Causal Inference in HIV Vaccine Trials: Comparing Outcomes in a Subset Chosen After Randomization

Bryan E. Shepherd

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

Doctor of Philosophy

University of Washington

2005

Program Authorized to Offer Degree: Public Health and Community Medicine – Biostatistics

University of Washington

Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Bryan E. Shepherd

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of Supervisory Committee:

Peter B. Gilbert

Reading Committee:

Peter B. Gilbert

Thomas Richardson

Thomas Lumley

Date: 7/28/05

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Bell and Howell Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature 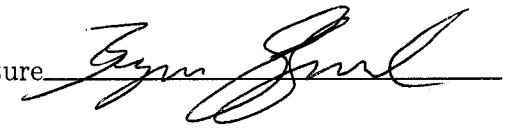

Date 7/28/05

University of Washington

Abstract

Causal Inference in HIV Vaccine Trials: Comparing Outcomes in a Subset Chosen After Randomization

by Bryan E. Shepherd

Chair of Supervisory Committee:

Professor Peter B. Gilbert
Biostatistics

In many experiments researchers would like to compare between treatments an outcome that only exists in a subset of participants selected after randomization. For example, in preventative HIV vaccine efficacy trials it is of interest to determine whether randomization to vaccine affects post-infection outcomes such as HIV viral load or the time from infection diagnosis to AIDS. This dissertation addresses some of the challenges of making causal comparisons conditioning on an event that occurs after randomization. Following the approach of Gilbert, Bosch, and Hudgens (2003) (GBH), I propose sensitivity analysis methods to estimate the average causal effect of treatment assignment on a post-infection outcome among those who would be infected whether randomized to vaccine or placebo. My key assumption is that subjects randomized to the vaccine arm who become infected would also have become infected if randomized to the placebo arm. It is not known which of those subjects infected in the placebo arm would have been infected if assigned vaccine, but this can be modeled using baseline covariates, the observed outcome variable, and a specified sensitivity parameter. In this dissertation I first construct a general likelihood. Using this likelihood, I then show that the method proposed by GBH yields a semiparametric maximum likelihood estimate of the average causal effect. Based on the likelihood, I extend GBH by including baseline covariates; allowing discrete, truncated, and censored

outcomes; and permitting more general selection bias functions. In particular, I study two different modeling approaches for estimating the average causal effect conditional on baseline covariates when the outcome of interest is continuous. In addition, I propose and study the behavior of semiparametric estimators of the causal effect of vaccination on an independently censored time to event outcome, deriving asymptotic properties and evaluating small sample behavior through simulations. I apply these methods to the first Phase III preventative HIV vaccine trial (VaxGen's trial of AIDSVAX B/B).

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank my committee for their help: Steve Self has provided good direction, Thomas Richardson has taught me much about causal inference, Donna Ankerst has given me several interesting ideas and I am excited to adjust these methods and apply them to some of her data, and Thomas Lumley has provided essential help with regards to asymptotic theory and computational complexities. I should also thank Michael Hudgens who gave me many excellent suggestions during the initial stages of this work. I would also like to thank Andrea Rotnitzky and Yannis Jemiai in the Department of Biostatistics at Harvard. Around Christmas of 2003, we discovered that we were working on similar problems and decided to collaborate, making this work much better. I should also thank VaxGen for allowing me to use their data. In particular, I would like to thank my advisor, Peter Gilbert. I have enjoyed working with Peter. He has always been helpful, patient, willing to meet, and insightful. Finally, I should write something cheezy here, acknowledging the love of my life, the blossom of my happiness, the chocolate chip of my cookies, my wife, Kristen. Oh yeah, and Sienna too.

Chapter 1

INTRODUCTION

Vaccines are being developed in an attempt to curb the spread of HIV. VaxGen recently completed the first Phase III preventive HIV vaccine trial. In this trial, over 5,000 HIV negative, high risk individuals were randomized to receive either the vaccine (AIDSVAX B/B) or placebo. Individuals were followed for three years and tested for HIV infection. With consent, those who became infected were enrolled in a post-infection phase of the study and were followed for an additional two years.

VaxGen's trial is the first of many HIV vaccine efficacy trials: other efficacy trials are in progress and many more are expected. Of primary interest in these trials is whether a vaccine protects against HIV infection; a perfect vaccine would eliminate infection. However, a less than perfect vaccine could also reduce morbidity and mortality by preventing infection for some and ameliorating disease progression or decreasing infectiousness among those who acquire HIV (Nabel, 2001; Graham, 2002; Gilbert et al., 2003b). In fact, the next wave of vaccines being tested are actually designed to induce a cellular response to destroy infected cells, not to prevent infection. (A short introduction to HIV and vaccines is found in Appendix A.)

To that end, investigators are very interested in the effect of candidate vaccines on post-infection outcomes. One such outcome is HIV viral load, a commonly used surrogate variable for measuring the extent of an infected individual's HIV disease and infectiousness (Mellors et al., 1996; O'Brien et al., 1996; Quinn et al., 2000). Viral load (typically $\log_{10}$-transformed) is a continuous measurement on the positive real line. (There are actually upper and lower quantification limits to assays measuring viral load, so in practice viral load is a continuous measurement on some positive interval. This distinction is addressed

in Chapter 3.) Other post-infection outcomes that may be of interest include CD4 count, genetic distance of the infecting HIV to the vaccine, time from HIV infection diagnosis to AIDS, time from HIV infection diagnosis to a composite endpoint (defined by viral load exceeding a certain value or onset of antiretroviral therapy), whether or not one's partner(s) become infected, and the number of infected partners over a fixed time, to name just a few possibilities. Some of these variables are continuous, others are discrete, some are time-to-event outcomes in the presence of censoring. Some of these outcomes exist whether or not someone is infected (e.g., CD4 count), others do not exist unless infected (e.g., genetic distance).

Two types of questions emerge from such investigations. Clinicians or epidemiologists are ultimately interested in whether they should recommend the vaccine to the general population (or a targeted subgroup therein). Scientists, however, may want to know whether there exists a mechanism through which the vaccine alters post-infection outcomes, perhaps leading to further scientific innovation and the elaboration of new vaccines.

For purpose of example, we will here assume that the outcome of interest is viral load. Naïvely comparing the distribution of viral loads between HIV infected individuals in the vaccine and placebo arms of a randomized experiment does not lead to unbiased estimates of the causal vaccine effect on viral load after infection because it improperly conditions on a post-randomization variable, HIV infection (Rosenbaum, 1984; Halloran and Struchiner, 1995). It addresses what may be an interesting question: among participants who become infected, do those in the vaccine arm have lower/higher viral loads than in the placebo arm? However, this is not a causal comparison and results of this analysis do not imply causation without making other assumptions. To illustrate this point, suppose half the study population have strong immune systems and half have weak immune systems. Suppose trial participants with a strong immune system are not infected if assigned vaccine but if assigned placebo they are infected and have a ($\log_{10}$) viral load value of 3. Suppose individuals with a weak immune system are infected regardless of randomization assignment and have a viral load of 5. (See Table 1.1.) One can see that the vaccine is partially effective to prevent infection, and has no effect on viral load. However, by comparing viral loads conditioning on infection status, the vaccine arm's mean viral load is 5, higher than the placebo arm

Table 1.1: Illustration of possible viral load values stratified by immune system.

	Immune System		
	Strong	Weak	
Vaccine	$--$	$Y = 5$	$\overline{Y} = 5$
Placebo	$Y = 3$	$Y = 5$	$\overline{Y} = 4$

mean of 4. If interpreted incorrectly, this suggests a detrimental effect of vaccine to increase viral load.

In the above example, there is an unmeasured variable, immune system, which is associated with viral load and with vaccine efficacy that is not allowing one to make a fair causal comparison of the effect of vaccine on viral load. This is related to the problem of confounding in observational studies. In an observational study, one often wants to see the effect of treatment on an outcome. However, because treatments are not randomly assigned, there could be unmeasured variables that are associated with both treatment and outcome that confound the causal effect of treatment on outcome. In our situation, we randomly assign treatment. However, by conditioning on infection, we are possibly biasing the estimate of the causal effect of vaccine on viral load because there may be other variables that are associated with vaccine efficacy and with viral load (see Figure 1.1). (This is more fully discussed in Section B.2 of the Appendix. Appendix B also contains a simple introduction to causal inference and discusses how this work relates to some of the existing causal literature.)

An alternative analytical strategy would assign a viral load value of 0 (or best rank) to all uninfected participants and perform an ITT analysis. This analysis addresses a causal question (i.e., does randomization to vaccine lower viral loads), but as verified in simulations, will lack power because most of the vaccine effect is washed out by the zeroes, which can occur in a large fraction of participants in an HIV vaccine trial (93% for the VaxGen trial). In addition, for other outcome variables there may not be a natural value to assign those who

Observational Study:

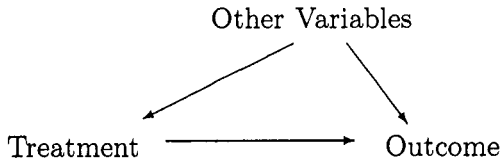

Vaccine Trial:

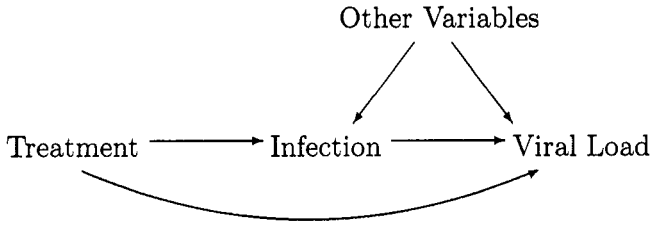

Figure 1.1: Causal diagrams of observational studies and vaccine trials

do not become infected. For example, suppose the outcome of interest is genetic distance, i.e. some measure of how much an HIV strain's genetic composition taken from an infected participant differs from the genetic blueprint of the vaccine. A genetic distance of 0 means that there is no difference between the two. There is no clear genetic distance value that one could assign to an individual who did not become infected. Naturally, one would want to assign a value far from zero, but how far? The choice would be arbitrary. In fact, the choice of viral load values of 0 for uninfected individuals is somewhat arbitrary. Since these measurements are in reality a $\log_{10}$ transformation, a $\log_{10}$ viral load value of 0 actually corresponds to someone being infected and having a very small viral load value (a value of 1). An uninfected individual truly has a viral load value of 0, which implies their $\log_{10}$ viral load value is actually $-\infty$. And in scientific importance, the difference between viral load values (untransformed) 0 and 1 is quite significant. (It should be noted that this may be less of a problem using rank-based ITT analyses, although the approach there will still have problems with low power.)

As pointed out by many authors (Kalbfleish and Prentice, 1980; Robins, 1995; Rubin, 2000; Robins and Greenland, 2000), a meaningful causal effect on viral load is defined in the subset of the population which would become infected under either placebo or vaccine. Specifically, each subject has a potential infection status if assigned vaccine and a potential infection status if assigned placebo, only one of which is observed during the trial. In addition, every subject who would be infected under a treatment (vaccine or placebo) also has a potential viral load under that treatment. Every subject can be classified into one of four possible combinations of the two potential infection status outcomes: never infected (not infected if assigned vaccine or placebo), harmed (infected if assigned vaccine but not infected if assigned placebo), protected (not infected if assigned vaccine but infected if assigned placebo), and always infected (infected regardless of assignment). This classification has been referred to as principal stratification by Frangakis and Rubin (2002). A comparison between treatment arms within a principal stratum is a valid causal comparison because principal stratum membership is unaffected by randomization and can be thought of as a comparison conditioning on a baseline covariate. Only in the always infected (ai) principal stratum do subjects have a potential viral load under both treatments. Therefore,

only in the *ai* principal stratum are causal comparisons meaningful. The type of questions addressed in this work are whether in individuals who would have been infected regardless of vaccine/placebo assignment, the vaccine alters viral load.

The vaccine trial cannot provide a definite answer to such questions because an individual's principal stratum membership can never be known and thus the distribution of viral load in the *ai* stratum under each treatment is not identified. One goal of this research is to derive a set of reasonable and easily understandable assumptions under which such distributions are identified by the clinical trial data. A second goal is to indicate how to conduct inference about the treatment effect in the *ai* stratum under such assumptions.

Hudgens, Hoering, and Self (2003) (HHS) and Gilbert, Bosch, and Hudgens (2003) (GBH) addressed the same question in the absence of baseline covariates. These authors assumed monotonicity which postulates that an individual who would get infected under vaccine would also get infected under placebo. Monotonicity appears to be a realistic assumption when comparing vaccine to placebo. These HIV vaccines are designed in such a way that the vaccine cannot mutate to become the virus. In addition, these trials are randomized, double blinded, and placebo controlled. Therefore, behavior should not change if one is randomized to the vaccine arm relative to the placebo arm. (It should be recognized that there are some hypothesized biological mechanisms which could possibly make monotonicity invalid. These issues are addressed at the end of Appendix A and in Section B.4.2 of Appendix B.)

Monotonicity implies that the vaccine effect on infection risk is either beneficial or harmless, and that all individuals infected in the vaccine group belong to the *ai* principal stratum. Thus, monotonicity and randomization are sufficient to identify the proportion of participants in the *ai* stratum (which can be estimated as simply the proportion of infected subjects randomized to vaccine) and the viral load distribution under vaccine in the *ai* stratum (which is simply the viral load distribution for individuals in the vaccine arm who become infected). These assumptions also allow us to compute the probability of being in the *ai* principal stratum given randomization to placebo and infection. This is equal to the probability of being infected if assigned vaccine divided by the probability of infection if assigned placebo. One minus this relative risk is a common measure of vaccine efficacy

to prevent infection (VE). This quantity can be estimated as one minus the proportion infected in the vaccine arm divided by the proportion infected in the placebo arm, or the observed proportions could be replaced by Kaplan-Meier estimates.

However, monotonicity is not sufficient to identify the viral load distribution under placebo in the ai stratum because infected individuals randomized to placebo can belong to either the ai or the protected principal strata. Therefore, the distribution of viral loads among participants who become infected in the placebo arm is a mixture of the viral load distributions in the ai and protected principal strata. As stated above, one can estimate the proportion of participants in the ai principal stratum, but one cannot identify this distribution because one does not know which of the subjects infected in the placebo arm would have been infected if randomized to vaccine. For example, consider the density of placebo viral loads shown in Figure 1.2. An estimable proportion of them are in the ai principal stratum (represented by the shaded regions of the density). However, even though one can consistently estimate the proportion in the ai principal stratum, from the observed data one cannot distinguish between any of the four plots.

Since this distribution is not identified, HHS proposed a method which estimated bounds on the estimate of the average causal effect (ACE) of vaccine on viral load in the ai principal stratum. Their method performs the analysis assuming that everyone infected after randomization to placebo with a viral load above the estimated VE^{th} quantile (in infected placebos) would have been infected if randomized to the vaccine. Visually, this corresponds to the lower-left plot in Figure 1.2. Under this assumption, the mean viral load for placebos in the ai stratum is computed and the ACE is estimated as this quantity subtracted from the mean viral load for participants infected in the vaccine arm. This represents the estimated lower bound of the ACE. The upper bound is estimated in a similar manner, assuming that everyone infected after randomization to placebo with a viral load below the estimated $(1 - VE)^{th}$ quantile would have been infected if randomized to the vaccine. In this manner, estimated upper and lower bounds for the ACE are computed.

Bounds are useful, but they represent extremes and do not allow the inclusion of additional scientific information to further refine estimation. The actual estimates that result from this type of analysis are biologically unlikely, and treat all values within the bounds

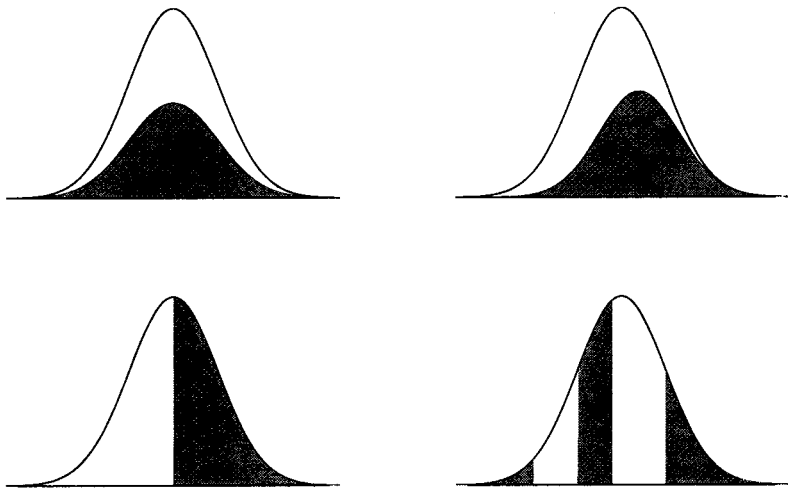

Figure 1.2: Density of placebo viral loads, with shaded area representing those in ai principal stratum. One can estimate the proportion in the ai stratum, but one cannot identify its location.

as equally probable. One may want to refine the analysis by performing it over a more restricted range. GBH advocated this approach.

In order to identify the distribution of viral loads for placebos in the ai principal stratum, GBH assumed that the probability that a subject who becomes infected under placebo is an always infected individual depends on his/her viral load y under placebo through the expit function $w(y) = e^{\alpha+\beta y}/\left(1+e^{\alpha+\beta y}\right)$. GBH showed that the parameter β is not identified by the clinical trial data but once it is specified, α is identified. This implies that one can identify the distribution of viral load under placebo in the ai stratum and therefore the ACE. Thus, similar to Scharfstein et al. (1999), GBH advocated carrying out a sensitivity analysis in which β is varied along a plausible range and inference about vaccine effects in the ai stratum is repeated for each value of β in the range. If results hold in one direction for a plausible range of the sensitivity parameter β, then a causal conclusion in that direction may be drawn. Otherwise, the analysis remains inconclusive.

An important extension of these methods would be to incorporate baseline covariates

into the analysis. This would allow one to investigate the conditional average causal effect of vaccination and possible interactions. These types of analyses may be of particular interest when studying the effect of vaccination. For example, in the VaxGen trial the vaccine did not appear to be effective at preventing infection from HIV in the main cohort. However, subgroup analyses were performed in which the vaccine appeared to be partially protective for non-whites and those involved in high risk behavior. These subgroup analyses were controversial and many are skeptical of their implications for the usual reasons (smaller samples, if one performs many subgroup analyses one is bound to find some false positive significants, etc.). If the vaccine is truly partially protective for these subpopulations, then it is quite possible that the same mechanism that is protecting individuals from infection may also lead to a beneficial post-infection immune response. Therefore, it is of particular interest to see if the vaccine is altering viral load and other post-infection outcomes in these subpopulations.

In addition, there are many diffferent types of post-infection outcomes. Ideally methods should be able to handle viral load measurements censored by assay detection limits, discrete outcomes, and time to event outcomes in the presence of censoring.

In this dissertation I extend the sensitivity analysis methods of HHS and GBH. In Chapter 2 I give notation and assumptions. My methods are likelihood based; in Chapter 3 I construct the likelihood and describe general estimation procedures. In Chapter 4 I show that in the special case where one does not include covariates and leaves the distribution of viral loads unspecified, the estimators of GBH are the semi-parametric maximum likelihood estimators. In Chapter 5 I study the method using baseline covariates in the continuous case under different parametric assumptions. In Chapter 6 I study methods where the outcome of interest is some time to event outcome in the presence of independent censoring. In Chapter 7 these methods are applied to the VaxGen trial, and Chapter 8 contains a discussion of this work and possible areas of future research. Most of the technical details and other information that did not fit in the main text of this work is found in the appendices.

