

2019

Bayesian Hierarchical Meta-Analysis of Asymptomatic Ebola Seroprevalence

Peter Brody-Moore

Recommended Citation

Brody-Moore, Peter, "Bayesian Hierarchical Meta-Analysis of Asymptomatic Ebola Seroprevalence" (2019). *CMC Senior Theses*. 2228.
https://scholarship.claremont.edu/cmc_theses/2228

This Open Access Senior Thesis is brought to you by Scholarship@Claremont. It has been accepted for inclusion in this collection by an authorized administrator. For more information, please contact scholarship@cuc.claremont.edu.

Claremont McKenna College

**Bayesian Hierarchical Meta-Analysis
of Asymptomatic Ebola Seroprevalence**

Submitted to
Professor Mark Huber

by
Peter Brody-Moore

for
Senior Thesis
Spring 2019
4/29/2019

Abstract

The continued study of asymptomatic Ebolavirus infection is necessary to develop a more complete understanding of Ebola transmission dynamics. This paper conducts a meta-analysis of eight studies that measure seroprevalence (the number of subjects that test positive for anti-Ebolavirus antibodies in their blood) in subjects with household exposure or known case-contact with Ebola, but that have shown no symptoms. In our two random effects Bayesian hierarchical models, we find estimated seroprevalences of 8.76% and 9.72%, significantly higher than the 3.3% found by a previous meta-analysis of these eight studies. We also produce a variation of this meta-analysis where we exclude two of the eight studies. In this model, we find an estimated seroprevalence of 4.4%, much lower than our first two Bayesian hierarchical models. We believe a random effects model more accurately reflects the heterogeneity between studies and thus asymptomatic Ebola is more seroprevalent than previously believed among subjects with household exposure or known case-contact. However, a strong conclusion cannot be reached on the seriousness of asymptomatic Ebola without an international testing standard and more data collection using this adopted standard.

Contents

1	Introduction	2
2	The Bayesian Paradigm	3
2.1	Bayesian vs. Frequentist	3
2.2	Bayesian Hierarchical Models	5
2.2.1	General Hierarchical Structure	6
2.2.2	Joint Posterior for General Hierarchical Structure	6
2.3	Meta-Analysis with Bayesian Hierarchical Models	7
2.3.1	Exchangeability	8
3	Probabilistic Programming and Stan	8
3.1	Hamiltonian Monte Carlo	9
4	Models	11
4.1	Replication of Bower & Glynn	11
4.2	Non-Hierarchical Stan Model	12
4.3	Beta-Binomial with Natural Parameterization of α and β	13
4.4	Beta-Binomial with Exponential Priors on α and β	16
4.5	Beta-Binomial Natural Parameterization without Gabon Data	17
5	Model Validation	19
5.1	Testing Model on Simulated Data	19
5.2	Model Diagnostics	20
6	Conclusion	21
7	References	22
8	Appendix	24
8.1	Theta Posteriors from Section 4.3	24
8.2	Theta 95% Credible Intervals from Section 4.3	25
8.3	Theta Markov Chain Traceplots from Section 5.2	25

1 Introduction

Since its first outbreak in present day Democratic Republic of Congo, in 1976, Ebolavirus Disease (Ebola) has been studied extensively. Research is relatively difficult because outbreaks are rare and fatality rates from the virus are high. There have been roughly 29 outbreaks since the first in 1976 [3] and the mean fatality rate is 50% [1]. Yet, significant progress has been made in understanding how Ebola originates, spreads, and affects its hosts. More research is needed, however, in understanding asymptomatic Ebola infection, or Ebola that manifests itself without symptoms. Previous studies have reached a wide range of conclusions on the seriousness of asymptomatic Ebola. Because there is no standard method to test for asymptomatic Ebola seropositivity, even though the World Health Organization has stated that there is “an urgent need for one,” [7] there is often criticism on the legitimacy of study results. During the recent outbreak in West Africa between 2013 and 2016, asymptomatic Ebola was not an input into disease transmission models or considered when projecting intervention effects [18]. Thus, the current general consensus does not consider asymptomatic Ebola to be a serious or relevant issue when considering how to deal with the diagnosis and spread of Ebola.

Conducting more research on asymptomatic Ebola is important because it has potentially relevant effects on Ebolavirus transmission. A better understanding of the seroprevalence of asymptomatic Ebola would improve mathematical models of disease transmission, could be used to calculate a more accurate vaccination threshold, and would inform outbreak management procedures [15] [11]. If asymptomatic Ebolavirus infection is sufficiently uncommon, then it is not a major factor in the spread of the disease during outbreaks and public health officials need not spend resources trying to identify it for outbreak control. However, if it is more frequent than previously thought, then its omission in transmission models is problematic and more resources need to be dedicated to identifying asymptomatic infection. Further, the evidence is inconclusive on whether asymptomatic Ebola is infectious to others or whether Ebolavirus remains in those with asymptomatic Ebola for a long period after the outbreak [11]. The prevalence of asymptomatic Ebola could also provoke the need for study in these areas.

This paper conducts a novel meta-analysis of asymptomatic Ebola through a Bayesian framework. Using a variety of Bayesian hierarchical models, we build on a previous meta-analysis done by Bower & Glynn [7]. The authors use eight studies in their meta-analysis [18] [11] [2] [19] [14] [4] [9] [16]. All eight studies consist of data from groups with household exposure or known case-contact, meaning each tested patient had comparable exposure to someone diagnosed with Ebola. Furthermore, each study tested for seroprevalence of Ebolavirus antibodies by ELISA (enzyme-linked immunosorbent assay) or IFA (immunofluorescence antibody tests). Studies that tested by IFA had a cut-off $\geq 1:64$, generally accepted as a rigorous threshold for serotesting. However, Bower & Glynn admit “there is no definitive evidence that this an appropriate threshold” [7].

This paper first explores the Bayesian statistical paradigm and Bayesian hierarchical

models to understand why they are powerful for meta-analyses. Then, it will replicate the result of Bower & Glynn, compare this result to a fixed effects Bayesian model, and extend this thinking to two random effects Bayesian hierarchical models. Next, we test the sensitivity of these results to removing two questionable studies from the meta-analysis. Finally, we do a deep dive into the model validation and convergence diagnostics for one of our Bayesian hierarchical models and conclude with implications of our results.

2 The Bayesian Paradigm

2.1 Bayesian vs. Frequentist

There are two main schools of thought in statistical inference: Bayesian and Frequentist. Frequentist statistics has traditionally dominated the field and is what students largely learn in their first introduction to statistics. A typical calculus-based introductory statistics course may explore Bayes' theorem and credible intervals, but learning about Bayesian inference often ends there. In recent years, however, rises in computing power and the development of Markov Chain Monte Carlo (MCMC) techniques have allowed Bayesian statisticians to build novel methods for inference. This has led to a revival of Bayesian statistics as more people are exposed to these methods.

