Published in final edited form as:

Biometrics. 2011 September ; 67(3): 1163–1170. doi:10.1111/j.1541-0420.2011.01561.x.

Bayesian Design of Non-Inferiority Trials for Medical Devices Using Historical Data

Ming-Hui Chen^{*}, Joseph G. Ibrahim[†], Peter Lam[§], Alan Yu[◇], and Yuanye Zhang^{*}

^{*} Department of Statistics, University of Connecticut, 215 Glenbrook Road, Storrs, CT 06269

[†] Department of Biostatistics, University of North Carolina, McGavran-Greenberg Hall, Chapel Hill, NC 27599

[§] Boston Scientific Corporation, 100 Boston Scientific Way M1, Marlborough, MA 01752

[◇] Independent Consultant, 15620 76th PI N, Maple Grove, MN 55311

Summary

We develop a new Bayesian approach of sample size determination (SSD) for the design of non-inferiority clinical trials. We extend the fitting and sampling priors of Wang and Gelfand (2002) to Bayesian SSD with a focus on controlling the type I error and power. Historical data are incorporated via a hierarchical modeling approach as well as the power prior approach of Ibrahim and Chen (2000). Various properties of the proposed Bayesian SSD methodology are examined and a simulation-based computational algorithm is developed. The proposed methodology is applied to the design of a non-inferiority medical device clinical trial with historical data from previous trials.

Keywords

Fitting prior; Hierarchical model; Power prior; Sampling prior; Simulation

1. Introduction

Recently, the FDA released “Guidance for the Use of Bayesian Statistics in Medical Device Clinical Trials” (February 5, 2010, www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm071072.htm). This document provides guidance on statistical aspects of the design and analysis of Bayesian clinical trials for medical devices. It lays out detailed guidance on the determination of sample size in a Bayesian clinical trial. This document also provides guidance on the evaluation of the operating characteristics of a Bayesian clinical trial design. Specifically, the evaluation of a Bayesian clinical trial design should include type I error (probability of erroneously approving an ineffective or unsafe device), type II error (probability of erroneously disapproving a safe and effective device), and power (the converse of type II error: the probability of appropriately approving a safe and effective device).

Sample size determination (SSD) is a crucial aspect of clinical trial design. In this paper, we are particularly interested in the design and analysis of non-inferiority trials. There is a vast literature on the frequentist methods of SSD in various non-inferiority trials, which includes,

^{*} mhchen@stat.uconn.edu.

for example, D'Agostino Sr. et al. (2003), Hung et al. (2003), Rothmann et al. (2003), Hung et al. (2005), Kieser and Friede (2007), and Fleming (2008). The literature on Bayesian SSD has been growing recently due to recent advances in Bayesian computation and Markov chain Monte Carlo sampling. Joseph et al. (1995), Lindley (1997), Rubin and Stern (1998), Katsis and Toman (1999), and Inoue et al. (2005) are the Bayesian SSD articles cited in the FDA 2010 Guidance. An early review of Bayesian SSD is given in Adcock (1997). The most recent work includes Rahme and Joseph (1998), Simon (1999), Wang and Gelfand (2002), De Santis (2007), and M'Lan et al. (2006, 2008). The existing literature on Bayesian SSD primarily focuses on simple normal, one or two sample binomial problems, standard normal linear regression, and generalized linear models. Although the literature on Bayesian SSD discusses a variety of performance criteria, the widely used ones include the Bayes factor (Weiss, 1997), the average posterior variance criterion (APVC) (see, for example, Wang and Gelfand, 2002), the average coverage criterion (ACC), the average length criterion (ALC), the worst outcome criterion (WOC) (e.g., Joseph et al., 1995 and Joseph and Bélisle, 1997), and the approach based on the range of equivalence (see, for instance, Spiegelhalter et al., 2004) for superiority/non-inferiority trials. Other criteria have also been considered in the literature, including Lindley (1997), Pham-Gia (1997), Lam and Lam (1997), and M'Lan et al. (2006, 2008). However, most of the aforementioned Bayesian articles do not directly address design and analysis of non-inferiority trials except for Spiegelhalter et al. (2004).

The rest of the paper is organized as follows. In Section 2, we present the design of a non-inferiority trial with two treatment arms for evaluating the performance of a new generation of medical devices in order to motivate the methodology developed in this paper. The availability of historical data from first generation medical devices is also discussed in detail. In Section 3, we propose a general framework of Bayesian SSD for designing a non-inferiority trial. Section 4 provides a detailed development of the incorporation of historical data via the hierarchical modeling approach as well as the power prior formulation. The posterior distribution is discussed and a simulation-based computational algorithm is developed in Section 5. In Section 6, we apply the proposed methodology to sample size determination of the non-inferiority medical device trial discussed in Section 2. The proposed Bayesian SSD method is compared to frequentist SSD methods. We show that Bayesian SSD yields a substantial reduction in the sample size compared to a frequentist design. We conclude the paper with some discussion and extension of the proposed Bayesian SSD method in Section 7.

2. Design of A Non-Inferiority Trial with Two Treatment Arms for Medical Devices

We consider designing a clinical trial to evaluate the performance of a new generation of drug-eluting stent (DES) ("test device") with a non-inferiority comparison to the first generation of DES ("control device"). Thus, the trial has two arms: test device and control device. The primary endpoint is the 12-month Target Lesion Failure (TLF) (binary) composite endpoint, which is an ischemia-driven revascularization of the target lesion (TLR), myocardial infarction (MI) (Q-wave and non-Q-wave) related to the target vessel, or (cardiac) death related to the target vessel. The secondary endpoint is the 9-month in-segment percent diameter stenosis (%DS) (continuous), which is the percentage of narrowing in the coronary artery caused by the plaque. Let $\mathbf{y}_t^{(n_t)} = (y_{t1}, y_{t2}, \dots, y_{tn_t})'$ and $\mathbf{y}_c^{(n_c)} = (y_{c1}, y_{c2}, \dots, y_{cn_c})'$ be the data corresponding to the test device and the control device, respectively, collected from this trial. Let $n = n_t + n_c$ denote the total sample size. Also, we write $\mathbf{y}^{(n)} = ((\mathbf{y}_t^{(n_t)})', (\mathbf{y}_c^{(n_c)})')'$. We assume that the ratio of two sample sizes, $r = \frac{n_c}{n_t}$ is fixed.

Thus, $n_t = \frac{n}{1+r}$ and $n_c = \frac{rn}{1+r}$. We choose r to be small, for example, $r = 1/4$, so that $n_t > n_c$. The goal of the trial is to show that the test device is non-inferior to the control device.