Chapter 2

THE SET-UP

2.1 Notation

Consider a study in which N subjects, independently and randomly selected from a given population of interest, are randomized to either placebo or vaccine. Let $Z_i = 1$ if subject i, $i = 1, ..., N$, is randomized to vaccine and $Z_i = 0$ if randomized to placebo. Trial participants are monitored for HIV infection for a predetermined period of time. The recorded data on subject i are composed of a vector of baseline covariates X_i, an indicator S_i of infection during the study follow-up period ($S_i = 1$ if infected and $S_i = 0$ if not), and, if infected, the viral load Y_i (measured on a $\log_{10}$ scale) shortly after diagnosis of infection.

Although variables are defined in terms of a vaccine study, they could generally apply to the case where Z is some random treatment assignment, S some post-randomization variable on which one would like to condition, and Y some outcome of interest. I define variables in the HIV vaccine trial context because I believe this makes the explanation more clear and obviously this work is motivated by the vaccine trials. Applications of these methods in other settings are discussed in Chapter 8. Along this same train of thought, although Y is defined as viral load, Y could be defined as any discrete or continuous post-infection outcome.

In order to define the estimand of interest, we use the concept of potential outcomes – also called counterfactuals (Rubin 1978; Robins 1986). Specifically, define $S_i(0)$ to be the infection status indicator if, possibly contrary to fact, subject i is assigned placebo. Define $S_i(1)$ to be the infection status indicator if subject i is assigned vaccine. Similarly, define $Y_i(0)$ to be the viral load if participant i is assigned placebo and $Y_i(1)$ the viral load if assigned vaccine, where for a subject who does not become infected, i.e., $S_i(j) = 0$, we define $Y_i(j) = *$, $j = 0, 1$. The notation implicitly assumes that the potential outcomes

of each trial participant are not influenced by the treatments of other participants. This assumption is known as the Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1978, 1986). It implies that

$$Y_i(j) = Y_i \text{ if } Z_i = j,\ j = 0, 1. \tag{2.1}$$

That is, the potential outcome under the treatment actually assigned is equal to the observed outcome. Assuming the study participants make up a random sample from a large population of interest, then the potential outcomes $W_i = (S_i(0), S_i(1), Y_i(1), Y_i(0), Z_i, X_i)$, $i = 1, ..., N$, are independent and identically distributed (i.i.d.) copies of a random vector

$$W = (S(0), S(1), Y(1), Y(0), Z, X)$$

and similarly the observed data vectors $O_i = (Z_i, X_i, S_i, Y_i)$, $i = 1, ..., N$, are i.i.d. copies of a vector $O \equiv (Z, X, S, Y)$, where we define $Y = *$ if $S = 0$.

Randomization, possibly depending on the baseline covariates X, ensures that

$$(S(0), S(1), Y(1), Y(0)) \amalg Z|X \tag{2.2}$$

because $(S(0), S(1), Y(1), Y(0))$ can, like genetic make-up, be considered an unobserved baseline characteristic of each subject. Here, for random variables A, B and C, $A \amalg B \mid C$ indicates conditional independence of A and B given C (Dawid, 1979).

The four principal strata alluded to in the introduction can be defined in terms of the counterfactual pair $(S(0), S(1))$: the never infected are those with $S(0) = S(1) = 0$, the harmed are those with $S(0) = 0$ and $S(1) = 1$, the protected are those with $S(0) = 1$ and $S(1) = 0$, and the always infected are those with $S(0) = S(1) = 1$.

For a subject i who is in the ai principal stratum, a causal effect measure on his/her viral load is some measure of discrepancy between the values of $Y_i(0)$ and $Y_i(1)$, for example the difference $Y_i(1) - Y_i(0)$. The average causal effect at covariate level $X = x$ in the always infected principal stratum is defined as

$$ACE(x) = E[Y(1) - Y(0) | S(0) = S(1) = 1, X = x].$$

The average causal effect, not conditioning on covariates, is similarly defined:

$$ACE = E[Y(1) - Y(0) | S(0) = S(1) = 1].$$

2.2 *Key Assumptions*

Our goal is to propose methods for estimating the function $ACE(\cdot)$ based on the observed data $O_i, i = 1, ..., N$. Unfortunately, if no assumptions are made on the distribution of W, $ACE(x)$ is not identified by the observed data because randomization alone does not suffice to determine the ai principal stratum in each arm. Arguing as in GBH, it can be shown that the following assumptions, together with (2.1) and (2.2), do identify $ACE(x)$.

- A.1: Monotonicity: $S_i(1) \leq S_i(0)$ for all $i = 1, \cdot, N$.

- A.2: A model for the mixing probabilities of the always infected and protected in the placebo arm. Specifically,

$$P(S(1) = 1 | S(0) = 1, Y(0), X) = w(X, Y(0); \beta); \tag{2.3}$$

$$w(x, y; \beta) = \Phi\{m(x) + g(x, y; \beta)\},$$

  where β is fixed and known, $\Phi(\cdot)$ is a known cdf, $m(\cdot)$ is an unspecified function of X, and for each β, $g(\cdot, \cdot; \beta)$ is a known function of X and Y.

Monotonicity is tantamount to assuming that if a subject is to be infected under vaccine then he/she is also to be infected under placebo and hence that the infected subjects in the vaccine arm are all members of the ai principal stratum. The assumption implies that the vaccine does not increase infection risk. The assumption is not unrealistic in our scenario comparing vaccine to placebo, but may be called into question when comparing two vaccines.

The parameter β is not identified by the observed data. We propose regarding β (and therefore, the function $g(\cdot, \cdot; \beta)$) as fixed and known and, as a form of sensitivity analysis, estimating $ACE(x)$ under different values of β. The range for β should be chosen independently from the data. Choosing β such that $g(X, Y(0); \beta) = 0$ is the same as assuming that $Y(0) \amalg S(1) | S(0) = 1, X$ (here referred to as assumption A.3), or equivalently that the distribution of viral loads under placebo is the same in the always infected and in the protected principal strata. This is illustrated by the upper-left plot in Figure 1.2.

2.3 Modeling Assumptions

If X is discrete and can take a small number of values, we can then estimate $ACE(x)$ by applying the methods of GBH within each level x of the covariate vector X. However, if the distribution of X is continuous or discrete with a large number of values in its support this approach is unfeasible because the data are too sparse to conduct cell-specific estimation. I will address this problem by imposing the following additional distributional assumptions on the law of W:

- M.1: The probability of infection given covariates when assigned placebo is known up to a finite dimensional parameter μ; that is,

$$P(S(0) = 1|X) = \theta_p(X; \mu)$$

  where μ is an unknown parameter vector and for each μ, $\theta_p(\cdot; \mu)$ is a known function.

- M.2: The function $m(x)$ in the mixing model of A.2 follows a parametric model

$$m(X) = m(X; \alpha)$$

  where α is an unknown parameter vector and $m(\cdot; \cdot)$ is a known function of X and α.

- M.3: The conditional distribution of viral load under vaccine given covariates X and infection (hence, by A.1 being in the always infected stratum) is known up to a finite dimensional parameter η_1; that is,

$$f_{Y(1)|S(1)=1,X}(y|S(1) = 1, x) = f_v(y|x; \eta_1)$$

  where η_1 is an unknown parameter vector and for each η_1, $f_v(\cdot|x; \eta_1)$ is a known conditional density.

We additionally make one of the following two assumptions on the distribution of viral loads under placebo:

- M.4a: The conditional distribution of viral load under placebo given covariates X in the population comprising both protected and always infected individuals, is known up to a finite dimensional parameter η_0^a; that is,

$$f_{Y(0)|S(0)=1,X}\left(y|S\left(0\right)=1,x\right)=f_p(y|x;\eta_0^a)$$

  where η_0^a is an unknown parameter vector and for each η_0^a, $f_p(\cdot|x;\eta_0^a)$ is a known conditional density.

- M.4b: The conditional distribution of viral load under placebo given covariates X in the always infected principal stratum is known up to a finite dimensional parameter η_0^b; that is,

$$f_{Y(0)|S(0)=S(1)=1,X}\left(y|S\left(0\right)=S\left(1\right)=1,x\right)=f_p^{ai}(y|x;\eta_0^b)$$

  where η_0^b is an unknown parameter vector and for each η_0^b, $f_p^{ai}(\cdot|x;\eta_0^b)$ is a known conditional density.

Notice that M.4b models the distribution of viral loads in the placebo arm for those in the ai stratum, whereas M.4a models the distribution of viral loads for all infected individuals in the placebo arm (both those in the ai and protected principal strata). For ease of reference, we call the model defined by assumptions (2.1), (2.2), A.1, A.2, M.1-M.3 and M.4a, model $\mathcal{M}_a$. We call $\mathcal{M}_b$ the model defined like $\mathcal{M}_a$ except that M.4a is replaced by M.4b and in A.2 we demand that $w(x,y;\beta)$, now written as $w\left(x,y;\beta,\alpha\right)$ because of M.2, satisfy $w\left(x,y;\beta,\alpha\right)>0$ for all $\left(x,y,\beta,\alpha\right)$. There are advantages and disadvantages to choosing between models $\mathcal{M}_a$ and $\mathcal{M}_b$, which we discuss in Chapter 5.

Chapter 3

MAXIMUM LIKELIHOOD ESTIMATION

3.1 Constructing the Likelihood

Under model $\mathcal{M}_a$, $ACE(x)$ is a function of the unknown parameters $(\alpha, \eta_1, \eta_0^a)$. Specifically, $ACE(x) = ACE_a(x; \alpha, \eta_1, \eta_0^a)$ where

$$ACE_a(x; \alpha, \eta_1, \eta_0^a) \equiv \int y f_v(y|x; \eta_1) dy - \frac{\int y w(x, y; \beta, \alpha) f_p(y|x; \eta_0^a) dy}{\int w(x, y; \beta, \alpha) f_p(y|x; \eta_0^a) dy(0)}.$$

Similarly, under model $\mathcal{M}_b$, $ACE(x)$ is a function of (η_1, η_0^b) since it equals $ACE_b(x; \eta_1, \eta_0^b)$ $\equiv \int y f_v(y|x; \eta_1) dy - \int y f_p^{ai}(y|x; \eta_0^b) dy$. This is shown in Section C.1 of the Appendix. The maximum likelihood estimators of $ACE(x)$ under models $\mathcal{M}_a$ and $\mathcal{M}_b$ are therefore equal to the functions $ACE_a(x; \cdot, \cdot, \cdot)$ and $ACE_b(x; \cdot, \cdot)$ evaluated at the ML estimators of $(\alpha, \eta_1, \eta_0^a)$ and (η_1, η_0^b) under models $\mathcal{M}_a$ and $\mathcal{M}_b$ respectively.

To derive the maximum likelihood estimator under model $\mathcal{M}_a$ we express the joint density of the observables O

$$f_O(O) = f_X(X) P_{Z|X}(Z|X) P_{S|Z,X}(S|Z, X) f_{Y|S,Z,X}(Y|S, Z, X),$$

in terms of the parameters of the model (and we do the same for model $\mathcal{M}_b$). Specifically, in Section C.2 of the Appendix we show that

$$f_{Y|S,Z,X}(y|S = 1, Z = 0, X = x) = \begin{cases} f_p(y|x; \eta_0^a) \\ f_p^*(y|x; \alpha, \eta_0^b) \equiv \frac{w^{-1}(x,y;\beta,\alpha) f_p^{ai}(y|x;\eta_0^b)}{\int w^{-1}(x,y;\beta,\alpha) f_p^{ai}(y|x;\eta_0^b) dy} \end{cases}$$

under $\mathcal{M}_a$ and $\mathcal{M}_b$, respectively,

$$f_{Y|S,Z,X}(y|S = 1, Z = 1, X = x) = f_v(y|x; \eta_1) \qquad \text{under } \mathcal{M}_a \text{ or } \mathcal{M}_b,$$

$$P_{S|Z,X}(S = 1|Z = 0, X = x) = \theta_p(x; \mu) \qquad \text{under } \mathcal{M}_a \text{ or } \mathcal{M}_b,$$

and

$$P_{S|Z,X}\left(S=1|Z=1,X=x\right)=\begin{cases}\theta_p\left(x;\mu\right)\int w\left(x,y;\beta,\alpha\right)f_p(y|x;\eta_0^a)dy & \text{under } \mathcal{M}_a\\ \theta_p\left(x;\mu\right)\int w\left(x,y;\beta,\alpha\right)f_p^*(y|x;\alpha,\eta_0^b)dy & \text{under } \mathcal{M}_b.\end{cases}$$

It follows that the likelihoods $\mathcal{L}_a\left(\rho^a\right)$ and $\mathcal{L}_b\left(\rho^b\right)$ for $\rho^a \equiv (\mu,\alpha,\eta_0^a,\eta_1)$ and $\rho^b \equiv (\mu,\alpha,\eta_0^b,\eta_1)$ under models $\mathcal{M}_a$ and $\mathcal{M}_b$ respectively, with β known, are

$$\begin{aligned}\mathcal{L}_a\left(\rho^a\right) \propto \prod_{i=1}^{N}\Bigg\{&\left[f_v\left(y_i|x_i;\eta_1\right)\theta_p\left(x_i;\mu\right)\int w\left(x_i,y;\beta,\alpha\right)f_p(y|x_i;\eta_0^a)dy\right]^{S_i}\\ &\times\left[1-\theta_p\left(x_i;\mu\right)\int w\left(x_i,y;\beta,\alpha\right)f_p(y|x_i;\eta_0^a)dy\right]^{1-S_i}\Bigg\}^{Z_i}\\ &\times\left\{\left[\theta_p\left(x_i;\mu\right)f_p(y_i|x_i;\eta_0^a)\right]^{S_i}\left[1-\theta_p\left(x_i;\mu\right)\right]^{1-S_i}\right\}^{1-Z_i},\end{aligned}\tag{3.1}$$

and $\mathcal{L}_b\left(\rho^b\right)$ is defined like $\mathcal{L}_a\left(\rho^a\right)$ but with $f_p^*(y|x;\alpha,\eta_0^b)$ replacing $f_p(y|x;\eta_0^a)$.

3.2 Estimation

When $w(\cdot)$ depends on y, obtaining the MLEs of ρ^a and ρ^b requires maximizing over an integral that is not in closed form, which makes estimation more cumbersome, but still possible. In Chapter 5 I discuss numerical methods that I used to maximize the likelihood under the specific parameterizations used in my simulations and examples.

Provided the protected principal stratum is non-empty, then under sufficiently smooth parameterizations the ML estimators of the model parameters are asymptotically normally distributed. The variance of the normal limiting distribution can be consistently estimated with either the observed or the (estimated) expected information. These, in turn, can be used in conjunction with the delta method to obtain consistent variance estimators of $ACE(x)$ for each fixed x. (Details are given in Section C.3 of the Appendix.)

In our simulation studies, variance estimates based on the observed and expected information yielded nearly identical inferences. Alternatively, variances may be calculated by bootstrapping in a manner similar to that of GBH. Bootstrap procedures are more fully discussed in Chapter 6.

3.3 *Extensions of the Likelihood*

The likelihood given by (3.1) is obviously also the likelihood when Y is a discrete outcome, with the different distributional form of Y reflected in modeling assumptions M.3 and M.4. In the simple case where Y is binary and there are no covariates, then as shown in Section C.4 of the Appendix, this likelihood is equivalent to the likelihood presented by Hudgens and Halloran (2005).

In some datasets, including that from the VaxGen trial, the outcomes Y_i may be censored either above or below certain detection limits, say U and L. The likelihood in such cases becomes

$$\begin{aligned}
\mathcal{L}_a\left(\rho^a\right) \propto \prod_{i=1}^{N} \Bigg\{ \Bigg[& f_v\left(y_i|x_i;\eta_1\right)^{I(L\le y_i\le U)} \theta_p\left(x_i;\mu\right) \int w\left(x_i,y;\beta,\alpha\right) f_p(y|x_i;\eta_0^a)dy \\
& \times F_v\left(L|x_i;\eta_1\right)^{I(y_i<L)} \left(1-F_v\left(U|x_i;\eta_1\right)\right)^{I(y_i>U)} \Bigg]^{S_i} \\
& \times \left[1-\theta_p\left(x_i;\mu\right) \int w\left(x_i,y;\beta,\alpha\right) f_p(y|x_i;\eta_0^a)dy\right]^{1-S_i} \Bigg\}^{Z_i} \\
& \times \Bigg\{ \Big[\theta_p\left(x_i;\mu\right) f_p(y_i|x_i;\eta_0^a)^{I(L\le y_i\le U)} \\
& \times F_p\left(L|x_i;\eta_0^a\right)^{I(y_i<L)} \left(1-F_p\left(U|x_i;\eta_0^a\right)\right)^{I(y_i>U)} \Big]^{S_i} \\
& \times \left[1-\theta_p\left(x_i;\mu\right)\right]^{1-S_i} \Bigg\}^{1-Z_i}
\end{aligned} \tag{3.2}$$

where $F_z(L|x_i) = \int_{-\infty}^{L} f_z(y|x_i)dy$ for $z = p, v$ and $I(\cdot)$ is the indicator function. The likelihood $\mathcal{L}_b\left(\rho^b\right)$ under model $\mathcal{M}_b$ is modified accordingly.

Extending the likelihood to handle the case where Y is a time-to-event outcome is discussed in Chapter 6.

Chapter 4

CONTINUOUS OUTCOME WITHOUT COVARIATES

4.1 Estimators of GBH are Semi-Parametric MLEs

In the absence of baseline covariates, with $w(y;\beta,\alpha) = exp(\alpha+\beta y)/(1+exp(\alpha+\beta y))$, and with M.3 and M.4 left unspecified, models $\mathcal{M}_a$ and $\mathcal{M}_b$ are the same model and the ML estimator of $ACE \equiv E[Y(1)-Y(0)|S(0)=S(1)=1]$ coincides with the estimator of ACE derived by GBH. The proof follows.

Proof. Let $\theta_p(x) = \theta_p$, $w(x,y;\beta,\alpha) = w(y;\beta,\alpha)$, and $1-VE = \int w(y;\beta,\alpha) f_p(y)dy$. From (3.1), the likelihood is

$$\begin{aligned}\mathcal{L} \;\propto\; & \prod_{i=1}^{N} \left\{ \left[f_v(y_i)\,\theta_p\,(1-VE) \right]^{s_i} \left[1-\theta_p\,(1-VE) \right]^{1-s_i} \right\}^{z_i} \\ & \times \left\{ \left[\theta_p f_p(y_i) \right]^{s_i} (1-\theta_p)^{1-s_i} \right\}^{1-z_i}. \end{aligned} \tag{4.1}$$

As before, $N_v = \sum_{i=1}^{N} z_i$, $N_p = \sum_{i=1}^{N}(1-z_i)$, $n_v = \sum_{i=1}^{N} z_i s_i$, and $n_p = \sum_{i=1}^{N}(1-z_i)s_i$. For notational convenience, we will order the data so that $Z_i = 0$ for $i = 1,\cdots,N_p$, $S_i = 1$ for $i = 1,\cdots,n_p$, and $S_i = 1$ for $i = N_p+1,\cdots,N_p+n_v$. For $S_i = 1$, assume a dominating measure which is a counting measure on the discrete values $\{y_1,\cdots,y_{n_p},y_{N_p+1},\cdots,y_{N_p+n_v}\}$. In other words, probability mass will only be put on this set of points. Without loss of generality we will assume there are no ties (the proof is easily extended if there are ties). Define $p_{pi} \equiv P(Y=y_i|Z=0,S=1)$ and $p_{vi} \equiv P(Y=y_i|Z=1,S=1)$, where $\sum_{i=1}^{n_z} p_{zi} = 1$ for $z = p,v$. Since $f_v(\cdot)$ factors out of the likelihood, its estimation can be considered separately: The ML estimates for p_{vi} (and hence $F_v(\cdot)$) are the standard nonparametric ML estimators, $1/n_v$ (and $\sum_{i=1}^{N} z_i s_i I(Y_i \le y)/n_v$). The log-likelihood for

what remains can be written as

$$\begin{aligned} l_p &= \sum_{i=1}^{n_p} log\,(p_{pi}) + n_p log(\theta_p) + (N_p - n_p) log(1-\theta_p) + n_v log(\theta_p) \\ &\quad + n_v log\,(1-VE) + (N_v - n_v) log\,(1-\theta_p(1-VE))\,, \end{aligned} \tag{4.2}$$

where $1 - VE = \sum_{i=1}^{n_p} w(y_i;\beta,\alpha)p_{pi}$.

