

## RESEARCH ARTICLE

# Bayesian sample size determination for diagnostic accuracy studies

Kevin J. Wilson<sup>1</sup>  | S. Faye Williamson<sup>2</sup>  | A. Joy Allen<sup>3,4</sup> |  
Cameron J. Williams<sup>1,3,4</sup> | Thomas P. Hellyer<sup>4</sup> | B. Clare Lendrem<sup>3,4</sup>

<sup>1</sup>School of Mathematics, Statistics & Physics, Newcastle University, Tyne and Wear, UK

<sup>2</sup>Biostatistics Research Group, Population Health Sciences Institute, Newcastle University, Tyne and Wear, UK

<sup>3</sup>NIHR Newcastle In Vitro Diagnostics Co-operative, Newcastle University, Tyne and Wear, UK

<sup>4</sup>Translational and Clinical Research Institute, Newcastle University, Tyne and Wear, UK

## Correspondence

Kevin J. Wilson, School of Mathematics, Statistics & Physics, Newcastle University, Herschel Building, Tyne and Wear NE1 7RU, UK.

Email: [kevin.wilson@ncl.ac.uk](mailto:kevin.wilson@ncl.ac.uk)

## Abstract

The development of a new diagnostic test ideally follows a sequence of stages which, among other aims, evaluate technical performance. This includes an analytical validity study, a diagnostic accuracy study, and an interventional clinical utility study. In this article, we propose a novel Bayesian approach to sample size determination for the diagnostic accuracy study, which takes advantage of information available from the analytical validity stage. We utilize assurance to calculate the required sample size based on the target width of a posterior probability interval and can choose to use or disregard the data from the analytical validity study when subsequently inferring measures of test accuracy. Sensitivity analyses are performed to assess the robustness of the proposed sample size to the choice of prior, and prior-data conflict is evaluated by comparing the data to the prior predictive distributions. We illustrate the proposed approach using a motivating real-life application involving a diagnostic test for ventilator associated pneumonia. Finally, we compare the properties of the approach against commonly used alternatives. The results show that, when suitable prior information is available, the assurance-based approach can reduce the required sample size when compared to alternative approaches.

## KEYWORDS

Bayesian assurance, binomial intervals, contingency tables, power calculations, sensitivity, specificity

## 1 | INTRODUCTION

Diagnostic accuracy studies evaluate the ability of a diagnostic test (the index test) to correctly identify patients with and without a target condition. This is typically achieved by prospectively comparing results from the index test to the true disease status obtained from the best available reference standard for a cohort of patients. The two main measures used to assess intrinsic diagnostic accuracy are sensitivity and specificity. For a test to proceed to the next stage of evidence development, it is important that these measures are estimated to an appropriate degree of accuracy. This hinges on the sample size chosen for the diagnostic accuracy study. Too small a sample size will lead to an imprecise estimate with wide corresponding intervals, which is non-informative to the decision maker, and contributes to research waste.<sup>1</sup> Alternatively, too large a sample size may delay the results of the study due to longer recruitment times and resource limitations,

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

in addition to financial and ethical implications.<sup>2</sup> Consequently, choosing a sample size which strikes a balance between accuracy and efficiency is a crucial step in the design of any diagnostic accuracy study.

Traditional sample size calculations are based on a hypothesis-testing framework. The idea is to choose a sample size such that the probability of rejecting the null hypothesis when there is a clinically relevant difference is greater than a required power (typically 80% or 90%) with a specified type I error rate (typically 5% for a two-sided test).<sup>3</sup> However, a sample size which captures the precision of the measure of interest, by targeting a desirable width of the corresponding confidence interval, can be more appropriate in certain circumstances.<sup>4,5</sup> This is pertinent in early clinical diagnostic studies, where the aim is to estimate test accuracy with sufficient precision, which is the approach adopted here.

In this article, we consider the sample size problem from a Bayesian perspective and propose a novel approach, referred to as the *Bayesian assurance method (BAM)*, to determine sample sizes for diagnostic accuracy studies. In doing so, we explore whether utilizing information from the preceding laboratory study will reduce the sample size in the diagnostic accuracy study, and thus lead to a more efficient development process. This may be important if there is need to deploy accurate diagnostic tests rapidly, such as in response to the COVID-19 pandemic, where early detection of infectious individuals is critical to outbreak containment.<sup>6</sup> Another relevant area is rare diseases, where there are a limited number of patients available, or where there are practical or ethical issues with conducting large studies. This extends to (rare) disease subgroups, in which the sensitivity and specificity of a diagnostic test can vary.<sup>7</sup>

The BAM shares similar characteristics to seamless and adaptive designs, in that it utilizes data from one stage to inform decisions in the subsequent stages in order to improve efficiency and flexibility. Seamless designs, which aim to combine separate studies, and adaptive designs, which allow for prespecified modifications to the design based on accruing data, are well-established in interventional studies, yet have received little attention in the context of diagnostics. However, the flexibility offered by these designs is just as important in diagnostic accuracy studies. Motivated by the desire to accelerate diagnostic research, Vach et al<sup>8</sup> and Zapf et al<sup>9</sup> discuss the utility of seamless and adaptive designs, respectively, in developing diagnostics. Zapf et al<sup>9</sup> advocate the development and implementation of adaptive designs for diagnostics, and highlight this as a promising area for future research, which this article contributes towards.

The BAM can be used to choose the sample size according to both sensitivity and specificity criteria simultaneously, rather than separately as in most existing methods. Criteria for combining sensitivity and specificity to define the success of a diagnostic test, and how this affects the sample size required, are discussed by Vach et al.<sup>10</sup> Korevaar et al<sup>11</sup> suggest specifying a joint hypothesis on the sensitivity and specificity based on predefined minimally acceptable criteria. Branscum et al<sup>12</sup> proposed an approach to choose the sample size based on the predictive probability that the posterior probability of the sensitivity and specificity both being within prespecified limits is high. Although the assurance approach in this article is related to that taken by Branscum et al,<sup>12</sup> there are some key differences. For example, they required the estimated sensitivity and specificity, along with the upper and lower limits for both intervals, to be specified in advance, and focused only on a two-sided approach, whereas we assure the widths of the intervals directly, requiring only the prior distributions for the parameters, and consider both the one- and two-sided cases.

Several existing approaches consider binomial confidence intervals based on a normal approximation to determine the sample size (referred to as the Wald interval)<sup>13</sup> or some adjustment to it, for example, the Agresti-Coull (AC) interval.<sup>14</sup> An alternative is to use an exact binomial interval (known as the Clopper-Pearson [CP] interval<sup>15</sup>). A description of commonly used intervals for proportions is provided in Newcombe.<sup>15(Chapter3)</sup> Zhou et al<sup>13(Chapter4)</sup> recommend the Zhou et al<sup>16</sup> interval for values of sensitivity or specificity close to zero or one. Another recommended interval is the equal-tailed Jeffreys interval,<sup>17</sup> constructed using a Bayesian approach with a non-informative Jeffreys prior (ie,  $\text{Beta}(1/2, 1/2)$ ) for the binomial proportion. Wei and Hutson<sup>18</sup> provide a sample size calculation based on the conditional expectation of interval width given a hypothesized proportion. We compare the BAM to some of these approaches in Section 6.

Sample size determination from a Bayesian perspective is typically based on *assurance*, which is considered an alternative to power.<sup>19</sup> Assurance, and modifications to it, can be referred to as the probability of success<sup>20</sup> and the expected/average power,<sup>21</sup> among others; a review is provided in Kunzmann et al.<sup>3(Section5)</sup> Unlike power, which is conditional on the true (but unknown) parameter value, the distinguishing property of assurance is that it is an *unconditional* probability which incorporates parameter uncertainty through a prior distribution and integration over the parameter range.<sup>22</sup> This is formally defined in Section 3.

The use of assurance for sample size calculations has occurred predominantly within clinical trials.<sup>3</sup> In this article, we use assurance to represent the probability of obtaining the desired accuracy (based on a target interval width) in our estimates of sensitivity and/or specificity. The sample size is then taken to be the minimum which yields the required assurance. We describe inference for a standard diagnostic accuracy study in Section 2. The BAM is presented and further described in Section 3, with issues such as prior sensitivity and prior-data conflict addressed in Section 4. As a motivating

**TABLE 1** (A) A  $2 \times 2$  contingency table for a typical diagnostic accuracy study. (B) The  $2 \times 2$  contingency table for the *biomarker selection study* based on the biomarker IL-1 $\beta$ . (C) The  $2 \times 2$  contingency table for the *diagnostic accuracy study* based on the biomarker IL-1 $\beta$ 

(A)	Disease	No disease	Total
Test positive	$n_{1,1}$	$n_{1,2}$	$n_{1,T}$
Test negative	$n_{2,1}$	$n_{2,2}$	$n_{2,T}$
Total	$n_{T,1}$	$n_{T,2}$	$n_T$
(B)	VAP	No VAP	Total
Test positive	16	35	51
Test negative	1	20	21
Total	17	55	72
(C)	VAP	No VAP	Total
Test positive	51	55	106
Test negative	2	42	44
Total	53	97	150

case study, we use the BAM to redesign a diagnostic accuracy study of a test for ventilator associated pneumonia (VAP) in Section 5, and assess the properties of the BAM, in comparison to some standard approaches, in Section 6.

