



Modelling freedom from disease - how do we compare between countries?

Rowland Raymond Kao based on reviews by Arata Hidano and 1 anonymous reviewer

Open Access

A recommendation of:

A modelling framework for the prediction of the herd-level probability of infection from longitudinal data

Published: 1st September 2021

Copyright: This work is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd/4.0/>

Aurélien Madouasse, Mathilde Mercat, Annika van Roon, David Graham, Maria Guelbenzu, Inge Santman Berends, Gerdien van Schaik, Mirjam Nielen, Jenny Frössling, Estelle Ågren, Roger Humphry, Jude Eze, George Gunn, Madeleine Henry, Jörn Gethmann, Simon J. More, Nils Toft, Christine Fourichon (2021), *bioRxiv*, 2020.07.10.197426, ver. 6 peer-reviewed and recommended by *PCI Animal Science* [10.1101/2020.07.10.197426](https://doi.org/10.1101/2020.07.10.197426)

Submitted: 23 July 2020, Recommended: 24 August 2021

In this paper, Madouasse et al. (2021) present a generalisable Bayesian method for calculating the probability that a herd is free from disease, based on its prior disease status, and using data (herd status over time over a sufficient number of herds to inform the model) and reasonable prior estimates of the sensitivity and specificity of tests being used to determine animal infection status.

Where available, the modelling approach can also include relevant additional risk factors.

By bringing all these factors together, it allows for most countries to use the same analytical approach on their data, with differences across datasets expressed in terms of the uncertainty around the central estimates.

Having a single methodology that generates both a central estimate of disease freedom, and uncertainty thus provides the opportunity (given typically available data) to compare the probability of freedom across different systems. This is relevant in terms of the context of trade (since international trade of livestock in many cases depends on disease freedom). It is also important when evaluating, for example, transnational burdens of disease - and with different regulations in place in different countries, this is invaluable and can be used, for example, to assess risks of zoonotic infection including for zoonotic infection emergence. In the BVD example provided, the point is made that, since regular testing would probably pick up infection rapidly, the addition of risk factors is most valuable where testing is infrequent. This emphasizes the advantages of direct incorporation of risk factors into a single modelling framework.

From a technical point of view, the analysis compares two different packages for the Markov Chain Monte Carlo (MCMC) implementation necessary to run the model. They show that, while there are some slight systematic differences, the estimates provided by the two methods are similar to each other; as one method is approximate but substantially more stable and generally much more computationally efficient, this is an important outcome. Both implementations are freely available and with relevant additional software made similarly available by the authors. This is extremely welcome and should encourage its general adoption across different countries.

No single model can of course account for everything. In particular, the reliance on past data means that there is an implicit assumption common to all purely statistical methods that the underlying risks have not changed. Thus projections to altered circumstances (changing underlying risk factors or systematic changes in testing or test performance) cannot so easily be incorporated, since these factors are complicated by the dynamics of infection that lie outside the modelling

approach. Of course the well known quote from George Box that "all models are wrong" applies here - the generality of approach, statistical robustness and open source philosophy adopted make this model very useful indeed..

References

Madouasse A, Mercat M, van Roon A, Graham D, Guelbenzu M, Santman Berends I, van Schaik G, Nielen M, Frössling J, Ågren E, Humphry RW, Eze J, Gunn GJ, Henry MK, Gethmann J, More SJ, Toft N, Fourichon C (2021) A modelling framework for the prediction of the herd-level probability of infection from longitudinal data. bioRxiv, 2020.07.10.197426, ver. 6 peer-reviewed and recommended by PCI Animal Science.
<https://doi.org/10.1101/2020.07.10.197426>

Cite this recommendation as:

Rowland Raymond Kao (2021) Modelling freedom from disease - how do we compare between countries?. *Peer Community in Animal Science*, 100007.
<https://doi.org/10.24072/pci.animsci.100007>

Reviews

Revision round #2

07 Jul 2021

Author's Reply

First of all, very sorry for the delays in replying to you on your paper.

However, I am pleased to say that we have now had replies from your two reviewers. While they have made some comments, and responding to them I think will improve the paper, there is nothing there to prevent recommendation - therefore I suggest that you take a look at those comments, and decide for yourselves how best to further to revise it. Once this is done, I will certainly be more than happy to recommend it without need for further review. It is a very useful addition to the literature.