One way to maximize the log-likelihood is to use a Lagrange multiplier. Consider the problem of minimizing the function:

$$\begin{aligned} \phi_\lambda &= -\sum_{i=1}^{n_p} log\,(p_{pi}) - n_p log(\theta_p) - (N_p - n_p) log(1-\theta_p) - n_v log(\theta_p) \\ &\quad - n_v log\,(1-VE) - (N_v - n_v) log\,(1-\theta_p(1-VE)) + \lambda\left(\sum_{i=1}^{n_p} p_{pi} - 1\right). \end{aligned}$$

Setting the partial derivatives equal to zero yields the equations

$$\frac{\partial \phi_\lambda}{\partial p_{pi}} = -\frac{1}{p_{pi}} + \lambda - n_v \frac{w(y_i;\beta,\alpha)}{1-VE} + (N_v - n_v)\frac{\theta_p w(y_i;\beta,\alpha)}{1-\theta_p(1-VE)} = 0, \tag{4.3}$$

$$\frac{\partial \phi_\lambda}{\partial \alpha} = -\frac{n_v - N_v\theta_p(1-VE)}{(1-VE)(1-\theta_p(1-VE))}\sum_{j=1}^{n_p} w'(y_j;\beta,\alpha)p_j = 0, \tag{4.4}$$

$$\frac{\partial \phi_\lambda}{\partial \theta_p} = -\frac{n_p - N_p\theta_p}{\theta_p(1-\theta_p)} - \frac{n_v - N_v\theta_p(1-VE)}{\theta_p(1-\theta_p(1-VE))} = 0, \tag{4.5}$$

with the ML estimators being the solutions to these equations. Note that $\sum_{i=1}^{n_p} p_{pi} = 1$, which taken with equation (4.3) implies that

$$\sum_{i=1}^{n_p}\left[\lambda - w(y_i;\beta,\hat{\alpha})\frac{n_v - N_v\hat{\theta}_p(1-\widehat{VE})}{(1-\hat{\theta}_p(1-\widehat{VE}))(1-\widehat{VE})}\right]^{-1} = 1. \tag{4.6}$$

In order for (4.4) to hold, one of the following two conditions must hold:

$$\sum_{j=1}^{n_p} w'(y_j;\beta,\hat{\alpha})\hat{p}_{pj} = 0, \tag{4.7}$$

$$\frac{n_v}{N_v} = \hat{\theta}_p(1-\widehat{VE}). \tag{4.8}$$

Suppose (4.7) does not hold. Then from (4.5), (4.6), and (4.8) it is easily seen that $\hat{\theta}_p = n_p/N_p$, $1-\widehat{VE} = n_v N_p/(N_v n_p)$, and $\hat{p}_{pi} = 1/n_p$. (Solving first for $\hat{\theta}_p$ and $\widehat{VE}$, in order for (4.6) to hold, λ must equal n_p, which implies $\hat{p}_{pi} = 1/n_p$.)

Now consider the case where (4.7) holds. In the special case where $w(y_j;\beta,\alpha)$ is the expit function, then (4.7) is equivalent to

$$\sum_{j=1}^{n_p} \frac{exp(\hat{\alpha}+\beta y_j)}{(1+exp(\hat{\alpha}+\beta y_j))^2}\hat{p}_{pj} = 0,$$

which implies $\hat{\alpha}=\infty$. Therefore, $w(y;\beta,\hat{\alpha})=1$ for all y and $1-\widehat{VE}=1$. Plugging this into (4.5) and (4.6), one learns that $\hat{\theta}_p=(n_p+n_v)/(N_p+N_v)$ and $\hat{p}_{pi}=1/n_p$.

Notice that if $n_v/N_v > n_p/N_p$ (which implies that A3 is suspect), then the estimator $1-\widehat{VE}=n_vN_p/(n_pN_v)$ yields an estimate for $1-VE$ outside the interval in which it is constrained to belong, $[0,1]$. Therefore, when $n_v/N_v > n_p/N_p$ the ML estimators must be such that (4.7) holds, and the ML estimators are

$$\begin{aligned}
&\hat{\theta}_p = \frac{n_p+n_v}{N_p+N_v}\\
&\hat{p}_{vj} = \frac{1}{n_v} \qquad \text{for } j=1,\cdots,n_v\\
&\hat{p}_{pi} = \frac{1}{n_p} \qquad \text{for } i=1,\cdots,n_p\\
&\hat{\alpha}=\infty \Rightarrow \widehat{VE}=0.
\end{aligned} \tag{4.9}$$

On the other hand, if $n_v/N_v < n_p/N_p$, the ML estimators are

$$\begin{aligned}
&\hat{\theta}_p = \frac{n_p}{N_p}\\
&\hat{p}_{vj} = \frac{1}{n_v} \qquad \text{for } j=1,\cdots,n_v\\
&\hat{p}_{pi} = \frac{1}{n_p} \qquad \text{for } i=1,\cdots,n_p\\
&\hat{\alpha} \text{ such that } \sum_{i=1}^{n_p} \frac{exp(\hat{\alpha}+\beta y_i)}{1+exp(\hat{\alpha}+\beta y_i)}\hat{p}_{pi} = \frac{n_vN_p}{N_vn_p} = 1-\widehat{VE},
\end{aligned} \tag{4.10}$$

the estimators used in GBH. And if $n_v/N_v = n_p/N_p$, then the two estimators are equivalent.

□

Some technical notes: The estimators (4.9) are also solutions to (4.3)-(4.5) when $n_v/N_v < n_p/N_p$. However, they do not maximize the log-likelihood. (This is easily seen by plugging in values for n_v, N_v, n_p, N_p. However, I have not been able to mathematically prove this for all n_v, N_v, n_p, N_p with $n_v/N_v < n_p/N_p$.)

Notice that the score equations (4.3)-(4.5) and hence the conditions (4.6)-(4.8) are for any general $w(y_i; \beta, \alpha)$, as defined by A.2. Therefore, maximum likelihood estimators based on different $w(\cdot)$ can be constructed based on these score equations in a similar manner.

4.2 Other Models

As pointed out by GBH, setting $\beta = \pm\infty$ is equivalent to performing the analysis of HHS. This can also be written by defining $w(y)$ as $I_{\{y \geq q^{VE}\}}$ and $I_{\{y \leq q^{1-VE}\}}$ for $\beta = \infty$ and $-\infty$, respectively; where q^{VE} designates the VE^{th} quantile of $Y(0)$.

One could model the distribution of viral loads by specifying M.3 and either M.4a or M.4b. Parametric approaches are discussed in Chapter 5, in the context of including baseline covariates in the analyses.

Chapter 5

CONTINUOUS OUTCOME WITH COVARIATES

In Chapter 3, I constructed a general likelihood. The purpose of this chapter is to study estimation under specific parameterizations when baseline covariates are included in the models. In Chapter 3, there were two general choices for parameterization, $\mathcal{M}_a$ and $\mathcal{M}_b$. Here I study both, examining advantages and disadvantages of both approaches. Although I focus on a continuous outcome and will make specific distributional assumptions, the general principles discussed in this chapter apply to other types of outcomes and other distributional assumptions.

For the simulations and example described in the following sections, in both models $\mathcal{M}_a$ and $\mathcal{M}_b$ we have considered

$$\theta_p(x;\mu) = \frac{exp(x^T\mu)}{1+exp(x^T\mu)} \tag{5.1}$$

$$w(x,y;\beta,\alpha) = \frac{exp(x^T\alpha+\beta y)}{1+exp(x^T\alpha+\beta y)} \tag{5.2}$$

where $x = (1, x_1, \cdots, x_q)^T$ and μ and α are parameter vectors of length $q+1$.

5.1 Parameterization under $\mathcal{M}_a$

$\mathcal{M}_a$ is a natural approach to modeling the potential outcomes. One assigns parametric distributions to the observed viral loads and then through $w(x,y;\beta,\alpha)$, models which portion of the infected placebos would have been infected regardless of treatment. Therefore the two modeling assumptions, M.2 and M.4a, together induce the distribution of placebo viral loads in the always infected principal stratum.

Model $\mathcal{M}_a$, however, has drawbacks. First, estimands of interest cannot be easily expressed with model parameters. For example, natural functional forms for $ACE\,(x)$ when H_0 fails cannot be expressed as simple restrictions on the parameters of the model. Similarly,

there are no parameters that describe the interaction between treatment and covariates on viral load in the *ai* stratum.

A second drawback of model $\mathcal{M}_a$ involves model compatibility. Under this parameterization there are difficulties choosing distributional models such that all parameters are identifiable and the null hypothesis can hold. Under model $\mathcal{M}_a$, in order to ensure that the null hypothesis $H_0 : ACE(x) = 0$, for all x is not a-priori excluded, $f_v(y|x;\eta_1)$ in M.3 must have a particular functional form. Specifically,

$$f_v(y|x;\eta_1) = \frac{w(x,y;\beta_v,\alpha_v)\, f_p(y|x;\eta_v)}{\int w(x,y;\beta_v,\alpha_v)\, f_p(y|x;\eta_v)dy}, \tag{5.3}$$

where $\eta_1 = (\alpha_v, \beta_v, \eta_v)$ is an unknown parameter vector.

The choice (5.3) brings some identification problems. Specifically, at the value of β_v that makes $w(x,y;\beta_v,\alpha_v)$ a function of x only, the parameter α_v is not identified. For example, under (5.2) the parameter α_v is not identified at $\beta_v = 0$. Also, depending on the functional forms of $w(x,y;\cdot,\cdot)$ and $f_p(y|x;\cdot)$, there may be additional non-identification problems. For example if, as we assume in our example and simulations, in addition to (5.2) we assume

$$f_p(y|x;\eta_0^a) = \phi(y;x^T\lambda,\sigma^2) \tag{5.4}$$

where $\eta_0^a = (\lambda,\sigma), \lambda = (\lambda_0,\lambda_1,\cdots,\lambda_q)^T$ and $\phi(y;x^T\lambda,\sigma^2)$ is a normal density with mean $x^T\lambda$ and variance σ^2, then with $\eta_v = (\lambda_v,\sigma_v)$, all values of $\eta_1 = (\alpha_v,\beta_v,\lambda_v,\sigma_v)$ that satisfy $\alpha_v - (-\beta_v\lambda_{v0} - \beta_v^2\sigma_v^2/2, -\beta_v\lambda_{v1},\cdots,-\beta_v\lambda_{vq}) = 0$ give the same function $f_v(y|x;\eta_1)$. This is shown in section D.1 of the Appendix.

The lack of identifiability of η_1 may not appear to be a problem because it does not rule out the identifiability of $f_v(y|x;\eta_1)$ and hence of $E(Y(1)|S(1)=1,X=x)$ which is the relevant term needed to compute $ACE(x)$. Indeed, in our example $ACE(x)$ remains identified even when η_1 is not. However, the lack of identification of η_1 does complicate inference. Specifically, standard asymptotic theory does not apply for the asymptotic distribution of ML estimators of η_1. (See Appendix C.2.) Consequently, using the delta method in conjunction with the usual calculations for the asymptotic variance of the ML estimator of η_1 to compute estimates of the variance of the ML estimator of $ACE(x)$ yields Wald confidence intervals whose coverage probability, even with large samples, is not close to the nominal level.

Even though standard asymptotic calculations do not result in consistent variance estimators, we suspect that boostrap estimators of the variance of the estimate of $E(Y(1)|S(1) = 1, X = x)$ can be used to compute valid confidence intervals. Our suspicion follows from the conjecture that the ML estimators of $E(Y(1)|S(1) = 1, X = x)$ may be regular and asymptotically normal even when η_1 is not identified. Our simulations support this conjecture.

Since $f_v(y|x;\eta_1)$ factors out of the likelihood, one may perform the bootstrap by resampling only from those infected in the vaccine arm. The resulting bootstrap variance estimator for the estimate of $E(Y(1)|S(1) = 1, X = x)$ can be added to the variance estimator for the estimate of $E(Y(0)|S(0) = S(1) = 1, X = x)$ obtained from the information matrix using the delta method, to estimate the variance of the ML estimator of $ACE(x)$. There are, of course, additional complications to modeling M.3 with (5.3) which we will discuss more fully in the simulations section. Briefly, we have found that it is often difficult and computationally time consuming to find the ML estimators of $E(Y(1)|S(1) = 1, X = x)$ under (5.3) due to there being many local maxima. And because of the many free parameters (twice as many as there would be if we modeled $f_v(y|x;\eta_1)$ with a normal linear model), estimates of $E(Y(1)|S(1) = 1, X = x)$ tend to over-fit the data.

One can use other models to estimate $E(Y(1)|S(1) = 1, X = x)$. The simplest is to set β_v in (5.3) equal to 0, assuming $Y(1)|S(1) = 1, X$ and $Y(0)|S(0) = 1, X$ have the same distributional form. With M.4a defined by (5.4) this models viral loads in the vaccine arm with a normal linear model. This approach eliminates problems of over-fitting and non-identifiability. Its disadvantage, however, is that the distributions to be compared, $F_p^{ai}(y|x;\cdot)$ and $F_v(y|x;\cdot)$, are not of the same form when $\beta \neq 0$. This means that for any given sensitivity parameter $\beta \neq 0$, there do not exist parameter values η_1 and η_0^a such that $ACE(x) = 0$ for all x, therefore forcing a non-null causal vaccine effect for at least one subgroup. In practice this may not be important; any a-priori bias in the null may be much smaller than stochastic variation. Indeed, in our simulations we found that even though the data were generated from (5.3), estimates obtained fixing $\beta_v = 0$ performed well – essentially demonstrating the robustness of normal linear regression. However, in large samples, if $\beta \neq 0$ the inferential procedures will always find an $ACE(x) \neq 0$, irrespective of the data. (A similar phenomenon, when using standard parametric families for estimating

the effects of sequential treatments on an outcome, was reported by Robins and Wasserman, 1997.)

Another approach is to fix $\beta_v = \beta$. This reduces the number of parameters to estimate by one and eliminates the problem of non-identifiability. This also allows the null hypothesis to be true for every β and every x. Since the other parameters remain free the model is still fairly flexible. However, one complaint against this parameterization is that it implies that the distribution of $Y(1)|S(1) = 1, X$ changes for every value of β even though this distribution is identifiable without β based on the standard assumptions (A.1-A.2). There is some justification for allowing β_v to change: One type of sensitivity analysis might be to change the functional form of $w(\cdot)$, doing one analysis using a logistic function and a second using, for example, a probit function. For the null hypothesis to be possible, the functional form of the distribution of $Y(1)|S(1) = 1, X$ given by (5.3) would have to change as the form of $w(\cdot)$ changes. Our approach is analagous, we are considering the value of β to be an intrinsic part of the functional form of $w(\cdot)$, and changing the form of the distribution of $Y(1)|S(1) = 1, X$ each time we change β.

One could take this approach to an extreme, fixing $(\alpha_v, \beta_v) = (\alpha, \beta)$. Since $f_v(\cdot)$ factors out of the likelihooed, this could be done by estimating α, ignoring $f_v(\cdot)$, and then setting α_v of (5.3) equal to $\hat{\alpha}$. However, this approach is somewhat ad-hoc: estimating a parameter and then fixing another parameter equal to this estimate. The other option would be to maximize the likelihood under the constraint that $(\alpha_v, \beta_v) = (\alpha, \beta)$ where β is still considered a known constant. However, this would make obtaining MLEs more difficult because the likelihood would not factor as nicely. More importantly, it does not make sense intuitively to use vaccine viral loads to help estimate the odds of being in the ai stratum given infection in the placebo arm, or vice versa. Therefore, we did not study this final approach.

There are no identifiability/compatibility problems in the placebo arm. Estimating the expected conditional viral load among always infected placebos, $E(Y(0)|S(0) = S(1) = 1, X = x)$, and the variance of this estimate is a little complex computationally (discussed below) but possible. Using the observed or expected information, one can estimate the variance of parameter estimates and then use the delta method to estimate the variance of

the ML estimator for $E(Y(0)|S(0) = S(1) = 1, X = x)$. Details are shown in section D.2 of the Appendix.

5.2 *Computation under* $\mathcal{M}_a$

Under $\mathcal{M}_a$ using (5.1)-(5.4), MLEs were obtained using quasi-newton methods implemented in R using the function optim(·). When $\beta \neq 0$, obtaining MLEs requires maximizing over an integral that is not in closed form. Numerical integration programs written in C were used to obtain these integrals, which were then called into R. For some simulations, there were local maximums. Analyses were run using multiple initial values in order to obtain the MLE. Computationally, it is important to specify good initial parameter values. For instance, if β is relatively large, for many choices of α, $w(x, y; \beta, \alpha)$ is very close to 1 for all plausible values of x and y. Optimization routines may have difficulty finding MLEs because minor shifts in the value of α may result in $w(x, y; \beta, \alpha)$ still approximately equal to 1 for all x and y, yielding an identical likelihood. In practice we recommend starting the estimation at $\beta = 0$, and then iteratively conduct estimation at increasing (or decreasing) values of β using as initial values the ML estimators obtained at the nearest neighbour of the current β.

5.3 *Simulations under* $\mathcal{M}_a$

To evaluate the small sample performance of our estimators of $ACE(x)$ under $\mathcal{M}_a$ we conducted a $2 \times 3 \times 3$ factorial simulation experiment, corresponding to generating data under $VE \equiv P(S(1) = 0|S(0) = 1) \approx 0.3$ or 0.6; $\beta = 0, 1$, or 3; $ACE(\cdot)$ constant and equal to $0, 1/3$, or $1/2$. Each simulation generated 1000 vectors W according to the following steps:

Step 1. The first 500 vectors were set at $Z = 0$, the second 500 were set at $Z = 1$.

Step 2. X was a single covariate generated according to the $N(38, 6^2)$ distribution (this distribution resembles the age distribution in the VaxGen trial).

Step 3. Given X, $S(0)$ was drawn from a Bernoulli($\theta_p(X;\mu)$) distribution where $\theta_p(X;\mu)$ was as in (5.1) with $\mu = (\log(1/3), 0)$ so that $\theta_p(X;\mu)$ was constant and equal to 0.25 (this choice yields an expected number of infections in the placebo arm of 125, which is typical for a Phase III vaccine trial).

Step 4. $Y(0)$ was generated for all realizations with $S(0) = 1$ according to the density $f_p(y|x;\eta_0^a)$ given in (5.4). The parameter $\eta_0^a = (\lambda_0, \lambda_1, \sigma) = (2.3, 0.05, 1.0)$ (which resembles the viral load distribution for infected placebos in the VaxGen trial, where the mean and variance were 4.2 and 1.0, respectively).