## 2 | INFERENCE IN A DIAGNOSTIC ACCURACY STUDY

We consider a diagnostic accuracy study to assess an index test under development. In the study, we observe the numbers of individuals in a  $2 \times 2$  contingency table (Table 1A).

The number of individuals with and without the disease is assumed to be known, based on a reference test. The intrinsic accuracy of the index test can be measured by its sensitivity and specificity, defined as the probability of a positive test given disease and the probability of a negative test given no disease, respectively.

There are two approaches used to model numbers of individuals in the cells of the  $2 \times 2$  table: assuming either binomial or multinomial likelihoods. In the first case,  $n_{1,1}|\lambda, n_{T,1} \sim \text{Bin}(n_{T,1}, \lambda)$  and  $n_{2,2}|\theta, n_{T,2} \sim \text{Bin}(n_{T,2}, \theta)$ , where  $\lambda$  is the sensitivity and  $\theta$  is the specificity of the index test. The conjugate prior distributions are  $\lambda \sim \text{Beta}(a_\lambda, b_\lambda)$  and  $\theta \sim \text{Beta}(a_\theta, b_\theta)$ . If we assume in the prior that the sensitivity and specificity are independent, then their posterior distributions are  $\lambda|\mathbf{n} \sim \text{Beta}(a_\lambda + n_{1,1}, b_\lambda + n_{2,1})$  and  $\theta|\mathbf{n} \sim \text{Beta}(a_\theta + n_{2,2}, b_\theta + n_{1,2})$ . The independence assumption will often be reasonable since the diagnostic thresholds for the test are fixed at this stage, and the sensitivity and specificity consider mutually exclusive populations of patients.

In the second case, we consider the vector  $\mathbf{n} = (n_{1,1}, n_{1,2}, n_{2,1}, n_{2,2})'$  and assume  $\mathbf{n}|\boldsymbol{\gamma} \sim \text{Multi}(\mathbf{n}, \boldsymbol{\gamma})$ , where  $\boldsymbol{\gamma} = (\gamma_{1,1}, \gamma_{1,2}, \gamma_{2,1}, \gamma_{2,2})'$  is a vector containing the probabilities of each cell of the contingency table. Here, the sensitivity and specificity are given by  $\lambda = \gamma_{1,1}/(\gamma_{1,1} + \gamma_{2,1})$  and  $\theta = \gamma_{2,2}/(\gamma_{1,2} + \gamma_{2,2})$ . A typical form for the prior distribution is a Dirichlet distribution, which provides conjugacy. That is,  $\boldsymbol{\gamma} \sim \text{Dir}(\boldsymbol{\alpha})$ , where  $\boldsymbol{\alpha} = (\alpha_{1,1}, \alpha_{1,2}, \alpha_{2,1}, \alpha_{2,2})'$ .

It can be shown that the two approaches are equivalent in terms of inference for the sensitivity and specificity (see the Appendix). In this article, we will use the binomial form as it allows for the direct specification of the priors for the sensitivity, specificity, and prevalence. We will assume conjugate beta priors, as detailed above, throughout the rest of the article.

## 3 | SAMPLE SIZE DETERMINATION

### 3.1 | Assurance

Assurance is a Bayesian alternative to power to choose a sample size. Consider a two-armed clinical trial in which a hypothesis test is to be conducted with  $H_0 : \delta = 0$  vs  $H_1 : \delta > 0$ , where  $\delta$  represents the difference in the effect of two treatments. A typical power calculation would choose a sample size to provide a certain statistical power at a

particular assumed value  $\delta_c$  for  $\delta$ , often taken to be the minimal clinically relevant difference. In this case, the power is  $\Pr(\text{Reject } H_0 | \delta = \delta_c)$  and would increase with sample size.

In practice, the choice of  $\delta_c$  is relatively arbitrary. As the true effect size  $\delta$  is unknown, this can result in conditioning on an event which is extremely unlikely. One approach to mitigate this is to conduct a sensitivity analysis, varying the value of  $\delta_c$  and choosing a sample size which is robust to small perturbations.<sup>23</sup> In the Bayesian context, we can take an alternative approach, and represent our uncertainty over  $\delta$  using a prior distribution  $\pi(\delta)$ . The assurance is the expected power of the hypothesis test with respect to this prior,  $A(n) = E_\delta [\Pr(\text{Reject } H_0 | \delta)] = \int \Pr(\text{Reject } H_0 | \delta) \pi(\delta) d\delta$ . We choose to make the dependence on the sample size  $n$  explicit for the assurance  $A(\cdot)$ .

Assurance is not restricted to the case where we will perform a hypothesis test at the end of a trial. If we perform a Bayesian analysis instead, then we may declare the trial a success and the new treatment superior if  $\Pr(\delta \leq 0) \leq 0.05$  in the posterior, for example. In this case,  $A(n) = E_\delta [\Pr(\text{Trial a success} | \delta)] = \int \Pr(\text{Trial a success} | \delta) \pi(\delta) d\delta$ . Thus, the assurance is the unconditional probability that the trial results in a successful outcome.

We use assurance to choose a sample size to estimate sensitivity, specificity, or both, of the index test to a certain degree of accuracy. We initially focus on *sensitivity* of the index test,  $\lambda$  and consider two cases: assuring the width of the posterior probability interval (two-sided), and assuring the width of the *lower half* of the posterior probability interval (one-sided).

### 3.1.1 | Two-sided case

Considering the inference from Section 2, a  $100(1 - \alpha)\%$  symmetric posterior probability interval for  $\lambda$  is  $(\lambda_L, \lambda_U)$ , where the limits of the interval are defined such that  $\Pr(\lambda \leq \lambda_L | \mathbf{n}) = \frac{\alpha}{2}$  and  $\Pr(\lambda \geq \lambda_U | \mathbf{n}) = \frac{\alpha}{2}$ . The accuracy of the estimation of  $\lambda$  can be considered as the width of this interval,  $\lambda_U - \lambda_L$ , and a successful diagnostic accuracy study would produce an interval with a width smaller than some target,  $\lambda_U - \lambda_L \leq w^*$ .

Suppose the number of individuals with the disease in the study,  $n_{T,1}$ , is fixed. There are three possibilities: no values of  $n_{1,1}$  lead to an interval with width smaller than  $w^*$ , all values of  $n_{1,1}$  lead to an interval with width smaller than  $w^*$ , or some values of  $n_{1,1}$  lead to an interval with width smaller than  $w^*$ . To investigate the third case, consider the posterior variance of  $\lambda$ ,  $\text{Var}(\lambda | \mathbf{n}) = \{(a_\lambda + n_{1,1})(b_\lambda + n_{2,1})\} / \{(a_\lambda + b_\lambda + n_{T,1})^2 (a_\lambda + b_\lambda + n_{T,1} + 1)\}$ . For a fixed sample size  $n_{T,1}$ , the denominator of this fraction is constant. That is,  $\text{Var}(\lambda | \mathbf{n}) \propto (a_\lambda + n_{1,1})(b_\lambda + n_{2,1}) \propto n_{1,1}(b_\lambda - a_\lambda + n_{T,1}) - n_{1,1}^2$ , substituting  $n_{2,1} = n_{T,1} - n_{1,1}$ . The variance is quadratic in  $n_{1,1}$  and the squared term has a negative coefficient. Thus, the posterior probability interval will be narrower than  $w^*$  when  $n_{1,1} \leq c_1$  and  $n_{1,1} \geq c_2$ , for two critical numbers of individuals  $c_1 < c_2$ . We define this set as  $\mathcal{N} = \{n_{1,1} : n_{1,1} \leq c_1 \text{ or } n_{1,1} \geq c_2\}$ .

### 3.1.2 | One-sided case

We consider a  $100(1 - \alpha)\%$  posterior probability interval for  $\lambda$  of the form  $(\lambda_L, 1)$ , where the lower limit of the interval is defined such that  $\Pr(\lambda \leq \lambda_L | \mathbf{n}) = \alpha$ . We consider the distance between the lower limit of the interval and a central point estimate of  $\lambda$ , that is,  $\lambda_{0.5} - \lambda_L$ , where  $\lambda_{0.5}$  is the posterior median. A successful diagnostic accuracy study would result in this interval having a width smaller than some target,  $\lambda_{0.5} - \lambda_L \leq w^*$ .

By the same logic as the two-sided case, the posterior probability interval will be narrower than  $w^*$  when  $n_{1,1} \leq c_1$  and  $n_{1,1} \geq c_2$ , for two critical numbers of individuals  $c_1 < c_2$ . Thus, we consider the set  $\mathcal{N} = \{n_{1,1} : n_{1,1} \leq c_1 \text{ or } n_{1,1} \geq c_2\}$  for the one-sided case, with  $c_1$  and  $c_2$  determined by the interval  $\lambda_{0.5} - \lambda_L$ .

