

Inferring vaccine efficacies and their uncertainties.

A simple model implemented in JAGS/rjags

Giulio D'Agostini¹ and Alfredo Esposito²

[19 November 2020, revised 25 November 2020]

Abstract

Taking the cue from the incredibly precise value of the efficacy of Moderna's COVID-19 vaccine candidate broadcasted by the media these days (94.5%, without any uncertainty attached to it, as instead it should always be the case for a scientific result) we try to get the probability distribution of such efficacy with the limited information available. The work has been done with the help of a simple Bayesian Network, processed by a Markov Chain Monte Carlo. The inferred efficacy results $(93.3 \pm 2.9)\%$ (mean \pm standard uncertainty) and a 95% credible interval of [86.6%, 97.6%]. We have also processed through the same model the new Pfizer results, claiming a 95% efficacy, getting $(94.4 \pm 1.9)\%$ with a 95% credible interval of [90.0%, 97.5%]. The efficacies reported by the two companies correspond indeed to the modal values of the distributions.

*“It is scientific only to say what is more likely
and what is less likely”
(R. Feynman)*

¹Università “La Sapienza” and INFN, Roma, Italia, giulio.dagostini@roma1.infn.it

²Retired, alfespo@yahoo.it

1 Introduction

Our perspectives about living with Covid-19 are changed dramatically in just a few days with the results from vaccine trials of the past days by Pfizer and Moderna. The former claimed a 90% efficacy (then updated to 95%); the latter 94.5%. Obviously, the media did not mention any uncertainty, so we understood that the initial Pfizer's number was the result of a rounding, with uncertainty of the order of the percent. We were more surprised by the Moderna's one, providing the tenths of the percent, as if it were much more precise. Looking around, we had the impression that the "point five" was taken very serious, not only by media speakers, who put the emphasis on the third digit, but also by experts from which we would have expected a phrasing implying some uncertainty in the result (see e.g. Ref.[1]).

A fast exercise showed that, in order to have an uncertainty of the order of a few tenths of percent, the number of vaccine-treated individuals that got the Covid-19 had to be at least of the order of *several hundreds*. But this was not the case. In fact, the actual numbers were indeed much smaller: "*This first interim analysis was based on 95 cases, of which 90 cases of COVID-19 were observed in the placebo group versus 5 cases observed in the mRNA-1273 group, resulting in a point estimate of vaccine efficacy of 94.5% ($p < 0.0001$)*" [2]. Now, it is a matter of fact that if a physicist reads a number like '5', she tends to associate to it, as a rule of thumb, an uncertainty of the order of its square root, that is ≈ 2.2 . Applied to the Moderna claims, this implies an *inefficacy* of $\approx (5.5 \pm 2.3)\%$, or an efficacy of $\approx (94.5 \pm 2.3)$. Another reason to worry about the scientific validity of the result [3] was not only the absence of an uncertainty associated to the result, but also the " $p < 0.0001$ " accompanying it, especially for those who are extremely critical against *p-values* and other *frequentist prescriptions* [4].

We then tried to see whether it was possible to get an idea of the possible values of efficacy consistent with the data, each associated with a degree of belief. In other words, we have tried to critically review the claims, on the basis of the scarce data made available, following a sound probabilistic approach, in order to arrive to a probability density function (pdf), although not obtained in closed form, of the quantity of interest. In these kinds of situations, we have learned (see e.g. [5]) that the most important starting point is to build up a graphical representation of the *causal model* relating the quantities of interest, some of them 'observed' and others 'unobserved', among the latter the quantities we want to infer. Also in this case, despite some initial skepticism about the possibility to get some reasonable results, because of the scarce information in our hands, once we have built up the model, very basic indeed, it was clear that the main result about the efficacy was not depending on the many details of the trials.




Figure 1: Simplified Bayesian network of the vaccine vs placebo experiment (see text).