Bayesian statistics is grounded in Bayes' Theorem:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}.$$

$P(A|B)$ will be referred to as the posterior distribution, $P(B|A)$ as the likelihood, $P(A)$ as the prior, and $P(B)$ as the marginal likelihood. Intuitively, the posterior is the probability of our hypothesis A given data B, the likelihood is the probability of observing the data B given hypothesis A, the prior is the probability of our hypothesis A before collecting any data, and the marginal normalizes the distribution with the probability of observing the data B under all hypotheses.

It is helpful to understand the key differences between the two schools before probing deeper into Bayesian methods. Philosophically, Bayesians seek to model uncertainty by placing a probability distribution on hypotheses, while Frequentists do not. Bayesians see outcomes as drawn from distributions, while Frequentists see them as fixed.

For example, suppose ten coins are flipped, and seven of them come up as heads. We want to estimate A, the probability of flipping a head with this coin. A Frequentist would look at the data and give a point estimate for A of 0.7, the mean. A Bayesian, however, would put a prior on A, say $\text{beta}(25,25)$. This encodes the relatively strong belief of a fair coin, because the expectation of a beta distribution is $\alpha/(\alpha + \beta)$.

We are interested in $P(A|B) \propto P(B|A)P(A)$. The likelihood is:

$$\binom{10}{7}A^7(1-A)^3$$

And the density of the beta prior is:

$$\frac{\Gamma(50)}{\Gamma(25)\Gamma(25)}A^{24}(1-A)^{24}$$

Combining these, we get:

$$\binom{10}{7}A^7(1-A)^3 \frac{\Gamma(50)}{\Gamma(25)\Gamma(25)}A^{24}(1-A)^{24}$$

Pulling out the constants to focus on the kernel:

$$P(A|B) \propto A^7(1-A)^3A^{24}(1-A)^{24}$$

And finally, combining like terms:

$$P(A|B) \propto A^{31}(1-A)^{27}$$

With this prior and likelihood, we get a posterior distribution of $A \sim \text{beta}(32,28)$. This is a special case of calculating a posterior distribution because a beta prior is conjugate with binomial data. This means that the likelihood and prior distribution match up such that we do not have to calculate the marginal likelihood. After 10 flips, we have a posterior mean estimate of 32/60, or approximately 0.53. Thus, the data moved the posterior mean to 0.53 from a prior mean of 0.5 after seeing the data. Notice this philosophical difference in that after 10 data points, the Frequentist estimates $A = .7$ and the Bayesian estimates A to be drawn from a beta distribution with parameters $\alpha = 32$ and $\beta = 28$.

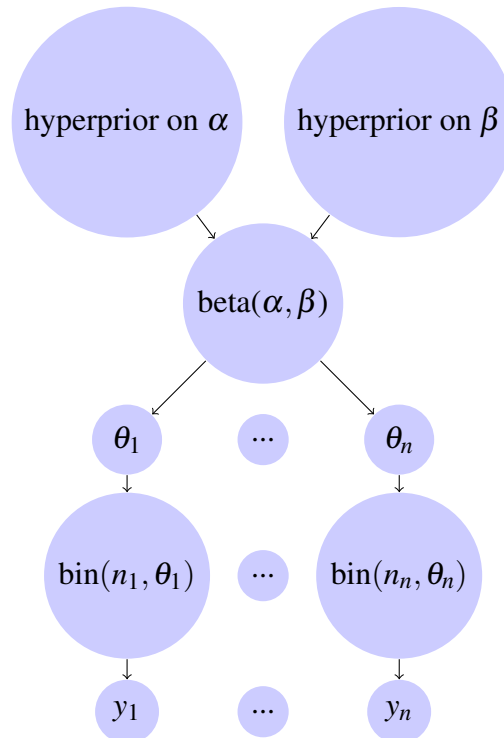
There are two distinct eras of Bayesian statistics: pre-computation and computation. Before the rise of computational power, Bayesian statistics was constrained to simple examples like the one above with a conjugate prior and likelihood. However, with the rise of computing power in recent decades, the development of MCMC techniques has given Bayesian statistics tremendous inferential power. Bayesian statistics stalled because you either had to use a conjugate prior to ignore the normalizing constant, or do increasingly complex multidimensional integrations to find the normalizing constant (because the denominator of Bayes' theorem is $\int p(\theta)p(x|\theta)d\theta$). The development of MCMC techniques such as Metropolis-Hastings and Gibbs Sampler allowed Bayesians to sample from conditional posterior distributions in order to create a stationary distribution. This stationary distribution hopefully converges to the posterior distribution of interest, often referred to as the target distribution.

Finally, one notable feature that will be relevant when interpreting model results is the difference between confidence intervals and credible intervals. A confidence interval

seeks to describe a range of values that will contain the true effect at least a certain percentage of the time. For example, we would hope a 95% confidence interval would contain the true value at least 95% of the time. If we replicated the experiment 100 times, then at least 95 of the experiments would contain the true value. This does not fit into the Bayesian paradigm because Bayesians see parameters as drawn from some probability distribution. A credible interval seeks to contain a certain percentage of the posterior probability distribution. If we want a 95% credible interval, then we give a lower and upper bound that contain 95% of the posterior distribution's probability. If the credible interval is equal-tailed, we can cut off 2.5% from either end. There are also variations such as the highest density region that tries to capture that 95% of the posterior in the smallest possible interval, but this can require a bit more work for complex posteriors.

2.2 Bayesian Hierarchical Models

One powerful method of Bayesian inference is the idea of Bayesian hierarchical modeling. This method uses multiple levels in order to develop estimates for parameters of interest. At their core, all hierarchical models do is draw from conditional probability distributions. A two-layer hierarchical model with a beta prior distribution on the thetas and binomial data can be visualized as follows:



The only data observed is y_i across i studies, where y_i is the number of successes out of n_i trials. The only items that are manually inputted are the hyperpriors on α and β . For example, if each α and β is drawn from an exponential distribution, then λ_1 and λ_2 are inputted as the respective exponential distribution parameters. Each θ_i that we are

estimating is considered to be independently drawn from the common beta distribution, whose parameters we estimate by drawing from conditional distributions. Even though we do not know the true theta values, we can estimate the parameters of the common beta distribution using the observed data.

2.2.1 General Hierarchical Structure

The general structure of our Bayesian hierarchical model is as follows:

$$y_i \sim p(y|\theta_i)$$

$$\theta_i \sim p(\theta|\phi)$$

$$\phi \sim p(\phi)$$

In the Bayesian hierarchical models that follow, we define these variables as:

- y_i are the number of patients that test seropositive in study i .
- θ_i is the estimate of seroprevalence in study i .
- ϕ is the common distribution that the θ 's are drawn from.
- $p(\phi)$ is the hyperprior we put on this common distribution.