### 3.1.3 | Evaluating the assurance

We can obtain an expression for the assurance for a sample size  $n_T$ , conditional on a fixed number of diseased individuals  $n_{T,1}$ . This is denoted by  $A_\lambda(n_T | n_{T,1})$  and defined as

$$\begin{aligned} A_\lambda(n_T | n_{T,1}) &= \int \Pr(\text{Accuracy achieved} | \lambda) \pi(\lambda) d\lambda, \\ &= \frac{\Gamma(a_\lambda + b_\lambda)}{\Gamma(a_\lambda) \Gamma(b_\lambda)} \sum_{n_{1,1} \in \mathcal{N}} \binom{n_{T,1}}{n_{1,1}} \frac{\Gamma(a_\lambda + n_{1,1}) \Gamma(b_\lambda + n_{2,1})}{\Gamma(a_\lambda + b_\lambda + n_{T,1})}, \end{aligned} \quad (1)$$

where  $\Gamma(\cdot)$  represents the gamma function. A derivation is given in Section A of the supplementary material.

As the number of individuals with the disease,  $n_{T,1}$ , will not be known in advance, we need to sum over the possible values  $n_{T,1}$  can take. If we have a random sample from the target population, then  $n_{T,1}|\rho \sim \text{Bin}(n_T, \rho)$ , where  $\rho$  is the prevalence of the disease. Let  $\rho \sim \text{Beta}(a_\rho, b_\rho)$  for some chosen values of  $(a_\rho, b_\rho)$ . The unconditional assurance is then

$$A_\lambda(n_T) = \sum_{n_{T,1}=0}^{n_T} \left[ \int \Pr(\text{Accuracy achieved}|\lambda)\pi(\lambda)d\lambda \times f(n_{T,1}) \right],$$

where  $f(n_{T,1}) = \int f(n_{T,1}|\rho)\pi(\rho)d\rho$  is the probability of observing  $n_{T,1}$  individuals in the disease group. The assurance can thus be expressed as

$$A_\lambda(n_T) = \frac{\Gamma(a_\lambda + b_\lambda)}{\Gamma(a_\lambda)\Gamma(b_\lambda)} \sum_{n_{T,1}=0}^{n_T} \left\{ \sum_{n_{1,1} \in \mathcal{N}} \left[ \binom{n_{T,1}}{n_{1,1}} \frac{\Gamma(a_\lambda + n_{1,1})\Gamma(b_\lambda + n_{2,1})}{\Gamma(a_\lambda + b_\lambda + n_{T,1})} \right] \times \binom{n_T}{n_{T,1}} \frac{\Gamma(a_\rho + b_\rho)}{\Gamma(a_\rho)\Gamma(b_\rho)} \frac{\Gamma(a_\rho + n_{T,1})\Gamma(b_\rho + n_{T,2})}{\Gamma(a_\rho + b_\rho + n_T)} \right\}. \quad (2)$$

This is derived in Section A of the supplementary material.

All that remains is to find the values of  $(c_1, c_2)$ . For each fixed sample size,  $n_T$ , and number of diseased individuals,  $n_{T,1}$ , the values of  $\lambda_L$ ,  $\lambda_{0.5}$ , and  $\lambda_U$  will depend only on  $n_{1,1}$  and, hence, the width of the interval will be a function of  $n_{1,1}$ ,  $W(n_{1,1})$ , in both cases. Therefore,  $c_1 = \text{argmin} \{W(n_{1,1}) \geq w^*\} - 1$  and  $c_2 = \text{argmax} \{W(n_{1,1}) \geq w^*\} + 1$  for  $n_1 < n_{T,1} < n_2$ , where  $n_1$  is a number below which the interval can never achieve the desired width and  $n_2$  is a number above which the width of the interval is always below  $w^*$ . Hence,  $A_\lambda(n_T|n_{T,1}) = 0$  for all  $n_{T,1} \leq n_1$  and  $A_\lambda(n_T|n_{T,1}) = 1$  for all  $n_{T,1} \geq n_2$ .

To estimate the *specificity* of the index test to a given accuracy of  $w^*$ , we can derive the assurance in the same way, which results in an assurance analogous to that in Equation (2). The details are given in Section A of the supplementary material.

Finally, suppose we wish to estimate *both* the sensitivity and specificity to a particular accuracy. Consider different accuracy targets,  $w_\lambda^*$  and  $w_\theta^*$ , for the sensitivity and specificity, respectively. In this case, the assurance for the sample size  $n_T$  conditional on  $n_{T,1}$  (and hence  $n_{T,2}$ , since  $n_{T,2} = n_T - n_{T,1}$ ) is given by

$$A_{\lambda,\theta}(n_T|n_{T,1}) = \int \int \Pr(\text{Accuracy achieved}|\lambda, \theta)\pi(\lambda)\pi(\theta)d\lambda d\theta, \\ = \frac{\Gamma(a_\lambda + b_\lambda)}{\Gamma(a_\lambda)\Gamma(b_\lambda)} \sum_{n_{1,1} \in \mathcal{N}_1} \binom{n_{T,1}}{n_{1,1}} \frac{\Gamma(a_\lambda + n_{1,1})\Gamma(b_\lambda + n_{2,1})}{\Gamma(a_\lambda + b_\lambda + n_{T,1})} \times \frac{\Gamma(a_\theta + b_\theta)}{\Gamma(a_\theta)\Gamma(b_\theta)} \sum_{n_{2,2} \in \mathcal{N}_2} \binom{n_{T,2}}{n_{2,2}} \frac{\Gamma(a_\theta + n_{2,2})\Gamma(b_\theta + n_{1,2})}{\Gamma(a_\theta + b_\theta + n_{T,2})},$$

where  $\mathcal{N}_1$  contains the values  $n_{1,1} \leq c_1$  and  $n_{1,1} \geq c_2$  that give a posterior interval narrower than  $w_\lambda^*$  for the sensitivity, and  $\mathcal{N}_2$  contains the values  $n_{2,2} \leq \tilde{c}_1$  and  $n_{2,2} \geq \tilde{c}_2$  that give a posterior interval narrower than  $w_\theta^*$  for the specificity.

To find the unconditional assurance, we sum over the possible values of  $n_{T,1}$  to give:

$$A_{\lambda,\theta}(n_T) = \frac{\Gamma(a_\lambda + b_\lambda)}{\Gamma(a_\lambda)\Gamma(b_\lambda)} \frac{\Gamma(a_\theta + b_\theta)}{\Gamma(a_\theta)\Gamma(b_\theta)} \sum_{n_{T,1}=0}^{n_T} \left\{ \sum_{n_{1,1} \in \mathcal{N}_1} \left[ \binom{n_{T,1}}{n_{1,1}} \frac{\Gamma(a_\lambda + n_{1,1})\Gamma(b_\lambda + n_{2,1})}{\Gamma(a_\lambda + b_\lambda + n_{T,1})} \right] \right. \\ \left. \times \sum_{n_{2,2} \in \mathcal{N}_2} \left[ \binom{n_{T,2}}{n_{2,2}} \frac{\Gamma(a_\theta + n_{2,2})\Gamma(b_\theta + n_{1,2})}{\Gamma(a_\theta + b_\theta + n_{T,2})} \right] \times \binom{n_T}{n_{T,1}} \frac{\Gamma(a_\rho + b_\rho)}{\Gamma(a_\rho)\Gamma(b_\rho)} \frac{\Gamma(a_\rho + n_{T,1})\Gamma(b_\rho + n_{T,2})}{\Gamma(a_\rho + b_\rho + n_T)} \right\}. \quad (3)$$

The proposed BAM is now summarized via the following steps:

1. Choose whether we wish to assure our estimate of sensitivity  $\lambda$ , specificity  $\theta$ , or both.
2. Choose a target width(s)  $w^*$  for the accuracy measure(s), a one- or two-sided posterior interval and a level  $\alpha$  for the interval.
3. Specify the prior distributions for the chosen accuracy measure(s) and the prevalence  $\rho$ . We detail how to do this in the next section.

4. Use Equation (2) or (3) (or see Section A of the supplementary material) to calculate the assurance for sample sizes  $n_T = 1, 2, \dots$ .
5. Choose the minimum sample size  $n_T^*$  to give the desired assurance.

*Example:* Suppose we wish to estimate both sensitivity and specificity to within 5%, with posterior probability 0.99 using a two-sided interval, that is,  $w^* = 0.05$  and  $\alpha = 0.01$ . We specify prior distributions for  $\lambda$ ,  $\theta$ , and  $\rho$ , and use Equation (3) to evaluate the assurance for sample sizes  $n_T = 1, 2, \dots$ . To achieve the desired accuracy with a probability of at least 0.9, say, we choose the smallest value of  $n_T$  which gives rise to an assurance greater than 0.9.