2 Model and analysis method

The causal model used is represented in the *Bayesian network* of Fig. 1. The top *nodes* n_V and n_P stand, respectively, for the number of individual in the vaccine and placebo groups, as the subscripts indicate, while the bottom ones (n_{V_I} and n_{P_I}) are the number of individuals of the two groups resulting infected after the trial period.

The sure data are $n_{V_I} = 5$ and $n_{P_I} = 90$ for Moderna [2] and $n_{V_I} = 8$ and $n_{P_I} = 162$ for Pfizer [6]. As far as the number of individual subject to the trials there were certainly some information in the press releases, but, fortunately, as we shall see, the exact number is not critical at all in regard to the value of efficacy and we can even change it by orders of magnitudes without affecting the results of interest.

Then, there was the question of how to relate the numbers of infected to the numbers of the participants in the trial. This depends in fact from several variables, like the prevalence of the virus in the population(s) of the involved people, their life-style, behavior, and so on, and, hopefully, from the fact that a person has been vaccinated or not. We simplified the model defining an *assault probability*, p_A , a catch-all term embedding the many real life variables, apart being vaccinated or not. Nodes n_{V_A} and n_{P_A} represent them the number of ‘assaulted individuals’ in each group, and they are modeled according to a binomials distributions, that is

$$n_{V_A} \sim \text{Binom}(n_V, p_A) \quad (1)$$

$$n_{P_A} \sim \text{Binom}(n_P, p_A), \quad (2)$$

represented in the graphical model by solid arrows.

The ‘assaulted individuals’ of the placebo group are then assumed to be all infected, and hence the deterministic link with dashed arrow relating the node n_{P_A} to the node n_{P_I} (in fact the two numbers are the same, and we make this graphical distinction only for symmetry with respect to the vaccine group).

Instead, the ‘assaulted individuals’ of the other group are ‘shielded’ by the vaccine, with probability of being infected equal to $1 - \epsilon$, where ϵ is the efficacy:

$$n_{V_I} \sim \text{Binom}(n_{V_A}, 1 - \epsilon). \quad (3)$$

At this point all the rest is a matter of calculations, that we do by Markov Chain Monte Carlo (MCMC) techniques with the help of the program JAGS [7] interfaced with R [8] via rjags [9].

The nice thing using such a tool is that we have to take care only to describe the model, with instructions whose meaning is rather transparent. Then we have to provide the data, in our case n_V , n_P , n_{V_I} and n_{P_I} . The program samples the space of possibilities and returns lists of numbers (a ‘chain’) for each ‘monitored variable’ such that the frequency of the values in each list is proportional to the probability of that values of the variable (Bernoulli’s theorem). Here is, verbatim, the model:

```

model {
  nP.I ~ dbin(pA, nP)           # 1.
  nV.A ~ dbin(pA, nV)           # 2.
  pA   ~ dbeta(1,1)             # 3.
  nV.I ~ dbin(fffe, nV.A)       # 4. [ ffe = 1 - eff ]
  ffe  ~ dbeta(1,1)             # 5.
  eff  <- 1 - ffe               # 6.
}

```

We easily recognize in lines 1. and 2. of the code Eqs. (1) and (2), while line 4. stands for Eq. (3). Line 6. is simply the transformation of ‘ $1 - \epsilon$ ’ (‘ffe’ in the code) to ϵ , the quantity we want to trace in the chain. Finally lines 3. and 5. describe the *priors* of the ‘unobserved nodes’ that have no ‘parents’, in this case p_A and $1 - \epsilon$. We use in both cases a *uniform prior*, modeled by a Beta distributions with both parameters equal to 1 (we cannot go into the details of this choice that we consider quite reasonable, given the information provided by the data, and refer for the details to Ref. [5] and references therein). Finally, those who have no experience with JAGS can find in Ref. [5] several ready-to-run R scripts.