Step 5. Given X and $Y(0)$, for each realization with $Z = 1$ and $S(0) = 1$, $S(1)$ was drawn from a Bernoulli($w(X, Y(0); \beta, \alpha)$) distribution with $w(x, y; \beta, \alpha)$ defined as in (5.2) and $\alpha = (\alpha_0, \alpha_1)$. In all simulations, $\alpha_1 = \log(2)/10$, so that for a 10 year increase in age, the odds of being in the always infected stratum doubled. The parameter α_0 was chosen so that $VE = 1 - \int\int w(x, y; \beta, \alpha) f_p(y|x;\eta_0^a)\varphi((x-38)/6^2)dxdy \approx 0.3$ or 0.6 (where $\varphi(\cdot)$ denotes the standard normal density). To ensure that $VE \approx 0.3$, α_0 was set at $-1.8, -5.8$, or -13.4, when β was set at $0, 1$, or 3 respectively; and to ensure that $VE \approx 0.6$, α_0 was set at $-3.1, -7.4$, or -16.3, when $\beta = 0, 1$, or 3.

Step 6. For the realizations with $Z = 1$ and $S(1) = 1$, $Y(1)$ was set equal to $Y(0) + \Delta$, with $\Delta = 0, 1/3$, or $1/2$. Note that with this choice, Δ is equal to $ACE(x)$.

It is easily shown that these steps generate data under model $\mathcal{M}_a$. We first investigate the different techniques for esimating $E(Y(1)|S(1) = 1, X = x)$.

As stated above, under (5.3), η_1 is not identifiable. Hence, within local neighborhoods of the regions of unidentifiability there may be several maxima, standard asymptotic theory may not apply, the variance of the MLE may not be the inverse information, and the Wald confidence intervals may not cover at the nominal level. Our simulations demonstrate many of these phenomena:

- There are several local maxima, so even after maximizing using many different starting values, we are not sure that our final estimate is the true MLE.

- At our sample sizes, the estimates of α and β_v are biased and are nowhere near converging to a multivariate normal distribution.

- Many of the simulations yielded information matrices that were either not invertible or else resulted in negative variances.

Figure 5.1 is an example of a simulated dataset with multiple local maxima. Each curve represents a different estimate for $E(Y(1)|S(1) = 1, X = x)$, obtained by starting optimization routines at different initial values. The true parameter values under which data were generated are $(\alpha_{v0}, \alpha_{v1}, \beta_v, \lambda_{v0}, \lambda_{v1}, \sigma_v)$ =(-16.3, 0.069, 3, 2.3, 0.05, 1.00), and the expected value under the truth is shown with the solid line. Estimated parameter values are A: (-39.2, 0.47, 4.07, -2.81, 0.15, 1.01), B: (-126, -0.35, 35.5, 5.94, -0.026, 0.73), C: (-217, 5.39, 0.083 ,5.59, -0.014, 0.62), D: (39.2, -0.47, -4.07, 1.36, 0.15, 1.01). Notice that some of the curves are very similar despite parameter estimates being quite different. The estimates A and D actually yield essentially identical estimates for $E(Y(1)|S(1) = 1, X = x)$ and are closely related to each other, with $(\hat{\alpha}_{v0}, \hat{\alpha}_{v1}, \hat{\beta}_v, \hat{\lambda}_{v1}, \hat{\sigma}_v^2)$ of A approximately equal to $(-\hat{\alpha}_{v0}, -\hat{\alpha}_{v1}, -\hat{\beta}_v, \hat{\lambda}_{v1}, \hat{\sigma}_v^2)$ of D. Parameter estimates B result in the largest value for the likelihood, so we take it to be the MLE.

The MLE under (5.3) with β_v unknown also has a tendency to over-fit the data. The three plots in Figure 5.2 show the MLE of different datasets simulated under the same distributions and parameters as Figure 5.1, and estimated with β_v free, $\beta_v = 0$, and $\beta_v = \beta = 3$. (Table 5.1 contains the parameter estimates for the plots in Figure 5.2.) Although $\beta \neq 0$, one can see that a straight line is a good approximation to the true value $E(Y(1)|S(1) = 1, X = x)$. On the other hand, the estimate based on leaving β_v free sees more in the data than is actually there. Naturally, as the sample size increases or as the variance of the data decreases, problems with over-fitting the data become smaller.

Although there is a tendency to over-fit the data, based on our simulations the estimate for $E(Y(1)|S(1) = 1, X = x)$ under (5.3) appears unbiased (see Table 5.2, columns A). The most serious limitation to this parameterization is that second moment estimation is not feasible using standard likelihood theory. If one improperly estimated the variance as the inverse of the information, then approximately 10% of simulations resulted in either negative variance estimates or had uninvertible information matrices. And if one removed the simulations with variance problems, the coverage probability was poor, between 82 and

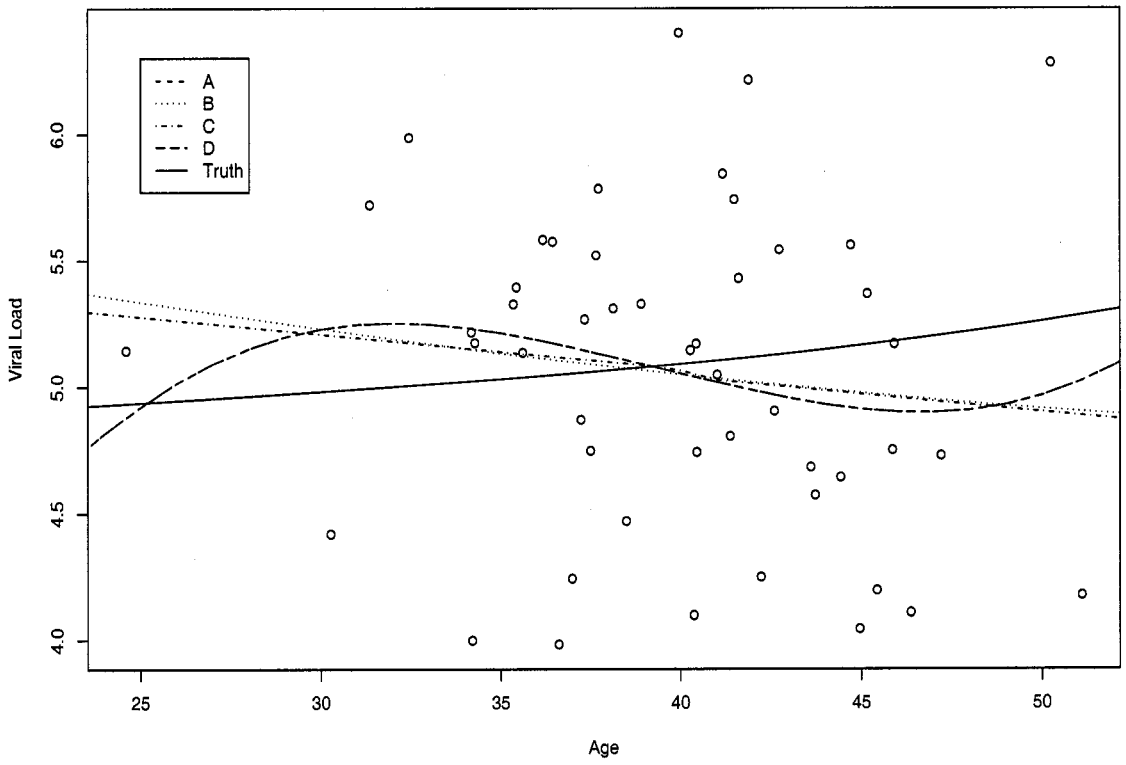

Figure 5.1: Different estimates for $E(Y(1)|S(1) = 1, X = x)$ under $\mathcal{M}_a$ using (5.1)-(5.4) with β_v unknown; X =Age.

93.5%, instead of the nominal 95% level.

Although variance estimation based on the information was never completely successful, it is interesting to notice the differences between rows. For example, for those simulations where $VE \approx 0.3$ and $\beta = 3$, the naïve variance estimate performed much better than simulations generated with $VE \approx 0.6$ and $\beta = 0$. As previously stated, data were generated setting $\alpha_{v1} = \log(2)/10 \approx 0.07$, $\lambda_{v1} = 0.05$, $\sigma_v = 1$, and $(\alpha_{v0}, \beta_v) = (-1.8, 0), (-5.8, 1), (-13.4, 3), (-3.1, 0), (-7.4, 1), (-16.3, 3)$. For the simulations with $\beta_v = 0$ we are in a region of non-identifiability. And for $\beta_v = 1$ and 3 there is non-identifiability if $\alpha_v = (-2.8, -0.05)$ and $(-11.4, -0.15)$, respectively. Our simulations do not appear to be far enough from this region to have performed adequately, although perhaps those that did perform better were further from the region of non-identifiability.

Because of the computational time it takes to calculate the bootstrap variance of the estimate of $E(Y(1)|S(1) = 1, X = x)$ under $\mathcal{M}_a$, it is not feasible to perform an extensive simulation study with M.3 correctly specified as (5.3). On a small subset of the simulated data used in Table 5.2 (the first 200 datasets generated with $VE \approx 0.3$ and $\beta = 1$) we examined the performance of the ML estimator of $E(Y(1)|S(1) = 1, X = x)$ with M.3 properly specified by (5.3) (leaving β_v free) and with its variance estimated from 100 bootstrap repetitions. The coverage probabilities of Wald-based 95% confidence intervals of $E(Y(1)|S(1) = 1, X = x)$ for $x = 30, 38$, and 55 using the bootstrap variance estimates and leaving β_v free were 0.950, 0.975, and 0.955, respectively. Coverage probabilities for the same 200 simulated datasets using linear regression ($\beta_v = 0$) were 0.960, 0.940, and 0.915, respectively.

Table 5.2 columns B and C also contain coverage probabilities fixing $\beta_v = 0$ and fixing $\beta_v = \beta$, respectively. For these simulations, simple linear regression ($\beta_v = 0$) performed very well, with good coverage probability and small bias. Of course, since setting $\beta_v = 0$ is actually a model misspecification, for certain simulation scenarios this model will not work as well. And, as the sample size increases, eventually one will have poor coverage probability. However, our choices for parameters from which to simulate the data were chosen a-priori and we feel they are fairly realistic of what one could see in practice. In many cases it may be reasonable to model $f_v(y|x; \eta_1)$ with the same distribution as $f_p(y|x; \eta_0^a)$.

In general, when fixing $\beta_v = \beta$, the coverage and bias were worse than the analyses assuming $\beta_v = 0$. Although not as extreme as leaving β_v free, there was still some over-fitting with $\beta_v = \beta$.

Table 5.3 reports Monte Carlo rejection probabilities, based on 1000 simulated datasets, of two-sided Wald tests of $H_0 : ACE\,(x) = 0$ at nominal 0.05 level for the values $x = 30, 38$, and 55 under model M_a using (5.1)-(5.4) with values of β set at 0, 1 or 3 (only one of them being the true value under which the data were generated) and β_v in (5.3) fixed at 0. Interestingly, even when the true value of β was different from 0 (and hence inference was conducted assuming an incorrect value of β_v in (5.3)) the type I error (i.e. the rejection probability under $\Delta = 0$) was quite close to the nominal 0.05 level and the tests had substantial power for detecting $\Delta = 1/2$ at $x = 38$ (the mean of X). The lower power of the tests at $x = 30$ and especially 55 is due to the fact that these values are quite distant from the mean of X (with 55 being about three standard deviations away from the mean).

Also, as predicted by theory, very poor inference is obtained when the values of β are incorrectly specified. The $ACE(x)$ is monotone decreasing in β, so if true $\beta = 0$ and one assumes $\beta = 1$, then the estimated $ACE(x)$ will be less than the truth. And if one incorrectly presumes that $\beta = 3$, then the $ACE(x)$ will be estimated as even smaller. This fact makes sense of the estimated powers reported in Table 5.3 when β is misspecified.

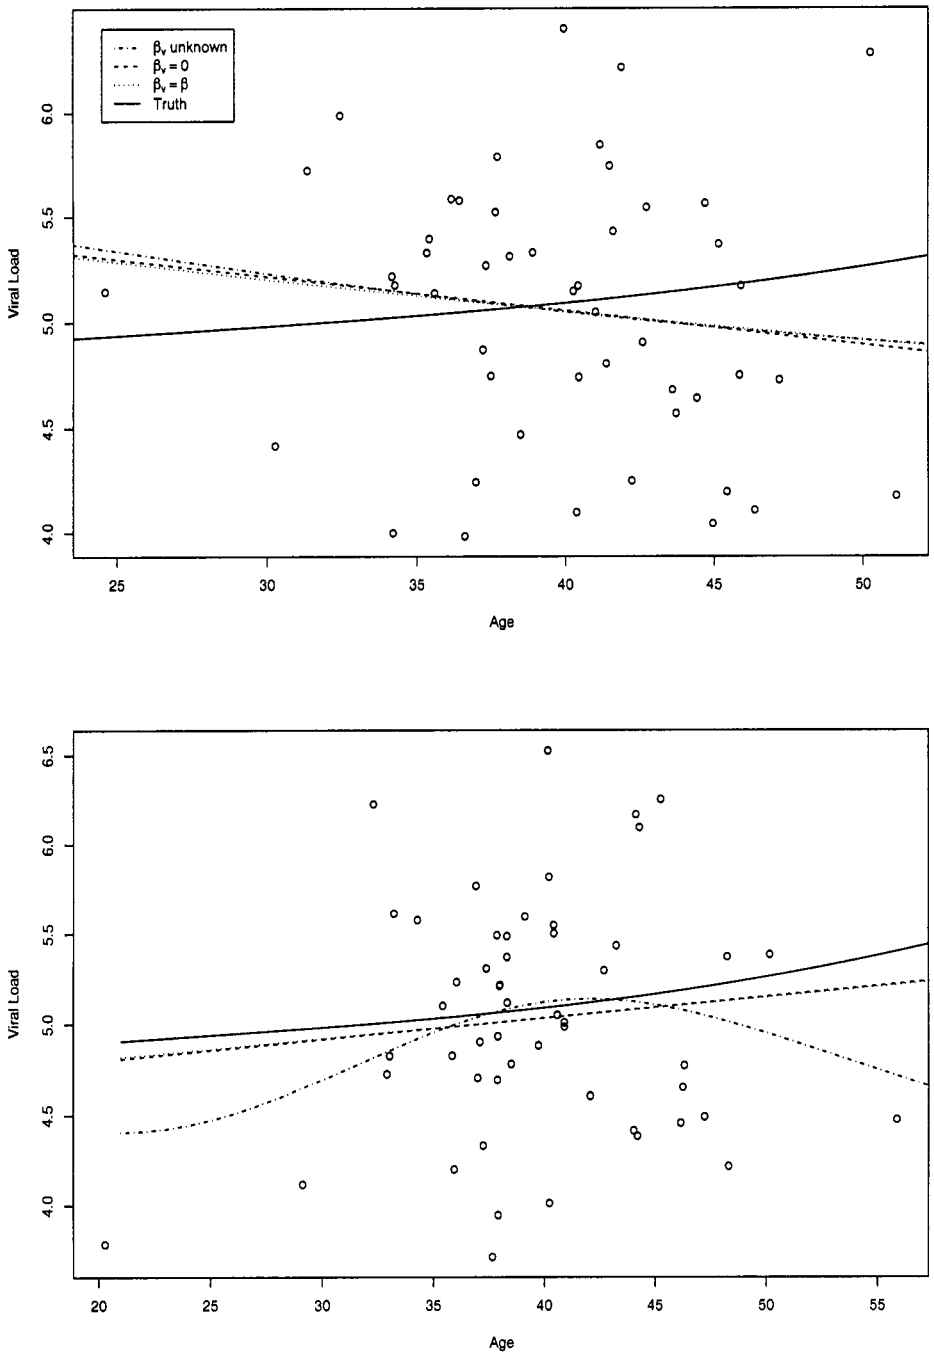

Figure 5.2: Estimates for $E(Y(1)|S(1) = 1, X = x)$ in a simulated dataset under $\mathcal{M}_a$ with β_v unknown, $\beta_v = 0$, and $\beta_v = \beta = 3$. (Labeling is the same for all three plots.)

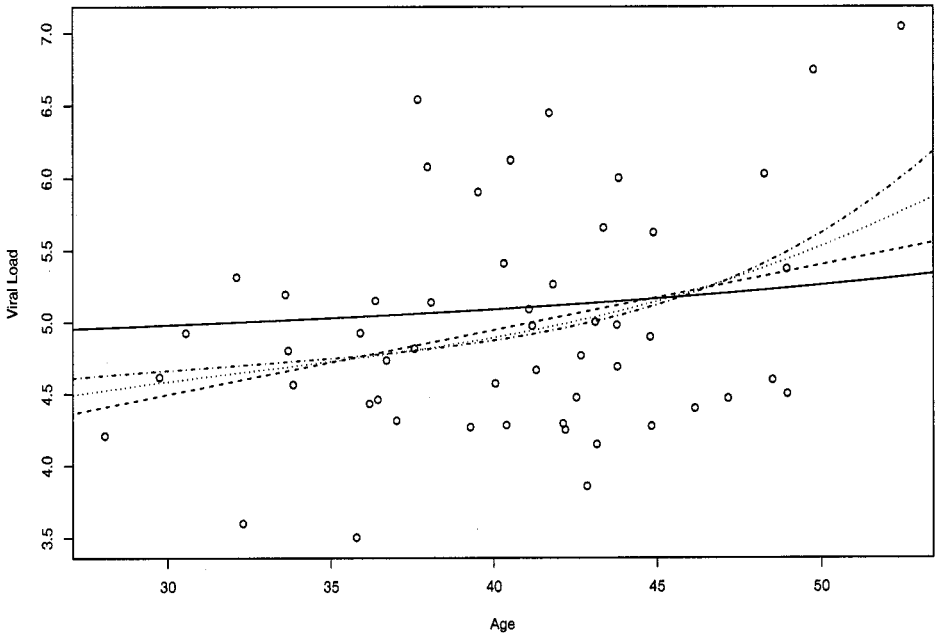

Figure 5.2: (continued)

Table 5.1: Estimated parameter values for Figure 5.2

Plot		α_{v0}	α_{v1}	β_v	λ_{v0}	λ_{v1}	σ_v
Top Left	β_v unknown	-126	-0.35	35.5	5.94	-0.026	0.73
	$\beta_v = 0$	-	-	-	5.70	-0.016	0.62
	$\beta_v = \beta$	-16.0	0.12	-	5.35	-0.01	0.68
Bottom Left	β_v unknown	0.23	0.50	-3.3	7.00	-0.04	0.75
	$\beta_v = 0$	-	-	-	4.56	0.012	0.66
	$\beta_v = \beta$	-9.9	-0.05	-	4.47	0.01	0.74
Right	β_v unknown	-27.1	0.15	4.79	-8.89	0.28	1.43
	$\beta_v = 0$	-	-	-	3.14	0.05	0.75
	$\beta_v = \beta$	-18.6	0.14	-	-0.91	0.13	0.99

Table 5.2: Coverage Probability and (Bias) of different methods for estimating $E(Y(1)|X)$; based on 95% Wald confidence intervals, the observed information, and 1000 simulations. Columns A: Estimating β_v; B: fixing $\beta_v = 0$; C: fixing $\beta_v = \beta$.

