A Time to Event Approach for Standard of Care Meta-Analysis in HCV Trials

Daniel G. Polhamus[1], Kyle T. Baron[1], Vivek S. Purohit[2], Patanjali Ravva[2], James A. Rogers[1], Marc R. Gastonguay[1]

[1]Metrum Research Group, Tariffville, CT; [2]Pfizer Inc., New London, CT

Abstract

The current standard of care for hepatitis C virus (HCV) is a combination of pegylated interferon with ribavirin (peg-IFNα+RBV) for 48 weeks in genotype 1 infections. When treated with peg-IFNα+RBV, sustained viral response (SVR) is observed in approximately 40-60% of previously untreated patients [Hadziyannis et al., 2004, McHutchison et al., 2009b, Fried et al., 2002]. SVR is strongly correlated with the rate at which the viral load initially becomes undetectable [Ferenci et al., 2005] (reduced viral load - RVL). In the context of a clinical trial concerning interim analysis, one may therefore use patient RVL event time to inform SVR prediction. Meta data (HCV-g1) as peg-IFNα+RBV treatment study response summaries are leveraged to inform both standard of care treatment arm RVL times and the corresponding SVR for patient level data. The proposed model utilizes Bayesian methodology for the inclusion of summary data through a combination of parametric time-to-event and logistic models. The model is easily modified to characterize additional or alternative treatments (e.g., direct acting agents) or pharmacometrics, such as exposure. To our knowledge this is the first attempt to describe these data using time-to-event distributions.

Introduction

Hepatitis C virus (HCV) is a leading cause of chronic liver disease, infecting an estimated 270–300 million people worldwide and progressing to cirrhosis (over a 20 year period) at a rate of roughly 20%. Several factors influence the rate of disease progression, e.g., age, gender, alcohol consumption, and HCV genetic subtype. There are six major genotypes of the HCV virus with HCV-g1 accounting for approximately 80% of US hepatitis C patients. Unfortunately, this subtype is also the most difficult to treat, with successful treatment in around half of those patients treated. Current treatment for HCV-g1 patients uses pegylated interferon-α plus ribavirin (peg-IFNα+RBV) for 48 weeks. Response to therapy is measured as change in viral load (IU/ml); undetectable levels at 24 weeks post-treatment indicate successful treatment application and the patient is said to have sustained viral response (SVR). The probability of a patient attaining SVR can be predicted by the rapidity of response to the treatment, motivating monitoring of patient viral load during treatment—typically at weeks 4, 12, 24, and 48 (EOT). Thus, metadata typically occurs in the form of proportions of the intent-to-treat population responding at one of the on-treatment weeks, or the 24-week post-treatment time (72 weeks). Non-responding patients are often dropped at 12 weeks (a less than 2 log reduction of viral load count) and 24 weeks (viral load still detectable).

New treatments brought to the clinical trial stage are benchmarked against peg-IFNα+RBV. In light of this, the desire to use a model based meta-analysis to both provide a realistic summary of expected SOC patient data as well as predict current SOC patient behavior given partial information (i.e., interim data) is obvious. We set out to address a solution for this problem using a Bayesian statistical paradigm.

Data

To warrant inclusion in our meta-analysis, we required published studies to have reported results on HCV-g1 patients receiving 180 μg/week peg-IFNα plus RBV 800-1200 mg/day for 48 weeks. Data was typically reported in two ways. The most common data format was to report the proportion of the ITT population with undetectable viral load at weeks 4, 12, 24, and 48 (on treatment) and then study week 72, which is the 24-week post-treatment measurement used to indicate patient sustained viral response. Occasionally, extra information was reported as the number of patients with a sustained viral response, *given that their viral count was first undetectable at week* t.

To summarize, the two data formats are as follows:

1. A sequence of time dependent binomials describing the number of patients that had sustained viral response, given that they were first non-detectable at week t. Let $n_i^{\text{SVR} \mid \text{RVL}}(t)$ be this number, and $n_i^{\text{RVL}}(t)$ be the number of patients with non-detectable viral load at time t. Figure 3(a) shows this data as, for studies $i = 1, ...,$ # of studies following format 1:

$$\frac{n_i^{\text{SVR} \mid \text{RVL}}(t)}{n_i^{\text{RVL}}(t)},$$

   which is an estimate of the probability of having sustained viral response given a patient was first non-detectable at time t.