We would like to thank you and the reviewers for helpful comments. Below are the responses to the second round of comments.

Reviews

Reviewed by Arata Hidano, 2021-06-22 06:47

My sincere apologies for this huge delay in my response. Thank you very much for addressing my previous comments. It is great to see that the Stan use (with huge efforts to account for latent discrete variables in Stan) has addressed many of the difficulties the authors faced previously.

My final minor comment is about L498 – 499 “This could explain the apparent cumulative effect of the number of introductions”. Does ‘number of introductions’ mean the number of animals introduced? Also I have probably missed somewhere, but did you try different functional forms of the variable representing the number of animals introduced (other than the raw number and the natural log)? Both the uses of raw number and natural log still assume a monotonous (increase or decrease) association between this variable and the outcome, hence, it will show the cumulative effect by specification, even if the true association between these two variables was not a monotonous one (e.g. quadratic, polynomial). This may happen, for instance, herds purchasing a large number of animals may have better biosecurity and scrutinize the source herd status, leading to a decreased risk of disease introductions from such herds. If this was not already addressed, please clarify this issue.

The cumulative effect of the number of animals refers to what is presented in the left-hand side panel of Figure 5. In this Figure, lag 1 refers to the end of the interval (closest to the current month) evaluated and lag 2 to its start (further in the past). For example, the cell on the bottom right of the plot represents the model evaluating the total number of animals between the current month (lag1 = 0) and 24 months before (lag 2 = 24). We have added an explanation of how to read Figure 5 in the caption which we hope will clarify this for readers.

Regarding the functional forms for the associations evaluated, up to now, we had only included either the untransformed number of animals introduced or its natural logarithm. As can be seen in Figure 7, the natural logarithm allows for an increase in the probability of becoming positive that is less steep as the number of animals introduced increases. As rightly pointed out, this probability could decrease when the number of animals introduced is high, which had not been tested. In response to this comment, we fitted GLMs with quadratic and cubic terms for the number of animals introduced and compared the AICs of these 2 models with AIC of the model with the logarithm of the number of animals introduced (reference model). The model with the cubic term had an AIC that is roughly the same as the AIC of the reference model (difference of 1.5 in favour of the model with the polynomials), but the coefficients for the quadratic and cubic terms were not significant. The model with a quadratic term had an AIC that was again close to the one of the reference model, even though the difference was slightly bigger (3.5). This time, the coefficient for the quadratic term was significant. The model indeed predicted a drop in the probability of becoming positive after 35 animals introduced. However, because the AICs are so close, we consider both models to be equivalent, probably because the number of instances with a high number of animals introduced does not bring enough evidence in favour of either hypothesis. Given this, we prefer to keep our model as is. Below are the details of the different models we tested.

Reference model

Call:

```
# glm(formula = formula, family = binomial(link = "logit"), data = data)
```

#

Deviance Residuals:

```
#   Min     1Q   Median     3Q    Max
```

```

# -1.2418 -0.5593 -0.5593 -0.5593 1.9660
#
# Coefficients:
#
#           Estimate Std. Error  z value Pr(>|z|)
# (Intercept)   -1.77610   0.04358  -40.757 < 2e-16 ***
# log_nAnim_8_8  0.37860   0.06122   6.184 6.24e-10 ***
# ---
# Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
#
# (Dispersion parameter for binomial family taken to be 1)
#
#   Null deviance: 3865.5  on 4507  degrees of freedom
# Residual deviance: 3831.2  on 4506  degrees of freedom
# AIC: 3835.2
#
# Number of Fisher Scoring iterations: 4

Model with quadratic term

# Call:
# glm(formula = formula, family = binomial(link = "logit"), data = data)
#
# Deviance Residuals:
#   Min       1Q   Median       3Q      Max
# -1.5560 -0.5615 -0.5615 -0.5615  1.9623
#

```

Coefficients:

#	Estimate	Std. Error	z value	Pr(> z)
# (Intercept)	-1.767740	0.043003	-41.107	< 2e-16 ***
# nAnim_8_8	0.138272	0.023710	5.832	5.48e-09 ***
# nAnim_squared_8_8	-0.002371	0.000823	-2.881	0.00397 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#