2.2.2 Joint Posterior for General Hierarchical Structure

The joint posterior distribution in Bayesian hierarchical models can be decomposed to full conditional distributions. This provides some intuition for why conditional distributions are the building blocks for hierarchical models. These derivations are adapted from Jarad Niemi's lecture notes on Bayesian statistics [17].

We are interested in the joint posterior: $p(\theta, \phi|y)$:

$$p(\theta, \phi|y) \propto p(y|\theta, \phi)p(\theta, \phi) = p(y|\theta)p(\theta|\phi)p(\phi)$$

Looking across i studies and factoring out $p(\phi)$ because it does not change across the i 's, we get:

$$p(y|\theta)p(\theta|\phi)p(\phi) = [\prod_{i=1}^n p(y_i|\theta_i)p(\theta_i|\phi)]p(\phi)$$

And connecting the two, we get:

$$p(\theta, \phi|y) \propto [\prod_{i=1}^n p(y_i|\theta_i)p(\theta_i|\phi)]p(\phi)$$

Next, we want to break down this joint posterior distribution into separate components:

$$p(\theta, \phi | y) = p(\theta | \phi, y) p(\phi | y)$$

Taking $p(\theta | \phi, y)$, we can further decompose it to:

$$p(\theta | \phi, y) \propto p(y | \theta) p(\theta | \phi) = \prod_{i=1}^n p(y_i | \theta_i) p(\theta_i | \phi) \propto \prod_{i=1}^n p(\theta_i | \phi, y_i)$$

And taking $p(\phi | y)$, we can first decompose it to:

$$p(\phi | y) \propto p(y | \phi) p(\phi)$$

And break down $p(y | \phi)$ to:

$$p(y | \phi) = \int p(y | \theta) p(\theta | \phi) d\theta$$

Integrating over the thetas, we get:

$$\int \dots \int \prod_{i=1}^n [p(y_i | \theta_i) p(\theta_i | \phi)] d\theta_1 \dots d\theta_n$$

which is equivalent to:

$$\begin{aligned} & \prod_{i=1}^n \int p(y_i | \theta_i) p(\theta_i | \phi) d\theta_i \\ & = \prod_{i=1}^n p(y_i | \phi) = p(y | \phi) \end{aligned}$$

We see in this derivation that the joint posterior can be decomposed into dependencies among the parameters, allowing parameter estimation using the hierarchical structure.

2.3 Meta-Analysis with Bayesian Hierarchical Models

A meta-analysis seeks to make statistical inferences by drawing on multiple studies of the same or similar topics. The idea behind a meta-analysis is that more robust inferences can be made when one aggregates information from multiple sources. There are many ways to conduct a quantitative meta-analysis. A meta-analysis can be explored using a fixed effects model or a random effects model. The key difference between these two types of models is how they deal with heterogeneity. A fixed effects model assumes that an effect is fixed across all the scientific studies, ignoring any heterogeneity between the studies. A random effects model assumes that each study's effect is drawn from the same distribution for all the studies, accounting for the heterogeneity between the studies [6].

A meta-analysis can also be conducted on a spectrum, ranging from completely separate to completely pooled, or somewhere in the middle. If the studies are completely different, such that the outcomes from one do not inform the outcomes of the other, then they need to be analyzed separately. If the studies are considered to be identical, then they can be pooled and analyzed together to get one common effect. However, if the studies

are similar (or, more specifically, exchangeable), then an analysis can borrow information from the others without completely combining them. This is the case when utilizing a Bayesian Hierarchical Model for meta-analysis.

Bayesian Hierarchical Models provide a natural framework to conduct a meta-analysis. It allows for the borrowing of strength, allowing studies with a larger sample size to inform the inference on smaller studies [10]. In the context of biostatistics, this is especially useful because many studies are conducted with small sample sizes. In the eight studies we look at in our meta-analysis, sample sizes range from 24 to 404, allowing the bigger studies to inform the smaller studies. Furthermore, Bayesian hierarchical modeling is a compromise between pooling all the data and running separate analyses, allowing for more flexible inference. Pooling data implies that it is sampled from the same model, ignoring variation among studies. Analyzing unpooled data does not allow for any borrowing of strength between studies. Thus, Bayesian hierarchical models provides an ideal compromise of partial pooling that addresses these concerns.

2.3.1 Exchangeability

A challenge in meta-analysis is choosing which studies to include and which to exclude. One of the theoretical underpinnings for Bayesian hierarchical models is exchangeability. In our model, we are trying to estimate θ_i , the Ebolavirus seroprevalence in each of our eight studies indexed by i . Intuitively, exchangeability means that there is no prior knowledge to reason that θ_1 is different from θ_2 , that θ_2 is different from θ_3 , etc. Mathematically, following Gelman et al., we define the parameters $\theta_1, \dots, \theta_i$ as exchangeable if the joint distribution $p(\theta_1, \dots, \theta_i)$ is invariant to permutations of the indexes $(1, \dots, i)$ [10]. This is important because in hierarchical models, we claim that the thetas are drawn independently from a common population distribution (or that the thetas are independent, identically distributed random variables). All independently and identically distributed random variables are exchangeable, so exchangeability is a necessary assumption for hierarchical models.

3 Probabilistic Programming and Stan

This paper implements our hierarchical model using the probabilistic programming language Stan. The goal of probabilistic programming languages is to allow a complex statistical analysis to be conducted without the need for extensive programming [21]. With the emergence of extremely large datasets, fields such as machine learning have exploded. However, it can often be burdensome to generate inferences from these large datasets because of the need for highly efficient programming. Probabilistic programming languages solve this problem because they are highly declarative and automate many of the complex mathematical operations needed for complex statistical inference.

One of the probabilistic programming languages at the forefront of this movement

is Stan. Stan is an imperative programming language, making it more flexible than typical declarative probabilistic programming languages such as BUGS and JAGS [8]. It uses Hamiltonian Monte Carlo (or the No-U-Turn Sampler, abbreviated NUTS, an extension of Hamiltonian Monte Carlo), a sampler that has been shown to be more efficient and robust than Gibbs Sampling or Metropolis-Hastings for models with complex posteriors [12] [5].

The Stan language is organized into six code blocks: data, transformed data, parameters, transformed parameters, model, and generated quantities. In the data block, the user declares variables that are read in from the data. The transformed data block is similar to the data block, but is used for variables that are transformed through some operation. The parameters block allows the user to input the parameters that are sampled through Hamiltonian Monte Carlo (or NUTS). In the transformed parameters block, the user can generate new parameters through operations on the parameters declared in the parameters block. The model block allows the user to make statements that define the model using the inputted parameters. Finally, the purpose of the generated quantities block is to encode options for posterior inference. It is not executed until after the sample has been generated, and can be used to calculate posterior expectations, generate simulated data, and many more useful methods of posterior inference.

