



This is a repository copy of *The utility of the Laplace effect size prior distribution in Bayesian fine-mapping studies*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/170033/>

Version: Accepted Version

Article:

Walters, K. orcid.org/0000-0002-5718-5734, Cox, A. and Yaacob, H. (2021) The utility of the Laplace effect size prior distribution in Bayesian fine-mapping studies. *Genetic Epidemiology*, 45 (4). pp. 386-401. ISSN 0741-0395

<https://doi.org/10.1002/gepi.22375>

This is the peer reviewed version of the following article: Walters, K, Cox, A, Yaacob, H. The utility of the Laplace effect size prior distribution in Bayesian fine-mapping studies. *Genetic Epidemiology*. 2021; 45: 386– 401, which has been published in final form at <https://doi.org/10.1002/gepi.22375>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley's version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

The utility of the Laplace effect size prior distribution in Bayesian fine-mapping studies

Kevin Walters^{1*}, Angela Cox², Hannuun Yaacob^{1,3},

¹ School of Mathematics and Statistics, University of Sheffield, Sheffield, UK

² Department of Oncology, Sheffield Cancer Research Centre, University of Sheffield Medical School, Sheffield, UK

³ Department of Applied Statistics, Faculty of Economy and Administration, University of Malaya, Malaysia

* E-mail: k.walters@sheffield.ac.uk, Telephone: +44 (0)114 2223720

Abstract

The Gaussian distribution is usually the default causal SNP effect size prior in Bayesian population-based fine-mapping association studies, but a recent study showed that the heavier-tailed Laplace prior distribution provided a better fit to breast cancer top hits identified in genome-wide association studies. We investigate the utility of the Laplace prior as an effect size prior in univariate fine-mapping studies. We consider ranking SNPs using Bayes factors and other summaries of the effect size posterior distribution, the effect of prior choice on credible set size based on the posterior probability of causality, and on the noteworthiness of SNPs in univariate analyses. Across a wide range of fine-mapping scenarios the Laplace prior generally leads to larger 90% credible sets than the Gaussian prior. These larger credible sets for the Laplace prior are due to relatively high prior mass around zero which can yield many non-causal SNPs with relatively large Bayes factors. If using conventional credible sets, the Gaussian prior generally yields a better trade off between including the causal SNP with high probability and keeping the set size reasonable. Interestingly when using the less well utilised measure of noteworthiness, the Laplace prior performs well, leading to causal SNPs being declared noteworthy with high probability, whilst generally declaring fewer than 5% of non-causal SNPs as being noteworthy. In contrast, the Gaussian prior leads to the causal SNP being declared noteworthy with very low probability.

Introduction

In Bayesian fine-mapping studies the Gaussian distribution is the default choice of prior for the effect size, or equivalently, the log-odds ratio (Wakefield [2009]; Spencer et al. [2016]; Benner et al. [2016]; Chen et al. [2015]). The zero-mean Gaussian distribution is a sensible choice of effect size prior in that it is symmetric around zero and has reasonably fast decaying tails. However the main reason for using the Gaussian prior is mostly due to be its computational simplicity. With an asymptotic Gaussian likelihood [Wakefield, 2009] the Gaussian distribution is a conjugate prior for the effect size leading to a Gaussian posterior distribution. The Bayes factor is easily calculated with a zero-centred Gaussian prior, giving the so-called Wakefield Bayes factor (WBF), even with a hierarchical Bayesian model [Spencer et al., 2015].

Other priors have been utilised in fine-mapping studies including the t distribution [Marchini and Howie, 2010], the normal-gamma prior (Boggis et al. [2016]; Alenazi et al. [2019]) and the Laplace prior [Hoggart et al., 2008]. There are also methods starting to become available that implement methods based on binned empirical effect sizes that do not assume specific parametric distributions [Vsevolozhskaya and Zaykin, 2019]. These latter approaches require careful choice of bin number and size but allow flexibility in the prior form and can use available known causal SNP effect sizes to provide a well-matched summary of the disease-specific effect sizes discovered to date. Walters et al. [2019] used a Bayesian model selection approach to show that, based on a large number of breast cancer genome-wide association study (GWAS) top hits, the Laplace prior showed a much better fit than the Gaussian prior (where the fit was determined by the posterior probability of the model given the GWAS top hits data). The only consideration given to the Laplace prior specifically in fine-mapping is in Hoggart et al. [2008] who used a penalised regression LASSO-type approach which is not truly Bayesian.

The Laplace prior shares many of the advantages of the Gaussian distribution as a prior on the effect size when used in conjunction with an asymptotic Gaussian likelihood: there is a closed-form expression for the posterior distribution which does not require MCMC as does the normal-gamma prior (Boggis et al. [2016]; Alenazi et al. [2019]); and calculating the Bayes factor does not require a Laplace approximation, as does the t distribution prior implemented in SNPTEST [Marchini and Howie, 2010]. The zero-mean Laplace prior has more mass close to zero than the Gaussian and

also has heavier tails. This makes it an excellent choice of prior distribution for complex diseases where there are likely to be many causal SNPs with small effect sizes but also a small number of relatively large effect sizes. Walters et al. [2019] demonstrated that for the GWAS breast cancer top hits, the Gaussian distribution couldn't simultaneously place a high mass around zero whilst retaining sufficiently heavy tails to provide a good fit.

A few previous studies have compared the fine-mapping performance of either frequentist or Bayesian approaches (Spencer et al. [2015]; Van de Bunt et al. [2015]; Spencer et al. [2014]). However, where Bayesian approaches have been considered, the prior is always assumed to be Gaussian. To the best of our knowledge no other authors have examined the fine-mapping performance of the Laplace prior.

The approaches used in this paper can be applied to the situation where there are confounding variables since a simple reparameterisation trick means they do not affect the Bayes factor or posterior distribution [Wakefield, 2009]. We only consider the univariate case, which is appropriate for fine-mapped regions with at most one causal SNP. A multi-SNP approach could readily be taken but this makes it harder to identify situations under which the choice of prior is important (since there become many more variables to vary).

The Laplace prior was shown to be a better fit in Walters et al. [2019] but the question of interest is really whether the choice of prior is important in fine-mapping, and if so, in which situations. For example in studies with large sample sizes it seems reasonable to hypothesize that the choice of prior would be less important than in small studies since the likelihood contains more information as the sample size increases. The Bayes factor (one interpretation of which is the ratio of the posterior odds of the SNP being causally associated to the prior odds of causal association) is the most widely used Bayesian statistic in fine-mapping studies because it compares the evidence under the null and the alternative hypothesis, but there are other statistics that could be used. We do three things in this paper: we derive and consider alternative common posterior summary statistics for ranking SNPs in fine-mapping studies and compare their ranking performance with that of the Laplace Bayes factor (LBF); we compare the ranking performance of the LBF and WBF; we compare the credible set sizes and noteworthiness (defined using Bayesian decision theory) of SNPs using the LBF and WBF. We make these comparisons via simulation and via fine-mapping of the CASP8 region using data from the iCOGS array developed by the Collaborative Oncological

Gene-Environment Study (COGS) Consortium (Michailidou et al. [2013]; Michailidou et al. [2017]).

The summaries we consider are based on the posterior distribution of the effect size. They do not have the interpretability of the Bayes factor but could be used if the interest is in ranking a set of SNPs to take forward for further functional analysis. It should be noted that the Bayes factor described here are approximate since they assume a Gaussian likelihood for the observed regression parameters rather than the true logistic likelihood for the binary response.

Rather than focussing solely on capturing the causal SNP, some authors consider whether any SNP highly correlated with the causal SNP ($r^2 > 0.99$ for example) has been identified as a possible causal SNP. In this paper, the true positive rates we report are for the causal SNP only and do not include those SNPs in very high linkage disequilibrium (LD) with it.

Materials and Methods

Summaries of the effect size posterior distribution

We assume that the prior for β (the effect size) is the zero-mean Laplace distribution with rate parameter λ , $\beta \sim La(\lambda)$, and further assume an asymptotic Gaussian distribution for the likelihood, $\hat{\beta} | \beta \sim N(\beta, V)$, as was first proposed in Wakefield [2009]. The Gaussian likelihood approximation is suitable for fine-mapping scenarios of sufficient sample size. Then using the same argument to that used in the reparameterisation proposed by Wakefield [2009] (which is independent of prior parametric form), we can ignore the prior on the intercept and any other confounders when deriving the posterior distribution of the effect size. Bayes theorem gives the posterior distribution on the effect size as

$$f(\beta | \hat{\beta}) = \frac{f(\hat{\beta} | \beta)\pi(\beta)}{\int f(\hat{\beta} | \beta)\pi(\beta) d\beta}. \quad (1)$$

The posterior distribution can be derived from Equation (1) by substituting in the expression for the prior and likelihood. Breaking the denominator integral in Equation (1) into positive and

negatives parts and completing the square gives, after a little algebra, the posterior density as

$$f(\beta | \hat{\beta}) = \begin{cases} \frac{E_-}{\sqrt{2\pi V}} \exp\left(-\frac{1}{2V}(\beta - Q_-)^2\right) & \text{if } \beta < 0 \\ \frac{E_+}{\sqrt{2\pi V}} \exp\left(-\frac{1}{2V}(\beta - Q_+)^2\right) & \text{if } \beta \geq 0 \end{cases} \quad (2)$$

where $Q_- = \hat{\beta} + V\lambda$, $Q_+ = \hat{\beta} - V\lambda$ and E_- and E_+ are normalising constants defined in the Appendix. We consider six statistics as possible ranking statistics: the posterior mean, median, mode, credible intervals (CIs), highest density intervals (HDIs) and Laplace Bayes factor. Derivations of these summary statistics for the Laplace prior are given in the Appendix. Examples of possible posterior probability density functions are shown in Figure 1. With a Laplace prior and Gaussian likelihood, the posterior density has a separate and scaled truncated Gaussian density on the negative and positive effect size support as can be seen in Figure 1 and in Equation (2). It isn't possible for the posterior distribution to be bi-modal in this case. A bi-modal posterior distribution requires that $Q_- < 0$ and $Q_+ > 0$. Both inequalities cannot be satisfied since $Q_+ > 0 \Rightarrow \hat{\beta} > V\lambda \Rightarrow Q_- > 0$. Thus the bi-modal requirement cannot be satisfied. Basic algebra shows that the posterior distribution is continuous at $\beta = 0$: $\lim_{\beta \rightarrow 0^-} = \lim_{\beta \rightarrow 0^+} = \exp(-\hat{\beta}^2/2V)/\sqrt{2\pi V D^2}$ where D is a constant defined in the Appendix. Note that with a Gaussian prior and Gaussian likelihood, the posterior distribution is a single non-truncated Gaussian distribution.

