Dear Dr Chen Liao,

Thanks for giving us the opportunity to improve the overall quality of our manuscript and the related code. We would also like to thank the reviewers for their detailed assessment of our work and thoughtful suggestions.

Below, we address the reviewers' recommendations in detail. In particular, we included a more thorough bibliography on the use of partial differential equations in the epidemiology field. We also addressed the main issues and most of the minor issues in the code/software to ensure better reproducibility and accessibility for newcomers.

Finally, we now provide a Zenodo repository (`https://doi.org/10.5281/zenodo.5549752`) containing all the scripts used within this study, following the *PCI Math & Comp Bio* guidelines.

We hereby resubmit our manuscript and hope it is now acceptable for recommendation.

Yours truly,
Bastien Reyné (for all co-authors)

Response to anonymous Reviewer #1

This article presents a method to model Covid-19 using a series of partial differential equations, which the authors claim is a novel method of modeling Covid, and apply the model to infection patterns in France. The benefit of the PDE approach is capturing multiple time bins, the authors track general time, the host age, time since infection, time since clearance, and time since vaccination. The R0 value is calculated using next generation operation, and investigates future trends in disease spread under future vaccination. Overall the paper was thorough, presented their methods well and were clear about the benefits of the PDE approach.

My main critique is the proposed uniqueness of the PDE approach. This work seems useful and the proposition of age structure being easier to understand in the context of a PDE system is likely true, however, other authors have PDE approaches already and this should be clearly acknowledged in the text with more citations to previous work.

> **Authors :** We clarified our text to remove the impression that our approach was unique.
> We enriched the bibliography with more references on the use of PDE approaches in epidemiology in general, and some examples of applications to the Covid-19 epidemic.

Introduction

Could add something on

> **Authors :** We enriched the part related to the use of PDEs.

Model

I am confused how the full version of your model, which is explicitly age-structured, can be applied to the French dataset if the hospital admission data is not age stratified? Doesn't that make the usefulness of the age structure in your model impossible to show via data?

> **Authors :** Unfortunately, the age-stratified model outputs related to the new hospital admissions could not be compared because these data are not publicly available.
> Furthermore, our main goal was not to make precise estimations of the age-related parameters. For this, we suggest the reading of the studies by Salje *et alli* and Lefrancq *et alli*, which provide age-stratified parameters of the trajectories of hospitalised patients estimated on French data.
> Nevertheless, modelling host age is important for two reasons. First, as we show in our sensitivity analysis using real data (namely the Covid-19 specific contact matrices), age-structure has a strong impact on the dynamics. Second, all estimates of Covid-19 severity show a huge effect on host age. More generally, our results confirm that age structure is a key component when modelling Covid-19 epidemics.

Results

The importance of variance from the contact matrix is interesting, would be worth discussing how you could include a time-varying version

> **Authors :** Varying the contact matrix over time is technically feasible but it comes across a limitation mentioned in the discussion: we do not know the underlying patterns of human behaviour making this matrix impossible to forecast. While we could investigate the effect of strong assumptions such as changes in patterns during the school holidays, it does not seem more relevant to set up those values arbitrarily. Hence, we choose the simplest option by assuming the matrix chosen for a run does not vary with time.

Overall

I think this paper is very solid. I could find no errors in the model derivation nor in the code implementation. I recommend this paper for publication.

> **Authors :** Thank you for the positive appreciation and also for checking out the model implementation. This is rare and we truly appreciated it!

Response to Facundo Muñoz (Reviewer #2)

General considerations

Main goal: Demonstrate the adaptation and extension of a recently published (by some of the same authors) approach which overcomes the Markovian hypothesis (lack of memory) of classical compartmental epidemiological models, with the objective of understanding the interplay between vaccination rates and age-structure in the Covid-19 pandemic in France.

Method: The methodology generalises the classical methodology based on Ordinary Differential Equations (ODEs) with respect to **time**, with Partial Differential Equations (PDEs) with respect to the **age** and to the **time since infection**, in addition to **time**. The authors deploy a specific compartment for vaccinated population and explore the use of alternative data sources to inform the age-structured contact patterns.

