

Statistica Sinica Preprint No: SS-2023-0208

Title	Optimal Priors for the Discounting Parameter of the Normalized Power Prior
Manuscript ID	SS-2023-0208
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202023.0208
Complete List of Authors	Yueqi Shen, Luiz M. Carvalho, Matthew A. Psioda and Joseph G. Ibrahim
Corresponding Authors	Yueqi Shen
E-mails	ys137@live.unc.edu
Notice: Accepted version subject to English editing.	

OPTIMAL PRIORS FOR THE DISCOUNTING PARAMETER OF THE NORMALIZED POWER PRIOR

Yueqi Shen¹, Luiz M. Carvalho², Matthew A. Psioda³ and Joseph G. Ibrahim¹

¹*University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, U.S.A.*

²*School of Applied Mathematics, Getulio Vargas Foundation, Rio de Janeiro, Brazil*

³*GSK, Philadelphia, Pennsylvania, U.S.A*

Abstract:

The power prior is a popular class of informative priors for incorporating information from historical data. It involves raising the likelihood for the historical data to a power, which acts as discounting parameter. When the discounting parameter is modelled as random, the normalized power prior is recommended. In this work, we prove that the marginal posterior for the discounting parameter for generalized linear models converges to a point mass at zero if there is any discrepancy between the historical and current data, and that it does not converge to a point mass at one when they are fully compatible. In addition, we explore the construction of optimal priors for the discounting parameter in a normalized power prior. In particular, we are interested in achieving the dual objectives of encouraging borrowing when the historical and current data are compatible and limiting borrowing when they are in conflict. We propose intuitive procedures for eliciting the shape parameters of a beta prior for the discounting parameter based on two minimization criteria, the Kullback-Leibler divergence and the mean squared error. Based on the proposed criteria, the optimal priors derived are often quite different from commonly used priors such as the uniform prior.

Key words and phrases: Bayesian analysis, Clinical trial, Normalized power prior, Power prior.

1. Introduction

The power prior (Ibrahim and Chen, 2000) is a popular class of informative priors that allow the incorporation of historical data through a tempering of the likelihood. It is constructed by raising the historical data likelihood to a power a_0 , where $0 \leq a_0 \leq 1$. The discounting parameter a_0 can be fixed or modelled as random. When it is modelled as random and estimated jointly with other parameters of interest, the normalized power prior (NPP) (Duan et al., 2006) is recommended as it appropriately accounts for the normalizing function necessary for forming the correct joint prior distribution (Neuenschwander et al., 2009). Many extensions of the power prior and the normalized power prior have been developed. Banbeta et al. (2019) develop the dependent and robust dependent normalized power priors which allow dependent discounting parameters for multiple historical datasets. When the historical data model contains only a subset of covariates currently of interest and the historical information may not be equally informative for all parameters in the current analysis, Boonstra and Barbaro (2020) propose an extension of the power prior that adaptively combines a prior based upon the historical information with a variance-reducing prior that shrinks parameter values toward zero.

The power prior and the normalized power prior have been shown to have several desirable properties. Ibrahim et al. (2003) show that the power prior defines an optimal class of priors in the sense that it minimizes a convex combination of Kullback-Leibler (KL) divergences between a

distribution based on no incorporation of historical data and a distribution based on completely pooling the historical and current data. Ye et al. (2022) prove that the normalized power prior minimizes the expected weighted KL divergence similar to the one in Ibrahim et al. (2003) with respect to the marginal distribution of the discounting parameter. They also prove that if the prior on a_0 is non-decreasing and if the difference between the sufficient statistics of the historical and current data is negligible from a practical standpoint, the marginal posterior mode of a_0 is close to one. Carvalho and Ibrahim (2021) show that the normalized power prior is always well-defined when the initial prior is proper, and that, viewed as a function of the discounting parameter, the normalizing function is a smooth and strictly convex function. Neelon and O'Malley (2010) show through simulations that for large datasets, the normalized power prior may result in more downweighting of the historical data than desired. Han et al. (2022) point out that the normalizing function might be infinite with improper initial priors on the parameters of interest for a_0 values close to zero, in which case the admissible set of the discounting parameter excludes values close to zero. Pawel et al. (2023) derive the marginal posterior distribution of a_0 when a beta prior is used for a_0 for *i.i.d.* normal and binomial models, and show that, under these model assumptions, the marginal posterior of a_0 does not converge to a point mass at one when the sample size becomes arbitrarily large. In this paper, we provide a formal proof of the limiting behavior of the marginal posterior of a_0 for generalized linear models (GLMs) when the historical and current datasets are fully compatible and when there is discrepancy.

Many empirical Bayes-type approaches have been developed to adaptively determine the dis-

counting parameter. For example, Gravestock and Held (Gravestock et al., 2017; Gravestock and Held, 2019) propose to set a_0 to the value that maximizes the marginal likelihood. Liu (2018) proposes choosing a_0 based on the p-value for testing the compatibility of the current and historical data. Bennett et al. (2021) propose using an equivalence probability weight and a weight based on tail area probabilities to assess the degree of agreement between the historical and current control data for cases with binary outcomes. Pan et al. (2017) propose the calibrated power prior, where a_0 is defined as a function of a congruence measure between the historical and current data. The function which links a_0 and the congruence measure is prespecified and calibrated through simulation. While these empirical Bayes approaches shed light on the choice of a_0 , there has not been any fully Bayesian approach based on an optimal prior on a_0 .

In this work, we first explore the asymptotic properties of the normalized power prior when the historical and current data are fully compatible (i.e., the sufficient statistics of the two datasets are equal) or incompatible (i.e., the sufficient statistics of the two datasets have some non-zero difference). We prove that for GLMs utilizing a normalized power prior, the marginal posterior distribution of a_0 converges to a point mass at zero if there is any discrepancy between the historical and current data. When the historical and current data are fully compatible, the asymptotic distribution of the marginal posterior of a_0 is derived for GLMs; we note that it does not concentrate around one. However, we prove an interesting finding that, for an i.i.d. normal model with finite samples, the marginal posterior of a_0 always has more mass around one when the datasets are fully compatible, compared to the case where there is any discrepancy. Secondly, we propose a novel fully

Bayesian approach to elicit the shape parameters of the beta prior on a_0 based on two optimality criteria, Kullback-Leibler (KL) divergence and mean squared error (MSE). For the first criterion, we propose as optimal the beta prior whose shape parameters result in a minimized weighted average of KL divergences between the marginal posterior for a_0 and user-specified target distributions based on hypothetical scenarios where there is no discrepancy and where there is a maximum tolerable discrepancy. This class of priors on a_0 based on the KL criterion is optimal in the sense that it is the best possible beta prior at balancing the dual objectives of encouraging borrowing when the historical and current data are compatible and limiting borrowing when they are in conflict. For the second criterion, we propose as optimal the beta prior whose shape parameters result in a minimized weighted average of the MSEs based on the posterior mean of the parameter of interest when its hypothetical true value is equal to its estimate using the historical data, or when it differs from its estimate by the maximum tolerable amount. We study the properties of the proposed approaches *via* simulations for the *i.i.d.* normal and Bernoulli cases as well as for the normal linear model. Two real-world case studies of clinical trials with binary outcomes and covariates demonstrate the performance of the optimal priors compared to conventionally used priors on a_0 , such as a uniform prior.

2. Asymptotic Properties of the Normalized Power Prior

Let D denote the current data and D_0 denote the historical data. Let θ denote the model parameters and $L(\theta|D)$ denote a general likelihood function. The power prior (Ibrahim and Chen, 2000) is

formulated as

$$\pi(\theta|D_0, a_0) \propto L(\theta|D_0)^{a_0} \pi_0(\theta),$$

where $0 \leq a_0 \leq 1$ is the discounting parameter which discounts the historical data likelihood, and $\pi_0(\theta)$ is the initial prior for θ . The discounting parameter a_0 can be fixed or modelled as random. Modelling a_0 as random allows researchers to account for uncertainty when discounting historical data and to adaptively learn the appropriate level of borrowing. Duan et al. (2006) propose the *normalized power prior*, given by

$$\pi(\theta, a_0|D_0) = \pi(\theta|D_0, a_0)\pi(a_0) = \frac{L(\theta|D_0)^{a_0} \pi_0(\theta)}{c(a_0)} \pi(a_0), \quad (2.1)$$

where $c(a_0) = \int L(\theta|D_0)^{a_0} \pi_0(\theta) d\theta$ is the normalizing function. The normalized power prior is thus composed of a conditional prior for θ given a_0 and a marginal prior for a_0 .

Ideally, the posterior distribution of a_0 with the normalized power prior would asymptotically concentrate around zero when the historical and current data are in conflict, and around one when they are compatible. In this section, we study the asymptotic properties of the normalized power prior for the exponential family of distributions as well as GLMs. Specifically, we are interested in exploring the asymptotic behaviour of the posterior distribution of a_0 when the historical and current data are incompatible and when they are compatible, respectively.

2.1 Exponential Family

First, we study the asymptotic properties of the normalized power prior for the exponential family of distributions. The density of a random variable Y in the one-parameter exponential family has

the form

$$p(y|\theta) = q(y) \exp(y\theta - b(\theta)), \quad (2.2)$$

where θ is the canonical parameter and $q(\cdot)$ and $b(\cdot)$ are known functions. Suppose $D = (y_1, \dots, y_n)$ is a sample of n *i.i.d.* observations from an exponential family distribution in the form of (2.2).