Prior to Stan, Hamiltonian Monte Carlo was considered to be effective at exploring high-dimensional probability spaces, but it was often difficult or inefficient to implement computationally. Because Hamiltonian Monte Carlo is rooted in analytic and differential geometry, it requires the computation of gradients and high-order derivatives [5]. As more dimensions are added to a model, these operations can take a long time to compute. Stan translates into C++ code before it is executed, allowing for quick and efficient model estimation without the need for C++ knowledge. Stan has made Hamiltonian Monte Carlo a practical and powerful Markov Chain Monte Carlo sampler.

3.1 Hamiltonian Monte Carlo

Much of the following information on Hamiltonian Monte Carlo (HMC) is adapted from Betancourt's, "A Conceptual Introduction to Hamiltonian Monte Carlo" [5]. The goal of MCMC is to approximate a posterior distribution from which we cannot sample directly. Traditional MCMC algorithms such as Metropolis-Hastings are effective, but have difficulty scaling as more dimensions are added to the model. With high-dimension models, more computational resources are needed to sufficiently explore the probability space and for the chain to converge. Metropolis-Hastings can reach a point where it is prohibitively difficult to run the chain long enough for it to converge. This is one of the main modeling improvements with the implementation of HMC through Stan.

The key to HMC's efficiency is making informed jumps, in contrast to Metropolis Hasting's random walk. Hamiltonian Monte Carlo generates a vector field using the probability space's geometry. It then uses this vector field to explore the space in a more deliberate manner than Metropolis Hastings. To build this vector field, the gradient of

the target density and an auxiliary momentum variable have to be computed. Betancourt develops an analogy where the target density's mode is a planet, the gradient is a gravitational force, and the space we want to explore is the orbit. Like gravity, the gradient points toward the mode, but the momentum variable prevents the exploration from going directly into the mode. Thus, exploration orbits around the mode in accordance with the gradient and momentum, leading to precise and efficient exploration of the probability space in an ideal model.

Before giving a general HMC algorithm, we need to define a few terms. In HMC, each dimension θ_j of the parameter space has an auxiliary momentum variable ϕ_j . By standard practice, we define ϕ with a multivariate normal distribution, $\phi \sim N(0, M)$. Here, M is a mass matrix of covariances that we input for the multivariate normal. Each iteration requires us to update θ and ϕ with L leapfrog steps. We also input ε , the scaling factor for the updates that affects step size. Each full update includes two half-updates of ϕ , with a full update of θ in between. Each half-update of ϕ requires the calculation of the gradient of the log-posterior density of θ given y .

Following Gelman et al., [10] we define the steps of an HMC iteration as:

1. Draw $\phi \sim N(0, M)$
2. Update (θ, ϕ) L times:
 - Make a first half-step for ϕ : $\phi = \phi + \frac{1}{2}\varepsilon \frac{d \log p(\theta|y)}{d\theta}$
 - Update θ : $\theta = \theta + \varepsilon M^{-1} \phi$
 - Make a second half-step of ϕ : $\phi = \phi + \frac{1}{2}\varepsilon \frac{d \log p(\theta|y)}{d\theta}$
3. Let θ^{t-1} and ϕ^{t-1} be the values of θ and ϕ before updating and let θ^* and ϕ^* be these values after updating. Calculate:

$$r = \frac{p(\theta^*|y)p(\phi^*)}{p(\theta^{t-1}|y)p(\phi^{t-1})}$$
4. Set $\theta_t = \theta^*$ with probability $\min(r, 1)$. Else, do not accept and let $\theta_t = \theta^{t-1}$

The Bayesian models in this paper use Stan to estimate our posterior distributions and the RStan interface to implement these models in the R programming language. We choose to use Stan's No-U-Turn Sampler (NUTS). As shown in the algorithm above, Hamiltonian Monte Carlo traditionally requires the statistician to input step size ε and number of steps L . Stan's NUTS makes it possible not to input any parameters by estimating likely candidate points across the target distribution (replacing the need to specify step size) and using primal-dual averaging to strategically adjust ε while the model explores the probability space [12]. We provide our Stan Code where relevant, and show how Stan models can be easily implemented directly in R.

4 Models

Following Bower & Glynn, we include data from eight different studies. Three of the studies are from the Democratic Republic of the Congo (DRC), two from Gabon, two from Sierra Leone, and one from Uganda. Observed estimates of seroprevalence range from 1% in Uganda during the 2007 outbreak, to 46% in Gabon during the 1996 outbreak, demonstrating significant heterogeneity between studies. Sample sizes range from 24 during the Gabon 1996 outbreak (the same study as the highest seroprevalence estimate) to 404 during the 1976 DRC outbreak. We consider these 8 studies to be exchangeable because all participants have household or known case-contact and because the varying testing mechanisms have sufficient rigor to identify seropositive subjects. There may be questions surrounding the two Gabon studies since they both took place during the same year, in the same country, and have by far the highest estimates (46% and 21%). We will address this in the last model when we remove these two studies and test how sensitive the parameter estimations are to this removal.

Table 1: Data

	samples	seropositive	country	start_year	end_year	perc
1	404	10	Congo (Democratic Republic of the)	1976	1976	0.02
2	101	4	Congo (Democratic Republic of the)	1995	1995	0.04
3	24	11	Gabon	1996	1996	0.46
4	56	12	Gabon	1996	1996	0.21
5	38	4	Congo (Democratic Republic of the)	2002	2002	0.11
6	210	2	Uganda	2007	2007	0.01
7	187	12	Sierra Leone	2015	2016	0.06
8	388	10	Sierra Leone	2015	2015	0.03

4.1 Replication of Bower & Glynn

We begin our analysis by replicating the results of Bower & Glynn’s paper, “A systematic review and meta-analysis of seroprevalence surveys of Ebolavirus infection” [7].

```
replication <- metaprop(data$seropositive,data$samples,  
                        sm="PFT",  
                        comb.random=FALSE,  
                        prediction=FALSE)
```

Table 2: Replication

	estimate	95% confidence interval
1	0.0	[.024,.044]

Bower & Glynn’s results show an overall estimate of seroprevalence in asymptomatic people with known case-contact of 3.3% with a 95% confidence interval of [2.4,4.4]. This fixed effects model uses the Freeman Tukey arcsine square root transformation method, which we replicate using the R package, ”Meta.” The authors also acknowledge the ”substantial heterogeneity due to three small studies with higher estimates” [7]. With this estimated seroprevalence of 3.3%, asymptomatic Ebola would not be a significant factor in disease transmission outbreak control.