## 4 | PRIOR SPECIFICATION AND MODEL CHECKING

A diagnostic accuracy study is part of an extensive development process for the diagnostic test, see Figure 1 in Reference 24. Its main purpose is to estimate performance characteristics of the test, particularly the sensitivity and specificity, in the target population in a clinically relevant setting. Prior to the diagnostic accuracy study is the analytical validity phase, in which the test may still be under development and the data generated may be used to support regulatory approvals.<sup>24</sup> The validation conducted during this stage may test individuals from the target population. Consequently, the data produced can be used to inform the prior distributions in the diagnostic accuracy study. This assumes that the observations in the two stages are exchangeable, which may not always be reasonable. Therefore, in Section B of the supplementary material, we detail how the BAM can be used under weaker assumptions.

### 4.1 | Specifying prior distributions

Consider the analytical validity testing. Suppose that a random sample of  $n_T^0$  individuals was taken and the numbers in the cells of the  $2 \times 2$  contingency table were  $\mathbf{n}^0 = (n_{1,1}^0, n_{1,2}^0, n_{2,1}^0, n_{2,2}^0)'$ . Using the inferential approach in Section 2, priors for the sensitivity, specificity, and prevalence would be  $\lambda \sim \text{Beta}(a_\lambda^0, b_\lambda^0)$ ,  $\theta \sim \text{Beta}(a_\theta^0, b_\theta^0)$ , and  $\rho \sim \text{Beta}(a_\rho^0, b_\rho^0)$ , respectively. The corresponding posterior distributions (excluding conditioning statements) would be  $\lambda \sim \text{Beta}(a_\lambda^1, b_\lambda^1)$ ,  $\theta \sim \text{Beta}(a_\theta^1, b_\theta^1)$  and  $\rho \sim \text{Beta}(a_\rho^1, b_\rho^1)$ , where  $a_\lambda^1 = a_\lambda^0 + n_{1,1}^0$ ,  $b_\lambda^1 = b_\lambda^0 + n_{2,1}^0$ ,  $a_\theta^1 = a_\theta^0 + n_{2,2}^0$ ,  $b_\theta^1 = b_\theta^0 + n_{1,2}^0$ ,  $a_\rho^1 = a_\rho^0 + n_{T,1}^0$ , and  $b_\rho^1 = b_\rho^0 + n_{T,2}^0$ .

These latter beta distributions can be used as priors for the diagnostic accuracy study. Although this does not negate the necessity of choosing the initial prior values  $(a_\lambda^0, b_\lambda^0)$ ,  $(a_\theta^0, b_\theta^0)$ , and  $(a_\rho^0, b_\rho^0)$ , these will have a small effect on the sample size chosen if sufficient data are available from the analytical validity stage. This is explored further in the next section. The approach taken here is equivalent to using a power prior with the parameter quantifying the heterogeneity between the diagnostic study population and analytic validity population set equal to one (representing homogeneous populations). In cases of heterogeneity between the two populations, a power prior could be used with this parameter taking a value in the range  $[0, 1]$ . For full details, see Reference 25.

In cases where it is controversial to use data from the analytical validity stage when inferring the sensitivity and specificity of the test, we could use a weaker prior in the analysis, but retain the original prior in the design to inform the sample size calculations. This is illustrated in Section B of the supplementary material

### 4.2 | Prior sensitivity

The choice of initial prior parameters,  $(a_\lambda^0, b_\lambda^0)$ ,  $(a_\theta^0, b_\theta^0)$ , and  $(a_\rho^0, b_\rho^0)$  may have little effect on the assurance if sufficient data are observed at the analytic validity stage. We explore this using local sensitivity analysis and investigate the following two questions:

1. How does the optimal sample size,  $n_T^*$ , change when varying the prior parameters?
2. How does the assurance at  $n_T^*$ ,  $A(n_T^*)$ , change when varying the prior parameters?

In particular, we vary the prior parameters  $(a_C^0, b_C^0)$  for  $C = \{\lambda, \theta, \rho\}$  in turn over a range of values around their initial values, and record the smallest and largest values of the optimal sample size  $(n_T^*, \bar{n}_T^*)$  and assurance  $(\underline{A}(n_T^*), \bar{A}(n_T^*))$ . If these values do not differ by much, then the optimal sample size is relatively robust to the initial prior choice.

Using the grid search approach<sup>26</sup> to determine an appropriate range of prior parameter values, we explore the sensitivity on a grid  $G_{a^0, b^0}(\epsilon)$ , where  $\epsilon$  represents the distance between a prior and the original prior with parameters  $(a^0, b^0)$ . That is,  $G_{a^0, b^0}(\epsilon) = \{(a, b) : d(\pi_{a, b}(\gamma), \pi_{a^0, b^0}(\gamma)) = \epsilon\}$ , where  $\pi_{a, b}(\gamma)$  represents the beta prior distribution with parameters  $(a, b)$  and  $\gamma$  is one of  $\lambda, \theta$ , and  $\rho$ . We use the Hellinger distance<sup>26</sup> which, for the beta distribution, can be expressed as

$$d(\pi_{a, b}(\gamma), \pi_{a^0, b^0}(\gamma)) = 1 - \frac{B([a + a^0]/2, [b + b^0]/2)}{\sqrt{B(a, b)B(a^0, b^0)}},$$

where  $B(a, b) = \Gamma(a)\Gamma(b)/\Gamma(a + b)$  is the beta function. To conduct the grid search, it is sensible to work in polar co-ordinates. Therefore, we set  $a = \exp(z) \cos(\varphi)$  and  $b = \exp(z) \sin(\varphi)$ , where  $z = \log(r)$ . We search in the range  $\varphi \in [-\pi, \pi]$ , solving for the value of  $r$  which gives the correct value of  $\epsilon$ .

To find the values of  $a$  and  $b$ , we convert back via  $a = a^0 + r \cos(\varphi)$  and  $b = b^0 + r \sin(\varphi)$ . From this grid search, we can then find the corresponding  $(n_T^*, \bar{n}_T^*)$  and  $(\underline{A}(n_T^*), \bar{A}(n_T^*))$  for this  $\epsilon$ . We suggest a sensible choice of  $\epsilon$  in Section 5.2.

### 4.3 | Prior-data conflict

Label the counts in the  $2 \times 2$  table from the diagnostic accuracy study  $\mathbf{n}^1 = (n_{1,1}^1, n_{1,2}^1, n_{2,1}^1, n_{2,2}^1)'$ . The posterior distributions for the sensitivity and specificity (omitting the conditioning) will be  $\lambda \sim \text{Beta}(a_\lambda^2, b_\lambda^2)$  and  $\theta \sim \text{Beta}(a_\theta^2, b_\theta^2)$ , respectively, where

$$\begin{aligned} a_\lambda^2 &= a_\lambda^1 + n_{1,1}^1 = a_\lambda^0 + n_{1,1}^0 + n_{1,1}^1, & b_\lambda^2 &= b_\lambda^1 + n_{2,1}^1 = b_\lambda^0 + n_{2,1}^0 + n_{2,1}^1, \\ a_\theta^2 &= a_\theta^1 + n_{2,2}^1 = a_\theta^0 + n_{2,2}^0 + n_{2,2}^1, & b_\theta^2 &= b_\theta^1 + n_{1,2}^1 = b_\theta^0 + n_{1,2}^0 + n_{1,2}^1. \end{aligned}$$

The inference for the sensitivity and specificity is in the form of a weighted average of the prior and the observations, with weights determined by the relative sample sizes of each. The prior is made up of a weighted average of the observations in the analytical validity stage and the original prior. If all of the elements are in broad agreement, then the posterior distribution will provide an accurate summary of the properties of the index test in the population of interest. However, it could be the case that the prior and observations are not in agreement, which is known as *prior-data conflict*.<sup>27,28</sup> For example, if the two studies are carried out at different times or in different locations, the spectrum of disease in the target population may not be the same. In this case, it is important to investigate why the differences are there and what action should be taken.

We can evaluate prior-data conflict by comparing the observations to the prior predictive distributions of the parameters. We consider the prior predictive distributions of the number of observations in the disease group,  $n_{T,1}$ , and, conditional on this, the number who test positive,  $n_{1,1}$  of those with the disease, and the number who test negative of those without the disease,  $n_{2,2}$ . These are given by  $f(y) = \binom{n}{y} \frac{B(a+y, b+n-y)}{B(a, b)}$ , where  $y$  is  $(n_{T,1}, n_{1,1}, n_{2,2})$  in turn,  $n$  is the corresponding sample size, that is,  $(n_T, n_{T,1}, n_{T,2})$ , and  $(a, b)$  are the beta distribution parameter values for the prevalence, sensitivity, and specificity, respectively. We can then plot the prior predictive distributions and calculate probabilities of the form  $\Pr(n \geq n_{\text{obs}})$ , for observed number of individuals  $n_{\text{obs}}$ . If the observed value lies in the body of the associated prior predictive distribution, then that prior is consistent with the data. Otherwise, this provides evidence of prior-data conflict.