For the posterior mean, median, and mode we used the absolute values of the summary as the ranking statistic. The LBF is strictly non-negative so we used its actual value. For the CIs and HDIs our ranking statistic was the size of the largest probability interval that did not contain $\beta = 0$ (e.g. an 83% or 72% interval). We used receiver operating characteristic (ROC) curves and area under the ROC curves (AUC) to compare the performance of the ranking statistics. Since we are really interested in the true positive rates (TPRs) at low false positive rates (FPRs) we report the partial AUC on FPRs < 0.1 given as an integer percentage. For example, a partial AUC of 70 means that the AUC for FPR < 0.1 was 0.07 (compared to a maximum partial AUC in this region of 0.10). There are several ways to aggregate results over multiple ROC curves. We chose to use vertical averaging [Fawcett, 2006] which calculates the average TPR across a range of FPRs and hence gives an average TPR at each FPR.

Comparing approximate Bayes factors

We also compared the ranking performance of the LBF with that of the WBF in our simulated scenarios. The WBF was derived in Wakefield [2009] and we derive the LBF in the Appendix. The expression for the LBF is not as concise as the WBF but it is just as easily calculated in practice since it has a closed form expression. To compare the ranks, we plotted ROC curves and report partial AUCs as we did when comparing the posterior summaries to the LBF.

Noteworthiness of SNPs

The interest in fine-mapping studies is often not simply ranking SNPs but in trying to identify a set of potentially interesting SNPs to take forward for further investigation. Wakefield [2008] recommended using a Bayesian decision theory approach to determine which SNPs in the fine-mapping region are to be declared noteworthy, in the sense that the posterior probability of the alternative hypothesis exceeded a threshold, determined by the specific costs of making the wrong decision. Noteworthy here is simply taken to mean worthy of further investigation. Wakefield showed how to formalise this by specifying the cost of false non-discovery (C_β) and the costs of false-discovery (C_α). If H_1 is the alternative hypothesis that $\beta \neq 0$, then a SNP is declared noteworthy if

$$\Pr(H_1 \mid \text{data}) = \frac{\phi(1 - \pi_0)}{\pi_0 + \phi(1 - \pi_0)} > \frac{C_\alpha}{C_\alpha + C_\beta} \quad (3)$$

where ϕ is the Bayes factor (here LBF or WBF) for the SNP and π_0 is the prior probability that the SNP is not causally associated. We use the shorthand ‘ratio of costs’ for $r = C_\beta/C_\alpha$. With r defined in this way, Equation (3) is equivalent to $\Pr(H_1 \mid \text{data}) > 1/(1 + r)$. In this paper we set the ratio of costs to be 4, meaning that the cost of false non-discovery is four times that of the cost of false-discovery. We also set π_0 to be 0.99 meaning that we, a priori, expect 99% of SNPs in the fine-mapping region not to be causally associated. Wakefield [2008] suggested that values of r in this region might be appropriate for fine-mapping scenarios where it would be considered relatively costly not to identify causally associated SNPs, relative to the cost of labelling non-causal SNPs as noteworthy. The higher the ratio of costs, the more SNPs will be identified as noteworthy in any given analysis. In his simulations, Wakefield [2008] demonstrated quickly diminishing returns

from increasing the ratio much beyond four. If the ratio of costs is four and the prior probability of the null is 0.99 then a SNP is declared noteworthy if the posterior probability of the alternative hypothesis, H_1 , is greater than $1/5$. With $\pi_0 = 0.99$ and $r = 4$, noteworthiness requires the SNP Bayes factor to exceed 25. Using noteworthiness as the statistic of interest, we consider the TPRs and FPRs derived from using the LBF and WBF in simulated datasets.

Credible set size

Although noteworthiness has the advantage of fitting into a Bayesian decision-theoretic approach, it is not widely used in practice. Typically posterior probabilities of causality (PPC) are calculated for each SNP. Under the assumption that there is exactly one causal SNP in the region the PPC for a given SNP is calculated as the Bayes factor for that SNP divided by the sum of the Bayes factors for all SNPs under consideration [Maller et al., 2012]. These are then ranked and SNPs are added to the set until the cumulative probability exceeds some specified threshold, say 0.90 or 0.95.

To form our credible sets we used 0.90 as the cumulative probability threshold. The values of W and λ used in this study are a subset of the values calculated from GWAS breast cancer top hits data. To allow more widely applicable comparisons of the effect of Gaussian versus Laplace priors we also consider $W = 0.2^2$ which is a more typical value used in many studies. To allow comparison with the Laplace prior we equate the prior variances giving $\lambda = \sqrt{2/W}$, which gives $\lambda = 7.1$. This final pair of prior hyperparameter values are calculated on a different basis to the preceding two hyperparameter pairs (which were based on maximum likelihood estimates obtained using the breast cancer GWAS top hits) but fixing the prior variance or assigning it a distribution is a standard approach in Bayesian variable selection problems [Alenazi et al., 2019].

Simulation details

We used Hapgen2 [Su et al., 2011] to simulate haplotypes in the CASP8 region between base pairs 201666128 and 201866128 of the Hg 19 build of chromosome 2. The reference data we used are the European haplotypes of the August 2010 release of the 1000 Genomes Project. We simulated haplotypes containing a single causal SNP. We converted the haplotype data to genotype data and then calculated the minor allele frequency (MAF) of each SNP identified. Hapgen2 identified 412 SNPs in this region however SNPs with MAF less than 0.01 were removed, which typically left

180 SNPs in the datasets. We simulated 100 data sets for each of eight scenarios. In all scenarios we fixed the odds ratio and MAF of the causal SNP and (by specifying equal numbers of cases and controls) determined the total sample size that gave 60% power to reject the null hypothesis of no association when the causal SNP had an effect size given by the specified odds ratios. The scenarios we considered are given in Table 1. In each scenario we selected a different causal SNP to allow for a variety of LD patterns in which some causal SNPs were in high LD (which we define as $r^2 > 0.8$) with several SNPs and some were not in high LD with any SNPs. Note that we also considered an odds ratio of 1.12 but found the results to be in line with those of the other values of the odds ratios, so omitted them for the sake of brevity.

Based on the 148 available breast cancer GWAS top hits effect sizes [Fachal and Dunning, 2015] and estimates of the number of yet-to-be-discovered (ytbd) SNPs, Walters et al. [2019] showed how they could be used to obtain maximum likelihood estimates of the Laplace and Gaussian prior hyperparameters λ and W . Walters et al. [2019] considered five different numbers of ytbd SNPs, ranging from zero to 1000 (the number estimated in breast cancer by Michailidou et al. [2013] using two independent GWAS data sets). We wanted to assess performance across a range of simulation scenarios and prior parameter settings. To reduce the number of combinations to examine we used only the extreme numbers of ytbd SNPs: zero and 1000. When no ytbd SNPs were assumed, Walters et al. [2019] obtained hyperparameter estimates of $\lambda = 18.3$ and $W = 6.92 \times 10^{-3}$ whereas when 1000 ytbd SNPs were assumed the authors calculated $\lambda = 60.5$ and $W = 1.21 \times 10^{-3}$. We tested all eight of our scenarios at both pairs of hyperparameter estimates. To indicate how a priori likely the odds ratios used in this study are, we calculated the two-tailed probabilities of exceeding the odds ratios used in this study for both the Laplace and Gaussian prior densities for both pairs of hyperparameter estimates. These two-tailed probabilities are presented in Table 2.

iCOGS analysis

We compared the number of noteworthy SNPs identified when using different priors in fine-mapping the CASP8 region between base positions 201,500,074 and 202,569,992 on chromosome 2. The data we used is from a custom-built iCOGS array developed by the Collaborative Oncological Gene-Environment Study (COGS) Consortium [Michailidou et al., 2013]. The genotype data consists of 1733 variants (501 genotyped and 1232 imputed using Impute2 [Marchini and Howie, 2010]) on

46,450 breast cancer cases and 42,600 controls.

We varied the ratio of costs of false-discovery to false non-discovery, r , between 1 and 4. We set the prior probability that a given SNP is causally associated to be $1/3000$. Assuming that the prior probability that a given SNP is causally associated is independent of the other 1732 SNPs in the region, this gives a prior probability of 0, 1, and 2 causally associated SNPs in the region of 0.56 0.32 0.09 respectively. With $\pi_0 = 2999/3000$, noteworthiness requires the SNP Bayes factor to exceed 2999, 1500 and 750 when $r = 1, 2$ and 4 respectively.