True		A*			B			C*	
β	X=30	X=38	X=55	X=30	X=38	X=55	X=30	X=38	X=55
$VE \approx 0.3$									
0	0.892	0.918	0.875	0.954	0.960	0.958	0.954	0.960	0.958
	(-0.000)	(-0.019)	(0.062)	(-0.012)	(-0.009)	(-0.003)	(-0.012)	(-0.009)	(-0.003)
1	0.928	0.930	0.912	0.951	0.954	0.939	0.937	0.902	0.892
	(0.019)	(-0.015)	(0.068)	(-0.014)	(0.010)	(-0.065)	(0.004)	(0.002)	(0.029)
3	0.920	0.935	0.927	0.949	0.952	0.925	0.946	0.940	0.912
	(0.022)	(-0.018)	(0.129)	(-0.033)	(0.014)	(-0.092)	(0.018)	(-0.015)	(0.105)
$VE \approx 0.6$									
0	0.865	0.886	0.820	0.942	0.951	0.934	0.942	0.951	0.934
	(-0.030)	(0.004)	(-0.033)	(-0.013)	(-0.005)	(0.012)	(-0.013)	(-0.005)	(0.012)
1	0.839	0.885	0.829	0.946	0.932	0.946	0.927	0.902	0.916
	(0.037)	(-0.016)	(0.110)	(0.002)	(-0.003)	(-0.038)	(-0.006)	(-0.002)	(0.008)
3	0.874	0.927	0.858	0.949	0.951	0.912	0.915	0.930	0.872
	(0.067)	(-0.018)	(0.117)	(-0.013)	(0.011)	(-0.074)	(0.019)	(-0.006)	(0.009)

* Coverage probability is after those with uninvertible informations or negative variances were removed. For $VE \approx 0.3, \beta = 0, 1, 3$ and $VE \approx 0.6, \beta = 0, 1, 3$, a total of 118, 100, 18, 217, 97, and 78 of the 1000 simulations in columns A were removed, respectively; a total of 0, 2, 1, 0, 5, and 1 of the 1000 simulations in columns C were removed, respectively.

Table 5.3: Power for detecting a 0, 1/3, and 1/2 mean shift alternative ($H_0 : ACE(x;\beta) = 0$) under $\mathcal{M}_a$ with β_v in (5.3) set equal to 0.

VE	True	Presumed		Δ=0			Δ=1/3			Δ=1/2	
	β	β	X=30	X=38	X=55	X=30	X=38	X=55	X=30	X=38	X=55
~ 0.3	**0**	**0**	**0.047**	**0.052**	**0.050**	**0.251**	**0.634**	**0.129**	**0.506**	**0.936**	**0.230**
	0	1	0.260	0.342	0.074	0.053	0.098	0.073	0.096	0.360	0.145
	0	3	0.520	0.588	0.086	0.137	0.097	0.065	0.066	0.080	0.106
	1	0	0.383	0.479	0.048	0.868	0.994	0.115	0.956	1	0.237
	1	**1**	**0.056**	**0.056**	**0.059**	**0.264**	**0.605**	**0.086**	**0.477**	**0.919**	**0.159**
	1	3	0.146	0.132	0.060	0.066	0.173	0.075	0.123	0.460	0.128
	3	0	0.832	0.951	0.067	0.993	1	0.161	0.998	1	0.288
	3	1	0.178	0.275	0.074	0.653	0.949	0.114	0.838	0.998	0.197
	3	**3**	**0.048**	**0.073**	**0.083**	**0.184**	**0.521**	**0.089**	**0.379**	**0.860**	**0.161**
~ 0.6	**0**	**0**	**0.058**	**0.042**	**0.058**	**0.193**	**0.462**	**0.128**	**0.350**	**0.807**	**0.195**
	0	1	0.462	0.752	0.085	0.144	0.122	0.061	0.073	0.049	0.086
	0	3	0.825	0.979	0.141	0.507	0.682	0.061	0.318	0.336	0.059
	1	0	0.483	0.813	0.062	0.828	0.997	0.193	0.921	1	0.303
	1	**1**	**0.058**	**0.071**	**0.043**	**0.213**	**0.471**	**0.107**	**0.357**	**0.775**	**0.167**
	1	3	0.217	0.391	0.051	0.057	0.109	0.171	0.102	0.060	0.096
	3	0	0.987	1	0.114	1	1	0.335	1	1	0.487
	3	1	0.433	0.635	0.068	0.999	1	0.334	1	1	0.485
	3	**3**	**0.049**	**0.053**	**0.067**	**0.190**	**0.492**	**0.115**	**0.339**	**0.807**	**0.189**

5.4 *Parameterization under* $\mathcal{M}_b$

Rather than modeling viral loads for infected placebos, a different approach assumes a distribution for viral loads in infected placebos in the always infected principle stratum. In other words, instead of specifying $f_p(y|x;\cdot)$, specifying $f_p^{ai}(y|x;\cdot)$. If $w(x,y;\alpha,\beta)$ is greater than 0 for all x and y, then one can write

$$f_p^*(y|x;\alpha,\eta_0^b) = \frac{w(x,y;\alpha,\beta)^{-1} f_p^{ai}(y|x;\eta_0^b)}{\int w(x,y;\alpha,\beta)^{-1} f_p^{ai}(y|x;\eta_0^b)dy},$$

and all components of the likelihood (3.1) are specified as shown in Appendix C.1.

A major advantage of model $\mathcal{M}_b$ is that, as opposed to model $\mathcal{M}_a$, the functional form for $f_v(y|x;\eta_1)$ can be easily chosen so as to ensure both that the null hypothesis H_0 can hold under the model and that η_1 is identified and $\sqrt{N}-$estimable. This approach also yields clearly interpretable parameter estimates and simple tests of null hypotheses.

We study estimation for this method of parameterization using logistic/Gaussian models. In our example and simulations we use

$$f_p^{ai}(y|x;\eta_0^b) = \phi(y;x^T\gamma_p,\sigma_p^2) \text{ and } f_v(y|x;\eta_1) = \phi(y;x^T\gamma_v,\sigma_v^2), \tag{5.5}$$

where $\eta_0^b = (\gamma_p,\sigma_p)$, $\eta_1 = (\gamma_v,\sigma_v)$, γ_p and γ_v are unknown vectors of the same dimension as x, $\phi(y;x^T\gamma_v,\sigma_v^2)$ is a normal density with mean $x^T\gamma_v$ and variance σ_v^2, and σ_p and σ_v are unknown. Note that under (5.5) $ACE(x)$ takes the simple and easily interpretable form $x^T(\gamma_v-\gamma_p)$. Interestingly, under (5.2) and (5.5) a closed form expression for the likelihood $\mathcal{L}_{b,n}(\rho^b)$ exists because

$$\begin{aligned} 1-VE(x;\beta,\alpha,\eta_0^b) &\equiv \int w(x,y;\beta,\alpha) f_p^*(y|x;\alpha,\eta_0^b)dy \\ &= \left[1+exp\left(-\beta\left(x^T\gamma_p-\frac{1}{2}\beta\sigma_p^2\right)-x^T\alpha\right)\right]^{-1}. \end{aligned}$$

In addition,

$$f_p^{prot}(y|x;\beta,\eta_0^b) = \phi(y;x^T\gamma_p-\beta\sigma_p^2,\sigma_p^2).$$

This is derived in section D.3 of the Appendix. The log-likelihood and score equations under these parameterizations are given in section D.4.

In addition to tests of $H_0 : ACE(x) = 0$, parameterization under $\mathcal{M}_b$ allows simple tests for interaction, i.e. a test of $H_0^{inter} : \gamma_{vj} = \gamma_{pj}$ is a test of no interaction between treatment and the jth covariate. A global test of the average causal effect of vaccine on viral load is a test of $H_0^{global} : \gamma_v = \gamma_p$. Specifically, a likelihood ratio statistic for testing H_0^{global} is

$$\lambda = \frac{\mathcal{L}_b(\hat{\rho}^b)}{\mathcal{L}_b(\hat{\rho}_0^b)}, \text{ with } 2\log\lambda \rightarrow_d \chi_q^2,$$

where $\hat{\rho}^b$ is the MLE of ρ^b and $\hat{\rho}_0^b$ is the MLE under the constraint that $\gamma_v = \gamma_p$. A similar likelihood ratio statistic could be used for testing H_0^{inter}, or alternatively one could use a Wald statistic.

Computation is also more simple under $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5) than under $\mathcal{M}_a$ using (5.1)-(5.4) because the likelihood can be written in closed form. Computationally, the only challenge to maximizing this likelihood is that $f_p^*(y|x;\alpha,\eta_0^b)$ is written as the mixture $(1-VE(x;\cdot))f_p^{ai}(y|x;\cdot)+VE(x;\cdot)f_p^{prot}(y|x;\cdot)$, which suggests employing an EM or Newton-Rapson algorithm. For our simulations and examples, we again have used quasi-newton methods implemented in R using the function optim(·).

One potential disadvantage of model $\mathcal{M}_b$ is that it requires that $w(x,y;\beta,\alpha)$ is nonzero for all x and y. This rules out sharp bound analyses (HHS) and is tantamount to assuming the vaccine is not 100% effective in any subpopulation. However, this is a plausible assumption for most candidate HIV vaccines (see Graham, 2002), and even if there were subpopulations where the vaccine was known to be 100% effective, participants in these subpopulations would not belong to the *ai* principal stratum and could therefore simply be removed before performing the analysis.

Note that by assuming a functional form for $f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1,x)$ and another for $P(S(1)=1|S(0)=1,Y(0),X=x)$ we are indirectly imposing functional form restrictions on the distribution $f_{Y(0)|S(0)=1,X}(y|S(0)=1,x)$, which is identified by the observed data without assumptions A.1 and A.2. Indeed, this functional form is strongly driven by the chosen value of β. In fact, β is identified under model $\mathcal{M}_b$. A consequence of this remark is that some choices of β can result in poor model fits. For example, extreme values of β may correspond to a bimodal distribution for $f_{Y(0)|S(0)=1,X}(y|S(0)=1,x)$ and this may be contradicted by the evidence in the data. This is demonstrated in Figures 5.3 and

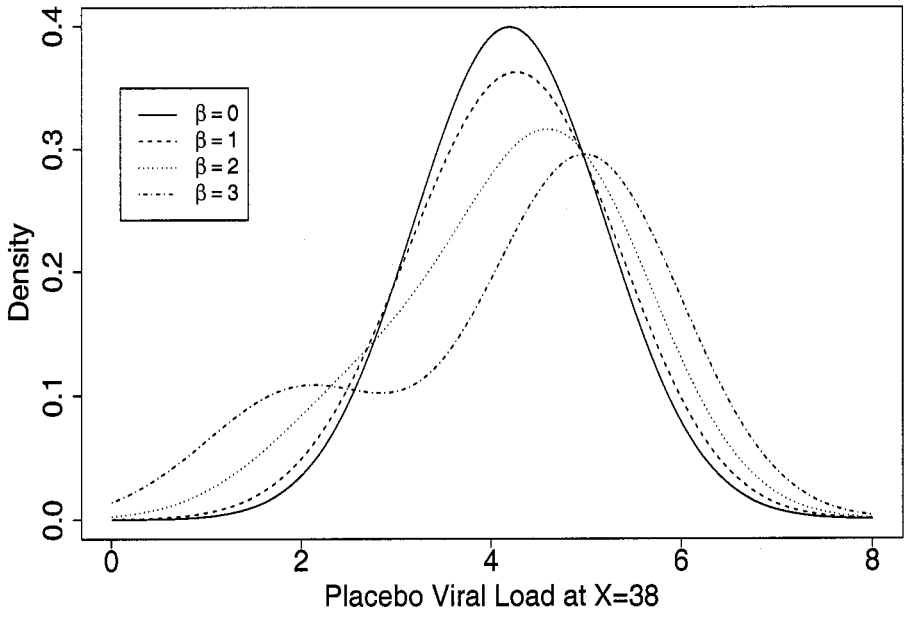

Figure 5.3: Induced placebo viral load densities at the mean covariate value for different values of the sensitivity parameter β under $\mathcal{M}_b$.

5.4, which show how the density and expected value of $Y(0)|X, S(0) = 1$ differ for different assumed values of β. These figures were created under model $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5) with each distribution set up so that $VE \approx 0.3$ and $E(Y(0)|X = 38, S(0) = 1) \approx 4.2$ (see Simulations Section). One can see that at these settings for a sensitivity analysis ranging from $\beta = 0$ to 3, one assumes widely different distributional forms for the observed viral loads in the placebo arm. In order for an analysis to be reasonable, the observed data must support the distributional assumptions. If extreme values of β were indeed regarded as plausible prior to assuming M.1-M.4b, then we recommend that the analyst consider more flexible distributional shape assumptions since poor model fits under plausible values of β suggest incorrect specification of at least one of the assumptions in M.1-M.4b. (More discussion of this point is found in Appendix D.5).

A second disadvantage of model $\mathcal{M}_b$, therefore, is that sensitivity analyses may be valid only over a smaller range of β. From simulations (see Section 5.5 for details on

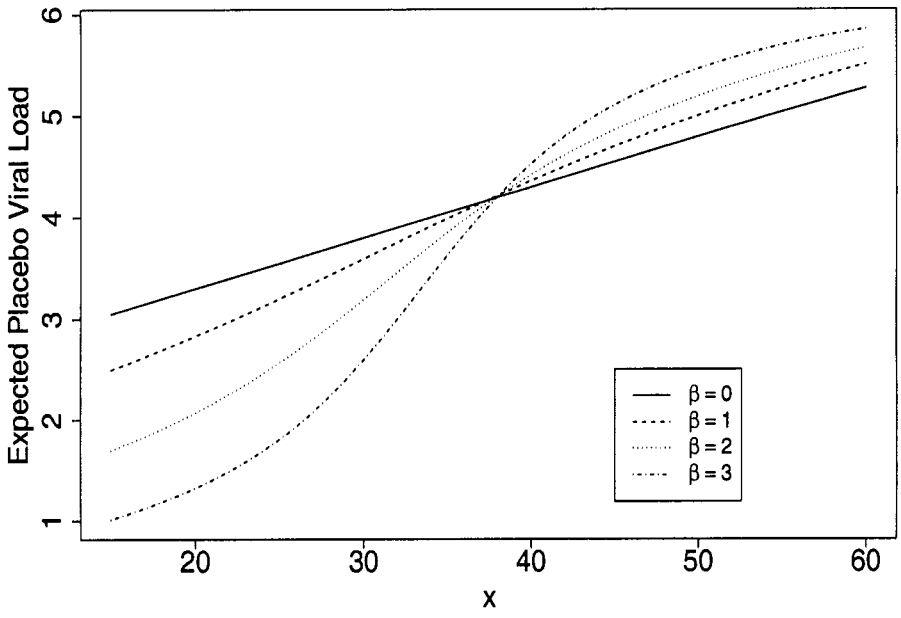

Figure 5.4: Expected placebo viral load as a function of X for different values of the sensitivity parameter β under $\mathcal{M}_b$.

how data were generated) we have found that when β is incorrectly set away from its true value, some parameter estimates are nonsensical because the model poorly fits the observed data. Figure 5.5 is a sensitivity analysis on data simulated under $\mathcal{M}_b$ (with $\beta = 0, VE \approx 0.3$, and $ACE(x) = 0$ for all x) where the model appears to be a good fit to the data for all $\beta \in [-3, 3]$. As one would expect, $\widehat{ACE}(x)$ is monotone in β and the model-based estimates for VE, $\sum_i^N VE(x_i; \beta, \hat{\alpha}_0, \hat{\eta}_0)/N$, written at the top of the plot for $\beta = (-3.0, -2.5, -2.0, \cdots, 2.5, 3.0)$ are similar to the non-parametric estimate of VE, $1 - n_v N_p/(n_p N_v)$. In contrast, Figure 5.6 is an example of a sensitivity analysis where this parameterization breaks down. The data used to construct Figure 5.6 were generated under the same model as the data used to construct Figure 5.5. However, as the assumed β moves farther away from the true β ($\beta = 0$), the $\widehat{ACE}(x)$ shifts unexpectedly. Notice that at the same time, the model-based estimate of VE becomes very different from the non-parametric estimate. In short, what is happening is that for large/small values of β the assumed model is forcing a bimodal distribution, and at some point the model becomes such a bad fit that the estimated mixing proportion $VE(x; \beta, \hat{\alpha}_0, \eta_0)$ is forced to be either much higher or much lower than the truth, implying that most of the observed viral loads in the placebo arm come from only one of the distributions, $f_p^{prot}(\cdot)$ or $f_p^{ai}(\cdot)$.

5.5 Simulations under $\mathcal{M}_b$

Simulation studies were performed to investigate finite sample properties and robustness of estimates to misspecifying β. Data were generated from the models (5.1)-(5.2) and (5.5) using a $2 \times 3 \times 3$ factorial simulation experiment, corresponding to $VE \equiv P(S(1) = 0 | S(0) = 1) \approx 0.3$ or 0.6; $\beta = 0, 1$, or 3; $ACE(\cdot)$ constant and equal to $0, 1/3$, or $1/2$. Each simulation generated 1000 vectors W according to the following steps (repeated from Section 5.3, with changes in *Steps 4* and *5*):

Step 1. The first 500 vectors were set at $Z = 0$, the second 500 were set at $Z = 1$.

Step 2. X was a single covariate generated according to the $N(38, 6^2)$ distribution.

Step 3. Given X, $S(0)$ was drawn from a Bernoulli$(\theta_p(X; \mu))$ distribution where $\theta_p(X; \mu)$ was as in (5.1) with $\mu = (\log(1/3), 0)$.

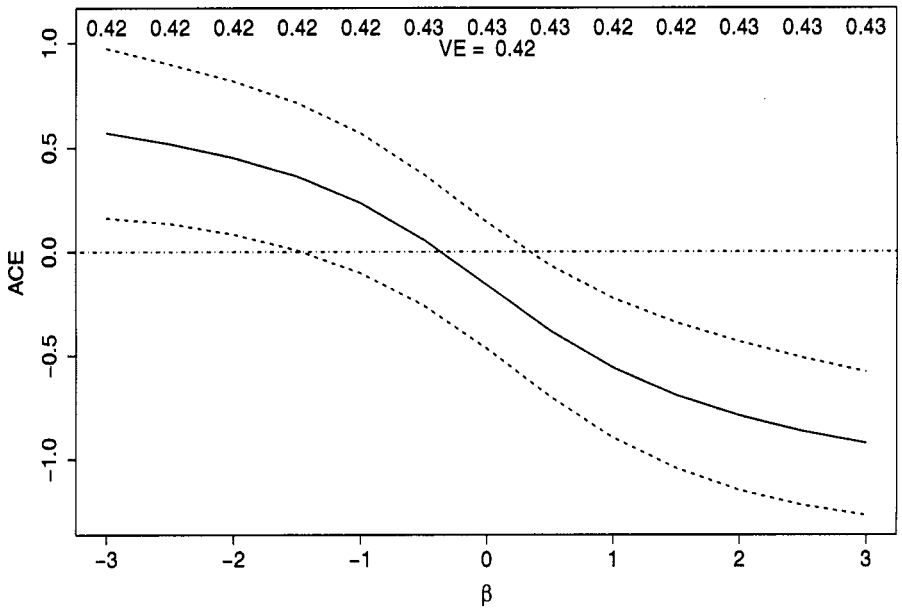

Figure 5.5: Sensitivity analysis of a simulated dataset under $\mathcal{M}_b$ where the model appears adequate for all β.

