

©Copyright 2013

Elizabeth A. Teeple

Adjusting for Misclassified Outcomes in a Multistate Model

Elizabeth A. Teeple

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

Elizabeth R. Brown, Chair

Barbra A. Richardson

Susanne May

Program Authorized to Offer Degree:
School of Public Health - Department of Biostatistics

University of Washington

Abstract

Adjusting for Misclassified Outcomes in a Multistate Model

Elizabeth A. Teeple

Chair of the Supervisory Committee:
Research Associate Professor Elizabeth R. Brown
Department of Biostatistics

In this dissertation, we present a new model that accounts for time-dependent sensitivity in multivariate outcome survival models. Our problem was motivated by an interest in estimating baseline risk and associated covariates of time to postnatal transmission of HIV from mother to child (MTCT). Studies of postnatal transmission may be subject to bias due to several factors, including informative censoring due to death, the competing risk of weaning, and imperfect sensitivity of diagnostic testing for HIV in infants. While these issues have been partially addressed in previous literature, no single approach addresses all three. We propose to jointly model time to postnatal infection, time to death, and time to weaning as three separate outcomes in a multistate model, which is a framework often used to describe the progression of a process through successive events. Conditional on the underlying multistate model, we define an additional random variable to measure the probability of detection as a function of time since infection, and the cumulative distribution of this random variable serves as a sensitivity function that increases over time.

We focus first on jointly modeling time to postnatal infection and time to death as two outcomes of an illness-death model. To incorporate sensitivity, we define an exponential random variable for the delay in detection, with the exponential rate determined a priori from the literature. We specify a full likelihood by parameterizing the baseline transition intensities as flexible penalized spline functions and numerically calculate maximum likelihood

estimates for all spline parameters and covariates of interest.

We next expand the illness-death model to include time to weaning as an additional outcome of the model. As before, we specify an exponential delay random variable and define spline-based baseline transition intensities. To estimate the parameters, we implement a Bayesian MCMC algorithm, using the prior for the spline coefficients to impose smoothness. Finally, we compare the extended illness-death model to the simpler illness-death model and to ad-hoc methods often used in the applied literature.

TABLE OF CONTENTS

	Page
List of Figures	ii
List of Tables	vi
Chapter 1: Introduction	1
Chapter 2: The illness-death model	4
2.1 Introduction	4
2.2 Methods	6
2.3 Simulations	17
2.4 Application	20
2.5 Discussion	33
2.6 Appendices	37
Chapter 3: An extended illness-death model	40
3.1 Introduction	40
3.2 Methods	42
3.3 Simulations	48
3.4 Application	51
3.5 Discussion	68
3.6 Appendices	71
Chapter 4: Comparison to univariate survival models	81
4.1 Introduction	81
4.2 Methods	83
4.3 Results	84
4.4 Discussion	87

LIST OF FIGURES

Figure Number	Page
2.1 The illness-death model for MTCT of HIV. There are three states (HIV-negative and breastfeeding, illness, death) and three possible transitions. . . .	7
2.2 Comparison of the Markov and semi-Markov assumptions.	9
2.3 Example of a sensitivity function. Given a true onset time $t_{12} = 0.5$, a diagnostic test given at 2 months is detectable if the delay in detection is less than 1.5 months. Assuming the delay is distributed as exponential with rate 1.5 months, this translates to 90% sensitivity.	12
2.4 Conditional on time of illness at t_{12} , this diagram depicts the probability that the delay is longer than $12 - t_{12}$ months. When the assumed mean delay of detection is long, the range of probable true illness times is extended earlier.	28
2.5 Coefficient estimates and 95% confidence intervals for $\beta_{12}, \beta_{13}, \beta_{23}$, across assumed lag time distributions. Within each lag time, level of smoothing increases from left to right. CV scores marked with “+” correspond to the plotted estimates.	30
2.6 Estimated smoothed mortality curves for the healthy to death transition, assuming 0, 3, and 6 week delay in detection, and estimated over several degrees of smoothing. Observed event times are indicated as points. The histogram represents the distribution of weaning times of those who do not test positive during the study period. The slight increase in mortality at 8 months could be a delayed result of infants weaned at 6 months.	31
2.7 Estimated smoothed mortality curves for the illness to death duration transition, assuming 0, 3, and 6 week delay in detection, and estimated over several degrees of smoothing. Observed event times (time of death - time of first positive test) are indicated as points.	31
2.8 Cumulative incidence of illness from 100 randomly generated data sets, with parameters estimated assuming a 3-week delay sensitivity curve. These plots are stratified by treatment group, with antibiotics on the left and placebo on the right. The true CIFs for both antibiotics and placebo are in red on both plots.	32

2.9	Cumulative incidence of healthy-to-death transition from 100 randomly generated data sets, with parameters estimated assuming a 3-week delay sensitivity curve. These plots are stratified by treatment group, with antibiotics on the left and placebo on the right. The true CIFs for both antibiotics and placebo are in red on both plots.	32
2.10	Estimated transition intensity coefficients, under the Markov assumption. The assumed delay in detectability follows an exponential distribution with mean parameters 0, 3, and 6 weeks, and the smoothing parameter increases from left to right, starting at no smoothing in the far left.	34
2.11	Example transition intensities from healthy to death, under the Markov assumption, over three levels of smoothing. The points are observed deaths for infants who did not test positive prior to death. The assumed delay in detectability follows an exponential distribution with mean parameters 0, 3, and 6 weeks. When a 0 or 3 week delay is assumed, the intensities are smoothed down to zero.	35
2.12	A sample of the generated random sensitivity curves. The gray lines are the sensitivity curves for a subset of infants in the study, with the pointwise mean curve in black. The red line is the assumed sensitivity curve used in the likelihood. Increasing variability is denoted A-D.	39
3.1	Possible transitions in the mother-to-child transmission of HIV.	42
3.2	Posterior medians and HPD intervals for each covariate estimate. Columns contain the transitions and rows contain the covariates intercept, randomization, maternal CD4 cell count, log viral load, and weight at delivery, infant weight at birth, and indicators for sites at Lusaka, Dar es Salaam, and Blantyre (with Lilongwe as the baseline site). CD4 counts, log viral load, and maternal/infant weights are centered and scaled, so that effect size is interpreted on the unit standard deviation scale. The intercept term for the weaned/non-breasted to death transition is the constant for proportionality to the healthy to death transition. The gray line indicates a null effect. Each plot contains four intervals, with the far left indication the estimation assuming perfect sensitivity, increasing to an assumed delay with a mean of 6 weeks on the far right. *Mothers at the Dar site were counseled to wean when their infants were 6 months old.	60
3.3	Estimated weaning intensities from the model fit in Section 3.4.4. In this fit, the assumed mean delay was 3 weeks.	61
(a)	Estimated weaning intensities for the three sites that did not counsel early weaning. The estimated spline coefficients were taken as the median of the posterior distribution for each coefficient.	61
(b)	Estimated weaning intensity for Dar, where mothers were counseled to wean at 6 months.	61

3.4	Mortality transition intensities for transitions from the healthy and death states.	62
(a)	Baseline healthy to death transition intensity, across levels of smoothing. The points indicate observed deaths among infants who did not test positive or wean before death. As the assumed mean delay of the sensitivity curve increases, early deaths are weighted as deaths from illness, resulting in a decreased estimate of the healthy to death transition intensity. Mortality rates are similar across sensitivity curves after 6 months.	62
(b)	Baseline illness to death transition intensity, across levels of smoothing. Time is duration of illness prior to death. The points indicate observed illness duration (measured as time of death - time of first positive test) among infants who tested positive before death. After accounting for misclassification, there is a spike in mortality rates at 6 months and 10 months after the onset of illness.	62
3.5	Non-parametric cumulative incidence functions for the time to first positive test in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained within in the CIFs for the simulated data sets across all sites.	64
3.6	Non-parametric cumulative incidence functions for the time to weaning in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained in the CIFs for the simulated data sets across all sites.	65
3.7	Non-parametric cumulative incidence functions for the time to death from the healthy state in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained in the CIFs for the simulated data sets across all sites.	66
3.8	Posterior medians and HPD intervals for each covariate estimate of the illness-death model. Most estimates were similar when compared to the full MTCT model (Figure 3.2). Columns contain the transitions and rows contain the covariates intercept, randomization, maternal CD4 cell count, log viral load, and weight at delivery, infant weight at birth, and indicators for sites at Lusaka, Dar es Salaam, and Blantyre (with Lilongwe as the baseline site). CD4 counts, log viral load, and maternal/infant weights are centered and scaled, so that effect size is interpreted on the unit standard deviation scale. The gray line indicates a null effect. Each plot contains four intervals, with the far left indication the estimation assuming perfect sensitivity, increasing to an assumed delay with a mean of 6 weeks on the far right. *Mothers at the Dar site were counseled to wean when their infants were 6 months old.	69

3.9	Example trace plots for the log baseline rate of illness and coefficients for the eight covariates of interest for the illness transition. All chains appear to have mixed well. Because of early weaning at the Dar site, there were no positive tests after 6 months, and the lack of information led to a skewed posterior for rate of illness in Dar.	79
3.10	Posterior and prior distributions for the log baseline rate of illness (left) and the maternal CD4 count (right) coefficients. In both cases, the resulting posterior distributions (represented by black solid line) appear to be data driven, with little influence from the non-informative priors (represented by the gray dotted line).	80
4.1	Point estimates and 95% HPDs/CIs from Bayesian estimation of the MSM and univariate models, respectively. The first four columns are results of the time to postnatal illness transition, and the remaining columns are results from the time from infection to death transition. Within models, from left to right, results are from the sensitivity-adjusted MSM that includes weaning, univariate model that censors at the first occurrence of weaning, death, or last visit, sensitivity-adjusted ID model, and univariate model that censors death or last visit only. Sensitivity-adjusted MSM/ID models are fit assuming no delay in detection and mean delays of 3, 4, and 6 weeks. Univariate models are fit conditional on observing a last negative test on or after 0, 2, 4, and 6 weeks of age. Estimates marked by a triangle are comparable across appropriate models.	88

LIST OF TABLES

Table Number	Page
<p>2.1 A summary of simulation results for the mean, standard error, and coverage of the illness transition intensity parameters. The mortality intensities are not smoothed. The far left column indicates the data were generated assuming mean delays of 3 and 6 weeks. The next column lists the assumed mean delays (0, 3, and 6 weeks) used in estimation.</p>	19
<p>2.2 Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (-0.69, 0, -0.22)$ and an average 3 week delay (in bold). Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t - t_{12}))$, respectively. The semi-Markov model is assumed.</p>	21
<p>2.3 Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (-0.69, 0, -0.22)$ and an average 6 week delay (in bold) Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.</p>	22
<p>2.4 Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (0.4, 0, 0)$ and an average 6 week delay (in bold) Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.</p>	23
<p>2.5 Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated individual sensitivity curves assuming an mean 3 week delay sensitivity curve. The case where all observations are generated assuming the same sensitivity curve is denoted “none”, and the cases A-D denote individual sensitivity curves with increasing variability. Estimated models assume a single sensitivity curve with a mean 3 week delay and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.</p>	24

2.6	Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated individual sensitivity curves assuming an mean 6 week delay sensitivity curve. The case where all observations are generated assuming the same sensitivity curve is denoted “none”, and the cases A-D denote individual sensitivity curves with increasing variability. Estimated models assume a single sensitivity curve with a mean 6 week delay and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.	25
3.1	Results from 100 simulated data sets examining the robustness of sensitivity curve specification. Each column contains bias (observed - expected), standard error, and 95% HPD interval coverage for data generated under a specific sensitivity curve. From left to right, data were generated assuming a 3, 6, and 6.5 mean delay in detection. Rows contain the estimated results for four parameters of interest (indicated by transition, type of data, and true value). The covariates represent baseline transition intensity, treatment on illness rate, CD4 on illness rate, and treatment on infected mortality rates. The models were all estimated with no delay assumed or an assumed mean delay of 3, 6, and 6.5 weeks. Blank entries were not estimated. Bold entries indicate the model was estimating assuming the true sensitivity curve.	52
3.2	Specification of the log transition intensity $\log \alpha_{hj}(t) = \sum_{k=1}^m \beta_k^{(hj)} z_k + \sum_{k=1}^{k+4} \theta_k^{(hj)} B_k^{(hj)}(t)$. The intercept term of the (34) transition is the proportionality constant so that $\alpha_{34}(t - t_\omega) = \theta_{34} \alpha_{14}(t)$, where t_ω is time of weaning.	55
3.3	Basic descriptive statistics for the observed transitions of the MTCT model, stratified by location. In the first transition (illness, weaning, or death), if an infant experienced more than one event, only their first event is counted. Blantyre has a low rate of postnatal first positive tests after 2 months of age, while Lusaka remains relatively higher over the course of the study. There are no first positive tests in Dar after 6 months of age, since most infants have weaned.	56
4.1	Sample sizes of the univariate models. In the univariate time to postnatal infection model, infants are included in the sample conditional on having a negative test at 0-6 weeks (left column), and the final sample size is listed in the <i>Number negative</i> column. Infants who test positive are categorized according to Becquet et al. [2012]. In the univariate elapsed time from HIV infection to death model, the sample is restricted to all infants who test positive during the study, and the unknown category is grouped with peripartum infections to define time of infection.	87

ACKNOWLEDGMENTS

I would like to sincerely thank my committee for their never-ending support. My advisor, Dr. Elizabeth Brown, was patient in our discussions whether I was stuck on a concept, slow in writing, and all times in between, and for that, I am appreciative. I look forward to our continued collaboration and your guidance through this new part of my career. My committee members, Dr. Barbra Richardson and Dr. Susanne May, constructively discussed my work with me from the big picture down to the little details. And to Barb, I especially appreciate your insights from your own experiences of the everyday matters in this field.

I would also like to thank the faculty of the Department of Biostatistics for their thoughtful instruction and variety of perspectives. It's just as important to effectively convey the story as it is to rigorously support the message. To Gitana and the staff of the department, I always knew you'd be there to make everything go. My thanks go also to my fellow students, family and friends.

DEDICATION

To my family.

Chapter 1

INTRODUCTION

In studies of disease progression, the effect that an intervention has on both the risk of illness and the risk of any unintended harmful outcomes, such as mortality, is often of interest. Survival analyses of clinical trials often combine several distinct outcomes for inference on a combined outcome, which may not be the true outcome of interest. Several recent examples include studies in the fields of hematology (recurrent blood clots or mortality [Gauthier et al., 2013]), oncology (recurrence of liver cancer or mortality following transplant [Menon et al., 2013]; recurrence of breast cancer or mortality across drug treatment [Hadji et al., 2013]), and cardiology (recurrence of cardiac events or mortality across drug treatment [Konishi et al., 2010]). Alternatively, continuous time multistate models (MSMs) provide a framework to describe the progression of a disease through successive events, and each event time is modeled as a separate outcome. However, if one or more of the event times are measured with an imperfect diagnostic test administered at discrete visit times, the time of progression from one disease state to the next cannot be observed exactly. In order to directly model disease progression, existing multivariate methods must be modified to account for the discrete testing schedule and the imperfect sensitivity of the test. In this dissertation, we introduce new methods that incorporate time-dependent sensitivity in MSMs. We motivate this model with a problem arising in HIV infection and disease progression in infants who acquire HIV from their mothers.

Mother-to-child transmission of HIV (MTCT) can occur while the infant is in utero, during delivery (intrapartum), or throughout the duration of breastfeeding. Most infants born to HIV-positive mothers will test positive using an antibody test because of maternal HIV antibodies that circulate in the infant for the first year, resulting in a high false positive

rate [Wessman, Theilgaard, and Katzenstein, 2012]. Infant testing must then be done using RNA or DNA assays that depend on circulating viral load. These assays may be unable to detect infection when the viral load is too low; however, we expect the sensitivity of these tests to increase over time as the virus replicates. Once maternal antibodies have waned, an antibody-based test can be considered as a gold-standard. Thus, it is possible to study the sensitivity of virological tests in non-breastfeeding populations, where exposure to HIV is clearly limited to the in utero/intrapartum period. Numerous studies have been able to establish a time-varying sensitivity in this particular subgroup. These studies have been conducted with various combinations of diagnostic test (for example, RNA, DNA, p24 antigen), how the sample is stored (serum sample or dried blood spot), prevalent circulating HIV type, and level of prophylaxis anti-retroviral use [Dunn et al., 2000, Lambert et al., 2003, Fiscus et al., 2007, Cachafeiro et al., 2009, Zhang et al., 2008]. The pattern was similar across all studies, where sensitivity was low at birth (anywhere from 10%-60%) and rising quickly to 80%-90% at 2-6 weeks after birth.

In breastfeeding populations a previously uninfected infant continues to be exposed to HIV through breastmilk. As a result, an infant who tests negative at birth and positive soon after could have been infected either by in utero/intrapartum or postnatal transmission. In practice, due to the low sensitivity of HIV-diagnostic tests immediately after birth, analyses that focus on transmission via breastfeeding either classify early infections as unknown timing [Becquet et al., 2012] or restrict the sample to infants who test negative at 4-6 weeks [see, for example, Charurat et al., 2009]. This results in the misclassification of some infants and precludes our ability to examine the early postnatal risk period. Since it is of interest to estimate risk over the full exposure period, adjusting for sensitivity is one approach to correct for misclassification. To overcome this deficiency, several authors have approached modeling transmission risk over the full exposure period by adjusting for sensitivity as a function of time elapsed since the true time of onset, imposing functional forms that increase over time [Balasubramanian and Lagakos, 2003, Gupte et al., 2007, Brown, 2010, Brown

and Chen, 2010]. In this dissertation, we extend their methods to the multistate model setting, assuming the sensitivity curve is known.

In Chapter 2, we present an illness-death model that adjusts for time-dependent sensitivity (Figure 2.1). To maintain a flexible, yet completely parametric model, we assume the baseline transition intensities are penalized cubic splines, estimated by maximizing the likelihood [Joly et al., 2002]. We assess the assumption of a known sensitivity function through simulations and demonstrate our approach using data from a clinical trial (HPTN 024) conducted in Africa in 2001-2003.

In Chapter 3, we extend the model to include weaning as an additional transition and adjust this extended illness-death model to account for time-dependent sensitivity (Figure 3.1). As before, we specify a fully parametric, yet flexible, model through the use of penalized cubic-spline baseline intensities. Here, we follow the Bayesian approach of Kneib and Hennerfeind [2008], which estimates the smoothing parameter as part of the MCMC process. Additionally, because our model is completely parametric, we are able to include discrete and continuous covariates. As in Chapter 2, we assess the assumption of a known sensitivity function through simulations and demonstrate our approach using the HPTN 024 data, comparing our results to both the illness-death model and simpler ad-hoc methods often employed in the applied literature. Finally, in Chapter 4, we compare the results of the Bayesian implementation of the illness-death and multistate models of Chapter 3 to results of the well-known univariate outcome survival techniques used in the applied MTCT literature that rely on ad-hoc approaches to account for low sensitivity, competing risks, and interval censoring.

Chapter 2

THE ILLNESS-DEATH MODEL**2.1 Introduction**

In this chapter, we introduce an illness-death model that adjusts for imperfect sensitivity. Our model is motivated by analyses of postnatal mother-to-child transmission of HIV (MTCT). There are two main aspects we address in our approach. First, it is well-documented in the scientific literature that the sensitivity of viral load-based HIV tests used in infants is low immediately after birth and rises to approximately 80%-90% at 2-6 weeks after birth [Dunn et al., 2000, Lambert et al., 2003, Fiscus et al., 2007, Cachafeiro et al., 2009, Zhang et al., 2008]. As a result, infants infected in utero or during delivery may be misclassified as a false negative when tested at birth or soon after infection. Several authors have accounted for a time-dependent structure by modeling sensitivity as a function of time elapsed since the true time of onset. Balasubramanian and Lagakos [2003] and Brown [2010] both incorporated a joint distribution for all observed tests, dependent upon true infection time and assuming independence of each test. Balasubramanian and Lagakos [2003] assumed stepwise and linear piecewise sensitivity functions, while Brown [2010] assumed an exponential form, with a shift parameter for in utero/intrapartum infections. Alternatively, Gupte et al. [2007] modeled the extra variability of the observation process as a univariate exponential random variable measuring the delay in detection, estimating the rate parameter. They allowed for different curves, depending on mode of transmission, but did not allow for shift parameter to account for nonzero sensitivity for detecting in utero/intrapartum infections at birth.

Secondly, an infant may die prior to HIV infection, and in populations with high infant mortality, death must be considered as a competing risk. There has been some expo-

ration of methods that jointly model infection and death event times in MTCT applications. Hughes and Richardson [2000] estimated a joint distribution for illness and death, where the distribution for death was nonparametric and the distribution for illness, conditional on death, was a continuous parametric density. The choice of a continuous density allowed the method to handle interval censoring of the illness observations. Frydman and Szarek [2010] estimated a fully nonparametric illness-death model, also allowing for interval censoring. Neither approach adjusted for imperfect sensitivity.

In the literature, there has been extensive development and use of a general class of MSMs that account for misclassification called hidden Markov models (HMMs). In medical and biological fields, these models have been applied in studies on progression of liver cancer [Bartolomeo et al., 2011] and HIV [Guihenneuc-Jouyaux et al., 2004], failure times after lung transplants [Jackson and Sharples, 2001], and genetic mapping [Li et al., 2005]. HMMs describe an underlying true continuous process, with observed states subject to error. The error process can vary over time through the inclusion of time-dependent covariates, but this approach does not explicitly account for sensitivity increasing over time. Additionally, constant or piecewise constant hazard rates are often assumed, which can be too restrictive and introduce bias if the assumptions are not met (Jackson et al., 2003, Bureau, Shiboski, and Hughes, 2003). To better fit the data, flexible parametric spline-based intensity functions have been used to parametrize MSM models, but they do not adjust for misclassification. Joly et al. [2002] introduced an illness-death model where the baseline transition intensities were modeled as penalized spline functions, while Titman [2011] discussed a general MSM with B-spline based transition intensities.

In this chapter, we adjust for time-dependent sensitivity in a spline-based illness death model to estimate the baseline intensities for postnatal transmission of HIV and death and include covariates to assess preventive interventions. Similar to Gupte et al. [2007], we adjust for time-varying sensitivity by introducing univariate exponential random variables to model sensitivity, with two exceptions. We do not estimate sensitivity, and we shift the sensitivity

curve to ensure sensitivity greater than 0 at birth. The rest of the chapter proceeds as follows. In Section 2.2, we define notation, develop the full likelihood for maximization, and discuss model selection and assessment. We evaluate the impact of model assumptions with a simulation study in Section 2.3. In Section 2.4, we apply our method to data collected in HPTN 024, a clinical trial of breastfeeding HIV-positive mothers in Africa conducted 2001-2003. Finally, we conclude with a discussion in Section 2.5.

2.2 Methods

2.2.1 Observed data notation

We begin by introducing some notation. For each of $i = 1, \dots, n$ infants, we collect data at pre-scheduled visit times $t_{1,i}, \dots, t_{m_i,i}$, where the number of visits, m_i , can differ across infants. In our example data, while infants were seen regularly, not all of the samples collected were tested. Most infants were seen at 4-8 weeks, and 6, 9, and 12 months. Following protocol, samples collected prior to the 12 month visit were tested only if the 12 month test was positive or unavailable, and the 12 month viral test was conducted only if the 12 month antibody test was positive. Thus, for all infants, the available visit information can be reduced to time of the last negative test and time of the first positive test, which we denote $(t_{l,i}, t_{r,i})$. Additionally, we observe the time of death or censoring for each infant, denoted by the pair (t_i, δ_i) . The t_i is an observation of the random variable $T_i = \min(T_{d,i}, C_i)$, where $T_{d,i}$ is time of death and C_i is time of censoring for infant i , and $\delta_i = I(T_{d,i} < C_i)$, where $I(\cdot)$ is an indicator function equal to 1 if the condition is true. Finally, because our goal is to understand the association between covariates of interest, such as treatment, and the risk of postnatal HIV infection, let \mathbf{z}_i be a vector of observed covariates for infant i .