## 5 | A BIOMARKER TEST FOR VAP

Using published results,<sup>29,30</sup> we consider the development of a biomarker test for VAP. The development of the test involved four stages; an exploratory study to look at possible biomarkers for VAP diagnosis, a single center observational study to choose suitable biomarkers, a multicenter diagnostic accuracy study to develop biomarker cut offs and validate accuracy and a randomized controlled trial of clinical utility. At each stage, the target population was patients on

a ventilator with suspected VAP. The reference standard test was the growth of pathogens at  $>10^4$  colony forming units per milliliter of bronchoalveolar fluid. All patients with suspected VAP receive antibiotics, although only 20% to 60% of patients will have VAP confirmed by the reference standard, leading to overuse of antibiotics. Microbiology culture and sensitivities take up to 72 hours to return results to clinicians, which delays the opportunity to discontinue antibiotics in patients who do not have infection. A rapid, highly sensitive biomarker test could allow for early stopping of antibiotics.

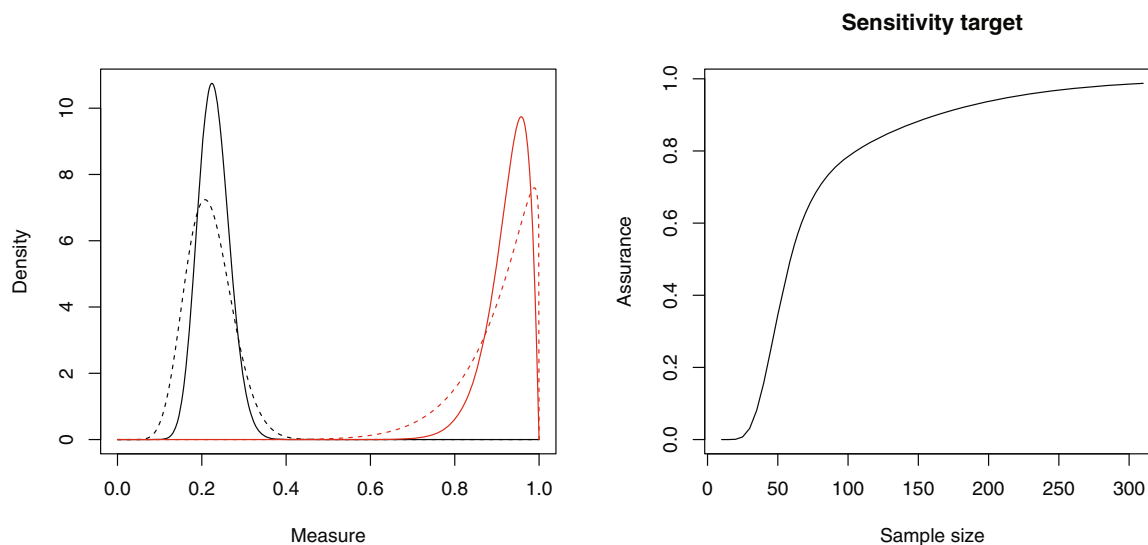
We consider planning the diagnostic accuracy study. The sample size was originally chosen to reduce the width of the 95% confidence interval for the post-test probability of VAP to 0.16, and resulted in  $n_T = 150$ . Estimates from the single center observational study were used to calculate the sample size. The estimated sensitivity and prevalence in the single center observational study were  $\hat{\lambda} = 0.94$  and  $\hat{\rho} = 0.24$ , respectively, for the most promising biomarker, IL-1 $\beta$ . If instead the sample size had been chosen based on a confidence interval for the sensitivity, using the Wald interval,<sup>13</sup> a larger sample size of 196 would have been required.

## 5.1 | Choosing the sample size using assurance

To use assurance to determine the sample size, we require the prior parameters for the sensitivity,  $(a_\lambda^0, b_\lambda^0)$ , and the prevalence,  $(a_\rho^0, b_\rho^0)$ , before the biomarker selection study. In the initial exploratory study, there were 55 patients, 12 of whom were confirmed by the reference test to have VAP. Assuming exchangeability, a suitable prior for the prevalence is  $\rho \sim \text{Beta}(12, 43)$ . The most promising biomarker gave an estimated sensitivity of 0.93. Since it was unclear which biomarker(s) would be used in the final test, it is not reasonable to make an exchangeability assumption for the test results in the two stages. A more suitable prior for the sensitivity is more diffuse but with a mean around this value, such as  $\lambda \sim \text{Beta}(9.9, 1.1)$ . These priors are represented by the dashed lines in Figure 1.

In the biomarker selection study, the  $2 \times 2$  contingency table is provided in Table 1B for the most promising biomarker, IL-1 $\beta$ .

We assume that these patients are exchangeable with those in the diagnostic accuracy study as they are randomly sampled from the same population. Therefore, the prior distributions for the diagnostic accuracy study are  $\lambda \sim \text{Beta}(25.9, 2.1)$  and  $\rho \sim \text{Beta}(29, 98)$  (see Section 2), illustrated by solid lines in the left-hand side of Figure 1. Suppose we would like to estimate the sensitivity of the test to within 0.16 in a 95% symmetric probability interval and choose a sample size to give 80% assurance. Based on the priors above, we use the BAM to obtain a sample size of  $n_T^* = 106$ . This is significantly smaller than the original sample size of  $n_T = 150$  (which would give an assurance of 88%). The full assurance curve for is provided in the right-hand side of Figure 1. Note that the assurance curve has a different shape to a power curve, and is monotonically increasing between 0 and 1.



**FIGURE 1** Left: The prior distributions for the sensitivity (red) and the prevalence (black) for the biomarker selection study (dashed lines) and the diagnostic accuracy study (solid lines). Right: The assurance curve showing the assurance achieved at different sample sizes for the diagnostic accuracy study



**TABLE 2** The smallest and largest values of the assurance,  $A(n_T)$ , and the smallest and largest sample sizes,  $n_T$ , found in the local sensitivity analysis

Measure	$\min\{A(n_T)\}$	$\max\{A(n_T)\}$
Sensitivity	0.73	0.86
Prevalence	0.80	0.81
Measure	$\min\{n_T\}$	$\max\{n_T\}$
Sensitivity	82	130
Prevalence	104	108

## 5.2 | Prior sensitivity

To assess the sensitivity of the sample size and assurance to the prior distribution, we use the approach outlined in Section 4.2. In particular, we conduct a grid search for both the sensitivity and prevalence priors using a value of  $\varepsilon = 0.00354$  (equivalent to a mean shift in a standard normal random variable of 0.1).

The resulting values of the beta distribution parameters ( $a, b$ ) are provided in Section B.3 of the supplementary material for the sensitivity and prevalence. The corresponding smallest and largest values of the assurance and sample size are provided in Table 2.

Changes to the prevalence prior have little effect on the sample size or the assurance at  $n_T = 106$ . The effect is slightly larger for the sensitivity prior but, even for the most extreme prior, a sample size of 130 would be sufficient (which is considerably less than the sample size of 150 in the study).

## 5.3 | Prior-data conflict

The results from the diagnostic accuracy study with the 150 patients are summarized in Table 1C for the biomarker IL-1 $\beta$ .

The resulting posterior distributions for the sensitivity of and prevalence are  $\lambda \sim \text{Beta}(76.9, 4.1)$  and  $\rho \sim \text{Beta}(82, 195)$ , respectively. The corresponding 95% posterior probability interval for the sensitivity is (0.893, 0.986), and so we meet the target of 0.16 on the width of the interval. To assess possible prior-data conflict, we use the approach detailed in Section 4.3 and compare the observations to the prior predictive distributions.


The prior predictive distributions of the number of patients with VAP (left) and the number of patients with VAP who tested positive (right) are provided in Figure 2, with the observation shown as a red dashed line. A color version of this figure can be found in the electronic version of the article.

We see the number of patients correctly diagnosed with VAP lies within the main body of the prior predictive distribution. The observed number of patients with VAP lies in the body of the distribution, but is closer to the upper tail, in the 99th percentile. The observed number of patients correctly diagnosed lies in the 76th percentile. This provides some evidence of prior-data conflict for the number of patients with VAP, so we may choose a prior on the prevalence which is not based on the single center observational study.

The posterior mean and 95% posterior probability interval for the prevalence are 0.296 and (0.244, 0.351), respectively. The same quantities using a flat prior with  $a_\rho = b_\rho = 1$  are 0.355 and (0.281, 0.433), respectively, which would not affect the inference on the sensitivity. However, if we believe the sub-populations with VAP are different between the two stages we may also consider an alternative prior for sensitivity.