(Dispersion parameter for binomial family taken to be 1)

#

Null deviance: 3865.5 on 4507 degrees of freedom

Residual deviance: 3825.7 on 4505 degrees of freedom

AIC: 3831.7

#

Number of Fisher Scoring iterations: 5

Model with cubic term

Call:

glm(formula = formula, family = binomial(link = "logit"), data = data)

#

Deviance Residuals:

#	Min	1Q	Median	3Q	Max
#	-1.5524	-0.5614	-0.5614	-0.5614	1.9624

#

Coefficients:

```

#               Estimate Std. Error    z value Pr(>|z|)
# (Intercept)   -1.768e+00  4.327e-02 -40.854 <2e-16 ***
# nAnim_8_8     1.390e-01  3.573e-02  3.890  0.0001 ***
# nAnim_squared_8_8 -2.442e-03  2.704e-03  -0.903  0.3664
# nAnim_cubed_8_8  1.085e-06  3.890e-05  0.028  0.9778
# ---
# Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
#
# (Dispersion parameter for binomial family taken to be 1)
#
# Null deviance: 3865.5 on 4507 degrees of freedom
# Residual deviance: 3825.7 on 4504 degrees of freedom
# AIC: 3833.7
#
# Number of Fisher Scoring iterations: 6

```

Line 548 should probably read 'despite the fact that only one...'.
 Changed

Reviewed by anonymous reviewer, 2021-05-27 13:19

Thanks for thoroughly considering my previous comments - I think the manuscript is now much improved. I particularly like the addition of the Stan version of the model, although this has given rise to an additional point that I would like to raise:

To what extent is the difference in performance and inference between the JAGS and Stan models due to (1) implementation of MCMC (plain old Gibbs vs HMC) and (2) the inherent difference between the forward algorithm and the use of discrete latent classes? I am well aware of the limitation that Stan cannot estimate discrete parameters, but have not (yet!) tried using the approach suggested by Damiano. But is the forward algorithm an exact re-parameterisation of the discrete latent class model, or is it an approximation that holds under

certain regulatory conditions? This may be answered by Damiano - I have only had time to skim that document unfortunately - but I think discussing it in a brief paragraph would be useful, particularly because figure 7 (and table 3) seems to show a slight discrepancy in the estimates for Se and Sp between the two versions of the model (lower Se and higher Sp for JAGS). If HMC was just giving a performance advantage then I would expect the posteriors to be more similar: this looks to me more like there is a more fundamental difference between the models. If the forward algorithm is using some approximation in place of the discrete latent states then that would explain both the discrepancy and the performance advantage. On the other hand, perhaps the difference is due to poor convergence (and therefore poor approximation to the posterior) for the JAGS model? I would certainly expect HMC to perform better for these types of models (with typically poor mixing), so that is no big surprise. To put this another way: could the forward algorithm be implemented in JAGS, and if so would you expect the difference to be reduced or remain the same?

It was hard to answer this because the gain in performance is due in part to the re-parameterisation and in part to the use of HMC, but it is not possible to quantify the relative parts from the present work. We found useful answers to these question in the following paper:

Yackulic, C.B., Dodrill, M., Dzul, M., Sanderlin, J.S., Reid, J.A., 2020. A need for speed in Bayesian population models: a practical guide to marginalizing and recovering discrete latent states. *Ecol. Appl.* 30, 1–19. <https://doi.org/10.1002/eap.2112>

In short, marginalisation is associated with a massive increase in speed, but marginalised versions of the same model are still estimated much faster in Stan than in JAGS (See Figure 1 of the paper). Both formulations should give the same results. To clarify this, we added the following paragraph to the discussion (starting L639):

When estimated in either JAGS or BUGS, discrete latent state models such as HMMs are known to converge slowly; and the autocorrelation in the draws from the posterior distributions is usually high. Yackulic et al. (2020) showed that the marginalisation of the latent states considerably reduces the time needed to estimate the parameters of such models while returning the same estimates. We did not implement this approach in JAGS, although this would have been possible using the ones trick, as explained in the article by Yackulic et al. (2020). The forward algorithm is a type of marginalisation that partly explains the better performance of the Stan version of the model. However, Yackulic et al. (2020) also compared the speed of the marginalised versions of their model in different programmes and observed that Stan was orders of magnitude faster than JAGS.