2.2.2 The illness-death model

In our approach, we assume the underlying disease process can be described as an illness-death model, which is a multistate random process $X(t)$, indexed by time t . In our motivating example, we consider only three states, denoted by $S = \{1, 2, 3\}$, where 1 is HIV-negative, 2 is HIV-positive, and 3 is death (Figure 2.1). For state-space S , the probability of being in state j by time t , given an initial state h at time s and all previous history \mathcal{H}_{s-} is the transition probability, defined as $P_{hj}(s, t) = \Pr(X(t) = j | X(s) = h, \mathcal{H}_{s-})$. The instantaneous risk or hazard of transitioning from state h to state j is the transition intensity

$$\alpha_{hj}(t | \mathcal{H}_{t-}) = \lim_{\Delta t \rightarrow 0} \frac{\Pr(T_{hj} \in (t, t + \Delta t) | T_{hj} > t, \mathcal{H}_{t-})}{\Delta t}. \quad (2.1)$$

We denote the cumulative intensity as $A_{hj}(t) = \int_0^t \alpha_{hj}(s) ds$.




Figure 2.1: The illness-death model for MTCT of HIV. There are three states (HIV-negative and breastfeeding, illness, death) and three possible transitions.

For mathematical convenience, MSMs are often fit under a Markov or semi-Markov assumption. Given state-space S , states $h, j \in S$, time $s \leq t$, a Markov process assumes $P_{hj}(s, t) = \Pr(X(t) = j | X(s) = h)$, so that the probability of a transition depends only on the current state. A semi-Markov process assumes the transition probability depends on the current state h and time of entry into the state t_h , so that $P_{hj}(s, t) = \Pr(X(t) = j | X(s) = h, t_h)$. From Equation 2.1, the transition intensity under the Markov assumption is the limit of $P_{hj}(t, t + \Delta t) = \Pr(X(t + \Delta t) = j | X(t) = h)$ as Δt goes to 0, a function that depends only on t , time since the start of the process. The Markov assumption implies


that the risk of death for an individual who transitions into illness is defined by a piecewise, discontinuous function. Specifically, let $\alpha_{13}(t)$ be the risk of death from the healthy state, and $\alpha_{23}(t)$ be the risk of death from the illness state. The risk of death for an individual who transitions into illness at time t_{12} is

$$\alpha_{.3}(t) = \begin{cases} \alpha_{13}(t) & \text{for } t < t_{12} \\ \alpha_{23}(t) & \text{for } t \geq t_{12}, \end{cases}$$


where we use $\alpha_{.3}(t)$ to indicate the risk of death from any state. In other words, as long as an infant remains in the healthy state, his/her risk of death at time t is $\alpha_{13}(t)$. Then immediately after the infant transitions into the illness state, the risk of death becomes $\alpha_{23}(t)$. Unless $\alpha_{13}(t) = \alpha_{23}(t)$ for all t , there is a jump in mortality risk at the time of transition into illness. See Figure 2.2a for a graphical depiction of the piecewise mortality curve.

In many medical settings, however, we expect the transition between the two intensities to be more subtle. For example, in infectious diseases such as HIV, it is biologically plausible that the rate of death increases as illness worsens. In this case, a semi-Markov duration model is appropriate. Under the semi-Markov model, the transition probability depends on time in the process (t) and time of transition to the current state (t_h), resulting in a transition intensity $\alpha_{hj}(t, t_h)$ that is a function of both times. Often, the dependency on (t, t_h) is simplified to a single time point, and $\alpha_{hj}(t, t_h) = \alpha_{hj}(t - t_h)$. This is known as the duration model, since the intensity depends on length of time in the current state, or the renewal model, since the time scale is reset at each transition. With a clock reset at the time of illness, the risk of death can increase with worsening disease after the time of transition (see Figure 2.2b). We assume that our model follows the semi-Markov assumption. See Andersen [2002], Meira-Machado et al. [2009] for an overall introduction to MSMs.

To represent the underlying stochastic process of the illness-death model, let (T_{12}, T_{13}, T_{23}) be a random vector of possible event times, where each T_{hj} is the time of



(a) Under the Markov assumption, transition intensities are a function of time since the start of the process. If an individual transitions to illness at time $t_{12} = 6$, they transition from healthy to the illness state, and their risk of death jumps to the transition intensity for the illness to death transition.



(b) In the duration formulation of the semi-Markov assumption, after the onset of illness at $t_{12} = 6$, the risk of death depends on time since entering the state. This allows for a gradual increase in mortality after illness. The dotted line is the risk of mortality for an infant who did not transition into illness.

Figure 2.2: Comparison of the Markov and semi-Markov assumptions.

transition from state h to state j . In our application, it is not possible to transition from illness to the healthy state, and death is considered an absorbing event time, as no event can occur after death.

Because an infant can be infected in utero/during delivery, the process can start either in the healthy or illness states. Thus, we define a baseline distribution where $\eta = \Pr(X(0) = 2) = \Pr(T_{12} \leq 0)$ and $1 - \eta = \Pr(X(0) = 1) = \Pr(T_{12} > 0)$. Then, the distribution $f(t_{12}, t_{13}, t_{23})$ is defined as

$$\begin{aligned}
 f(t_{12}, t_{13}, t_{23}) &= \eta \exp \{-A_{23}(t_{23})\} \alpha_{23}(t_{23}) && \text{for } T_{12} \leq 0 \\
 &= (1 - \eta) \exp \{-A_{12}(t_{12}) - A_{13}(t_{12})\} \\
 &\quad \times \alpha_{12}(t_{12}) \exp \{-A_{23}(t_{23} - t_{12})\} \alpha_{23}(t_{23} - t_{12}) && \text{for } 0 < T_{12} < T_{13} \\
 &= (1 - \eta) \exp \{-A_{12}(t_{13}) - A_{13}(t_{13})\} \alpha_{13}(t_{13}) && \text{for } T_{12} > T_{13}.
 \end{aligned} \tag{2.2}$$

We are interested in estimating covariate coefficients for each of the transitions of the distribution in Equation 2.2. To do so, let hj denote the transition from state h to state j and let $\mathbf{Z}^{(hj)}$ be a matrix of covariates we hypothesize are associated with transition hj . We specify a proportional hazards model, so that

$$\alpha_{hj}(t|\mathbf{Z}^{(hj)}) = \alpha_{hj,0}(t) \exp \left(\boldsymbol{\beta}^{(hj)} \mathbf{Z}^{(hj)} \right), \tag{2.3}$$

where $\boldsymbol{\beta}^{(hj)}$ is the vector of coefficients to be estimated and $\alpha_{hj,0}(t)$ is the baseline transition intensity for the transition from state h to state j . Thus, the full set of parameters for estimation in the illness-death model are $\theta = \left\{ \eta, \alpha_{12}(t), \boldsymbol{\beta}^{(12)}, \alpha_{13}(t), \boldsymbol{\beta}^{(13)}, \alpha_{23}(t), \boldsymbol{\beta}^{(23)} \right\}$.

In the next sections, we develop the likelihood for estimation.

2.2.3 Time-varying sensitivity

To incorporate time-varying sensitivity, we define a conditional random variable $W | T_{12} \sim g_W(w)$ to be the length of time after postnatal transmission until detectability. More specifically, the event time $V = T_{12} + W$ is a random variable we call time of detectability, when an infant's HIV infection progresses past a point where it can be detected by assay. The distribution of (V, T_{13}, T_{23}) is a straightforward convolution of the MSM distribution in Equation 2.2. Assuming T_{12} and W are independent, for T_{12} ,

$$f_{(V, T_{13}, T_{23})}(v, t_{13}, t_{23}) = \int_0^v f(t_{12}, t_{13}, t_{23}) g_W(v - t_{12}) dt_{12}. \quad (2.4)$$

Equation 2.4 demonstrates that the distribution for time to detectability can be interpreted as a weighted mean, where all true onset times less than the time of detectability v are assigned a weight defined by the sensitivity curve. Given a particular true onset t_{12} , this weight describes how likely it is that the illness is detectable by time v .

Conditional on the data and true time of illness, the random variable W is interval censored with $w_i \in (t_{l,i} - t_{12,i}, t_{r,i} - t_{12,i})$ for infant i , where $t_{l,i}$ is the time of the last negative test and $t_{r,i}$ is the time of the first positive test. To account for the interval censoring, we integrate over the range of W in the likelihood, and the resulting cumulative distribution $G_W(w)$ acts as a sensitivity function that increases over time. See Figure 2.3 for an example.

2.2.4 Likelihood

To formulate the likelihood of the model, we account for imperfect sensitivity and interval censoring in the measurement of time to illness. Time of death is subject to right censoring, and the uncertainty in the true time of illness means that it may not be known if an observed death is a transition from the healthy or illness state. In Section 2.2.3, we defined $W | T_{12} > 0 \sim g_W(w)$ as the delay in detection of postnatal infections. In addition, we

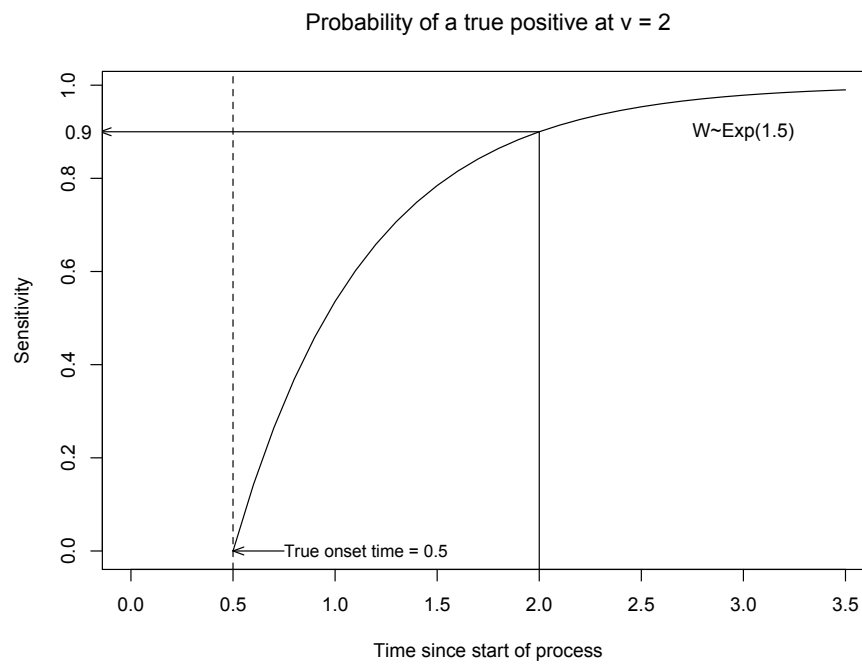


Figure 2.3: Example of a sensitivity function. Given a true onset time $t_{12} = 0.5$, a diagnostic test given at 2 months is detectable if the delay in detection is less than 1.5 months. Assuming the delay is distributed as exponential with rate 1.5 months, this translates to 90% sensitivity.

specify another random variable $U | T_{12} \leq 0 \sim g_U(u)$ to be the length of time after birth until detectability for in utero/during delivery transmissions to allow for a non-zero sensitivity at birth.

Starting from the distribution in Equation 2.2, we derive the likelihood in two cases. First, if infant i tests positive during the study, the range of $T_{12,i}$ is restricted to $T_{12,i} \leq 0$ or $0 < T_{12,i} < t_{r,i}$. Then

$$\begin{aligned}
L_i^{\text{pos}}(\theta; t_{l,i}, t_{r,i}, t_i, \delta_i, t_{r,i} < t_i) &= \eta [G_U(t_{r,i}) - G_U(t_{l,i})] \exp \{-A_{23}(t_i)\} \alpha_{23}(t_i)^{\delta_i} \\
&+ (1 - \eta) \left[\int_0^{t_{l,i}} [G_W(t_{r,i} - t_{12}) - G_W(t_{l,i} - t_{12})] \exp \{-A_{12}(t_{12}) - A_{13}(t_{12})\} \right. \\
&\quad \times \alpha_{12}(t_{12}) \exp \{-A_{23}(t_i - t_{12})\} \alpha_{23}(t_i - t_{12})^{\delta_i} dt_{12} \\
&\quad + \int_{t_{l,i}}^{t_{r,i}} G_W(t_{r,i} - t_{12}) \exp \{-A_{12}(t_{12}) - A_{13}(t_{12})\} \\
&\quad \left. \times \alpha_{12}(t_{12}) \exp \{-A_{23}(t_i - t_{12})\} \alpha_{23}(t_i - t_{12})^{\delta_i} dt_{12} \right] \quad (2.5)
\end{aligned}$$

Second, if infant i does not test positive during the study, the range of $T_{12,i}$ is partitioned into $T_{12,i} \leq 0$, $0 < T_{12,i} < t_{l,i}$, and $T_{12,i} > t_{l,i}$, which correspond to a missed in utero/intrapartum transmission, a missed postnatal transmission, or a true negative at the time of the last death or censoring. Then

$$\begin{aligned}
L_i^{\text{neg}}(\theta; t_{l,i}, t_{r,i}, t_i, \delta_i, t_{r,i} > t_i) &= \eta [1 - G_U(t_{l,i})] \exp \{-A_{23}(t_i)\} \alpha_{23}(t_i)^{\delta_i} \\
&+ (1 - \eta) \left[\int_0^{t_{l,i}} [1 - G_W(t_{l,i} - t_{12})] \exp \{-A_{12}(t_{12}) - A_{13}(t_{12})\} \right. \\
&\quad \times \alpha_{12}(t_{12}) \exp \{-A_{23}(t_i - t_{12})\} \alpha_{23}(t_i - t_{12})^{\delta_i} dt_{12} \\
&\quad + \int_{t_{l,i}}^{t_i} \exp \{-A_{12}(t_{12}) - A_{13}(t_{12})\} \alpha_{12}(t_{12}) \\
&\quad \left. \times \exp \{-A_{23}(t_i - t_{12})\} \alpha_{23}(t_i - t_{12})^{\delta_i} dt_{12} \right] \\
&+ (1 - \eta) \exp \{-A_{12}(t_i) - A_{13}(t_i)\} \alpha_{13}(t_i)^{\delta_i} \tag{2.6}
\end{aligned}$$

From the proportional hazards model of Equation 2.3, $\alpha_{hj}(t|\mathbf{Z}^{(hj)}) = \alpha_{hj,0}(t) \exp(\boldsymbol{\beta}^{(hj)} \mathbf{Z}^{(hj)})$, where $\mathbf{Z}^{(hj)}$ is the matrix of covariates hypothesized to be associated with transition hj , $\boldsymbol{\beta}^{(hj)}$ is the vector of coefficients to be estimated, and $\alpha_{hj,0}(t)$ is the baseline transition intensity for the transition from state h to state j . Then the full set of parameters to be estimated is $\theta = \{\eta, \alpha_{12}(t), \boldsymbol{\beta}^{(12)}, \alpha_{13}(t), \boldsymbol{\beta}^{(13)}, \alpha_{23}(t), \boldsymbol{\beta}^{(23)}\}$. Each column of $\mathbf{Z}^{(hj)}$ is a vector of observed covariate values for infant i associated with the hj transition, and for infant i , the full set of observed covariates is represented as $\mathbf{z}_i = \{\mathbf{z}_i^{(hj)}, (hj) = (12), (13), (23)\}$. The set of included covariates can differ across transitions. In our implementation, we focus on covariates that do not depend on time, though we could extend our approach to include time-dependent covariates.

Let $j = 1, \dots, n_j$ index participants who test positive at some point during the study and $k = 1, \dots, n_k$ index those with all negative tests over the course of the study. For a sample of $n = n_j + n_k$ independent observations, the likelihood is the product $L(\theta) = \prod_{j=1}^{n_j} L_j^{\text{pos}}(\theta; t_{l,j}, t_{r,j}, t_j, \delta_j, t_{r,j} < t_j, \mathbf{z}_j) \times \prod_{k=1}^{n_k} L_k^{\text{neg}}(\theta; t_{l,k}, t_{r,k}, t_k, \delta_k, t_{r,k} > t_k, \mathbf{z}_k)$.

2.2.5 Spline baseline transition intensities

To maximize the likelihood of the illness-death model, we need to specify a parametric form for the baseline transition intensities. For our example data, we assume a constant illness rate, α_{12} , since previous breastfeeding studies have determined a constant illness rate over the first year is a reasonable assumption in postnatal mother-to-child transmission of HIV [Kourtis et al., 2006].

However, we expect the mortality rates to vary over time. For a flexible, yet fully parametric approach, we specify the non-constant mortality intensities as penalized spline functions, which allow the shape of the baseline intensity to be data-driven and impose a smoothness by penalizing the curvature of the function [Commenges, 2002]. Let

$$\alpha_{13,0}(t) = \sum_{k=1}^{\kappa+4} f\left(b_k^{(13)}\right) B_k(t) \quad (2.7)$$

$$\alpha_{23,0}(t; t_{12}) = \sum_{k=1}^{\kappa+4} f\left(b_k^{(23)}\right) B_k(t - t_{12}), \quad (2.8)$$

for transitions from healthy to death and illness to death, respectively, where $\{b_k^{(13)}, b_k^{(23)}, k = 1, \dots, \kappa + 4\}$ are spline coefficients, $\{B_k(t), B_k(t - t_{12}), k = 1, \dots, \kappa + 4\}$ are cubic B-spline bases at time t or $t - t_{12}$, and $\kappa + 4$ is the number of B-spline bases, defined by the number of internal knots (κ) plus order of the cubic function [de Boor, 2001]. The function f is a transformation to constrain the curve to be positive. In practice, we use the same number and placement of knots for both mortality curves. The penalty term is based on total curvature $\int_c^d \left(\alpha''_{ij}(s)\right)^2 ds$, where (c, d) are the boundary knots of the B-spline basis. Let $\boldsymbol{\beta} = \{\boldsymbol{\beta}^{(12)}, \boldsymbol{\beta}^{(13)}, \boldsymbol{\beta}^{(23)}\}$ and $\boldsymbol{\theta} = \{\eta, \alpha_{12}, \mathbf{b}, \boldsymbol{\beta}\}$ be the full set of parameters to be estimated. Each penalty term is subtracted from the log likelihood,

$$\log L_{\text{pen}}(\boldsymbol{\theta}) = \log L(\boldsymbol{\theta}) - \lambda_{13} \int_c^d \left(\alpha''_{13,0}(s)\right)^2 ds - \lambda_{23} \int_c^d \left(\alpha''_{23,0}(s)\right)^2 ds,$$

for fixed smoothing parameters $\lambda_{13}, \lambda_{23}$. Larger smoothing parameters result in smoother

baseline intensities.

Because the integral of a B-spline basis is known exactly, this parameterization avoids numerical integration of the cumulative transition intensity curve, reducing computation time. We use $f(a_k) = a_k^2$ to constrain the curve to be positive. Similar to Joly, Commenges, and Letenneur [1998], we found in practice that this transformation was better than other transformations (such as $f(a_k) = \exp\{a_k\}$) when the curve is near 0 and more computationally stable than methods that directly constrained $a_k > 0$.

In estimation, the smoothing parameters are fixed, but unknown, resulting in a family of estimators $\{\hat{\theta}_\lambda; \lambda \in \Lambda\}$. Minimization of the expected Kullback-Leibler (EKL) risk is a common criterion for selecting an estimator. Commenges et al. [2007] derived an approximation to the EKL for penalized likelihoods using a leave-one-out cross-validation (LCV) approximation, where $LCV_\lambda \approx -n^{-1} [l(\hat{\theta}) - \text{tr}(H_{l_{\text{pen}}}^{-1} H_l)]$, l is the unpenalized log-likelihood, l_{pen} is the penalized log-likelihood, $\hat{\theta}$ are the estimated parameters of l_{pen} (including spline and covariate coefficients), $H_{l_{\text{pen}}}$ is the Hessian of the penalized log-likelihood evaluated at $\hat{\theta}$, and H_l is the Hessian of the unpenalized log-likelihood evaluated at $\hat{\theta}$. While overall minimization of the LCV is possible using standard search algorithms such as a line or grid search, differences in the LCV can also be used to directly compare different parameterizations of the model [Commenges et al., 2007]. In particular, in our application, we propose to use the LCV criterion to compare estimators across fixed assumed sensitivity curves and levels of smoothing.

The log-likelihood can be maximized with quasi-Newton methods, such as BFGS [Press et al., 2007]. For inference, we use the estimated Hessian of the penalized likelihood. This approach was justified by O’Sullivan [1988] from a Bayesian perspective, where the spline coefficients were considered as random variables with a prior distribution defined by the penalty term. In this case, the inverse of the negative Hessian from the penalized likelihood is an approximate variance for the posterior distribution of the spline coefficient estimates.

2.3 Simulations

In this section, we conducted several simulations to assess how robust our covariate coefficient estimates are to misspecification of the sensitivity curve. Our model relies on the assumption that the sensitivity curve for HIV testing in infants is known exactly and is the same for all infants in the study, and violations of this assumption may result in biased estimates. To examine this, we fixed the covariate effects at known values and generated data sets under two possible violations of the sensitivity curve assumption. In the first set of simulations, we assumed that the same sensitivity curve applies to all infants in the study, but that we incorrectly specified in estimation how quickly the curve increases over time. Second, to explore the effect of individual-specific sensitivity curves, we generated data where the sensitivity curves varied across observations and fitted the model with a constant sensitivity curve. We compared the model estimates to true covariate effects to determine bias and coverage.

In the following simulations, we assumed the delay in detectability for in utero/intrapartum transmissions, U , is a shifted exponential with mean μ_U , shift τ , and the delay in detectability for postnatal transmissions, W , is exponential with mean μ_W . The shifted exponential reflects the possibility that some observations have reached detectability by the start of the process. We set $\tau = 0.5$ months. This is similar to the estimated shifted exponential in Balasubramanian and Lagakos [2003] and Brown [2010]. We assumed a constant illness rate, α_{12} , and spline-based mortality intensity curves under the semi-Markov assumption, denoted $\alpha_{13}(t)$ and $\alpha_{23}(t - t_{12})$. Sample size for all simulations was 1500 participants. All data sets were generated using the process outlined in Appendix 2.A.

2.3.1 Incorrect specification of the mean sensitivity curve

We ran two sets of simulations to assess the robustness of model estimates to incorrect specification of the mean sensitivity curve. In the first simulation, we generated data assuming

a mean delay in detectability of $\mu_U = \mu_W = 0.75$ per month, corresponding to an average 3 week delay. In the second, we generated data assuming a mean delay in detectability of $\mu_U = \mu_W = 1.5$ per month, for an average 6 week delay. We set the covariate coefficients $(\beta_{12}, \beta_{13}, \beta_{23}) = (-0.69, 0, -0.22)$ for α_{12} , $\alpha_{13}(t)$, and $\alpha_{23}(t)$, respectively. For simplicity, we denote the scalar $\beta^{(hj)}$ as β_{hj} . In both simulations, we estimated the model assuming sensitivity functions with means 0, 0.75, or 1.5 (per month) and over increasing levels of smoothing. We parameterized the spline-based mortality intensity curves with 10 equally spaced knots. In practice, we found that numerical optimization algorithms led to non-positive definite estimated variance-covariance matrices in approximately 5% of the models fit. The negative variance estimates occurred primarily in the spline coefficients of the healthy-death intensity for a period early in the study period where there were few events and so may have been a result of a lack of information. In these cases, we used the nearest positive definite matrix algorithm to find the variance-covariance matrix [Higham, 1988].