4.2 Non-Hierarchical Stan Model

Our first Stan model is a non-hierarchical fixed effects model that provides a global estimate for θ , the estimate of seroprevalence. The main purpose of this model is to show that it is possible to run a fixed effects model in Stan and to give a direct comparison to the fixed effects model chosen by Bower & Glynn. We put a beta(1,19) prior on the θ value, giving a prior expectation for θ of .05. However, there are likely enough observations in the meta-analysis data such that this relatively weak prior will not hold much weight in the posterior estimates.

```
a <- 1
b <- 19

stan_model_non_hierarchical = "
data {
  int<lower=0> N; // the number of studies
  int<lower=0> samples[N]; // the number of samples in each of the N
                          studies
  int<lower=0> seropositive[N]; // the number of seropositive samples
                              in each of the N studies
  int a; // inputted parameter alpha of the beta hyperprior
  int b; // inputted parameter beta of the beta hyperprior
}

parameters {
  real theta; // the parameter of interest is the global estimate for
              seroprevalence
}

model {
  theta ~ beta(a,b); // declares the theta parameter to be distributed
                    with a beta prior
  for (n in 1:N) // loops over each study and declares the seropositive
                 count to be binomially distributed
    seropositive[n] ~ binomial(samples[n], theta);
}
"
```


”

Table 3: Non-Hierarchical Model

	mean	se mean	2.5%	97.5%
theta	0.046	0.000	0.036	0.058

After running the model, we get a global estimate of seroprevalence of 4.6% with a 95% credible interval of [3.6,5.8]. The posterior mean of θ is slightly lower than our prior expectation of 0.05, demonstrating that our data moved the posterior parameter estimates lower than our prior expectation. Also note that we are now using a credible interval because it is a Bayesian model. Thus, 95% of our posterior probability is captured between 3.6% and 5.8%. Again, we do not want to place too much emphasis on this model because it is not our focus, but it is interesting to see the variation among fixed effects models. These differences can be caused by the priors we set in a Bayesian setting, or by weighting the studies differently to compute the fixed effects estimate.

4.3 Beta-Binomial with Natural Parameterization of α and β

In our first random effects Bayesian hierarchical model, the population beta distribution $P(\theta|\phi)$ is conjugate to the binomial likelihood $P(y|\theta)$. We believe a random effects model more accurately captures the situation because the true estimate of seroprevalence is presumably different across studies. There are likely factors that affect seroprevalence differently in each study, although they are not immediately known (not harming exchangeability). We expect estimates of seroprevalence to be similar but not identical in each study, which is captured in our model.

We parameterize the beta distribution in a natural way, with the prior expectation as μ and the prior sample size as η . We put a hyperprior beta distribution on μ and a hyperprior exponential distribution on η . So, our only chosen inputs to the model are the hyperparameters α_0 and β_0 that parameterize the hyperprior beta distribution and λ that parameterizes the hyperprior exponential distribution. We choose $\alpha_0 = 1$ and $\beta_0 = 19$, so $\mathbb{E}[\mu] = .05$ and $\lambda = .05$ so $\mathbb{E}[\eta] = 20$.

In the Stan code, we include a transformed parameter block that computes α_1 and β_1 from η and μ . These α_1 and β_1 are the parameters for the beta prior on the θ 's. The expectation of a beta distribution is $\alpha/(\alpha + \beta)$. So, $\mu = \alpha_1/(\alpha_1 + \beta_1)$. The prior sample size of a beta distribution is $\alpha_1 + \beta_1$. We can substitute η to get $\mu = \alpha_1/\eta$ or $\alpha_1 = \mu \cdot \eta$. Similarly, we get $\beta_1 = \alpha_1 + \beta_1 - \alpha_1 = \alpha_1 + \beta_1 \cdot (1 - \alpha_1/(\alpha_1 + \beta_1)) = \eta \cdot (1 - \mu)$

Summarized, we have:

$$\begin{aligned}y_i &\sim \text{bin}(n_i, \theta_i) \\ \theta_i &\sim \text{beta}(\alpha_1, \beta_1) \\ \alpha_1 &= \eta \cdot \mu \\ \beta_1 &= \eta \cdot (1 - \mu) \\ \mu &\sim \text{beta}(1, 19) \\ \eta &\sim \text{exp}(.05)\end{aligned}$$

```
a = 1
b = 19
l = .05

stan_model_natural = "
data {
  // input data
  int<lower=0> N;
  int<lower=0> samples[N];
  int<lower=0> seropositive[N];

  // Inputted priors
  real<lower=0> a; // hyperparameter alpha
  real<lower=0> b; // hyperparameter beta
  real<lower=0> l; // hyperparameter lambda
}
parameters {
  real<lower=0,upper=1> mu; // prior expectation
  real<lower=0> eta; // prior sample size
  real<lower=0,upper=1> theta[N]; // estimate of seroprevalence
                                for each study
}
transformed parameters{
  real<lower=0> alpha;
  real<lower=0> beta;
  alpha = eta* mu ; // calculates alpha
  beta = eta*(1-mu); // calculates beta
}
model {
  mu ~ beta(a,b); // hyperprior on prior expectation
  eta ~ exponential(l); // hyperprior on prior sample size
```

```

theta ~ beta(alpha,beta); // thetas drawn from
                           beta distribution
seropositive ~ binomial(samples,theta); // seropositive number
                                         isbinomially distributed
}
"

```

Table 4: Beta-Binomial Natural Parameterization

	mean	se mean	2.5%	97.5%	observed probability
mu	0.09	0.00	0.05	0.17	-
eta	10.62	0.03	3.13	25.10	-
alpha	0.93	0.00	0.31	2.03	-
beta	9.69	0.03	2.70	23.30	-
theta[1]	0.03	0.00	0.01	0.04	0.02
theta[2]	0.04	0.00	0.01	0.09	0.04
theta[3]	0.35	0.00	0.19	0.54	0.46
theta[4]	0.20	0.00	0.11	0.30	0.21
theta[5]	0.10	0.00	0.03	0.20	0.11
theta[6]	0.01	0.00	0.00	0.03	0.01
theta[7]	0.07	0.00	0.04	0.10	0.06
theta[8]	0.03	0.00	0.01	0.05	0.03

Each of our Stan Models runs for 10,000 iterations, which includes a burn-in (referred to as warmup in Stan) of 1,000 iterations. The purpose of the burn-in is to deal with the situation where the Markov chain starts in a poor location. In this case, the chain may initially explore a low probability area, leading to poorly representative samples. Thus, we discard the burn-in iterations, then start collecting samples after the burn-in has ended and the chain has neared its equilibrium distribution.

After running the model, we see shrinkage, or the pulling of estimates towards the population mean. For example, in study 3, 46% of the study participants were seropositive. However, our posterior mean for θ_3 is .35, demonstrating that the other studies informed the estimate, resulting in a decreased estimate from the observed data. Posterior visualization and credible intervals for the eight theta parameters can be viewed in Sections 8.1 and 8.2 of the appendix.