I also have the following minor comments:

- The first sentence of the abstract is phrased awkwardly: "there are existing" should be "there are", but it would be even better to turn the sentence around e.g. "The collective control programmes (CPs) that exist for many infectious diseases of farm animals rely ... "

Thanks for the suggestion. This has been changed.

- Table 2: I would not present these p-values with such precision! Maybe either <0.001 or if you prefer e.g. $<10^{-6}$

Changed.

Decision round #2

First of all, very sorry for the delays in replying to you on your paper.

However, I am pleased to say that we have now had replies from your two reviewers. While they have made some comments, and responding to them I think will improve the paper, there is nothing there to prevent recommendation - therefore I suggest that you take a look at those comments, and decide for yourselves how best to further to revise it. Once this is done, I will certainly be more than happy to recommend it without need for further review. It is a very useful addition to the literature.

Preprint DOI: <https://doi.org/10.1101/2020.07.10.197426>

Reviewed by [Arata Hidano](#), 22 Jun 2021 06:47

My sincere apologies for this huge delay in my response. Thank you very much for addressing my previous comments. It is great to see that the Stan use (with huge efforts to account for latent discrete variables in Stan) has addressed many of the difficulties the authors faced previously.

My final minor comment is about L498 – 499 “This could explain the apparent cumulative effect of the number of introductions”. Does ‘number of introductions’ mean the number of animals introduced? Also I have probably missed somewhere, but did you try different functional forms of the variable representing the number of animals introduced (other than the raw number and the natural log)? Both the uses of raw number and natural log still assume a monotonous (increase or decrease) association between this variable and the outcome, hence, it will show the cumulative effect by specification, even if the true association between these two variables was not a monotonous one (e.g. quadratic, polynomial). This may happen, for instance, herds purchasing a large number of animals may have better biosecurity and scrutinize the source herd status, leading to a decreased risk of disease introductions from such herds. If this was not already addressed, please clarify this issue.

Line 548 should probably read ‘despite the fact that only one...’.

Reviewed by anonymous reviewer, 27 May 2021 13:19

Thanks for thoroughly considering my previous comments - I think the manuscript is now much improved. I particularly like the addition of the Stan version of the model, although this has given rise to an additional point that I would like to raise:

To what extent is the difference in performance and inference between the JAGS and Stan models due to (1) implementation of MCMC (plain old Gibbs vs HMC) and (2) the inherent difference between the forward algorithm and the use of discrete latent classes? I am well aware of the limitation that Stan cannot estimate discrete parameters, but have not (yet!) tried using the approach suggested by Damiano. But is the forward algorithm an exact re-parameterisation of the discrete latent class model, or is it an approximation that holds under certain regulatory conditions? This may be answered by Damiano - I have only had time to skim that document unfortunately - but I think discussing it in a brief paragraph would be useful, particularly because figure 7 (and table 3) seems to show a slight discrepancy in the estimates for Se and Sp between the two versions of the model (lower Se and higher Sp for

JAGS). If HMC was just giving a performance advantage then I would expect the posteriors to be more similar: this looks to me more like there is a more fundamental difference between the models. If the forward algorithm is using some approximation in place of the discrete latent states then that would explain both the discrepancy and the performance advantage. On the other hand, perhaps the difference is due to poor convergence (and therefore poor approximation to the posterior) for the JAGS model? I would certainly expect HMC to perform better for these types of models (with typically poor mixing), so that is no big surprise. To put this another way: could the forward algorithm be implemented in JAGS, and if so would you expect the difference to be reduced or remain the same?

I also have the following minor comments:

- The first sentence of the abstract is phrased awkwardly: "there are existing" should be "there are", but it would be even better to turn the sentence around e.g. "The collective control programmes (CPs) that exist for many infectious diseases of farm animals rely ... "
- Table 2: I would not present these p-values with such precision! Maybe either <0.001 or if you prefer e.g. $<10^{-6}$

Revision round #1

01 Dec 2020

Author's Reply

[Download author's reply \(PDF file\)](#)

Dear Rowland,

Apologies for the long delay in replying. We have made some major changes to the model which have made it better, but which have taken a long time.

