Contemporary Clinical Trials xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Contemporary Clinical Trials

journal homepage: www.elsevier.com/locate/conclintrial

Bayesian clinical trial design using Markov models with applications to autoimmune disease

Barry S. Eggleston [a], Joseph G. Ibrahim [b,*], Diane Catellier [a]

[a] RTI International, 3040 East Cornwallis Road, Research Triangle Park, NC 27709, USA
[b] Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA

ARTICLE INFO

Article history:
Received 23 July 2016
Received in revised form 9 December 2016
Accepted 4 February 2017
Available online xxxx

Keywords:
Markov models
Illness-death model
Bayesian methods
Clinical trial design

ABSTRACT

Immune Thrombocytopenia is an autoimmune disease associated with bleeding that is treated by increasing the platelet count to a level where the chance of uncontrollable bleeding is low. Failure occurs when platelet counts are not raised sufficiently (initial failure), or when high platelet counts are not maintained after initial success (relapse). In this paper, we propose a Bayesian clinical trial design that uses a Markov multistate model along with a power prior for the parameters which incorporates historical control data to estimate transition rates among two randomized groups as defined by the model. A detailed simulation is carried out to examine the operating characteristics of a trial to test whether a new treatment reduces the relapse rate by 40% relative to standard care when data from 60 historical controls treated with standard care is available. We also use simulated data to demonstrate effects of discordance between historical and randomized controls on the estimated hazard ratios. Finally, we use a simulated trial to demonstrate briefly what type of results the model can give and how those results can be used to address hypotheses regarding treatment effects. Using simulated data, we show that the model yields good operating characteristics when the historical and randomized controls are from the same population, and demonstrate how discordance between the control groups affects the operating characteristics.

1. Introduction

Immune Thrombocytopenia (ITP) is an autoimmune disease associated with bleeding. The bleeding is the result of impaired platelet production that is often persistent and chronic and due to platelet destruction by the immune system [5]. ITP treatment consists of two objectives: quickly raising the platelet count to a therapeutic level to stop the ITP relating bleeding and maintaining the therapeutic platelet level to prevent future ITP-related bleeding. Current treatment uses established treatments to address the first objective, but clinicians need improved treatments to reduce the risk of relapse among initial treatment responders. To identify these improved treatments, clinical researchers need efficient effect estimation models to properly account for failure to meet either objective.

This article proposes a three-state Markov Bayesian model that can incorporate historical control data if necessary to evaluate experimental ITP treatments while taking into account the two possible failure modes mentioned above. The multi-state nature of the model accounts for these two types of failure. The Bayesian model allows for inclusion of prior information and historical control data via the power prior of Ibrahim and Chen [7]. Using simulation, we demonstrate how the proposed Bayesian model can be used to estimate hazard rates and ratios from ITP trial data, its power and type 1 error characteristics, and how the model responds to differences between randomized and historical control relapse rates.

The article is organized as follows. In Section 2, we describe a clinical trial design that corresponds to the proposed three-state Markov model, we describe the Bayesian model used to estimate the parameters of the model, we describe the structure of the analysis dataset used in the model, and we illustrate the simulation process used to study and illustrate the Bayesian model. In Section 3, we describe the simulation results, and in Section 4 we conclude the article with a discussion. Financial funding for the research reported in this publication was given by the National Heart, Lung, And Blood Institute of the National Institutes of Health under Award Number U24HL114577. The sponsor had no other role in the research beyond funding.

2. Methods

2.1. Trial design

The ITP trial design, for which the Bayesian model is used to estimate treatment effects, consists of one active arm and one control arm, into which patients are randomized at the time of enrollment using a 1:1 ratio. After enrollment, both arms are treated with a selected therapy

* Corresponding author.
E-mail address: jibrahim@email.unc.edu (J.G. Ibrahim).

http://dx.doi.org/10.1016/j.cct.2017.02.004
1551-7144/

intended to raise the platelet counts to therapeutic levels. One possible outcome of this initial treatment is failure to respond. After this initial treatment period, initial responders in both arms have their treatment regimens modified to allow for a comparison of at least one experimental therapy with respect to ITP relapse rates. For example, responders in the control arm might receive standard relapse prevention care while the active arm might begin to take an experimental therapy intended to reduce relapse rates relative to standard care. Using data collected from the trial design, the Bayesian model is used to estimate initial and relapse failure rates and hazard ratios, as well as response rates to the initial treatment.