I think the title and the introduction could be improved to better identify the main topic of the manuscript. The title focus on the scientific questions about some mechanisms at play in the French Covid-19 pandemic ("the importance of the population age-structure"), without any mention to the methodology. By contrast, the abstract and introduction focus on the methodology, presenting the lack-of-memory issue as the knowledge gap to be addressed: "... we introduce an alternate formalism relying on partial differential equations..." (l. 35).

In my first read of the manuscript, I thought that the main topic was the methodology and the title was highlighting an application. It took me a second, more careful read to understand that the methodology had been introduced previously in Richard et al. (2021) and the present paper demonstrated how to tailor it to address different questions. But this is not clearly conveyed by neither the title nor the introduction (or the abstract).

> **Authors :** We agree with this remark and decided to make the following changes.
> First, we modified the title and the abstract to better reflect the contents of the manuscript.
> Second, we made some changes in the introduction to state that we a) use a (fairly uncommon) existing formalism rather than introducing a new one; b) better acknowledge the influence from Richard *et alli*; c) specify we use this model to address questions related to the past French Covid-19 epidemic.

Introduction

The section is very well structured, first stating the context and quickly identifying the knowledge gap. Namely, the need to model memory effects. It then explains the limits of two prior alternative approaches and makes a good case by arguing that multiplying compartments do not scale well and models quickly become very difficult to parameterise and interpret.

However, the first method is dismissed as a "workaround" which "artificially" increases the number of compartments. I think these are inappropriate dismissive qualifiers. Every model could be ultimately considered as *artificial* and used to *work around* reality. The question is how *useful* they are.

Authors : It was not our intention to be dismissive about ODEs models and we remove the term *artificially* which might have induced this impression.
We also deleted the term *workaround* from the abstract but kept it in the introduction, while underlying the simplicity of use of this solution. We hope this will not read as dismissive since our message is now more elaborated.

More specific statements about the relative merits would be much more informative. For instance, the authors could rather argue that modelling heterogeneities by age **continuously** is more *parsimonious* than introducing **artificial** boundaries between age groups. This formulation explicitly specifies what exactly is being considered *artificial*, by contrast to the current proposal.

I would have appreciated further introductory references to epidemiological modelling with PDEs. The only reference is Richard et al. (2021), which in turn says that it is a "less common and much more challenging" approach, without further references.

Authors : We enriched the sections about the use of PDEs and better explain their relative merits. More specifically, we explicitly mention non-linear properties of PDEs over ODEs, which are very relevant in our case.
We also added references related to the use of PDEs approaches in epidemiology in general, and to model the Covid-19 pandemic in particular.

Materials and methods

The presentation of the model is condensed, but well structured, rigorous and sufficiently detailed. Especially given that the main ideas were presented previously in some more detail.

I have only missed one or two sentences to discuss the recovery rate $\gamma^{mv}(a, i)$ from compartment I^{mv}_{aik}, about line 70, where the need for this compartment is introduced. In particular, justifying the choice for recovered individuals returning back to the compartment V_{ak} rather than R_{aj}. Stating explicitly that, in so doing, the time since vaccination k is preserved, and possibly other consequences of the choice.

Authors : Thank you for the positive feedback!
We now introduce the recovery rate $\gamma^{mv}(a, i)$ in this paragraph to acknowledge explicitly that mildly infected vaccinated individuals might recover differently (*e.g.* faster) than unvaccinated individuals.

l. 75: « ... the number of [+newly] severely infected individuals of age a at time t [-is][+are] given by the boundary condition[+s] »

Authors : We corrected it. Thanks for pointing that out.

In point 6 of Assumption S1, I think it is missing the case $l \in d$, or is there a reason for leaving it out?

Authors : Yes, it was indeed a mistake, we corrected it. Thanks for noticing.

I must confess that I could not quite follow the demonstration of the well-posedness of the system in appendix A.2, nor the derivation of the basic reproduction number in appendix A.3. It's been a long time since I last revisited Banach spaces, and I am not familiar with the utilised methods and results. Nevertheless, both sections provide enough references and pointers for interested readers.

Authors : Fortunately, at least two of the authors are experts on the topic and can vouch for the robustness of these results

Results

All the data and code were appropriately available for reproducing the results. Providing cached intermediate results which are lengthy to compute is very much appreciated. However, the documentation and comments are not sufficiently detailed.