The likelihood is then given by

$$L(\theta|D) = Q(D) \exp\left(\sum_{i=1}^n y_i\theta - nb(\theta)\right),$$

where $Q(D) = \prod_{i=1}^n q(y_i)$. Suppose $D_0 = (y_{01}, \dots, y_{0n_0})$ is a sample of n_0 *i.i.d.* observations from the same exponential family. The likelihood for the historical data raised to the power a_0 is

$$[L(\theta|D_0)]^{a_0} = Q(D_0)^{a_0} \exp\left(a_0 \left[\sum_{i=1}^{n_0} y_{0i}\theta - n_0b(\theta)\right]\right),$$

where $Q(D_0) = \prod_{i=1}^{n_0} q(y_{0i})$. Using the normalized power prior defined in (2.1), the joint posterior of θ and a_0 is given by

$$\pi(\theta, a_0|D, D_0) \propto L(\theta|D)\pi(\theta, a_0|D_0) = L(\theta|D) \frac{L(\theta|D_0)^{a_0} \pi_0(\theta)}{c(a_0)} \pi(a_0).$$

The marginal posterior of a_0 is given by

$$\pi(a_0|D, D_0) = \int \pi(\theta, a_0|D, D_0) d\theta \propto \int L(\theta|D) \frac{L(\theta|D_0)^{a_0} \pi_0(\theta)}{c(a_0)} \pi(a_0) d\theta. \quad (2.3)$$

With these calculations in place, the question now arises as to what prior should be given to a_0 . One commonly used class of priors on a_0 is the beta distribution (Ibrahim and Chen, 2000). Let

α_0 and β_0 denote the shape parameters of the beta distribution. We first prove that the marginal posterior of a_0 (2.3) with $\pi(a_0) = \text{beta}(\alpha_0, \beta_0)$ converges to a point mass at zero for a fixed, non-zero discrepancy between \bar{y} and \bar{y}_0 .

Theorem 2.1. Suppose y_1, \dots, y_n and y_{01}, \dots, y_{n_0} are independent observations from the same exponential family distribution (2.2). Let $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ and $\bar{y}_0 = \frac{1}{n_0} \sum_{i=1}^{n_0} y_{0i}$. Suppose also that the difference in the estimates of the canonical parameter θ is fixed and equal to δ , i.e., $|\dot{b}^{-1}(\bar{y}) - \dot{b}^{-1}(\bar{y}_0)| = \delta$, and $\frac{n_0}{n} = r$, where $\delta > 0$ is finite and $r > 0$ is a constant, and $\dot{b}(\cdot) = \partial_\theta b(\cdot)$. Then, the marginal posterior of a_0 using the normalized power prior (2.3) with a $\text{beta}(\alpha_0, \beta_0)$ prior on a_0 , where $\alpha > 0$ and $\beta > 0$, converges to a point mass at 0. That is, $\lim_{n \rightarrow \infty} \frac{\int_0^\epsilon \pi(a_0|D, D_0, \alpha_0, \beta_0) da_0}{\int_0^1 \pi(a_0|D, D_0, \alpha_0, \beta_0) da_0} = 1$ for any $\epsilon > 0$.

Proof. See Supplemental Materials S1.2. □

In Theorem 2.1, we fix the ratio of n_0 to n , denoted as r , and let n go to infinity. This is because if n_0 is fixed while n is allowed to go to infinity, then the historical likelihood in the normalized power prior becomes irrelevant for inference on θ , regardless of the prior on a_0 . Therefore, n_0 must go to infinity as well. Theorem 2.1 asserts that the normalized power prior is sensitive to any discrepancy between the sufficient statistics in large samples, as the mass of the marginal distribution of a_0 will concentrate near zero as the sample size increases for any fixed difference δ . Figure 2 in Supplemental Materials S2 shows that the marginal posterior of a_0 converges to a point mass at zero rather quickly as the sample size grows when there is discrepancy.

The natural question to then ask is whether Theorem 2.1 has a sort of converse in that the posterior should concentrate around one under compatibility. We derive the asymptotic marginal posterior distribution of a_0 when $\bar{y} = \bar{y}_0$ and show that it does not converge to a point mass at one.

Corollary 2.1. Suppose y_1, \dots, y_n and y_{01}, \dots, y_{n_0} are independent observations from the same exponential family distribution (2.2). Let $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ and $\bar{y}_0 = \frac{1}{n_0} \sum_{i=1}^{n_0} y_{0i}$. Suppose $\bar{y} = \bar{y}_0$ and $\frac{n_0}{n} = r$ where $r > 0$ is a constant. The marginal posterior of a_0 using the normalized power prior, as specified in (2.3), converges to

$$\tilde{\pi}(a_0|D, D_0) = \frac{\sqrt{\frac{ra_0}{ra_0+1}} \pi(a_0)}{\int_0^1 \sqrt{\frac{ra_0}{ra_0+1}} \pi(a_0) da_0}$$

as $n \rightarrow \infty$.

Proof. See Supplemental Materials S1.3. □

Corollary 2.1 shows that the normalized power prior fails to fully utilize the historical data when the means of the historical data and the current data are equal for a generic, non-degenerate prior on a_0 . It is worth noting that the adjustment to the prior $\sqrt{\frac{ra_0}{ra_0+1}}$ is maximized at $a_0 = 1$. If $\pi(a_0)$ is chosen to be concentrated near one, then the marginal posterior of a_0 may be concentrated near one.

We recognize that, in reality, $\bar{y} = \bar{y}_0$ will occur with probability zero. However, the point of this theorem is to show that even in the most extreme case where we have identical data, the marginal posterior of a_0 with the normalized power prior does not concentrate around one. Relaxing the equality will lead to the same conclusion. In fact, Theorem 2.1 indicates that when the difference

between \bar{y} and \bar{y}_0 is fixed at some $\delta \neq 0$, the posterior of a_0 asymptotically converges to a point mass at zero.

Even though the marginal posterior of a_0 does not concentrate under one when $\bar{y} = \bar{y}_0$, we find that for the i.i.d. normal model, the cumulative density function when $\bar{y} = \bar{y}_0$ is dominated by every other cumulative density function, as demonstrated by the following theorem. This is a novel result which establishes a desirable borrowing property of the normalized power prior for i.i.d. normal data for finite samples.

Theorem 2.2. Suppose $y_1, \dots, y_n \sim N(\theta, \sigma^2)$ and $y_{01}, \dots, y_{n_0} \sim N(\theta, \sigma_0^2)$, where σ^2 , σ_0^2 , n and n_0 are fixed. Let $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ and $\bar{y}_0 = \frac{1}{n_0} \sum_{i=1}^{n_0} y_{0i}$. Let $|\bar{y} - \bar{y}_0| = d$. Let $F_d(a_0)$ denote cumulative density function of the marginal posterior of a_0 using the normalized power prior in (2.3), and let $w_d(a_0) = F_d(a_0) - F_0(a_0)$. The function $w_d(a_0) > 0$ for all $d > 0$ and $0 < a_0 < 1$.

Proof. See Supplemental Materials S1.4. □

Theorem 2.2 shows that for i.i.d. normal data, for any fixed sample sizes, $1 - F_0(a_0)$ will always be greater than $1 - F_d(a_0)$ for any $d > 0$. This means the posterior of a_0 always has more mass around one when the datasets are fully compatible, compared to the case where there is any discrepancy, which is a desirable property.

2.2 Generalized Linear Models

The ability to deal with non *i.i.d.* data and incorporate covariates is crucial to the applicability of the normalized power prior; we thus now extend these results to generalized linear models (GLMs).

We first define the GLM with a canonical link and fixed dispersion parameter. Let y_i denote the response variable and x_i denote a p -dimensional vector of covariates for subject $i = 1, \dots, n$. Let $\beta = (\beta_1, \dots, \beta_p)'$ be a p -dimensional vector of regression coefficients. The GLM with a canonical link is given by

$$p(y_i|x_i, \beta, \phi) = q(y_i, \phi) \exp\{\phi^{-1}[y_i x_i' \beta - b(x_i' \beta)]\}. \quad (2.4)$$

Without loss of generality, we assume $\phi = 1$. Let $D = \{(y_i, x_i), i = 1, \dots, n\} \equiv (n, Y_{n \times 1}, X_{n \times p})$ where $Y = (y_1, \dots, y_n)'$ and $X = (x_1, \dots, x_n)'$. Assuming the y_i 's are (conditionally) independent, the likelihood is given by

$$L(\beta|D) = Q(Y) \exp\left(\sum_{i=1}^n y_i x_i' \beta - \sum_{i=1}^n b(x_i' \beta)\right),$$

where $Q(Y) = \prod_{i=1}^n q(y_i, 1)$. Let $\hat{\beta}$ denote the posterior mode of β obtained by solving $\partial_\beta \log L(\beta|D) = 0$. Let $D_0 = \{(y_{0i}, x_{0i}), i = 1, \dots, n_0\} \equiv (n_0, Y_{0n_0 \times 1}, X_{0n_0 \times p})$ where $Y_0 = (y_{01}, \dots, y_{0n_0})'$ and $X_0 = (x_{01}, \dots, x_{0n_0})'$. Assuming the y_{0i} 's are (conditionally) independent, the historical data likelihood raised to the power a_0 is given by

$$[L(\beta|D_0)]^{a_0} = Q(Y_0)^{a_0} \exp\left(a_0 \left[\sum_{i=1}^{n_0} y_{0i} x_{0i}' \beta - \sum_{i=1}^{n_0} b(x_{0i}' \beta)\right]\right),$$

where $Q(Y_0) = \prod_{i=1}^{n_0} q(y_{0i}, 1)$. Let $c^*(a_0) = \int L(\beta|y_0)^{a_0} \pi_0(\beta) d\beta$. Using the normalized power prior defined in (2.1), the joint posterior of β and a_0 is given by

$$\pi(\beta, a_0|D, D_0) \propto L(\beta|D) \pi(\beta, a_0|D_0) = L(\beta|D) \frac{L(\beta|D_0)^{a_0} \pi_0(\beta)}{c^*(a_0)} \pi(a_0).$$

Let $\hat{\beta}_0$ denote the posterior mode of β obtained by solving $\partial_{\beta} \log \left[\frac{L(\beta|D_0)^{a_0} \pi_0(\beta)}{c^*(a_0)} \right] = 0$. Note that $\hat{\beta}_0$ is independent of a_0 after we take the limit. The marginal posterior of a_0 is given by

$$\pi(a_0|D, D_0) = \int \pi(\beta, a_0|D, D_0) d\beta \propto \int L(\beta|D) \frac{L(\beta|D_0)^{a_0} \pi_0(\beta)}{c^*(a_0)} \pi(a_0) d\beta. \quad (2.5)$$

Now we extend Theorem 2.1 to GLMs.