## 6 | ALTERNATIVE APPROACHES

In this section, we compare properties of the proposed BAM to alternative commonly used methods. Assume we wish to obtain the number of individuals with the disease,  $n_{T,1}$ , required to estimate the sensitivity to within a particular degree of accuracy. The alternative methods are based on a hypothesis test of  $H_0 : \lambda = \lambda_0$  against the two-tailed alternative  $H_1 : \lambda \neq \lambda_0$  conducted at a significance level of  $\alpha$ . We take the value of  $\lambda_0$  to be  $\hat{\lambda}$ , that is, the maximum likelihood estimate of sensitivity using the analytical validity data. The sample size can be chosen according to a desired power



**FIGURE 2** The prior predictive distributions of the number of patients with VAP (left) and the number of VAP patients who test positive (right) together with the observations (red)

of  $\beta$  to detect a difference of size  $w^*$ . As discussed in Section 1, there are several possible approaches; we consider the following.

The first is based on a normal approximation. In this case, to achieve a power of  $\beta$  we choose the sample size in the disease group as  $n_{T,1} = \left\lceil \left\{ \left( z_{\alpha/2} + z_{\beta} \right) \sqrt{\hat{\lambda}(1 - \hat{\lambda})} \right\}^2 / (w^*/2)^2 \right\rceil$ , where  $z$  is the upper percentile of a standard normal distribution. We construct a  $100(1 - \alpha)\%$  confidence interval based on this normal approximation, known as the *Wald interval*.

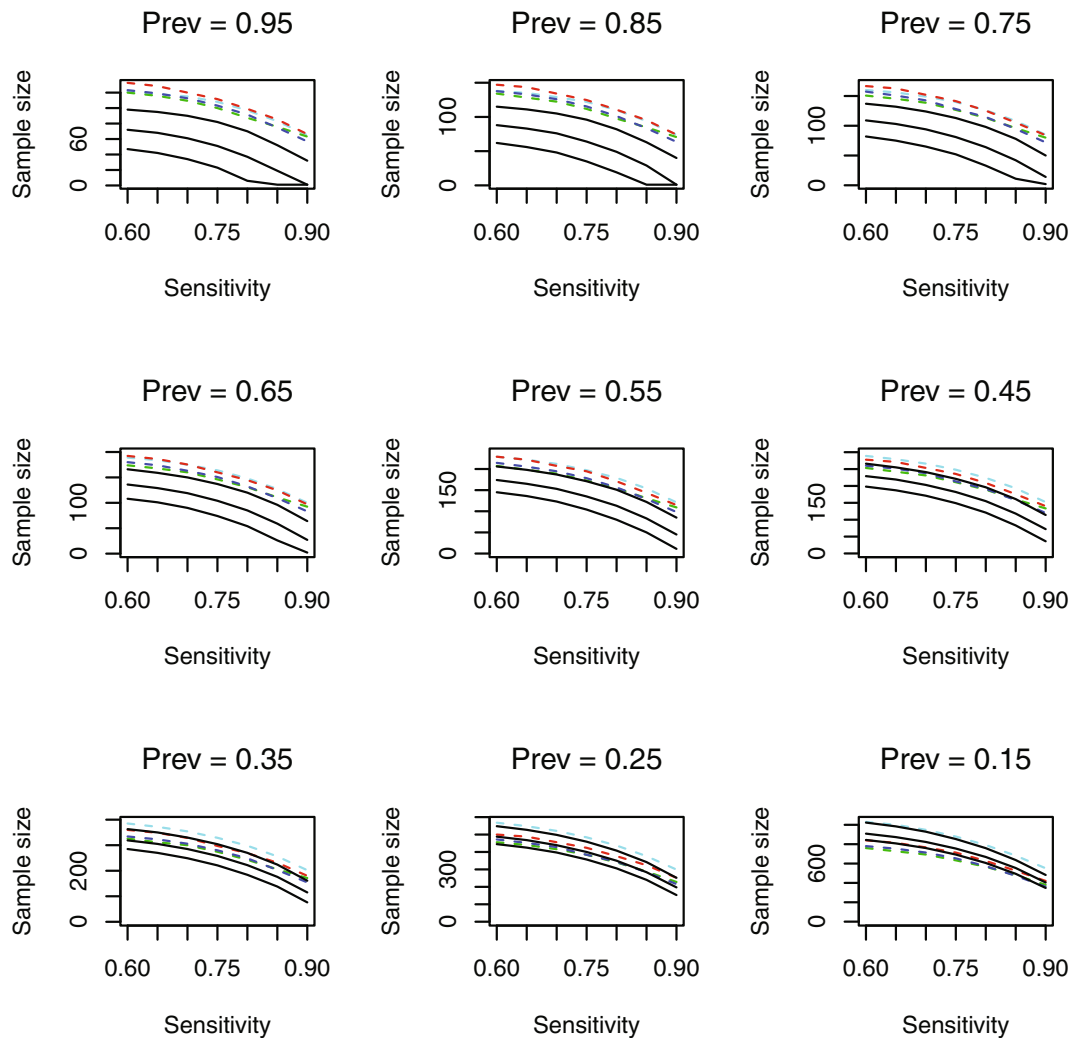
The second approach is based on an exact binomial test to give the CP interval. The third approach combines the normal approximation with an adjustment to the hypothesized value as the center of the interval to give the AC interval.

In practice, the standard way of obtaining the required sample size is to use the appropriate sample size formula (if available), or in-built functions within statistical software (eg, the `binDesign` function from the `binGroup` R package<sup>31</sup>). However, these often give rise to unreliable sample sizes and, in our investigation, are shown to perform poorly over the range of parameter values considered; see Section E of the supplementary material. We instead rely on simulation. That is, we choose the smallest sample size  $n_{T,1}$  to give the correct proportion of intervals below the desired target width  $w^*$ , based on simulating confidence intervals repeatedly and finding the power empirically. The *total* number of individuals to recruit,  $n_T$ , is found by scaling with respect to the estimated prevalence  $\hat{\rho}$ , that is,  $n_T = n_{T,1}/\hat{\rho}$ . The same procedure is used to obtain the number of individuals without the disease,  $n_{T,2}$ , required to estimate the specificity to a certain degree of accuracy. In this case,  $n_T = n_{T,2}/(1 - \hat{\rho})$ .

### 6.1 | Comparison of sample sizes

In this section, we compare the sample sizes required for a diagnostic accuracy study using the methods outlined above. We consider a significance level of  $\alpha = 0.05$ , a power/assurance of  $\beta = 0.8$ , and aim to estimate sensitivity to within 0.18 in a two-sided interval. We vary the sensitivity over the range [0.6, 0.9] and the prevalence over the range [0.15, 0.95]. For the proposed BAM, we consider three prior sample sizes of  $n_T^0 = 25, 50,$  and  $75$  to represent “small,” “medium,” and “large” analytical validity studies. The results for all scenarios and methods are illustrated in Figure 3.

Note that the power calculations are based on the true parameter values. The assurance calculation, however, uses beta priors with parameters  $(n_T^0 \lambda \rho, n_T^0 \rho [1 - \lambda])$  for the sensitivity and  $(n_T^0 \rho, n_T^0 [1 - \rho])$  for the prevalence. An assurance

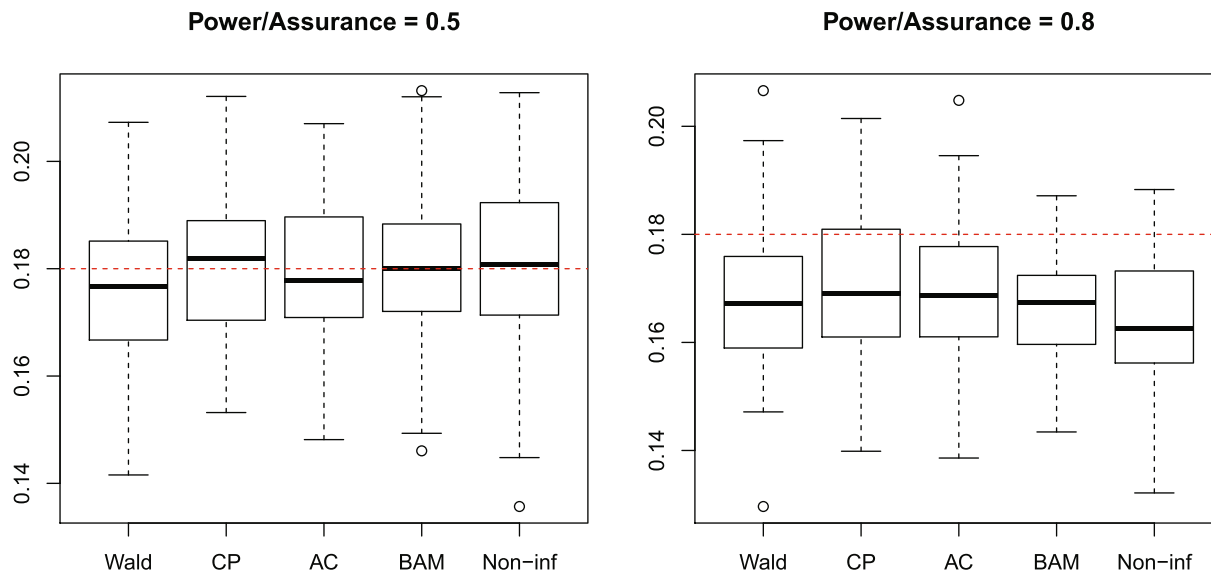


**FIGURE 3** A comparison of the sample sizes required based on power calculations (dashed) using a Wald interval (dark blue), Clopper-Pearson (red), Agresti-Coull (green), assurance (black, solid), and assurance based on non-informative analysis priors (light blue). In each plot, there are three black curves relating to prior sample sizes of (from top to bottom) 25, 50, and 75

calculation with non-informative priors for the analysis is also considered. This is based on a design prior from the “small” analytical validity study to represent a reasonable “worst case” scenario.