For instance, the first script (`1_fit_vaccionation.R`) performs a calculation in parallel, which seems computationally demanding (I stopped it after a few minutes). It stores the results into an object called `results`. Coincidentally, there is a cached data file named `results.RData` which, judging by the name, seems to correspond with said computation. Yet there is no comment or indication confirming this, and loading such data file brings in a number of objects, none of them called *results*. It takes some more investigation to figure out that `results.RData` is created by the 5th script, and used in the 7th. So, it seems related to something else.

> **Authors :** This is accurate and the first script indeed did not save any results. These were only printed and to be reported in the `model.R` script.
> We now save the results within a cached file.
> We changed the name of the variable `results` into `vacc_params` and added comments within the script to remove any ambiguities.

Next, the second script warns from the beginning that it takes a few hours to run. Yet, it does not provide any pointer to the generated object (called `best` in the script), for which there is no cached results.

> **Authors :** Thank you for noticing this. We corrected this issue.

I don't pretend to be overly critic. It is apparent that the authors put some effort in cleaning up and commenting their code, and I truly appreciate it. Still, making code available and **accessible** to other people is difficult and takes a lot of time. Sometimes as much as producing the code itself.

The R package `modelvacc` is a wrapper around a set of C++ functions that implement the model equations and procedures. However, its complete lack of documentation (code comments, help pages) and tests somewhat hampers its reliability and re-usability by other researchers. I believe that this package is of considerable scientific value as a companion to the paper and can be instrumental in the adoption and improvement of the approach proposed by the authors. As such, it should be subject to the same high standards as the manuscript itself.

In summary, I would encourage the editors and the authors to improve the code a bit before, or after, publication.

> **Authors :** As suggested by another reviewer, we added comments in the `C++` files to link each function to the relevant equations in the manuscript. We added the same comments within the `R` script that used these functions for the ones using directly the `modelvacc` package we provide.

Discussion

The discussion is well structured, placing the results in context, and stating the relevant scientific conclusions given the strengths and limitations of the approach.

> **Authors :** Thanks!

Responses to Kevin Bonham (Reviewer #3)

Review

In *The importance of the population age-structure: insights from Covid-19 dynamics model structured by age, time since infection and acquired immunity*, Reyné and colleagues present a SIR model based on partial differential equations (PDE) as opposed to the typical ODE-based models. The authors state that this provides the ability to more faithfully capture the time that individuals spend within model compartments (memory) without the need to artifically inflate the number of compartments modeled. This has the advantage of increasing interpretability and flexibility of the model at the cost of more up-front effort at parameterization.

Unfortunately, I fear I lack the mathematical expertise to comment directly on the construction of the model and on its outputs. I will instead focus on the clarity of the writing and on the software, in the hopes that this will be useful.

Writing

The authors do an admirable job explaining the construction of their PDE model, including how individual terms relate to real-world scenarios and the source of values for initial parameterization. Though I am not able to readily follow the math, the descriptions in the text are clear and sensible. Figure 1 provides a useful reference for the modeled compartments, and the pathways between them.

Many of the limitations that I perceived are mentioned in the main text or in the discussion, and are adequately explained. One exception here is regarding the waning of immunity after vaccination (lines 104-105).

> Regarding the modelling of vaccine efficacy, for simplicity, we neglect immune waning, i.e. the decrease of immunity with time

The time-dependent changes in vaccine effectiveness strike me as a major source of uncertainty in this pandemic, and something for which models of this sort are well-suited to address (as claimed by the authors on line 34 as one motivation for this approach). In other portions of the manuscript, the authors imply that they are modeling this waning (eg ln 67 and ln 286). Perhaps it is clear from the equations, but I find myself unclear on whether this is actually accounted for or not.