We assume that $\mathbf{y}_t^{(n)}$ and $\mathbf{y}_c^{(n)}$ are two independent random samples. For the primary endpoint, we assume that y_{ti} (y_{ci}) follows a Bernoulli distribution $\text{Ber}(p_t)$ ($\text{Ber}(p_c)$). Let $\mu_t = \log\left(\frac{p_t}{1-p_t}\right)$ and $\mu_c = \log\left(\frac{p_c}{1-p_c}\right)$. For the secondary endpoint, we assume that $y_{ti} \sim N(\mu_t, \sigma^2)$ and $y_{ci} \sim N(\mu_c, \sigma^2)$ independently. Let $\theta = (\mu_t, \mu_c)$ for the primary endpoint and $\theta = (\mu_t, \mu_c, \sigma^2)$ for the secondary endpoint. Then, the joint distribution of $\mathbf{y}^{(n)}$ for the primary endpoint is given by

$$f(\mathbf{y}^{(n)}|\theta) = \prod_{i=1}^{n_t} \frac{\exp(y_{ti}\mu_t)}{1+\exp(\mu_t)} \times \prod_{i=1}^{n_c} \frac{\exp(y_{ci}\mu_c)}{1+\exp(\mu_c)}, \quad (2.1)$$

and the joint distribution of $\mathbf{y}^{(n)}$ for the secondary endpoint is given by

$$f(\mathbf{y}^{(n)}|\theta) = \prod_{i=1}^{n_t} \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{1}{2\sigma^2}(y_{ti} - \mu_t)^2\right\} \times \prod_{i=1}^{n_c} \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{1}{2\sigma^2}(y_{ci} - \mu_c)^2\right\}. \quad (2.2)$$

The design parameter is the difference between μ_t and μ_c , namely, $\mu_t - \mu_c$, and the hypotheses for non-inferiority testing are

$$H_0: \mu_t - \mu_c \geq \delta \text{ versus } H_1: \mu_t - \mu_c < \delta, \quad (2.3)$$

where δ is a prespecified non-inferiority margin. The trial is successful if H_1 is accepted.

Historical data are available from two previous trials on the first generation of DES. The first trial conducted in 2002 evaluated the safety and effectiveness of the slow release paclitaxel-eluting stent for treatment of de novo coronary artery lesions. The second trial conducted in 2004 expanded on the first trial, studied more complex de novo lesions, and involved multiple overlapping stents and smaller and larger diameter stents. Our historical data based on lesion size matched criteria are subsets of the data published in Stone et al. (2004, 2005). A summary of the historical data is given in Table 1. In Table 1, SD stands for standard deviation.

In the next two sections, we will develop the general methodology for Bayesian SSD and elicit priors via historical data.

3. The General Methodology

We first develop a new but general method to determine Bayesian sample size for a non-inferiority trial. Denote the data associated with a sample size of n by $\mathbf{y}^{(n)}$ and let θ be the vector of all the model parameters. Then, the joint distribution of $\mathbf{y}^{(n)}$ and θ is written as $f(\mathbf{y}^{(n)}|\theta)\pi(\theta)$, where $\pi(\theta)$ denotes the prior distribution. Let $h(\theta)$ be a scalar function that measures the “true” size of the treatment effect. Then, let δ denote the non-inferiority margin. Similar to Hung et al. (2003), we assume that the hypotheses for non-inferiority testing can be formulated as follows:

$$H_0: h(\theta) \geq \delta \text{ versus } H_1: h(\theta) < \delta. \quad (3.1)$$

Consequently, we let Θ_0 and Θ_1 denote the parameter spaces corresponding to H_0 and H_1 . For the hypotheses given in (2.3), $h(\theta) = \mu_t - \mu_c$; $\Theta_0 = \{\theta = (\mu_t, \mu_c): \mu_t - \mu_c \geq \delta\}$ and $\Theta_1 = \{\theta: \mu_t - \mu_c < \delta\}$ for the primary endpoint; and $\Theta_0 = \{\theta = (\mu_t, \mu_c, \sigma^2): \mu_t - \mu_c \geq \delta, \sigma^2 > 0\}$ and $\Theta_1 = \{\theta: \mu_t - \mu_c < \delta, \sigma^2 > 0\}$ for the secondary endpoint.

Following Wang and Gelfand (2002), let $\pi^{(s)}(\theta)$ denote the sampling prior and also let $\pi^{(f)}(\theta)$ denote the fitting prior. The sampling prior, which captures a certain specified portion of the parameter space in achieving a certain level of performance in SSD, is used to generate the data while the fitting prior is used to fit the model once the data is obtained. We note that $\pi^{(f)}(\theta)$ may be improper as long as the resulting posterior, $\pi^{(f)}(\theta|\mathbf{y}^{(n)}) \propto f(\mathbf{y}^{(n)}|\theta)\pi^{(f)}(\theta)$, is proper. Further we let $f^{(s)}(\mathbf{y}^{(n)})$ denote the marginal distribution that is induced from the sampling prior. Now, we introduce the key quantity

$$\beta_s^{(n)} = E_s \left[1\{P(h(\theta) < \delta | \mathbf{y}^{(n)}, \pi^{(f)}) \geq \gamma\} \right], \quad (3.2)$$

where the indicator function $1\{A\}$ is 1 if A is true and 0 otherwise, $\gamma > 0$ is a prespecified quantity, the probability is computed with respect to the posterior distribution given the data $\mathbf{y}^{(n)}$ and the fitting prior $\pi^{(f)}(\theta)$, and the expectation is taken with respect to the marginal distribution of $\mathbf{y}^{(n)}$ under the sampling prior $\pi^{(s)}(\theta)$.

Now, we propose a new Bayesian SSD algorithm as follows. Let $\bar{\Theta}_0$ and $\bar{\Theta}_1$ denote the closures of Θ_0 and Θ_1 . Let $\pi_0^{(s)}(\theta)$ denote a “sampling prior” with support $\Theta_B = \bar{\Theta}_0 \cap \bar{\Theta}_1$. Also let $\pi_1^{(s)}(\theta)$ denote a “sampling prior” with support $\Theta_1^* \subset \Theta_1$. For given $\alpha_0 > 0$ and $\alpha_1 > 0$, we compute

$$n_{\alpha_0} = \min\{n: \beta_{s0}^{(n)} \leq \alpha_0\} \text{ and } n_{\alpha_1} = \min\{n: \beta_{s1}^{(n)} \geq 1 - \alpha_1\}, \quad (3.3)$$

where $\beta_{s0}^{(n)}$ and $\beta_{s1}^{(n)}$ given in (3.2) corresponding to $\pi^{(s)} = \pi_0^{(s)}$ and $\pi^{(s)} = \pi_1^{(s)}$ are the Bayesian type I error and power, respectively. Then, the Bayesian sample size is given by $n_B = \max\{n_{\alpha_0}, n_{\alpha_1}\}$. According to the FDA 2010 Guidance, we choose $\gamma \geq 0.95$. Common choices of α_0 and α_1 include $\alpha_0 = 0.05$ and $\alpha_1 = 0.20$ so that the Bayesian sample size n_B guarantees that the type I error rate is less than or equal to 0.05 and the power is at least 0.80. In addition, for a given sample size n_B , the operating characteristic curve can be constructed by varying Θ_1^* inside of Θ_1 . If $h(\theta)$ is a monotonic function of the distance between Θ_1^* and Θ_B , then the further Θ_1^* is away from Θ_B , the higher the power will be.