We report summary results for the illness transition intensity α_{12} and log hazard ratio (HR) β_{12} in Table 2.1. Full results are reported in Tables 2.2-2.3. When the model was fit under the correct sensitivity curve, mean estimates were least biased and coverage was good. Both adjusted models resulted in a large improvement in coverage and bias when compared to the unadjusted model. Comparing across adjusted models, estimates for the HRs were affected by misspecification of sensitivity curves. When the true mean delay in detectability is 3 weeks, the 6 week assumed mean delay results in an approximate 30% downward bias and 96% - 97% coverage of the illness log HR estimate (Table 2.2). When the true mean delay in detectability is 6 weeks, the 3 week assumed mean delay resulted in an approximate 30% upward bias and 85% - 86% coverage of the illness log HR estimate (Table 2.3). Thus, specifying a sensitivity curve with a longer delay will indicate a stronger treatment effect than there actually is, but the increased uncertainty of the longer delay leads to wider confidence intervals and reasonable coverage. Alternatively, erring on the side of a shorter delay sensitivity curve biases the estimate toward a null effect. Bias toward the

null when assuming a short delay also occurs when treatment is associated with an increased risk of illness. To verify this, we generated 250 data sets where treatment increased the risk of illness by 50%, but did not affect either mortality curve. In generating the data, we assumed a true sensitivity curve with a mean 6 week delay. When the assumed sensitivity curve for estimation had a 3 week mean delay, estimates were biased down toward a null effect (see Table 2.4).

Misspecification of the sensitivity curve results in biased estimates and poor coverage for the proportion starting in the HIV-positive state, η . When the data were generated assuming a 3 week mean delay, and the model was fit assuming a 6 week delay, the estimate for η was biased upward by 13%, with 46% coverage. When the data were generated assuming a 6 week mean delay, and the model was fit assuming a 3 week delay, the estimate for η was biased downward by 20%, with only 15% coverage. If this parameter is of interest, another approach may result in a better estimate. For example, as discussed in Chapter 1, the approach often used is to assume all postnatal positive tests prior to 4 or 6 weeks were truly infected in utero or intrapartum. While some early postnatal infections will be misclassified, this approach will result in a reasonable approximation for proportion of infants infected in utero or during delivery.

Table 2.1: A summary of simulation results for the mean, standard error, and coverage of the illness transition intensity parameters. The mortality intensities are not smoothed. The far left column indicates the data were generated assuming mean delays of 3 and 6 weeks. The next column lists the assumed mean delays (0, 3, and 6 weeks) used in estimation.

Mean delay for generation (wks)	Mean delay for estimation (wks)	$\alpha_{12} = 0.0070$		$\beta_{12} = -0.69$	
		Mean	(SE, Cov)	Mean	(SE, Cov)
3	0	0.0163	(0.0016, 0)	-0.23	(0.14, 11)
	3	0.0069	(0.0011, 97)	-0.70	(0.29, 94)
	6	0.0056	(0.0011, 86)	-0.90	(0.37, 97)
6	0	0.0174	(0.0017, 0)	-0.18	(0.14, 5)
	3	0.0091	(0.0014, 61)	-0.47	(0.24, 85)
	6	0.0069	(0.0013, 96)	-0.70	(0.36, 95)

2.3.2 *Person-specific sensitivity curves*

In this section, we examine the impact of misspecification of the variability of sensitivity curves. It has been documented in the literature that the performance of HIV DNA assays may depend on population characteristics, such as particular subtype of circulating virus [Zhang et al., 2008]. In addition to the population level characteristics, there may also be differences due to person-level characteristics. For example, there may be some individual-specific factor, such as a slower replicating virus, that prevents detection of the virus for a longer period of time than would be expected. Ignoring this possible variability of sensitivity curves may result in bias or incorrect standard errors of the covariate coefficient estimates.

Through a simulation study, we examined how robust our model is to variations of the sensitivity across individuals. To do so, we generated an random sensitivity curve for each participant in the study, ensuring that the mean of the random sensitivity curves is the same as the assumed sensitivity curve used in estimation, as detailed in Appendix 2.B. For each of five increasing levels of variability in the sensitivity curves, we generated 250 random data sets, with the curves centered at mean delay $\mu_U = \mu_W = 0.75$ or 1.5 (per month). We estimated the parameters assuming the appropriate sensitivity curve. We parameterized the mortality intensities with 10 equally spaced knots and estimated with two levels of smoothing parameters: $(0,0)$ and $4.25 \times (10^5, 10^4)$. We found that when the model was fit under the correct mean sensitivity curve, there was no difference in bias or coverage across increasing variability of the sensitivity curves. We report the full results in Tables 2.5 - 2.6.

2.4 *Application*

We illustrate the model above using a double-blind, placebo-controlled, multi-site randomized trial conducted through the HIV Prevention Trials Network (HPTN) to determine the efficacy of antibiotics in reducing perinatal MTCT of HIV [Taha et al., 2006]. The study enrolled women at 20-24 weeks gestation in Lilongwe and Blantyre (Malawi), Lusaka

Table 2.2: Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (-0.69, 0, -0.22)$ and an average 3 week delay (in **bold**). Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t - t_{12}))$, respectively. The semi-Markov model is assumed.

		$\eta = 0.15$			$\alpha_{12} = 0.007$		
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	0.055	(0.005, 0)	0.0163	(0.0016, 0)	0.0163	(0.0016, 0)
	2.75	0.055	(0.005, 0)	0.0164	(0.0016, 0)	0.0164	(0.0016, 0)
	6.5	0.055	(0.005, 0)	0.0164	(0.0016, 0)	0.0164	(0.0016, 0)
3 weeks	0	0.150	(0.010, 97)	0.0069	(0.0011, 97)	0.0069	(0.0011, 97)
	2.75	0.150	(0.010, 97)	0.0069	(0.0011, 97)	0.0069	(0.0011, 97)
	6.5	0.150	(0.010, 97)	0.0069	(0.0011, 97)	0.0069	(0.0011, 97)
6 weeks	0	0.170	(0.011, 46)	0.0056	(0.0011, 86)	0.0056	(0.0011, 86)
	2.75	0.171	(0.011, 44)	0.0057	(0.0011, 88)	0.0057	(0.0011, 88)
	6.5	0.171	(0.011, 44)	0.0057	(0.0011, 89)	0.0057	(0.0011, 89)

		$\beta_{12} = -0.69$		$\beta_{13} = 0$		$\beta_{23} = -0.22$	
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	-0.23	(0.14, 11)	-0.03	(0.29, 94)	-0.23	(0.18, 96)
	2.75	-0.23	(0.14, 11)	0.07	(0.32, 91)	-0.24	(0.18, 96)
	6.5	-0.23	(0.14, 11)	0.07	(0.34, 91)	-0.24	(0.18, 96)
3 weeks	0	-0.70	(0.29, 94)	-0.03	(0.29, 94)	-0.24	(0.18, 96)
	2.75	-0.71	(0.29, 94)	0.08	(0.32, 90)	-0.24	(0.18, 96)
	6.5	-0.71	(0.29, 95)	0.08	(0.33, 90)	-0.24	(0.18, 96)
6 weeks	0	-0.90	(0.37, 97)	0.03	(0.35, 93)	-0.22	(0.18, 96)
	2.75	-0.92	(0.37, 96)	0.18	(0.41, 90)	-0.23	(0.18, 96)
	6.5	-0.92	(0.37, 96)	0.19	(0.44, 91)	-0.23	(0.18, 96)

Table 2.3: Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (-0.69, 0, -0.22)$ and an average 6 week delay (in **bold**) Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.

		$\eta = 0.15$		$\alpha_{12} = 0.007$	
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	0.032	(0.004, 0)	0.0174	(0.0017, 0)
	2.75	0.032	(0.004, 0)	0.0175	(0.0017, 0)
	6.5	0.032	(0.004, 0)	0.0175	(0.0018, 0)
3 weeks	0	0.120	(0.010, 14)	0.0091	(0.0014, 61)
	2.75	0.120	(0.010, 15)	0.0092	(0.0014, 59)
	6.5	0.120	(0.010, 15)	0.0092	(0.0014, 59)
6 weeks	0	0.150	(0.011, 92)	0.0069	(0.0013, 96)
	2.75	0.150	(0.011, 92)	0.0070	(0.0013, 96)
	6.5	0.150	(0.011, 91)	0.0070	(0.0013, 96)

		$\beta_{12} = -0.69$		$\beta_{13} = 0$		$\beta_{23} = -0.22$	
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	-0.18	(0.14, 5)	-0.03	(0.24, 97)	-0.23	(0.18, 98)
	2.75	-0.19	(0.14, 5)	0.08	(0.26, 94)	-0.24	(0.18, 98)
	6.5	-0.19	(0.14, 5)	0.09	(0.28, 94)	-0.24	(0.18, 98)
3 weeks	0	-0.47	(0.24, 85)	-0.02	(0.24, 98)	-0.24	(0.18, 98)
	2.75	-0.48	(0.24, 86)	0.08	(0.25, 94)	-0.25	(0.18, 98)
	6.5	-0.48	(0.24, 85)	0.08	(0.26, 94)	-0.25	(0.18, 98)
6 weeks	0	-0.70	(0.36, 95)	0.03	(0.28, 97)	-0.24	(0.18, 98)
	2.75	-0.72	(0.36, 95)	0.16	(0.32, 94)	-0.25	(0.18, 97)
	6.5	-0.72	(0.36, 95)	0.17	(0.35, 94)	-0.25	(0.18, 97)

Table 2.4: Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated assuming $(\beta_{12}, \beta_{13}, \beta_{23}) = (0.4, 0, 0)$ and an average 6 week delay (in **bold**) Estimated models assume average 0, 3, and 6 week delays, and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.

		$\eta = 0.15$		$\alpha_{12} = 0.007$	
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	0.031	(0.004, 0)	0.0174	(0.0016, 0)
	2.75	0.031	(0.004, 0)	0.0175	(0.0016, 0)
	6.5	0.031	(0.004, 0)	0.0175	(0.0016, 0)
3 weeks	0	0.115	(0.009, 6)	0.0092	(0.0013, 49)
	2.75	0.115	(0.009, 6)	0.0093	(0.0013, 48)
	6.5	0.115	(0.009, 6)	0.0093	(0.0013, 48)
6 weeks	0	0.149	(0.011, 97)	0.0070	(0.0013, 97)
	2.75	0.149	(0.011, 97)	0.0071	(0.0013, 97)
	6.5	0.149	(0.011, 97)	0.0071	(0.0013, 97)

		$\beta_{12} = 0.4$		$\beta_{13} = 0$		$\beta_{23} = 0$	
Mean assumed delay	λ_c	Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0 weeks	0	0.17	(0.14, 53)	-0.01	(0.27, 95)	-0.02	(0.17, 96)
	2.75	0.17	(0.13, 52)	0.09	(0.28, 93)	-0.02	(0.17, 96)
	6.5	0.17	(0.13, 52)	0.09	(0.28, 93)	-0.02	(0.17, 96)
3 weeks	0	0.34	(0.20, 92)	-0.01	(0.26, 95)	-0.01	(0.17, 96)
	2.75	0.33	(0.20, 92)	0.08	(0.27, 92)	-0.02	(0.17, 97)
	6.5	0.33	(0.20, 92)	0.08	(0.28, 92)	-0.02	(0.17, 97)
6 weeks	0	0.43	(0.26, 95)	-0.03	(0.30, 95)	-0.01	(0.17, 96)
	2.75	0.42	(0.26, 95)	0.09	(0.34, 92)	-0.03	(0.17, 96)
	6.5	0.42	(0.26, 95)	0.08	(0.35, 93)	-0.03	(0.17, 96)

Table 2.5: Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated individual sensitivity curves assuming an mean 3 week delay sensitivity curve. The case where all observations are generated assuming the same sensitivity curve is denoted “none”, and the cases A-D denote individual sensitivity curves with increasing variability. Estimated models assume a single sensitivity curve with a mean 3 week delay and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.

λ_c		$\eta = 0.15$		$\alpha_{12} = 0.007$	
		Mean	(SE, Cov)	Mean	(SE, Cov)
0	none	0.150	(0.010, 96)	0.0070	(0.0012, 97)
	A	0.150	(0.010, 96)	0.0071	(0.0012, 96)
	B	0.149	(0.010, 94)	0.0069	(0.0011, 97)
	C	0.149	(0.010, 94)	0.0070	(0.0012, 97)
	D	0.149	(0.010, 94)	0.0070	(0.0012, 94)
4.25	none	0.149	(0.010, 94)	0.0071	(0.0012, 94)
	A	0.150	(0.010, 95)	0.0070	(0.0012, 94)
	B	0.150	(0.010, 95)	0.0070	(0.0012, 95)
	C	0.149	(0.010, 95)	0.0071	(0.0012, 92)
	D	0.149	(0.010, 94)	0.0072	(0.0012, 93)

λ_c		$\beta_{12} = -0.69$		$\beta_{13} = 0$		$\beta_{23} = -0.22$	
		Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0	none	-0.70	(0.30, 96)	0.05	(0.27, 96)	-0.22	(0.18, 95)
	A	-0.71	(0.30, 96)	0.17	(0.30, 90)	-0.23	(0.18, 94)
	B	-0.67	(0.32, 91)	-0.01	(0.28, 94)	-0.22	(0.20, 93)
	C	-0.68	(0.32, 92)	0.11	(0.31, 91)	-0.23	(0.20, 93)
	D	-0.69	(0.30, 94)	-0.01	(0.26, 96)	-0.23	(0.19, 95)
4.25	none	-0.69	(0.30, 94)	0.10	(0.28, 93)	-0.23	(0.19, 95)
	A	-0.70	(0.29, 96)	0.00	(0.27, 96)	-0.22	(0.20, 94)
	B	-0.71	(0.29, 96)	0.11	(0.29, 92)	-0.23	(0.20, 93)
	C	-0.72	(0.30, 95)	0.02	(0.27, 93)	-0.24	(0.19, 95)
	D	-0.72	(0.30, 95)	0.13	(0.30, 91)	-0.25	(0.19, 94)

Table 2.6: Mean estimates (standard error, 95% CI coverage) from 250 simulated data sets, generated individual sensitivity curves assuming an mean 6 week delay sensitivity curve. The case where all observations are generated assuming the same sensitivity curve is denoted “none”, and the cases A-D denote individual sensitivity curves with increasing variability. Estimated models assume a single sensitivity curve with a mean 6 week delay and smoothing parameters $\lambda_c \times (10^5, 10^4)$ for $(\alpha_{13}(t), \alpha_{23}(t))$, respectively. The semi-Markov model is assumed.

λ_c		$\eta = 0.15$		$\alpha_{12} = 0.007$	
		Mean	(SE, Cov)	Mean	(SE, Cov)
0	none	0.150	(0.011, 94)	0.0070	(0.0013, 96)
	A	0.151	(0.011, 95)	0.0071	(0.0013, 95)
	B	0.150	(0.011, 95)	0.0070	(0.0013, 95)
	C	0.151	(0.011, 95)	0.0072	(0.0014, 94)
	D	0.148	(0.010, 94)	0.0070	(0.0012, 97)
4.25	none	0.149	(0.010, 96)	0.0071	(0.0012, 97)
	A	0.148	(0.010, 94)	0.0069	(0.0013, 95)
	B	0.148	(0.010, 94)	0.0070	(0.0013, 95)
	C	0.147	(0.010, 93)	0.0070	(0.0014, 95)
	D	0.147	(0.010, 93)	0.0071	(0.0014, 95)

λ_c		$\beta_{12} = -0.69$		$\beta_{13} = 0$		$\beta_{23} = -0.22$	
		Mean	(SE, Cov)	Mean	(SE, Cov)	Mean	(SE, Cov)
0	none	-0.72	(0.37, 95)	-0.02	(0.28, 96)	-0.21	(0.19, 95)
	A	-0.74	(0.38, 95)	0.11	(0.32, 94)	-0.23	(0.19, 95)
	B	-0.71	(0.35, 95)	-0.01	(0.29, 95)	-0.23	(0.19, 95)
	C	-0.73	(0.35, 96)	0.14	(0.35, 91)	-0.24	(0.19, 95)
	D	-0.68	(0.33, 95)	-0.03	(0.31, 93)	-0.22	(0.21, 93)
4.25	none	-0.70	(0.34, 96)	0.10	(0.34, 93)	-0.24	(0.21, 94)
	A	-0.72	(0.35, 96)	-0.02	(0.32, 93)	-0.22	(0.20, 94)
	B	-0.73	(0.35, 97)	0.12	(0.37, 92)	-0.23	(0.20, 94)
	C	-0.72	(0.39, 95)	-0.03	(0.26, 96)	-0.25	(0.21, 93)
	D	-0.74	(0.39, 96)	0.10	(0.29, 97)	-0.26	(0.21, 92)

(Zambia), and Dar es Salaam (Tanzania) during 2001-2003. A short course of antibiotics or placebo was administered to the mother at enrollment and during labor. Additionally, mothers and infants were also given prophylactic nevirapine at birth. Our data set included 1499 live-born infants, with HIV-1 RNA PCR assays conducted on blood samples collected at 4-8 weeks, 3, 6, 9, and 12 months of age. Infants were followed until death or their last visit, and event times were administratively censored at 13 months. Across all included sites, approximately 20% of infants had been weaned by 6 months. By 12 months, weaning rates were around 70% at both sites in Malawi and 57% in Lusaka. We excluded the site at Dar es Saalam, which counseled early weaning, resulting in a weaning rate of 95% by 6 months. We did not consider weaning as a censoring event, and so infants who weaned during the study remained in the denominator for illness. A total of 190 infants (13%) died during the study. Of the 190 infants, 135 (71%) tested positive for HIV prior to death. Overall, 322 infants (21%) tested positive for HIV by the end of the study period. The trial was stopped early when the treatment failed to show significant benefit in reducing HIV transmission. There was no significant difference across treatment group in either overall infection rates or disease-free survival.

To implement the model, we first determined an appropriate sensitivity curve from the scientific literature. A recent review of diagnostic testing of infants approximates sensitivity of RNA PCR tests as 25-50% at birth, 55-66% at 1-3 weeks of age, and near 100% at one month. DNA PCR tests were 10-25% at birth, 75% at 2 weeks, and 95% at 1 month of age [Wessman et al., 2012]. Assuming a shifting parameter of 0.5 months, we found while no one rate parameter represented the estimates from the literature well, a mean parameter of 0.75 months reflected the steeper rise usually seen just after birth/onset and a mean parameter of 1.5 months represented a low sensitivity at birth. Thus, we specified the following sensitivity curves for estimation: $U \sim \text{Shifted Exp}(\mu_u = 0.75, \delta = 0.5)$, $W \sim \text{Exp}(\mu_w)$, with $\mu_w = 0.75, 1.5$, and additionally fit the model that assumes perfect sensitivity.

Next, we defined the transition intensities assuming a semi-Markov model with identical

bases for $\alpha_{13}(t)$ and $\alpha_{23}(t - t_{12})$. We used $\kappa = 10$ equally spaced knots, with smoothing parameters ranging from 5000 to 100,000 for $\alpha_{13}(t)$ and 500 to 10,000 for $\alpha_{23}(t)$. We assumed a constant baseline intensity for illness. We parameterized the treatment effect, with $\beta = \{\beta_{12}, \beta_{13}, \beta_{23}\}$ as the intensity ratios for the transition from healthy to illness, healthy to death, and illness to death, respectively, under the proportional hazards assumption. While estimating the β is our primary question of interest, we also examine how model choice impacts the baseline intensities of the MSM.

The results are reported in Figures 2.5-2.7. Approximately 17.5% of infants were estimated to experience in utero/intrapartum infections (95% confidence interval: 15.5%, 19.6%). Cumulative incidence of postnatal infections was approximately 7% by 12 months. Choice of sensitivity function influenced the estimated event rates for transitions from healthy to illness and healthy to death. Illness rates are higher and uninfected mortality rates are lower when comparing a 6 week mean delay to a 3 week mean delay for the assumed sensitivity curve. This is because assuming a longer delay results in a longer range of probable true onset times. To demonstrate this, we plotted the range of probable true onset times for an individual who tests negative at 12 months and dies at 12.5 months (Figure 2.4). For a specific example, the probability that an infant is infected at 8 months and tests negative at 12 months is 0, 0.005, and 0.07, for respective assumed mean delays of 0 weeks, 3 weeks, and 6 weeks (corresponding to perfect sensitivity, and delay rates of 1.33 and 0.67). When applied to our example data, this translates to approximately 0%, 6% and 40% of infants estimated as undetected infections at the end of the study, when assuming the mean delays of 0 weeks, 3 weeks, and 6 weeks sensitivity curves, respectively, leading to a higher estimated illness rate in the 6 week delay sensitivity curve (demonstrated in Figure 2.5b).

Choice of sensitivity function does not affect the illness mortality rate. Although, as expected, adjusting for sensitivity does increase the duration in the illness state prior to death. In uninfected infants, mortality rates were highest between 2 and 5 months of

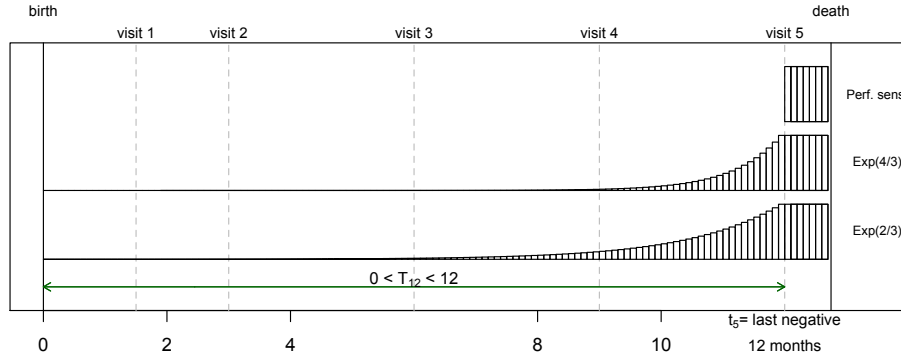


Figure 2.4: Conditional on time of illness at t_{12} , this diagram depicts the probability that the delay is longer than $12 - t_{12}$ months. When the assumed mean delay of detection is long, the range of probable true illness times is extended earlier.

age, with a second slight increase between 8 and 10 months that may be associated with early weaning. In infected infants, mortality peaked at about 4 months after infection, followed by a steady decline. It has been shown in the literature that infants with early infection are more likely to develop a rapidly progressive disease course than those with a late infection [Newell et al., 2004]. Because we did not stratify by timing of infection (in utero/intrapartum versus postnatal), this result may be driven by the infants with early transmission.

The log proportional hazards (PH) coefficient estimates for the treatment covariate were relatively stable across sensitivity curves and smoothing parameters (Figure 2.5a.) After adjustment for sensitivity, the PH coefficient for illness (β_{12}) had a slightly larger point estimate (0.1 in the unadjusted to 0.2 in the adjusted models), but much wider confidence intervals (Figure 2.5a, top.) The mortality curve estimates, β_{13} and β_{23} , do increase slightly as the baseline curves become smoother, but this increase is minor compared to the confidence intervals of the estimates (Figure 2.5a, middle and bottom.) The confidence intervals for β_{13} are large compared to the other two PH estimates (Figure 2.5a, middle.) As the

assumed delay in detection increases, early observed positive tests are weighted as more likely to be true postnatal infections. This is demonstrated by the higher point estimates of the rate of postnatal infection under the assumed 6-week delay, when compared to the shorter assumed 3-week delay (Figure 2.5b, top). The LCV scores are plotted in Figure 2.5b (bottom), where level of smoothing increases from left to right. These scores demonstrate that the least smoothed models fit as well, or better, than more smoothed models. However, as seen in Figures 2.6 and 2.7, the more smoothed models better reflect an overall trend of the mortality risks than the least smoothed models. In this case, if interpretability of the spline-based transition intensities is important, we can use the LCV to select the best fitting model among a subset of best descriptive models.

The model showed a slight insignificant trend indicating that the antibiotics had a harmful effect on both HIV infection and risk of mortality in HIV-uninfected patients, with point estimates of log proportional hazard coefficient of 0.2 and 0.5, respectively (Figure 2.5a.) One possible explanation is that treatment may have prevented some pre-term deaths, leading to baseline differences in the overall health of the infants across the two groups. In other words, the infants that had benefitted from an in utero treatment could be more susceptible to harmful events after birth, when the treatment was no longer given.