2. A sequence of time dependent binomials describing the proportion of the intent-to-treat population having non-detectable viral load at the indicated week. Let n_i^{ITT} be the number of patients in the intent-to-treat population for study i, then figure 3(b) shows this second subgroup of data as:

$$\frac{n_i^{\text{RVL}}(t)}{n_i^{\text{ITT}}},$$

   which is an estimate of the probability of a patient being non-detectable by time t. The 24-week sustained viral response counts are reported at week 72 for all but one study reporting data in this format.

Data reported in the first format is the most useful to the clinician, giving predictive capability of the endpoint (SVR) given the week at which a patient's viral load first becomes non-detectable. This information can be leveraged to predict endpoints for the SOC arm in trials comparing the SOC to novel treatments. Data following the second format also provides less (yet still useful) information, instead giving only an idea of the expected *overall* SVR rates for a studied group relative to how rapidly the group achieved non-detectable viral load status.

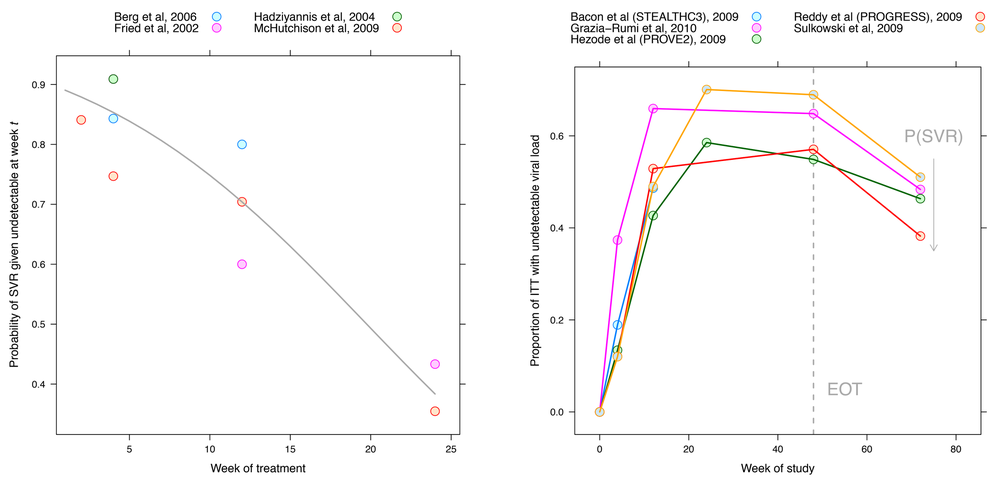

(a) Proportion of SVR patients, given first RVL observed at t.

(b) Proportion of ITT with RVL at t.

Figure 1: *Data: The published data are typically reported in one of two formats: 1) The number of patients with RVL at time t and 2) The number of patients showing sustained viral response (RVL at 24 weeks post-treatment), given that the patient first showed RVL at time t.*

Model

In order to use the data reported in both formats we propose the use of two submodels with the end goal being predictive power for sustained viral response given a patients time to non-detectable or reduced viral load. A logistic model is used to relate sustained viral response to the RVL event times, which are themselves assumed to behave according to a Weibull distribution.

Submodel 1: Logistic regression on probability of SVR given RVL at t

Using $n_i^{\text{SVR} \mid \text{RVL}}(t)$ and $n_i^{\text{RVL}}(t)$ (as defined earlier) to be counts of patients in study i from the first data group, the first submodel uses simple logistic regression to relate sustained viral response to the time to non-detectable viral load.

$$\begin{aligned} n_i^{\text{SVR} \mid \text{RVL}}(t) &\sim \text{Binomial}\left(\Delta(t), n_i^{\text{RVL}}(t)\right), \text{ where} \\ n_i^{\text{RVL}}(t) &= \#\text{ patients undetectable at } t \text{ and} \\ \text{logit}\left(\Delta(t)\right) &= \beta_0 + \beta_1 t, \end{aligned}$$

where $\Delta(t)$ is the probability of a patient having SVR, given they were first non-detectable at time t.

The grey line shown in figure 3(a) shows the predicted value of $\Delta(t)$ from a logistic model fit on just the data indicated in the figure (data following the first format). We will improve upon this by including data reported in the second format.