A simple illustration: i.i.d. normal case

Suppose $y_1, y_2, \dots, y_n$ are i.i.d. $N(\theta, \tau^{-1})$, where τ is a known precision parameter. Suppose the hypotheses for non-inferiority testing are formulated as follows: $H_0: \theta \geq \delta$ versus $H_1: \theta < \delta$. We specify an improper uniform fitting prior for θ , i.e., $\pi^{(f)}(\theta) \propto 1$. In addition, we specify two point mass sampling priors for θ such that $\pi_0^{(s)}(\theta) = 1$ if $\theta = \delta$ and $\pi_1^{(s)}(\theta) = 1$ if $\theta = 0$. After some algebra, we can show that (i) a necessary condition for achieving a type I error rate of α_0 is $1 - \gamma \leq \alpha_0$ and (ii) if $1 - \gamma \leq \alpha_0$, the Bayesian sample size is the smallest integer n_B

satisfying $n_B \geq \frac{1}{\tau^2} \left[\Phi^{-1}(1 - \alpha_1) + \Phi^{-1}(\gamma) \right]^2$, where Φ denotes the $N(0, 1)$ cumulative distribution function. It is interesting to note that for this simple case, $\beta_0^{(n)} \leq \alpha_0$ always holds for all n when $1 - \gamma \leq \alpha_0$. We also note that the Bayesian sample size n_B is identical to the classical sample size formulation for a one-sided alternative hypothesis when $\alpha_0 = 1 - \gamma$.

4. The Incorporation of Historical Data in Bayesian SSD

Historical data are often available *only* for the control medical device. Now suppose that there are K historical datasets for the control device, denoted by $\mathbf{y}_{c0k} = (y_{c0k1}, \dots, y_{c0kn_{0k}})'$ for $k = 1, \dots, K$. Let $\mathbf{y}_{c0} = (\mathbf{y}_{c01}', \dots, \mathbf{y}_{c0K}')'$ denote all K historical datasets. We develop two approaches, namely, the hierarchical prior and the power prior, to incorporate the historical data $\mathbf{y}_{c0}$.

4.1 Hierarchical Priors

Under the hierarchical Bernoulli/normal model, we assume that $\mathbf{y}_{c0k}$ follows the same model given in either (2.1) or (2.2). Let $\boldsymbol{\theta}_0 = (\mu_{c01}, \dots, \mu_{c0K})'$ (or $\boldsymbol{\theta}_0 = (\mu_{c01}, \dots, \mu_{c0K}, \sigma^2)'$) for the primary (or secondary) endpoint. Then, the joint distribution of $\mathbf{y}_{c0}$ is given by

$$f(\mathbf{y}_{c0}|\boldsymbol{\theta}_0) = \prod_{k=1}^K \prod_{i=1}^{n_{0k}} \frac{\exp(y_{c0ki}\mu_{c0k})}{1 + \exp(\mu_{c0k})} \text{ for the primary endpoint and}$$

$$f(\mathbf{y}_{c0}|\boldsymbol{\theta}_0) = \prod_{k=1}^K \prod_{i=1}^{n_{0k}} \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{1}{2\sigma^2}(y_{c0ki} - \mu_{c0k})^2\right\} \text{ for the secondary endpoint. We further assume } \mu_c \sim N(\mu_{c0}, \tau^2), \text{ where } \tau^2 > 0, \text{ and independently } \mu_{c0k} \sim N(\mu_{c0}, \tau^2) \text{ for } k = 1, \dots, K.$$

Let $\boldsymbol{\theta}^* = (\mu_t, \mu_c, \boldsymbol{\theta}_0, \mu_{c0}, \tau^2)'$. Then, the hierarchical prior for $\boldsymbol{\theta}^*$ is given by

$$\pi(\boldsymbol{\theta}^*|\mathbf{y}_{c0}) \propto f(\mathbf{y}_{c0}|\boldsymbol{\theta}_0) \varphi(\mu_c|\mu_{c0}, \tau^2) \prod_{k=1}^K \varphi(\mu_{c0k}|\mu_{c0}, \tau^2) \pi_0(\mu_t, \sigma^2, \mu_{c0}, \tau^2), \quad (4.1)$$

where $\varphi(\cdot|\mu_{c0}, \tau^2)$ denotes the probability density function (pdf) of a $N(\mu_{c0}, \tau^2)$ distribution. In (4.1), $\pi_0(\mu_t, \sigma^2, \mu_{c0}, \tau^2)$ is the initial prior, which is specified as

$\pi_0(\mu_t, \sigma^2, \mu_{c0}, \tau^2) \propto \frac{1}{\sigma^2} (\tau^2)^{-(\xi_0+1)} \exp(-\eta_0/\tau^2)$, where $\xi_0 > 0$ and $\eta_0 > 0$ are two prespecified hyperparameters. The joint prior in (4.1) is improper since an improper uniform prior is assumed for μ_t and the historical data are borrowed for μ_c and σ^2 via the hierarchical model. Finally, the fitting prior is obtained after integrating out $\mu_{c01}, \dots, \mu_{c0K}, \mu_{c0}$, and τ^2 from (4.1). Specifically, we have

$$\pi^{(f)}(\boldsymbol{\theta}|\mathbf{y}_{c0}) \propto \int \pi(\boldsymbol{\theta}^*|\mathbf{y}_{c0}) d\mu_{c01} \dots d\mu_{c0K} d\mu_{c0} d\tau^2. \quad (4.2)$$

To specify the sampling prior $\pi^{(s)}(\boldsymbol{\theta})$, we assume μ_t, μ_c , and σ^2 are independent and then specify point mass priors for μ_t and μ_c and use the historical data to specify the sampling prior for σ^2 . Specifically, we take

$$\pi^{(s)}(\boldsymbol{\theta}) = \pi^{(s)}(\mu_t) \pi^{(s)}(\mu_c) \text{ or } \pi^{(s)}(\boldsymbol{\theta}) = \pi^{(s)}(\mu_t) \pi^{(s)}(\mu_c) \pi^{(s)}(\sigma^2), \quad (4.3)$$

where

$\pi^{(s)}(\sigma^2) \propto \int f(\mathbf{y}_{c0}|\theta_0) \left[\prod_{k=1}^K \varphi(\mu_{c0k}|\mu_{c0}, \tau^2) \times \frac{1}{\sigma^2} (\tau^2)^{-(\xi_0^{(s)}+1)} \exp(-\eta_0^{(s)}/\tau^2) \right] d\mu_{c01} \dots d\mu_{c0K} d\mu_{c0} d\tau^2$, and $\xi_0^{(s)} > 0$ and $\eta_0^{(s)} > 0$ are prespecified hyperparameters, which may be different from (ξ_0, η_0) . As discussed in Section 3, the sampling prior must be proper. We can show that under very mild conditions, the sampling prior $\pi^{(s)}(\sigma^2)$ is proper.

We note that under the normal model, the hierarchical prior (4.1) for θ^* reduces to

$$\pi(\theta^*|\mathbf{y}_{c0}) \propto (\sigma^2)^{-\frac{1}{2}\sum_{k=1}^K n_{0k}-1} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{k=1}^K \left[n_{0k}(\mu_{c0k} - \bar{y}_{c0k})^2 + (n_{0k} - 1)S_{0k}^2 \right] \right\} \\ \times (\tau^2)^{-[\xi_0 + (K+1)/2+1]} \exp \left\{ -\frac{1}{\tau^2} \left[\eta_0 + \frac{1}{2}(\mu_c - \mu_{c0})^2 + \frac{1}{2} \sum_{k=1}^K (\mu_{c0k} - \mu_{c0})^2 \right] \right\}, \quad (4.4)$$

where $\bar{y}_{c0k} = (1/n_{0k}) \sum_{i=1}^{n_{0k}} y_{c0ki}$ and $S_{0k}^2 = [1/(n_{0k} - 1)] \sum_{i=1}^{n_{0k}} (y_{c0ki} - \bar{y}_{c0k})^2$ for $k = 1, \dots, K$. Thus, the fitting prior and the sampling prior depend only on the sufficient statistics $\{(\bar{y}_{0k}, S_{0k}^2), k = 1, \dots, K\}$ from the historical data.