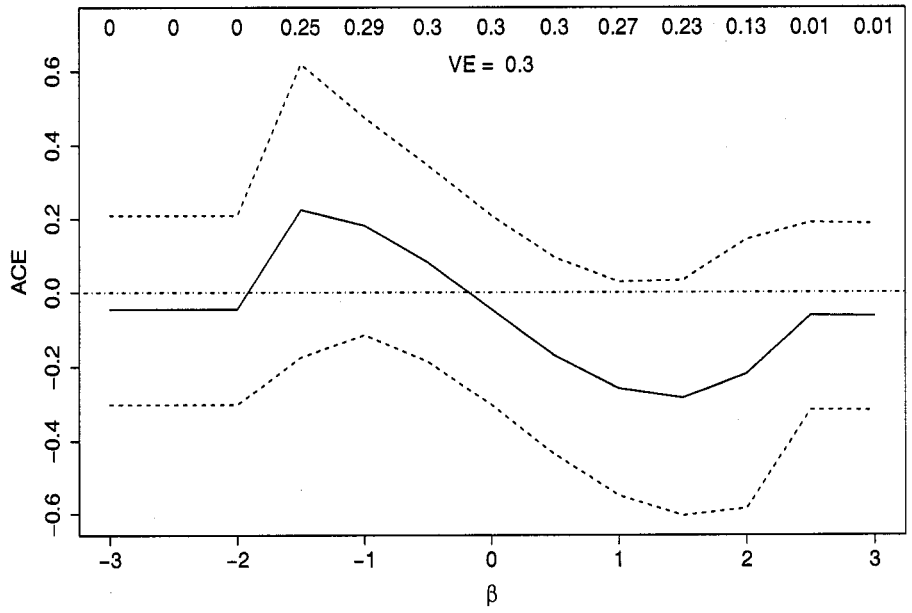

Figure 5.6: Sensitivity analysis of a simulated dataset under $\mathcal{M}_b$ where the model is not adequate for all β.

Step 4. $Y(0)$ was generated for all realizations with $S(0) = 1$ according to the density $f_p^*(y|x;\alpha,\eta_0^b)$ induced by $f_p^{ai}(y|x;\eta_0^b)$ in (5.5) and $w(x,y;\beta,\alpha)$ in (5.2). We detail our choices of α and η_0^b in *Step 5.*

Step 5. Given X and $Y(0)$, for each realization with $Z = 1$ and $S(0) = 1$, $S(1)$ was drawn from a Bernoulli$(w(X, Y(0); \beta, \alpha))$ distribution with $w(x,y;\beta,\alpha)$ defined as in (5.2). In all simulations, $\alpha_1 = \log(2)/10$, so that for a 10 year increase in age, the odds of being in the ai stratum doubled. Parameter values, α_0 and η_0^b were chosen together so that $E(Y(0)|S(0) = 1) \approx 4.2$ and $VE \approx 0.3$ or 0.6. For $VE \approx 0.3$ we set $(\alpha_0, \gamma_{p0}, \gamma_{p1}, \sigma_p)$ as $(-1.8, 2.3, 0.05, 1.0), (-5.7, 2.6, 0.05, 1.0)$, or $(-12.1, 3.1, 0.05, 1.0)$, when β was equal to $0, 1$, or 3, respectively; for $VE \approx 0.6$ we set $(\alpha_0, \gamma_{p0}, \gamma_{p1}, \sigma_p)$ as $(-3.1, 2.3, 0.05, 1.0)$, $(-7.4, 2.9, 0.05, 1.0)$, or $(-17, 4.2, 0.05, 1.0)$, when $\beta = 0, 1$, or 3, respectively.

Step 6. For the realizations with $Z = 1$ and $S(1) = 1$, $Y(1)$ was set equal to $Y(0) + \Delta$, with $\Delta = 0, 1/3$, or $1/2$.

The sensitivity analyses shown in Figures 5.5 and 5.6 were performed on data generated using $VE \approx 0.3, \beta = 0$, and $\Delta = 0$. Based on informal observation, only one out of every four or five simulated datasets generated under these specifications resulted in sensitivity analyses that were valid (i.e. the data fit the model, as in the lower right plot of Figure 5.7) for all β from $\beta = [-3, 3]$. Figure 5.8 shows a typical sensitivity analysis generated with true $\beta = 3, VE \approx 0.3$, and $\Delta = 0$. In this sensitivity analysis, the data appear to fit the model for β as low as -1 or -2.

It is impossible to specify a range of β where the method under $\mathcal{M}_b$ always performs well. In general, the more narrow the range of β the more likely $\mathcal{M}_b$ is reasonable. The modeling will actually perform better for smaller sample sizes: with smaller samples it is harder to see models and values of β that are incompatible with the observed data because the less data you have the easier it is to use whatever you want to describe it. In practice, I recommend computing the empirical and model-based estimates of VE. Although not a formal test, when these two quantities are clearly different, the model is obviously misspecified.

Table 5.4 reports rejection probabilities for the same test under the same settings as in Table 5.3 except that data were generated under model $\mathcal{M}_b$ and ML estimation was conducted under model $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5). Simulation results are similar,

although it should be pointed out that because data were generated under different models, Tables 5.3 and 5.4 are not directly comparable. Table 5.5 reports rejection probabilities for H_0^{inter} and H_0^{global}, as defined in Section 4.2, based on the data and estimators used in Table 5.4. Finally, Table 5.6 shows the power to reject H_0 assuming no interaction and $\sigma_v = \sigma_p$. In all cases, when β is properly specified the tests have adequate power and size.

Table 5.4: Size/power for detecting a 0, 1/3, and 1/2 mean shift alternative of the Wald test of $H_0 : ACE(x) = 0$ under $\mathcal{M}_b$.

VE	True	Presumed		Δ=0			Δ=1/3			Δ=1/2	
	β	β	$x = 30$	$x = 38$	$x = 55$	$x = 30$	$x = 38$	$x = 55$	$x = 30$	$x = 38$	$x = 55$
~ 0.3	**0**	**0**	**0.058**	**0.055**	**0.047**	**0.308**	**0.673**	**0.124**	**0.559**	**0.949**	**0.211**
	0	1	0.232	0.363	0.051	0.054	0.084	0.096	0.110	0.350	0.157
	0	3	0.492	0.710	0.153	0.253	0.323	0.133	0.213	0.272	0.159
	1	0	0.471	0.573	0.048	0.881	0.988	0.095	0.956	1	0.172
	1	**1**	**0.056**	**0.061**	**0.050**	**0.269**	**0.547**	**0.121**	**0.449**	**0.853**	**0.197**
	1	3	0.314	0.432	0.096	0.384	0.439	0.140	0.441	0.579	0.200
	3	0	1	1	0.391	1	1	0.173	1	1	0.092
	3	1	0.422	0.422	0.051	0.767	0.945	0.032	0.888	0.994	0.057
	3	**3**	**0.056**	**0.052**	**0.058**	**0.186**	**0.502**	**0.121**	**0.351**	**0.831**	**0.201**
~ 0.6	**0**	**0**	**0.057**	**0.058**	**0.069**	**0.209**	**0.507**	**0.143**	**0.387**	**0.824**	**0.212**
	0	1	0.466	0.742	0.117	0.139	0.136	0.066	0.063	0.050	0.098
	0	3	0.797	0.988	0.210	0.492	0.697	0.127	0.326	0.382	0.120
	1	0	0.644	0.900	0.083	0.896	0.998	0.174	0.964	1	0.257
	1	**1**	**0.061**	**0.049**	**0.062**	**0.180**	**0.388**	**0.113**	**0.307**	**0.686**	**0.167**
	1	3	0.284	0.460	0.129	0.132	0.110	0.121	0.124	0.170	0.164
	3	0	1	1	0.062	1	1	0.048	1	1	0.070
	3	1	0.505	0.655	0.038	0.750	0.935	0.076	0.838	0.982	0.116
	3	**3**	**0.063**	**0.064**	**0.056**	**0.145**	**0.301**	**0.105**	**0.240**	**0.545**	**0.168**

Table 5.5: Size of the Wald test for $ACE(x)$ independent of x (H_0^{inter}: $\gamma_{v1} = \gamma_{p1}$), and size/power of the likelihood ratio test of $ACE(x) = 0$ for all x (H_0^{global}: $\gamma_v = \gamma_p$), under $\mathcal{M}_b$.

VE	True β	Presumed β	H_0^{inter}	H_0^{global} $\Delta = 0$	$\Delta = 1/3$	$\Delta = 1/2$
~ 0.3	**0**	**0**	**0.063**	**0.051**	**0.561**	**0.907**
	0	1	0.068	0.283	0.093	0.289
	0	3	0.128	0.569	0.460	0.530
	1	0	0.162	0.491	0.983	1
	1	**1**	**0.057**	**0.052**	**0.456**	**0.780**
	1	3	0.135	0.261	0.311	0.453
	3	0	0.948	1	1	1
	3	1	0.187	0.307	0.874	0.979
	3	**3**	**0.055**	**0.057**	**0.420**	**0.769**
~ 0.6	**0**	**0**	**0.075**	**0.056**	**0.430**	**0.766**
	0	1	0.087	0.662	0.124	0.076
	0	3	0.139	0.975	0.635	0.368
	1	0	0.146	0.824	0.994	1
	1	**1**	**0.051**	**0.059**	**0.317**	**0.599**
	1	3	0.118	0.383	0.104	0.151
	3	0	0.878	1	1	1
	3	1	0.170	0.507	0.901	0.971
	3	**3**	**0.057**	**0.056**	**0.275**	**0.528**

5.6 Summary

In this chapter I have presented two different parametric approaches to estimating $ACE(x)$. The defining difference between models $\mathcal{M}_a$ and $\mathcal{M}_b$ is whether one specifies the distributional form for viral load for infected placebos ($\mathcal{M}_a$) or for always infected placebos ($\mathcal{M}_b$). There are advantages and disadvantages to both methods. Under $\mathcal{M}_b$, parameters are easily interpretable and lead to simple tests for interaction and tests of H_0^{global}. However, β is identifiable under $\mathcal{M}_b$, which may mean that it is difficult to choose parametric distributions that are compatible with the observed data over the range of β that is scientifically interesting. This is not a problem under $\mathcal{M}_a$. However, $\mathcal{M}_a$ has problems of its own, particularly the fact that it may not be possible to choose models that have identifiable parameters but that allow the null hypothesis, $ACE(x) = 0$, to hold for all x as sample sizes go to infinity. For the above reasons, one would probably prefer model $\mathcal{M}_b$ if the scientifically plausible range for β is narrow. For a large range, $\mathcal{M}_a$ might be preferable. In practice, I would recommend doing both (and reporting both). These are, of course, sensitivity analyses – which I define as analyses performed over a range of assumptions to investigate the sensitivity of conclusions to those assumptions. Why not investigate the sensitivity of results to distributional assumptions?

These sensitivity analyses can be thought of as examining departures from assumption A.3, that $Y(0) \amalg S(1)|S(0) = 1, X$. If A.3 holds, then one can simply perform standard regression analyses on the infected subjects. Adding to X additional covariates thought to be associated with both vaccine efficacy and viral load may make A.3 more believable. However, this does not imply that adding covariates to the analysis restricts the range of values of β considered plausible (Scharfstein et al., 1999).

Table 5.6: Size/power for detecting a 0, 1/3, and 1/2 mean shift alternative of the Wald test of $H_0 : ACE(x) = 0$ under $\mathcal{M}_b$ assuming $(\gamma_{v1}, \sigma_v) = (\gamma_{p1}, \sigma_p)$.

VE	True	Presumed		Δ	
	β	β	0	1/3	1/2
~ 0.3	**0**	**0**	**0.049**	**0.672**	**0.938**
	0	1	0.366	0.090	0.350
	0	3	0.216	0.575	0.774
	1	0	0.484	0.990	0.999
	1	**1**	**0.057**	**0.544**	**0.860**
	1	3	0.371	0.645	0.728
	3	0	0.992	1	1
	3	1	0.408	0.959	0.994
	3	**3**	**0.046**	**0.564**	**0.872**
~ 0.6	**0**	**0**	**0.058**	**0.491**	**0.819**
	0	1	0.800	0.171	0.058
	0	3	0.966	0.836	0.634
	1	0	0.865	0.998	1
	1	**1**	**0.056**	**0.412**	**0.723**
	1	3	0.629	0.195	0.161
	3	0	1	1	1
	3	1	0.998	1	1
	3	**3**	**0.045**	**0.871**	**0.996**

Chapter 6

TIME-TO-EVENT OUTCOMES

6.1 Additional Motivation

Some of the most important outcomes in HIV vaccine trials may be post-infection time-until-event outcomes (Gilbert et al., 2003b). A couple of examples include the time from HIV infection diagnosis until the onset of AIDS or the time from infection diagnosis until the viral load passes a specific threshold.

In Phase III vaccine trials, such as VaxGen's trial of AIDSVAX B/B, several thousands of individuals are randomized to either the vaccine or placebo arm. As in a typical clinical trial, they are followed for a certain amount of time, with the primary endpoint being HIV infection. Those that get infected are then enrolled in a post-infection study, and monitored for several additional years. Those who are not infected by the initial stopping time are no longer followed.

In this situation one could make a strong argument for performing an intention-to-treat (ITT) analysis: i.e., using all randomized individuals, ignoring infection status and defining the outcome of interest as the post-infection outcome (e.g., time from randomization until AIDS). However, this is not an ideal analysis because the vast majority of trial participants will not get infected and are therefore not followed past the time that the initial study concludes, t_1. If one censors these uninfected individuals at t_1 (which we refer to as analysis ITT_a), then one induces dependent censoring, leading to a biased estimate that under-estimates the survival probability for $t > t_1$. On the other hand, if one assumes that the post-infection outcome occurs in none of the individuals who are not infected at time t_1 and censors these individuals at the time when the post-infection study ends, $t_1 + t_2$, (we refer to this analysis as ITT_b) then one is most likely over-estimating the survival probability at $t > t_1$, particularly for large t_2.

A simple illustration: Define X_1 as the time from randomization to HIV infection diagnosis and X_2 as the time from HIV infection diagnosis to AIDS. The ITT outcome of interest is therefore $X = X_1 + X_2$. For simplicity, assume that all uninfected individuals are censored at the end of the first phase of the trial (we assume this throughout). This is stated by defining C_1 as the time from randomization to pre-infection censoring and setting $C_1 = t_1$ for all participants. Define C_2 as the time from HIV infection diagnosis to post-infection censoring. Under the ITT_a analysis, the time from randomization to censoring, C, equals $\min(X_1, t_1) + I(X_1 \leq t_1)\min(C_2, t_2)$. Under the ITT_b analysis, $C = I(X_1 > t_1)(t_1 + t_2) + I(X_1 \leq t_1)(X_1 + \min(C_2, t_2))$ and X is assumed to be greater than C if $X_1 > t_1$. Under either analysis, X and C are not independent, even if X_1 and X_2 are independent of C_1 and C_2 respectively.

For example, consider the simple scenario where $t_1 = t_2$, X_1 and X_2 are independent and both are distributed as $\text{EXP}(\lambda)$ with cdf F. This implies that the true survival probability at time t is $1 - G(t)$, where G is the cdf of a $\text{GAM}(2, \lambda)$ distribution. Under this model the analysis ITT_a yields a Kaplan-Meier estimate of $S(t)$ for $t > t_1$ that converges to $(1 - F(t))(1 - G(t_1))$. For analysis ITT_b, for $t > t_1$ the estimate of $S(t)$ converges to $1 - G(t_1) - F(t - t_1)[F(t_1) - G(t_1)]$. Both estimates are clearly biased as shown in Figure 6.1 for $t_1 = t_2 = 3$ and $\lambda = 0.1$.

The only unbiased ITT analysis would be to censor everyone who does not experience the post infection event at t_1. This is, of course, unattractive because much information is thrown away. From a statistical standpoint, the ideal trial design would not have an initial stopping point, but would follow all individuals for the entire study time, $t_1 + t_2$. However, resource constraints make a Phase III trial of this type unfeasible.

In this chapter I will extend the methods of the previous chapters to the case where the outcome of interest is a right-censored time-to-event outcome. Therefore, my methods will answer the following type of question: among those who would have been infected regardless of treatment assignment, does assignment to the vaccine increase/decrease the probability of being AIDS free t months post-infection? This question allows us to specifically look at the causal effect of vaccination on the time from infection diagnosis to some post-infection outcome. It should be emphasized that my methods address a different question than

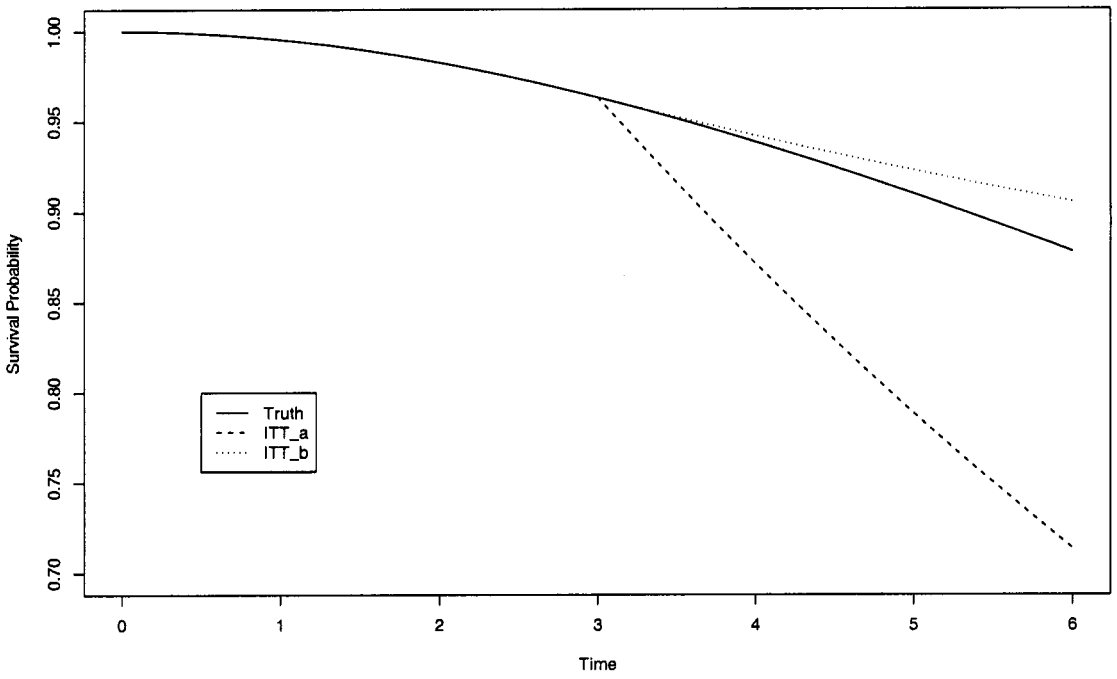

Figure 6.1: Example of survival distribution and biased estimation of that distribution using different ITT analyses.

the question answered by an ITT analysis. An ITT analysis addresses whether or not vaccination increases the probability of being AIDS free t months *post-randomization.* One could legitimately argue that despite the problems mentioned in the preceding paragraphs, an ITT analysis better addresses a more important question for public health purposes because it is more interpretable for future subjects and because it includes information with regards to both infection and the post-infection outcome. In a sense, an ITT analysis combines the primary and secondary analyses into one analysis. However, many researchers want an analysis that specifically addresses the post-infection questions. This is especially true for the next round of candidate HIV vaccines (Graham (2002); Gilbert et al. (2003b); www.hvtn.org; see also Appendix A). The analyses discussed in this chapter allow one to look directly at the causal effect of vaccination on post-infection time-to-event outcomes. In practice, I advocate using and presenting results from both ITT and causal analyses.

6.2 Notation and Assumptions

I slightly adjust my notation here from that presented in Chapter 2 to address the time-to-event problem. Let T_i be the time from infection diagnosis until some event for subject i, and define C_i as the time from infection diagnosis until censoring. As with typical time-to-event outcomes, we do not observe T_i but rather we observe $Y_i = min(T_i, C_i)$ and $\Delta_i = I(Y_i = T_i)$. Notice that T_i, C_i, Y_i, and Δ_i only exist if $S_i = 1$; otherwise they are assigned the value $*$ with probability 1.