Theorem 2.3. Suppose X is $n \times p$ of rank p and X_0 is $n_0 \times p$ of rank p . Suppose $\hat{\beta} - \hat{\beta}_0 = \delta$ where $\delta \neq 0$ is a finite vector, and $\frac{n_0}{n} = r$ where $r > 0$ is a constant scalar. Assume $n \left[\frac{\partial^2 \log[L(\beta|D)]}{\partial \beta_i \partial \beta_j} \right]^{-1}$ and $n_0 a_0 \left[\frac{\partial^2 \log[L(\beta|D_0)^{a_0} \pi_0(\beta)]}{\partial \beta_i \partial \beta_j} \right]^{-1}$ do not depend on n and a_0 . Then, the marginal posterior of a_0 using the normalized power prior (2.5) with a beta(α_0, β_0) prior on a_0 , where $\alpha > 0$ and $\beta > 0$, converges to a point mass at zero. That is, $\lim_{n \rightarrow \infty} \frac{\int_0^\epsilon \pi(a_0|D, D_0, \alpha_0, \beta_0) da_0}{\int_0^1 \pi(a_0|D, D_0, \alpha_0, \beta_0) da_0} = 1$ for any $\epsilon > 0$.

Proof. See Supplemental Materials S1.5. □

Theorem 2.3 asserts that the normalized power prior is sensitive to discrepancies in the historical and current data in the presence of covariates. The mass of the marginal distribution of a_0 will concentrate near zero as the sample size increases for any fixed discrepancy between the historical and current data, assuming $\frac{1}{n} X'X$ and $\frac{1}{n_0} X_0'X_0$ are fixed, i.e., $n \left[\frac{\partial^2 \log[L(\beta|D)]}{\partial \beta_i \partial \beta_j} \right]^{-1}$ and $n_0 a_0 \left[\frac{\partial^2 \log[L(\beta|D_0)^{a_0} \pi_0(\beta)]}{\partial \beta_i \partial \beta_j} \right]^{-1}$ do not depend on n and a_0 .

Next, we derive the asymptotic marginal posterior distribution of a_0 when the sufficient statistics and covariate (design) matrices of the historical and current data equal.

Corollary 2.2. Suppose X is $n \times p$ of rank p and X_0 is $n_0 \times p$ of rank p . Let $Y = (y_1, \dots, y_n)'$ and $Y_0 = (y_{01}, \dots, y_{0n_0})'$. Consider the GLM in (2.4). If $n = n_0$, $X = X_0$, and $X'Y = X_0'Y_0$, then

the marginal posterior of a_0 using the normalized power prior, as specified in (2.5), converges to

$$\tilde{\pi}(a_0|X, Y, X_0, Y_0) = \frac{\left(\frac{a_0}{a_0+1}\right)^{\frac{p}{2}} \pi(a_0)}{\int_0^1 \left(\frac{a_0}{a_0+1}\right)^{\frac{p}{2}} \pi(a_0) da_0},$$

as $n \rightarrow \infty$.

Proof. See Supplemental Materials S1.6. □

Corollary 2.2 states that the marginal posterior of a_0 using the normalized power prior does not converge to a point mass at one when the sufficient statistics and the covariates of the historical and current data are equal. We also observe that as p approaches infinity, the marginal posterior of a_0 specified above converges to a point mass at one. The form of the asymptotic marginal posterior of a_0 suggests that the normalized power prior may be sensitive to overfitting when the historical and current datasets are compatible. Figure 1 in Supplemental Materials S2 presents numerical results corroborating Corollaries 2.1 and 2.2. Therein we present histograms of posterior samples of a_0 that faithfully recapitulate the theoretical density functions.

In Theorem 2.4 we also relax the previous result by deriving the asymptotic marginal posterior distribution of a_0 assuming only that the sufficient statistics of the historical and current data are equal. This means that the covariate matrices need not be equal so long as the sufficient statistics $X'Y$ and X'_0Y_0 are, increasing the applicability of the result.

Theorem 2.4. Suppose X is $n \times p$ of rank p and X_0 is $n_0 \times p$ of rank p . Let $Y = (y_1, \dots, y_n)'$ and $Y_0 = (y_{01}, \dots, y_{0n_0})'$. Consider the GLM in (2.4), where $\frac{n_0}{n} = r$ and $r > 0$ is a constant. If

$X'Y = X'_0Y_0$ and $X \neq X_0$, then the marginal posterior of a_0 using the normalized power prior, as specified in (2.5), is asymptotically proportional to

$$\pi(a_0) \cdot \frac{|\hat{\Sigma}_g|^{1/2}}{|\tilde{\Sigma}_k|^{1/2}} \exp \left\{ -n[g_n(\hat{\beta}) - k_n(\tilde{\beta})] \right\},$$

where the definitions of $g_n(\beta)$, $k_n(\beta)$ and $\frac{|\hat{\Sigma}_g|^{1/2}}{|\tilde{\Sigma}_k|^{1/2}}$ can be found in Supplemental Materials S1.7.

Proof. See Supplemental Materials S1.7. □

Corollary 2.2 and Theorem 2.4 show that, for GLMs, the marginal posterior of a_0 using the normalized power prior does not converge to a point mass at one when the sufficient statistics of the historical and current data are equal. From Theorems 2.1-2.4, we conclude that, asymptotically, the normalized power prior is sensitive to discrepancies between the historical and current data, but cannot fully utilize the historical information when there are no discrepancies. However, we show that the posterior of a_0 always has the most mass around one when the datasets are fully compatible for finite i.i.d. normal observations.

We highlight the differences between the theorems above and the results presented in Pawel et al. (2023). Pawel et al. (2023) derive the marginal posterior distribution of a_0 when a beta prior is used for a_0 for *i.i.d.* normal and binomial models and show through graphical approaches that the distribution shifts toward zero as the standard error of the current data converges to zero, when $\hat{\theta} \neq \hat{\theta}_0$, and the standard error of the historical data is fixed. We prove this phenomenon much more generally and analytically for the exponential family of distributions as well as GLMs.

3. Optimal Beta Priors for a_0

3.1 Kullback-Leibler Divergence Criterion

In this section, we propose a prior based on minimizing the KL divergence of the marginal posterior of a_0 to two reference distributions. This resulting prior is optimal in the sense that it is the best possible beta prior at balancing the dual objectives of encouraging borrowing when the historical and current data are compatible and limiting borrowing when they are in conflict.

Let \bar{y}_0 denote the mean of the historical data and \bar{y} denote the mean of the hypothetical current data. Let $\pi_1(a_0) \equiv \text{beta}(c, 1)$ ($c \gg 1$ is fixed) and $\pi_2(a_0) \equiv \text{beta}(1, c)$. The distributions $\pi_1(a_0)$ and $\pi_2(a_0)$ represent two ideal scenarios, where $\pi_1(a_0)$ is concentrated near one and $\pi_2(a_0)$ is concentrated near zero. The KL-based approach computes the hyperparameters (α_0 and β_0) for the beta prior on a_0 that will minimize a convex combination of two KL divergences; one is the KL divergence between $\pi_1(a_0)$ and the marginal posterior of a_0 when $\bar{y} = \bar{y}_0$, while the other is the KL divergence between $\pi_2(a_0)$ and the marginal posterior of a_0 when there is a user-specified difference between \bar{y} and \bar{y}_0 .

Let $d = \bar{y} - \bar{y}_0$, representing the difference between the means of the hypothetical current data and the historical data. Our approach is centered on a user-specified **maximum tolerable difference** (MTD), d_{MTD} . Let $\pi^*(a_0)$ denote the marginal posterior of a_0 when $d = 0$. Let $\pi_{\text{MTD}}(a_0)$ denote the marginal posterior of a_0 when $d = d_{\text{MTD}}$. For $d = 0$, we want $\pi^*(a_0)$ to resemble $\pi_1(a_0)$ and for $d = d_{\text{MTD}}$, we want $\pi_{\text{MTD}}(a_0)$ to resemble $\pi_2(a_0)$. The distributions

$\pi_1(a_0)$ and $\pi_2(a_0)$ have been chosen to correspond to cases with substantial and little borrowing, respectively. Therefore, our objective is to solve for $\alpha_0 > 0$ and $\beta_0 > 0$ to minimize

$$K(\alpha_0, \beta_0) = wKL(\pi^*(a_0), \pi_1(a_0)) + (1 - w)KL(\pi_{\text{MTD}}(a_0), \pi_2(a_0)).$$

Here $0 < w < 1$ is a scalar and $KL(p, q)$ for distributions P and Q with P as reference is defined as

$$KL(p, q) = \int \log \left(\frac{p(x)}{q(x)} \right) dP(x) = E_p[\log(p)] - E_p[\log(q)].$$

The scalar w weights the two competing objectives. For $w > 0.5$, the objective to encourage borrowing is given more weight, and for $w < 0.5$, the objective to limit borrowing is given more weight. Even though this approach requires specifying w , c and d_{MTD} , these parameters are much more intuitive to choose than the hyperparameters of the prior for a_0 .