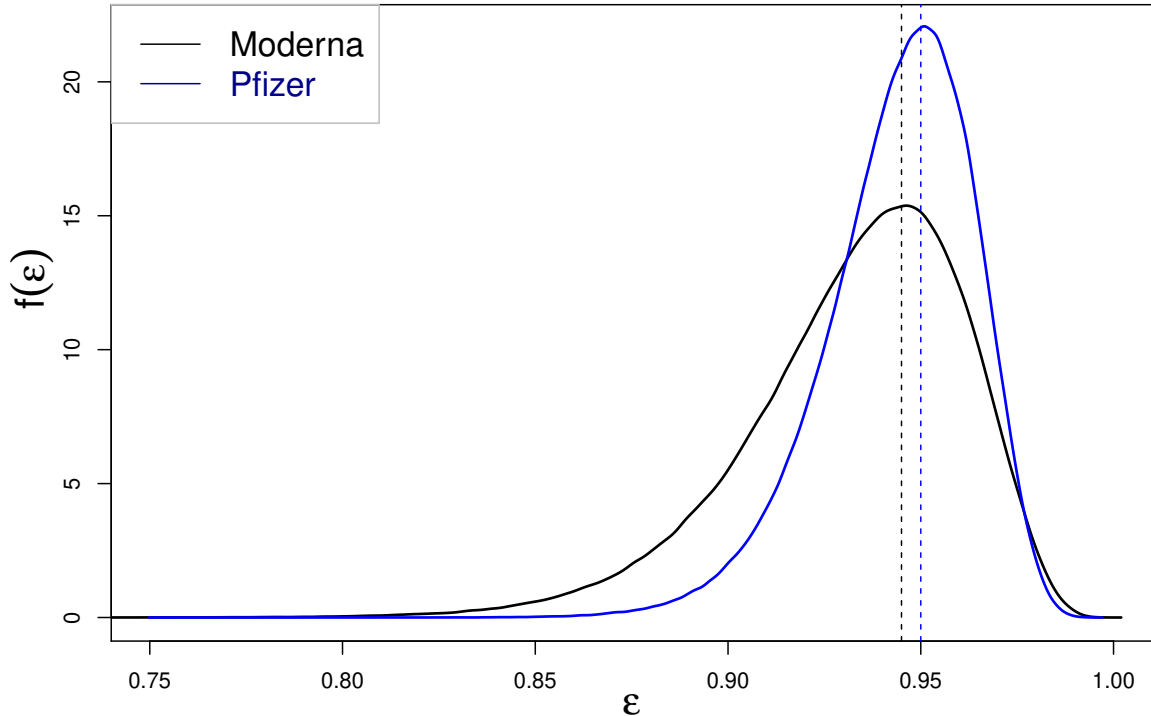


Figure 2: MCMC results for the vaccine efficacies (see text).

3 Results

We run the model assuming that the Moderna claims come from the entire sample, i.e. about thirty thousand people, equally divided between placebo and vaccine groups. This assumption is, anyway, not relevant for the efficacy estimates, as we have verified running the same model for samples one-tenth and one-hundredth of the full size and getting practically the same results for the efficacy.

The results of the MCMC sampling are reported in Fig. 2 with smooth curves that follow the profile of the histograms of the MCMC ‘data’ (one million steps have been chosen in order to reduce the sampling fluctuations): the black one (a bit broader) for Moderna; the blue one (a bit narrower) for Pfizer. The vertical dashed lines show the press release results of the two companies, that is 0.945 and 0.95, respectively. Indeed they correspond ‘practically exactly’ to the *modal values* of the distributions. But this is only one possible summary of a distribution, and not always the best one, especially if not associated to an *uncertainty* (and, certainly, the ‘ $p < 0.0001$ ’ given by both companies does not provide such information).