To examine how well the estimated model fits the observed data, we simulated 100 data sets using the estimated parameters from the model that assumed a 3-week delay sensitivity curve with smoothing parameters (42500, 4250) for $\alpha_{13}(t)$ and $\alpha_{23}(t)$. We simulated both illness and death times following the algorithm discussed in Appendix 2.A. We report the resulting non-parameteric observed cumulative incidence functions (CIF) for illness and death in Figures 2.8 and 2.9 (see Andersen, Abildstrom, and Rosthøj [2002] for more on CIFs). The observed CIFs fell within the variability of the randomly generated CIFs for both illness and death, and across both randomization groups, indicating our model provides a reasonable approximation to the observed outcomes of the underlying disease process. There is a clear difference across the CIFs for death (though with some overlap), which is reflective

of the confidence interval for β_{13} being nearly significant at the 95% confidence level.




Figure 2.5: Coefficient estimates and 95% confidence intervals for $\beta_{12}, \beta_{13}, \beta_{23}$, across assumed lag time distributions. Within each lag time, level of smoothing increases from left to right. CV scores marked with “+” correspond to the plotted estimates.

2.4.1 The Markov assumption

We compared our results, based on a semi-Markov assumption, to the model fit under the more common Markov assumption. Under the Markov assumption, the probability of the transitioning into a new state depends only on the current state, and no other patient history, and the rate of transition is a function of time since the start of the process (see Section 2.2.2). Overall, estimation under the Markov assumption was more numerically unstable in maximization, leading to wider confidence intervals and more instances of non-positive definite variance matrices (Figure 2.10). In particular, when a perfect sensitivity

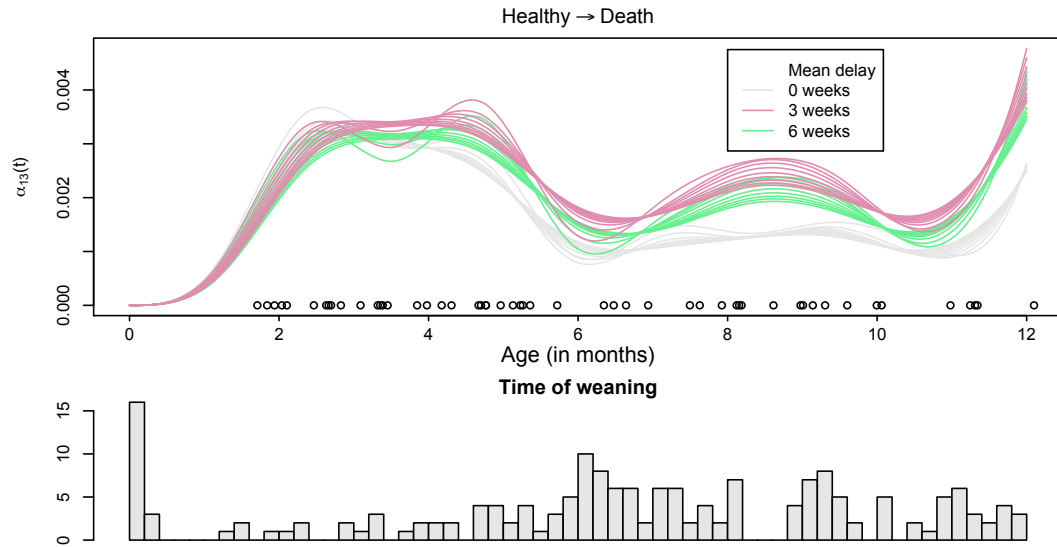


Figure 2.6: Estimated smoothed mortality curves for the healthy to death transition, assuming 0, 3, and 6 week delay in detection, and estimated over several degrees of smoothing. Observed event times are indicated as points. The histogram represents the distribution of weaning times of those who do not test positive during the study period. The slight increase in mortality at 8 months could be a delayed result of infants weaned at 6 months.

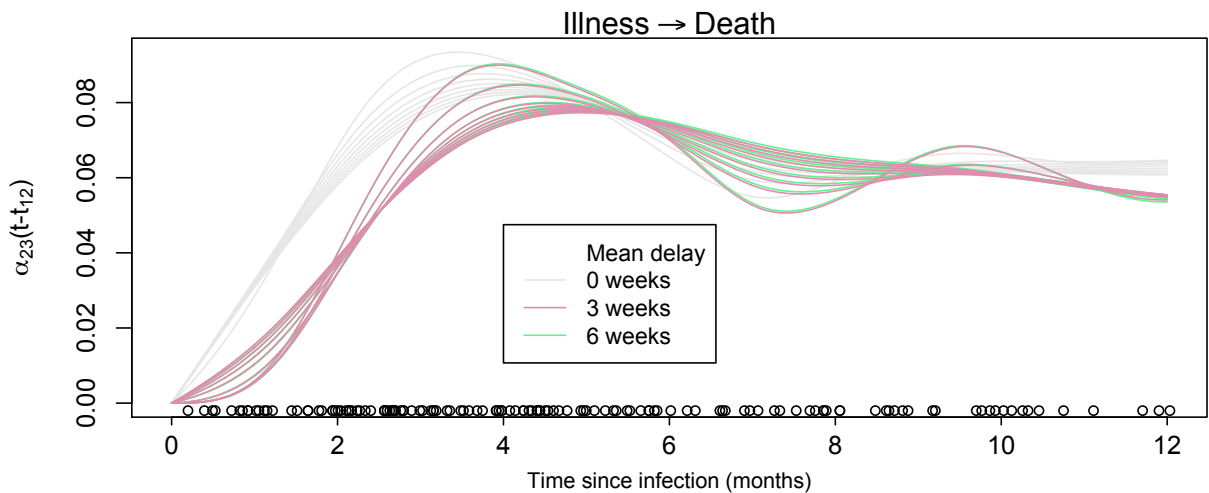


Figure 2.7: Estimated smoothed mortality curves for the illness to death duration transition, assuming 0, 3, and 6 week delay in detection, and estimated over several degrees of smoothing. Observed event times (time of death - time of first positive test) are indicated as points.

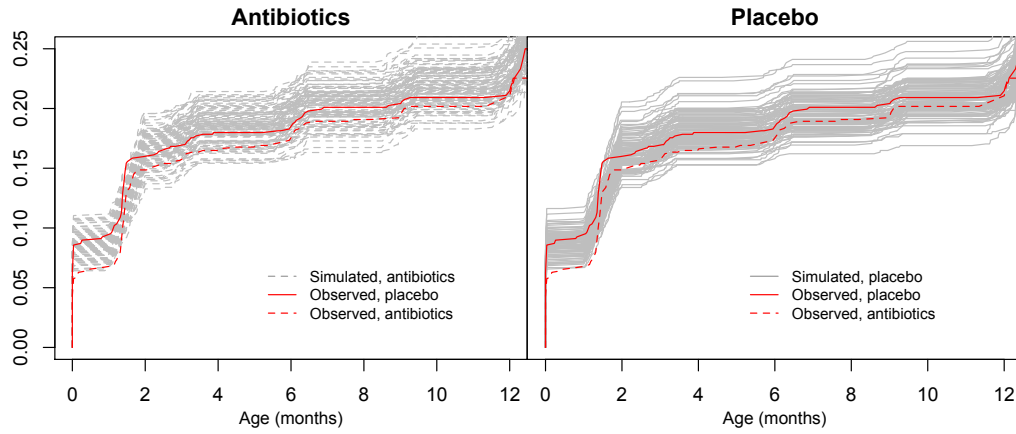


Figure 2.8: Cumulative incidence of illness from 100 randomly generated data sets, with parameters estimated assuming a 3-week delay sensitivity curve. These plots are stratified by treatment group, with antibiotics on the left and placebo on the right. The true CIFs for both antibiotics and placebo are in red on both plots.

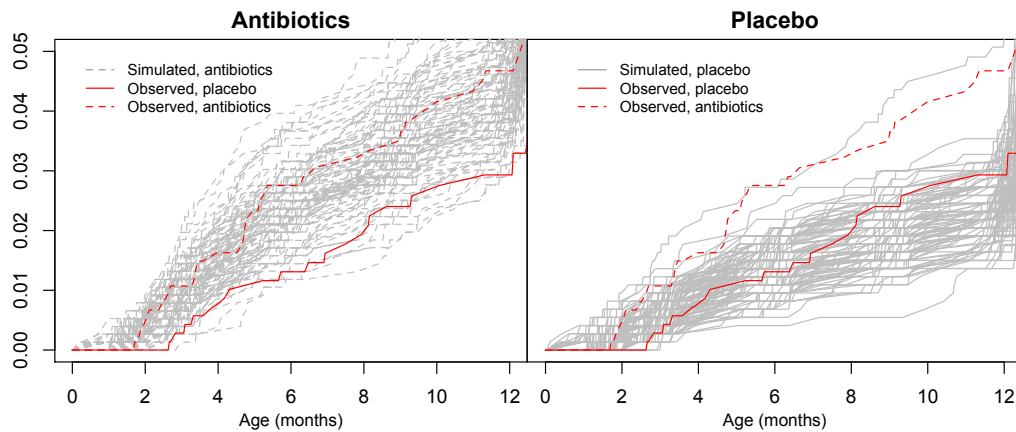


Figure 2.9: Cumulative incidence of healthy-to-death transition from 100 randomly generated data sets, with parameters estimated assuming a 3-week delay sensitivity curve. These plots are stratified by treatment group, with antibiotics on the left and placebo on the right. The true CIFs for both antibiotics and placebo are in red on both plots.

or a 6 week mean delay was assumed, increasing the level of smoothing resulted in intensities that were smoothed down to zero (Figure 2.11), which lead to the wide confidence intervals for β_{13} . Because the model under the Markov assumption is sensitive to both over adjustment (assuming long delay times) and under adjustment (assuming no delay) for misclassification, the effect of smoothing toward zero is likely due to some combination of the jump in mortality curves (Figure 2.2a) and the lack of events at the beginning of the study. While determining the exact mechanism for why the smoothing down to 0 occurs in these two circumstances requires more exploration, the result suggests that a Markov assumption is not appropriate for this analysis.

2.5 Discussion

In this chapter, we introduced an illness-death model that accounts for time-dependent sensitivity. By parameterizing all baseline intensities as penalized cubic B-spline functions, our model can easily incorporate methods for interval censoring, while maintaining flexibility in the shape of the intensity curves. This model addresses several shortcomings in the illness-death model, and MSMs in general, that make it more suitable for application to infectious disease problems. In the multistate model literature, hidden Markov and semi-Markov models are currently used to incorporate misclassification into MSMs, but the adjustment for error is not necessarily increasing over time, and in practice, baseline intensities are often restricted to constant or piecewise constant. In this aspect, our model introduces a continuous time approach to adjustment for error in MSMs that incorporates an increasing sensitivity curve, while allowing for flexibility in the baseline intensities.

Our model can be generalized to many disease settings where sensitivity is thought to be a function of disease severity. It may be possible to use new technologies of diagnostic testing to establish disease status and estimate the sensitivity curve of an older diagnostic test that




Figure 2.10: Estimated transition intensity coefficients, under the Markov assumption. The assumed delay in detectability follows an exponential distribution with mean parameters 0, 3, and 6 weeks, and the smoothing parameter increases from left to right, starting at no smoothing in the far left.

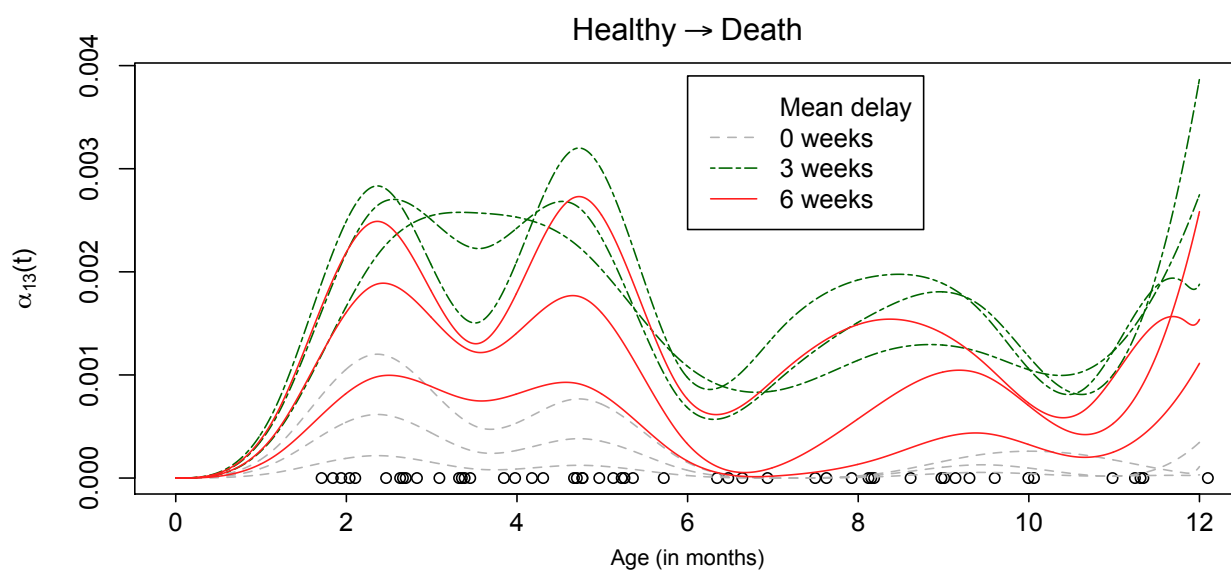


Figure 2.11: Example transition intensities from healthy to death, under the Markov assumption, over three levels of smoothing. The points are observed deaths for infants who did not test positive prior to death. The assumed delay in detectability follows an exponential distribution with mean parameters 0, 3, and 6 weeks. When a 0 or 3 week delay is assumed, the intensities are smoothed down to zero.

may be well-established or easier to administer, and thus, more readily collected in a wider range of studies. For example, recent advancements in the diagnosis of glaucoma have been used to estimate sensitivity curves of the older technology [Leite et al., 2010]. In modeling the progression of cystic fibrosis, there has been growing interest in testing procedures that are more sensitive than the commonly used forced expiratory volume per second, such as a test discussed in Pittman et al. [2012]. This particular proposed newer diagnostic test is not yet widespread, but could be used to establish an approximate sensitivity curve.

In opting to maintain flexible baseline intensities, we have limited our ability to estimate the sensitivity curve. Instead, we specify the curve from the literature, where we could find a reasonable range for plausible curves in studies that have documented sensitivity curves in non-breastfeeding populations. Alternatively, one could argue that results in the literature also demonstrate that differences in sensitivity curves depend of specific characteristics of the population, test used, or some unknown other factor. In this case, we could easily misspecify the rate the curve increases or the shape the curve takes over time for our particular sample. We note, however, that we fit the model under two sensitivity curves, with similar results in the coefficient estimates. Several papers discuss identifiability issues when both event rates and misclassification rates are being estimated [Balasubramanian and Lagakos, 2001, Bureau, Shiboski, and Hughes, 2003].

In our application, we chose to model the baseline transition intensities under the semi-Markov assumption based on an understanding of the biological process. However, a Markov assumption is commonly made in the literature, though we found our model harder to fit under this assumption. While the difficulty in estimation was evident in our particular application, indicating that the semi-Markov model was a better fit in this case, this difficulty may not occur in other applications. Commenges et al. [2007] present a more formal approach to comparing Markov and semi-Markov models, which may be helpful if it is unclear which assumption is more appropriate.

Our model could be extended in several ways. If justified by evidence, we could specify

a different sensitivity curve with 1) a more flexible distribution, such as the Weibull, or 2) a dependency on categorical covariates, such as treatment. Additionally, we could relax the assumption of perfect specificity. Finally, we could extend the underlying illness-death model to include more states. For example, we hypothesize from Figure 2.6 that the mortality rate of uninfected infants increases soon after weaning. By adding in a state for weaning, we could specify a duration model with a flexible baseline intensity to see if there is a spike in the weaning mortality curve soon after entering the state.

2.6 Appendices

2.A Generating the simulated data

To generate the data sets, for each observation, we simulate all possible event times as latent events, the delay time from infection to detectability, and a schedule of visit times. When the hazard curve is spline-based, we solve for the onset time from a randomly generated probability of survival.

1. Simulate visit schedule. In Section 2.3, to generate data for simulations, we drew from $U(v - \epsilon, v + \epsilon)$ for each pre-determined visit time v , where ϵ is a reasonable span around visit time. We chose a visit schedule of 1.5, 3, 6, 9, and 12 months, and $\epsilon = 0.5$ to allow for a flexibility of two weeks on either side of the planned visit time. Alternatively, in Section 2.4, to better mimic the observed visit schedule of the application, we generated an empirical frequency distribution of visit times and randomly drew visit times according to the observed frequency.
2. Simulate the entering state is illness with probability η . When all observations are assumed to enter the process as healthy, $\eta = 0$.
3. For all observations that enter the process as disease negative, independently generate time of illness t_{12} and time of death from healthy t_{13} using transition intensities $\alpha_{12}(t)$

and $\alpha_{13}(t)$, respectively. When the transition intensities are spline based, times are randomly generated through the survival function. For example,

- a. Set the range of observable times $(0, t_\kappa)$ where t_κ is the upper boundary knot (or maximum observable time) of the spline-based hazard.
 - b. Determine the probability of survival at t_κ : $p = S_{13}(t_\kappa)$
 - c. Generate $u \sim U(0, 1)$.
 - d. If $u > p$, set $t_{13} = A_{13}^{-1}[-\log u]$ for time of death t_{13} . If $u < p$, censor at t_m , which is the last observed visit time for each participant.
4. If $t_{12} < t_{13}$ and $t_{12} < t_m$, generate the time of death while in the illness state. Following steps **a - c** above, set $t_{23} = A_{23}^{-1}[A_{23}(t_{12}) - \log u]$ if a Markov model is assumed, and set $t_{23} = A_{23}^{-1}(-\log u) + t_{12}$ if a semi-Markov model is assumed. In the latter case, if $t_{23} > t_m$, censor t_{23} at t_m .
 5. Generate the delay in detectability w for $t_{12} > 0$ from $\text{Exp}(\mu_w)$, and set the time illness is first able to be detected $v = t_{12} + w$. For an observation that enters as disease positive, sample a delay in detectability u from $\text{Shifted Exp}(\mu_u, \tau)$, with shift parameter τ .
 6. Determine the observed values, time of last negative test, time of first positive test, and time of death or time censored. If there are no positive tests observed, set the time of the first positive test as $t_{k+1} = \infty$.

2.B Generating the individual-specific sensitivity curves

We used the following algorithm to generate random sensitivity curves with a specific mean curve.

1. Generate 25,000 sensitivity curves parameterized by exponential rates drawn from a gamma distribution with a mean μ_u . When the prior variance of the gamma is larger, there is a stronger within individual correlation.
2. Calculate the Hellinger distance between each curve and the assumed mean sensitivity curve. The squared Hellinger distance between two exponential curves with rate parameters α and β is

$$1 - 2 \frac{\sqrt{\alpha\beta}}{\alpha + \beta}.$$

3. Randomly choose N rates from the empirical sample generated above, weighted by the inverse of the Hellinger distance.




Figure 2.12: A sample of the generated random sensitivity curves. The gray lines are the sensitivity curves for a subset of infants in the study, with the pointwise mean curve in black. The red line is the assumed sensitivity curve used in the likelihood. Increasing variability is denoted A-D.

Chapter 3

AN EXTENDED ILLNESS-DEATH MODEL**3.1 Introduction**

In studies on mother-to-child transmission of HIV, it is well-established that an infant remains at risk for infection throughout the duration of breastfeeding [Miotti et al., 1999, Becquet et al., 2009]. Yet, in resource-limited settings, it is recommended that an infant is breastfed to benefit from circulating maternal antibodies and reduce the risk of mortality. Analyses of the association between HIV infection, infant mortality, and weaning practices are faced with several challenges, including imperfect sensitivity in the detection of disease and interval censoring due to a discrete visit schedule. Additionally, after the cessation of breastfeeding, the infant is no longer at risk for mother-to-child transmission (MTCT) of HIV. Thus, in studies where infants are followed beyond the duration of breastfeeding, weaning must be considered as a competing risk [Alioum et al., 2003]. In Chapter 2, we introduced an illness-death model that adjusted for misclassification and interval censoring, and in this chapter, we expand the model to include weaning as an additional event.

Methods that address both misclassification and weaning in the estimation of HIV infection rates in infants have had limited treatment in the literature. Balasubramanian and Lagakos [2003] defined a joint distribution of time to illness and time to weaning, constraining the risk of time to illness to be 0 after time of weaning. This approach is similar to the multistate model, where an infant's time to illness is censored at the time of weaning if weaning precedes infection. To adjust for misclassification, Balasubramanian and Lagakos defined a step function for sensitivity that depends on time since infection and incorporated a constant specificity to allow for false positives. Brown and Chen [2010] found estimates of cumulative incidence differed when comparing an approach that censors at time of weaning

to an approach that ignores weaning as a competing risk. They accounted for interval censoring in their estimation algorithm by incorporating an imputation step that adjusts for increasing sensitivity. Neither method included time to death as an outcome in the model.

In this chapter, we propose a multistate model to examine the outcomes of time to illness, death, and weaning (Figure 3.1), where time to illness is not observed exactly. Commenges [2002] developed a frequentist approach to estimate an illness-death model where illness was interval censored, but did not adjust for misclassification. While their method could be extended to include weaning as another state in the MSM, the numerical optimization routines needed for estimation can become cumbersome as the number of parameters increase. As an alternative, complex likelihoods may be simplified under a Bayesian formulation, where it is possible to impute time to illness by sampling from its posterior distribution. In similar applications, conditional on the posterior samples of the unobserved event times, authors have then implemented multiple imputation [Brown and Chen, 2010], Monte Carlo EM algorithms [Sutradhar and Cook, 2008], or MCMC sampling algorithms [Yu, 2010] to estimate model parameters.

To estimate the multistate model, we propose a Bayesian MCMC algorithm, where we sample time of illness from a posterior distribution that adjusts for time-dependent sensitivity. Then, conditional on the sampled time, we implement the MCMC algorithm introduced by Kneib and Hennerfeind [2008] to estimate an exactly observed flexibly parametric multistate model, with spline-based transition intensities. The rest of this chapter proceeds as follows. In Section 3.2, we develop the likelihood and introduce the two-stage MCMC sampling algorithm. In Section 3.3, we conduct a simulation study to explore how robust our estimates are to misspecification of the time-dependent sensitivity curve. In Section 3.4, we apply our method to the HPTN 024 clinical trial conducted in Africa in 2001-2003. We conclude with a discussion in Section 3.5.

3.2 Methods

3.2.1 The multistate model likelihood

We consider a disease process as a path through a series of states, where a transition from one state to the next is defined by an event such as a new stage of a worsening disease or death. Progression through the disease process results in an accrued history of event times, and the collection of these events can be described by a multistate model (MSM). In an MSM, each event is denoted by a state $h \in \mathcal{S}$, where \mathcal{S} is the full state-space of all possible events in the process. Let T_{hj} represent the transition time from state h to state j , $h, j \in \mathcal{S}$. Similar to the hazard function of a univariate survival model, the instantaneous risk or hazard of transitioning from state h to state j is the transition intensity

$$\alpha_{hj}(t | \mathcal{H}_{t-}) = \lim_{\Delta t \rightarrow 0} \frac{\Pr(T_{hj} \in (t, t + \Delta t) | T_{hj} > t, \mathcal{H}_{t-})}{\Delta t},$$

where \mathcal{H}_{t-} is the history of the process up to time t . We denote the cumulative intensity as $A_{hj}(t) = \int_0^t \alpha_{hj}(s) ds$. We represent the underlying true disease process of mother-to-child transmission of HIV with are four possible states (HIV-negative/breastfeeding, HIV-negative/non-breastfeeding, HIV-positive, death) and five possible transitions (see Figure 3.1). We may observe this process with error, and our approach to account for this error is discussed in Section 3.2.2.