Several other authors have fine-mapped the CASP8 region using the iCOGS data (Spencer et al. [2015], Alenazi et al. [2019]). We compare the results of using the Laplace and Wakefield Bayes factors, the power prior Bayes factor of Spencer et al. [2015], the normal-gamma prior of Alenazi et al. [2019] and the normal-gamma prior that incorporates functional information in the prior Alenazi et al. [2019]. The power prior Bayes factor (PPBF) assumes a $N(0, W)$ prior on the effect size but allows for uncertainty in W by putting a power prior on it. For the PPBF analysis we used the exponent hyperparameter $k = -1.66$ as used in Spencer et al. [2015]. The normal-gamma prior (NG) is a scale mixture of normals prior that assumes a $N(0, W)$ prior on the effect size but allows the prior effect size variance, W , to be different for each SNP; the SNP-specific variances are assumed to follow a gamma distribution with specified shape and rate parameters. Alenazi et al. [2019] also considered a normal-gamma prior that incorporates functional information in the prior (NGFS). This approach groups SNPs based on functional significance (FS) scores [Lee and Shatkay, 2009] and allows the gamma rate and shape parameter to be group-specific. To compare the performance of the different methods we report the ranks of the top 20 SNPs for each method.

Results

Comparing ranks using Bayesian summaries

The posterior mean, median and mode were all very similar in terms of partial AUC in all cases. The AUCs for the posterior median and mode were identical (to 2 s.f.) in the eight scenarios and were always at least as large as the AUC for the posterior mean. In addition, the LBF was universally better than the posterior CI and HDI in terms of partial AUC in all scenarios. As a result we only present the results for the posterior median and LBF. Figures 2 and 3 show the ROC

curves for FPRs less than 0.1 when no further ytbd SNPs are assumed and when 1000 ytbd SNPs are assumed respectively. The results indicate that the ranking performance of the LBF is mostly superior to that of the median (and hence the mean and mode) except when the odds ratio is 1.15 and the MAF is 0.09 with no assumed ytbd SNPs, and when the odds ratio is 1.15 or 1.18 and the MAF is 0.09 with 1000 assumed ytbd SNPs. In the second of these three exceptions, the TPR for the median is around 0.2 greater than the TPR for the LBF for FPRs less than 0.05, which is a substantial improvement. In the second exception the partial AUCs are 96 and 83 for the LBF and posterior median respectively. So it appears that if the aim of a fine-mapping study is simply to rank SNPs then the posterior median is worth considering along with the LBF.

Comparing ranks using Bayes Factor

Figures 4 and 5 show the ROC curves for FPRs less than 0.1 when no further ytbd SNPs are assumed and when 1000 are assumed respectively. When comparing the two Bayes factors we observe a similar pattern as for the median versus LBF ROC curves. Generally the LBF has a higher partial AUC than the WBF except for the same three cases that showed a higher partial AUC for the median compared to the LBF, although the advantage of the WBF over the LBF in these three cases is less than when comparing the median to the LBF. These plots clearly demonstrate that there could be occasionally meaningful differences in the ranks of causal SNPs in Bayesian fine-mapping studies when using the LBF compared to the WBF, but that generally ranking performance using these Bayes factors might be expected to be rather similar.

Comparing noteworthiness using different priors

Figures 6 and 7 show box plots of the FPR (for both LBF and WBF) using noteworthiness of the SNP as the classifier for the eight scenarios in Table 1, with 0 and 1000 ytbd SNPs respectively. The values of the TPRs are also shown on the box plots. The comparison of the noteworthiness of the SNPs reveals substantial differences when modelling the effect sizes with a zero-centred Laplace prior compared to using a zero-centred Gaussian prior (LBF versus WBF). The ROC curves in Figures 4 and 5 show that there were some differences between the ranks of LBF and WBF. Essentially Figures 4 and 5 compare average ranks of the causal SNP Bayes factors for a given scenario. However the noteworthiness takes into account the actual size of the Bayes factor,

so that a causal SNP could have the same rank when ranked using LBF or WBF but be clearly noteworthy using LBF but not using WBF (because the LBF for the causal SNP is substantially higher than the WBF). Figures 6 and 7 indicate that, when using the WBF, the causal SNP is rarely identified as being noteworthy in any of the eight scenarios with either 0 or 1000 ytbD SNPs. The largest TPR for the WBF in any of the sixteen plots is 0.3 whilst it is 0.91 for the LBF. There is a concomitant increase in the FPR for the LBF but the values observed are generally at an acceptable level in most scenarios. The median noteworthiness FPR for the LBF is 0.05 or less in 14 of the 16 scenarios in Figures 6 and 7. Therefore, in these simulations, using the LBF declares the true causal SNP to be noteworthy with high probability whilst generally declaring no more than 5% of the non-causal SNPs to be noteworthy. This is particularly noticeable in Figure 6 where the TPR is 0.78 or more in 7 of the eight scenarios. Using the WBF declares very few SNPs to be noteworthy when the prior probability of causal association is 0.01 and the ratio of costs is 4. For the WBF the TPR is 0.15 or less in 13 out of the 16 scenarios in Figures 6 and 7.

Rather than fixing the power, it is often of interest in simulation studies to increase the odds ratio of the causal SNP (and hence the power) for a fixed sample size. We chose to consider the three sample sizes and MAFs from scenarios 2, 3 and 5 in Table 1. We then chose the odds ratios to give 40, 60 and 90% power. We essentially have versions of Figures 2 - 7 where the sample size is fixed in each of three rows and the odds ratio varies in each of three columns. The results are presented in Supplementary File 1.

Comparing credible set size using different priors

Table 3 shows the median 90% credible set sizes for the eight scenarios in Table 1 for three pairs of hyperparameter values. Table 3 also shows the percentage of 90% credible sets that contained the simulated causal SNP and the median posterior probability of causality (PPC) of the causal SNPs across all 100 datasets.

The results for all three pairs of hyperparameter values exhibit similar patterns. Using the Laplace prior leads to larger credible sets. The credible set size can be a lot larger for the Laplace prior than for the Gaussian prior, particularly when the effect size prior variance is extremely small. The causal SNP PPC are generally higher with a Gaussian prior although this is much less pronounced when the effect size prior variance is large. As a result of larger sizes, the credible sets

formed using the Laplace prior are always at least as likely to contain the causal SNP compared to the sets based on the Gaussian prior. Occasionally this difference is reasonably large, for example in Scenario 5 when $\lambda = 18.3$ the credible set size is 24 for the Laplace prior versus 18 for the Gaussian prior but 94 versus 87 sets contain the causal SNP for the Laplace and Gaussian priors respectively.

Noteworthiness and SNP ranks in the iCOGS data

We assessed the effect of using Laplace and Gaussian priors when fine-mapping the CASP 8 region. We assumed either 0 or 1000 ytbcd SNPs and varied the ratio of costs of false non-discovery to false-discovery between $r = 1$ and 4. Table 4 shows the number of noteworthy CASP8 SNPs when the prior probability that a given SNP is causally associated is $1/3000$.

The results for the iCOGS data show a similar pattern to the results of the noteworthiness of SNPs in the simulated data: the Laplace prior identifies a much larger set of noteworthy SNPs. In some scenarios, the analysis using the Gaussian effect size prior doesn't identify any SNPs as being noteworthy, whilst the Laplace prior for the same parameter settings yields between 7 and 9 noteworthy SNPs. This reiterates the message from the analysis of the simulated data, that the choice of prior parametric form potentially has significant implications on the set of SNPs declared noteworthy in univariate fine-mapping. It should be noted that the SNP identified as noteworthy using the Gaussian prior is also identified as being noteworthy in all analyses using the Laplace effect size prior. The rs numbers of the SNPs identified in Table 4 are given in the second supplementary file.

We also considered the similarity of the ranks of the SNPs using the posterior median, the LBF and WBF. When there were assumed to be 1000 ytbcd SNPs there was good agreement between the top 10 ranked SNPs using the posterior median and the top ranking SNPs using WBF and LBF. The top ten SNPs using the median were all in the top 18 SNPs using either LBF or WBF. When there were assumed to be no ytbcd SNPs the disparity increased. There was still generally good agreement between the top 10 SNPs using the posterior median and the two Bayes factor but there were a couple of exceptions. In particular, the top ranking SNP by posterior median was ranked 107th using LBF and 33rd using WBF.

Table 5 reports the SNP ranks for the Laplace and Wakefield Bayes factors with 0 and 1000

ytbd SNPs and the ranks for the PPBF, the NG and the NGFS. Ranks 1 to 9 using each of the Bayes factors are identical. The SNPs in the top 20 are mostly the same for all the Bayes factor methods. The NG and NGFS rank very different SNPs in the top 20 compared to the Bayes factor methods. SNP 33 which is ranked top by all the Bayes factor methods is in the top 8 of both the NG and NGFS but otherwise there is little agreement between the Bayes factor methods and the shrinkage-based methods that try to shrink non-causal SNP effect sizes more than causal SNP effect sizes. Alenazi et al. [2019] also used FINEMAP [Benner et al., 2016] to finemap the CASP8 region in the iCOGS data. The only SNP identified in the top 20 by NG, NGFS and FINEMAP was SNP 31 (rs2540050) which was the top ranking SNP in all the methods using Bayes factors. The consensus is that the top ranking SNP is SNP31. The difference between the posterior probability of causal association when using the Laplace Bayes factor compared to the Wakefield Bayes factor is stark. Assuming, as before, that the prior probability that a given SNP is causally associated is $1/3000$ and that there are 1000 ytbd SNPs, the posterior probability of causal association using the Gaussian prior is 0.23 whereas it is 0.98 using the Laplace prior. Although the ranks are the same, the evidence is considerably stronger using the Laplace prior.