> **Authors :** This is a very accurate point. We designed the model such that it can account for a decrease in vaccine efficacy (as suggested on line 67 in the model construction section). However, in the model parameterization we used to analyse the model, we neglected immune waning (as stated in lines 104–105). Our motivation to do so was double. First, we only consider a medium-term scenario and, hence, do not expect immune waning to cause any qualitative changes in the epidemic dynamics. This was justified on line 105
>
> > *This assumption is motivated by the fact that we consider a medium-term scenario and it could readily be modified.*
>
> Second, parametrizing correctly the immune waning requires data that was barely available and still remains preliminary.
> For these reasons, we choose to make the simplest choice and neglect it in our case.
> Nevertheless, as stated in the discussion (on line 286), this should not be neglected in a waning immune context, *e.g.* when considering longer-term scenarios.

Software

The authors make their software (written in C++ and R) available via an institutional gitlab repository. I was able to download a tarball of this code and follow the instructions to install dependencies on my laptop (Ubuntu xenial, R v4.0.1) Though (as mentioned in the README) many of the scripts take a long time to run, intermediate results are helpfully provided, and all of the code that I tried ran without errors until I interrupted it. The R portions of the code contain many helpful comments.

There are a few places where values that should perhaps be determined programmatically are hard-coded (eg here), and it might be nice if the parameters described in the paper could be found in a single configuration file (or something) rather than sprinkled throughout the scripts, as this would make it easier to tweak the assumptions of the model to see their effects, but these are very minor gripes.

> **Authors :** We now define (almost) all parameters –their baseline at least– in the `model.R` script, which can be seen as the configuration file.
> Moreover, we added (mainly for illustration purposes) the `run_example.R` script that shows how to change some inputs parameters and how to retrieve some outputs from a model run execution. We hope this will be enough to increase accessibility and usability for interested readers.

I find it quite admirable that code is provided in a runnable state for review. A few additional steps could make this code availability even stronger (though I hesitate to demand any of these steps as necessary).

1. Register / archive the code via an independent institutional repository such as zenodo.org or osf.io. Especially one that provides a digital object identifier (DOI). As it stands, there is no guarantee that this code won't disappear tomorrow.

   > **Authors :** We now provide a Zenodo link in order to satisfy *PCI Math & Comp Bio* guidelines.

2. Provide additional instructions for installing specific versions of packages. The provided *session_info.txt* file is a great start - using the renv package allowed me to reproduce the environment, at least as regards R dependencies. Additional information about C++ versions and compilation would also be welcome.

   > **Authors :** We added the C++ compiler version within the `README` file.

3. Provide some kind of indication within the scripts (just comments would be fine) which portions of the code take approximately what amount of time. I would have like to try to run the code that only takes minutes or hours so that I could inspect the output, but without knowing which parts might take days (or be infeasible on my laptop), this isn't practically possible.

   > **Authors :** We added this information through comments. We also specify which scripts needed a lot of RAM when a lot of runs were made.

4. Descriptions in the code that reference specific parts of the paper. Especially given my difficulty with understanding the math, being able to link the code directly to descriptions in the paper would be immensely helpful.

   > **Authors :** As suggested, we now indicate the corresponding equations from the manuscript each time it is appropriate, both in `R` and `C++` scripts.

The commit history on the publicly available project looks like it starts when the project was basically complete. Many people are uncomfortable sharing in-progress code (it's possible that version tracking was not even done earlier), and I don't think anything different is expected, which is why I'm not including it in my list of suggestions. But it's a shame.

> **Authors :** This could be interesting in a training or a history of science perspective but the interest seems limited to a general audience...

Results

The results, so far as I understand them, are impressive and on the whole, clearly presented. I am a bit unclear about figure 2 - in particular, I wonder if it would make more sense to split the ages into more plausible units, rather than just 10 year increments. For example, infants and toddlers are likely to be more dissimilar from school age kids than eg 9 vs 11 year olds. One might also consider breaking up based on availability of vaccine (eg the youngest kids still can't get vaccinated).

> **Authors :** We completely agree that the age groups used in the contact matrixes do not represent society contact patterns in the best way, especially for the youngest ages.
> This is cause the different data sources did not have homogeneous age groups and overlapped. For instance, the age contact matrices were given by 5 years bins, while the vaccination data were given with more realistic separations (5–11 years old corresponding to the primary school, 11–18 y.o. for (junior) high school students, 18–24 y.o. for college students, …).
> We eventually decided to use the 10 years bins, which was the one used for the age-differentiated disease severity, with the idea that it required little data modification compared to others delineations.
> We now clearly state our assumption in the Appendix.