Figure 3.1: Possible transitions in the mother-to-child transmission of HIV.

To estimate the MSM, it is common to assume that the process is either Markov or

semi-Markov. Under a Markov assumption, the transition intensity is conditional only on the current state and no other part of the history, so that $\alpha_{hj}(t|\mathcal{H}_{t-}) = \alpha_{hj}(t)$. Under a semi-Markov assumption, the transition probability depends on both the current state and time entered into the state, t_h , so that $\alpha_{hj}(t|\mathcal{H}_{t-}) = \alpha_{hj}(t|t_h)$. We proceed under the semi-Markov assumption with a duration model, defined as $\alpha_{hj}(t_{hj}|t_h) = \alpha_{hj}(t_{hj} - t_h)$. See Andersen et al. [1993], Meira-Machado et al. [2009] for more a thorough introduction to MSMs.

Under the semi-Markov assumption, if the current state and time of entry into the state are known, each transition can be estimated as an independent survival model, conditional on covariates, where the outcome is time in the current state [Putter et al., 2007]. More specifically, for individual $i = 1, \dots, N$, let $T_{hj,i}, j \neq h$ index all transition times from state h and C_i be the censoring time, where censoring is independent of $T_{hj,i}$. Define the observed event time as $t_{hj,i} = \min(T_{hj,i}, C_i)$ and the indicator for a transition into state j as $\delta_{hj,i} = 1(T_{hj,i} = \min(T_{hj,i}, C_i))$, with the indicator function equalling one if the condition is true. For the set of all transition times $\mathbf{t} = \{t_{hj,i}, i = 1, \dots, N, h \neq j\}$ and indicators $\boldsymbol{\delta} = \{\delta_{hj,i}, i = 1, \dots, N, h \neq j\}$, the joint distribution is

$$f(\mathbf{t}, \boldsymbol{\delta}) = \prod_{i=1}^N \prod_{h,j} \alpha_{hj}(t_{hj,i} - t_{h,i})^{\delta_{hj,i}} \exp\{-A_{hj}(t_{hj,i} - t_{h,i})\},$$

where $t_{h,i}$ is the time of transition into state h for individual i . Additionally, it is possible that $t_{h,i} \leq 0$, meaning that individual i transitioned into state h prior to the start of the study. We define a distribution of initial states $\boldsymbol{\eta} = \{\eta_h : \Pr(X(0) = h) > 0, \sum \eta_h = 1\}$, and the joint distribution is

$$f(\mathbf{t}, \boldsymbol{\delta}) = \prod_{i=1}^N \left[\prod_{h,j:t_{h,i} \leq 0} \eta_h \alpha_{hj}(t_{hj,i})^{\delta_{hj,i}} \exp\{-A_{hj}(t_{hj,i})\} \right. \\ \left. \times \prod_{h,j:t_{h,i} > 0} \alpha_{hj}(t_{hj,i} - t_{h,i})^{\delta_{hj,i}} \exp\{-A_{hj}(t_{hj,i} - t_{h,i})\} \right]. \quad (3.1)$$

3.2.2 Accounting for imperfect sensitivity and interval censoring

In some cases, the true transition time may not be observed exactly. For example, time to illness, depicted as T_{12} in Figure 3.1, is often assessed through imperfect tests administered at pre-specified visit times, resulting in misclassification and interval censoring. To account for imperfect sensitivity, we define the random variable $T_{12}^d = T_{12} + D$, where D is a random variable measuring the time after infection needed for a disease to reach detectable levels. T_{12}^d is the point in time after which a diagnostic test results in a true positive and the process has reached a state of detection. We refer to the event time T_{12}^d as time to detectability. Conditional on the true illness time, we assume duration until detectability is independent of all other events, so that the joint distribution is the product $f(t_{12}, d) = f(t_{12})f(d)$. Additionally, time of detectability is interval censored between the last negative test and the first positive test, which we denote as $t_{12}^d \in (t_l, t_r)$. We define $t_r = \infty$ to denote the case where no first positive test is observed. Thus, $d \in (t_l - t_{12}, t_r - t_{12})$. We follow the typical approach to interval censored data and integrate over the unknown d to find the likelihood for the observed illness range

$$f(t_l, t_r | t_{12}) = \int_{t_l - t_{12}}^{t_r - t_{12}} f(d | t_{12}) f(t_{12}) dd.$$

Conditional on the observed data, the distribution for the latent illness time is

$$f(t_{12} | t_l, t_r) \propto \Pr(t_l - t_{12} < d < t_r - t_{12}) f(t_{12}),$$

where $\Pr(t_l - t_{12} < d < t_r - t_{12})$ can be viewed as a sensitivity function that weights the underlying true density for illness.

Incorporating illness into the joint likelihood for all events (Equation 3.1),

$$\begin{aligned}
 f(\mathbf{t}, \boldsymbol{\delta}; d) = & \prod_{i=1}^N \left[\Pr(t_{l,i} - t_{12,i} < d < t_{r,i} - t_{12,i}) \right. \\
 & \times \eta_1 \alpha_{12}(t_{12,i}) \exp \{-A_{12}(t_{12,i} - t_{2,i})\} \\
 & \times \prod_{\substack{hj: hj \neq 12 \\ t_{h,i} \leq 0}} \eta_h \alpha_{hj}(t_{hj,i})^{\delta_{hj,i}} \exp \{-A_{hj}(t_{hj,i})\} \\
 & \left. \times \prod_{\substack{hj: hj \neq 12 \\ t_{h,i} > 0}} \alpha_{hj}(t_{hj,i} - t_{h,i})^{\delta_{hj,i}} \exp \{-A_{hj}(t_{hj,i} - t_{h,i})\} \right]. \quad (3.2)
 \end{aligned}$$

Here, $t_{l,i}, t_{r,i}$ are the time of the last negative and first positive tests for individual i , respectively. We use the notation $hj : hj \neq 12$ to index the set of all transitions, excluding the unobserved transition into the illness state. We set $t_{12,i} = 0$ and $t_{2,i} = 0$ if the individual enters the process already in the illness state. To ensure the sensitivity function is increasing over time, we select a random variable D with a monotonically increasing cumulative distribution function. In our application, we assume $D \sim \text{Exp}(\mu_d)$ for $t_{12,i} > 0$ and $D \sim \text{Shifted Exp}(\mu_d, \tau)$ for the case $t_{12,i} \leq 0$, where μ_d is fixed and known and τ is a shift parameter that allows for a non-zero sensitivity at birth. The simple parametric form of an exponential facilitates sensitivity analyses that vary over the one parameter.

3.2.3 Parameterizing a flexible, smooth transition intensity

In this section, we establish a fully parametric and flexible form for the baseline transition intensities and add explanatory covariates under a proportional hazards model. First, for a flexible smooth function, we parameterize the log baseline intensities as cubic B-splines. Under this approach, the log-baseline intensity for the transition from state h to state j is defined as a linear combination of B-spline bases, so that under the proportional hazards

assumption

$$\log \alpha_{hj}(t) = \sum_{k=1}^{m^{(hj)}} \beta_k^{(hj)} z_k^{(hj)} + \sum_{k=1}^{\kappa+4} b_k^{(hj)} B_k^{(hj)}(t), \quad (3.3)$$

where for transition hj , $z_k^{(hj)}, \beta_k^{(hj)}$, $k = 1, \dots, m^{(hj)}$ represent explanatory covariates and corresponding coefficients, respectively, with the number of covariates $m^{(hj)}$ possibly differing across transitions, and $B_k^{(hj)}(t), b_k^{(hj)}$, $k = 1, \dots, \kappa + 4$ are the cubic B-spline bases at time t and corresponding spline coefficients, respectively. Here, $\kappa + 4$ is the number of B-spline bases, which is determined by the number of internal knots (κ) plus order of the cubic function [de Boor, 2001]. While the specified knots can differ across transitions, in practice, we have been using the same spline bases for all transitions, and in this case, we simply designate the set of bases as $B_k(t), k = 1, \dots, \kappa + 4$. Depending on the availability of sufficient data, a large number of knots are specified so that the shape of the intensity function can better fit the data. This allows for flexibility in estimation but can result in widely varying B-spline coefficients, leading to overfitting and highly oscillating transition intensities. One approach to ensure sufficient smoothness is to penalize adjacent B-spline coefficients. Eilers and Marx [1996] proposed a simple penalty based on the second order differences of adjacent coefficients $\sum_k \left(b_k^{(hj)} - 2b_{k+1}^{(hj)} + b_{k+2}^{(hj)} \right)^2$. Lang and Brezger [2004] extended this penalty to the Bayesian analogue where a random walk prior describes the relationship between adjacent coefficients. As in Lang and Brezger, we define a second-order random walk prior for the coefficients of the baseline intensity for transition hj as

$$b_k^{(hj)} = 2b_{k-1}^{(hj)} - b_{k-2}^{(hj)} + \epsilon_k^{(hj)}, \quad k = 3, \dots, \kappa + 4,$$

with Gaussian errors $\epsilon_k^{(hj)} \sim N(0, \sigma_{hj}^2)$. Then, with $b_1^{(hj)}, b_2^{(hj)} \propto 1$, the joint prior distribution for the set of coefficients is

$$\begin{aligned} p(\mathbf{b}^{(hj)} | \sigma_{hj}^2) &\propto (\sigma_{hj}^2)^{\frac{\kappa+2}{2}} \exp \left\{ -\frac{1}{2\sigma_{hj}^2} \sum_{k=3}^{\kappa+4} \left(b_k^{(hj)} - 2b_{k-1}^{(hj)} + b_{k-2}^{(hj)} \right)^2 \right\} \\ &\propto (\sigma_{hj}^2)^{\frac{\kappa+2}{2}} \exp \left\{ -\frac{1}{2\sigma_{hj}^2} \mathbf{b}'^{(hj)} \mathbf{K} \mathbf{b}^{(hj)} \right\}, \end{aligned} \quad (3.4)$$

where $\mathbf{b}^{(hj)} = (b_1^{(hj)}, \dots, b_{\kappa+4}^{(hj)})$ is the vector of spline coefficients and \mathbf{K} is a less than full rank precision matrix of a multivariate Gaussian distribution. Under this parameterization, the cumulative intensity is intractable, but can be numerically approximated using Gaussian quadrature. The second order differences penalty term is reflected in the precision matrix and is controlled by the variance of the random walk σ_{hj}^2 , with smaller variances leading to smoother curves. We estimate the level of smoothing as part of the MCMC algorithm and place an inverse gamma prior on σ_{hj}^2 , with hyperpriors shape a and scale c . Finally, we assume non-informative Gaussian priors for all covariate parameters with $\beta_k^{(hj)}$, for $k = 0, \dots, m$.

3.2.4 Estimation and the MCMC algorithm

Let data = $\{\mathbf{t}_l, \mathbf{t}_r, \mathbf{t}_{hj:hj \neq 12}, \boldsymbol{\delta}_{hj:hj \neq 12}\}$ denote the set of observed data, where $\mathbf{t}_l, \mathbf{t}_r$ represent the last negative and first positive visits over all individuals and $\mathbf{t}_{hj:hj \neq 12}, \boldsymbol{\delta}_{hj:hj \neq 12}$ are event times and indicators for each transition hj , with the exception of the unobservable illness time, over all individuals. The joint posterior combines Equations 3.2 and 3.4,

$$\begin{aligned} p(\boldsymbol{\beta}, \mathbf{b}, \boldsymbol{\sigma}^2, \mathbf{t}_{12}, \boldsymbol{\eta} | \text{data}) &\propto \\ &\prod_{i=1}^N \left[\prod_{hj:hj \neq 12} f(t_{hj,i}, \eta_h | \boldsymbol{\beta}^{(hj)}, \mathbf{b}^{(hj)}, \sigma_{hj}^2) f(t_{12,i}, d_i, \eta_1 | \boldsymbol{\beta}^{(12)}, \mathbf{b}^{(12)}, \sigma_{12}^2, \mu_d) \right. \\ &\quad \left. \times \prod_{hj} p(\mathbf{b}^{(hj)} | \sigma_{hj}^2) IG(\sigma_{hj}^2 | a, c) \right]. \end{aligned} \quad (3.5)$$

The β , \mathbf{b} , and σ^2 are indexed over all transitions, and the $\boldsymbol{\eta}$ are indexed over initial states.

We use a Gibbs sampler to sample from the full conditionals for each of the parameters [Gelfand and Smith, 1990]. To sample coefficients of the spline-based intensity function, we implement a conditional prior proposal Metropolis-Hastings algorithm introduced by Knorr-Held [1999] and applied in the context of smoothing B-spline based functions by Fahrmeir and Lang [2001]. We sample the covariates using an automatically tuned random-walk Metropolis-Hastings algorithm [Çetinyürek and Lambert, 2011]. Due to the points of discontinuity in the support for the underlying illness time, sampling each $t_{12,i}$ involves sampling first from a mixed discrete-continuous posterior to sample within a subset of the support. Specifically, if individual i tests positive, the true time of illness could be in utero/during delivery, between birth and the first negative test, or between the last negative test and the first positive test. If no first positive test is observed, the true time of illness could be in utero/during delivery, between birth and the first negative test, between the last negative test and time of death or time of censoring, or after death or censoring. Thus, we first sample the appropriate range. If either point mass is sampled (in utero/during delivery or after death/censoring), no further sampling is needed. However, to sample within the constrained continuous segments of the distribution, we used a slice sampler [Neal, 2003]. Further details of the sampling algorithm are outlined in Appendix 3.A at the end of this chapter. We implemented the code in C and R.

3.3 Simulations

In this section, we conducted a simulation study to assess how robust our covariate coefficient estimates are to misspecification of the sensitivity curve. In the derivation of the model introduced in Section 3.2, we assumed that the true sensitivity curve was known. In practice, we would specify a sensitivity curve based on previous studies found in the literature, which may result in a misspecification of the sensitivity curve and biased estimates. To examine any potential biases, we fixed the covariate effects at known values and generated data sets assuming an exponentially distributed delay in detection with rate $\mu_d = 1.5$ months,

for an approximate mean 3 weeks delay. We estimated the model assuming no delay in detection, and delay rates $\mu_d = 1.5, 0.75, 0.67$, which correspond to mean delays of 3, 6, and 6.5 weeks, respectively. In all cases, we specified the distribution of D as a shifted exponential with rate μ_d and shift $\tau = 0.5$ when $t_{12,i} \leq 0$ and an exponential with rate μ_d when $t_{12,i} > 0$. The specification for τ is similar to the estimated shifted exponential indicated in Balasubramanian and Lagakos [2003] and Brown [2010]. We repeated the above simulations twice, assuming a generating delay rate $\mu_d = 0.75$ and a generating delay rate $\mu_d = 0.67$.

In generating the data sets, we used true baseline intensities similar to the estimated intensities of our motivating example. Because previous breastfeeding studies have determined a constant illness rate over the first year is a reasonable assumption in postnatal mother-to-child transmission of HIV [Coutsoudis et al., 2004], we specify the illness rate $\alpha_{12}(t) = \alpha_{12} = 0.0065$ per month, roughly based on illness rates estimated below in Figure 3.2. The remaining baseline intensities are spline-based, parameterized with 12 knots each. For the rate of weaning, $\alpha_{13}(t)$, we specified that all participants weaned at the rate estimated for one site of our motivating example, Lusaka, Zambia (see Figure 3.3a). We based the mortality rate for all uninfected infants on the estimates of the healthy to death transition and the mortality rate for infected infants on the estimates of illness to death transition. We generated the data assuming a rate of illness associated with treatment and CD4 counts with log intensity ratios -0.2 and -1.32, respectively. We assumed no covariate association on any of the other transition intensities. The process we used to simulate data sets is outlined in Appendix 3.B.

For each simulation, we generated 100 data sets. We fit the model using the MCMC process outlined in Appendix 3.A. For each data set, we ran a single chain for 60,000 iterations, and removed 10,000 for burn in. Bias of the posterior mean, and posterior standard errors and coverage for 95% highest posterior density (HPD) intervals are presented in Table 3.1. To preserve efficiency, we did not thin the chains for inference [Link and Eaton,

2011].

Estimates were least biased when the correct sensitivity curve was assumed. Standard errors increased as the assumed mean delay increased because larger delays introduce additional uncertainty about the true time of illness. The estimates of the treatment coefficient for the rate of illness were fairly robust to misspecification, with a similar amount of bias and 89% or higher coverage across all estimates from the adjusted models. On the other hand, the estimates for the continuous covariates, representing baseline illness intensity and CD4 count, performed the worst, with higher bias and lower coverage rates, when the sensitivity curve was misspecified. The parameter for treatment and mortality rates of the infected had coverage rates of 100%, resulting from large confidence intervals.

We next compared the three cases where models were fit assuming the correct sensitivity curve (highlighted in bold in Table 3.1). As the mean delay increased, estimates for the rate of illness had lower bias (-0.248 for mean 3 weeks compared to -0.072 for mean 6.5 weeks) and higher coverage (74% for mean 3 weeks compared to 91% for mean 6.5 weeks). However, the direction was reversed for estimates of the CD4 covariate coefficient (denoted $\beta_2^{(12)}$ in Table 3.1). As mean delay increased, the estimates of $\beta_2^{(12)}$ increased in bias (0.009 for mean 3 weeks compared to 0.145 for mean 6.5 weeks) and had decreased coverage (95% for mean 3 weeks compared to 81% for mean 6.5 weeks). The estimates were fairly robust across specification of the sensitivity curve when the effect was small or there was no effect at all.

Overall, the simulation results demonstrate that failure to adjust for misclassification results in large biases and poor coverage, and incorporation of any sensitivity curve improves estimates. However, some bias remains after adjustment for misclassification when the sensitivity curve is misspecified. The degree of bias depends on the magnitude of the coefficient and the level of misspecification of the sensitivity curve. For example, the mild association of $\beta_1^{(12)} = -0.2$ (standard error ≈ 0.2) was fairly robust to misspecification of the sensitivity curve, with bias that was under 20% of the truth and coverage of at least

85% for each of the assumed sensitivity curves used in the simulation study. On the other hand, for the stronger association of $\beta_2^{(12)} = -1.32$ (standard error ≈ 0.2), the estimate was estimated with 30% bias and coverage of 34% when the sensitivity curve was specified with a shorter delay than the true delay. Thus, our simulations showed that if there is a concern that the curve could be highly misspecified, erring on the side of specifying a shorter delay than true may be an appropriate choice because the resulting estimates are attenuated toward the null and can be considered as a conservative statement of the level of association between the covariate and outcome.

3.4 Application

3.4.1 Data

In this section, we illustrate our method using data collected as part of the HIV Prevention Trials Network (HPTN) 024 study, conducted in 2001-2003 to assess if the administration of antibiotics to the mother would reduce MTCT. When the study was conceived, antiretroviral use at birth had not yet been established as standard of care, and one thought was that administration of antibiotics would reduce the rate of two comorbid conditions associated with an increased risk of HIV transmission, leading to a corresponding decrease rate of transmission. Early in the study, however, the study protocol was changed to reflect an updated standard of care, and all mothers and infants were given nevirapine at the onset of labor and within 72 hours of birth, respectively. In addition to nevirapine, mothers were randomized to receive antibiotics or placebo at the onset of labor. The trial was stopped early when treatment did not decrease the rate of infection [Taha et al., 2006].

Our sample included 1845 women and infant pairs, across four sites in Africa (Lusaka, Zambia; Lilongwe and Blantyre, Malawi; Dar-es-Salaam, Tanzania). Mothers were enrolled at 20-24 weeks gestation and followed until delivery. Follow-up visits for infants occurred at 4-8 weeks, and 6, 9, and 12 months. Samples collected prior to the 12 month visit were tested only if 12 month sample was positive or missing. Thus, per protocol, a majority of

Table 3.1: Results from 100 simulated data sets examining the robustness of sensitivity curve specification. Each column contains bias (observed - expected), standard error, and 95% HPD interval coverage for data generated under a specific sensitivity curve. From left to right, data were generated assuming a 3, 6, and 6.5 mean delay in detection. Rows contain the estimated results for four parameters of interest (indicated by transition, type of data, and true value). The covariates represent baseline transition intensity, treatment on illness rate, CD4 on illness rate, and treatment on infected mortality rates. The models were all estimated with no delay assumed or an assumed mean delay of 3, 6, and 6.5 weeks. Blank entries were not estimated. Bold entries indicate the model was estimating assuming the true sensitivity curve.

Parameter	Est.	Assumed mean delay in data generation (truth)					
	Delay	3 weeks		6 weeks		6.5 weeks	
	(wks)	Bias	(SE, Cov.)	Bias	(SE, Cov.)	Bias	(SE, Cov.)
Illness rate	0	0.521	(0.133, 1)	0.772	(0.124, 0)	0.827	(0.109, 0)
$\log(\alpha_{12})$	3	-0.248	(0.223, 74)	0.179	(0.233, 72)	0.269	(0.222, 59)
truth: -5	6	-0.496	(0.307, 46)	-0.191	(0.291, 85)		
	6.5	-0.591	(0.316, 38)			-0.072	(0.256, 91)
Discrete 0,1	0	0.096	(0.156, 86)	0.095	(0.143, 84)	0.092	(0.147, 88)
$\beta_1^{(12)}$	3	0.010	(0.214, 91)	0.037	(0.200, 89)	0.028	(0.178, 93)
truth: -0.2	6	-0.009	(0.248, 92)	0.011	(0.232, 94)		
	6.5	-0.018	(0.262, 94)			0.004	(0.206, 97)
Continuous	0	0.625	(0.113, 0)	0.772	(0.102, 0)	0.806	(0.095, 0)
$\beta_2^{(12)}$	3	0.009	(0.188, 95)	0.314	(0.193, 47)	0.369	(0.202, 34)
truth: -1.32	6	-0.114	(0.250, 90)	0.074	(0.250, 86)		
	6.5	-0.163	(0.269, 88)			0.145	(0.233, 81)
Discrete 0,1	0	-0.003	(0.051, 100)	0.005	(0.049, 100)	-0.006	(0.049, 100)
$\beta_1^{(24)}$	3	-0.002	(0.048, 100)	0.008	(0.048, 100)	-0.005	(0.048, 100)
truth: 0	6	0.000	(0.048, 100)	0.010	(0.046, 100)		
	6.5	0.001	(0.047, 100)			-0.003	(0.045, 100)

infants had a test result only at the last visit seen. Our study period was birth through 12.5 months of age. Weaning and death were considered to be observed exactly. While it is possible that the time of weaning could be interval censored between the last visit and time of death, in the case where an infant was breastfed through the time of the last visit, we assumed that the infant was still being breastfed at the time of death/censoring. This is a simplifying assumption used to reduce the number possible states of the interval censored period between the last visit and death. In other words, suppose an infant dies without testing positive and no weaning is reported prior to the last visit. In the case where the true time of illness does not occur prior to death, without further assumptions, we would not be able to identify if the infant transitioned into weaning before death or remained uninfected and breastfeeding before death. While there are MCMC algorithms that allow for sampling without making this additional assumption, they are more complex. As a result of this simplifying assumption, we may be estimating the mortality rate from the weaned state. For reference, there were 52 infants who died without testing positive or reporting a time of weaning. Of the 52, 40% died within 2 months of their last visit, 65% within 3 months, and 79% within 4 months.

We were primarily interested in understanding the association between treatment and illness, within the context of the four-state MSM. Secondly, we were interested to determine if treatment was associated with the other four event times (weaning, death prior to illness and weaning, death after weaning, and death after illness). Finally, we note that several covariates (maternal CD4 count, log viral load, and weight at delivery, and infant weight at birth) were included as precision variables for all transition intensities to improve sampling of the underlying illness time, and as such, were standardized to have mean 0 and standard deviation 1 and should be interpreted on this scale.