Below we demonstrate the simulation results using this method for the *i.i.d.* normal case, the *i.i.d.* Bernoulli case and the normal linear model. We compare the marginal posterior of a_0 using the KL-based optimal prior with that using the uniform prior. For all simulations in this section, we choose $w = 0.5$ so that the two competing objectives are given equal weight. We choose $c = 10$ so that $\pi_1(a_0)$ and $\pi_2(a_0)$ represent cases with substantial and little borrowing, respectively. We vary the choice of d_{MTD} and examine each resulting optimal prior. In practice, we recommend that d_{MTD} be no larger than the treatment effect of the historical study.

3.1.1 Normal *i.i.d.* Case

We assume y_1, \dots, y_n and y_{01}, \dots, y_{0n_0} are *i.i.d.* observations from $N(\mu, \sigma^2)$ where $\sigma^2 = 1$. We choose $\bar{y}_0 = 1.5$ and $n = n_0 = 30$. The objective function $K(\cdot, \cdot)$ is computed using numerical integration and optimization is performed using the `optim()` function in (base) R (R Core Team, 2022).

In Figure 1, the first figure of each row plots the historical and current data likelihoods if the hypothetical degree of conflict is equal to d_{MTD} . For each row of the figure below, the maximum tolerable difference d_{MTD} is chosen to be 0.5, 1 and 1.5, and the corresponding optimal prior is derived for each value of d_{MTD} . For each optimal prior, we vary the observed sample mean, denoted by \bar{y}_{obs} , to evaluate the posterior based on the optimal prior for different observed current data. We use $d_{\text{obs}} = \bar{y}_{\text{obs}} - \bar{y}_0$ to represent the difference between the means of the observed current data and the historical data. For columns 2-4, d_{obs} is chosen to be 0, 1 and 1.5, respectively. Note that the values of d_{MTD} and d_{obs} are relative to the choices of σ^2 , n and n_0 . For example, for larger n , d_{MTD} would need to be decreased to produce a similar plot to Figure 1.

From columns 2-4, we observe that when $d_{\text{MTD}} = 0.5$, very little conflict is tolerated, and the resulting optimal prior does not strongly encourage either borrowing substantially or borrowing little. As d_{MTD} becomes larger, larger conflict is allowed and the optimal prior shifts more towards $\pi_1(a_0)$. We also observe that when $d_{\text{MTD}} = 1$ (the optimal hyperparameters are $\alpha_0 = 1$ and $\beta_0 = 0.4$) and $d_{\text{MTD}} = 1.5$ (the optimal hyperparameters are $\alpha_0 = 2.6$ and $\beta_0 = 0.5$), the marginal posterior of a_0 with the optimal prior more closely mimics the target distribution when

Table 1: Posterior mean (variance) of μ for the normal *i.i.d.* case

	$d_{\text{obs}} = 0$	$d_{\text{obs}}=0.5$	$d_{\text{obs}}=1$	$d_{\text{obs}}=1.5$
$d_{\text{MTD}}=0.5$	1.5 (0.022)	1.85 (0.026)	2.32 (0.036)	2.87 (0.036)
$d_{\text{MTD}}=1$	1.5 (0.019)	1.82 (0.026)	2.36 (0.042)	2.93 (0.035)
$d_{\text{MTD}}=1.5$	1.5 (0.018)	1.78 (0.020)	2.19 (0.040)	2.84 (0.039)

$d_{\text{obs}} = 0$, i.e., the observed current and historical data are fully compatible. As d_{obs} increases, the marginal posterior shifts toward zero. This behaviour is highly desirable as it achieves both goals of encouraging borrowing when the datasets are compatible and limiting borrowing when they are incompatible.

We can compare the marginal posterior of a_0 using the optimal prior with that using a uniform prior in Figure 1. We observe that while the marginal posterior on a_0 with the uniform prior is very responsive to conflict, it does not concentrate around one even when the datasets are fully compatible. We conclude that when d_{MTD} is chosen to be reasonably large, the optimal prior on a_0 achieves a marginal posterior that is close to the target distribution when the datasets are fully compatible, while remaining responsive to conflict in the data.

Table 1 shows the posterior mean and variance of the mean parameter μ for various combinations of d_{MTD} and d_{obs} values corresponding to the scenarios in Figure 1. The posterior mean and variance

of μ with a normalized power prior are computed using the R package *BayesPPD* (Shen et al., 2022). Again, \bar{y}_0 is fixed at 1.5. Since $\bar{y}_{\text{obs}} \geq \bar{y}_0$, within each row, the posterior mean of μ is always smaller than \bar{y}_{obs} due to the incorporation of \bar{y}_0 . We can also compare the results by column. For fixed d_{obs} (or equivalently \bar{y}_{obs}), if more historical information is borrowed, we expect the posterior mean of μ to be smaller. When $d_{\text{obs}} = 0$, the posterior mean stays constant while the variance decreases as d_{MTD} increases. If the maximum tolerable difference is large, more historical information is borrowed, leading to reduced variance. When $d_{\text{obs}} = 0.5$, the posterior of μ decreases as more borrowing occurs when d_{MTD} increases. When $d_{\text{obs}} = 1$ or 1.5, the posterior of μ first increases and then decreases, as d_{MTD} increases. This is a result of two competing phenomena interacting; as d_{MTD} increases, the optimal prior gravitates towards encouraging borrowing; however, since d_{obs} is very large, the marginal posterior of a_0 moves toward zero even though the prior moves toward one. In conclusion, we argue that the posterior estimates of μ with the optimal prior respond in a desirable fashion to changes in the data. In Supplemental Materials S3, we have included numerical experiments that demonstrate the reliability of the optimization scheme.

3.1.2 Bernoulli Model

For the Bernoulli model, we assume y_1, \dots, y_n and y_{01}, \dots, y_{0n_0} are *i.i.d.* observations from a Bernoulli distribution with mean μ . Again, we choose $n = n_0 = 30$ and optimization is performed analogously to the normal case.

The resulting optimal priors and posteriors are shown in Figure 2. For each row of Figure 2

below, the maximum tolerable difference d_{MTD} is chosen to be 0.2, 0.4 and 0.6, and the corresponding optimal prior is derived for each value of d_{MTD} . For each optimal prior, we vary the observed \bar{y}_{obs} to evaluate the performance of the optimal prior for different observed data. For columns 2-4, $d_{\text{obs}} = \bar{y}_{\text{obs}} - \bar{y}_0$ is chosen to be 0, 0.4 and 0.6, respectively. Values of \bar{y}_0 and \bar{y}_{obs} are chosen so that the variance stays constant for different values of d_{MTD} or d_{obs} .

The optimal marginal prior and posterior of a_0 for Bernoulli data are similar to those of the normal model. We observe that when the datasets are perfectly compatible, i.e., $d_{\text{obs}} = 0$, the marginal posterior of a_0 with the optimal prior concentrates around one when d_{MTD} is relatively large. When d_{obs} increases to 0.4 or 0.6, the marginal posterior of a_0 concentrates around zero when d_{MTD} is relatively large. The optimal prior becomes increasingly concentrated near one as d_{MTD} increases. Compared to the marginal posterior with the uniform prior, the optimal prior on a_0 achieves a marginal posterior that closely mimics the target distribution when the datasets are fully compatible, while remaining responsive to conflict in the data.

3.1.3 Normal Linear Model

Suppose y_{01}, \dots, y_{0n_0} are independent observations from the historical data where $y_{0i} \sim N(\beta_0 + \beta_1 x_{0i}, \sigma^2)$ for the i -th observation and x_{0i} is a single covariate. Also suppose y_1, \dots, y_n are independent observations from the current data where $y_j \sim N(\beta_0 + \beta_1 x_j + d_{\text{MTD}}, \sigma^2)$ for the j -th observation and x_j is a single covariate. We vary d_{MTD} to represent different degrees of departure of the intercept of the simulated current data to the intercept of the historical data. We choose

$\beta_0 = 1.5$, $\beta_1 = -1$, $\sigma^2 = 1$ and $n = n_0 = 30$. We choose $d_{\text{MTD}} = 0.1, 0.5$, and 1 and $d_{\text{obs}} = 0, 0.5$, and 1 . The objective function K is computed using Monte Carlo integration and optimization is performed using the `optim()` function in R.

Figure 3 shows the optimal prior and optimal posterior for a_0 as well as the posterior of a_0 with the uniform prior for various d_{MTD} and d_{obs} values. We observe that when the datasets are perfectly compatible, i.e., $d_{\text{obs}} = 0$, the marginal posterior of a_0 with the optimal prior concentrates around one when d_{MTD} is relatively large. When d_{obs} increases to 1, the marginal posterior of a_0 concentrates around zero. The optimal prior becomes increasingly concentrated near one as d_{MTD} increases. Compared to the marginal posterior with the uniform prior, the optimal prior for a_0 achieves a marginal posterior that closely mimics the target distribution when the datasets are fully compatible, while remaining responsive to conflict in the data.