Furthermore, we estimate θ 's to be drawn from a common beta population distribution with parameters $\alpha = .93$ and $\beta = 9.69$, making $\mathbb{E}[\theta] = .93/ (.93+9.69) = .0876$, or 8.76% in this model. We see a significantly higher estimate than the 3.3% estimated by Bower and Glynn and the 4.6% estimated by our fixed effects model. This is because the random effects model gives more weight to the smaller sample size studies that have high observed seroprevalence. Finally, our standard errors (posterior standard deviation divided by square root of effective sample size) are all small, signaling that we expect our estimate to be close to the true value.

4.4 Beta-Binomial with Exponential Priors on α and β

Here, we again have a conjugate population distribution and likelihood. However, we now place exponential hyperpriors on α and β . Our only set inputs to this model are the hyperparameters $\lambda_1 = 1$ and $\lambda_2 = 19$. With these hyperpriors, we get $\mathbb{E}[\alpha] = 1$ and $\mathbb{E}[\beta] = 19$. This gives the same prior expectation (.05) and sample size (20) for the population distribution as the model in Section 4.3, but with a different hierarchical structure and no transformed parameters.

```
lambda1 = 1
lambda2 = 1/19

stan_model_exponential = "
data {
  // input data
  int<lower=0> N;
  int<lower=0> samples[N];
  int<lower=0> seropositive[N];

  // Inputted priors
  real<lower=0> lambda1; // hyperparameter input into
                        // exponential hyperprior
  real<lower=0> lambda2; // hyperparameter input into
                        // exponential hyperprior
}
parameters {
  real<lower=0> alpha;
  real<lower=0> beta;
  real<lower=0,upper=1> theta[N];
}
model {
  alpha ~ exponential(lambda1); // exponential hyperprior on alpha
  beta ~ exponential(lambda2); // exponential hyperprior on alpha

  theta ~ beta(alpha,beta);
  seropositive ~ binomial(samples,theta);
}
"
```

Table 5: Beta-Binomial with Exponential Priors

	mean	se mean	2.5%	97.5%	observed probability
alpha	0.94	0.00	0.35	1.90	-
beta	8.73	0.03	2.20	21.00	-
theta[1]	0.03	0.00	0.01	0.04	0.02
theta[2]	0.04	0.00	0.01	0.09	0.04
theta[3]	0.36	0.00	0.19	0.55	0.46
theta[4]	0.20	0.00	0.11	0.31	0.21
theta[5]	0.10	0.00	0.03	0.21	0.11
theta[6]	0.01	0.00	0.00	0.03	0.01
theta[7]	0.07	0.00	0.04	0.10	0.06
theta[8]	0.03	0.00	0.01	0.05	0.03

We obtain very similar results to the previous model. The only change in the posterior means for the thetas is θ_3 increasing from .35 to .36. The posterior means for the α and β parameters are also similar, with $\alpha = .94$ and $\beta = 8.73$. With an estimated population distribution of $\text{beta}(.94, 8.73)$, we get $\mathbb{E}[\theta] = .0972$, slightly higher than the previous estimate of .0876. Again, our standard errors are low meaning we expect our estimates to be close to the true value. These results demonstrate that changes to the hierarchical structure do not significantly change the model results. Given expected seroprevalence estimates of 8.76% in our previous model and 9.72% in this model, asymptomatic Ebola is extremely relevant to disease transmission models.

4.5 Beta-Binomial Natural Parameterization without Gabon Data

In our final model, we remove two studies from the eight selected by Glynn & Bower. Both of the removed studies took place during the 1996 outbreak in Gabon with the first having 45.9% seroprevalence [14] and the second having 21.4% seroprevalence [4]. These two were much higher than the third highest calculation of 10.5%. Given that they were both taken during the same outbreak, there may be a factor that would cause these estimates to be so much higher and thus violate exchangeability. One possible explanation for this could be poor testing procedures that resulting in an unreasonably high observed seroprevalence. However, given our lack of scientific background, we will not make any claims to predict this factor.

If these two studies are not exchangeable with the other six, then they would be ineligible for inclusion in the meta-analysis. Furthermore, both these studies have small sample sizes, with 24 and 56 participants respectively. Since a random effects model gives more weight to smaller studies than a fixed effects model, the decision to include or exclude these studies is particularly relevant. We use the same natural hyperparameterization as the model in Section 4.3, but test to see how removing these two studies affects the results.

```

a = 1
b = 19
l = .05

stan_model_natural = "
// Same Stan Code as previous natural parameterization model with
  different data inputted
data {
// input data
int<lower=0> N;
int<lower=0> samples[N];
int<lower=0> seropositive[N];

// Inputted priors
real<lower=0> a;
real<lower=0> b;
real<lower=0> l;
}
parameters {
real<lower=0,upper=1> mu;
real<lower=0> eta;
real<lower=0,upper=1> theta[N];
}
transformed parameters{
real<lower=0> alpha;
real<lower=0> beta;
alpha = eta* mu ;
beta = eta*(1-mu);
}
model {
mu ~ beta(a,b);
eta ~ exponential(l);

theta ~ beta(alpha,beta);
seropositive ~ binomial(samples,theta);
}
"

```

Table 6: Natural Parameterization without Gabon Data

	mean	se mean	2.5%	97.5%	observed probability
mu	0.05	0.00	0.02	0.09	-
eta	38.60	0.14	8.50	97.16	-
alpha	1.70	0.01	0.46	4.02	-
beta	36.90	0.13	7.88	93.62	-
theta[1]	0.03	0.00	0.01	0.04	0.02
theta[2]	0.04	0.00	0.01	0.08	0.04
theta[3]	0.08	0.00	0.03	0.16	0.11
theta[4]	0.01	0.00	0.00	0.03	0.01
theta[5]	0.06	0.00	0.03	0.10	0.06
theta[6]	0.03	0.00	0.01	0.04	0.03

In these model results, we see two of our parameter estimates change. First, θ_3 decreases from .1 to .08. The observed probability is .11, so removing the two Gabon studies forced this parameter to shrink to an even lower estimate. Second, θ_5 in this model decreased from .07 to .06. In the model with Gabon data, it was actually forced upward from its observed value of .06, but removing the Gabon data brought it back down to its observed value of .06. Further, we estimate posterior means of 1.7 for α and 36.90 for β . With these values, the thetas have an expectation of $1.7/(1.7+36.9) = .044$, significantly lower than our previous two models. With the expected value of theta being 4.4% in this model, the results are more in line with the two fixed effects models. With this estimate, if there is suitable scientific rationale to remove the Gabon studies from the dataset, then asymptomatic Ebola is a much less significant concern than estimated in our two random effects models.