3.4.2 Model definition

We applied our method to the four-state multistate model depicted in Figure 3.1. In the scientific literature, infection rates have been shown to be constant after 4-6 weeks of age [Kourtis et al., 2006]. Thus, we assumed a constant intensity function for illness. The log transition intensity for each transition, other than illness, is defined as a flexible spline-based function as in Equation 3.3. We stratified the transition for weaning by site and used a cubic B-spline function with 5 knots for each baseline weaning transition intensity. For the mortality transition intensities, we used cubic B-splines with 12 knots. As in Equation 3.3, covariates were included under the proportional hazards assumption, and based on a priori beliefs, we selected which covariates included for each transition in the model. With only 15 infant deaths observed after the infant had weaned, we assumed that the weaning deaths and uninfected deaths shared a baseline transition intensity. The specification for each log transition intensity is outlined in Table 3.2.

Based on the scientific literature, we specified the delay parameter $D \sim \text{Exp}(\mu_d)$ for in utero/intrapartum transmissions and $D \sim \text{Shifted Exp}(\mu_d, \tau = 0.5)$ for postnatal transmissions, with $\mu_d = 0.75, 1, 1.33$ months [Wessman et al., 2012]. Respectively, these correspond to 95% sensitivity at approximately 1.75, 2.5, and 3.5 months after birth. Additionally, we estimated the model assuming perfect sensitivity.

3.4.3 Estimation and Convergence

We fit the model following the sampling algorithm outlined in Section 3.2.3. For transition hj , the prior for the spline coefficients was a second-order random walk with variance σ_{hj}^2 . We specified a noninformative inverse gamma (0.005, 0.005) conjugate prior on σ_{hj}^2 . We specified noninformative proper priors for the remaining coefficients, placing $N(0, 4)$ priors on all covariates. Finally, we placed a $N(\log(0.0065), 4)$ prior on the log of the

Table 3.2: Specification of the log transition intensity $\log \alpha_{hj}(t) = \sum_{k=1}^m \beta_k^{(hj)} z_k + \sum_{k=1}^{\kappa+4} b_k^{(hj)} B_k^{(hj)}(t)$. The intercept term of the (34) transition is the proportionality constant so that $\alpha_{34}(t - t_\omega) = \theta_{34} \alpha_{14}(t)$, where t_ω is time of weaning.

Transition	(hj)	Covariates ($\beta_k^{(hj)}$)	Spline parameters ($b_k^{(hj)}$)
Healthy \rightarrow illness	(12)	Intercept Treatment Maternal CD4 Maternal viral load Maternal weight Infant weight Site [†] = Lusaka, Dar, Blantyre	all set to 0
Healthy \rightarrow weaning	(13)	Maternal CD4 Maternal viral load Maternal weight Infant weight	Stratified by site Each site with 12 knots
Healthy \rightarrow death	(14)	Treatment Maternal CD4 Maternal viral load Maternal weight Infant weight Site [†] = Lusaka, Dar, Blantyre	A single intensity function with 12 knots
Illness \rightarrow death (time since onset)	(24)	Treatment Maternal CD4 Maternal viral load Maternal weight Infant weight	A single intensity function with 20 knots
Weaning \rightarrow death (time since weaned)	(34)	Intercept (for PH* with $\alpha_{14}(t)$) Maternal CD4 Maternal viral load Maternal weight Infant weight	Same coefficients as (14)


*PH = proportional hazards; [†]Site covariate is relative to baseline site Lilongwe

Table 3.3: Basic descriptive statistics for the observed transitions of the MTCT model, stratified by location. In the first transition (illness, weaning, or death), if an infant experienced more than one event, only their first event is counted. Blantyre has a low rate of postnatal first positive tests after 2 months of age, while Lusaka remains relatively higher over the course of the study. There are no first positive tests in Dar after 6 months of age, since most infants have weaned.

	Site			
	Blantyre	Dar es Saalam	Lilongwe	Lusaka
	Count (%)	Count (%)	Count (%)	Count (%)
Pos. at birth	11 (3)	14 (4)	44 (7)	24 (5)
Pos. before end of study, after neg. at birth	80 (20)	52 (15)	74 (12)	90 (19)
Weaned, not pos.	59 (15)	248 (76)	55 (9)	62 (13)
Death, not pos.	16 (4)	15 (4)	17 (3)	19 (4)
Death, test pos.*	43 (48)	29 (44)	48 (41)	42 (37)

*Denominator is infants who test positive.

Kernel smoothed densities of time to first positive test, after testing negative at birth



constant illness intensity parameter, which reflects a rate found in similar studies (see, for example, Miotti et al. [1999], Nduati et al. [2000], Coovadia et al. [2007]).

We ran two chains, each with 200,000 iterations, with different starting values and different seeds. After removing a 75,000 burn-in from each, and thinning by 25, our final posterior sample size for inference was 10,000 samples. Trace plots for the covariate coefficients showed the chains mixed well across all four assumed sensitivity curves. The Gelman-Rubin statistic compares within chain variances to the combined chain variance, and a statistic near 1 indicates good mixing [Brooks and Gelman, 1998]. The Gelman-Rubin statistics for our estimated covariate coefficients were all near 1, with the exception of the intercept term of the weaning to death transition intensity. With only 15 deaths after weaning, this was due to a lack of power, rather than a lack of convergence, confirmed by widely variable trace plots (see Figure 3.9 in Appendix 3.C). Finally, we compared the posterior density estimates of the covariates to their prior distributions and found that posterior estimation was strongly driven by the data. Example trace plots (Figure 3.9) and posterior density plots (Figure 3.10) are included in Appendix 3.C.

3.4.4 Results

Descriptive statistics for the observed event time data are included in Table 3.3. We plotted posterior medians and posterior 95% HPD intervals for log rate of illness and all covariates in Figure 3.2. Treatment had no statistically significant effect on any of the transitions we tested (healthy to illness, healthy to death, and illness to death). The transition intensity for illness was associated with maternal health at delivery, with both lower CD4 cell counts and higher viral load levels associated with higher rates of MTCT.

Adjusting for sensitivity had the largest impact on the rate of illness. As the lag in detection increased, there was a decline in the log baseline illness rate (the rate for Lilongwe), with posterior medians ranging from -5.4 in the unadjusted model to -6.9 in the model that assumed a 6 week mean delay. This is to be expected since a longer assumed mean lag time

weights the early postnatal infections more toward to in utero/intrapartum infections. There was a similar decline for the log transition intensity ratios for both Dar and Blantyre. In the unadjusted model, both Dar and Blantyre had estimated illness rates that were significantly higher (with posterior median log transition intensities 2.1 and 0.6, respectively) than the baseline site, Lilongwe. But as the assumed mean delay increased, these estimates were adjusted toward the null and were insignificant under the 6 week delay assumption (with posterior median log transition intensities 0.7 and 0.07, respectively), meaning the estimates for the illness rate were not significantly different across the three sites for the 6 week delay model. From the descriptive statistics in Table 3.3, this is likely because Lilongwe had a higher proportion of infants test positive at birth than either Dar or Blantyre. But after adjusting for misclassification, the estimated proportion of infants truly born positive was not significantly different across the three sites, suggesting that the statistical difference of the unadjusted model might be an artifact of the observation process.

The significant associations, CD4 count, viral load and Lusaka site associated with rate of illness, were adjusted away from the null as the assumed mean delay increased (Figure 3.2). The HPD intervals for the illness transition covariates were larger in the models that assumed a longer delay in detection, reflecting uncertainty around true illness time is larger in this case. In the transitions other than illness, all covariate coefficient estimates were consistent across specified sensitivity curve. Any potential trends in covariate effects for transitions other than illness were small compared to the variability of the HPD intervals. This suggests that choice of sensitivity curve has the largest impact on estimates of the illness transition and the associated covariates, but little impact on covariate effects on weaning and death transition intensities. Thus, it is important to estimate the model assuming sensitivity curves that well-represent the range of sensitivities found in the literature.

We plotted the baseline weaning intensities for all four sites in Figure 3.3, assuming a mean 3 week delay in detection, though other assumed sensitivity curves resulted in similar estimated weaning intensities. The three sites that did not counsel early weaning had

similar weaning rates in general, though the curves were all shaped differently. For example, mothers in Lusaka weaned earlier, with an early peak at 6 months, while mothers in Lilongwe weaned later, with a peak at around 10 months. Dar es Saalam, which did counsel early weaning, saw a large increase in weaning rates between 2 and 4 months after birth. We also plotted the mortality curves from the healthy and illness states across sensitivity curves in Figure 3.4. As the assumed delay in detection increased, the early deaths in infants who did not test positive prior to death were weighted higher as deaths from illness, resulting in a decreased estimate of the healthy to death intensity for the first 6 months after birth. After 6 months, rates were similar across sensitivity curves (Figure 3.4a). In the illness to death transition intensity, the estimated mortality was higher around the 6 month and 10 month periods after adjusting for misclassification, when compared to no adjustment for misclassification. It is unclear why the increased mortality estimate was localized to two time spans. One hypothesis is that the time of death is not observed exactly, and in this case, approximating the time of death may result in clusters around early and mid-month dates.

3.4.5 *Checking the model fit*

To check the fit of the data, we simulated 100 data sets using the posterior median for each parameter as estimated in Section 3.4.4, using the model that assumed a mean 3 week delay in detection. We simulated transition times from each of the five conditionally independent survival models of the semi-Markov MSM. We generated a delay in detection under the assumed sensitivity curve and a set of visit times for each individual following the empirical frequency of visits in HPTN 024. We then determined the observed pathway for each individual. The full algorithm for data generation is outlined in Appendix 3.B.

As a summary measure for the simulated data and for the motivating data set, HPTN 024, we plotted the non-parameteric cumulative incidence functions (CIF) for each event of the competing risks of time to first positive, time of weaning, and time of death from the


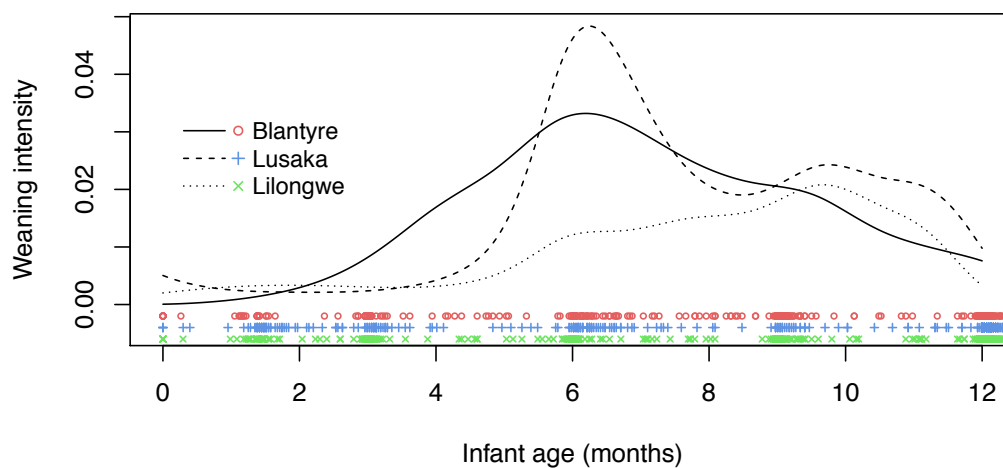



Figure 3.2: Posterior medians and HPD intervals for each covariate estimate. Columns contain the transitions and rows contain the covariates intercept, randomization, maternal CD4 cell count, log viral load, and weight at delivery, infant weight at birth, and indicators for sites at Lusaka, Dar es Salaam, and Blantyre (with Lilongwe as the baseline site). CD4 counts, log viral load, and maternal/infant weights are centered and scaled, so that effect size is interpreted on the unit standard deviation scale. The intercept term for the weaned/non-breasted to death transition is the constant for proportionality to the healthy to death transition. The gray line indicates a null effect. Each plot contains four intervals, with the far left indication the estimation assuming perfect sensitivity, increasing to an assumed delay with a mean of 6 weeks on the far right. *Mothers at the Dar site were counseled to wean when their infants were 6 months old.




(a) Estimated weaning intensities for the three sites that did not counsel early weaning. The estimated spline coefficients were taken as the median of the posterior distribution for each coefficient.

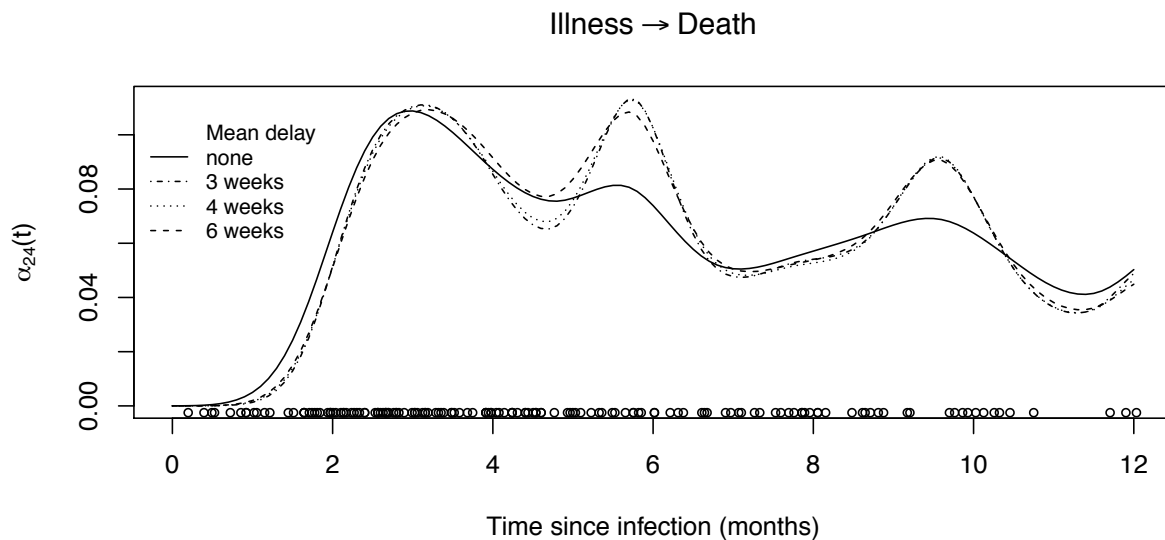


(b) Estimated weaning intensity for Dar, where mothers were counseled to wean at 6 months.

Figure 3.3: Estimated weaning intensities from the model fit in Section 3.4.4. In this fit, the assumed mean delay was 3 weeks.



(a) Baseline healthy to death transition intensity, across levels of smoothing. The points indicate observed deaths among infants who did not test positive or wean before death. As the assumed mean delay of the sensitivity curve increases, early deaths are weighted as deaths from illness, resulting in a decreased estimate of the healthy to death transition intensity. Mortality rates are similar across sensitivity curves after 6 months.



(b) Baseline illness to death transition intensity, across levels of smoothing. Time is duration of illness prior to death. The points indicate observed illness duration (measured as time of death - time of first positive test) among infants who tested positive before death. After accounting for misclassification, there is a spike in mortality rates at 6 months and 10 months after the onset of illness.

Figure 3.4: Mortality transition intensities for transitions from the healthy and death states.

healthy state (see Andersen, Abildstrom, and Rosthoj [2002] for more on CIFs). The results for each outcome, stratified by site, are found in Figures 3.5 - 3.7. The gray lines represent the distribution of expected observed CIFs. In other words, if our model accurately captures the underlying true distribution of the three event times and misclassification due to a delay in detection, our generated data sets should produce CIFs similar to the observed CIF of HPTN 024 once interval censoring is imposed.

Overall, the CIFs for HPTN 024 were contained within the CIFs of the simulated data sets, though with some minor differences in shape. For example, in Figure 3.5, the HPTN 024 CIF for Blantyre was flatter than the simulated CIFs. This may be a result of using a simulated visit schedule in data generation. Because the complete visit schedule for each participant was unavailable, we simulated visit times in each data set according to an empirical distribution of visits established from the example data set, and we assumed that there were no missed visits. In the case of the site Blantyre, if a large number of 9 month visits were missed, most infections over the 6-12 month range would not be detected until the 12 month visit, leading to a generally flat observed CIF, with a spike at the end. However, in generating the data, we assumed no missed visits, and so infections were detected throughout the study and the CIF increased over the duration of the study. As another example, the model over-estimated early events for Lilongwe and Lusaka, and underestimated later events in Lusaka. This may be because the rate of illness is not constant in these two cases, though if that is the case, we can consider a constant illness rate as a reasonable approximation, since the HPTN 024 curve is contained within the variability of the simulated data CIFs. Finally, we note that the weaning pattern in Dar is reflected in the variability of the CIFs. Because approximately 75% of the infants had weaned by 12 months (prior to testing positive), the low rates of illness and death from the healthy state led to widely variable CIFs after 6 months and less variability in the weaning CIF, as compared to the other three sites.




Figure 3.5: Non-parametric cumulative incidence functions for the time to first positive test in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained within in the CIFs for the simulated data sets across all sites.




Figure 3.6: Non-parametric cumulative incidence functions for the time to weaning in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained in the CIFs for the simulated data sets across all sites.




Figure 3.7: Non-parametric cumulative incidence functions for the time to death from the healthy state in 100 simulated data sets (in gray) and HPTN 024 (in black), stratified by site. The CIF for HPTN 024 was contained in the CIFs for the simulated data sets across all sites.

3.4.6 Comparison to the illness-death model

In the MTCT MSM depicted in Figure 3.1, the first transition is to one of three possible states - illness, death, or weaning. If most infants are breastfed over the duration of the study or only early outcomes are of interest, the risk of weaning is low and weaning can be excluded as a competing risk. However, weaning should be considered as a competing risk if most infants wean during the duration of the study period. This may occur if early weaning is recommended as part of a particular study protocol or is the norm in the population of interest (i.e. Becquet et al. [2009]) or if long term efficacy of a particular treatment is the study question of interest (i.e. Leroy et al. [2002]). If weaning is ignored as a competing risk in these cases, an infant who is weaned will continue to contribute to the denominator of the estimate of illness rate, but cannot contribute to the numerator. To see how ignoring a competing risk might influence model estimates, we refit the model of Section 3.4.4 as an illness-death model. The results are in Figure 3.8.

The estimates from the illness-death were similar to those from the model that included weaning as an event, with one exception. The rate of illness for Dar was estimated lower in the illness-death model than in the full model estimated in Section 3.4.4. For example, assuming a 4 week mean delay sensitivity curve, the posterior median (95% HPD interval) log transition intensity was 1.2 (-1.3, 2.3) in the full model and -1.6 (-4.6, 0.3) in the illness-death model. This is because with approximately 73% of infants in Dar weaning by 6 months, weaning is a measurable competing risk. In the illness-death model, infants who are weaned remain in the denominator, leading to the lower estimate for rate of illness in Dar. Additionally, when weaning is ignored, infants who have weaned are included in the denominator for the transition from healthy to death. Thus, the confidence intervals are much narrower for the site=Dar coefficient when compared to confidence intervals from the model that includes weaning. In the remaining three sites, only 11% - 15% of infants had weaned by 6 months, and 29% - 37% had weaned by the end of the study, and so most infants remained at risk for illness over the duration of the study. As a result of

the low risk of weaning, the illness-death model and the model that includes weaning are approximately equivalent for the three sites, leading to similar estimates for the rate of illness, randomization, maternal CD4 count, viral load, and weight at delivery, and infant weight at birth across both models (Figures 3.2 and 3.8).

Finally, we compared this version of the illness-death model to the illness-death model fit in Chapter 2. In the Chapter 2 model, we estimated only one covariate effect (treatment/randomization) for each of the three transitions (healthy to illness, healthy to death, and illness to death; see Chapter 2 Figure 2.5a). In both models, the covariate estimates for treatment were insignificant, though the point estimates for the Chapter 2 were positive, trending toward a harmful treatment effect with an increased risk of illness. Alternatively, the point estimates were negative in this application, trending toward a protective effect with a decreased risk of illness. At this point, it is not clear why this association changed direction and these differences require further investigation.

3.5 Discussion

In this chapter, we introduced a multistate model that describes the three main outcomes of MTCT, time to illness, time to death, and time to weaning, and accounts for imperfect sensitivity of diagnostic testing for HIV in infants and interval censoring of detectability. We approached estimation using a Bayesian MCMC algorithm where we sample an exact time of illness, accounting for misclassification and interval censoring. Then, conditional on the sampled time of illness, we have a multistate model where all event times are observed exactly, and the model parameters can be estimated with well-established methods for estimating MSMs.

By including weaning in our model, we can better understand how weaning intensities influence the estimates of baseline illness and mortality rates, as well as estimates of covariate effects on these intensities. For example, in estimating the model, we found that while the coefficient estimates are robust to misspecification of the weaning transition intensities, an incorrectly specified weaning intensity can lead to incorrect cumulative incidence esti-




Figure 3.8: Posterior medians and HPD intervals for each covariate estimate of the illness-death model. Most estimates were similar when compared to the full MTCT model (Figure 3.2). Columns contain the transitions and rows contain the covariates intercept, randomization, maternal CD4 cell count, log viral load, and weight at delivery, infant weight at birth, and indicators for sites at Lusaka, Dar es Salaam, and Blantyre (with Lilongwe as the baseline site). CD4 counts, log viral load, and maternal/infant weights are centered and scaled, so that effect size is interpreted on the unit standard deviation scale. The gray line indicates a null effect. Each plot contains four intervals, with the far left indication the estimation assuming perfect sensitivity, increasing to an assumed delay with a mean of 6 weeks on the far right. *Mothers at the Dar site were counseled to wean when their infants were 6 months old.

mates. In particular, when we assumed the three sites that did not counsel early weaning (Lusaka, Blantyre, and Lilongwe) were proportional in baseline weaning intensities, the illness cumulative incidence function for Lilongwe was overestimated. This is because mothers in Lilongwe that weaned their infants in the first year tended to wean in the 8-10 month range, while the other two sites had some early weaning around 6 months (Figure 3.3). While the magnitude of the weaning intensity was similar, the difference in shape lead to an erroneous conclusion about the incidence of illness. This demonstrates the importance of allowing flexibility in the baseline transition intensity for weaning.

As another example, we found that the parameter estimates for the covariate effects were similar when comparing the full model of Section 3.4.4 to the illness-death model of Section 3.4.6. The only differences were in the estimate for the rate of illness and the coefficient for the uninfected mortality curve in the site with early weaning, Dar es Saalam. This is because across all sites, 23% of infants had weaned by the end of the study, prior to testing positive, while 76% of infants in Dar had weaned by the end of the study, prior to testing positive. Thus, if the main interest is in understanding covariate effects on transition intensities, estimates may be robust to ignoring a competing risk if the overall competing risk is low. However, further work is needed to determine if this robustness holds when the covariate effects are larger (i.e. further from the null) or what is considered ignorable as a competing risk.

Our model improves upon methods used in MTCT studies. In the applied literature, early methods that have accounted for weaning as a competing risk generally relied on nonparametric methods that do not adjust for misclassification (see Leroy et al. [2002] and Leroy et al. [2003]). In our method, we adjust for time-dependent sensitivity, allowing us to estimate the risk of postnatal transmission using the full observed data set. Additionally, by specifying our baseline intensities as B-spline functions, we estimate a smoothed function of event risk, which may be of interest in epidemiological applications.

Finally, we conclude with one limitation of our current model. In implementation, we

have specified a duration model for the risk of mortality from the weaning state. However, a Markov model is more plausible, because the risk of death is likely to change once the infant is weaned and no longer protected by maternal antibodies. Thus, in future work, we will specify the weaning mortality intensity as a function of time since the start of the process. Though with only 15 events, the change in specification is unlikely to change the covariate coefficient estimates.

3.6 Appendices

3.A MCMC sampling algorithm

We sample from the set of full conditionals as follows.