4.2 Power Priors

We extend the power priors of Ibrahim and Chen (2000) to build the prior distribution for μ_c or (μ_c, σ^2) when multiple historical datasets are available. For the primary endpoint, we consider the following *normalized power prior* for μ_c given multiple historical data $\mathbf{y}_{c0}$,

$$\pi(\mu_c|\mathbf{y}_{c0}, \mathbf{a}_0) = \frac{1}{C(\mathbf{a}_0)} \prod_{k=1}^K \left[\prod_{i=1}^{n_{0k}} \frac{\exp(y_{c0ki}\mu_c)}{1 + \exp(\mu_c)} \right]^{a_{0k}} \pi_0(\mu_c), \quad (4.5)$$

where $\mathbf{a}_0 = (a_{01}, \dots, a_{0K})'$, $0 \leq a_{0k} \leq 1$ for $k = 1, 2, \dots, K$, $\pi_0(\mu_c)$ is an initial prior, and

$C(\mathbf{a}_0) = \int_0^\infty \prod_{k=1}^K \left[\prod_{i=1}^{n_{0k}} \frac{\exp(y_{c0ki}\mu_c)}{1 + \exp(\mu_c)} \right]^{a_{0k}} \pi_0(\mu_c) d\mu_c$. When $\pi_0(\mu_c) \propto 1$, (4.5) reduces to

$$\pi(\mu_c|\mathbf{y}_{c0}, \mathbf{a}_0) = \frac{\exp \left\{ \mu_c \sum_{k=1}^K a_{0k} n_{0k} \bar{y}_{c0k} \right\}}{B \left(\sum_{k=1}^K a_{0k} n_{0k} \bar{y}_{c0k}, \sum_{k=1}^K a_{0k} n_{0k} (1 - \bar{y}_{c0k}) \right) [1 + \exp(\mu_c)]^{n_0(\mathbf{a}_0)}},$$

where $B(., .)$ denotes the complete beta function, $n_0(\mathbf{a}_0) = \sum_{k=1}^K a_{0k} n_{0k}$, and $\bar{y}_{c0k}$ is defined in (4.4). For the secondary endpoint, the normalized power prior for μ_c and σ^2 is given by

$$\pi(\mu_c, \sigma^2|\mathbf{y}_{c0}, \mathbf{a}_0) = \frac{1}{C(\mathbf{a}_0)} \prod_{k=1}^K \left[\prod_{i=1}^{n_{0k}} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left\{ -\frac{1}{2\sigma^2} (y_{c0ki} - \mu_c)^2 \right\} \right]^{a_{0k}} \pi_0(\mu_c, \sigma^2), \quad (4.6)$$

where $\pi_0(\mu_c, \sigma^2)$ is an initial prior and $C(\mathbf{a}_0)$ is the normalizing constant, which is similar to the one in (4.5). Let $\bar{y}_{c0}(\mathbf{a}_0) = \frac{\sum_{k=1}^K a_{0k} n_{0k} \bar{y}_{c0k}}{n_0(\mathbf{a}_0)}$ and

$S_0^2(\mathbf{a}_0) = \sum_{k=1}^K a_{0k} n_{0k} (\bar{y}_{c0k} - \bar{y}_{c0}(\mathbf{a}_0))^2 + \sum_{k=1}^K a_{0k} (n_{0k} - 1) S_{0k}^2$, where $\bar{y}_{c0k}$ and S_{0k}^2 are defined in (4.4). When $\pi_0(\mu_c, \sigma^2) \propto 1/\sigma^2$, (4.6) reduces to

$$\pi(\mu_c, \sigma^2 | \mathbf{y}_{c0}, \mathbf{a}_0) = \left(\frac{n_0(\mathbf{a}_0)}{2\pi\sigma^2} \right)^{1/2} \exp \left\{ -\frac{n_0(\mathbf{a}_0)}{2\sigma^2} [\mu_c - \bar{y}_{c0}(\mathbf{a}_0)]^2 \right\} \\ \times \left\{ \frac{S_0^2(\mathbf{a}_0)/2}{\Gamma([n_0(\mathbf{a}_0)+1]/2)} (\sigma^2)^{-\frac{1}{2}[n_0(\mathbf{a}_0)+1]} \exp \left\{ -S_0^2(\mathbf{a}_0)/(2\sigma^2) \right\} \right\}.$$
 To complete the specification of the power prior, we assume that the a_{0k} 's are independent and distributed as $a_{0k} \sim \text{beta}(b_{01}, b_{02})$, where $b_{01} > 0$ and $b_{02} > 0$ are prespecified hyperparameters. We mention that the normalized power prior is also considered by Duan et al. (2006), Neuenschwander et al. (2009), and Hobbs et al. (2009).

Using (4.5) or (4.6), the fitting prior of θ is of the form

$$\pi^{(f)}(\theta | \mathbf{y}_{c0}) \propto \int \left[\pi(\tilde{\theta} | \mathbf{y}_{c0}, \mathbf{a}_0) \prod_{k=1}^K a_{0k}^{b_{01}-1} (1 - a_{0k})^{b_{02}-1} \right] d\mathbf{a}_0 \pi_0(\mu_c), \quad (4.7)$$

where $\pi(\tilde{\theta} | \mathbf{y}_{c0}, \mathbf{a}_0) = \pi(\mu_c | \mathbf{y}_{c0}, \mathbf{a}_0)$ defined in (4.5) and $\tilde{\theta} = \mu_c$ for the primary end-point, and $\pi(\theta | \mathbf{y}_{c0}, \mathbf{a}_0) = \pi(\mu_c, \sigma^2 | \mathbf{y}_{c0}, \mathbf{a}_0)$ defined in (4.6) and $\theta = (\mu_c, \sigma^2)$ for the secondary endpoint. Similar to the hierarchical prior, the sampling prior $\pi^{(s)}(\theta)$ under the normalized power prior is specified as follows: $\pi^{(s)}(\theta) = \pi^{(s)}(\mu_t) \pi^{(s)}(\mu_c)$ or $\pi^{(s)}(\theta) = \pi^{(s)}(\mu_t) \pi^{(s)}(\mu_c) \pi^{(s)}(\sigma^2)$, where $\pi^{(s)}(\mu_t)$ and $\pi^{(s)}(\mu_c)$ are two prespecified proper priors,

$$\pi^{(s)}(\sigma^2) \propto \int \pi(\tilde{\theta} | \mathbf{y}_{c0}, \mathbf{a}_{0s}) \pi_0^{(s)}(\sigma^2) d\mu_c, \quad (4.8)$$

$\mathbf{a}_{0s}$ is prespecified and $\pi_0^{(s)}(\sigma^2)$ may be an improper initial prior such as $\pi_0^{(s)}(\sigma^2) \propto 1/\sigma^2$.