Discussion

The results in this paper indicate that the parametric form of the prior could potentially have a considerable effect on the choice of SNPs selected for further consideration across a wide range of fine-mapping settings with a single causal SNP if noteworthiness is used as SNP selection method. The LBF is rarely surpassed by any posterior summaries in terms of causal SNP ranking and there seems to be relatively small difference in ranking performance between the WBF and the LBF. The clear difference in this study was found to be in the differences in size of the LBF and WBF which manifested itself as dramatic differences in the TPRs using the LBF compared to the WBF when noteworthiness is used as SNP selection method. The WBF rarely identified the causal SNP (or indeed any non-causal SNPs) as being noteworthy. The credible sets formed using the posterior probability of causality indicate the Laplace prior leads to larger credible sets that are generally more likely to contain the causal SNP than the sets formed using the Gaussian prior. The larger set sizes may be due to the higher Bayes factors for the non-causal SNPs that have estimated

odds ratios quite close to zero. The Laplace prior has a high probability density near to zero and so the Bayes factors for these SNPs will be considerably larger than those calculated under the Gaussian prior. This will tend to make the posterior probabilities of causality more uniform under the Laplace prior and so lead to larger credible set sizes.

The differences observed in this paper would seem to be purely related to the ratio of the probability densities of the Laplace versus the Gaussian Prior. The Bayes factor can be thought of as a sum of weighted likelihoods, where the weights are the prior probability densities. Thus it is to be expected that alternative effect size priors could also lead to Bayes factor showing strong evidence of association. There has been relatively little attention paid to the choice of prior in fine-mapping setting, with the Gaussian effect size prior almostly universally being the prior of choice. This seems to be an oversight that could potentially reduce the effectiveness of Bayesian fine-mapping in identifying causally associated disease SNPs. The results in this paper are applicable to univariate analyses in fine-mapping regions with a single causal SNP which is likely to be a common scenario but is clearly not true of all fine-mapping regions. It isn't immediately obvious how applicable the results in this paper are to regions with multiple causal SNPs however given that the results really relate to the relative prior densities of the competing priors, there seems little reason to think that they are not at least indicative of the results that might be obtained in a multi-SNP analysis.

The results presented here will of course be affected by the underlying LD structure in the genomic region simulated. The region containing the CASP8 region is known to be a region with moderate LD [Spencer et al., 2014]. We selected different causal SNPs with a range of LD patterns as described in Table 1. Some causal SNPs were in strong LD with several or a large number of SNPs, others were not in strong LD with any other SNPs in the region. In univariate analyses the effect of having multiple SNPs in strong LD with the causal SNP will depend on the sample size, MAF, causal SNP effect size and the strength of the LD. If V for the causal SNP is low (large sample size, common causal SNP) it is likely, unless the prior is really inappropriate for the true effect size, that the posterior probability of causal association will be reasonably high for all SNPs in LD with the causal SNP. This may mean that the 90% credible sets are small but it may mean that they don't contain the causal SNP (especially if the LD with the causal SNP is really strong). If V is large (small sample size and/or low MAF causal SNP) then the posterior probabilities are likely to be small for all the SNPs in LD with the causal SNP and the credible sets are likely to be

larger but the concomitant chance that it contains the causal SNP is likely to be higher.

Deciding on the choice of prior is not straightforward. Walters et al. [2019] showed how this could be done using the posterior probability of the prior form where GWAS top hits are used as the data in a Bayesian analysis. This could conceivably be extended relatively easily to take into account other parametric prior forms beyond the Laplace and Gaussian however the choice of priors that give closed-form expressions for the Bayes factors is relatively limited. Deriving Monte Carlo estimates of the Bayes factors in univariate fine-mapping is, however, trivial so this should not be a barrier to considering a wider range of prior forms.

The data used to inform the effect size prior distribution are likely to be GWAS top hits and possibly some functionally verified SNPs. It may also be augmented by an estimate of the number of ytbds disease-specific SNPs. Walters et al. [2019] showed how to use GWAS top hits and estimates of the number of ytbds breast cancer SNPs to calculate maximum-likelihood estimates of the prior hyperparameters. The breast cancer data the authors used included 148 SNPs and there are now an additional 65 breast cancer-associated SNPs [Michailidou et al., 2017]. There may be many complex diseases which are less well studied and so have fewer GWAS top hits. In such diseases the estimates of the number of ytbds SNPS will therefore be important as they will carry relatively more information compared to diseases with many GWAS top hits. In breast cancer the number of ytbds SNPs was estimated to be around 1000 [Michailidou et al., 2013] but no idea of the uncertainty on this estimate was provided. Further work is needed to enable some form of uncertainty to be associated with this point estimate, especially if these estimates are to be used to inform effect size priors in diseases with relatively few GWAS top hit.

Data availability statement

The data needed to produce the results are hosted on zenodo with DOI: 10.5281/zenodo.4147536

R code

The R code needed to produce the results are hosted on zenodo with DOI:10.5281/zenodo.4147481

Acknowledgements

Some of this was carried out as part of a PhD funded by the Ministry of Education Malaysia (Higher Education) and University of Malaya. We would like to thank the two anonymous referees who made valuable suggestions that substantively improved the manuscript.

References

- Alenazi, A. A., Cox, A., Juarez, M., Lin, W.-Y., and Walters, K. (2019). Bayesian variable selection using partially observed categorical prior information in fine mapping association studies. *Genetic epidemiology*, 43(6):690–703.
- Benner, C., Spencer, C. C., Havulinna, A. S., Salomaa, V., Ripatti, S., and Pirinen, M. (2016). Finemap: Efficient variable selection using summary data from genome-wide association studies. *Bioinformatics*, 32(10):1493–1501.
- Boggis, E., Milo, M., and Walters, K. (2016). equips: eqtl analysis using informed partitioning of snps—a fully bayesian approach. *Genet Epidemiol*, 40(4):273–283.
- Chen, W., Larrabee, B. R., Ovsyannikova, I. G., Kennedy, R. B., Haralambieva, I. H., Poland, G. A., and Schaid, D. J. (2015). Fine mapping causal variants with an approximate bayesian method using marginal test statistics. *Genetics*, 200(3):719–736.
- Fachal, L. and Dunning, A. M. (2015). From candidate gene studies to gwas and post-gwas analyses in breast cancer. *Curr Opin Genet Dev*, 30:32–41.
- Fawcett, T. (2006). An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874.
- Hoggart, C. J., Whittaker, J. C., De Iorio, M., and Balding, D. J. (2008). Simultaneous analysis of all snps in genome-wide and re-sequencing association studies. *PLoS genetics*, 4(7):e1000130.
- Lee, P. H. and Shatkay, H. (2009). An integrative scoring system for ranking snps by their potential deleterious effects. *Bioinformatics*, 25(8):1048–1055.
- Marchini, J. and Howie, B. (2010). Genotype imputation for genome-wide association studies. *Nature Reviews Genetics*, 11(7):499–511.

- Maller, J. B., McVean, G., Byrnes, J., Vukcevic, D., Palin, K., Su, Z., ... Donnelly, P. (2012). Bayesian Refinement of Association Signals for 14 Loci in 3 Common Diseases. *Nature Genet*, 44(12):1294–1301.
- Michailidou, K., Hall, P., Gonzalez-Neira, A., Ghoussaini, M., Dennis, J., Milne, R. L., ... Easton, D. F. (2013). Large-scale genotyping identifies 41 new loci associated with breast cancer risk. *Nat Genet*, 45(4):353–361.
- Michailidou, K., Lindstrom, S., Dennis, J., Beesley, J., Hui, S., Kar, S., ... Rostamianfar, A. (2017). Association analysis identifies 65 new breast cancer risk loci. *Nature*, 551(7678):92–94.
- Spencer, A. V., Cox, A., Lin, W.-Y., Easton, D. F., Michailidou, K., and Walters, K. (2015). Novel bayes factors that capture expert uncertainty in prior density specification in genetic association studies. *Genet Epidemiol*, 39(4):239–248.
- Spencer, A. V., Cox, A., Lin, W.-Y., Easton, D. F., Michailidou, K., and Walters, K. (2016). Incorporating functional genomic information in genetic association studies using an empirical bayes approach. *Genet Epidemiol*, 40(3):176–187.
- Spencer, A. V., Cox, A., and Walters, K. (2014). Comparing the efficacy of snp filtering methods for identifying a single causal snp in a known association region. *Ann Hum Genet*, 78(1):50–61.
- Su, Z., Marchini, J., and Donnelly, P. (2011). Hapgen2: simulation of multiple disease snps. *Bioinformatics*, 27(16):2304–2305.
- Van de Bunt, M., Cortes, A., Brown, M. A., Morris, A. P., and McCarthy, M. I. (2015). Evaluating the performance of fine-mapping strategies at common variant gwas loci. *Plos Genet*, 11(9):e1005535.
- Vsevolozhskaya, O. A. and Zaykin, D. V. (2019). Quantifying posterior effect size distribution of susceptibility loci by common summary statistics. *Genet Epidemiol*. 10.1002/gepi.22286.
- Wakefield, J. (2008). Reporting and interpretation in genome-wide association studies. *Int J Epidemiol*, 37(3):641–653.