Please find our replies to the reviewers' comments in the attached pdf.

Kind regards,

Aurélien

Decision round #1

Two reviews have been received for this preprint. Both reviewers have highlighted the importance of the problem - i.e. the estimation of herd "freedom from infection" status based on a series of partial observations with imperfect testing, which is particularly a problem where disease incidence is long, and the disease itself with long and variable latent periods. This is the case for BVD (the disease observed here) but also problematic long duration bacterial diseases such as paratuberculosis and bovine Tuberculosis. The development of a general framework, described here, is therefore welcome. Both reviewers have indicated the importance of the problem and relevance of the approach. Both have also been very positive about the execution

of the approach, but both have also highlighted methodological issues that, while they do not fundamentally compromise this analysis, do require addressing. A particular issue is the choice and use of priors, highlighted by both reviewers though with different perspectives. It also seems to me that the choice of prior is one area where other adopters of the framework would have to exercise the greatest individual judgement (also, as indicated in their discussion of this from lines 635 onwards, they are very aware of this) and therefore the reviewers comments require considerable attention. I also fully agree with one reviewer on the justification of the use of narrow priors - it suggests possible issues with the model specification itself or with the data, and is not necessarily an issue of the prior. While some of the other reviewer comments are outside my expertise, they all seem sensible and should be responded to by the authors. I also include here some comments from a 3rd reviewer who did not feel it sufficiently within their expertise to offer a full review: "The submitted manuscript aims to predict herd status using all prior information over time. This is an interesting new approach and worth being published. However, author imply in the introduction that they consider the number of animals tested (a frequent feature of herd testing) whereas they don't consider this option. Only tank milk testing is included. Moreover, I believe there is a mistake in formula 15 for the probability of infection for test-negative herds (I think this should read: $(1-T+)[Sp.(1-pS+) / (Sp.(1-pS+) + (1-Se).pS+)$ i.e. the negative predictive value for negative test results). Having said that, my advice is to consult a statistician for quality checking of the code. I recommend consulting my colleagues Prof Geoff Jones (g.jones@massey.ac.nz) or Prof Wesley Johnson (wjohnson@ics.uci.edu). The text requires an overhaul for clarity and typos. The manuscript is definitely worth reviewing as it is topical at a time of emerging infectious diseases, so it would likely attract citations. I cannot review this submission myself as it would need considerable more time and more statistical expertise than I can offer." Overall, this is likely to be a very good contribution to the literature and I look forward to seeing a revision.

Reviewed by [Arata Hidano](#), 21 Sep 2020 03:14

Please find enclosed review comments.

[Download the review \(PDF file\)](#)

Reviewed by anonymous reviewer, 30 Nov 2020 19:00

The authors present an MCMC modelling framework intended for use with serial observations of disease status, with a particular focus on the use of imperfect diagnostic tests. I find the motivation to be clear and compelling, the methodological approach to be sound, and the manuscript generally well written. However, I do have a number of comments that I think may improve the manuscript:

1. I am concerned about the apparent re-use of the same data for both informing the priors for Se and Sp (lines 454-456; 488-490) and presumably also contributing to the likelihood function of the fitted model, presuming that these same data points are used for model fitting? In that case the data will have a double contribution to the posterior, which is not a good idea. Can you either reassure the reader that the same data points are not used twice (either directly or indirectly) or base the priors on other information that is not already used in the model?
2. There seem to be considerable challenges with convergence with the model, presumably resulting from high autocorrelation. However, the brief mention of MCMC convergence diagnostics on lines 468-473 does not specifically mention the degree of autocorrelation