In (4.5) or (4.6), the parameter a_{0k} controls the influence of the k^{th} historical dataset on $\pi(\tilde{\theta} | \mathbf{y}_{c0}, \mathbf{a}_0)$. The parameter a_{0k} can be interpreted as a relative precision parameter for the k^{th} historical dataset. One of the main roles of $\mathbf{a}_0$ is that it controls the heaviness of the tails of the prior for μ_c (or (μ_c, σ^2)). As all of the a_{0k} 's become smaller, the tails of (4.5) or (4.6) become heavier. When $a_{0k} = 1$ for all k with probability 1, (4.7) corresponds to the update of $\pi_0(\theta)$ using Bayes theorem based on the historical data. When $\mathbf{a}_0 = \mathbf{0}$ with probability 1, then the power prior does not depend on the historical data. That is, $\mathbf{a}_0 = \mathbf{0}$ is equivalent to a prior specification with no incorporation of historical data. Thus, the a_{0k} 's control the influence of the multiple historical datasets on the current study. Such control is important in cases where there is heterogeneity among the historical studies, or heterogeneity between the historical and current studies, or when the sample sizes of the historical and current studies are quite different.

We note that the use of historical data via the power priors for Bayesian sample size determination is also considered by De Santis (2007). We also note that $\mathbf{a}_0$ may be considered to be fixed instead of random. For ease of exposition, we consider the primary endpoint. When $\mathbf{a}_0$ is fixed, the fitting prior of $\theta = (\mu_t, \mu_c)'$ is of the form $\pi^{(f)}(\theta | \mathbf{y}_{c0}) \propto \pi(\mu_c | \mathbf{y}_{c0}, \mathbf{a}_0) \pi_0(\mu_t)$, where $\pi(\mu_c | \mathbf{y}_{c0}, \mathbf{a}_0)$ is given by (4.5), $\pi_0(\mu_t)$ is an initial prior for μ_t , and $\mathbf{a}_0$ is fixed. When $\mathbf{a}_0$ is fixed, we know exactly how much historical data are incorporated in the new trial, and in addition, there is a theoretical connection between the power prior formulation and the hierarchical prior specification as established in Chen and Ibrahim (2006). De Santis (2006) also provides some useful comments on the fixed- $\mathbf{a}_0$ case as well as on power priors for the exponential family. On the other hand, when $\mathbf{a}_0$ is random, the

amount of incorporation of historical data is determined by the data and hence not prespecified by the data analyst.

5. Posteriors and Computations

For ease of exposition, we only consider the primary endpoint. Instead of directly sampling from $\pi^{(f)}(\theta|\mathbf{y}^{(n)}, \mathbf{y}_{c0}) \propto f(\mathbf{y}^{(n)}|\theta)\pi^{(f)}(\theta|\mathbf{y}_{c0})$, where $f(\mathbf{y}^{(n)}|\theta)$ is given by (2.1) and $\pi^{(f)}(\theta|\mathbf{y}_{c0})$ is defined in (4.2) or (4.7), we consider the augmented fitting posterior distribution parameters θ^* , where $\theta^* = (\mu_t, \mu_c, \mu_{c01}, \dots, \mu_{c0K}, \mu_{c0}, \tau^2)'$ for the hierarchical prior and $\theta^* = (\mu_t, \mu_c, \mathbf{a}_0)$ for the normalized power prior. Then, the augmented fitting posterior distribution of θ^* is given by $\pi^{(f)}(\theta^*|\mathbf{y}^{(n)}, \mathbf{y}_{c0}) \propto f(\mathbf{y}^{(n)}|\theta) \pi(\theta^*|\mathbf{y}_{c0})$, where $\pi(\theta^*|\mathbf{y}_{c0})$ is defined in (4.1) under the

hierarchical prior, and $\pi(\theta^*|\mathbf{y}_{c0}) \propto \pi(\mu_c|\mathbf{y}_{c0}, \mathbf{a}_0) \left[\prod_{k=1}^K a_{0k}^{b_{01}-1} (1 - a_{0k})^{b_{02}-1} \right] \pi_0(\mu_t)$ with $\pi(\mu_c|\mathbf{y}_{c0}, \mathbf{a}_0)$ defined in (4.5) under the normalized power prior. Although the posterior distribution $\pi^{(f)}(\theta^*|\mathbf{y}^{(n)}, \mathbf{y}_{c0})$ is analytically intractable, sampling from this distribution via the Gibbs sampler is quite straightforward, because the conditional posterior distributions of the components of θ^* (except for $\mathbf{a}_0$) are either known distributions or log-concave. For $\mathbf{a}_0$, we use the localized Metropolis algorithm discussed in Chen et al. (2000) to sample from its conditional posterior distribution.

Let $\{\theta^{*(m)}, m = 1, 2, \dots, M\}$ denote a Gibbs sample from the augmented fitting posterior distribution $\pi^{(f)}(\theta^*|\mathbf{y}^{(n)}, \mathbf{y}_{c0})$. As θ is a subvector of θ^* , let $\theta^{(m)}$ denote the corresponding components of $\theta^{*(m)}$ from the m^{th} Gibbs iteration. Then, it is easy to show that $\{\theta^{(m)}, m = 1, 2, \dots, M\}$ is a Gibbs sample from the fitting posterior distribution $\pi^{(f)}(\theta|\mathbf{y}^{(n)}, \mathbf{y}_{c0})$. Using this Gibbs sample, a Monte Carlo estimate of $P(h(\theta) < \delta|\mathbf{y}^{(n)}, \pi^{(f)})$ is given by

$$\hat{P}_f = \frac{1}{M} \sum_{m=1}^M 1\{h(\theta^{(m)}) < \delta\}. \quad (5.1)$$

To compute $\beta_s^{(n)}$ in (3.2), we propose the following computational algorithm: Step 0: Specify n_t, n_c, δ, γ , and N ; Step 1: Generate $\theta \sim \pi^{(s)}(\theta)$; Step 2: Generate $\mathbf{y}^{(n)} \sim f(\mathbf{y}^{(n)}|\theta)$; Step 3: Run the Gibbs sampler to generate a Gibbs sample $\{\theta^{(m)}, m = 1, 2, \dots, M\}$ of size M from the fitting posterior distribution $\pi^{(f)}(\theta|\mathbf{y}^{(n)}, \mathbf{y}_{c0})$; Step 4: Compute $\hat{P}_f$ via (5.1); Step 5: Check whether $\hat{P}_f \geq \gamma$; Step 6: Repeat Steps 1–5 N times; and Step 7: Compute the proportion of $\{\hat{P}_f \sim \gamma\}$ in these N runs, which gives an estimate of $\beta_s^{(n)}$.

6. Applications to Medical Device Trials

We apply the proposed methodology in designing the non-inferiority clinical trial for medical devices discussed in Section 2. We use the historical datasets given in Table 1 to construct our priors in Bayesian SSD. We set $\gamma = 0.95$, which implies a target type I error of 0.05. We notice that the same γ value was also used in Allocco et al. (2010). In all of the computations below, $N = 10,000$ and $M = 20,000$ were used.