- Wakefield, J. (2009). Bayes factors for genome-wide association studies: comparison with p-values. *Genet Epidemiol*, 33(1):79–86.
- Walters, K., Cox, A., and Yaacob, H. (2019). Using gwas top hits to inform priors in bayesian fine-mapping association studies. *Genet Epidemiol*, 43(6):675–689.

Tables

scenario	odds ratio	MAF	total sample size	number of SNPs in high LD with causal SNP ($r^2 > 0.8$)
1	1.05	0.3	65000	0
2	1.05	0.09	167000	1
3	1.08	0.3	26000	1
4	1.08	0.09	67000	3
5	1.15	0.3	7900	9
6	1.15	0.09	20300	0
7	1.18	0.3	5660	1
8	1.18	0.09	14500	1

Table 1. Simulation scenarios used in comparing the various Bayesian summary statistics in a univariate Bayesian fine-mapping analysis.

odds ratio	0 ytdb SNPs		1000 ytdb SNPs	
	$\lambda = 18.3$	$W = 6.92 \times 10^{-3}$	$\lambda = 60.5$	$W = 1.21 \times 10^{-3}$
1.05	0.409	0.558	0.052	0.161
1.08	0.245	0.355	0.010	0.027
1.15	0.077	0.093	2.13×10^{-4}	5.87×10^{-5}
1.18	0.048	0.047	4.48×10^{-5}	1.95×10^{-6}

Table 2. Two-tailed probabilities of exceeding the given odds ratios for both the Laplace and Gaussian prior densities with given hyperparameter estimates.

scenario	median 90% credible set size	percentage of 90% sets containing causal SNP	median PPC of causal SNP	median 90% credible set size	percentage of 90% sets containing causal SNP	median PPC of causal SNP
Laplace ($\lambda = 60.5$)			Gaussian ($W = 1.21 \times 10^{-3}$)			
1	108	99	0.4488	20	96	0.7032
2	88	92	0.0686	24	88	0.1005
3	90	99	0.0786	20	99	0.1051
4	85	98	0.1232	11	94	0.1812
5	141	99	0.0395	18	87	0.1236
6	123	97	0.0669	17	96	0.5732
7	150	100	0.0290	20	98	0.1376
8	155	99	0.0175	23	94	0.0940
Laplace ($\lambda = 18.3$)			Gaussian ($W = 6.92 \times 10^{-3}$)			
1	35	96	0.6296	20	96	0.7032
2	28	89	0.0971	24	88	0.1005
3	24	99	0.1057	20	99	0.1051
4	16	96	0.1751	11	94	0.1811
5	24	94	0.0965	18	87	0.1236
6	38	97	0.3632	17	96	0.5731
7	40	99	0.1037	20	98	0.1376
8	83	99	0.0763	23	94	0.0940
Laplace ($\lambda = 7.1$)			Gaussian ($W = 4 \times 10^{-2}$)			
1	25	96	0.6761	20	96	0.7032
2	24	88	0.1018	24	88	0.1005
3	20	99	0.1069	20	99	0.1051
4	13	95	0.1758	11	94	0.1811
5	18	88	0.1156	18	87	0.1235
6	23	96	0.5188	17	96	0.5728
7	20	99	0.1274	20	98	0.1375
8	31	97	0.0932	23	94	0.0940

Table 3. 90% credible set sizes for the eight scenarios in Table 1 and for three combinations of λ and W values. The three combinations correspond to 1000 ytbD SNPs, 0 ytbD SNPs and the value of $W = 0.04$ commonly used in fine mapping studies (with λ calculated by equating prior effect size variances). The table gives the median 90% credible set size; the percentage of 90% credible sets that contained the simulated causal SNP; the median posterior probability of causality (PPC) of the causal SNPs across all 100 data sets.

number of ytbd SNPs	effect size prior	ratio of costs (r)		
		1	2	4
0	La(λ)	9	9	9
0	N(0, W)	1	1	1
1000	La(λ)	7	9	9
1000	N(0, W)	0	0	1

Table 4. The number of noteworthy SNPs in the iCOGS data using both a Laplace and Gaussian effect size prior for three different ratio of costs (1,2, and 4). The prior probability that a given SNP is causally associated is assumed to be 1/3000.

Rank	LBF0	WBF0	LBF1000	WBF1000	PPBF	NG	NGFS
1	31	31	31	31	31	765	765
2	2	2	2	2	2	342	342
3	1	1	1	1	1	589	589
4	3	3	3	3	3	1101	1177
5	16	16	16	16	16	1177	31
6	24	24	24	24	24	1244	795
7	7	7	7	7	7	1245	33
8	29	29	29	29	29	31	1274
9	27	27	27	27	27	795	897
10	10	602	10	602	602	33	1253
11	602	10	14	10	31	1253	1165
12	23	23	23	23	2	1274	1639
13	14	14	6	14	1	897	1146
14	6	6	9	6	3	1639	1701
15	9	9	8	9	16	916	1101
16	8	8	15	8	24	1701	740
17	15	15	4	15	7	740	920
18	4	1639	602	1639	29	1056	916
19	1639	4	1639	4	27	932	1671
20	681	681	1096	681	602	1671	934

Table 5. The top 20 ranked SNPs by method in the iCOGS data. LBF0 and WBF0 are the Laplace and Wakefield Bayes factors assuming 0 ytbd SNPs; LBF1000 and WBF10000 are the Laplace and Wakefield Bayes factors assuming 1000 ytbd SNPs; PPBF is the power prior Bayes factor of Spencer et al. [2015]; NG is the normal-gamma prior of Alenazi et al. [2019]; NGFS is the normal-gamma prior of Alenazi et al. [2019] that incorporates functional genomic scores into the effect size prior.

Figure Captions

Figure 1: Examples of possible posterior probability density functions with a Laplace(λ) prior on β and a $N(\beta, V)$ likelihood. In both plots $\lambda = 60.5$ and the posterior density depends on the estimated effect size, $\hat{\beta}$, and the variance of the effect size estimate, V .

Figure 2: ROC curves comparing the ranking performance of the posterior median and Bayes factor when the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 18.3$ and the likelihood is an asymptotic Gaussian. This case corresponds to assuming there are no ytdb SNPs. For each ranking statistic the curve represents the average TPR for a given FPR, averaged over 100 data sets, simulated from Hapgen2. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively. The sample sizes are calculated to yield 60% power and are given in Table 1.

Figure 3: ROC curves comparing the ranking performance of the posterior median and Bayes factor when the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 60.5$ and the likelihood is an asymptotic Gaussian. This case corresponds to assuming there are 1000 ytdb SNPs. For each ranking statistic the curve represents the average TPR for a given FPR, averaged over 100 data sets, simulated from Hapgen2. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively. The sample sizes are calculated to yield 60% power and are given in Table 1.

Figure 4: ROC curves comparing the ranking performance of the Laplace and Wakefield Bayes factor. For the Laplace Bayes factor the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 18.3$ whilst for the Wakefield Bayes factor the prior distribution is the zero-centred Gaussian distribution with variance $W = 692 \times 10^{-5}$. This case corresponds to assuming there are no ytdb SNPs. In both cases the likelihood is an asymptotic Gaussian. For each ranking statistic the curve represents the average TPR for a given FPR, averaged over 100 data sets, simulated from Hapgen2. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively. The sample sizes are calculated

to yield 60% power and are given in Table 1.

Figure 5: ROC curves comparing the ranking performance of the Laplace and Wakefield Bayes factor. For the Laplace Bayes factor the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 60.5$ whilst for the Wakefield Bayes factor the prior distribution is the zero-centred Gaussian distribution with variance $W = 121 \times 10^{-5}$. This case corresponds to assuming there are 1000 ytdb SNPs. In both cases the likelihood is an asymptotic Gaussian. For each ranking statistic the curve represents the average TPR for a given FPR, averaged over 100 data sets, simulated from Hapgen2. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively. The sample sizes are calculated to yield 60% power and are given in Table 1.

Figure 6: Boxplots comparing the distribution of the false positive rate, for 100 data sets simulated from Hapgen2,(FPR) of the noteworthiness of SNPs for the Laplace and Wakefield Bayes factor. For the Laplace Bayes factor the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 18.3$ whilst for the Wakefield Bayes factor the prior distribution is the zero-centred Gaussian distribution with variance $W = 692 \times 10^{-5}$. This case corresponds to assuming there are no ytdb SNPs. In both cases the likelihood is an asymptotic Gaussian. The ratio of costs of false non-discovery to false-discovery is 4. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively. The sample sizes are calculated to yield 60% power and are given in Table 1.

Figure 7: Boxplots comparing the distribution of the false positive rate, for 100 data sets simulated from Hapgen2,(FPR) of the noteworthiness of SNPs for the Laplace and Wakefield Bayes factor. For the Laplace Bayes factor the prior distribution is the zero-centred Laplace distribution with parameter $\lambda = 60.5$ whilst for the Wakefield Bayes factor the prior distribution is the zero-centred Gaussian distribution with variance $W = 121 \times 10^{-5}$. This case corresponds to assuming there are 1000 ytdb SNPs. In both cases the likelihood is an asymptotic Gaussian. The ratio of costs of false non-discovery to false-discovery is 4. The simulated odds ratios in rows 1 to 4 are 1.05, 1.08, 1.15 and 1.18 respectively. The causal SNP MAFs in columns 1 and 2 are 0.3 and 0.09 respectively.