In Figure 3, we observe similar patterns across the frequentist approaches (represented by the colored lines) for each prevalence. CP always results in the largest sample size, with Wald and AC giving similar, slightly smaller, sample sizes. In comparison to assurance, the frequentist methods produce larger sample sizes when the prevalence is high. In some scenarios, they result in smaller sample sizes. For example, when the prior sample size is 25 below a prevalence of 0.5, when the prior sample size is 50 below a prevalence of 0.3, and when the prior sample size is 75 around a prevalence of 0.2. However, as the sensitivity increases, the required sample size based on assurance reduces quicker than the frequentist approaches, which are known to perform poorly as the sensitivity approaches one.

Further details are provided in Section C of the supplementary material, including an assessment of different target interval widths. The message is consistent across the parameter combinations considered: assurance for the sensitivity reduces the required sample size in the majority of cases, particularly in moderate to high prevalence populations and when a highly accurate test is required. High prevalence situations are common in secondary care, where patients have already been triaged (such as in a suspected stroke<sup>32</sup>), or in cancer pathways by the time an invasive test, such as a biopsy, is used. When the BAM is applied to even lower prevalences of 0.1, 0.05, and 0.01, the sample sizes required for a sensitivity of 0.9, and based on a medium analytic validity study, are 681, 1643, and 2770, respectively. Such low prevalences may be the case in large-scale geographic prevalence surveys, for example.



**FIGURE 4** The width of 95% confidence or posterior probability intervals based on 100 simulations for the Wald interval (Wald), Clopper-Pearson (CP), Agresti-Coull (AC), Assurance (BAM), and Assurance using a non-informative analysis prior (Non-inf). The power/assurance used to choose the sample size was 0.5 (left) and 0.8 (right). The horizontal line is at the desired width of  $w^* = 0.18$

## 6.2 | Comparison of interval widths

A smaller sample size will not be useful if the corresponding interval estimates are very wide. Therefore, we conduct a simulation study, outlined below, to assess the width of the intervals resulting from each approach.

First, we sample values of the sensitivity and prevalence from uniform distributions. These are used to sample analytical validity results,  $n_{1,1}^0$  and  $n_{T,1}^0$ , from their respective binomial distributions based on a “medium” total sample size of  $n_T = 50$ . From these data, we find estimates of the sensitivity and prevalence for the power calculations and set the prior distributions for the assurance calculations. We then find the required sample size for each method. We sample the results of the diagnostic accuracy study,  $n_{1,1}^1$  and  $n_{T,1}^1$ , from their respective binomial distributions and use these to calculate  $100(1 - \alpha)\%$  intervals for the sensitivity. Finally, we calculate the width of the intervals. By repeating this process 100 times, we consider the distributions of widths of the intervals, which are shown in Figure 4 for a power/assurance,  $\beta$ , of 0.5 (left) and 0.8 (right). In all cases,  $\alpha = 0.05$  and  $w^* = 0.18$ .

For both  $\beta = 0.5$  and  $\beta = 0.8$ , the approaches produce intervals with a similar distribution of widths. When  $\beta = 0.5$ , the median width of each approach lies approximately at the target width. When  $\beta = 0.8$ , the target width is around, or slightly above, the upper quartile for each method. Thus, the different sample sizes observed in the previous section do not come at the expense of less precision in inference.

The simulations were repeated with interval widths of  $w^* = 0.22$  and  $w^* = 0.14$ . The corresponding results are provided in Section D of the supplementary material. The main conclusions remain: for a power/assurance of 0.5, all of the distributions are approximately centered on the target width, and for a power/assurance of 0.8, each approach produces intervals which include the target width in the upper 25% of its empirical distribution.

In addition, we have investigated the properties of the BAM when assuring both sensitivity and specificity together, in terms of the sample size required and the resulting interval widths. This is provided in Section F of the supplementary material.

## 7 | DISCUSSION

In this article, we have proposed the novel BAM to determine sample sizes for diagnostic accuracy studies. Bayesian assurance fulfills a similar role to power and, as we have shown, can offer benefits when suitable prior information is available. In particular, representing uncertainty in unknown test characteristics using prior distributions, and utilizing information from different stages of the development pathway, allows for a wider range of evidence to be seamlessly

incorporated in the design and analysis of a diagnostic accuracy study. Consequently, we have shown that this has the potential to reduce the sample size, thus increasing efficiency in evidence development.

If no prior information is available, or accessible, from earlier stages of development, expert elicitation can be used to form the necessary prior distributions. Elicited distributions can include opinions from multiple experts, or be combined with data from other sources.<sup>33</sup> The larger the prior sample size, the more informative the prior distribution will be which, as shown in Figure 3, typically corresponds to a smaller sample size in the diagnostic accuracy study. If it is not appropriate to use an informative prior for the analysis (eg, to mitigate researcher bias), a skeptical or flat prior can be used instead. The BAM has the flexibility to allow for distinct prior distributions in the design and analysis stages, as illustrated in Section B of the supplementary material.

The proposed BAM can be used regardless of whether the final analysis is frequentist or Bayesian. Some assurance calculations may not result in closed form solutions (eg, if a Bayesian analysis uses a non-conjugate analysis prior), in which case, simulation and numerical methods are required. Thus, calculating assurance can be challenging and, unlike power, is not available in standard software packages. To increase accessibility of the BAM, R code is provided and an R Shiny application is currently under development.

This work focuses on assuring sensitivity and specificity as measures of diagnostic accuracy. We have also shown how the BAM can be used to assure sensitivity and specificity jointly, for which no existing approaches are available, to our knowledge. The assurance calculations can be modified to obtain sample sizes for other quantities, such as positive and negative predictive values or the area under the curve. Moreover, the assurance calculations could be extended to allow for multiple categorical results, or results in the form of continuous measures, which is area of further work. In this article, we considered the evaluation of a single diagnostic test, but further work could explore how the proposed method extends to multiple tests.

To reflect standard practice in diagnostic accuracy studies, we have inherently assumed that the sampling plan will be produced prior to the study, carried out accordingly and then the data analyzed at the end of the study. Future work could extend the approach so that it can be applied sequentially, participant-by-participant (or in blocks), to monitor the width of the posterior interval until the desired value is attained, at which point the study would terminate. This would reduce the sample size required. However, it would require a change in the way that diagnostic accuracy studies are routinely implemented.

## ACKNOWLEDGEMENTS

The authors thank Prof. James Wason and Prof. John Simpson for useful discussions. AJA, CJW and BCL are supported by the NIHR Newcastle In Vitro Diagnostics Co-operative. The funders had no role in the preparation or decision to publish this manuscript. The views are those of the authors and not necessarily those of the NHS, NIHR, or the Department of Health and Social Care. The VAPrapid trial was supported by the Department of Health and Wellcome Trust via the Health Innovation Challenge Fund.

## DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

## ORCID

Kevin J. Wilson  <https://orcid.org/0000-0001-9450-862X>

S. Faye Williamson  <https://orcid.org/0000-0002-1800-6625>

## REFERENCES

1. Ioannidis JPA, Greenland S, Hlatky MA, et al. Increasing value and reducing waste in research design, conduct, and analysis. *Lancet*. 2014;383(9912):166-175.
2. Altman DG. Statistics and ethics in medical research: III how large a sample? *BMJ*. 1980;281(6251):1336-1338. doi:10.1136/bmj.281.6251.1336
3. Kunzmann K, Grayling MJ, Lee KM, Robertson DS, Rufibach K, Wason JMS. A review of Bayesian perspectives on sample size derivation for confirmatory trials. *Am Stat*. 2021;75(4):424-432. doi:10.1080/00031305.2021.1901782
4. Jiroutek MR, Muller KE, Kupper LL, Stewart PW. A new method for choosing sample size for confidence interval-Based inferences. *Biometrics*. 2003;59(3):580-590.
5. Jia B, Lynn HS. A sample size planning approach that considers both statistical significance and clinical significance. *Trials*. 2015;16:213. doi:10.1186/s13063-015-0727-9