3.2 Mean Squared Error Criterion

In this section, we derive the optimal prior for a_0 based on minimizing the MSE. This prior is optimal in the sense that it minimizes the weighted average of the MSEs of the posterior mean of the parameter of interest when its hypothetical true value is equal to its estimate using the historical data, or when it differs from its estimate by the maximum tolerable amount. Suppose y_1, \dots, y_n and y_{01}, \dots, y_{0n_0} are observations from a distribution with mean parameter μ . Let μ^* denote the true value of μ . Let $\bar{\mu}$ denote the posterior mean of μ using the normalized power prior.

Then, the MSE of $\bar{\mu}$ is

$$\text{MSE}(\mu^*) = \int [\bar{\mu}(y) - \mu^*]^2 p(y|\mu^*) dy.$$

In the regression setting, μ is replaced by the regression coefficients β .

Let \bar{y}_0 denote the mean of the historical data. We aim to find the hyperparameters, α_0 and β_0 , for the beta prior for a_0 that will minimize

$$w\text{MSE}(\mu^* = \bar{y}_0) + (1 - w)\text{MSE}(\mu^* = \bar{y}_0 + d_{\text{MTD}}),$$

where d_{MTD} is the maximum tolerable difference. Again, we use $d_{\text{obs}} = \bar{y}_{\text{obs}} - \bar{y}_0$ to represent the difference between the means of the observed current data and the historical data.

3.2.1 Normal *i.i.d.* Case

We demonstrate the use of this criterion for the normal *i.i.d.* case. Suppose y_1, \dots, y_n and y_{01}, \dots, y_{0n_0} are *i.i.d.* observations from $N(\mu, \sigma^2)$ where $\sigma^2 = 1$ and $n = n_0 = 30$. In this example, we fix μ^* and y_{MTD} at 1.5, and define $d_{\text{MTD}} = \bar{y}_0 - \bar{y}$ and $d_{\text{obs}} = \bar{y}_0 - \bar{y}$. The posterior mean of μ is computed using Monte Carlo integration and optimization is performed using a grid search. The optimal prior, optimal posterior, and the posterior using the uniform prior for a_0 are plotted in Figure 4. When $d_{\text{MTD}} = 0.5$, the optimal prior is unimodal with mode around 0.3. When $d_{\text{MTD}} = 1$, the optimal prior is concentrated near zero. When $d_{\text{MTD}} = 1.5$, the optimal prior is U-shaped and favouring either strong or weak borrowing. When d_{MTD} is small, the algorithm cannot distinguish between the two competing scenarios in the objective function and the resulting

optimal prior concentrates around 0.5. When d_{MTD} is large, the optimal prior will favour the two scenarios equally. For columns 2-4, d_{obs} is chosen to be 0, 1 and 1.5. The marginal posterior using the optimal prior concentrates more around zero as d_{obs} increases for a given d_{MTD} . Comparing Figures 4 and 1, we observe that the optimal prior derived using the MSE criterion is more conservative in the sense that it tends to discourage borrowing than that derived using the KL criterion. Note that one might obtain very different optimal priors using the KL criterion versus the MSE criterion, since these metrics have different objectives.

Table 2 shows the MSE for the optimal prior, $\text{beta}(1, 1)$ (uniform) and $\text{beta}(2, 2)$ as well as the percent reduction of MSE of the optimal prior compared to the uniform prior. We can see that the percent reduction of MSE increases as d_{MTD} increases. The table in Supplemental Materials S4 displays the decomposition of MSE into bias squared and variance for the three choices of priors. When $d_{\text{MTD}} = 0.5$ or 1, the prior discourages borrowing which results in smaller bias and larger variance. When $d_{\text{MTD}} = 1.5$, the model can distinguish easily between the two contrasting objectives, leading to smaller bias and smaller variance. In Supplemental Materials S5, we compare the MSE of the posterior mean based on the normalized power prior using the optimal prior for a_0 and the robust mixture prior (Schmidli et al., 2014) for the case where the data are i.i.d. normal. A key contribution of this paper is to identify the optimal single beta prior for a_0 in an NPP. The comparison in Supplementary Material S5 is designed to help the reader put this notion of optimality in context by comparing the performance of the posterior mean point estimator from several families of priors. Additional simulation results for varying choices of n and n_0 are provided

Table 2: MSE for different prior choices and percent reduction of MSE of the optimal prior compared to the uniform prior

	Optimal Prior	beta(1, 1)	beta(2, 2)	Percent Reduction of MSE, Optimal Prior vs. beta(1, 1)
$d_{\text{MTD}} = 0.5$	0.054	0.057	0.057	5%
$d_{\text{MTD}} = 1$	0.063	0.069	0.079	9%
$d_{\text{MTD}} = 1.5$	0.052	0.059	0.067	12%

in the Supplemental Materials S6.

4. Case Studies

We now illustrate the proposed methodologies by analysing two clinical trial case studies. First, we study an important application in a pediatric trial where historical data on adults is available. This constitutes a situation of increased importance due to the difficulty in enrolling pediatric patients in clinical trials (U.S. Food and Drug Administration, 2016). Then, we study a classical problem in the analysis clinical trials: using information from a previous study. This is illustrated with data on trials of interferon treatment for melanoma.

4.1 Pediatric Lupus Trial

Enrolling patients for pediatric trials is often difficult due to the small number of available patients, parental concern regarding safety and technical limitations (Psioda and Xue, 2020). For many pediatric trials, additional information must be incorporated for any possibility of establishing efficacy (Psioda and Xue, 2020). The use of Bayesian methods is natural for extrapolating adult data in pediatric trials through the use of informative priors, and is demonstrated in FDA guidance on complex innovative designs (U.S. Food and Drug Administration, 2019).

Belimumab (Benlysta) is a biologic for the treatment of adults with active, autoantibody-positive systemic lupus erythematosus (SLE). It was proposed that the indication for Belimumab can be expanded to include the treatment of children (Psioda and Xue, 2020). The clinical trial PLUTO (Brunner et al., 2020) has been conducted to examine the effect of Belimumab on children 5 to 17 years of age with active, seropositive SLE who are receiving standard therapy. The PLUTO study has a small sample size due to the rarity of childhood-onset SLE. There have been two previous phase 3 trials, BLISS-52 and BLISS-76 (Furie et al., 2011; Navarra et al., 2011), which established efficacy of belimumab plus standard therapy for adults. The FDA review of the PLUTO trial submission used data from the adult trials to inform the approval decision (Psioda and Xue, 2020). All three trials employ the same composite primary outcome, the SLE Responder Index (SRI-4).

We conduct a Bayesian analysis of the PLUTO study incorporating information from the adult studies, BLISS-52 and BLISS-76, using a normalized power prior. We derive the optimal priors on

a_0 based on the KL criterion and the MSE criterion.

Our parameter of interest is the treatment effect of Belimumab for children, denoted by β . The total sample size of the pooled adult data (BLISS-52 and BLISS-76) is $n_0 = 1125$ and the treatment effect is 0.481. We choose $d_{\text{MTD}} = 0.481$ which equals the treatment effect of the historical study. The pediatric data has a sample size of 92 and the estimated treatment effect is 0.371. We use the asymptotic normal approximation to the logistic regression model (see equation (1.4) in Supplemental Materials) with one covariate (the treatment indicator). We choose $w = 0.5$ and $n = 100$ (sample size of the simulated current dataset). For the KL criterion, the objective function K is computed using Monte Carlo integration and optimization is performed using the `optim()` function in R. For the MSE criterion, the posterior mean of β is computed using the R package *BayesPPD* and optimization is performed using a grid search where the values of α_0 and β_0 range from 0.5 to 6 with an increment of 0.5. The optimal priors derived using the KL criterion and MSE criterion are displayed in Figure 5. Table 3 (left) provides the posterior mean, standard deviation and 95% credible interval for β using the optimal priors and several other beta priors for comparison. We observe that the optimal prior derived using the KL criterion leads to a lower posterior standard deviation compared to the uniform prior because more historical information is borrowed. Note that when a $\text{beta}(1, 10)$ prior is used for a_0 , i.e., very little historical information is borrowed, the 95% credible interval of β includes zero. In Supplemental Materials S7, we demonstrate using the proposed optimal priors in a clinical trial design application for the pediatric lupus trial.

Table 3: Pediatric lupus trial (left) and melanoma trial (right): posterior mean, standard deviation, and 95% credible interval for β

Prior for a_0	Lupus Trial			Prior for a_0	Melanoma Trial		
	Mean	SD	95% CI		Mean	SD	95% CI
beta(5.5, 5.5) (optimal by KL)	0.47	0.16	(0.15, 0.79)	beta(0.7, 1.5) (optimal by KL)	0.05	0.19	(-0.33, 0.43)
beta(2, 5) (optimal by MSE)	0.47	0.21	(0.04, 0.89)	beta(5.5, 3) (optimal by MSE)	-0.01	0.17	(-0.35, 0.33)
beta(1, 1)	0.47	0.18	(0.12, 0.83)	beta(1, 1)	0.04	0.19	(-0.32, 0.42)
beta(2, 2)	0.47	0.17	(0.13, 0.79)	beta(2, 2)	0.03	0.18	(-0.34, 0.40)
beta(0.5, 0.5)	0.47	0.17	(0.12, 0.81)	beta(0.5, 0.5)	0.05	0.19	(-0.33, 0.41)
beta(10, 1)	0.48	0.12	(0.24, 0.71)	beta(10, 1)	-0.08	0.16	(-0.38, 0.23)
beta(1, 10)	0.44	0.28	(-0.13, 1)	beta(1, 10)	0.06	0.19	(-0.3, 0.44)

4.2 Melanoma Trial

Interferon Alpha-2b (IFN) is an adjuvant chemotherapy for deep primary or regionally metastatic melanoma. IFN was used in two phase 3 randomized controlled clinical trials, E1684 and E1690 (Kirkwood et al., 1996). In this example, we choose overall survival (indicator for death) as the primary outcome. We conduct a Bayesian analysis of the E1690 trial incorporating information from the E1684 trial, using a normalized power prior. We include three covariates in the analysis, the treatment indicator, sex and the logarithm of age. As before, we obtain the optimal priors for α_0 based on both the KL criterion and the MSE criterion.