Usually our preference goes to the *mean* and the *standard deviation* because of rather general Probability Theory theorems, which make their use convenient for

further evaluations (to this standard deviation is related the concept of *standard uncertainty* [3]). Other ways to summarize with just a couple of number a probability distribution is the ‘central’ interval which contain the uncertain variable of interest at a given probability level *credible interval*. We report in the following table these summaries, that provide a quantitative evaluation of the uncertainty, together to the probability that the ‘true value’ of the efficacy is larger than 90%, reminding however that the most complete, quantitative information of the inference is contained in the curves of Fig. 2:

| | mean \pm stand. unc. | centr. 95% cred. int. | $P(\epsilon \geq 0.9)$ |
|---------|------------------------|-----------------------|------------------------|
| Moderna | 0.933 ± 0.029 | [0.866, 0.976] | 0.872 |
| Pfizer | 0.944 ± 0.019 | [0.900, 0.975] | 0.976 |

We would like to remind that this results do not depend on the exact values of n_V and n_P , provided they are enough larger than n_{P_I} .

As it is easy to expect, the MCMC also provides results on the other ‘unobserved’ nodes of the causal model, in our case p_A and n_{V_A} . We refrain to quote results on the ‘assault probability’, because they could easily be misunderstood, as they strongly depend, contrary to ϵ , on the precise values of n_V and n_P , being p_A a catch-all quantity embedding several real life variables, *including* the virus prevalence. We give, instead, the results concerning n_{V_A} , weakly dependent on p_A and n_{V_A} and that we expect to be *of the order of magnitude* of n_{V_I} . We get, in fact, respectively for Moderna and Pfizer, 89 ± 13 and 161 ± 18 (note that the standard uncertainties are not simply $\sqrt{n_{V_I}}$, as a rule of thumb would suggest).

4 Conclusion

The very recent announcements by Moderna and Pfizer give some hope of coping effectively with the pandemic. Their performances in terms of efficacy do not differ much, with the Pfizer’s one performing somehow better, definitely more that the bare numbers spread around by the media (“94.5% vs 95%”) would suggest.

Obviously, our results and our distributions depend on the very few data publicly available and on very simplifying assumptions, since there are many other variables that play important roles in considering which of the two is ‘better’ or ‘more promising’ in order to fight the pandemic, and on which we do not even try to enter, because they go far beyond our field of expertise. In fact, we wish to stress that our contribution is mainly methodological, even considering that further announcements will follow, as we strongly believe that a correct communication of a scientific result *must clearly state its uncertainties*, both in the official publication and (especially) in the press releases, avoiding to add excessive decimal figures that can be misinterpreted

by lay people (and not only) as a very sharp result. Moreover, as far as questions of practical interest are concerned, the uncertainty on the efficiency, possibly reported in detail by a probability distribution of its possible values, is without doubt important to develop realistic quantitative models and what-if scenarios of the development of the pandemic, once a part of the population has been vaccinated.

References

- [1] A. Fauci at NBC News, November 16, 2020, <https://youtu.be/8CG8aI4XCGw?t=32>.
- [2] Moderna Inc., Press Release, November 16, 2020, <https://investors.modernatx.com/news-releases/news-release-details/modernas-covid-19-vaccine-candidate-meets-its-primary-efficacy>.
- [3] International Organization for Standardization (ISO), *Guide to the expression of uncertainty in measurement*, Geneva, Switzerland, 1993.
- [4] G. D'Agostini, *The Waves and the Sigmas (To Say Nothing of the 750 GeV Mirage)*, arXiv:1609.01668 [physics.data-an].
- [5] G. D'Agostini and A. Esposito, *Checking individuals and sampling populations with imperfect tests*, arXiv:2009.04843 [q-bio.PE].
- [6] Pfizer Inc., Press Release, November 18, 2020, <https://www.pfizer.com/news/press-release/press-release-detail/pfizer-and-biontech-conclude-phase-3-study-covid-19-vaccine>.
- [7] M. Plummer, *JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling*, Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003), March 20–22, Vienna, Austria. ISSN 1609-395X, <http://mcmc-jags.sourceforge.net/>.
- [8] R Core Team (2018), *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- [9] M. Plummer, *rjags: Bayesian Graphical Models using MCMC*. R package version 4-10, <https://CRAN.R-project.org/package=rjags>.