The sample sizes are calculated to yield 60% power and are given in Table 1.

Appendix

Posterior expected value

The posterior expected value is given by

$$\mathbb{E}(\beta | \hat{\beta}) = \int_{-\infty}^0 \beta f(\beta | \hat{\beta}) \, d\beta + \int_0^{\infty} \beta f(\beta | \hat{\beta}) \, d\beta \quad (4)$$

where $f(\beta | \hat{\beta})$ is defined in Equation (2) in the main text. Basic integration gives

$$\mathbb{E}(\beta | \hat{\beta}) = E_+ \left[\sqrt{\frac{V}{2\pi}} \left[\exp\left(-\frac{Q_+^2}{2V}\right) \right] + Q_+ \left[1 - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right] \right] - \quad (5)$$

$$E_- \left[\sqrt{\frac{V}{2\pi}} \left[\exp\left(-\frac{Q_-^2}{2V}\right) \right] - Q_- \left[\Phi\left(\frac{-Q_-}{\sqrt{V}}\right) \right] \right] \quad (6)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard Gaussian distribution, Q_- and Q_+ are defined in the main text and

$$E_- = \frac{1}{D} \exp\left(-\frac{1}{2V}(\hat{\beta}^2 - Q_-^2)\right); \quad (7)$$

$$E_+ = \frac{1}{D} \exp\left(-\frac{1}{2V}(\hat{\beta}^2 - Q_+^2)\right); \quad (8)$$

$$D = \exp\left(-\frac{1}{2V}(\hat{\beta}^2 - Q_-^2)\right) \left[\Phi\left(\frac{-Q_-}{\sqrt{V}}\right) \right] + \exp\left(-\frac{1}{2V}(\hat{\beta}^2 - Q_+^2)\right) \left[1 - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right]. \quad (9)$$

Posterior median

We have to take into consideration whether the median (m) falls in the negative or positive support of the posterior distribution. Using the definition of the median applied to Equation (2) we see that the median is positive if $2E_+ \left[1 - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right] > 1$ and is negative otherwise. The value of the

median is

$$m = \begin{cases} \Phi^{-1}\left(1 - \frac{1}{2E_+}\right)\sqrt{V} + Q_+ & \text{if } m > 0 \\ \Phi^{-1}\left(\frac{1}{2E_-}\right)\sqrt{V} + Q_- & \text{if } m \leq 0 \end{cases} \quad (10)$$

Posterior mode

The value of the posterior mode, $\arg \max_{\beta} f(\beta | \hat{\beta})$, depends on the signs of the modes of the two underlying Gaussian distributions defining the posterior distribution. The mode is

$$\arg \max_{\beta} f(\beta | \hat{\beta}) = \begin{cases} Q_- & \text{if } Q_+ < Q_- \leq 0 \\ Q_+ & \text{if } 0 \leq Q_+ < Q_- \\ 0 & \text{if } Q_+ \leq 0 \text{ and } Q_- \geq 0 \end{cases} \quad (11)$$

Posterior credible interval

A Bayesian credible interval has equal tail probabilities so the $100(1 - \alpha)\%$ credible interval for β , $[\beta_L, \beta_U]$ is obtained by solving

$$\int_{-\infty}^{\beta_L} f(\beta | \hat{\beta}) \, d\beta = \int_{\beta_U}^{\infty} f(\beta | \hat{\beta}) \, d\beta = \frac{\alpha}{2}$$

There are three cases to consider according to whether β_L and β_U are both positive, both negative or of different sign. We just have to integrate the relevant part of the posterior density in Equation (2). For example, consider the case where $\beta_L < 0$ and $\beta_U > 0$ which is guaranteed to be true when both $E_- \left[\Phi\left(\frac{-Q_-}{\sqrt{V}}\right) \right]$ and $E_+ \left[1 - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right]$ are greater than $\alpha/2$. In this case we equate the integral of the posterior density on the negative reals between $-\infty$ and β_L to $\alpha/2$ and solve. This gives

$$\beta_L = \Phi^{-1}\left(\frac{\alpha}{2E_-}\right)\sqrt{V} + Q_- \quad (12)$$

We use a similar approach to find β_U and repeat for the other two cases. β_L and β_U for all three cases are given in Table 6.

$\text{sign}(\beta_L)$	$\text{sign}(\beta_U)$	Lower limit (β_L)	Upper limit (β_U)
< 0	> 0	$\Phi^{-1}\left(\frac{\alpha}{2E_-}\right)\sqrt{V} + Q_-$	$\Phi^{-1}\left(\frac{2E_+ - \alpha}{2E_+}\right)\sqrt{V} + Q_+$
< 0	< 0	$\Phi^{-1}\left(\frac{\alpha}{2E_-}\right)\sqrt{V} + Q_-$	$\Phi^{-1}\left(\frac{2-\alpha}{2E_-}\right)\sqrt{V} + Q_-$
> 0	> 0	$\Phi^{-1}\left(\frac{2E_+ - 2 + \alpha}{2E_+}\right)\sqrt{V} + Q_+$	$\Phi^{-1}\left(\frac{2E_+ - \alpha}{2E_+}\right)\sqrt{V} + Q_+$

Table 6. Upper and lower 95% posterior credible interval limits for the log odds ratio.

Highest density posterior interval

The Highest density posterior interval (HDI) does not have an equal tail probabilities but instead has equal probability density at the upper and lower limits. A $100(1 - \alpha)\%$ HDI therefore satisfies

$$\int_{\beta_L}^{\beta_U} f(\beta|\hat{\beta}) d\beta = 1 - \alpha \quad \text{and} \quad f(\beta_U|\hat{\beta}) = f(\beta_L|\hat{\beta})$$

As for the credible intervals, the limits depend on whether the signs of β_L and β_U are the same or different.

Case 1: We first consider the case where $\beta_L < 0$ and $\beta_U \leq 0$. Clearly in this case $Q_- < 0$ since $Q_- = (\beta_L + \beta_U)/2$. Using this, and the result that $\Phi(x) + \Phi(-x) = 1 \forall x \in \mathbb{R}$, we have that

$$\begin{aligned} \int_{\beta_L}^{\beta_U} f(\beta|\hat{\beta}) d\beta &= 1 - \alpha \\ E_- \left[\Phi\left(\frac{\beta_U - Q_-}{\sqrt{V}}\right) - \Phi\left(\frac{\beta_L - Q_-}{\sqrt{V}}\right) \right] &= 1 - \alpha \\ E_- \left[2\Phi\left(\frac{-\beta_L + Q_-}{\sqrt{V}}\right) - 1 \right] &= 1 - \alpha \end{aligned}$$

which gives

$$\beta_L = Q_- - \sqrt{V} \Phi^{-1}\left(\frac{1 - \alpha + E_-}{2E_-}\right). \quad (13)$$

and $\beta_U = 2Q_- - \beta_L$.

Case 2: If $\beta_L \geq 0$ and $\beta_U > 0$ (and hence $Q_+ > 0$). The solution is derived in an identical manner to that of Case 1 so the solutions are the same as in Equation(13) but E_- is replaced with E_+ and Q_- is replaced with Q_+ .

Case 3: The final case is where $\beta_L < 0$ and $\beta_U > 0$. We first need to find the relationship between β_L and β_U . Equating the posterior probability densities at $\beta_L (< 0)$ and $\beta_U (> 0)$ gives

$$\frac{E_-}{\sqrt{2\pi V}} \exp\left(-\frac{1}{2V}(\beta_L - Q_-)^2\right) = \frac{E_+}{\sqrt{2\pi V}} \exp\left(-\frac{1}{2V}(\beta_U - Q_+)^2\right) \quad (14)$$

Substituting our previously derived expressions for E_- and E_+ into Equation (14), we get, after a little algebra $\beta_L^2 - 2Q_- \beta_L - (\beta_U^2 - 2Q_+ \beta_U) = 0$. Since $\beta_L < Q_-$ it follows that the only solution to this quadratic is the smaller one and therefore

$$\beta_L = Q_- - \sqrt{Q_-^2 + \beta_U^2 - 2Q_+ \beta_U}. \quad (15)$$

For a $100(1 - \alpha)\%$ HDI for β we have

$$\int_{\beta_L}^0 f(\beta_- | \hat{\beta}) d\beta + \int_0^{\beta_U} f(\beta_+ | \hat{\beta}) d\beta = 1 - \alpha$$

$$E_- \left[\Phi\left(\frac{-Q_-}{\sqrt{V}}\right) - \Phi\left(\frac{\beta_L - Q_-}{\sqrt{V}}\right) \right] + E_+ \left[\Phi\left(\frac{\beta_U - Q_+}{\sqrt{V}}\right) - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right] = 1 - \alpha. \quad (16)$$

Substituting Equation(15) into Equation (16) gives

$$E_- \left[\Phi\left(\frac{-Q_-}{\sqrt{V}}\right) - \Phi\left(-\sqrt{\frac{Q_-^2 + \beta_U^2 - 2Q_+ \beta_U}{V}}\right) \right] + E_+ \left[\Phi\left(\frac{\beta_U - Q_+}{\sqrt{V}}\right) - \Phi\left(\frac{-Q_+}{\sqrt{V}}\right) \right] = 1 - \alpha \quad (17)$$

Equation (17) cannot be solved analytically for β_U so we use the `uniroot` function in the R 'stats' package. Having found β_U , we then solve for β_L using Equation (15).