Our parameter of interest is the treatment effect of IFN, denoted by β . The total sample size of the E1684 trial is $n_0 = 285$ and the treatment effect is -0.423 . We choose $d_{\text{MTD}} = 0.423$ which equals the treatment effect of the historical study. The E1690 trial has a sample size of 427 and the treatment effect is 0.098. We use the asymptotic normal approximation to the logistic regression model (see equation (1.4) in Supplemental Materials) with three covariates. We choose $w = 0.5$, $d_{\text{MTD}} = 0.423$ and $n = 400$ (sample size of the simulated current dataset). For the KL criterion, the objective function K is computed using Monte Carlo integration and optimization is performed using the `optim()` function in R. For the MSE criterion, the posterior mean of β is computed using the R package *BayesPPD* and optimization is performed using a grid search where the values of α_0 and β_0 range from 0.5 to 6 with an increment of 0.5. The optimal priors derived using the KL criterion and MSE criterion are displayed in Figure 6. The optimal prior derived using the KL criterion is $\text{beta}(0.7, 1.5)$, which has density around zero. For the MSE criterion, the optimal

prior derived is $\text{beta}(5.5, 3)$, which is unimodal with mode around 0.7. This is likely due to the fact that d_{MTD} is small relative to the total sample size of 712 – see also simulations in Supplemental Materials. Because the observed difference is larger than d_{MTD} , the marginal posterior of a_0 has mode around 0.4, which discourages more strongly than the prior. Table 3 (right) provides the posterior mean, standard deviation and 95% credible interval for β using the optimal priors and several other beta priors for comparison. Compared to the uniform prior, the optimal prior derived using the KL criterion results in a larger posterior mean, indicating that less historical information is borrowed. Compared to the uniform prior, the optimal prior derived using the MSE criterion borrows more historical information, resulting in a smaller posterior mean and a smaller variance.

5. Discussion

In this paper, we have explored the asymptotic properties of the normalized power prior when the historical and current data are compatible and when they are incompatible. Our results demonstrate that there is a fundamental asymmetry in the adaptive borrowing properties of the normalized power prior: while for any discrepancy between the historical and current data sets one can expect the marginal posterior of the discounting parameter a_0 to concentrate around zero, under complete compatibility this distribution does not *concentrate* around 1 – although the mode is at 1 under mild conditions (Han et al., 2022). This has been seen as flaw with the NPP by some authors (e.g. Pawel et al. (2023)), but we argue this view is misguided: expecting the posterior of a_0 to concentrate around 1 is effectively expecting two data sets to support exchangeability. This is a notoriously

hard task, as exemplified by the difficulty in estimating the population variance in hierarchical models. Moreover, we have shown that for an i.i.d. normal model with finite samples, the marginal posterior of a_0 always has more mass around one when the datasets are fully compatible, compared to the case where there is any discrepancy.

We have proposed two criteria based on which the optimal hyperparameters of the prior for a_0 can be derived. While the exact values of the hyperparameters can be obtained using our objective functions, we suggest the following rules of thumb for estimating the optimal prior given different choices of the maximum tolerable difference. When the KL criterion is used, a beta distribution centered around 0.5, such as the $\text{beta}(2, 2)$, is optimal for small values (when plots of the current and historical data likelihoods substantially overlap) of maximum tolerable difference, while a beta distribution with mean close to 1, such as the $\text{beta}(2, 0.5)$, should be used for large values of maximum tolerable difference. When the MSE criterion is used, a beta distribution with mean less than 0.5, such as the $\text{beta}(3, 6)$, is optimal for small values of maximum tolerable difference, while a beta distribution with modes at zero and one, as for example a $\text{beta}(0.5, 0.5)$, should be used for large values of maximum tolerable difference. The MSE criterion is a more conservative criterion, in the sense that it tends to discourage borrowing, than the KL criterion.

We observe that in Figures 1-3, the marginal posterior of a_0 is bimodal, with modes at zero and one, in some cases. Future work should further examine the issue of multi-modality. Other potential future work includes extending our method to survival and longitudinal outcomes, as well as accommodating dependent discounting parameters when multiple historical datasets are

available.

Supplementary Materials

The reader is referred to the online Supplementary Materials for proofs of the theorems in section 2, additional simulations and results for section 2 and 3, and a clinical trial design application for section 4.

Acknowledgements

We thank Professor Yuri Saporito and Rodrigo Alves for valuable discussions.

References

- Banbeta, A., J. van Rosmalen, D. Dejardin, and E. Lesaffre (2019, Mar). Modified power prior with multiple historical trials for binary endpoints. *Statistics in Medicine* 38(7), 1147–1169.
- Bennett, M., S. White, N. Best, and A. Mander (2021). A novel equivalence probability weighted power prior for using historical control data in an adaptive clinical trial design: A comparison to standard methods. *Pharmaceutical Statistics* 20(3), 462–484.
- Boonstra, P. S. and R. P. Barbaro (2020, Apr). Incorporating historical models with adaptive Bayesian updates. *Biostatistics* 21(2), e47–e64.
- Brunner, H. I., C. Abud-Mendoza, D. O. Viola, I. Calvo Penades, D. Levy, J. Anton, J. E. Calderon, V. G. Chasnyk, M. A. Ferrandiz, V. Keltsev, and et al. (2020, Jul). Safety and efficacy of intravenous belimumab in children with

- systemic lupus erythematosus: results from a randomised, placebo-controlled trial. *Annals of the Rheumatic Diseases* 79, 1340–1348.
- Carvalho, L. M. and J. G. Ibrahim (2021, Jul). On the normalized power prior. *Statistics in Medicine* 40(24), 5251–5275.
- Duan, Y., K. Ye, and E. P. Smith (2006, feb). Evaluating water quality using power priors to incorporate historical information. *Environmetrics (London, Ont.)* 17(1), 95–106.
- Furie, R., M. Petri, O. Zamani, R. Cervera, D. J. Wallace, D. Tegzová, J. Sanchez-Guerrero, A. Schwarting, J. T. Merrill, W. W. Chatham, and et al. (2011, Dec). A phase III, randomized, placebo-controlled study of belimumab, a monoclonal antibody that inhibits B lymphocyte stimulator, in patients with systemic lupus erythematosus. *Arthritis and Rheumatism* 63(12), 3918–3930.
- Gravestock, I. and L. Held (2019). Power priors based on multiple historical studies for binary outcomes. *Biometrical Journal* 61(5), 1201–1218.
- Gravestock, I., L. Held, and C.-N. consortium (2017, Jun). Adaptive power priors with empirical Bayes for clinical trials. *Pharmaceutical Statistics* 16(5), 349–360.
- Han, Z., K. Ye, and M. Wang (2022, Mar). A study on the power parameter in power prior Bayesian analysis. *The American Statistician*, 1–8.
- Ibrahim, J. G. and M.-H. Chen (2000, feb). Power prior distributions for regression models. *Statistical Science* 15(1), 46–60.
- Ibrahim, J. G., M.-H. Chen, and D. Sinha (2003, mar). On optimality properties of the power prior. *Journal of the*

- American Statistical Association* 98(461), 204–213.
- Kirkwood, J. M., M. H. Strawderman, M. S. Ernstoff, T. J. Smith, E. C. Borden, and R. H. Blum (1996). Interferon alfa-2b adjuvant therapy of high-risk resected cutaneous melanoma: the eastern cooperative oncology group trial est 1684. *Journal of Clinical Oncology* 14, 7–17.
- Liu, G. F. (2018). A dynamic power prior for borrowing historical data in noninferiority trials with binary endpoint. *Pharmaceutical Statistics* 17(1), 61–73.
- Navarra, S. V., R. M. Guzmán, A. E. Gallacher, S. Hall, R. A. Levy, R. E. Jimenez, E. K.-M. Li, M. Thomas, H.-Y. Kim, M. G. León, and et al. (2011, Feb). Efficacy and safety of belimumab in patients with active systemic lupus erythematosus: a randomised, placebo-controlled, phase 3 trial. *The Lancet* 377(9767), 721–731.
- Neelon, B. and A. J. O’Malley (2010). Bayesian analysis using power priors with application to pediatric quality of care. *J Biomet Biostat* 1(1), 1–9.
- Neuenschwander, B., M. Branson, and D. J. Spiegelhalter (2009). A note on the power prior. *Statistics in Medicine* 28, 3562–3566.
- Pan, H., Y. Yuan, and J. Xia (2017, Nov). A calibrated power prior approach to borrow information from historical data with application to biosimilar clinical trials. *Journal of the Royal Statistical Society. Series C, Applied statistics* 66(5), 979–996.
- Pawel, S., F. Aust, L. Held, and E.-J. Wagenmakers (2023). Normalized power priors always discount historical data. *Stat* 12(1), e591.
- Psioda, M. A. and X. Xue (2020). A Bayesian adaptive two-stage design for pediatric clinical trials. *Journal of*

Biopharmaceutical Statistics 30(6), 1091–1108.