For final publication, it might be nice to extend figures 5 and 6 with the most recent available data, as it currently ends in August. I don't know how feasible this is given the run-time of the code. Any further deviations from reality would not necessarily change the utility of this paper, but could be interesting fodder for discussion.

> **Authors :** While we understand the reviewer's motivation, this raises practical issues.
> We could just add the true data on the already done simulations, as done in Figure 1 and 2 hereafter. But it would require further discussion on why the simulations do not match with what was observed (namely changes in the contact rates partly induced by the implementation of an "Health Pass" by the French government disallowing for untested and unvaccinated individuals to go into public life places such as restaurants, theatres, concerts…). However, that would require adding more elements to the discussion that would depart from the discussion main points.
> An alternative would be to re-run the scripts to fit the contact rates till today and make some projections for the next months. But there are also many caveats in doing so since it would require to take into account the newly dominant omicron variant of concern (for which we know little and we are not even talking about immune escape possibilities) along with the waning immune (for which we still do not know that much) along with the boosting vaccination campaigns (for which we do not know the population overall adhesion and we do not have implemented). This option would require many efforts for results that would be nevertheless very uncertain given the current situation.
> For those reasons, we propose to stay with the current version of the figures, allowing for the paper to remain focused on the points already made.

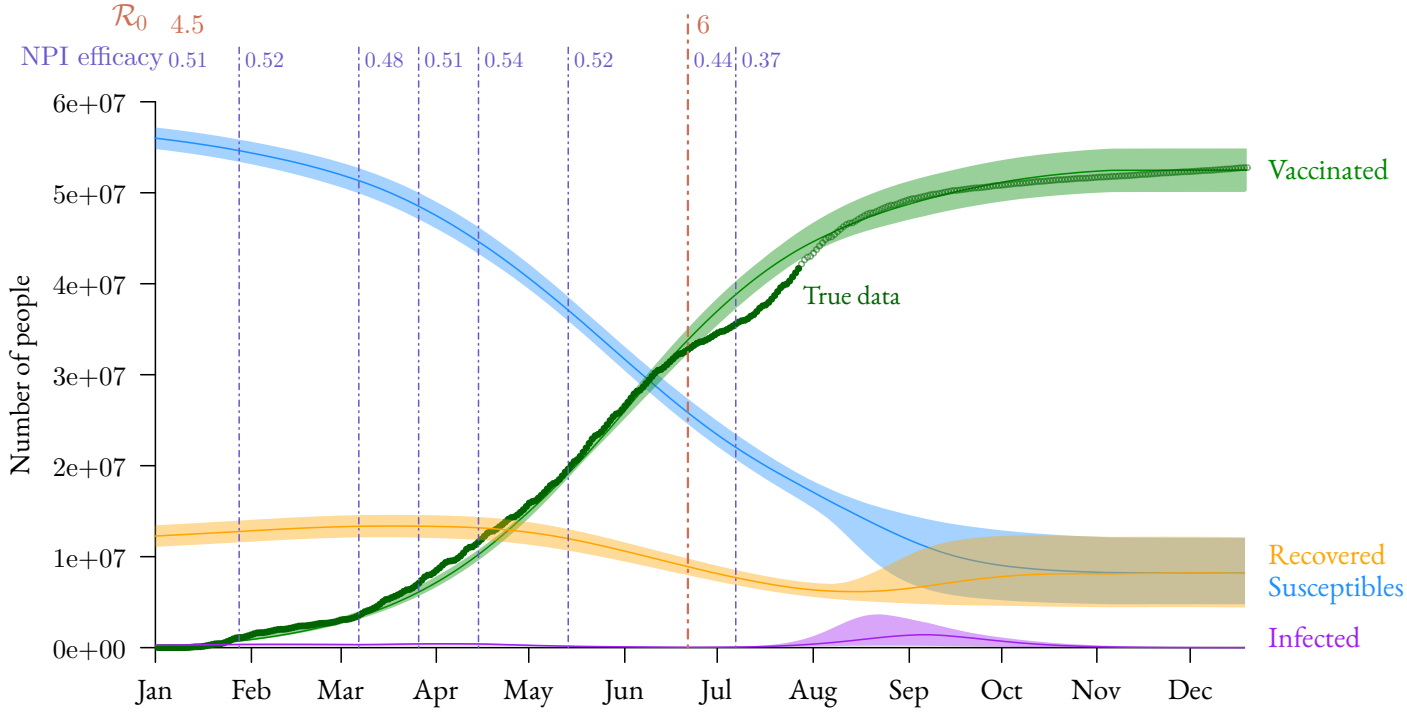

Figure 1: **Epidemic dynamics of the French SARS-CoV-2 epidemic in 2021.** We updated Figure 5 from the manuscript. Plain dots represent the data actually used, while empty dots represent data available after the analysis. The implementation in the summer of the "Health Pass" affected newly vaccinated individuals along with the contact rates but which were not fitted.

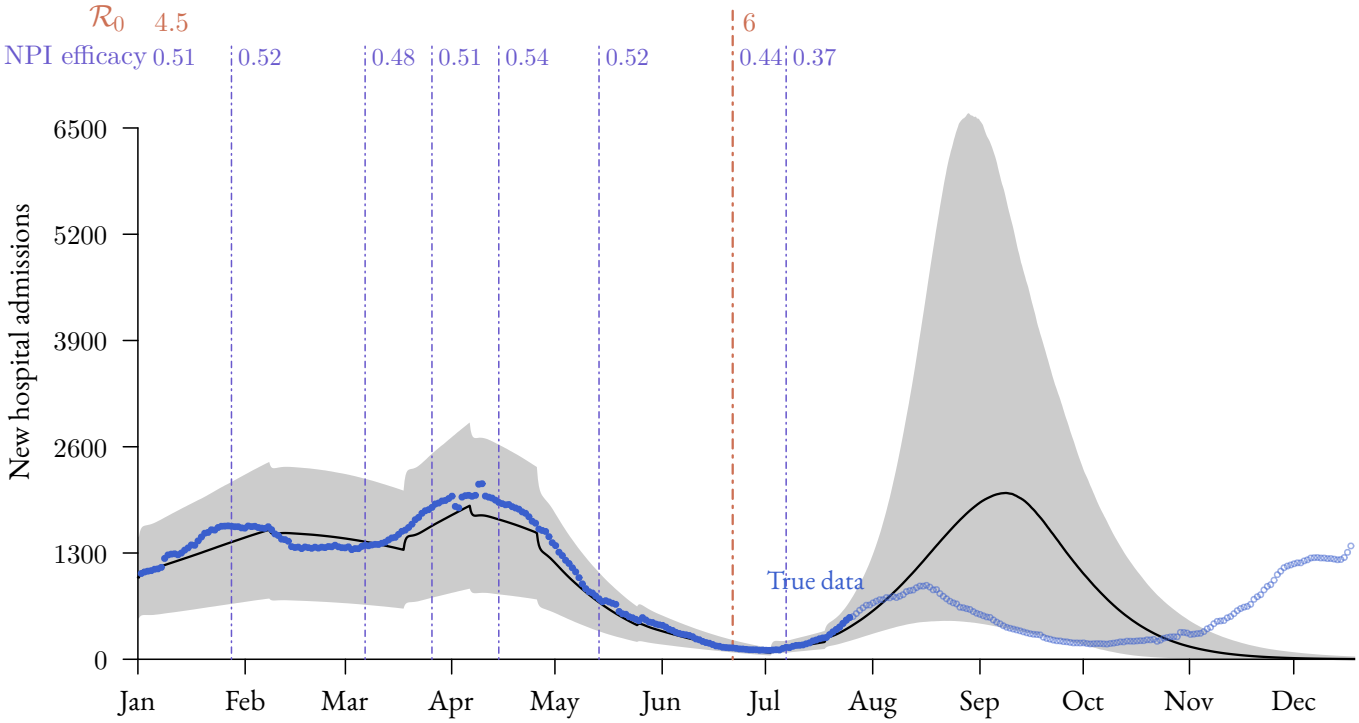

Figure 2: **Daily admissions for Covid-19 in French hospitals in 2021. Plain dots represent the data actually used, while empty dots represent data available after the analysis.** We updated Figure 6 from the manuscript