6. Rosenthal PJ. The importance of diagnostic testing during a viral pandemic: early lessons from novel coronavirus disease (COVID-19). *Am J Trop Med Hygiene*. 2020;102(5):915-916. doi:10.4269/ajtmh.20-0216
7. Bachmann LM, Puhan MA, Riet G, Bossuyt PM. Sample sizes of studies on diagnostic accuracy: literature survey. *BMJ*. 2006;332(7550):1127-1129.
8. Vach W, Bibiza E, Gerke O, Bossuyt PM, Friede T, Zapf A. A potential for seamless designs in diagnostic research could be identified. *J Clin Epidemiol*. 2021;129:51-59. doi:10.1016/j.jclinepi.2020.09.019
9. Zapf A, Stark M, Gerke O, et al. Adaptive trial designs in diagnostic accuracy research. *Stat Med*. 2020;39(5):591-601. doi:10.1002/sim.8430
10. Vach W, Gerke O, Høilund-Carlsen PF. Three principles to define the success of a diagnostic study could be identified. *J Clin Epidemiol*. 2012;65(3):293-300. doi:10.1016/j.jclinepi.2011.07.004
11. Korevaar DA, Gopalakrishna G, Cohen JF, Bossuyt PM. Targeted test evaluation: a framework for designing diagnostic accuracy studies with clear study hypotheses. *Diagn Progn Res*. 2019;3(22):1-10.
12. Branscum AJ, Johnson WO, Gardner IA. Sample size calculations for studies designed to evaluate diagnostic test accuracy. *J Agric Biol Environ Stat*. 2007;12(1):112-127.
13. Zhou X-H, Obuchowski NA, McClish DK. *Statistical Methods in Diagnostic Medicine*. 2nd ed. Hoboken, NJ: John Wiley & Sons; 2011.
14. Agresti A, Coull BA. Approximate is better than exact for interval estimation of binomial proportions. *Am Stat*. 1998;52:119-126.
15. Newcombe RG. *Confidence Intervals for Proportions and Related Measures of Effect Size*. Hoboken, NJ: CRC Press; 2012.
16. Zhou XH, Li CM, Yang Z. Improving interval estimation of binomial proportions. *Philos Trans*. 1874;2008(366):2405-2418.
17. Brown LD, Cai TT, DasGupta A. Interval Estimation for a Binomial Proportion. *Stat Sci*. 2001;16(2):101-117.
18. Wei L, Hutson AD. A comment on sample size calculations for binomial confidence intervals. *J Appl Stat*. 2013;40(2):311-319. doi:10.1080/02664763.2012.740629
19. O'Hagan A, Stevens JW. Bayesian assessment of sample size for clinical trials of cost-effectiveness. *Med Decis Mak*. 2001;21:219-230.
20. Bertsche A, Fleischer F, Beyersmann J, Nehmiz G. Bayesian Phase II optimization for time-to-event data based on historical information. *Stat Methods Med Res*. 2019;28(4):1272-1289. doi:10.1177/0962280217747310
21. Kowalski KG. Integration of Pharmacometric and Statistical Analyses Using Clinical Trial Simulations to Enhance Quantitative Decision Making in Clinical Drug Development. *Statistics in Biopharmaceutical Research*. 2019;11(1):85-103. doi:10.1080/19466315.2018.1560361
22. Chen D-G, Ho S. From statistical power to statistical assurance: It's time for a paradigm change in clinical trial design. *Commun Stat Simul Comput*. 2017;46(10):7957-7971. doi:10.1080/03610918.2016.1259476
23. Matthews JNS. *Introduction to Randomized Controlled Clinical Trials*. 2nd ed. Hoboken, NJ: Chapman & Hall/CRC Press; 2006.
24. Graziadio S, Winter A, Lendrem BC, et al. How to ease the pain of taking a diagnostic point of care test to the market: a framework for evidence development. *Micromachines*. 2020;11(3):291. doi:10.3390/mi11030291
25. Ibrahim JG, Chen M-H, Gwon Y, Chen F. The power prior: theory and applications. *Stat Med*. 2015;34(28):3724-3749.
26. Roos M, Martins TG, Held L, Rue H, others. Sensitivity analysis for Bayesian hierarchical models. *Bayesian Anal*. 2015;10(2):321-349.
27. Box GEP. Sampling and Bayes' inference in scientific modelling and robustness. *J Royal Stat Soc Ser A*. 1980;143:383-430.
28. Schmidli H, Gsteiger S, Roychoudhury S, O'Hagan A, Spiegelhalter D, Neuenschwander B. Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics*. 2014;70(4):1023-1032.
29. Conway MA, Kefala K, Wilkinson TS, et al. Diagnostic importance of pulmonary interleukin-1 $\beta$  and interleukin-8 in ventilator-associated pneumonia. *Thorax*. 2010;65(3):201-207.
30. Hellyer TP, Morris AC, McAuley DF, et al. Diagnostic accuracy of pulmonary host inflammatory mediators in the exclusion of ventilator-acquired pneumonia. *Thorax*. 2015;70(1):41-47.
31. Zhang B, Bilder C, Biggerstaff B, Schaarschmidt F, Hitt B. binGroup: evaluation and experimental design for binomial group testing online. R package version 2.2-1; 2018.
32. Shaw L, Graziadio S, Lendrem C, et al. Purines for rapid identification of stroke mimics (PRISM): study protocol for a diagnostic accuracy study. *Diagn Progn Res*. 2021;5(1):1-19.
33. Williams CJ, Wilson KJ, Wilson N. A comparison of prior elicitation aggregation using the classical method and SHELF. *J R Stat Soc Ser A*. 2021;184(3):920-940.

## SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

**How to cite this article:** Wilson KJ, Williamson SF, Allen AJ, Williams CJ, Hellyer TP, Lendrem BC. Bayesian sample size determination for diagnostic accuracy studies. *Statistics in Medicine*. 2022;41(15):2908-2922. doi:10.1002/sim.9393

## APPENDIX . DERIVATION OF THE EQUIVALENCE OF THE BETA-BINOMIAL AND MULTINOMIAL-DIRICHLET APPROACHES

Based on the multinomial-Dirichlet prior structure, the sensitivity and specificity are independent in the prior. To see this, we note that the parameter vector  $\boldsymbol{\gamma}$  can be reordered so that  $\tilde{\boldsymbol{\gamma}} = (\gamma_{1,1}, \gamma_{2,1}, \gamma_{1,2}, \gamma_{2,2})'$ , which has a Dirichlet prior with reordered parameter vector  $\tilde{\boldsymbol{\alpha}}$ . Then, by the properties of neutrality and aggregation of the Dirichlet distribution, we see that  $(\gamma_{1,1}/[\gamma_{1,1} + \gamma_{2,1}], \gamma_{2,1}/[\gamma_{1,1} + \gamma_{2,1}])$  and  $(\gamma_{1,2}/[\gamma_{1,2} + \gamma_{2,2}], \gamma_{2,2}/[\gamma_{1,2} + \gamma_{2,2}])$  are mutually independent. Hence  $\gamma_{1,1}/[\gamma_{1,1} + \gamma_{2,1}]$  and  $\gamma_{2,2}/[\gamma_{1,2} + \gamma_{2,2}]$  are independent. Now, since  $\boldsymbol{\gamma}$  has a Dirichlet prior distribution, this means that  $\gamma_{ij}\gamma_{\Sigma} \sim \text{Gamma}(\alpha_{ij}, 1)$ , where  $\gamma_{\Sigma} = \sum_{i=1}^2 \sum_{j=1}^2 \gamma_{ij}$ . We can re-express the sensitivity as  $\gamma_{1,1}/(\gamma_{1,1} + \gamma_{2,1}) = \gamma_{1,1}\gamma_{\Sigma}/(\gamma_{1,1}\gamma_{\Sigma} + \gamma_{2,1}\gamma_{\Sigma})$  and, by the properties of the gamma distribution, we have  $\lambda \sim \text{Beta}(\alpha_{1,1}, \alpha_{2,1})$ . By similar reasoning,  $\theta \sim \text{Beta}(\alpha_{2,2}, \alpha_{1,2})$  for specificity.

When the number of individuals in each cell of the contingency table is observed, the posterior distributions are

$$\boldsymbol{\gamma}|\mathbf{n} \sim \text{Dir}(\alpha_{1,1} + n_{1,1}, \alpha_{1,2} + n_{1,2}, \alpha_{2,1} + n_{2,1}, \alpha_{2,2} + n_{2,2}),$$

$$\lambda|\mathbf{n} \sim \text{Beta}(\alpha_{1,1} + n_{1,1}, \alpha_{2,1} + n_{2,1}),$$

$$\theta|\mathbf{n} \sim \text{Beta}(\alpha_{2,2} + n_{2,2}, \alpha_{1,2} + n_{1,2}).$$

Setting  $\alpha_{1,1} = a_{\lambda}$ ,  $\alpha_{1,2} = b_{\theta}$ ,  $\alpha_{2,1} = b_{\lambda}$ , and  $\alpha_{2,2} = a_{\theta}$ , we see that the two approaches are equivalent for the sensitivity and specificity.