R Core Team (2022). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Schmidli, H., S. Gsteiger, S. Roychoudhury, A. O’Hagan, D. Spiegelhalter, and B. Neuenschwander (2014, Dec). Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics* 70(4), 1023–1032.

Shen, Y., M. A. Psioda, and J. G. Ibrahim (2022). *BayesPPD: Bayesian Power Prior Design*. R package version 1.1.0.

U.S. Food and Drug Administration (2016). *Leveraging existing clinical data for extrapolation to pediatric uses of medical devices: guidance for industry and Food and Drug Administration Staff*.

U.S. Food and Drug Administration (2019). *Interacting with the FDA on complex innovative trial designs for drugs and biological products: Draft guidance for industry*.

Ye, K., Z. Han, Y. Duan, and T. Bai (2022). Normalized power prior Bayesian analysis. *Journal of Statistical Planning and Inference* 216, 29–50.

Department of Biostatistics, University of North Carolina at Chapel Hill

E-mail: ys137@live.unc.edu

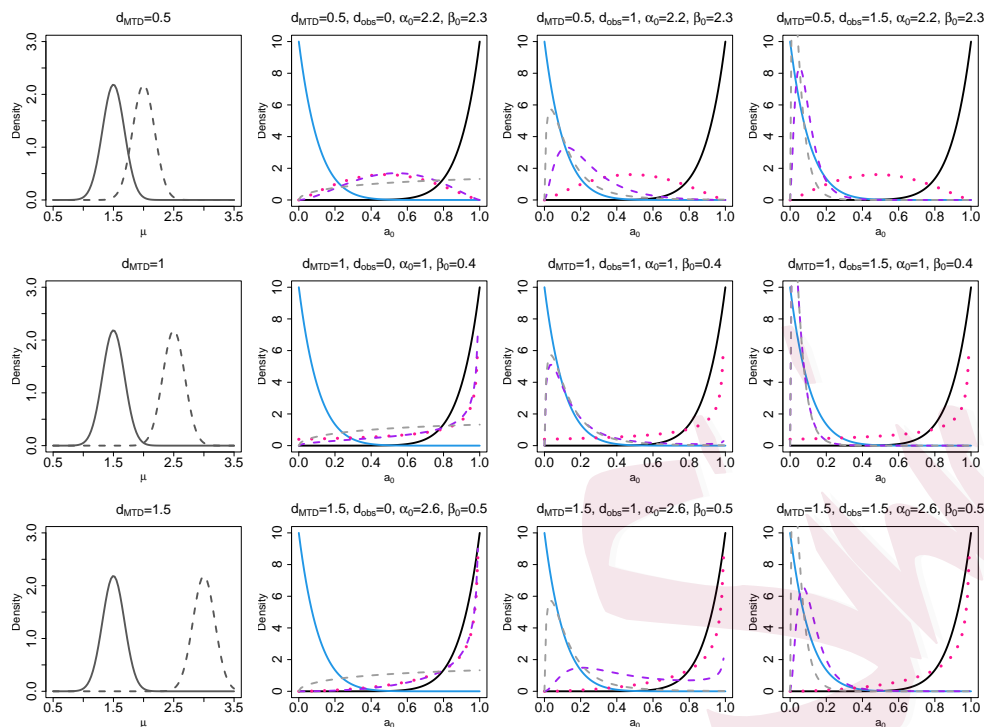


Figure 1: Simulation results for the normal *i.i.d.* case, where $\sigma^2 = 1$, $\bar{y}_0 = 1.5$ and $n = n_0 = 30$. The first figure of each row plots the historical (black solid line) and current (black dashed line) data likelihoods if the hypothetical degree of conflict is equal to d_{MTD} . For each row of the figure, the maximum tolerable difference d_{MTD} is chosen to be 0.5, 1 and 1.5, and the corresponding optimal prior (pink dotted line) is derived for each value of d_{MTD} . For each optimal prior, we vary $d_{\text{obs}} = \bar{y}_{\text{obs}} - \bar{y}_0$ to evaluate the performance of the optimal prior for different observed data. For columns 2-4, d_{obs} is chosen to be 0, 1 and 1.5, respectively. The black and blue curves correspond to $\pi_1(a_0) \equiv \text{beta}(10, 1)$ and $\pi_2(a_0) \equiv \text{beta}(1, 10)$, respectively. The purple dashed line represents the marginal posterior of a_0 with the optimal prior for a given d_{obs} . The grey dashed line plots the marginal posterior of a_0 with the uniform prior.

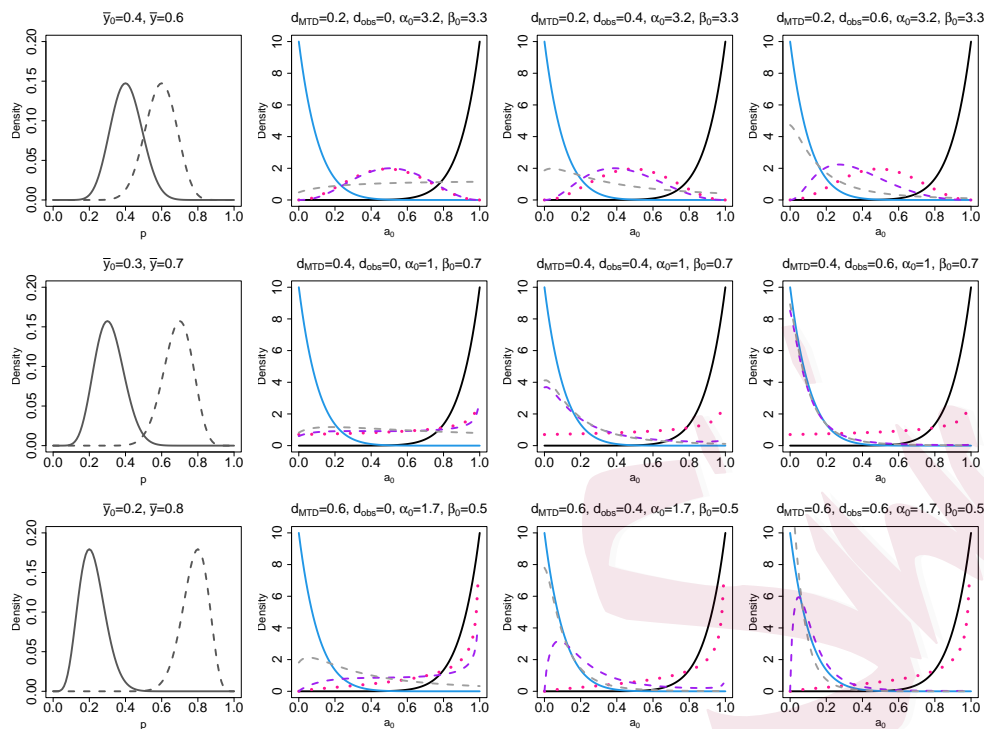


Figure 2: Simulation results for the Bernoulli *i.i.d.* case, where $\sigma^2 = 1$ and $n = n_0 = 30$. The first figure of each row plots the historical (black solid line) and current (black dashed line) data likelihoods if the hypothetical degree of conflict is equal to d_{MTD} . For each row of the figure, the maximum tolerable difference d_{MTD} is chosen to be 0.5, 1 and 1.5, and the corresponding optimal prior (pink dotted line) is derived for each value of d_{MTD} . For each optimal prior, we vary $d_{\text{obs}} = \bar{y}_{\text{obs}} - \bar{y}_0$ to evaluate the performance of the optimal prior for different observed data. For columns 2-4, d_{obs} is chosen to be 0, 1 and 1.5, respectively. The black and blue curves correspond to $\pi_1(a_0) \equiv \text{beta}(10, 1)$ and $\pi_2(a_0) \equiv \text{beta}(1, 10)$, respectively. The purple dashed line represents the marginal posterior of a_0 with the optimal prior for a given d_{obs} . The grey dashed line plots the marginal posterior of a_0 with the uniform prior.