As before, I use potential outcomes and assume SUTVA. The outcomes $W_i = (Z_i, S_i(0), S_i(1), T_i(0), T_i(1), C_i(0), C_i(1), X_i)$ for $i = 1, \cdots, N$ are i.i.d. copies of a random vector $W = (Z, S(0), S(1), T(0), T(1), C(0), C(1), X)$ and similarly the observed data $O_i = (Z_i, S_i, Y_i, \Delta_i, X_i)$ are i.i.d. copies of $O = (Z, S, Y, \Delta, X_i)$. As before, one could evaluate the average causal effect of vaccination in the ai stratum defined as $ACE(x) = E[T(1) - T(0) | S(0) = S(1) = 1, X = x]$. With a time-to-event outcome, it may be of greater interest to compare within the always infected stratum the probablities of being event free at time t, defined as $1 - P(T_i(z) \leq t | S_i(0) = S_i(1) = 1, X = x)$ for $z = 0, 1$.

Hence, my estimand of interest is the "survival" causal effect in the ai stratum, defined as

$$\begin{aligned} SCE(t,x) &\equiv P(T_i(0) \leq t | S_i(0) = S_i(1) = 1, X = x) - P(T_i(1) \leq t | S_i(0) = S_i(1) = 1, X = x) \\ &= F_p^{ai}(t|x) - F_v^{ai}(t|x). \end{aligned}$$

My goal is to estimate the function $SCE\,(t,x)$ based on the observed data $O_i, i = 1, ..., N$. As in Chapter 2, in addition to assuming SUTVA, I assume

- Randomization:

$$(S\,(0)\,, S\,(1)\,, T\,(0)\,, T\,(1)\,, C\,(0)\,, C\,(1)) \amalg Z|X \tag{6.1}$$

- Monotonicity:

$$S_i(1) \leq S_i(0), \text{ for } i = 1, \cdots, N \tag{6.2}$$

- A form for $P(S(1) = 1|S(0) = 1, T(0), X)$:

$$P(S(1) = 1|S(0) = 1, T(0), X) = w(T(0), X; \beta); \tag{6.3}$$

  where $w(t, x; \beta) = \Phi\left\{m(x) + g\,(t, x; \beta)\right\}$, β is fixed and known, $\Phi\,(\cdot)$ is a known cdf, $m(\cdot)$ is an unspecified function of X, and for each β, $g\,(\cdot; \beta)$ is a known function of X and T.

Our only additional assumption is

- Independent Censoring:

$$C(z) \amalg T(z)|S(z) = 1, X. \tag{6.4}$$

This final assumption is equivalent to the standard assumption of independent censoring often made when analyzing time-to-event data. I will refer to the collective assumptions of SUTVA, (6.1), (6.2), (6.3), and (6.4) as assumptions $\mathcal{A}$.

6.3 *Maximum Likelihood Estimation*

In chapter 3, I constructed the likelihood for the case where there is no censoring. One can easily extend the likelihood to account for independent censoring. I express the joint density of the observables as

$$f_O(O) = f_X(X)P_{Z|X}(Z|X)P_{S|Z,X}(S|Z,X)f_{T|S,Z,X}(T|S,Z,X)f_{C|T,S,Z,X}(C|T,S,Z,X).$$

Notice that this is the same form as given in chapter 3, only now I have replaced Y with T (I redefined the outcome of interest to be T) and included the density of the time until censoring, conditional on all other variables. Under $\mathcal{A}$ the restrictions presented in chapter 3 and derived in appendix C.1 are identical here (of course replacing Y with T). In section E.1 of the Appendix I show that

$$f_{C|T,S,Z,X}(y|T, S = 1, Z = z, X) = g_z(y|x),$$

where $g_z(\cdot|\cdot)$ is the density of $C(z)$ given $S(z) = 1$ and $X = x$, for $z = 0, 1$.

It then follows that the likelihood given by (3.1) can be extended to give

$$\begin{aligned}
\mathcal{L}(\rho) \propto \prod_{i=1}^{N} \Big\{ \Big[& f_v(y_i|x_i;\eta_1)^{\delta_i} (1 - F_v(y_i|x_i;\eta_1))^{1-\delta_i} g_1(y_i|x_i)^{1-\delta_i} \\
& \times (1 - G_1(y_i|x_i))^{\delta_i} \theta_p(x_i;\mu) \int w(x_i, y;\beta,\alpha) f_p(y|x_i;\eta_0) dy \Big]^{s_i} \\
& \times \left[1 - \theta_p(x_i;\mu) \int w(x_i, y;\beta,\alpha) f_p(y|x_i;\eta_0) dy \right]^{1-s_i} \Big\}^{z_i} \\
& \times \Big\{ \Big[f_p(y_i|x_i;\eta_0)^{\delta_i} (1 - F_p(y_i|x_i;\eta_0))^{1-\delta_i} g_0(y_i|x_i)^{1-\delta_i} \\
& \times (1 - G_0(y_i|x_i))^{\delta_i} \theta_p(x_i;\mu) \Big]^{s_i} [1 - \theta_p(x_i;\mu)]^{1-s_i} \Big\}^{1-z_i},
\end{aligned}$$

where $G(y|x) \equiv \int_0^y g(s|x)ds$ and $F(y|x) \equiv \int_0^y f(s|x;\eta_0)ds$. We are not interested in estimating the censoring distribution, so it can be factored out of the product leaving us with

$$\begin{aligned}
\mathcal{L}(\rho) \propto & \prod_{i=1}^{N} \Big\{ \Big[f_v(y_i|x_i;\eta_1)^{\delta_i} (1 - F_v(y_i|x_i;\eta_1))^{1-\delta_i} \\
& \times \theta_p(x_i;\mu) \int w(x_i, y; \beta, \alpha) f_p(y|x_i;\eta_0) dy \Big]^{s_i} \\
& \times \Big[1 - \theta_p(x_i;\mu) \int w(x_i, y; \beta, \alpha) f_p(y|x_i;\eta_0) dy \Big]^{1-s_i} \Big\}^{z_i} \\
& \times \Big\{ \Big[f_p(y_i|x_i;\eta_0)^{\delta_i} (1 - F_p(y_i|x_i;\eta_0))^{1-\delta_i} \\
& \times \theta_p(x_i;\mu)]^{s_i} [1 - \theta_p(x_i;\mu)]^{1-s_i} \Big\}^{1-z_i}. \qquad (6.5)
\end{aligned}$$

The manner in which this likelihood is written implicitly assumes some parametric models. Specifically, distributional assumptions M.1-M.4 are made. (One could either assume M.4a or M.4b. And, of course, Y is replaced by T in our distributional assumptions. Standard choices for the conditional distribution of T include Weibull and exponential distributions. Notice that no assumptions are made with regards to the censoring distribution other than its conditional independence from T.)

One can maximize the likelihood with respect to parameters to obtain maximum likelihood estimates and then estimate the average causal effect of vaccine on the post-infection time-to-event outcome in the ai principal stratum in the same manner described in Chapter 3. And, as in Chapter 3, since F_p^{ai} and F_v^{ai} are fully specified, one can make other types of comparisons: for example, one can compare medians of the time-to-event outcomes at certain covariate levels or one could estimate $SCE(t, x)$ for some post-infection time t and some covariate level $X = x$.

As stated in Chapter 2, the primary reason for making parametric assumptions M.1-M.4 is to allow the inclusion of baseline covariates in the analysis. Of course, if one does not want to condition on baseline covariates, these parametric models still may be useful. The general strengths and weaknesses of parametric methods seen in Chapter 5 are similar here, although I have not formally studied parametric methods for time-to-event outcomes. The remaining sections of this chapter discuss estimation of $SCE(t)$ not conditioning on covariate values, making less distributional assumptions.

6.4 *Non-parametric Estimation: Sharp Bounds*

Perhaps we are unwilling to make any modeling assumptions, but would like to estimate bounds on $SCE(t)$. One can think of this as extending HHS to the time-to-event scenario. It is easy to show that under (6.2), bounds on $F_p^{ai}(t)$ are

$$F_p^{ai,U}(t) = min\left\{\frac{F_p(t)}{1-VE}, 1\right\}, \tag{6.6}$$

$$F_p^{ai,L}(t) = max\left\{\frac{F_p(t)-VE}{1-VE}, 0\right\}. \tag{6.7}$$

These expressions are equivalent to those given in HHS.

If in addition to (6.2), we also assume SUTVA, (6.1), and (6.4), then bounds for $F_p^{ai}(t)$ for $t < \tau$ can be consistently estimated with

$$\hat{F}_p^{ai,U}(t) = min\left\{\frac{\hat{F}_p(t)}{1-\widehat{VE}}, 1\right\},$$

$$\hat{F}_p^{ai,L}(t) = max\left\{\frac{\hat{F}_p(t)-\widehat{VE}}{1-\widehat{VE}}, 0\right\}$$

where $\widehat{VE} = min\{1 - n_v N_p/(n_p N_v), 0\}$ and $\hat{F}_p(t)$ is the Kaplan-Meier estimator of the distribution of the time-to-event outcome among those in the placebo arm who become infected. Since under monotonicity $F_v^{ai}(t) = F_v(t)$, estimated sharp bounds for $SCE(t)$ are $\hat{F}_p^{ai,L}(t) - \hat{F}_v(t)$ and $\hat{F}_p^{ai,U}(t) - \hat{F}_v(t)$, where $\hat{F}_v(t)$ is the Kaplan-Meier estimator among infected participants in the vaccine arm.

Under certain conditions, $\hat{F}_p^{ai,U}(t)$ and $\hat{F}_p^{ai,L}(t)$ are asymptotically normal. Specifically, $\widehat{VE}$ is asymptotically normal if $0 < VE < 1$, or equivalently $0 < P(S(0) = 1)$ and $P(S(1) = 1) < P(S(0) = 1)$. The Kaplan-Meier estimate, $\hat{F}_p(t)$ is asymptotically normal under the usual standard conditions (Fleming and Harrington). Therefore, $\hat{F}_p^{ai,U}(t)$ and $\hat{F}_p^{ai,L}(t)$ are asymptotically normal if, in addition to these conditions, $0 < F_p(t) < 1 - VE$ and $VE < F_p(t) < 1$, respectively. (These latter conditions are a result of (6.6) and (6.7). If these conditions are violated then by definition, $F_p^{ai,U}(t)$ and $F_p^{ai,L}(t)$ are 1 and 0, respectively, and hence estimates will not be asymptotically normal.) This follows from the asymptotic normality of $(\hat{F}_p(\cdot), \widehat{VE})$ (Appendix E.3), the Hadamard differentiability of the

maps $\hat{F}_p(t)/(1-\widehat{VE})$ and $(\hat{F}_p(t)-\widehat{VE})/(1-\widehat{VE})$, and the functional delta method (Andersen et al., 1992). Under these conditions, expressions for the asymptotic variance obtained via the functional delta method are

$$var\left(\hat{F}_p^{ai,U}(t)\right) = \left(\frac{p_0}{p_1}\right)^2 \sigma^2(t) + \left(\frac{F_p(t)}{p_1}\right)^2 \frac{p_0(1-p_0)}{N_p} + \left(\frac{F_p(t)p_0}{p_1^2}\right)^2 \frac{p_1(1-p_1)}{N_v}$$

$$var\left(\hat{F}_p^{ai,L}(t)\right) = \left(\frac{p_0}{p_1}\right)^2 \sigma^2(t) + \left(\frac{1-F_p(t)}{p_1}\right)^2 \frac{p_0(1-p_0)}{N_p} + \left(\frac{(1-F_p(t))p_0}{p_1^2}\right)^2 \frac{p_1(1-p_1)}{N_v},$$

where $\sigma^2(t)$ is the variance of the Kaplan-Meier estimate, $p_0 \equiv P(S=1|Z=0)$, and $p_1 \equiv p_0(1-VE) = P(S=1|Z=1)$. From these equations one may estimate variances in the usual manner, by plugging in parameter estimates.

Variances may alternatively be estimated using a standard bootstrap procedure. Specifically, from $(O_1, \cdots, O_N)$ sample with replacement N vectors O_i, creating $(O_1^*, \cdots, O_N^*)$. Compute $\widehat{SCE}^*(t)$ (or any other relevant function) based on the bootstrap sample $(O_1^*, \cdots, O_N^*)$. Repeat this process K times to obtain K bootstrapped estimates $\widehat{SCE}^*(t)$. The variance of $\widehat{SCE}(t)$ can be estimated by the variance of $\widehat{SCE}^*(t)$, which can then be used to obtain Wald-based confidence intervals of $SCE(t)$. $100(1-\alpha)\%$- level confidence intervals may also be constructed using the $\alpha/2$- and $1-\alpha/2$- quantiles of $\widehat{SCE}^*(t)$.

6.5 Semi-parametric Estimation: Extension of GBH

One can think of the previous section as an extension of HHS to time-to-event outcomes. Still leaving the distribution of the time-to-event outcome unspecified, rather than constructing bounds, one might want to perform sensitivity analyses over a more limited, plausible range of scenarios similar to the approach of GBH. One can think of this section as extending GBH to time-to-event outcomes.

The estimating equations of GBH when $n_v/N_v < n_p/N_p$ are

$$\psi(p_0, \alpha) = \begin{cases} \sum_{i=1}^N (1-Z_i)(S_i - p_0) \\ \sum_{i=1}^N Z_i \left(S_i - p_0 \int_0^\infty w(t;\alpha,\beta) d\hat{F}_p(t)\right). \end{cases} \tag{6.8}$$

Similar to section 6.4, a natural approach would be to use these same equations, only now estimating $F_p(t)$ with the Kaplan-Meier estimate. In practice, however, this approach

may not be feasible because $\hat{F}_p(t)$ may not be well defined for large t. Specifically, suppose data are only collected up to time τ after infection, at which time point every remaining subject is censored. (This will almost always be the scenario in these vaccine trials, because resources are not available to follow participants indefinitely.) In this situation we cannot estimate $F(t)$ for $t > \tau$. This is problematic because we estimate α using an estimating equation that requires computing the integral

$$\int_0^\infty w(y;\alpha,\beta)d\hat{F}_p(y) = \int_0^\tau w(y;\alpha,\beta)d\hat{F}_p(y) + \int_\tau^\infty w(y;\alpha,\beta)d\hat{F}_p(y).$$

This second integral cannot be computed if $\hat{F}_p(\tau) < 1$.

One way to fix this problem would be to assume some distributional form for the tail of $F_p(\cdot)$. (Using the parametric methods of Section 6.3, because $F_p(\cdot)$ is fully specified we do not have this problem.) However, we want to leave $F_p(\cdot)$ unspecified. Therefore, another approach would be to change the form of $w(\cdot)$, making it constant after time τ. For example, consider $w(\cdot)$ defined as follows:

$$w(t;\alpha,\beta) = \begin{cases} (1+exp(-\alpha-\beta t))^{-1} \text{ for } t \leq \tau \\ (1+exp(-\alpha-\beta\tau))^{-1} \text{ for } t > \tau. \end{cases} \tag{6.9}$$

Another choice for $w(\cdot)$ could be

$$w(t;\alpha,\beta) = \left(1+exp(-\alpha-\beta I_{\{t>t_0\}})\right)^{-1} \tag{6.10}$$

for some $t_0 \leq \tau$. Both (6.9) and (6.10) define $w(\cdot)$ with the expit function, but do so in a manner that $w(\cdot)$ is constant for $t > \tau$. If $w(\cdot)$ is constant for $t > \tau$, then one can write:

$$\int_0^\infty w(y;\alpha,\beta)d\hat{F}_p(y) = \int_0^\tau w(y;\alpha,\beta)d\hat{F}_p(y) + w(\tau;\alpha,\beta)\left(1-\hat{F}_p(\tau)\right).$$

Of course, these choices of $w(\cdot)$ have implications with regards to interpretation. Under (6.9) the interpretation of β is technically the following: Given infection in the placebo arm, the odds of infection if randomized to the vaccine arm for $T = t_1$ versus $T = t_2$ are $exp\left\{\beta\left[min\left(t_1,\tau\right) - min\left(t_2,\tau\right)\right]\right\}$. This more complex interpretation might appear troublesome. However, if τ is chosen as the maximum study follow-up time it is important to realize

that over the range of t for which there is data, β has the usual odds ratio interpretation. Under (6.10), β has a standard odds ratio interpretation, except now we have dichotomized failure times, assigning $w(\cdot)$ a particular value for $t > t_0$ and another for $t \leq t_0$.

With $w(\cdot)$ modeled by either (6.9) or (6.10), an extension of GBH using the Kaplan-Meier estimates for $F_p(t)$ is the semi-parametric MLE under $\mathcal{A}$. The proof is a simple extension of that given in Chapter 4, and is found in Appendix E.2. In Appendix E.3 I show that under these same assumptions and $0 < VE < 1$, $p_0 > 0$, and for a properly specified well-behaved $w(\cdot)$ (i.e., constant for $t > \tau$, bounded, and twice differentiable), $\hat{F}_p^{ai}(t)$ is consistent and asymptotically normal for $t \in (0, \tau]$. Therefore, $\widehat{SCE}(t)$ is also consistent and asymptotically normal (of course, this is also under the standard conditions that guarantee the asymptotic normality of $\hat{F}_v(t)$ and $\hat{F}_p(t)$). In fact, the results in the appendix actually show that $\hat{F}_p^{ai}(\cdot)$ converges to a Gaussian process, implying that so does $\widehat{SCE}(\cdot)$.

Because $\widehat{SCE}(t)$ is asymptotically normal, pointwise confidence intervals based on the bootstrap will be valid for large sample sizes. It is also possible to obtain an analytic form for the asymptotic variance of $\widehat{SCE}(t)$. This variance estimate relies on being able to approximate $\hat{F}_p(t)$ with a sum of i.i.d. random variables. Such an approximation can be obtained from Stute (1995). Using this result, one can augment (6.8) by including additional estimating equations:

$$\begin{aligned} &\sum_{i=1}^{N}(1 - Z_i)S_i\left(V_{1i} - F_p(t_1)\right) \\ &\qquad\qquad \vdots \\ &\sum_{i=1}^{N}(1 - Z_i)S_i\left(V_{ki} - F_p(t_k)\right), \end{aligned}$$

where k is the number of distinct failure times in the placebo arm, t_j is the jth ordered failure time, and for a specific j, V_{ji} are i.i.d. random variables for $i = 1, \cdots, N$. One can then estimate the variance of parameter estimates using a sandwich estimator type approach, and from there one can estimate the variance of $\widehat{SCE}(t)$ using the delta method. Details are given in the Appendix E.4.

6.6 *Simulations*

To evaluate the small sample performance of our estimators of $SCE(t)$ we conducted a 2×4 factorial simulation experiment, corresponding to generating data under $VE \equiv P(S(1) = 0|S(0) = 1) \approx 0.3$ or 0.6, and $\beta = 0.1, 0.2, 1$, or ∞. Each simulation generated 1000 vectors W according to the following steps:

Step 1. The first 500 vectors were set at $Z = 0$, the second 500 were set at $Z = 1$.

Step 2. $S(0)$ was drawn from a Bernoulli(p_0) distribution with $p_0 = 0.25$ (this choice yields an expected number of infections in the placebo arm of 125, which is typical for a Phase III vaccine trial).

Step 3. $T(0)$ was generated for all realizations with $S(0) = 1$ according to the distribution $F_p(t; \eta)$ with $F_p(\cdot)$ a Weibull distribution and $\eta =$ (shape = 0.5, scale = 25). This distribution was chosen to reflect the distribution of the time from infection to initiation of antiretroviral therapy (ART) in the VaxGen trial, for which approximately 50% of infected participants started ART by 24 months post-infection diagnosis.