Submodel 2: Weibull time-to-event time model for patient RVL times

The data reported in the second format will be modeled using a Weibull distribution with a "cure" rate, here reflecting the fraction of the population that responds to treatment—the maximum treatment effect (MTE). For the responder population (say the MTE is π), the times to RVL will be assumed to behave according to the Weibull distribution parameterized with scale ϕ and shape γ. Thus, given that a patient responds to treatment:

$$\begin{aligned} t &\sim \text{Weibull}(\phi, \gamma), \\ f(t) &= \phi\gamma t^{\gamma-1} e^{-\phi t^{\gamma}} \\ F(t) &= 1 - e^{-\phi t^{\gamma}} \\ h(t) &= \frac{f(t)}{1-F(t)} = \phi\gamma t^{\gamma-1} \end{aligned}$$

where t is the time at which the responding patients viral load becomes non-detectable, $f(t)$ is the probability distribution function (pdf), $F(t)$ is the cumulative distribution function (CDF), and $h(t)$ is known as the hazard function. It is through the hazard function and the MTE that we will model inter-study variability, and these are also where patient characteristics (e.g., race, exposure) may be entered as well. In the metadata, we include only inter-study variability. It is modeled through the hazard function and the cure rate using hierarchical Bayesian prior distributions. The Weibull parameters are specified as follows (for study $i = 1, ...,$ # of studies in submodel 2):

- Variability through the scale parameter: $\phi_i \sim \text{Gamma}(\phi_a, \phi_b), \quad i = 1, ...,$ # of studies following submodel 2
- Variability through the cure parameter: $\pi_i \sim \text{Beta}(\pi_a, \pi_b), \quad i = 1, ...,$ # of studies following submodel 2.

Figure 2 highlights the motivation for using a cure rate model to adjust the limits of $F(t)$. The limiting probability of reaching non-detectable viral load asymptotes earlier than the probabilistic limit of a Weibull CDF, $\lim_{t\to\infty} F(t) = 1$ as there is a subgroup of patients that will never show response to the SOC. With response probability, π, then the probability of having RVL by time t in the responder group is described by:

$$\begin{aligned} F^{\star}(t) &= \pi F(t) \text{ so that} \\ \lim_{t\to\infty} F^{\star}(t) &= \pi, \text{ the MTE.} \end{aligned}$$

Thus, for study i (following submodel 2):

$$n_i^{RVL}(t) \sim \text{Binomial}\left(\pi_i F(t|\phi_i, \gamma), n_i^{\text{ITT}}\right).$$

The corresponding probability of sustained viral response is (linking back to submodel 1):

$$\text{P(SVR)}_i = \pi_i \times \text{E}_T\left[\Delta(t)\right] \text{ (see the appendix for derivation).}$$

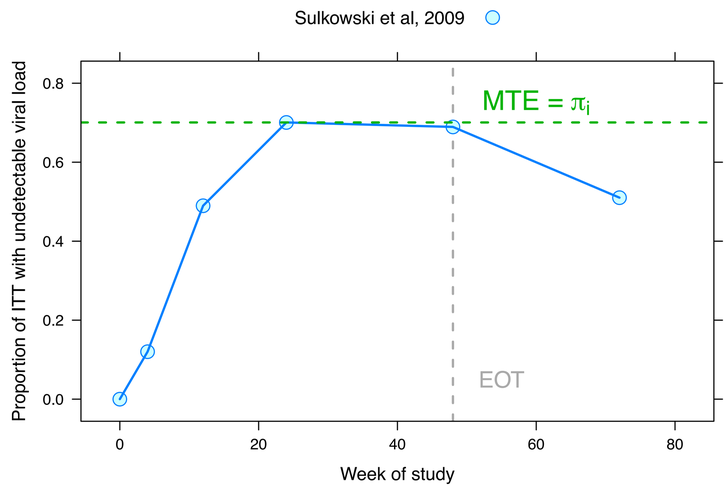

Figure 2: *Second submodel: Data following the second submodel follows a cumulative distribution function, conditioned on probability of response for study i.*

Patient data and covariate inclusion

Patient level data arises in the form of (t_{ij}, x_{ij}), an RVL time and covariates for patient j in study i. If t_{ij} is non-missing and less than 24 weeks (the standard trial censoring time) then the patient is considered a responder. Covariate inclusion occurs through the scale parameter of the current study, ϕ_i with the restriction that $\phi_i > 0$ for all i. To illustrate:

$$t_{ij} \sim \text{Weibull}(\phi_{ij}, \gamma) \text{ where } \phi_{ij} = \phi_i e^{\beta x_{ij}}.$$