5 Model Validation

5.1 Testing Model on Simulated Data

Table 7: Credible Interval and True Value

	mean	2.5%	97.5%	true value
mu	0.055	0.030	0.093	0.054
eta	34.237	8.543	82.361	25.372
alpha	1.814	0.491	4.423	1.361
beta	32.423	7.879	78.321	24.011
theta[1]	0.049	0.028	0.076	0.047
theta[2]	0.033	0.015	0.058	0.031
theta[3]	0.012	0.000	0.039	0.004
theta[4]	0.103	0.067	0.146	0.110
theta[5]	0.042	0.023	0.065	0.039
theta[6]	0.046	0.007	0.115	0.038
theta[7]	0.045	0.016	0.088	0.047
theta[8]	0.079	0.052	0.112	0.082

To test the robustness of our modeling, we run the model on simulated data in which we know the true values. We choose the hierarchical model from Section 4.3, where we parameterize the population beta distribution in a natural way, with the prior expectation as μ and the prior sample size as η . We want to test our model’s ability to recover simulated values. To do this, we calculate a 95% credible interval and see if it contains the true parameter value.

After running the model on simulated data, we see that the 95% credible intervals cover the true parameter value for the 12 parameters of interest. Since we successfully covered these simulated values, the model passes the first test of model validation.

5.2 Model Diagnostics

Table 8: Beta-Binomial Natural Parameterization Full

	mean	se_mean	sd	X2.5.	X97.5.	n_eff	Rhat
mu	0.09	0.00	0.03	0.05	0.17	36000.00	1.00
eta	10.62	0.03	5.79	3.13	25.10	29544.06	1.00
alpha	0.93	0.00	0.45	0.31	2.03	36000.00	1.00
beta	9.69	0.03	5.43	2.70	23.30	27981.94	1.00
theta[1]	0.03	0.00	0.01	0.01	0.04	36000.00	1.00
theta[2]	0.04	0.00	0.02	0.01	0.09	36000.00	1.00
theta[3]	0.35	0.00	0.09	0.19	0.54	36000.00	1.00
theta[4]	0.20	0.00	0.05	0.11	0.30	36000.00	1.00
theta[5]	0.10	0.00	0.04	0.03	0.20	36000.00	1.00
theta[6]	0.01	0.00	0.01	0.00	0.03	36000.00	1.00
theta[7]	0.07	0.00	0.02	0.04	0.10	36000.00	1.00
theta[8]	0.03	0.00	0.01	0.01	0.05	36000.00	1.00

Here, we add in two model diagnostics to our output table: Rhat and n_eff. Following John Howell’s definitions in his Stan Best Practices Guide, Rhat is a statistic that quantifies the consistency of an ensemble of Markov chains [13]. An important distinction to make is that Rhat identifies when there are problems with convergence. It does not confirm that an MCMC algorithm has converged. Our models each run four Markov chains (the default for Stan). Consistency means that each chain explores similar regions of the space for each parameter. If there are inconsistencies among the space explored, then the Rhat statistic will increase, signaling problems with the estimation of the respective parameter. Generally, any Rhat < 1.1 signals a valid model in terms of consistent chains. The Rhat statistics for the 12 parameters are all one, signaling the four chains are exploring similar regions. This can be visualized in Section 8.3 of the appendix, where we display the traceplots of the eight theta parameters. We include the burn-in iterations in these plots to show that the initial movement is often far away from the typical space that the chains explore.

The Stan Reference Manual defines effective sample size as the number of independent

samples with the same estimation power as N autocorrelated samples, where N is the number of iterations for which our chain runs [20]. This is why to calculate standard errors, we divide posterior standard deviation by the square root of effective sample size, as opposed to dividing it by the square root of N . Looking at the number of effective samples is a measure of how efficiently a chain explores a parameter space. If a chain is slow, then the number of effective samples is low, which becomes a problem as one adds dimensions to a space because it requires more computation to explore. Generally, the statistic we use to gauge a chain's effectiveness is number of samples per iteration or N_{eff} / N . For every parameter except for η and β , our number of effective samples equals 36,000, the number of iterations (10,000 iterations per chain, 4 chains, minus 1,000 burn-in for each chain). Generally, this means that our chain is exploring the parameter space very rapidly. Although this is somewhat unusual, it is not sufficient reason to disregard the model results.

6 Conclusion

Our three random effects models yield mixed results. When we include the Gabon data, we have expected seroprevalence estimates of 8.76% in our previous model and 9.72% (remembering that these are the expected values of the beta distributions that we consider each θ to be draw from). With these estimates, asymptomatic Ebola is a serious and relevant issue. It provokes the need for more research to understand whether asymptomatic Ebola is infectious to others and to determine how long Ebolavirus remains in those with asymptomatic Ebola after the outbreak. This would further inform understanding in why outbreaks may begin and also lead to more accurate transmission models to predict the spread of Ebola. Once this is known, further research can be done for the best way to prevent and contain Ebola outbreaks.

However, if the Gabon data is invalid, or at least not exchangeable, then our results are significantly impacted. When we remove the Gabon data, we have an expected seroprevalence estimate of 4.4%. This is more in line with previous estimates. And, given those estimates, researchers have essentially decided to exclude asymptomatic Ebola from their transmission dynamic models. Thus, if there is sufficient scientific reason to exclude the Gabon data, asymptomatic Ebola is sufficiently negligible that it can continue to be excluded.

Given these two scenarios, the next step needs to be developing an international standard testing mechanism and using this mechanism for more data collection. This will prevent or reduce disputes about the quality of the data. Ideally, data could be collected using this new standard during outbreaks in different countries. Then, a future meta-analysis could be conducted in which we could get a more accurate idea of asymptomatic Ebola seroprevalence without the limitations stated in this paper. This will take time, effort, and resources, but it is necessary to further our knowledge and understanding of Ebola.