1. To sample from $p(\mathbf{b}|\boldsymbol{\beta}, \boldsymbol{\sigma}^2, \mathbf{t}_{12}, \boldsymbol{\eta}, \text{data})$, for each transition hj , we use a blocked Metropolis-Hastings algorithm with a conditional prior proposal [Knorr-Held, 1999, Fahrmeir and Lang, 2001]. From Equation 3.4, the prior for the full set of \mathbf{b}_{hj} is an improper multivariate Gaussian distribution. Suppose we are interested in updating the block $\mathbf{b}_{r:s}^{hj} = b_r^{(hj)}, \dots, b_s^{(hj)}$, with $\mathbf{b}_{-r:s}^{(hj)}$ denoting the remaining parameters. Then the distribution for the block $\mathbf{b}_{r:s}^{hj} = b_r^{(hj)}, \dots, b_s^{(hj)} | \mathbf{b}_{-r:s}^{(hj)}$ is a proper multivariate Gaussian for block sizes greater than or equal to the degree of the spline basis. Applying standard properties for a conditional multivariate Gaussian distribution with p covariates, the mean is

$$\begin{aligned} \boldsymbol{\mu}_{r:s} &= -K_{r:s}^{-1} K_{1:r-1} \mathbf{b}_{1:r-1}^{(hj)} && \text{if } s = p \\ &= -K_{r:s}^{-1} K_{s+1:p} \mathbf{b}_{s+1:p}^{(hj)} && \text{if } r = 1 \\ &= -K_{r:s}^{-1} (K_{1:r-1} \mathbf{b}_{hj,1:r-1}^{(hj)} + K_{s+1:p} \mathbf{b}_{s+1:p}^{(hj)}) && \text{otherwise,} \end{aligned}$$

and the variance matrix is $\Sigma_{r,s} = \sigma_{hj}^2 K_{r,s}^{-1}$. Notationally, the sub-matrices of K are given by

$$K = \begin{pmatrix} & & K'_{1:r-1} & & \\ & K_{1,r-1} & K_{rs} & K_{s+1,p} & \\ & & K'_{s+1,p} & & \end{pmatrix}.$$

Given the current full set of coefficients $\mathbf{b}^{(hj)}$, we draw a proposal block of parameters $\mathbf{b}_{r:s}^{*(hj)}$ from this conditional Gaussian. Since the proposed samples are drawn from the conditional prior, the acceptance probability reduces to the likelihood ratio

$$\min \left\{ 1, \frac{\prod_{i=1}^N p(\mathbf{t} | \mathbf{b}_{r:s}^{*(hj)}, \mathbf{b}_{-r:s}^{(hj)})}{\prod_{i=1}^N p(\mathbf{t} | \mathbf{b}^{(hj)})} \right\}.$$

If the block size remains constant, there may be mixing problems at the boundaries r and s . To avoid this, we follow Fahrmeir and Lang [2001] and at each iteration, randomly sample a block size g from a range of reasonable block sizes. Then, in the current iteration, all block sizes are size g , with the exception of the last block, which can be of size less than g . In general, acceptance rates between 20% and 80% are recommended, and the range of block sizes can be adjusted to result in an appropriate acceptance rate.

2. To sample from $p(\boldsymbol{\sigma}^2 | \boldsymbol{\beta}, \mathbf{b}, \mathbf{t}_{12}, \boldsymbol{\eta}, \text{data})$, for each transition hj , we specified a conjugate inverse-gamma prior. Thus, we sample directly from the inverse-gamma posterior with shape $\text{rank}(K)/2 + a$ and scale $\frac{1}{2} \mathbf{b}'^{(hj)} \mathbf{K} \mathbf{b}^{(hj)} + c$. In practice, the prior is often weakly informative, say $a = 1$ and $c = 10^{-4}$ [Çetinyürek and Lambert, 2011].
3. For each participant i , we sample $t_{12,i}$ from $p(t_{12} | \boldsymbol{\beta}, \mathbf{b}, \boldsymbol{\sigma}^2, \boldsymbol{\eta}, \text{data})$. The support for each $t_{12,i}$ has points of discontinuity at 0, the last negative test for individual i , $t_{l,i}$, and the minimum of the first positive test, $t_{r,i}$, weaning time $t_{\omega,i}$, or end of follow-up $t_{\max,i}$. In our case, $t_{\max,i}$ is time of death or the last negative test if death was not

observed. If a positive test is observed, $t_{r,i}$ is also the upper bound of the support for $t_{12,i}$. Because of the discontinuities, we introduce a random variable $R_i = r_i$ to indicate which range of values contains the underlying $t_{12,i}$. For example, for a participant that tests positive at $t_{r,i}$, we say $r_i \in 1, 2, 3$, which corresponds to $t_{12,i} \leq 0$, $t_{12,i} \in (0, t_{l,i})$, and $t_{12,i} \in (t_{l,i}, t_{r,i})$, respectively. The first range is a point mass at $t_{12,i} \leq 0$ and the last range is a point mass at $t_{\omega,i}$ or $t_{\max,i}$, if applicable. For the remaining two ranges, we implement a slice sampler to sample a true $t_{12,i}$ with the constraints of the sampled range [Neal, 2003].

With only two possible points of discontinuity, it is computationally feasible calculate the normalizing constant for the posterior distribution for t_{12} . Then r_i is easily sampled from the resulting multinomial distribution.

Let $\boldsymbol{\theta} = \{\boldsymbol{\beta}^{(12)}, \mathbf{b}^{(12)}, \boldsymbol{\beta}^{(13)}, \mathbf{b}^{(13)}, \boldsymbol{\beta}^{(14)}, \mathbf{b}^{(14)}, \boldsymbol{\beta}^{(24)}, \mathbf{b}^{(24)}, \boldsymbol{\eta}\}$ and $\text{data}_i = \{t_{l,i}, t_{r,i}, \dots\}$. Assuming a non-informative prior $p(t_{12}) \propto 1$, we find from Equation 3.2 the posterior for an individual $t_{12,i}$ is

$$\begin{aligned} p(t_{12,i} | \boldsymbol{\theta}, \text{data}_i, \mu_d) &\propto \int_{t_{l,i} - t_{12,i}}^{t_{r,i} - t_{12,i}} g(d_i | t_{12,i}, \mu_d) dd \\ &\times \exp \{-(A_{12}(t_{12,i}) + A_{13}(t_{12,i}) + A_{14}(t_{12,i}))\} \alpha_{12}(t_{12,i}) \\ &\times \exp \{-A_{24}(t_{24,i} - t_{12,i})\} \end{aligned}$$

We use Gaussian quadrature with 10 nodes to numerically integrate over the range of t_{12} for the normalizing constant, leading to the following three cases. For convenience, because the originating state for transition into death is clear from the context, we use T to be time of death for all transitions and δ to be an indicator that death was observed. Additionally, we drop the notation i and describe the sampling procedure for a single individual.

A. If the participant tests positive at time t_{k+1} and exposure to illness ends at t_ω ,

$$\begin{aligned}
& p(t_{12}|\boldsymbol{\theta}, \text{data}, \mu_d) \\
& \propto \eta (G_W(t_{k+1} + \tau) - G_W(t_k + \tau)) \alpha_{24}(t)^\delta \exp \{-A_{24}(t)\} \quad \text{for } t_{12} \leq 0 \\
& \propto (1 - \eta) \int_0^{\min(t_k, t_\omega)} (G_W(t_{k+1} - s) - G_W(t_k - s)) \alpha_{12}(s) \\
& \quad \times \exp \left\{ - \sum_{j=2,3,4} A_{1j}(s) \right\} \alpha_{24}(t - s)^\delta \exp \{-A_{24}(t - s)\} ds \\
& \propto 1(t_\omega > t_k) (1 - \eta) \int_{t_k}^{\min(t_{k+1}, t_\omega)} G_W(t_{k+1} - s) \alpha_{12}(s) \\
& \quad \times \exp \left\{ - \sum_{j=2,3,4} A_{1j}(s) \right\} \alpha_{24}(t - s)^\delta \exp \{-A_{24}(t - s)\} ds
\end{aligned}$$

B. If the participant is weaned prior to the end of the study and tests negative throughout, there are two possible pathways are the illness pathway and the

weaning pathway. Then

$$\begin{aligned}
& p(t_{12}|\boldsymbol{\theta}, \text{data}, \mu_d) \\
& \propto \eta (1 - G_W(t_k + \tau)) \alpha_{24}(t)^\delta \exp \{-A_{24}(t)\} \quad \text{for } t_{12} \leq 0 \\
& \propto (1 - \eta) \int_0^{\min(t_\omega, t_k)} (1 - G_W(t_k - s)) \alpha_{12}(s) \exp \left\{ - \sum_{j=2,3,4} A_{1j}(s) \right\} \\
& \quad \times \alpha_{24}(t - s)^\delta \exp \{-A_{24}(t - s)\} ds \\
& \propto 1(t_\omega > t_k) (1 - \eta) \int_{t_k}^{t_\omega} \alpha_{12}(s) \exp \left\{ - \sum_{j=2,3,4} A_{1j}(s) \right\} \\
& \quad \times \alpha_{24}(t - s)^\delta \exp \{-A_{24}(t - s)\} ds \\
& \propto (1 - \eta) \alpha_{13}(t_\omega) \exp \left\{ - \sum_{j=2,3,4} A_{1j}(t_\omega) \right\} \\
& \quad \times \alpha_{34}(t - t_\omega)^\delta \exp \{-A_{34}(t - t_\omega)\} \quad \text{for } t_{12} > \min(t_\omega, t_k).
\end{aligned}$$

- C.** The participant breastfeeds throughout the study and does not test positive prior to the the last negative visit. Here, we assume that when a participant is negative at t_k and dies at t , no weaning has occurred in that range. Then the two possible

pathways are the illness pathway and the healthy-to-death pathway. Then

$$\begin{aligned}
& p(t_{12}|\boldsymbol{\theta}, \text{data}, \mu_d) \\
& \propto \eta (1 - G_W(t_k + \tau)) \alpha_{24}(t)^\delta \exp\{-A_{24}(t)\} \quad \text{for } t_{12} \leq 0 \\
& \propto (1 - \eta) \int_0^{t_k} (1 - G_W(t_k - s)) \alpha_{12}(s) \exp\left\{-\sum_{j=2,3,4} A_{1j}(s)\right\} \\
& \quad \times \alpha_{24}(t - s)^\delta \exp\{-A_{24}(t - s)\} ds \\
& \propto (1 - \eta) \int_{t_k}^t \alpha_{12}(s) \exp\left\{-\sum_{j=2,3,4} A_{1j}(s)\right\} \\
& \quad \times \alpha_{24}(t - s)^\delta \exp\{-A_{24}(t - s)\} ds \\
& \propto (1 - \eta) \alpha_{14}(t)^\delta \exp\left\{-\sum_{j=2,3,4} A_{1j}(t)\right\}.
\end{aligned}$$

4. For $p(\boldsymbol{\eta}|\boldsymbol{\beta}, \mathbf{b}, \boldsymbol{\sigma}^2, \mathbf{t}_{12}, \text{data})$, we estimate each $\eta_h = \frac{\sum_i 1(X_i(0)=h)}{N}$ for each entering state h .

5. To sample from $p(\boldsymbol{\beta}|\mathbf{b}, \boldsymbol{\sigma}^2, \mathbf{t}_{12}, \boldsymbol{\eta}, \text{data})$, for each transition hj , we implement an automatically-tuned univariate random-walk Metropolis-Hastings for each individual parameter [Çetinyürek and Lambert, 2011]. After updating the $(k - 1)^{\text{th}}$ component of the the m^{th} iteration of the chain, the state of the p parameters is $\boldsymbol{\beta}_{hj}^{(m)} = \{\beta_{hj,1}^{(m)}, \dots, \beta_{hj,k-1}^{(m)}, \beta_{hj,k}^{(m-1)}, \dots, \beta_{hj,p}^{(m-1)}\}$. Draw a proposal $\beta_{hj,k}^*$ for the k^{th} term from $N(\beta_{hj,k}^{(m-1)}, \lambda_k^2)$, and set

$$\boldsymbol{\beta}_{hj}^* = \{\beta_{hj,1}^{(m)}, \dots, \beta_{hj,k-1}^{(m)}, \beta_{hj,k}^*, \dots, \beta_{hj,p}^{(m-1)}\}.$$

Accept the proposal with probability

$$\psi(\boldsymbol{\beta}_{hj,k}^*, \boldsymbol{\beta}_{hj,k}^{(m+1)}) = \min\left\{1, \frac{\prod_{i=1}^N p(\mathbf{t}|\boldsymbol{\beta}_{hj,k}^*)}{\prod_{i=1}^N p(\mathbf{t}|\boldsymbol{\beta}_{hj,k}^{(m-1)})}\right\}.$$

The λ_k^2 is a tuning parameter that is adjusted during the burn-in period to achieve an acceptance probability of approximately 0.44. At each iteration, set

$$\lambda_k^{(n+1)} = h \left(\lambda_k^{(m)} + \gamma_m \left(\psi(\beta_{hj}^*, \beta_{hj}^{(m)}) - 0.44 \right) \right),$$

with

$$h(x) = \begin{cases} \epsilon & \text{if } x < \epsilon \\ x & \text{if } \epsilon < x < M \\ M & \text{if } x > M, \end{cases}$$

for small ϵ (say 10^{-6}) and large M (say 10^5). The γ_{m-1} represent a series such that $|\gamma_m - \gamma_{m-1}| \leq m^{-1}$. Our goal acceptance rate, 0.44, was shown optimal by Roberts and Rosenthal [2001] and further discussed in Roberts and Rosenthal [2009]. Though this rate was derived as an approximation under certain assumptions, we found our chains to have good properties when we used this rate in practice.

3.B Generating the simulated data

To generate the data sets, we simulate each of the latent event times and censor as appropriate. When the hazard curve is spline-based, we solve for the onset time from a randomly generated probability of survival.

1. Simulate visit schedule. To mimic the visit schedule of our motivating data set, we generated an empirical frequency distribution of visit times and randomly draw each visit schedule according to this observed frequency.
2. Simulate the entering state is illness with probability η . When all observations are assumed to enter the process as healthy, $\eta = 0$.

3. For all observations that enter the process as disease negative, independently generate time of illness t_{12} , time of death from healthy t_{13} , and time of weaning t_{14} , using transition intensities $\alpha_{12}(t)$, $\alpha_{13}(t)$, and $\alpha_{14}(t)$, respectively. When the transition intensities are spline based, times are randomly generated through the survival function. For example,
 - a. Set the range of observable times $(0, t_\kappa)$ where t_κ is the upper boundary knot (or maximum observable time) of the spline-based hazard.
 - b. Determine the probability of survival at t_κ : $p = S_{13}(t_\kappa)$
 - c. Generate $u \sim U(0, 1)$.
 - d. If $u > p$, set $t_{13} = A_{13}^{-1}[-\log u]$ for time of death t_{13} . If $u < p$, censor at t_m , which is the last observed visit time for each participant.
4. If $t_{12} < t_{13}$, $t_{12} < t_{14}$, and $t_{12} < t_m$, generate the time of death while in the illness state. Following steps **a** - **c** above, set $t_{23} = A_{23}^{-1} [A_{23}(t_{12}) - \log u]$ if a Markov model is assumed, and set $t_{23} = A_{23}^{-1}(-\log u) + t_{12}$ if a semi-Markov model is assumed. In the latter case, if $t_{23} > t_m$, censor t_{23} at t_m . The process is similar for generating time of death from the weaning state.
5. Generate the delay in detectability w for $t_{12} > 0$ from $\text{Exp}(\mu_w)$, and set the time illness is first able to be detected $v = t_{12} + w$. For an observation that enters as disease positive, sample a delay in detectability w from $\text{Shifted Exp}(\mu_w, \tau)$, with shift parameter τ .
6. Determine the observed values, time of last negative test, time of first positive test, and time of death or time censored. If there are no positive tests observed, set the time of the first positive test as $t_{k+1} = \infty$.

3.C Convergence diagnostics




Figure 3.9: Example trace plots for the log baseline rate of illness and coefficients for the eight covariates of interest for the illness transition. All chains appear to have mixed well. Because of early weaning at the Dar site, there were no positive tests after 6 months, and the lack of information led to a skewed posterior for rate of illness in Dar.




Figure 3.10: Posterior and prior distributions for the log baseline rate of illness (left) and the maternal CD4 count (right) coefficients. In both cases, the resulting posterior distributions (represented by black solid line) appear to be data driven, with little influence from the non-informative priors (represented by the gray dotted line).

Chapter 4

COMPARISON TO UNIVARIATE SURVIVAL MODELS

4.1 Introduction

In Chapters 2 and 3, we developed new methods that adjust for the imperfect sensitivity of HIV diagnostic testing in infants, account for weaning as a competing risk, incorporate death as an informative censoring event, and consider the interval censoring that results from a discrete testing schedule. In the applied literature, each of these challenges has often been treated with ad-hoc solutions to be able implement widely available and more easily understood methods of univariate survival models.

However, for two common main questions of interest, time to postnatal infection and time from infection to death, these ad-hoc approaches are not ideal for several reasons. First, because of low sensitivity of the diagnostic test immediately after birth, the univariate survival models are often restricted to infants who test negative at 4-6 weeks after birth. In addition to reducing the overall sample size, this approach precludes us from analysis of postnatal infections that occur immediately after birth. Additionally, the delay in detection of disease due to low sensitivity could result in the observation of event times that are considerably later than the true time of infection and possible bias in the estimation of covariate coefficients. Second, in practice, the timing of infection is often classified as known and occurring in utero/intrapartum if the infant tests positive at birth or occurring postnatally if the infant is negative at 4-6 weeks after birth and tests positive sometime later. However, because of the low sensitivity after birth, infants who test negative at birth but positive 4-6 weeks later are classified as unknown timing and treated as a separate category in mortality analyses [Newell et al., 2004, Marinda et al., 2007, Becquet et al., 2012]. Scientifically, it is of more interest to include those infants so that we can also draw

conclusions on the risk of mortality in infants who are infected via breastmilk in the early postnatal period. Finally, postnatal infection is informatively censored by death. An infant who dies is more likely to have been at a higher risk for infection, and because of imperfect sensitivity, it is also likely that an infant who dies without testing positive may be a missed diagnosis of HIV. To avoid bias from informative censoring, a combined outcome of time to disease-free survival is often modeled, which no longer directly addresses the true event of interest.

In this dissertation, we introduced new methods that aim to address the above possible sources of bias while focusing on time to postnatal infection. We based our approach on the multistate model (MSM) framework, which models the progression of disease as a series of transitions and defines a multivariate distribution where each outcome is the time of transition. In our application, we estimated the illness-death model (ID) (Figure 2.1) and the ID model extended to include weaning (Figure 3.1). In both cases, death is incorporated as an additional event to capture the information available in a dependent censoring event. Additionally, we included a new random variable in the MSM to measure the delay in detection that results from imperfect diagnostic testing. This accounts for the downward bias that is caused by using an observed event time that is later than the true onset and the uncertainty of an infant's true status prior to death.

While our method accounts for several sources of bias and better represents the true underlying disease process of mother-to-child transmission of HIV, it is more complex than the familiar univariate survival models. In this chapter, we compare our results from our application of the extended ID, which we refer to as MSM below (Section 3.4.4), and the simpler ID models (Section 3.4.6), implemented as a Bayesian MCMC algorithm, to univariate survival models for two outcomes of interest. In the ID model, we focus on the approach where the infant is censored at the last negative test if the infant does not test positive or die before the end of the study. This requires no information on weaning. The rest of the chapter proceeds as follows. In Section 4.2, we specify the univariate models and method

of estimation, and present the results in Section 4.3. Finally, we end with a discussion in Section 4.4.

4.2 *Methods*

In this section, we describe the two univariate models used for comparison. We define the samples and specify the approach used in estimation of the parameters.

In the first model, we compare a proportional hazards survival model of time to postnatal illness to the healthy to illness transition of the model estimated in Section 3.4.4. We fit four variations of the univariate survival model sample, restricting the denominator to infants who have a negative test at 0, 2, 4, or 6 weeks. As in the implementation in Section 3.4.4, we fit the model assuming an exponential hazard rate and include randomization, maternal CD4 count, log viral load, and weight at delivery, infant weight at birth, and site. Time of infection is defined as the midpoint between the time of the last negative and first positive tests.

In the second model, we estimate a proportional hazards survival model, with an unspecified baseline hazard and the outcome defined as the elapsed time between time of illness and time of death. This model estimates the risk of transitioning from illness to death. We first restrict the data to infants who test positive during the study period. Following Becquet et al. [2012], we classify infants who test positive as infected perinatally if the first positive occurs prior to a specified age we refer to as the conditional age, postnatally if the first positive occurs after a negative test at the conditional age or later, and unknown if the first positive test occurs prior to the conditional age. We apply this definition using conditional ages 0, 2, 4, and 6 weeks old. To better compare to our model, we treat both the perinatal and unknown timing groups as in utero/intrapartum infections and set the time of infection as 0. Thus, these infants are used only in estimating the time to death after infection model. As above, the time of postnatal infection is defined as the midpoint between the time of the last negative and first positive tests. The univariate model event is then the elapsed time between time of infection and time of death for all infants who test

positive.

4.3 Results

We report the number of infants meeting the inclusion criteria for the univariate models in Table 4.1. The parameter estimates are reported in Figure 4.1. In this study, most infants had their first postnatal visit by 6 weeks, and making 6 weeks of age an appropriate cut point for postnatal infections. In our approach, an assumed 3-week mean delay sensitivity curve results in 49% sensitivity at birth, 86% at 1 month, and 93% sensitivity at 6 weeks. These are in line with sensitivity estimates from a 2002 - 2006 study of French infants on postnatal prophylaxis antiretrovirals, which reported 55% - 58% sensitivity at birth and 89% sensitivity at 1 month [Burgard et al., 2012]. Thus, we focus on comparing the univariate model conditional on a negative test at 6 weeks to the MSM assuming a sensitivity curve with a mean of 3 weeks.

When comparing the time to postnatal infection parameters, point estimates and confidence intervals were similar for the risk of infection at the Lilongwe, Blantyre, and Lusaka sites and for the non-significant parameters of maternal weight at delivery and infant weight at birth. The estimates for two significant covariates, maternal CD4 counts and viral load at delivery, were attenuated toward the null in the univariate models, when compared to our sensitivity-adjusted MSMs. Additionally, the univariate confidence intervals were slightly narrower than those of our approach, most likely because the univariate approach does not incorporate the additional error of the observation process. Finally, we note that although the randomization covariate estimate is not significant, the estimated coefficient does switch direction when comparing the sensitivity-adjusted MSMs to the univariate model.

There was a difference in the point estimates and confidence intervals across models for the risk of infection at the Dar site, where most mothers had weaned their infants by 6 months of age. These differences are due in part to how the transition intensity of healthy to illness is interpreted when weaning is included as a state in the model (MSM) versus ignoring weaning completely (ID model). When weaning is included in the MSM,

the risk of transition to postnatal illness is conditional on the infant remaining in the state of healthy and breastfeeding (i.e. not transitioning to death or weaning), and the resulting estimate can be thought of as the estimate of one individual's risk of transitioning through the disease process. In the ID model, where weaning status is ignored, the infant remains in the denominator of the risk estimate until censored at the last visit or death, even though he/she is no longer at risk of disease. This estimates the overall population risk of disease, taking into account those infants who will never become infected. Intuitively, in a population where infants are weaned early, the proportion of all infants infected will be lower when compared to a population that breastfeeds for a longer period of time because of the differing periods of exposure. This idea of differing exposure periods is inherently captured in the ID model because we ignore weaning. Thus, we can think of the risk of illness in MSM and ID models as depending on different denominators, and this becomes evident when comparing the sensitivity-adjusted MSM to the sensitivity-adjusted ID model. In the MSM, as assumed mean delay increases, the width of the confidence interval increases as a reflection of the increasing uncertainty. However, this phenomenon does not carry over to the ID model, where any adjustment for misclassification results in wider confidence intervals. As expected, the ID point estimates were lower than the MSM point estimates. Comparing across the two univariate models, the point estimates were only slightly lower and the confidence intervals were similar. When compared to the univariate models, the point estimates were similar for both the MSM and ID models, while the univariate confidence intervals were much narrower. Overall, these differences demonstrate that in the case where the risk of disease is small compared to the risk of a competing event, as it is at the Dar site, it is possible that failing to account for decreased sensitivity in measuring the event of interest could lead to anti-conservative confidence intervals. More research is needed to explore how adjustment for misclassification impacts estimates under different degrees of competing risks in the MSM and ID model.