Bayesian SSD for TLF

For the primary endpoint, the margin was set to be $\delta = \text{logit}(4.1\%) = \log\left\{\frac{0.041}{1-0.041}\right\}$. We took $(\xi_0, \eta_0) = (0.01, 0.01)$ or $(\xi_0, \eta_0) = (0.001, 0.001)$ for the initial prior of γ in the fitting prior (4.2), $\pi_0(\mu_c) \propto 1$ and $b_{01} = b_{02} = 1$ for the initial priors of μ_c and a_{0k} in (4.7). We computed the powers at $\mu_t = \mu_c$ and the type I error at $\frac{\exp(\mu_t)}{1+\exp(\mu_t)} = \frac{\exp(\mu_c)}{1+\exp(\mu_c)} + \frac{\exp(\delta)}{1+\exp(\delta)}$. In other words, we convert μ_j back to p_j in the Bernoulli case. In the sampling prior (4.3), we assumed a point mass prior

at $\mu_c = \text{logit}(9.2\%)$ for $\pi^{(s)}(\mu_c)$, where 9.2% was the pooled proportion for the two historical control datasets, and a point mass prior at $\mu_t = \mu_c$ or $\mu_t = \text{logit}\left[\frac{\exp(\mu_c)}{1+\exp(\mu_c)} + \frac{\exp(\delta)}{1+\exp(\delta)}\right]$ for $\pi^{(s)}(\mu_t)$. We first computed the powers and the type I errors for various sample sizes based on the proposed Bayesian SSD without the incorporation of historical data. Table 2 shows the results. Table 2 also presents the powers of the two frequentist methods, namely, the z -test with unpooled variances and the score test (Farrington and Manning, 1990) for non-inferiority trials. For both frequentist methods, the target type I error was 0.05. In all calculations, the margin $\delta = 0.041$, $p_c = 9.2\%$, and a 3:1 sample size ratio were used. PASS 2008 (Hintze, 2008) was used for computing the powers for the two frequentist SSD methods. We see from Table 2 that the proposed Bayesian SSD without incorporation of historical data gives very similar powers compared to the score test for the frequentist SSD, while the type I errors of the Bayesian SSD are controlled at or below 5%. Both the score test and Bayesian SSD yield slightly higher powers than the z -test. In order to achieve 80% power, the z -test requires a total sample size of 1636 with $n_t = 1227$ and $n_c = 409$.

Table 2 also shows the powers and the type I errors of the Bayesian SSD procedure with hierarchical priors and power priors with fixed and random $\mathbf{a}_0$. The hierarchical prior with $(\xi_0, \eta_0) = (0.001, 0.001)$ leads to higher powers than the one with $(\xi_0, \eta_0) = (0.01, 0.01)$. In addition, the powers based on the power prior with $\mathbf{a}_0$ random are comparable to those based on the hierarchical prior with $(\xi_0, \eta_0) = (0.001, 0.001)$ and the power prior with $\mathbf{a}_0$ fixed at $\mathbf{a}_0 = (0.3, 0.3)$. These results imply that the power prior with random $\mathbf{a}_0$ and the hierarchical prior with $(\xi_0, \eta_0) = (0.001, 0.001)$ borrow approximately 30% of the historical data. With incorporation of the historical data, a sample size of $(n_t, n_c) = (810, 270)$ achieves 80% power. However, based on the frequentist SSD or the Bayesian SSD without incorporation of historical data, a sample size of 1480 with $n_t = 1110$ and $n_c = 370$ is required to achieve 80% power. Thus, the Bayesian SSD with incorporation of historical data leads to a substantial reduction in the sample size.

Bayesian SSD for %DS

For the secondary endpoint, the margin was set to be $\delta = 0.20$. We compute the power at $\mu_t = \mu_c$ and the type I error at $\mu_t = \mu_c + \delta$. In the sampling prior (4.3), we assume a point mass prior at $\mu_c = 3.15$ for $\pi^{(s)}(\mu_c)$ and a point mass prior $\mu_t = \mu_c$ or $\mu_t = \mu_c + \delta$ for $\pi^{(s)}(\mu_t)$. PASS 2008 (Hintze, 2008) was used to compute the powers of the frequentist SSD based on the pooled SD = 0.607. In the Bayesian SSD procedure which does not use any historical data, we used the same pooled SD for σ in generating the data. For the hierarchical prior, we took $(\xi_0, \eta_0) = (0.01, 0.01)$ or $(\xi_0, \eta_0) = (0.001, 0.001)$ for the initial prior of τ in the fitting prior (4.2) and $\xi_0 = 0.01$ and $\eta_0 = 0.01$ for the initial prior of τ in the sampling prior (4.3). For the power priors, we used (4.8) with $\mathbf{a}_{0s} = (0.05, 0.05)$ for the sampling prior $\pi^{(s)}(\sigma^2)$. Using the same sampling prior, we also computed the powers and type I errors with a fixed $\mathbf{a}_0 = (0.08, 0.08)$ in the fitting prior. The results are shown in Table 3.

Similar to TLF, the Bayesian SSD procedure with no incorporation of historical data yields similar powers, with the type I errors controlled at the 5% level, and the hierarchical prior with $(\xi_0, \eta) = (0.001, 0.001)$ yields higher powers than the one with $(\xi_0, \eta) = (0.01, 0.01)$. From Table 3, we also see that the power prior with random $\mathbf{a}_0$ leads to slightly higher powers than the hierarchical prior with $(\xi_0, \eta_0) = (0.001, 0.001)$, and the powers based on the power prior with random $\mathbf{a}_0$ are comparable to the power prior with a fixed $\mathbf{a}_0 = (0.08, 0.08)$. These results imply that the hierarchical prior borrows less than 8% of historical data, while the power prior with random $\mathbf{a}_0$ borrows about 8% of the historical data. Similar to TLF, the Bayesian SSD with incorporation of historical data again leads to a substantial reduction in the sample size compared to the frequentist design.

7. Discussion

In this paper, we have developed a general methodology of Bayesian SSD, which is particularly suitable for designing a non-inferiority clinical trial. We have discussed two types of priors, namely, the hierarchical prior and the normalized power prior, to incorporate historical data. We have shown that Bayesian SSD leads to a substantial reduction in the sample size compared to frequentist SSD. One unique feature of the proposed Bayesian SSD methodology is that we use the historical data only from the control device but not from the test device. This feature is desirable, since for the test device, historical data are often not available. Although we primarily focus on the Bernoulli and normal models in this paper, our methodology is applicable to other models in the exponential family. In addition, the proposed methodology can also be extended to generalized linear models (GLMs). The computational algorithm given in Section 5 for these two extensions is basically the same. However, there may be two potential complications. First, a closed-form expression of the normalized power prior under GLMs may not be available. Therefore, an efficient Markov chain Monte Carlo sampling algorithm needs to be developed to sample from the fitting posterior distribution in Step 3 of the computational algorithm in Section 5. Second, the determination of the non-inferiority margin may be more difficult for some GLMs than the situation without covariates. For example, for binomial regression models, the non-inferiority margin based on the difference in two proportions may not be easily converted to the margin on the regression coefficient corresponding to the treatment effect. However, this may not be an issue for other GLMs such as the normal linear regression model.