Results

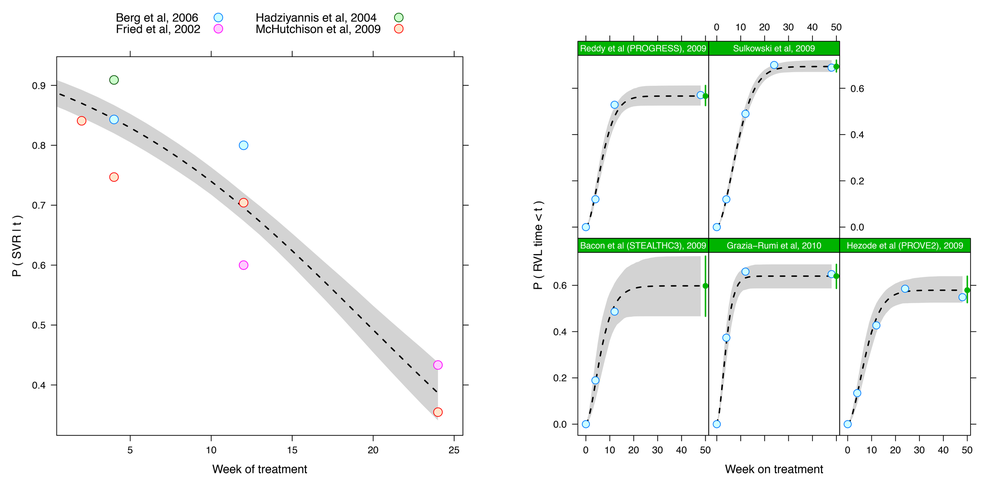

(a) Proportion of SVR patients, given first RVL observed at t.

(b) Proportion of ITT with RVL at t.

Figure 3: *Model predictions: The published data following the two submodels are overlayed on the model predicted mean (black) and 90% credible intervals, shown as the grey band. The 90% credible interval and mean for each study's MTE is shown as the green line and point on the right edge of panel b.*

In order to best demonstrate the utility of the model, we fit the model including a set of hypothetical patients, capturing their predicted SVR probability on the SOC. Figure 3 shows the model fit, demonstrating the predicted model behavior (fitted mean and 90% credible intervals) with the meta data. The density histograms describing the SVR probability for the hypothetical patients are shown in figure 4. The four hypothetical patients are as follows:

1. A patient who has their viral load drop below detectable levels at 4 weeks
2. A patient who has their viral load drop below detectable levels at 12 weeks
3. A patient who has their viral load drop below detectable levels at 24 weeks
4. A pretrial patient, with all response data missing.

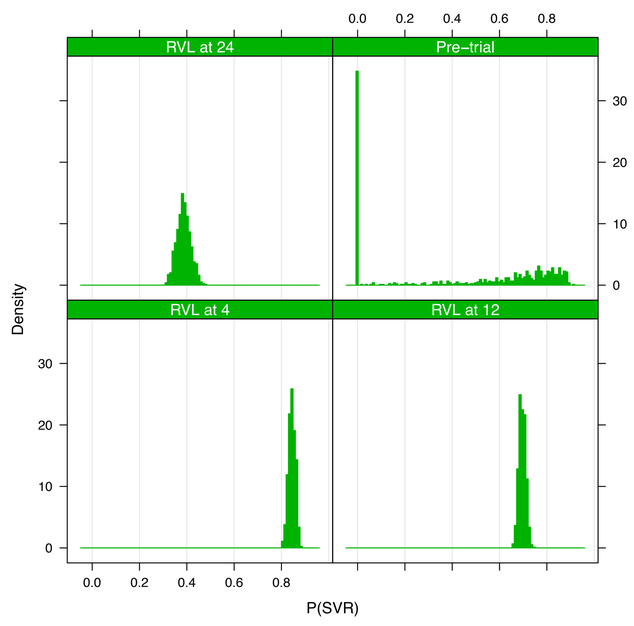

Figure 4: *Results: Predictive distributions for hypothetical patients.*

The model was fit using WinBUGS with 3 MCMC chains of 250 iterations, after burn-in and thinning. Using non-informative priors, the influence of the data on the model is apparent through convergence of the posterior to a distribution differing from the prior (see table 1). The logistic parameters (β_0, β_1), the shape parameters in the Weibull (γ), and hyperparameter ϕ_a all show high precision in their posteriors. The MTE hyperparameters (π_a, π_b) and the Weibull scale hyperparameter (ϕ_b) do not and indicate possible candidates for re-parameterization, but will likely benefit from additional data; we are including only four hypothetical patients. Nevertheless, the model captures the metadata well, as evidenced by Figure 3.