- was it not a problem after all? Also, can you give details on the effective sample size based on your 4 chains of 5000 correlated samples with thin of 20? Even if correlation in the 1000 samples is extremely low I would still worry about a minimum effective sample size of 400 being reached for all parameters.
3. A mixture of normals is assumed for the observed OD values, but (to me) figure 5 indicates a somewhat imperfect fit to the data: the left peak seems right-skewed and the right peak seems left-skewed, i.e. there are more observations in the middle than expected. This may have the effect of exaggerating the diagnostic capability of the test, as the "crossover" between the assumed normal distributions would be a lot less than that in the real data. How might this have affected your results?
 4. Model 1 uses a beta distribution for T1, but Models 2-4 use a logistic regression for T1 (Table 3). However, you have not fully considered the effect of the priors being on different scales (i.e. probability scale vs logit scale) - an excellent example of the problems this can give can be found in "Gelman A, Simpson D, Betancourt M. The prior can often only be understood in the context of the likelihood. *Entropy*. 2017;19:1–13.". Please re-run Model 1 using an intercept-only model with prior equivalent to the other models so that we can be sure the difference between the models is really due to inclusion of the risk factor, and not due to an indirect effect of the prior for the "average" probability being on different scales. Also, I would strongly advise to center the predictors to zero so that the intercepts can be directly compared between models - this should also have a beneficial effect on convergence.
 5. I found figures 7, 9 and 10 to be difficult to interpret due to overlapping densities - is there some way of presenting these results in a different manner, perhaps using a caterpillar plot / horizontal boxplot, or maybe an ECDF plot?

I also have a number of more minor comments:

- Line 175: I agree it that missing observations are workable within the framework, but should you mention that the fewer missing observations the better (both in terms of convergence and precision of inference)?
- Line 195: I think the preferred reference for JAGS is "1. Plummer M. JAGS : A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling JAGS : Just Another Gibbs Sampler. Proc 3rd Int Work Distrib Stat Comput (DSC 2003) [Internet]. 2003. p. March 20–22, Vienna, Austria. ISSN 1609-395X. Available from: <http://www.ci.tuwien.ac.at/Conferences/DSC-2003/Drafts/Plummer.pdf>
- Line 217: The observed distributions are never identical; maybe "are drawn from the same distribution" or "have close to the same distribution".
- Equation 4: Obviously the effect of equation 4 is that $P(s_i+)$ is EITHER equal to T1 or T2 conditional on the value of the dichotomous variable $St+$ - but I think this might it be clearer to the reader if the equation is written using definition by cases notation (<https://planetmath.org/definitionbycases>) rather than the implicit multiplication by zero trick (of course they give equivalent results, but this way reinforces to the reader that $St+$ is dichotomous).
- Line 298-333: The discussion of the defined latent state is important but perhaps out of place in the materials and methods section?
- Lines 369-370: What about values between 35 and 60?
- Lines 384-388: Technically it is not being "Bayesian" that takes the time, but relying on computationally intensive fitting methods (i.e. MCMC). Also, could you briefly mention the more Bayesian methods of approaching this e.g. stochastic variable selection - see "O'Hara, R.B., Sillanpää, M.J., 2009. A review of Bayesian variable selection methods: what, how and which. *Bayesian Anal.* 4, 85–117. <https://doi.org/10.1214/09-BA403>." I

don't think these would be computationally feasible in your case, but it is still a potential alternative to resorting to frequentists methods of fitting a simplified version of the model.

- Line 401-402: But as I understand you are using a log link function within the GLM, so by including an additive effect of $\log(N)$ on the log scale you still end up with a multiplicative effect of N on the scale of the response distribution, so I am not entirely sure that this "allows a decreasing effect of each animal" as you state. If you wanted to do this perhaps you could include a 2nd degree polynomial of $\log(N)$? Or have I misunderstood something?
- Line 640: Perhaps "failure" rather than "absence" of convergence?
- Lines 643-647: Would multiple imputation (i.e. running the MCMC sampler 20 times using 20 different fixed values of Se/Sp from across the prior) be an alternative to allow these models to be fit? Or perhaps using a different sampling strategy that is more efficient in these situations such as Hamiltonian Monte Carlo? Obviously the latter is not possible in JAGS, but multiple imputation has been done in these situations.
- Line 720: Please do not use the word "uninformative" in relation to priors - there is no such thing. Either "minimally informative" or "non-informative" is fine.
- Could you give some JAGS code for (possibly simplified versions of) models 1-4 as an appendix? I realise that this is contained within the GitHub repo, but as far as I can see an interested reader needs to mentally paste together the R code that generates the JAGS model to be able to assess the code. It would be easier if the model representation was available directly.

Finally, I would like to apologise for the lengthy delay in providing this review, which was due to an unexpectedly increased workload at my end - sorry.