Laplace Bayes factor

The Laplace Bayes factor is given by

$$\text{Bayes factor} = \frac{f(\text{data} | H_1)}{f(\text{data} | H_0)} = \frac{\int f(\hat{\beta} | \beta) \pi(\beta) d\beta}{f(\hat{\beta} | \beta = 0)}, \quad (18)$$

which, after a little algebra, can be shown to be

$$\lambda \sqrt{\frac{\pi V}{2}} \left[\exp\left(\frac{Q_+^2}{2V}\right) \Phi\left(\frac{Q_+}{\sqrt{V}}\right) + \exp\left(\frac{Q_-^2}{2V}\right) \Phi\left(\frac{-Q_-}{\sqrt{V}}\right) \right] \quad (19)$$

Figures

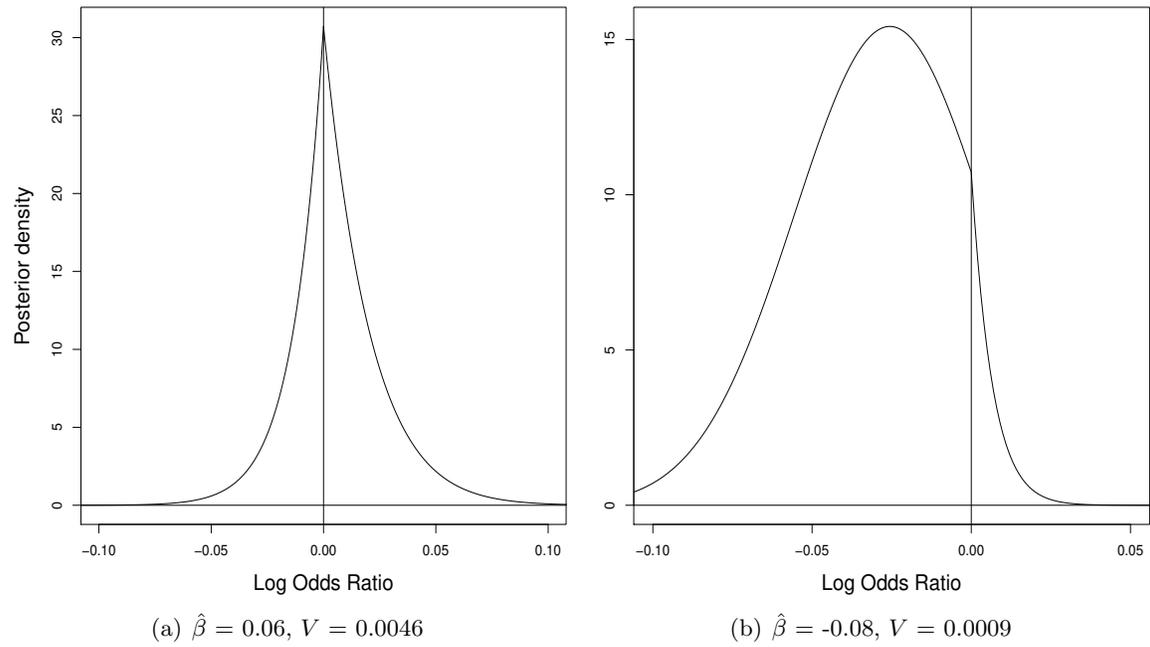


Figure 1

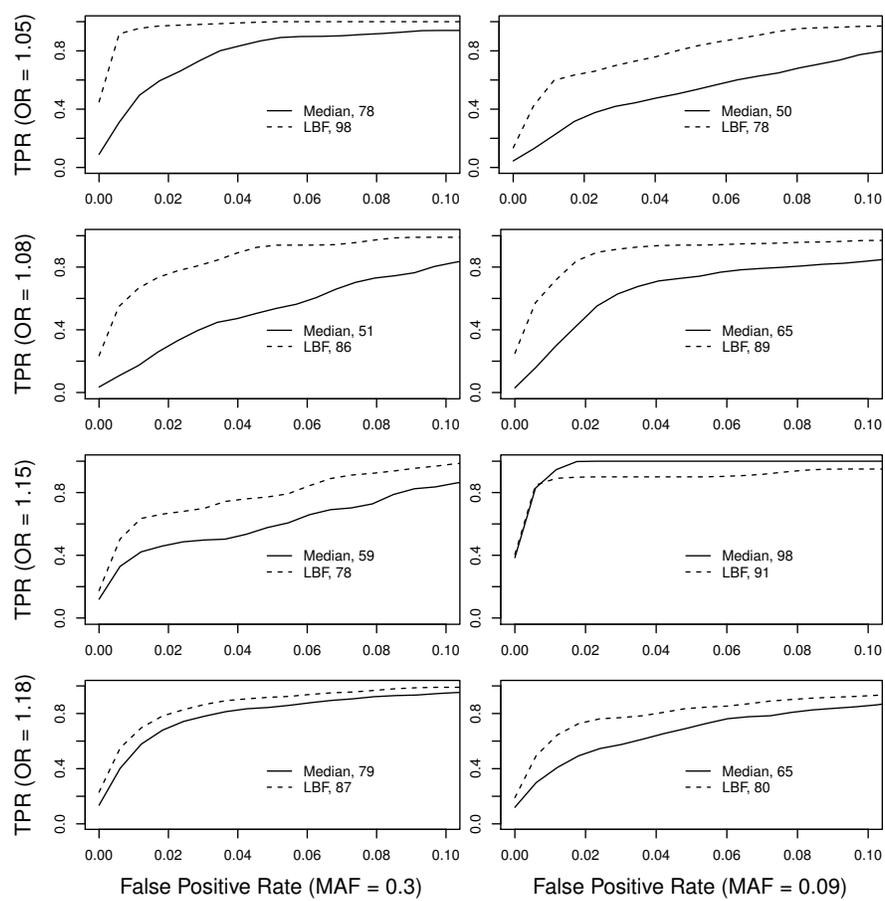


Figure 2

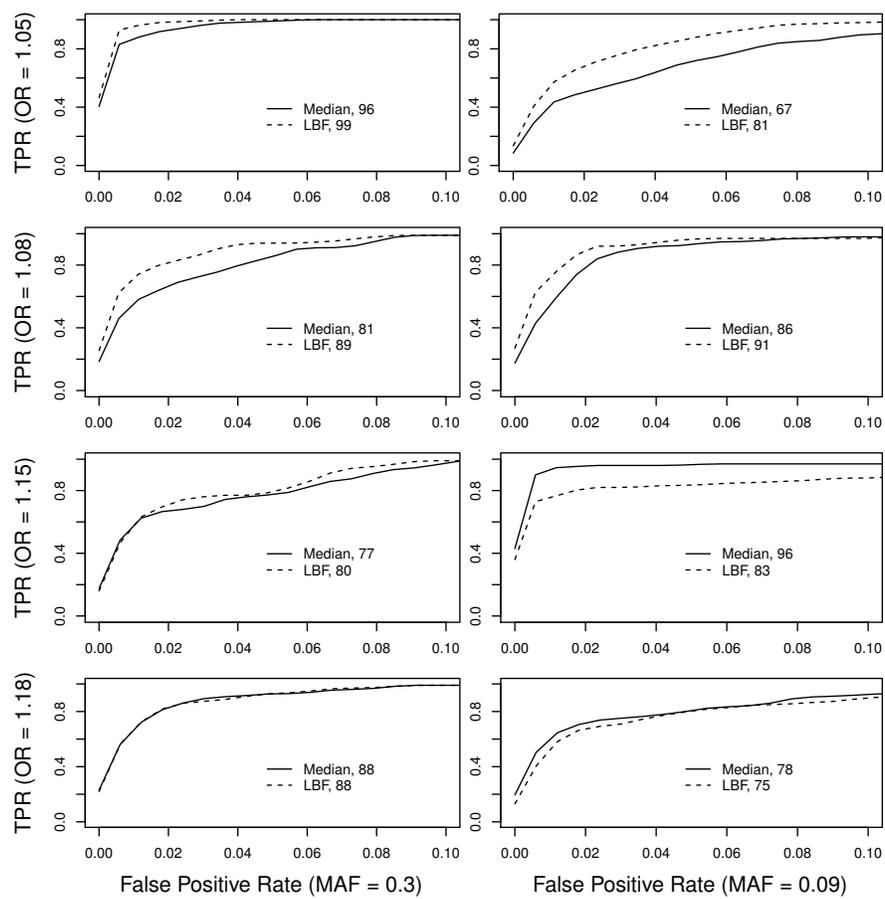


Figure 3

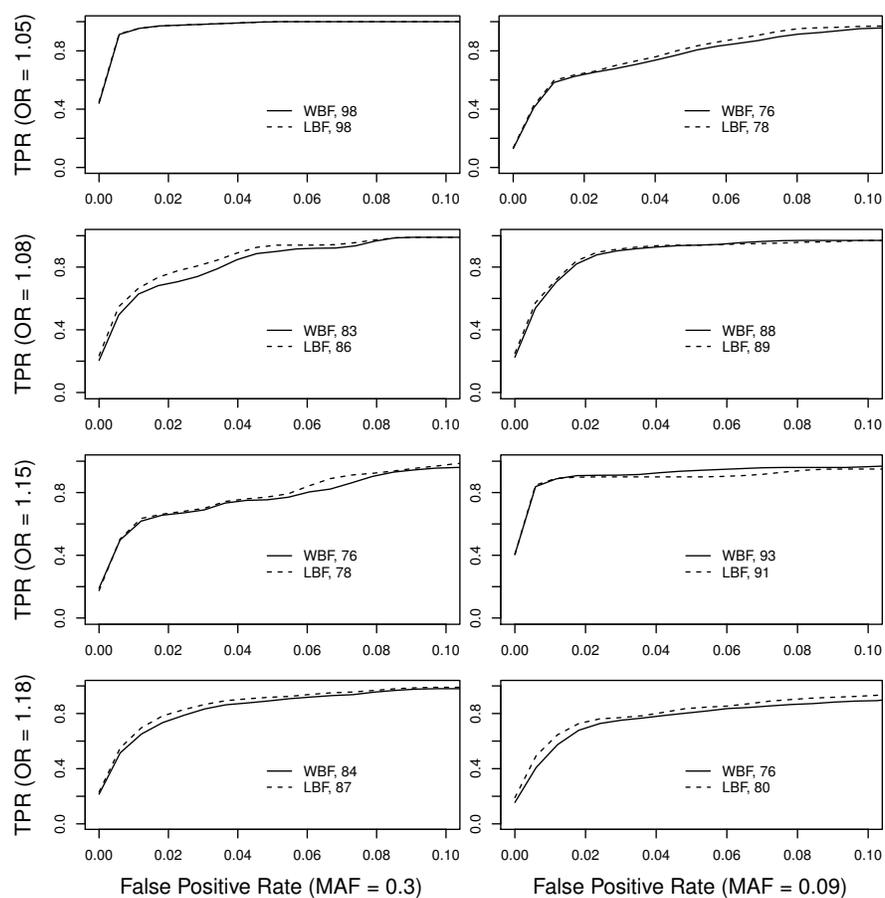


Figure 4

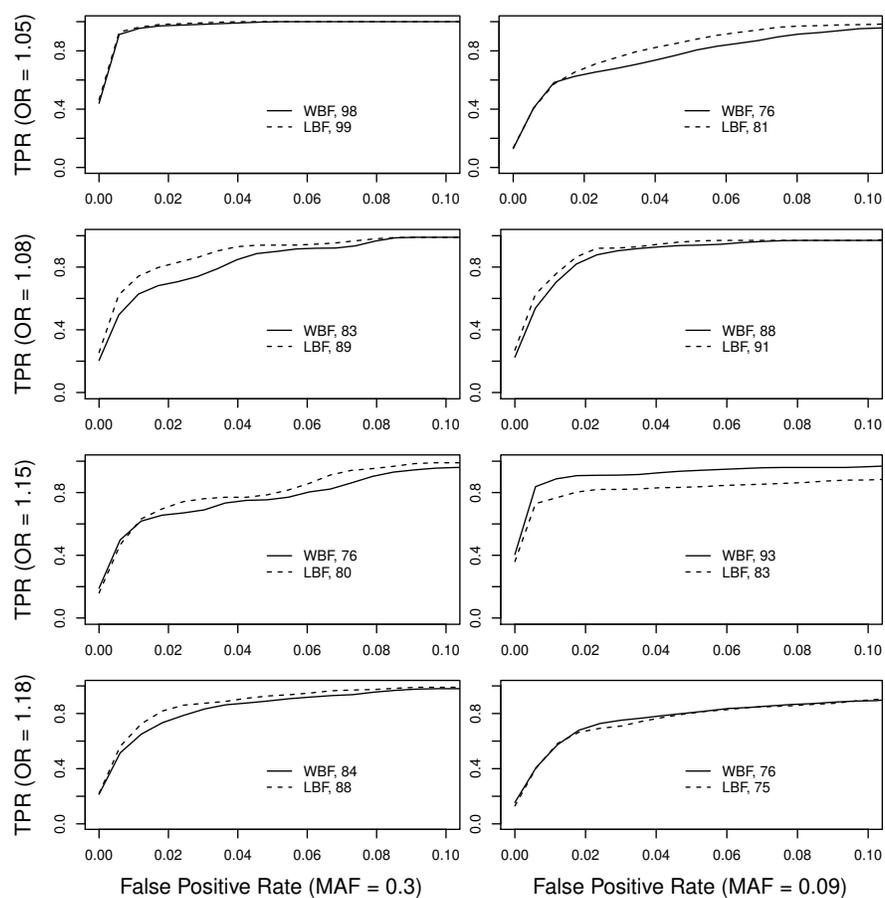


Figure 5

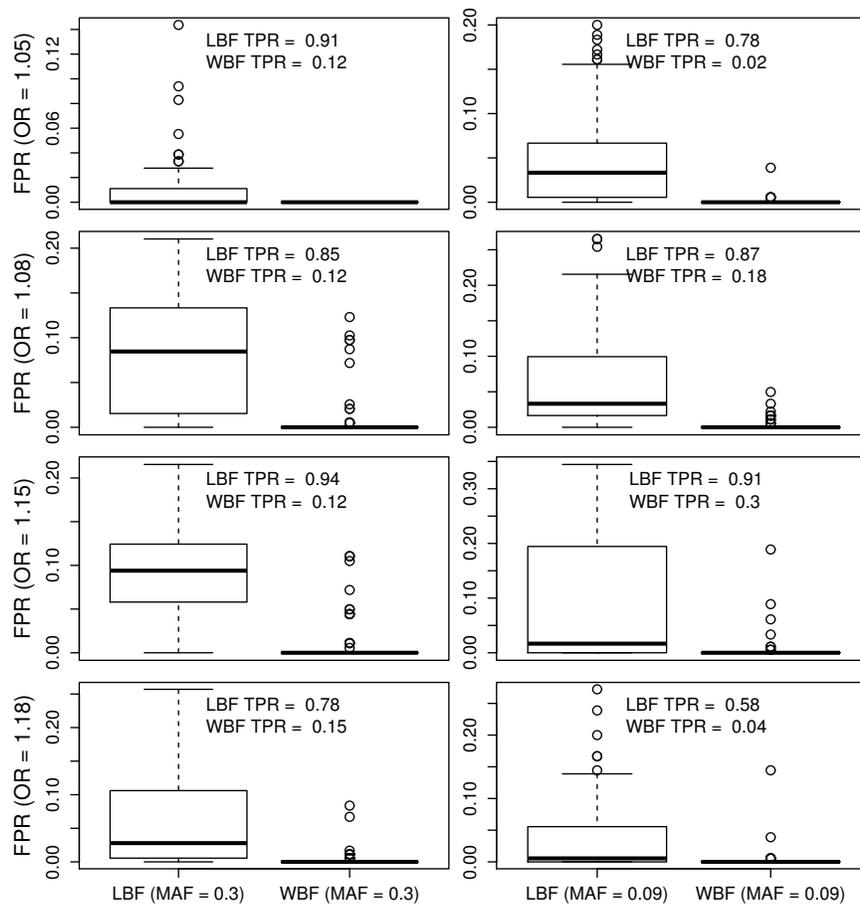


Figure 6

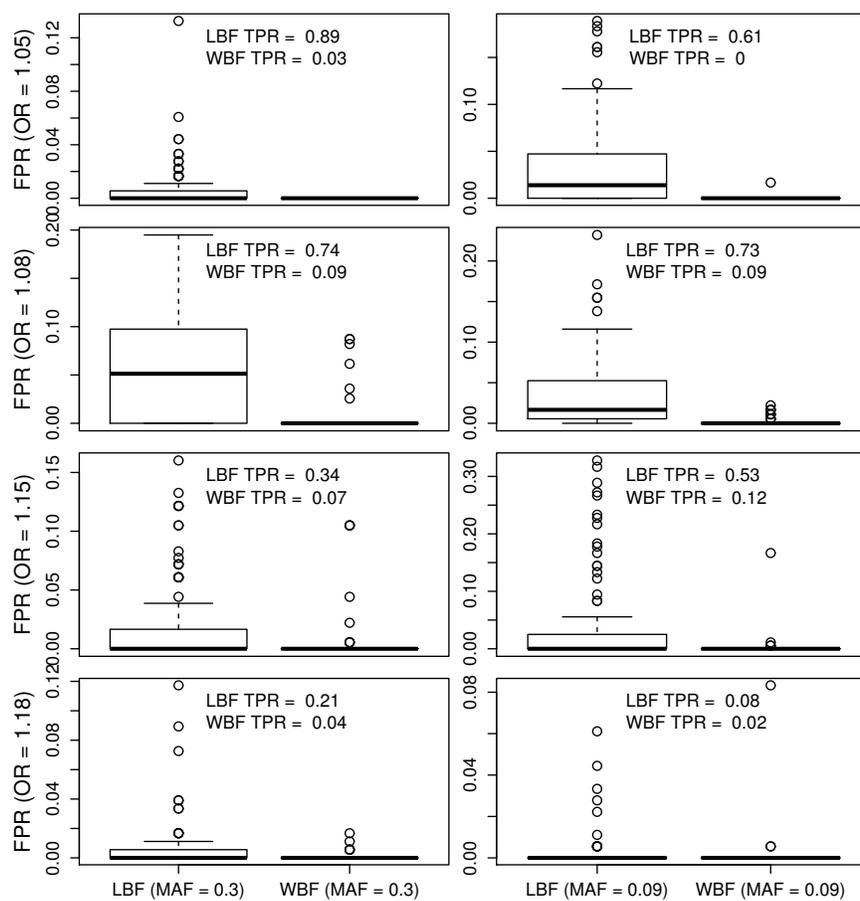


Figure 7