In the healthy to illness transition, we performed two sensitivity analyses, though neither

resulted in large differences from the results presented. First, in the univariate models, we also fit the time to illness proportional hazards model leaving the baseline hazard unspecified, which is more common than specifying the baseline hazard as exponential. Second, in the above analyses, we fit the univariate models conditional on slightly different samples. In the above analyses, we included infants who were observed to have a negative test on or after the conditional age, which excludes all infants with a positive test after the conditional age, but no observed negative test. Thus, as a sensitivity analysis, we included infants who had unknown timing of infection in that they test positive on or after the conditional age, but could not be confirmed as truly negative at birth because they did not have a negative test prior to the first positive test. As an example, when the conditional age is 6 weeks, our sample size increases by 50 by including this subset of infants.

In the infected to death transition, we estimated the univariate models conditional on observing a positive test after a negative test at 0, 2, 4, and 6 weeks (Figure 4.1). Because our approach adjusts for possible false negatives, the denominators could be slightly different for the MSM and ID models, and we present both results. Within model type (sensitivity-adjusted MSM/ID models and univariate model), the point estimates and confidence intervals were consistent. However, the coefficient estimates for the maternal (CD4, viral load, and weight) and infant covariates were farther away from the null in the univariate model when compared to the estimates from the sensitivity-adjusted models. Additionally, maternal viral load and weight and infant weight were significant at the 95% confidence level. It is not clear why these differences occur. It may be due to some interaction that the illness to death transition has with the other transition intensities in the multi-outcome models, the impact of adjusting for sensitivity in our model, or how interval censoring is handled in estimation. Finally, for the randomization covariate coefficient, none of the models resulted in significant estimates, but the estimates did change direction when comparing the MSM/ID models to the univariate model. This is similar to what we found in the healthy to illness model comparisons.

Table 4.1: Sample sizes of the univariate models. In the univariate time to postnatal infection model, infants are included in the sample conditional on having a negative test at 0-6 weeks (left column), and the final sample size is listed in the *Number negative* column. Infants who test positive are categorized according to Becquet et al. [2012]. In the univariate elapsed time from HIV infection to death model, the sample is restricted to all infants who test positive during the study, and the unknown category is grouped with peripartum infections to define time of infection.

Condition (Age)	Number negative	Timing of infection ($n = 388$)		
		Peripartum	Postnatal	Unknown
0 weeks	1752	133	255	0
3 weeks	1540	145	84	155
4 weeks	1540	148	84	152
5 weeks	1530	159	80	156
6 weeks	1474	221	67	159

4.4 Discussion

The models introduced in Chapter 3 were developed to adjust for time-dependent sensitivity in multistate models. Though more complicated to implement than the standard univariate survival models, our approach may provide new insights in certain instances. In particular, we have shown differences in model estimates in two instances: in the model where the competing risk is large compared to the event risk, as in the Dar site in the time to postnatal illness model, and in the model for time elapsed from onset of illness to death, where illness is subject to competing risks, misclassification, and interval censoring. In the latter case, it is important to note that our model did not result in significant estimates, while the simpler model did. Though it is unclear why the two approaches resulted in different conclusions, we would be concerned if this difference is caused by an incorrect assumption of the univariate model.




Figure 4.1: Point estimates and 95% HPDs/CIs from Bayesian estimation of the MSM and univariate models, respectively. The first four columns are results of the time to postnatal illness transition, and the remaining columns are results from the time from infection to death transition. Within models, from left to right, results are from the sensitivity-adjusted MSM that includes weaning, univariate model that censors at the first occurrence of weaning, death, or last visit, sensitivity-adjusted ID model, and univariate model that censors death or last visit only. Sensitivity-adjusted MSM/ID models are fit assuming no delay in detection and mean delays of 3, 4, and 6 weeks. Univariate models are fit conditional on observing a last negative test on or after 0, 2, 4, and 6 weeks of age. Estimates marked by a triangle are comparable across appropriate models.

BIBLIOGRAPHY

- A. Alioum, M. Cortina-Borja, F. Dabis, L. Dequae-Merchadou, G. Haverkamp, J. Hughes, J. Karon, V. Leroy, M.L. Newell, B.A. Richardson, L. van Weert, and G.J. Weverling. Estimating the efficacy of interventions to prevent mother-to-child transmission of human immunodeficiency virus in breastfeeding populations: comparing statistical methods. *American Journal of Epidemiology*, 158(6):596–605, 2003.
- P. Andersen. Multi-state models. *Statistical Methods in Medical Research*, 11(2):89–90, 2002.
- P. Andersen, O. Borgan, R. Gill, and N. Keiding. *Statistical models based on counting processes*. Springer-Verlag, 1993.
- P. Andersen, S. Abildstrom, and S. Rosthøj. Competing risks as a multi-state model. *Statistical Methods in Medical Research*, 11(2):203–215, 2002.
- R. Balasubramanian and S.W. Lagakos. Estimation of the timing of perinatal transmission of HIV. *Biometrics*, 57:1048–1058, 2001.
- R. Balasubramanian and S.W. Lagakos. Estimation of a failure time distribution based on imperfect diagnostic tests. *Biometrika*, 90(1):171–182, 2003.
- N. Bartolomeo, P. Terrotoli, and G. Serio. Progression of liver cirrhosis to HCC: an application of hidden Markov model. *BMC Medical Research Methodology*, 11(1):38, 2011.
- R. Becquet, R. Bland, V. Leroy, N.C. Rollins, D.K. Ekouevi, A. Coutoudis, F. Dabis, H.M. Coovadia, R. Salamon, and M.L. Newell. Duration, pattern of breastfeeding and postnatal transmission of HIV: pooled analysis of individual data from West and South African cohorts. *PLoS One*, 4(10):e7397, 2009.

- R. Becquet, M. Marston, F. Dabis, L.H. Moulton, G. Gray, H.M. Coovadia, M. Essex, D.K. Ekouevi, D. Jackson, A. Coutoudis, et al. Children who acquire HIV infection perinatally are at higher risk of early death than those acquiring infection through breastmilk: A meta-analysis. *PloS one*, 7(2):e28510, 2012.
- S.P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, 1998.
- E. Brown. Bayesian estimation of the time-varying sensitivity of a diagnostic test with application to mother-to-child transmission of HIV. *Biometrics*, 66(4):1266–1274, 2010.
- E.R. Brown and Y.Q. Chen. An imputation method for interval censored time-to-event with auxiliary information: analysis of the timing of mother-to-child transmission of HIV. *Statistical Communications in Infectious Diseases*, 2(1), 2010.
- A. Bureau, S. Shiboski, and J.P. Hughes. Applications of continuous time hidden Markov models to the study of misclassified disease outcomes. *Statistics in Medicine*, 22(3):441–462, 2003.
- M. Burgard, S. Blanche, C. Jasseron, P. Descamps, M.C. Allemon, N. Ciraru-Vigneron, C. Floch, B. Heller-Roussin, E. Lachassinne, F. Mazy, J. Warszawski, and C. Rouzioux. Performance of HIV-1 DNA or HIV-1 RNA tests for early diagnosis of perinatal HIV-1 infection during anti-retroviral prophylaxis. *The Journal of Pediatrics*, 160(1):60–66, 2012.
- A. Cachafeiro, G. Sherman, A. Sohn, C. Beck-Sague, and S. Fiscus. Diagnosis of HIV-1 in infants using dried blood spots and an ultrasensitive p24 antigen assay. *Journal of Clinical Microbiology*, 47(2):459–462, 2009.
- A. Çetinyürek and P. Lambert. Smooth estimation of survival functions and hazard ratios from interval-censored data using Bayesian penalized B-splines. *Statistics in Medicine*, 30(1):75–90, 2011.

- M. Charurat, P. Datong, B. Matawi, A. Ajene, W. Blattner, and A. Abimiku. Timing and determinants of mother-to-child transmission of HIV in Nigeria. *International Journal of Gynecology and Obstetrics*, 106:8–13, 2009. doi: 10.1016/j.ijgo.2009.02.017.
- D. Commenges. Inference for multi-state models from interval-censored data. *Statistical Methods in Medical Research*, 11(2):167–182, 2002.
- D. Commenges, P. Joly, A. Gégout-Petit, and B. Liqueur. Choice between semi-parametric estimators of Markov and non-Markov multi-state models from coarsened observations. *Scandinavian Journal of Statistics*, 34:33–52, 2007.
- H. M. Coovadia, N. C. Rollins, R. M. Bland, K. Little, A. Coutsooudis, M. L. Bennish, and M. L. Newell. Mother-to-child transmission of HIV-1 infection during exclusive breast-feeding in the first 6 months of life: an intervention cohort study. *The Lancet*, 369(9567):1107–1116, 2007.
- A. Coutsooudis, F. Dabis, W. Fawzi, P. Gaillard, G. Haverkamp, D.R. Harris, J.B. Jackson, V. Leroy, N. Meda, P. Msellati, et al. Late postnatal transmission of HIV-1 in breast-fed children: an individual patient data meta-analysis. *The Journal of Infectious Diseases*, 189(12):2154–2166, 2004.
- CR de Boor. *A Practical Guide to Splines*. Springer, 2001.
- D.T. Dunn, R.J. Simonds, M. Bulterys, L.A. Kalish, J. Moye Jr, A. Maria, et al. Interventions to prevent vertical transmission of HIV-1: effect on viral detection rate in early infant samples. *AIDS*, 14(10):1421, 2000.
- P.H. Eilers and B.D. Marx. Flexible smoothing with B-splines and penalties. *Statistical Science*, 11(2):89–121, 1996.
- L. Fahrmeir and S. Lang. Bayesian inference for generalized additive mixed models based on Markov random field priors. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 50(2):201–220, 2001.

- S.A. Fiscus, J. Wiener, E.J. Abrams, M. Bulterys, A. Cachafeiro, and R.A. Respass. Ultra-sensitive p24 antigen assay for the diagnosis of perinatal HIV-1 infection. *Journal of Clinical Microbiology*, pages 2274–2277, 2007.
- H. Frydman and M. Szarek. Estimation of overall survival in an “illness-death” model with application to the vertical transmission of HIV-1. *Statistics in Medicine*, 29:2045–2054, 2010.
- K. Gauthier, M. Kovacs, P. Wells, G. Le Gal, and M. Rodger. Family history of venous thromboembolism (VTE) as a predictor for recurrent VTE in unprovoked VTE patients. *Journal of Thrombosis and Haemostasis*, 11(1):200–203, 2013.
- A.E. Gelfand and A.F.M. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.
- C. Guihenneuc-Jouyaux, S. Richardson, and I. Longini. Modeling markers of disease progression by a hidden Markov process: application to characterizing CD4 cell decline. *Biometrics*, 56(3):733–741, 2004.
- N. Gupte, R. Brookmeyer, R. Bollinger, and G. Gray. Modeling maternal-infant HIV transmission in the presence of breastfeeding with an imperfect test. *Biometrics*, 63:1189–1197, 2007.
- P. Hadji, M. Frank, A. Jakob, and J.W. Siebers. Effect of adjuvant bisphosphonates on disease-free survival in early breast cancer: Retrospective analysis results in an unselected single-center cohort. *Journal of Bone Oncology*, 2013.
- N.J. Higham. Computing a nearest symmetric positive semidefinite matrix. *Linear algebra and its applications*, 103:103–118, 1988.
- J.P. Hughes and B.A. Richardson. Analysis of a randomized trial to prevent vertical transmission of HIV-1. *Journal of the American Statistical Association*, 95(452):1032–1043, 2000.

- C.H. Jackson and L.D. Sharples. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients. *Statistics in Medicine*, 21(1):113–128, 2001.
- C.H. Jackson, L.D. Sharples, S.G. Thompson, S.W. Duffy, and E. Couto. Multistate Markov models for disease progression with classification error. *The Statistician*, 5(2):193–209, 2003.
- P. Joly, D. Commenges, and L. Letenneur. A penalized likelihood approach for arbitrarily censored and truncated data: application to age-specific incidence of dementia. *Biometrics*, 54:185–194, 1998.
- P. Joly, D. Commenges, C. Helmer, and L. Letenneur. A penalized likelihood approach for an illness-death model with interval-censored data: application to age-specific incidence of dementia. *Biostatistics*, 3(3):433–443, 2002.
- T. Kneib and A. Hennerfeind. Bayesian semi-parametric multi-state models. *Statistical Modelling*, 8(2):169–198, 2008.
- L. Knorr-Held. Conditional prior proposals in dynamic models. *Scandinavian Journal of Statistics*, 26(1):129–144, 1999.
- M. Konishi, G. Haraguchi, S. Kimura, H. Inagaki, M. Kawabata, H. Hachiya, K. Hirao, and M. Isobe. Comparative effects of carvedilol vs. bisoprolol for severe congestive heart failure. *Circulation journal: Official journal of the Japanese Circulation Society*, 74(6):1127, 2010.
- A.P. Kourtis, F.K. Lee, E.J. Abrams, D.J. Jamieson, and M. Bulterys. Mother-to-child transmission of HIV-1: timing and implications for prevention. *The Lancet*, 6(11):726–732, 2006.
- J.S. Lambert, D.R. Harris, E.R. Stiehm, J. Moye, M.G. Fowler, W.A. Meyer, et al. Performance characteristics of HIV-1 culture and HIV-1 DNA and RNA amplification assays

- for early diagnosis of perinatal HIV-1 infection. *Journal of Acquired Immune Deficiency Syndromes*, 34(5):512–519, 2003.
- S. Lang and A. Brezger. Bayesian P-splines. *Journal of computational and graphical statistics*, 13(1):183–212, 2004.
- M.T. Leite, L.M. Zangwill, R.N. Weinreb, H.L. Rao, L.M. Alencar, P.A. Sample, et al. Effect of disease severity on the performance of Cirrus spectral-domain OCT for glaucoma diagnosis. *Investigative Ophthalmology & Visual Science*, 51(8):4104–4109, 2010.
- V. Leroy, J.M. Karon, A. Alioum, E.R. Ekpini, N. Meda, A.E. Greenberg, P. Msellati, M. Hudgens, F. Dabis, and S.Z. Wiktor. Twenty-four month efficacy of a maternal short-course zidovudine regimen to prevent mother-to-child transmission of HIV-1 in West Africa. *AIDS*, 16(4):631, 2002.
- V. Leroy, J.M. Karon, A. Alioum, E.R. Ekpini, P. van de Perre, A.E. Greenberg, P. Msellati, M. Hudgens, F. Dabis, S.Z. Wiktor, et al. Postnatal transmission of HIV-1 after a maternal short-course zidovudine peripartum regimen in West Africa. *AIDS*, 17(10):1493–1501, 2003.
- W. Li, C.A. Meyer, and X.S. Liu. A hidden Markov model for analyzing ChIP-chip experiments on genome tiling arrays and its application to p53 binding sequences. *Bioinformatics*, 21(suppl 1):i274–i282, 2005.
- W.A. Link and M.J. Eaton. On thinning of chains in MCMC. *Methods in Ecology and Evolution*, 3(1):112–115, 2011.
- E. Marinda, J.H. Humphrey, P.J. Iliff, K. Mutasa, K.J. Nathoo, E.G. Piwoz, L.H. Moulton, P. Salama, B.J. Ward, et al. Child mortality according to maternal and infant HIV status in Zimbabwe. *The Pediatric Infectious Disease Journal*, 26(6):519, 2007.
- L. Meira-Machado, J. de Una-Alvarez, C. Cadarso-Suarez, and P. Andersen. Multi-state

- models for the analysis of time-to-event data. *Statistical Methods in Medical Research*, 18(2):195–222, 2009.
- K.V. Menon, A.R. Hakeem, and N.D. Heaton. Meta-analysis: recurrence and survival following the use of sirolimus in liver transplantation for hepatocellular carcinoma. *Alimentary Pharmacology & Therapeutics*, 37(4):411–419, 2013.
- P.G. Miotti, T.E.T. Taha, N.I. Kumwenda, R. Broadhead, L.A.R. Mtimavalye, L. Van der Hoeven, J.D. Chipangwi, G. Liomba, and R.J. Biggar. HIV transmission through breastfeeding. *JAMA: the Journal of the American Medical Association*, 282(8):744–749, 1999.
- R. Nduati, G. John, D. Mbori-Ngacha, B. Richardson, J. Overbaugh, A. Mwatha, J. Ndinya-Achola, J. Bwayo, F.E. Onyango, J. Hughes, et al. Effect of breastfeeding and formula feeding on transmission of hiv-1. *JAMA: the Journal of the American Medical Association*, 283(9):1167–1174, 2000.
- R.M. Neal. Slice sampling. *Annals of Statistics*, pages 705–741, 2003.
- M.L. Newell, H. Coovadia, M. Cortina-Borja, N. Rollins, P. Gaillard, F. Dabis, et al. Mortality of infected and uninfected infants born to HIV-infected mothers in Africa: a pooled analysis. *Lancet*, 364(9441):1236–1243, 2004.
- F. O’Sullivan. Fast computation of fully automated log-density and log-hazard estimators. *Journal on Scientific and Statistical Computing*, 9(2):363–379, 1988.
- J.E. Pittman, R.C. Johnson, P.W. Jones, and S.D. Davis. Variability of a closed, rebreathing setup for multiple breath wash-out testing in children. *Pediatric Pulmonology*, 2012. doi: 10.1002/ppul.22531.
- W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, et al. *Numerical recipes: the Art of Scientific Computing*. Cambridge University Press, 2007.
- H. Putter, M. Fiocco, and R.B. Geskus. Tutorial in biostatistics: competing risks and multi-state models. *Statistics in Medicine*, 26(11):2389–2430, 2007.

- G.O. Roberts and J.S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351–367, 2001.
- G.O. Roberts and J.S. Rosenthal. Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.
- R. Sutradhar and R.J. Cook. Analysis of interval-censored data from clustered multistate processes: application to joint damage in psoriatic arthritis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 57(5):553–566, 2008.
- T. Taha, E. Brown, I. Hoffman, W. Fawzi, J. Read, M. Sinkala, et al. A phase III clinical trial of antibiotics to reduce chorioamnionitis-related perinatal HIV-1 transmission. *AIDS*, 20(9):1313, 2006.
- A. C. Titman. Flexible Nonhomogeneous Markov Models for Panel Observed Data. *Biometrics*, 67(3):780–787, 9 2011.
- M.J. Wessman, Z. Theilgaard, and T.L. Katzenstein. Determination of HIV status of infants born to HIV-infected mothers: A review of the diagnostic methods with special focus on the applicability of p24 antigen testing in developing countries. *Scandinavian Journal of Infectious Diseases*, 44(3):1–7, 2012.
- B. Yu. A Bayesian MCMC approach to survival analysis with doubly-censored data. *Computational Statistics & Data Analysis*, 54(8):1921–1929, 2010.
- Q. Zhang, L. Wang, Y. Jiang, L. Fang, P. Pan, S. Gong, et al. Early infant human immunodeficiency virus type 1 detection suitable for resource-limited settings with multiple circulating subtypes by use of nested three-monoplex DNA PCR and dried blood spots. *Journal of Clinical Microbiology*, 46(2):721–726, 2008.

CURRICULUM VITAE

EDUCATION

Ph.D., Biostatistics, University of Washington, Seattle, WA, 2013.

M.S., Biostatistics, University of Washington, Seattle, WA, 2006.

M.S., Statistics, University of Connecticut, Storrs, CT, 2001.

A.B., Applied Mathematics, Bryn Mawr College, Bryn Mawr, PA, 1998.

PROFESSIONAL EXPERIENCE

Research Assistant, Fred Hutchinson Cancer Research Center, 2013.

Research Assistant, Harborview/University of Washington, 2010 - 2012.

Research Assistant, Group Health Research Institute, 2009 - present.

Research Assistant, Department of Biostatistics, University of Washington, 2006 - 2009.

Research Assistant, School of Dentistry, University of Washington, 2005 - 2006.

Research Assistant, Seattle Childrens Hospital, 2004 - 2005.

Health Information Analyst. Qualidigm, Middletown, CT, 2001 - 2003.

TEACHING EXPERIENCE

Teaching Assistant Experience:

Department of Biostatistics, University of Washington

Courses: Medical Biometry I, II, III (introductory course for public health graduate students - 4 quarters), Correlated Data (1 quarter)

Department of Statistics, University of Connecticut

Courses: Introduction to Statistics (undergraduate level - 4 semesters)

HONORS, AWARDS, SCHOLARSHIPS

NIH Training Grant in cardiovascular health

PUBLICATIONS

1. J Wakefield, S Haneuse, A Dobra, and **E Teeple**. Bayes computation for ecological inference. *Statistics in Medicine*, 30(12):1381–1396, 2011
2. C Lewis, **E Teeple**, A Robertson, and A Williams. Preventive dental care for young, Medicaid-insured children in Washington State. *Pediatrics*, 124(1):e120–e127, 2009
3. HG Grundman and **EA Teeple**. Iterated sums of fifth powers of digits. *The Rocky Mountain Journal of Mathematics*, 38(4), 2008
4. HG Grundman and **EA Teeple**. Sequences of consecutive happy numbers. *The Rocky Mountain Journal of Mathematics*, 37(6), 2007a
5. HG Grundman and **EA Teeple**. Sequences of generalized happy numbers with small bases. *J. Integer Sequences*, 10, 2007b
6. TP Meehan, JP Tate, ES Holmboe, **EA Teeple**, A Elwell, RR Meehan, MK Petrillo, and SJ Huot. A collaborative initiative to improve the care of elderly Medicare patients with hypertension. *American Journal of Medical Quality*, 19(3):103–111, 2004
7. HG Grundman and **EA Teeple**. Heights of happy numbers and cubic happy numbers. *Fibonacci Quarterly*, 41(4):301–306, 2003
8. HG Grundman and **EA Teeple**. Generalized happy numbers. *Fibonacci Quarterly*, 39(5):462–466, 2001

MANUSCRIPTS IN PREPARATION

1. **EA Teeple** and ER Brown. Adjusting for time-dependent sensitivity in an illness-death model. In preparation, 2013A
2. **EA Teeple** and ER Brown. Bayesian estimation of a multistate model adjusted for time-dependent sensitivity. In preparation., 2013B

3. **EA Teeple**, JACD Delaney, ER Brown, C Grunfeld, M Mugavero, JH Willig, GA Burkholder, SL Boswell, WC Mathews, DR Drozd, PW Hunt, B Rodriguez, M Saag, MM Kitahata, and HM Crane. Impact of initial antiretroviral therapy on blood pressure. In preparation., 2013C

CONFERENCES, TALKS, AND POSTERS

1. **EA Teeple** and ER Brown. Adjusting for time-dependent sensitivity in an illness-death model. Presented at a University of Washington Department of Biostatistics student seminar, 2012A
2. **EA Teeple** and ER Brown. Adjusting for time-dependent sensitivity in an illness-death model. Presented at annual conference of the Western North American Region of The International Biometric Society, 2012B
3. **EA Teeple** and ER Brown. Adjusting for misclassified outcomes in multi-state models. Presented at a University of Washington Department of Biostatistics student seminar, 2011
4. **EA Teeple**, JACD Delaney, ER Brown, C Grunfeld, M Mugavero, JH Willig, GA Burkholder, SL Boswell, WC Mathews, DR Drozd, PW Hunt, B Rodriguez, M Saag, MM Kitahata, and HM Crane. Impact of antiretroviral medications on blood pressure in antiretroviral nave patients initiating their first regimen in the CNICS cohort. Presented at the HIV Primary Care Poster Session at the annual meeting for the Infectious Diseases Society of America, 2011

PROFESSIONAL ACTIVITIES

Member of the American Statistical Association