Another part of the design involves the platelet assessment schedule. Platelet levels are assessed periodically according to a prespecified schedule. If the primary outcome is time-to-relapse and relapse is defined as time until platelet levels are less than some threshold, then the time-to-relapse will be interval censored. Therefore, the Bayesian model also takes interval censoring into account.

2.2. Three-state model

As noted earlier, ITP treatment failure occurs in two ways. A subject is an initial non-responder (an initial failure), if he or she does not experience therapeutically high platelet levels during initial treatments to raise platelet levels. On the other hand, an initial responder can experience relapse if the initial responder experiences critically low platelet levels at some later time. In order to accurately estimate the relapse failure rate, modelling needs to account for the possibility of initial failure, initial nonresponse. If these two paths to failure are combined with an initial state of critically low platelet counts, and a second state of initial treatment response, then a three-state model identical to an Illness-Death model can be used to accurately estimate the relapse failure rate without confounding from the initial failure rate, Fig. 1 [3,13].

State 1 of the three-state model represents a patient's initial condition of critically low platelet levels. State 2 represents initial treatment responders who have obtained a therapeutically high platelet count. A subject can reach State 2 only from State 1, a "1–2" transition. State 3 represents treatment failure, and subjects can reach State 3 from States 1 or 2. First, a subject with a critically low platelet count could fail to respond to initial treatment within a pre-specified period. This subject is an initial non-responder and reaches State 3 from State 1, a "1–3" transition. Second, a subject could fail by relapsing to a critically low platelet count after reaching State 2, a "2–3" transition. In an ITP trial that is evaluating a treatment for relapse, transitions rates between States 1 and 2 should be high and transition rates between States 1 and 3 may be driven mostly by treatment protocol with respect to the definition of an initial non-responder. A subject must be observed to occupy State 2 before a "1–2" transition can be declared, and only subjects who are observed to occupy State 2 can contribute to the estimation of 2–3 transition model parameters. The proposed Bayesian model presented in this paper can be used to estimate all six transition rates, three for each arm, and all three hazard ratios.

2.3. Likelihood and priors for the randomized component of the Bayesian trial design

The MCMC models used in this article are based on exponential distributions with the mean transition time equal to a function of treatment. The likelihood function component for the 1–2 transition is given in Eq. (1), where $\delta_j = 1$ indicates whether or not the transition time is interval censored, $(t_{12L}k, t_{12U}k)$, $\delta_j = 2$ indicates a right censored at an unknown value, $[\hat{t}_{13L}k, \infty]$, or $\delta_j = 3$ indicates a right censored at a known value $[t_{12L}k, \infty]$. An unknown 1–2 right censored time occurs when an observed but interval censored 1–3 transition occurs. The values $t_{12L}k$ and $t_{12U}k$ are left and right interval time points for an interval censored 1–2 transition time similar to SAS's input structure for interval censoring in PROC LIFEREG. In case of a right censored time point, only one value is needed, so when $\delta_j = 3$ the value of $t_{12L}k$ represents the exact timing of a right censored 1–2 transition time. The value of $t_{12U}k$ can be any value allowable by the estimation algorithm. Similarly, when $\delta_j = 2$ the value of $\hat{t}_{13L}k$ represents the predicted value of a 1–3 transition time, but due to competing risk between 1–2 and 1–3 transitions $\hat{t}_{13L}k$ also is equal to the predicted time when the 1–2 transition was right censored, A predicted right censored 1–2 transition time is needed when a subject transitions out of State 1 into State 3, because the transition timing into State 2 is right censored at an unknown value. The 1–3 transition time is interval censored, so the right censored transition time for the 1–2 transition is only known to occur within an interval. In such a case the present model predicts the unknown right censoring time for the 1–2 transition by making a random draw from a truncated version of 1–3 transition time distribution. The μ_{12k} is the exponential mean for the kth subject [11]. The parameter μ_{12k} is linked with the treatment variable by the log link, $\log(\mu_{12k}) = \beta_{120} + \beta_{121}Trt_k$, where Trt_k equals 1, if subject k is in the active arm and 0 otherwise.