7 References


- [1] Ebola virus disease. World Health Organization. <https://www.who.int/news-room/fact-sheets/detail/ebola-virus-disease>, Feb 2018.
- [2] No authors listed. Ebola haemorrhagic fever in Zaire, 1976. *Bull. World Health Organ.*, 56(2):271–293, 1978.
- [3] Mackenzie Bean. Timeline of global Ebola outbreaks, 1976-present, Dec 2018.
- [4] E. Bertherat, A. Renault, R. Nabias, G. Dubreuil, and M. C. Georges-Courbot. Leptospirosis and Ebola virus infection in five gold-panning villages in northeastern Gabon. *Am. J. Trop. Med. Hyg.*, 60(4):610–615, Apr 1999.
- [5] Michael Betancourt. A Conceptual Introduction to Hamiltonian Monte Carlo, 2017.
- [6] Michael Borenstein. *Introduction to Meta-Analysis*. Wiley, 2009.
- [7] H. Bower and J. R. Glynn. A systematic review and meta-analysis of seroprevalence surveys of ebolavirus infection. *Sci Data*, 4:160133, 01 2017.
- [8] Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1), 2017.
- [9] D. V. Clark, H. Kibuuka, M. Millard, S. Wakabi, L. Lukwago, A. Taylor, M. A. Eller, L. A. Eller, N. L. Michael, A. N. Honko, G. G. Olinger, R. J. Schoepp, M. J. Hepburn, L. E. Hensley, and M. L. Robb. Long-term sequelae after Ebola virus disease in Bundibugyo, Uganda: a retrospective cohort study. *Lancet Infect Dis*, 15(8):905–912, Aug 2015.
- [10] Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin. *Bayesian Data Analysis*. CRC Press, 3 edition, 2014.
- [11] J. R. Glynn, H. Bower, S. Johnson, C. F. Houlihan, C. Montesano, J. T. Scott, M. G. Semple, M. S. Bangura, A. J. Kamara, O. Kamara, S. H. Mansaray, D. Sesay, C. Turay, S. Dicks, R. E. G. Wadoum, V. Colizzi, F. Checchi, D. Samuel, and R. S. Tedder. Asymptomatic infection and unrecognised Ebola virus disease in Ebola-affected households in Sierra Leone: a cross-sectional study using a new non-invasive assay for antibodies to Ebola virus. *Lancet Infect Dis*, 17(6):645–653, 06 2017.
- [12] Matthew D. Hoffman and Andrew Gelman. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo, 2011.
- [13] John Howell. Stan Best Practices, Oct 2017.

- [14] E. M. Leroy, S. Baize, V. E. Volchkov, S. P. Fisher-Hoch, M. C. Georges-Courbot, J. Lansoud-Soukate, M. Capron, P. Debre, J. B. McCormick, and A. J. Georges. Human asymptomatic Ebola infection and strong inflammatory response. *Lancet*, 355(9222):2210–2215, Jun 2000.
- [15] P. Mbala, M. Baguelin, I. Ngay, A. Rosello, P. Mulembakani, N. Demiris, W. J. Edmunds, and J. J. Muyembe. Evaluating the frequency of asymptomatic Ebola virus infection. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.*, 372(1721), May 2017.
- [16] S. Mulangu, M. Borchert, J. Paweska, A. Tshomba, A. Afounde, A. Kulidri, R. Swanepoel, J. J. Muyembe-Tamfum, and P. Van der Stuyft. High prevalence of IgG antibodies to Ebola virus in the Ef pygmy population in the Watsa region, Democratic Republic of the Congo. *BMC Infect. Dis.*, 16:263, 06 2016.
- [17] Jarad Niemi. STAT 544 Bayesian Statistics. <http://www.jarad.me/courses/stat544/>. [Online; accessed 2-April-2019].
- [18] E. T. Richardson, J. D. Kelly, M. B. Barrie, A. W. Mesman, S. Karku, K. Quiwa, R. H. Marsh, S. Koedoyoma, F. Daboh, K. P. Barron, M. Grady, E. Tucker, K. L. Dierberg, G. W. Rutherford, M. Barry, J. H. Jones, M. B. Murray, and P. E. Farmer. Minimally Symptomatic Infection in an Ebola 'Hotspot': A Cross-Sectional Serosurvey. *PLoS Negl Trop Dis*, 10(11):e0005087, 11 2016.
- [19] A. K. Rowe, J. Bertolli, A. S. Khan, R. Mukunu, J. J. Muyembe-Tamfum, D. Bressler, A. J. Williams, C. J. Peters, L. Rodriguez, H. Feldmann, S. T. Nichol, P. E. Rollin, and T. G. Ksiazek. Clinical, virologic, and immunologic follow-up of convalescent Ebola hemorrhagic fever patients and their household contacts, Kikwit, Democratic Republic of the Congo. Commission de Lutte contre les Epidemies Kikwit. *J. Infect. Dis.*, 179 Suppl 1:28–35, Feb 1999.
- [20] Stan Development Team. Stan Modeling Language Users Guide and Reference Manual. <http://mc-stan.org>, 2018.
- [21] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. An Introduction to Probabilistic Programming, 2018.

8 Appendix


8.1 Theta Posteriors from Section 4.3

Posterior for Theta 1




Asymptomatic Ebola Seroprevalence

Posterior for Theta 2




Asymptomatic Ebola Seroprevalence

Posterior for Theta 3




Asymptomatic Ebola Seroprevalence

Posterior for Theta 4




Asymptomatic Ebola Seroprevalence

Posterior for Theta 5




Asymptomatic Ebola Seroprevalence

Posterior for Theta 6




Asymptomatic Ebola Seroprevalence

Posterior for Theta 7



Asymptomatic Ebola Seroprevalence

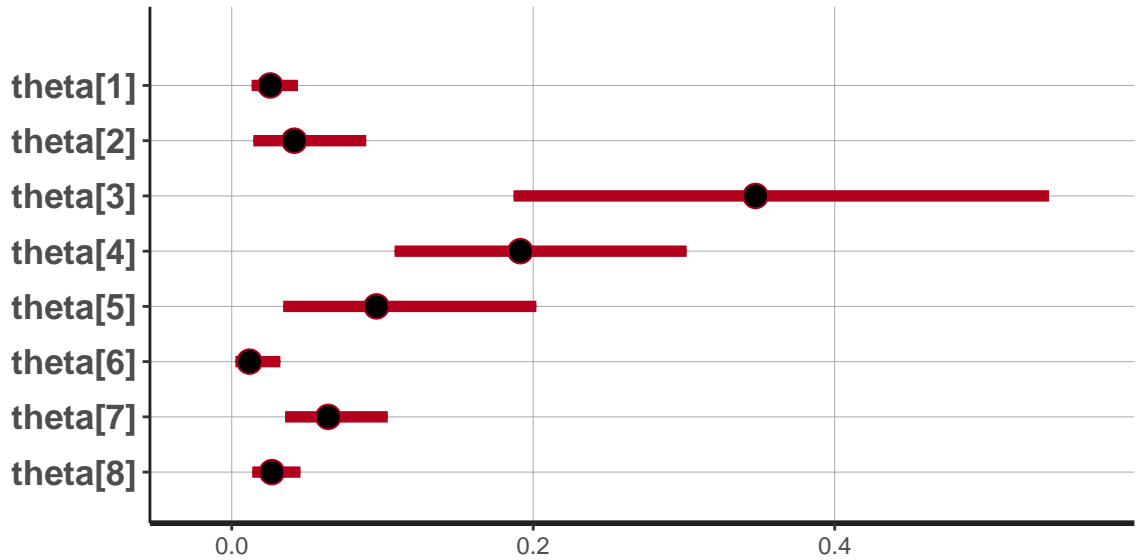
Posterior for Theta 8



Asymptomatic Ebola Seroprevalence

8.2 Theta 95% Credible Intervals from Section 4.3

```
## ci_level: 0.95 (95% intervals)
## outer_level: 0.95 (95% intervals)
```



8.3 Theta Markov Chain Traceplots from Section 5.2