Conclusion

Figure 4 clearly demonstrates both the utility and behavior of the model. Given *only subject RVL time*, we are able to leverage the published data into an estimate of the subject SVR probability. The model predicts that a patient showing non-detectable viral load by 4 weeks will have approximately an 80–90% chance of SVR, dropping to around 65–75% for a patient with RVL by 12 weeks, and then down to 30–40% for an RVL by 24 weeks. Pretrial, we expect that around 65% of the subjects in the trial will respond to treatment (as estimated from the posterior population MTE in table 1, $\frac{\pi_a}{\pi_a+\pi_b}$). Hence, at least 35% have no chance of attaining SVR, shown as the degenerate mode at P(SVR)$= 0$ in the "pretrial" panel of figure 4. Of the remaining set of patients not responding to treatment, their probability of attaining SVR varies according to the distribution of RVL times—a Weibull. The aforementioned process of covariate inclusion will further improve predictive precision by utilizing patient RVL times and additional information related to patient treatment response characteristics.

Thus far the utility of the model has been discussed from a patient-by-patient viewpoint. It is the extension of this to entire control arms that we anticipate driving practical use of the model. Summary statistics representative of trial endpoints can be easily estimated from the predictive SVR distributions given only vectors of patient RVL times. Many of the new HCV treatments (direct acting agents) rapidly decrease patient viral load (i.e., any increased effect over SOC is seen by week 4) which results in the use of shorter treatment durations. With these shorter treatment durations, SVR information is available earlier making the predictive methods shown here even more useful in exploring early trial efficacy while requiring minimal amounts of data, due to the utilization of the metadata.

Appendix: General derivation of modeling approach

Aggregate SVR probability

Data following the second submodel also can contribute information to the SVR probability. To utilize the second data type in the model for prediction of SVR, we first introduce some more notation: let Z indicate responder status, so that $z_j = 1$ indicates that patient j experiences RVL at some point in the trial—hence $t_j \leq 24$ weeks where T is the time to earliest non-detectable viral load. The overall probability of sustained viral response can then be written as follows:

$$\begin{aligned} \text{E}\left[\text{E}\left[SVR|Z\right]\right] &= \pi \times \text{E}\left[SVR|Z=1\right] + (1-\pi) \times \text{E}\left[SVR|Z=0\right] \\ \Rightarrow \text{P}(SVR) &= \pi \times \text{E}\left[\text{E}\left[SVR|T, Z=1\right]\right] \\ \Rightarrow \text{P}(SVR) &= \pi \times \text{E}\left[\Delta(T)\right] \\ \Rightarrow &= \pi \times \text{E}\left[\text{logit}^{-1}\left(\beta_0 + \beta_1 T\right)\right] \end{aligned}$$

where $\Delta(T)$ is defined from the first submodel as the probability of SVR as a function of time to non-detectable viral load. The overall probability of SVR is the expected value of the conditional probability over the first RVL times, so for studies reporting just the overall probability of SVR the number of patients with sustained viral response in the intent-to-treat group is binomially distributed with probability $\pi_i \text{E}\left[\Delta(t)\right]$. We use MC integration within each iteration of the Gibbs sampler to calculate the expected value of the inverse logit from study i's Weibull parameter configuration.

Table 1: *Parameter elicitation: Parameters used, summarizing aggregate data characteristics.*

Parameter	Distributional Characteristics				
	Prior	Posterior Mean	Posterior SD	Lower HPD 5%	Upper HPD 95%
Logistic intercept, β_0	N(0, 100)	2.12	0.14	1.90	2.36
Logistic slope, β_1	N(0, 10)	−0.11	0.01	−0.12	−0.09
MTE $\sim$ Beta(π_a, π_b)	Gamma(0.01, 0.01)	31.87	24.69	5.79	81.65
MTE $\sim$ Beta(π_a, π_b)	Gamma(0.01, 0.01)	19.57	14.68	3.65	49.50
Weibull shape, γ	Gamma(0.01, 0.01)	1.70	0.11	1.52	1.89
Weibull scale, $\phi \sim$ Gamma(ϕ_a, ϕ_b)	Gamma(0.01, 0.01)	2.07	1.48	0.54	5.10
Weibull scale, $\phi \sim$ Gamma(ϕ_a, ϕ_b)	Gamma(0.01, 0.01)	62.35	48.46	10.56	162.91

References

* References are included in a two-page version of the poster available online at *www.metrumrg.com*.

Acknowledgments

The authors would like to thank Katey Durham and Neal Thomas of Pfizer for their insights during the formulation of the model.

 Copies of this poster are available at www.metrumrg.com/publications