The proposed Bayesian SSD works best if the historical data from the control device are compatible to the data from the current trial. However, the target type I error and power may not be well maintained when the data from the historical and current trials are not compatible. For non-inferiority trials, we have empirically observed that (i) the type I errors are controlled but the powers are lower when the true proportions or means in the control devices from the current trial are greater than those in the historical data; and (ii) the type I errors tend to be larger, but the powers tend to be higher when the true proportions or the true means for the control devices in the current trial are less than those in the historical data. For illustrative purposes, we consider $n = 1200$ with $n_t = 900$ and $n_c = 300$ for the primary endpoint TLF and $n = 280$ with $n_t = 210$ and $n_c = 70$ for the secondary endpoint %DS. For %DS, if a point mass sampling prior at $\mu_c = 3.10$ is assumed and $\gamma = 0.95$, the powers and type I errors are 0.836 and 0.047 for the hierarchical priors with $(\xi_0, \eta_0) = (0.01, 0.01)$ and 0.936 and 0.049 for the power prior with random a_0 ; and if a point mass sampling prior at $\mu_c = 3.20$ is assumed, the powers and type I errors are 0.792 and 0.030 for the hierarchical priors with $(\xi_0, \eta_0) = (0.01, 0.01)$ and 0.815 and 0.004 for the power prior with random a_0 . In all cases, the type I errors are still controlled at 0.05. However, for TLF, the type I error is not controlled as shown in Table 4. Specifically, if a point mass sampling prior at $\mu_c = \text{logit}(8.0\%)$ is assumed, the type I errors are 0.068 for the hierarchical priors with $(\xi_0, \eta_0) = (0.01, 0.01)$ and 0.07 for the power prior with $a_{0k} \sim \text{beta}(1, 1)$ and $\gamma = 0.95$. There are two approaches for resolving this type I error problem. One approach is to change the initial prior $\text{beta}(b_{01}, b_{02})$ for a_{0k} in (4.7) to down-weight the historical control data as suggested by an anonymous Associate Editor. Another approach is to increase the value of γ , which is recommended in the FDA 2010 Guidance. As shown in Table 4 for TLF, if a point mass sampling prior at $\mu_c = \text{logit}(8.0\%)$ is assumed, the type I error decreases in b_{02} when b_{01} is fixed at 1. When $(b_{01}, b_{02}) = (1, 10)$, which gives an initial prior weight of 10% to the historical control data, the type I error is 0.053. Also, we see from Table 4 that for a fixed initial prior $\text{beta}(b_{01}, b_{02})$, the type I error decreases in γ . In particular, when $(b_{01}, b_{02}) = (1, 1)$ and $\gamma = 0.97$, the type I error is 0.041 if a point mass sampling prior at $\mu_c = \text{logit}(8.0\%)$ is assumed. A combination of these two approaches is also quite effective in controlling the

type I error while maintaining good power as shown in Table 4. Further methodological approaches for controlling the type I error are currently under investigation.

Finally, we briefly discuss how to determine whether the trial is successful after it is completed. The computational algorithm developed in Section 5 can still be used for this purpose. Specifically, the following algorithm can be used to determine the outcome of the trial: Step 0: Use the same γ and the same fitting prior specified at the design stage; Step 1: Obtain the data $y^{(n)}$ at the completion of the trial; Step 2: Run the Gibbs sampler to generate a Gibbs sample $\{\theta^{(m)}, m = 1, 2, \dots, M\}$ of size M from the fitting posterior distribution $\pi^{(f)}(\theta|y^{(n)}, y_{c0})$; Step 3: Compute $\hat{P}_f$ via (5.1); and Step 4: Declare a success of the trial if $\hat{P}_f \geq \gamma$.

Acknowledgments

Dr. Chen is a statistical consultant for Boston Scientific Corporation. This research was partially supported by Boston Scientific Corporation. The conclusions in this paper are entirely those of the authors and do not necessarily represent the views of Boston Scientific Corporation. No conflict of interest exists among the authors. In addition, Dr. Chen and Dr. Ibrahim's research was partially supported by NIH grants #GM 70335 and #CA 74015.

References

- Adcock CJ. Sample size determination: a review. *The Statistician*. 1997; 46:261–283.
- Allocco, DJ.; Cannon, LA.; Britt, A.; Heil, JE.; Nersesov, A.; Wehrenberg, S.; Dawkins, KD.; Kereiakes, DJ. A prospective evaluation of the safety and efficacy of the TAXUS Element paclitaxel-eluting coronary stent system for the treatment of de novo coronary artery lesions: design and statistical methods of the PERSEUS clinical program; *Trials*. 2010. p. 1
<http://www.trialsjournal.com/content/11/1/1>
- Chen MH, Ibrahim JG. The relationship between the power prior and hierarchical models. *Bayesian Analysis*. 2006; 1:551–574.
- Chen, M-H.; Shao, Q-M.; Ibrahim, JG. *Monte Carlo Methods in Bayesian Computation*. New York: Springer-Verlag; 2000.
- D'Agostino RB Sr, Massaro JM, Sullivan LM. Non-inferiority trials: design concepts and issues – the encounters of academic consultants in statistics. *Statistics in Medicine*. 2003; 22:169–186. [PubMed: 12520555]
- De Santis F. Using historical data for Bayesian sample size determination. *Journal of the Royal Statistical Society, Series A*. 2007; 170:95–113.
- De Santis F. Power priors and their use in clinical trials. *The American Statistician*. 2006; 60:122–129.
- Duan Y, Ye K, Smith EP. Evaluating water quality using power priors to incorporate historical information. *Environmetrics*. 2006; 17:95–106.
- Farrington CP, Manning G. Test statistics and sample size formulae for comparative binomial trials with null hypothesis of non-zero risk difference or non-unity relative risk. *Statistics in Medicine*. 1990; 9:1447–1454. [PubMed: 2281232]
- Fleming TR. Current issues in non-inferiority trials. *Statistics in Medicine*. 2008; 27:317–332. [PubMed: 17340597]
- Hintze, J. PASS 2008. NCSS, LLC; Kaysville, Utah, USA: 2008. www.ncss.com
- Hobbs, BP.; Carlin, BP.; Mandrekar, S.; Sargent, D. Technical Report 2009-017. Division of Biostatistics, University of Minnesota; 2009. Hierarchical commensurate prior models for adaptive incorporation of historical information in clinical trials.
- Hung HMJ, Wang SJ, O'Neill RT. A regulatory perspective on choice of margin and statistical inference issue in non-inferiority trials. *Biometrical Journal*. 2005; 47:28–36. [PubMed: 16395994]
- Hung HMJ, Wang SJ, Tsong Y, Lawrence J, O'Neill RT. Some fundamental issues with non-inferiority testing in active controlled trials. *Statistics in Medicine*. 2003; 22:213–225. [PubMed: 12520558]