$$L_{12} = \prod_{k=1}^{n}\left[\exp\left(-\frac{t_{12L}k}{\mu_{12k}}\right) - \exp\left(-\frac{t_{12U}k}{\mu_{12k}}\right)\right]^{\delta_k=1}\left[\exp\left(-\frac{\hat{t}_{13L}k}{\mu_{12k}}\right)\right]^{\delta_k=2}\left[\exp\left(-\frac{t_{12L}k}{\mu_{12k}}\right)\right]^{\delta_k=3} \quad (1)$$

In summary: using a convention similar to the SAS convention for representing interval censoring, if the subject experienced a 1–2 transition within the interval $(t_{12L}k, t_{12U}k)$, then δ_j equals 1. If the subject was right censored at an unknown value due to an observed but interval censored 1–3 transition, then the transition interval is represented by $(\hat{t}_{13L}k, C)$ and δ_j will equal 2, where $\hat{t}_{13L}k$ is a predicted value of when the 1–3 transition occurred and C is some arbitrary value larger than the possible maximum follow-up time. This C value was necessary given the JAGS 4.2 fitting algorithm used to estimate parameter posterior distributions. If the subject was right censored at a known value, then the transition interval is represented by $(t_{12L}k, C)$ and δ_j equals 3, where C is some arbitrary value larger than the possible maximum follow-up time.

Due to competing risks among 1–2 and 1–3 transitions, the likelihood function component for the 1–3 transition is similar to the likelihood function for the 1–2 transition and is given in Eq. (2), where δ_j indicates transition type, $\log(\mu_{13k}) = \beta_{130} + \beta_{131}Trt_k$, and $t_{13L}k$, $t_{13U}k$, and $\hat{t}_{12L}k$ are defined in a similar manner as $t_{12L}k$, $t_{12U}k$, and $\hat{t}_{13L}k$. In the case of interval censored 1–3 transition, the values $t_{13L}k$ and $t_{13U}k$, are left and right endpoints of a 1–3 transition. In the case of right censoring at a known value, $t_{13L}k$ is equal to the known timing of the right censoring. In the case of right censoring at an unknown value, $\hat{t}_{12L}k$ is equal to the predicted timing of a 1–2 transition which due to competing risks is equal to the predicted time of a right censored 1-3 transition time.

$$L_{13} = \prod_{k=1}^{n}\left[\exp\left(-\frac{t_{13L}k}{\mu_{13k}}\right) - \exp\left(-\frac{t_{13U}k}{\mu_{13k}}\right)\right]^{\delta_k=1}\left[\exp\left(-\frac{\hat{t}_{12L}k}{\mu_{13k}}\right)\right]^{\delta_k=2}\left[\exp\left(-\frac{t_{13L}k}{\mu_{13k}}\right)\right]^{\delta_k=3} \quad (2)$$

To estimate the 1–2 transition and 1–3 transition rates using the Bayesian model, priors need to be placed on β_{120}, β_{121}, β_{130}, and β_{131}. In this paper, non-informative $N(\mu=0, \sigma=100)$ priors are used.

Since $\hat{t}_{12L}k$ and $\hat{t}_{13L}k$ are predicted values, predictive distributions need to be identified for $\hat{t}_{12L}k$ and $\hat{t}_{13L}k$. The values of $\hat{t}_{12L}k$ will be drawn from a truncated version of the model's estimated exponential distribution for the 1-2 transition times, since the censoring mechanism for the 1-3 transition is equal to the 1-2 transition mechanism. For $\hat{t}_{12L}k$ the truncation will be driven by subject specific left and right values of the observed interval censored 1-2 transition. For $\hat{t}_{13L}k$ the truncation will be driven by subject specific left and right values of the observed interval censored 1-3 transition. Given the assumption of exponential

Please cite this article as: B.S. Eggleston, et al., Bayesian clinical trial design using Markov models with applications to autoimmune disease, Contemp Clin Trials (2017), http://dx.doi.org/10.1016/j.cct.2017.02.004