron's article "Why Isn't Everyone a Bayesian?" (Efron 1986), to demonstrate the flexibility of the Bayesian method.