- Ibrahim JG, Chen MH. Power prior distributions for regression models. *Statistical Science*. 2000; 15:46–60.
- Inoue LYT, Berry DA, Parmigiani G. Relationship between Bayesian and frequentist sample size determination. *The American Statistician*. 2005; 59:79–87.
- Joseph L, Bélisle P. Bayesian sample size determination for normal means and differences between normal means. *The Statistician*. 1997; 46:209–226.
- Joseph L, Wolfson DB, Du Berger R. Sample size calculations for binomial proportions via highest posterior density intervals. *The Statistician: Journal of the Institute of Statisticians*. 1995; 44:143–154.
- Katsis A, Toman B. Bayesian sample size calculations for binomial experiments. *Journal of Statistical Planning and Inference*. 1999; 81:349–362.
- Kieser M, Friede T. Planning and analysis of three-arm non-inferiority trials with binary endpoints. *Statistics in Medicine*. 2007; 26:253–273. [PubMed: 16685703]
- Lam Y, Lam CV. Bayesian double-sampling plans with normal distributions. *The Statistician*. 1997; 46:193–207.
- Lindley DV. The choice of sample size. *The Statistician*. 1997; 46:129–138.
- M'Lan CE, Joseph L, Wolfson DB. Bayesian sample size determination for binomial proportions. *Bayesian Analysis*. 2008; 3:269–296.
- M'Lan CE, Joseph L, Wolfson DB. Bayesian sample size determination for case-control studies. *Journal of the American Statistical Association*. 2006; 101:760–772.
- Neuenschwander B, Branson M, Spiegelhalter DJ. A note on the power prior. *Statistics in Medicine*. 2009; 28:3562–3566. [PubMed: 19735071]
- Pham-Gia T. On Bayesian analysis, Bayesian decision theory and the sample size problem. *The Statistician*. 1997; 46:139–144.
- Rahme E, Joseph L. Exact sample size determination for binomial experiments. *Journal of Statistical Planning and Inference*. 1998; 66:83–93.
- Rothmann M, Li N, Chen G, Chi GYH, Temple R, Tsou HH. Design and analysis of non-inferiority mortality trials in oncology. *Statistics in Medicine*. 2003; 22:239–264. [PubMed: 12520560]
- Rubin DB, Stern HS. Sample size determination using posterior predictive distributions. *Sankhyâ, Series B*. 1998; 60:161–175.
- Simon R. Bayesian design and analysis of active control clinical trials. *Biometrics*. 1999; 55:484–487. [PubMed: 11318204]
- Spiegelhalter, DJ.; Abrams, KR.; Myles, JP. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. New York: Wiley; 2004.
- Stone GW, Ellis SG, Cannon L, et al. Comparison of a polymer-based paclitaxel-eluting stent with a bare metal stent in patients with complex coronary artery disease: a randomized controlled trial. *Journal of the American Medical Association*. 2005; 294:1215–1223. [PubMed: 16160130]
- Stone GW, Ellis SG, Cox DA, et al. A polymer-based, paclitaxel-eluting stent in patients with coronary artery disease. *The New England Journal of Medicine*. 2004; 350:221–231. [PubMed: 14724301]
- Wang F, Gelfand AE. A simulation-based approach to Bayesian sample size determination for performance under a given model and for separating models. *Statistical Science*. 2002; 17:193–208.
- Weiss R. Bayesian sample size calculations for hypothesis testing. *The Statistician*. 1997; 46:185–191.

Table 1

Historical Data

	12-Month TLF	log(9-month %DS)
	% TLF (# of failure/n_{0k})	mean $\pm$ SD (n_{0k})
Historical Trial 1	8.2% (44/535)	3.0891 $\pm$ 0.6315 (242)
Historical Trial 2	10.9% (33/304)	3.1849 $\pm$ 0.5811 (263)

Table 2

Powers and Type I Errors for 12-Month TLF

Total Sample Size		1000	1080	1200	1280	1480
n_t		750	810	900	960	1110
		250	270	300	320	370
Frequentist SSD						
z Test (Unpooled)	Power	0.617	0.646	0.685	0.710	0.764
	Score Test	0.672	0.699	0.736	0.758	0.807
Bayesian SSD						
No Borrowing	Power	0.648	0.676	0.718	0.738	0.800
	Type I Error	0.049	0.048	0.048	0.050	0.044
$a_0 = (0, 0)$						
Hierarchical Prior	Power	0.796	0.820	0.841	0.863	0.894
	Type I Error	0.044	0.045	0.044	0.049	0.048
$(\xi_0, \eta_0) = (0.01, 0.01)$						
Hierarchical Prior	Power	0.839	0.860	0.882	0.900	0.922
	Type I Error	0.038	0.042	0.039	0.040	0.041
$(\xi_0, \eta_0) = (0.001, 0.001)$						
Power Prior	Power	0.840	0.856	0.884	0.892	0.923
	Type I Error	0.030	0.027	0.028	0.030	0.032
Fixed $a_0 = (0.3, 0.3)$						
Power Prior	Power	0.843	0.878	0.897	0.902	0.914
	Type I Error	0.038	0.031	0.029	0.036	0.039
Random a_0						

Table 3

Powers and Type I Errors for 9-Month %DS

Total Sample Size		200	240	260	280	308
n_t		150	180	195	210	231
	n_c	50	60	65	70	77
Frequentist SSD	Power	0.639	0.709	0.739	0.767	0.801
Bayesian SSD						
No Borrowing	Power	0.644	0.699	0.747	0.769	0.805
$a_0 = (0, 0)$	Type I Error	0.051	0.049	0.051	0.050	0.048
Hierarchical Prior	Power	0.710	0.773	0.800	0.820	0.847
$(\zeta_0, \eta_0) = (0.01, 0.01)$	Type I Error	0.037	0.038	0.040	0.038	0.039
Hierarchical Prior	Power	0.791	0.837	0.871	0.877	0.899
$(\zeta_0, \eta_0) = (0.001, 0.001)$	Type I Error	0.023	0.024	0.025	0.028	0.027
Power Prior	Power	0.812	0.864	0.880	0.899	0.918
Fixed $a_0 = (0.08, 0.08)$	Type I Error	0.022	0.023	0.026	0.027	0.028
Power Prior	Power	0.805	0.857	0.878	0.893	0.913
Random a_0	Type I Error	0.013	0.014	0.017	0.015	0.015

Table 4

Powers and Type I Errors under Three Sampling Priors for 12-Month TLF with $(n_t, n_c) = (900, 300)$

Fitting Prior		Point Mass Sampling Prior at $\mu_c = \text{logit}(p_c^*)$					
		$p_c^* = 8.0\%$		$p_c^* = 9.2\%$		$p_c^* = 10.0\%$	
	γ	Power	Type I Error	Power	Type I Error	Power	Type I Error
Hierarchical Prior							
$(\xi_0, \eta_0) = (0.01, 0.01)$	0.95	0.894	0.068	0.841	0.044	0.788	0.032
	0.96	0.880	0.058	0.816	0.037	0.757	0.027
	0.97	0.854	0.046	0.782	0.027	0.714	0.020
Power Prior with $a_{0k} \sim \text{beta}(b_{01}, b_{02})$ in (4.7)							
$(b_{01}, b_{02}) = (1, 1)$	0.95	0.945	0.070	0.882	0.039	0.799	0.034
	0.95	0.916	0.061	0.832	0.033	0.760	0.026
	0.95	0.868	0.053	0.791	0.038	0.728	0.032
$(b_{01}, b_{02}) = (1, 1)$	0.96	0.935	0.055	0.880	0.022	0.765	0.026
	0.97	0.917	0.041	0.848	0.015	0.719	0.009
$(b_{01}, b_{02}) = (1, 5)$	0.96	0.899	0.047	0.803	0.027	0.722	0.021