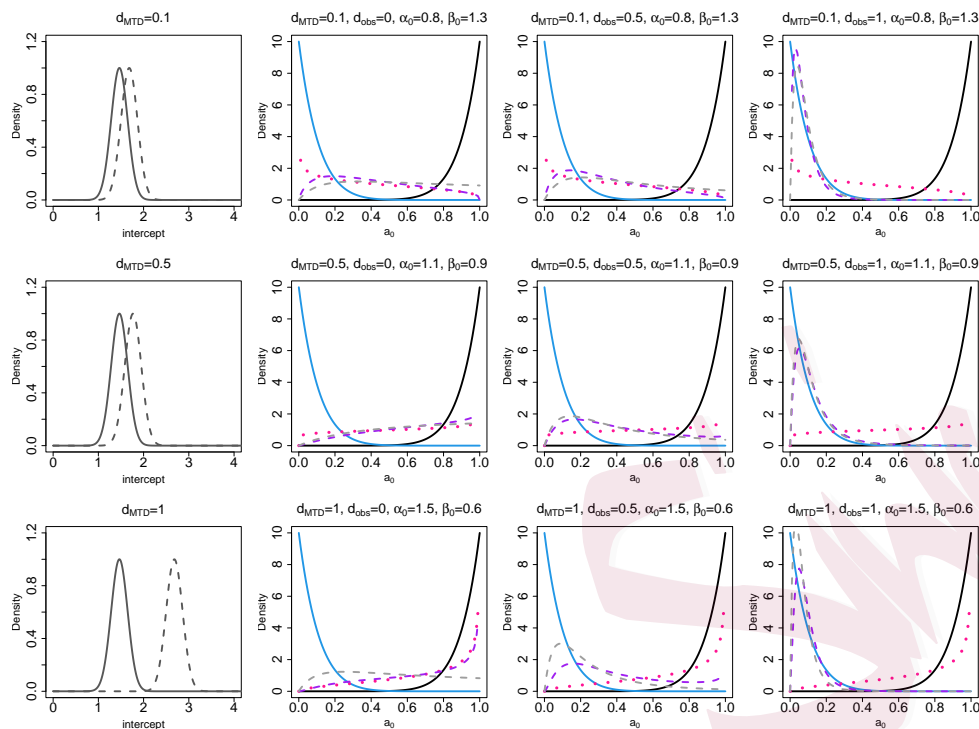


Figure 3: Simulation results for the normal linear model with one covariate where $\beta_0 = 1.5$, $\beta_1 = -1$, $\sigma^2 = 1$ and $n = n_0 = 30$. The first figure of each row shows the historical (black solid line) and current (black dashed line) data likelihoods as a function of the intercept if the hypothetical degree of conflict is equal to d_{MTD} . For each row, d_{MTD} is chosen to be 0.1, 0.5 and 1, and the corresponding optimal prior (pink dotted line) is derived for each value of d_{MTD} . For each optimal prior, we vary d_{obs} to represent different degrees of departure of the intercept of current data to that of historical data. For columns 2-4, d_{obs} is chosen to be 0, 0.5 and 1, respectively. The black and blue curves correspond to $\pi_1(a_0) \equiv \text{beta}(10, 1)$ and $\pi_2(a_0) \equiv \text{beta}(1, 10)$, respectively. The purple dashed line represents the marginal posterior of a_0 with the optimal prior for a given d_{obs} . The grey dashed line plots the marginal posterior of a_0 with the uniform prior.

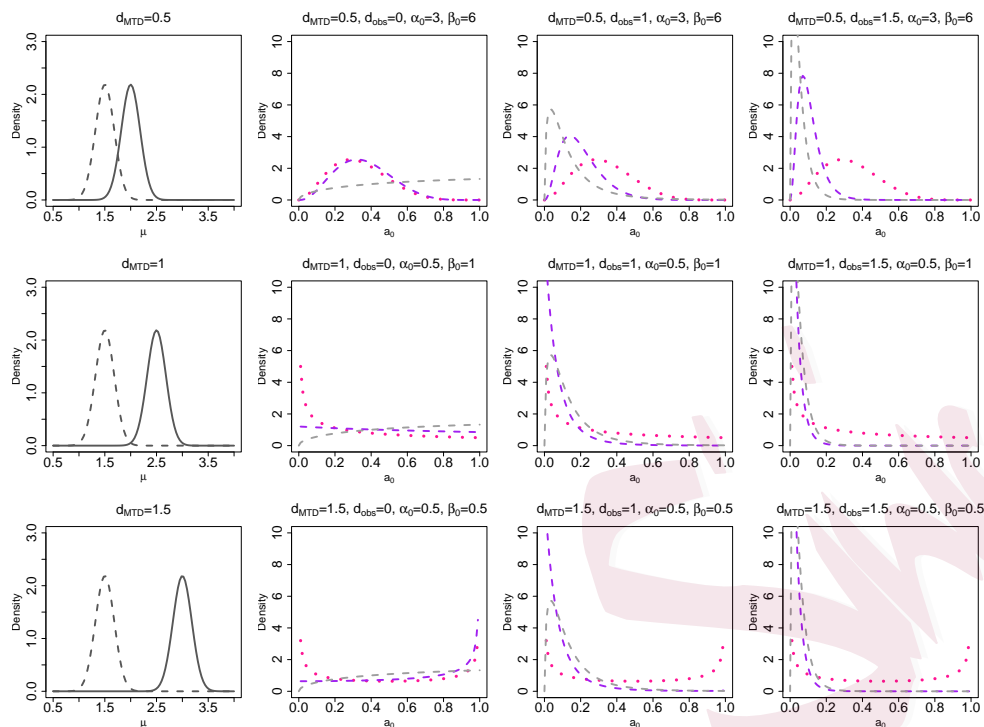


Figure 4: Simulation results for the normal *i.i.d.* case when minimizing a convex combination of MSEs when $n = n_0 = 30$. The first figure of each row shows the historical (black solid line) and current (black dashed line) data likelihoods if the hypothetical degree of conflict is equal to d_{MTD} . The mean of the hypothetical current data is fixed at 1.5. For each row of the figure, the maximum tolerable difference d_{MTD} is chosen to be 0.5, 1 and 1.5, and the corresponding optimal prior (pink dotted line) is derived for each value of d_{MTD} . For each optimal prior, we vary $d_{\text{obs}} = \bar{y}_0 - \bar{y}_{\text{obs}}$. For columns 2-4, d_{obs} is chosen to be 0, 1 and 1.5, respectively. The purple dashed line represents the marginal posterior of a_0 with the optimal prior for a given d_{obs} . The grey dashed line plots the marginal posterior of a_0 with the uniform prior.

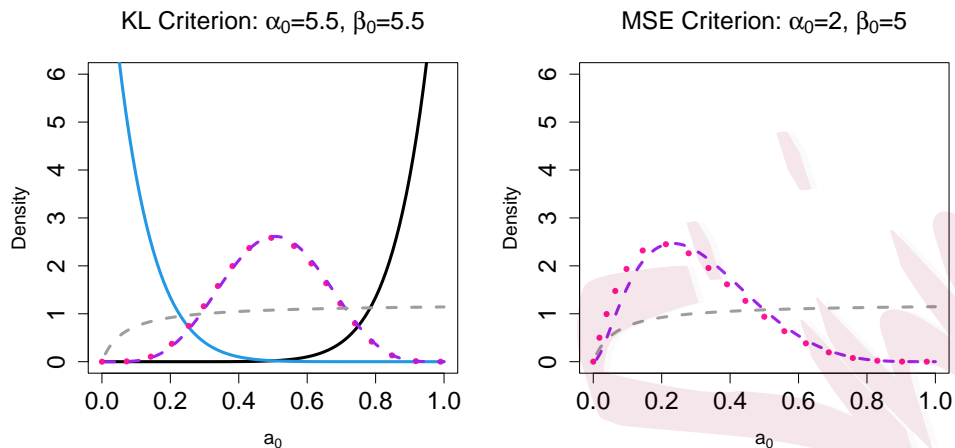


Figure 5: After combining studies BLISS-52 and BLISS-76 for adults, the total sample size is $n_0 = 1125$ and log odds ratio for treatment vs. control group is 0.481. We choose $d_{\text{MTD}} = 0.481$ to be the maximum tolerable difference. The pediatric data has a sample size of $n = 92$. The actual observed log odds ratio is 0.371. The figure on the left displays the optimal prior (pink dotted line) and posterior (purple dashed line) derived using the KL criterion. The figure on the right displays the optimal prior for a_0 and the posterior derived using the MSE criterion. The posterior of a_0 using the uniform prior (grey dashed line) is also shown.

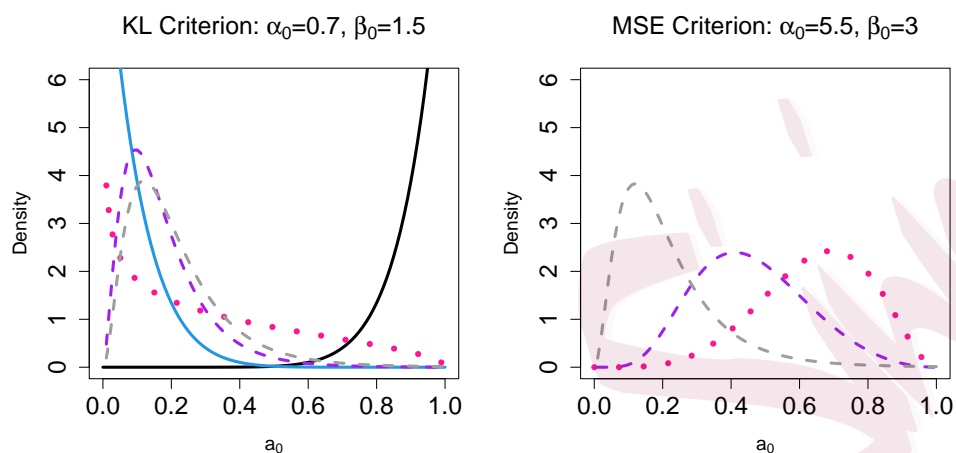


Figure 6: The total sample size of the E1684 trial is $n_0 = 285$ and log odds ratio for treatment vs. control group is -0.423 . We choose $d_{\text{MTD}} = 0.423$ to be the maximum tolerable difference. The E1690 trial has sample size $n = 427$. The observed log odds ratio is 0.098 . The figure on the left displays the optimal prior (pink dotted line) and posterior (purple dashed line) derived using the KL criterion. The figure on the right displays the optimal prior for a_0 and the posterior derived using the MSE criterion. The posterior of a_0 using the uniform prior (grey dashed line) is also shown.