Step 4. Given $T(0)$, for each realization with $Z = 1$ and $S(0) = 1$, $S(1)$ was drawn from a Bernoulli($w(T(0); \beta, \alpha)$) distribution. For $\beta = 0.1, 0.2$, and 1, $w(t; \beta, \alpha)$ was defined as in (6.9) with $\tau = 24$ months. To ensure that $VE \approx 0.3$, α was set at $-0.2, -0.9$, or -3.6, when β was set at $0.1, 0.2$, or 1 respectively; and to ensure that $VE \approx 0.6$, α was set at $-1.8, -3.4$, or -20, when $\beta = 0.1, 0.2$, or 1. For $\beta = \infty$, $w(t) = I_{\{t \geq q^{VE}\}}$ as discussed in Section 4.2, where $q^{0.3} = 3.18$ and $q^{0.6} = 21.0$.

Step 5. For the realizations with $Z = 1$ and $S(1) = 1$, $T(1)$ was set equal to $T(0)$.

Step 6. For the realizations with $S = 1$, C_1 was generated from a Weibull distribution with shape and scale parameters 3 and 35. Then C was set as $min(\tau, C_1)$, with $\tau = 24$ months.

Step 7. For all realizations with $S = 1$, Y was chosen as $min(C, T)$ and $\delta = I_{\{Y=T\}}$.

It is easily verifed that these steps result in simulating data under assumptions $\mathcal{A}$ and model (6.9).

For each simulated dataset, I computed $\hat{F}_p^{ai}(t)$ and $\widehat{SCE}(t)$ for $t = 24$ months, assuming the proper model for $w(\cdot)$ and the true value for β. Wald-based 95% confidence intervals

were constructed by estimating the standard error of estimates using both the bootstrap and asymptotic variance estimates for those simulations generated under $\beta = 0.1, 0.2, 1$. The variance was estimated using only the bootstrap for those simulations with $\beta = \infty$. Note that data were generated using time on a months scale. This is important for interpretation purposes, particularly with regards to β.

Table 6.1 reports the performance of $\widehat{SCE}(t)$ based on 1000 simulation iterations for $t = 24$ months. Since $\widehat{SCE}(t)$ is the difference between $\hat{F}_p^{ai}(t)$ and $\hat{F}_v(t)$, it is also sensible to look at the performance of the estimates of $F_p^{ai}(t)$. Table 6.2 does this using the same simulations and analyses reported in Table 6.1. In addition to presenting the coverage of the untransformed 95% confidence intervals for $F_p^{ai}(t)$, Table 6.2 also presents confidence intervals using the log-log transformation, which often results in better small sample coverage and ensures that confidence intervals do not extend beyond $[0, 1]$ (Dorey and Korn (1987); see Appendix E.4 for details).

In most cases, bias is minimal and coverage is good – perhaps just slightly below the nominal level using either the bootstrap or the asymptotic variance estimate. The only exceptions are when $VE \approx 0.6$ and β is large. The poor coverage probabilities here are essentially due to a boundary issue. Consider first the simulations with $\beta = \infty$. As discussed in Section 6.4, as $N \to \infty$, under the usual assumptions and if $VE < F_p(t)$ then $\hat{F}_p^{ai,L}(t)$ (which is equivalent to $\hat{F}_p^{ai}(t)$ with $\beta = \infty$) will be asymptotically normal and Wald-based confidence intervals will cover at their nominal level. However, in these simulations $F_p(t) = 0.6246$ for $t = 24$ months, which is very close to $VE = 0.6001$. Therefore, due to stochastic variation and our relatively small sample size, $\widehat{VE}$ is often greater than $\hat{F}_p(t)$ resulting in $\hat{F}_p^{ai}(t) = 0$ in nearly half of the simulations. Consequently, the distribution of $\hat{F}_p^{ai}(t)$ is far from normal; therefore these confidence intervals that assume normality have poor coverage. (It is worth noting that in this particular setting, Wald-based confidence intervals extended outside the $[0, 1]$ range; and log-log transformed confidence intervals could not be computed because the estimated value of $F_p^{ai,L}(t)$ was often 0 (see Appendix E.4). Interestingly, if one instead uses the 2.5- and 97.5- bootstrapped percentiles to construct confidence intervals for $F_p^{ai}(t)$ and $SCE(t)$, the coverage probabilities for the simulations with $VE \approx 0.6, \beta = \infty$ were 0.938 and 0.964, respectively. Analytic coverage probabilities when $VE \approx 0.6, \beta = \infty$

Table 6.1: Bias of estimates and coverage probability of Wald-based 95% confidence intervals for $SCE(t)$ with $t = 24$ months.

VE	β	Bias		Coverage Probability	
		Mean	Median	Bootstrap	Analytic
~ 0.3	0.1	-0.002	0.000	0.946	0.948
	0.2	-0.007	-0.007	0.943	0.949
	1	-0.006	-0.003	0.943	0.946
	∞	-0.007	0.005	0.948	0.953
~ 0.6	0.1	0.003	0.006	0.939	0.940
	0.2	0.013	0.015	0.945	0.945
	1	0.042	0.020	0.933	0.912
	∞	0.035	0.000	0.935	–

are not given in Table 6.1 because in many simulations, there were no failures in the vaccine arm.) The same logic explains why coverage and bias were also poor for the simulations with $VE \approx 0.6, \beta = 1$. Under these settings a value of $\beta = 1$ is quite large: for example, the odds of being infected in the vaccine arm given infection in the placebo arm from $t = 12$ and $t = 24$ (a difference of one year) multiplicatively increase $exp(12) \approx 163,000$! Hence, analyses based on the assumption that $\beta = 1$ are not too different from analyses assuming $\beta = \infty$. Again, the distribution of $\hat{F}_p^{ai}(t)$ is far from normal. Figure 6.2 shows a histogram of $\hat{F}_p^{ai}(t = 24)$ for the simulations with $VE \approx 0.6, \beta = 1$, as well as a similar histogram with $VE \approx 0.3, \beta = 1$ (where the method worked well) for purpose of comparison. Figure 6.2 also shows the true values of $F_p(t)$ and $F_p^{ai}(t)$ under these simulation settings. Notice how close $F_p^{ai}(t)$ is to 0 at $t = 24$ months.

Table 6.2: Bias of estimates and coverage probability of Wald-based 95% confidence intervals for $F_p^{ai}(t)$ with $t = 24$ months.

VE	β	Bias		Coverage Probability			
				Standard		Log-log transformation	
		Mean	Median	Bootstrap	Analytic	Bootstrap	Analytic
~ 0.3	0.1	-0.001	-0.002	0.939	0.941	0.946	0.942
	0.2	-0.006	-0.002	0.932	0.939	0.938	0.949
	1	0.001	-0.005	0.940	0.953	0.959	0.973
	∞	-0.007	0.002	0.934	0.947	0.957	0.967
~ 0.6	0.1	0.001	0.003	0.939	0.942	0.947	0.947
	0.2	0.015	0.011	0.929	0.927	0.938	0.937
	1	0.045	0.010	0.903	0.819	0.902	0.859
	∞	0.036	-0.016	0.889	0.948	–	–

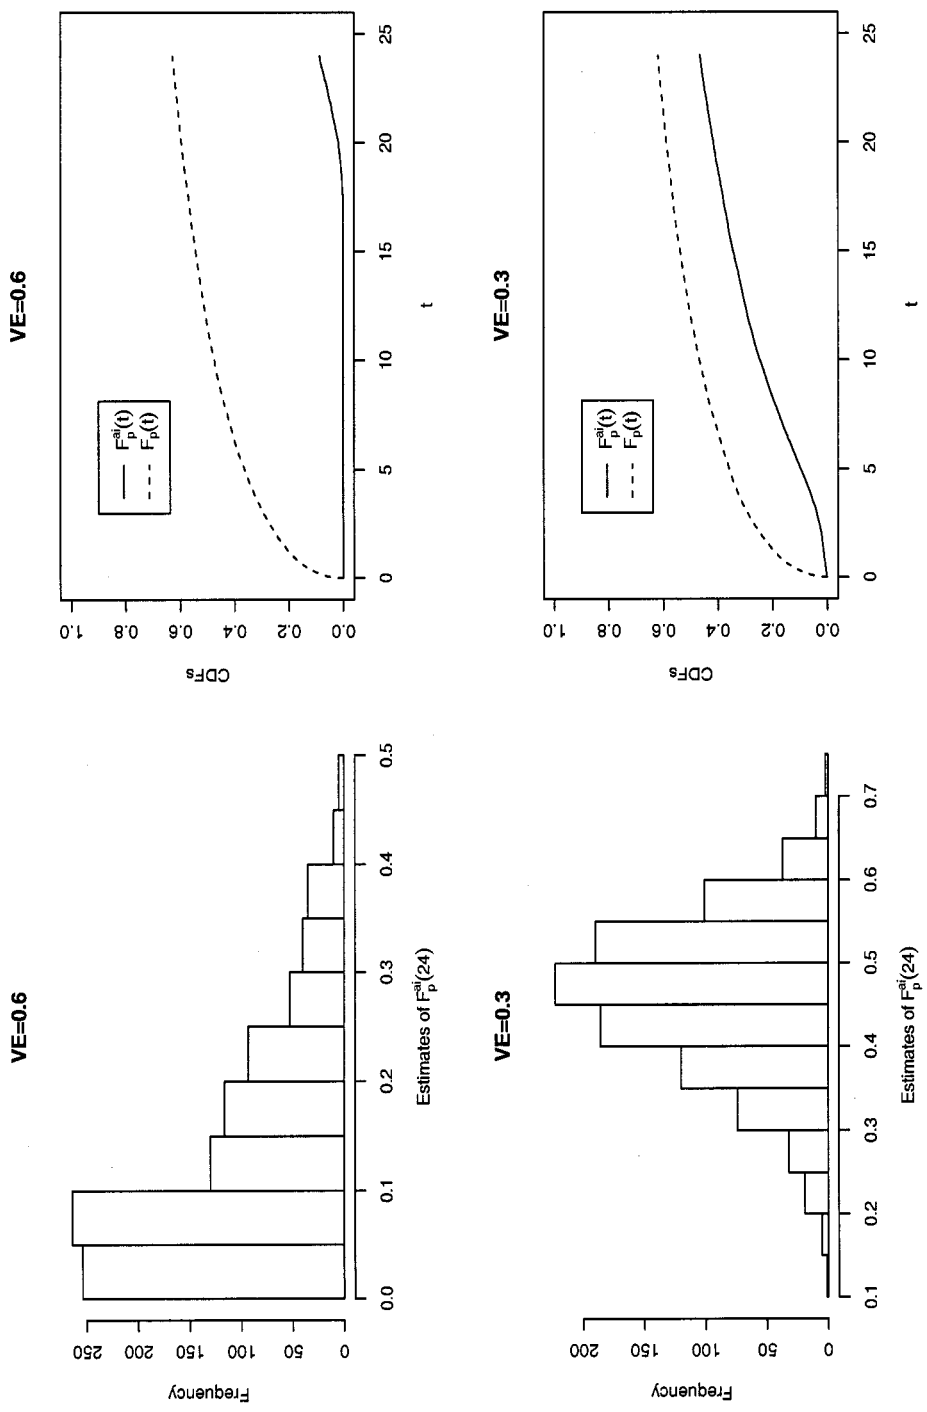

Figure 6.2: Histograms of $\hat{F}_p^{ai}(t)$ for $t = 24$ and plots of $F_p(t)$ and $F_p^{ai}(t)$ at different levels of VE for $\beta = 1$.

Chapter 7

EXAMPLE ANALYSIS

7.1 The VaxGen Trial

We illustrate our methods using data from VaxGen's Phase III trial of AIDSVAX B/B. This was a randomized, double-blind, placebo controlled trial conducted between 1998 and 2003. This study recruited 5,403 HIV negative, at risk individuals from 61 sites spanning large cities of North America and the Netherlands. The ratio of vaccine to placebo assignment was 2:1. The primary objective of the trial was to assess whether vaccination reduced the incidence of HIV infection, whereas the secondary objective of the trial was to see whether vaccination delayed disease progression for participants who became infected (Francis et al., 1998). Overall, the vaccine was not found to protect against HIV infection, although interaction tests suggested that the vaccine might partially prevent infection for non-whites and high risk subjects. Detailed study results are found in Flynn et al. (2005). Secondary analyses, looking at post-infection outcomes are reported in Gilbert et al. (in press, 2005). Most of these analyses were either of the intent-to-treat type (including all 5,403 participants) or analyses conditioning on infection (although the analyses that I will present in Section 7.2 were referenced in that paper). In this chapter we perform sensitivity analyses to compare post-infection outcomes between the vaccine and placebo arms among participants (overall and within covariate subgroups) who would have been infected regardless of randomization assignment.

There was presumably little interaction among study participants because the study was performed in many cities throughout North America and the Netherlands. Therefore SUTVA was thought to be reasonable. Because this trial was randomized and double blinded, individuals' behavior and exposure to HIV were expected to be the same regardless of treatment assignment. In addition, this vaccine was designed in such a way that it could

not mutate to become the virus. These facts justify A.1, although the assumption could be violated if the blinding was broken or if the vaccine induced susceptibility-enhancing immune responses for certain subjects (see Appendix A). Assumption A.1 implies that the rates of infection under vaccine cannot be greater than under placebo. Since the rates of infection were 241 of 3598 (6.7%) in the vaccine arm and 127 of 1805 (7.0%) in the placebo arm, a Wald test of $H_0^* : P(S = 0|Z = 1) \geq P(S = 0|Z = 0)$ fails to reject. Thus, monotonicity seems reasonable, although of course this is a population level test whereas monotonicity is assumed at the individual level (see Appendix B.4.2).

Of the 368 subjects who were infected during the trial, 347 enrolled in the post-infection phase of the study (225 in the vaccine arm). Infected subjects participating in this second phase of the study were followed at months 0, < 1, 1, 2, 4, 8, 16, 20, and 24 post-infection diagnosis. At each visit, the participants' viral load, CD4 count, and whether or not they had initiated antiretroviral therapy (ART) was recorded. Before randomization, the following baseline covariates were collected for each trial participant: race (white or non-white), risk score (an indicator of behavioral risk for becoming infected, taking integer values from 0-9), education (high school or less, more than high school), and age (18-30, 31-40, 41 and older).

7.2 Analysis of Set-Point Viral Load

We defined each participant's set-point viral load as the median of all $\log_{10}$ viral load measurements taken by the Month 2 visit and prior to initiation of antiretroviral therapy. (Results were comparable when means were used.) The viral load assay had lower and upper quantitative limits of 400 and 750,000 copies/ml, respectively. A subject's median viral load was defined exactly if the number of detectable viral load values exceeded the number undetectable. Otherwise, the median was left- or right-censored; 23 subjects had a left-censored median and 7 had a right-censored median.

We first estimated the average causal effect of vaccine on viral load (ACE), defined in Chapter 2. Figure 7.1 shows the estimated ACE for β in $[-3, 3]$ using A) the method of GBH with censored median viral load values set to either the lower or upper detection limit, and B) the ML estimators for the censored likelihood under $\mathcal{M}_a$ using (5.1)-(5.4) for the relevant

distributions and probabilities except not conditional on covariates. Results based on the two procedures are similar, with the method of GBH yielding slightly narrower confidence intervals presumably because censored values were truncated. The range of $[-3, 3]$ for β was chosen to reflect various possibilities about the relationship between the HIV viral load distributions in the always infected and protected strata. When β is negative, the always infected distribution is tilted to the left of the protected distribution and the opposite happens when β is positive. The tilting is more marked the larger the absolute value of β. Absolute values of β as large as 3 correspond to pronounced tilting (for example, with $\beta = 3$ the odds of being in the ai stratum versus the protected stratum for a one unit increase in viral load multiplicatively increase $exp(3) \approx 20$). Thus, for example, β would likely be positive if individuals with relatively strong immune systems tend to have lower viral loads when infected and if the vaccine is more likely to protect these individuals from infection. On the other hand, postulating negative values of β would be reasonable if it is believed that the vaccine prevents infection with relatively strong/virulent viruses better than it prevents infection with weaker/a-virulent viruses. For all β in $[-3, 3]$, the null hypothesis, H_0: $ACE = 0$, was not rejected. The ML estimator of ACE under $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5) was also computed for β in $[-3, 3]$. Results using this parameterization were similar for β in $[-1, 1.5]$. However, outside this range the assumed distribution of viral load in the infected placebos under $\mathcal{M}_b$ was not compatible with the observed distribution of viral loads among people infected in the placebo arm. Hence, presumably for β outside this range, the assumed model is misspecified (as discussed in Section 5.4). The test described in HHS (which corresponds to GBH setting $\beta = \pm\infty$) also did not reject H_0.

It is more compelling and interesting to estimate $ACE(x)$ for covariate levels under which vaccine efficacy is large; it is hypothesized that vaccine-induced antibodies that partially protect against infection also may have a beneficial effect to lower viral load. One can initially examine the relationship between baseline covariates and infection, vaccine efficacy, and viral load by putting all covariates in the model and running the analysis under $\mathcal{M}_a$ using (5.1)-(5.4) with $\beta = \beta_v = 0$. (Since this analysis was performed at $\beta = 0$, we have made the assumption that $Y(0) \amalg S(1)|S(0) = 1, X$; a no unmeasured confounders-type assumption. Sensitivity analyses to violations of this assumption are reported in the next

several paragraphs.) The resulting analysis is shown in Table 7.1. Results agree with those seen in Flynn et al. (2005): risk score is highly associated with the probability of infection and race and risk score tend to modify vaccine efficacy. There is an hypothesized biological mechanism for why the vaccine's ability to prevent infection and lower viral load might vary by risk behavior: natural exposure to HIV may "prime" the immune system, which is "boosted" by the vaccine to provide extra protection (Rowland-Jones et al., 1998). Hence, it is of interest to apply our methods to estimate the $ACE(x)$ at different risk levels and also in the non-white cohort.

The first row of Figure 7.2 contains sensitivity analyses of the estimated $ACE(x)$ at different risk scores based on the non-white cohort data only, under distributional assumptions $\mathcal{M}_a$ using (5.1)-(5.4) with x defined as risk score. For non-whites with risk scores of 2, if $\beta < -1/2$ then there is evidence that $ACE(x) > 0$. The second row of Figure 7.2 is similar to the first, only this analysis was performed under $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5) for the relevant distributions and probabilities. Under $\mathcal{M}_b$, if $\beta \geq 1$, there is evidence that the vaccine causes lower viral loads among non-whites with risk scores of 2, 3, and 4. In addition, based on the analysis under $\mathcal{M}_b$, if $\beta \geq 1.5$, H_0^{global} is rejected at the 0.05 level. There is insufficient evidence, however, to conclude that $ACE(x)$ varies by risk score; under $\mathcal{M}_b$, at all values of β in $[-3, 3]$, we fail to reject H_0^{inter} (P-values > 0.18).

The discrepancies between the first and second rows of Figure 7.2 are primarily due to different model choices for $f_v(y|x;\eta_1)$. In the first row under $\mathcal{M}_a$ using (5.3), $f_v(y|x;\eta_1)$ is modeled with six parameters, leading to a much more flexible estimate of $E(Y(1)|S(1) = 1, X = x)$ than under $\mathcal{M}_b$ using (5.5) (see Figure 7.3). When we perform the sensitivity analyses under $\mathcal{M}_a$ except setting β_v in (5.3) equal to 0 (as in Table 5.3 of our simulations, which is equivalent to modeling $f_v(y|x;\eta_1)$ with the normal linear model (5.5)), then the estimated $ACE(x)$ is very similar to the estimate under $\mathcal{M}_b$ (see Figure 7.2, row 3).

It should be noted that for large negative values of β, our parameterization choices under $\mathcal{M}_b$ are probably inadequate. The sharp shifts in the $ACE(x)$ (seen in the plot with Risk=4) for $\beta < -1.5$ are presumably due to either M.2 or M.4b being misspecified.

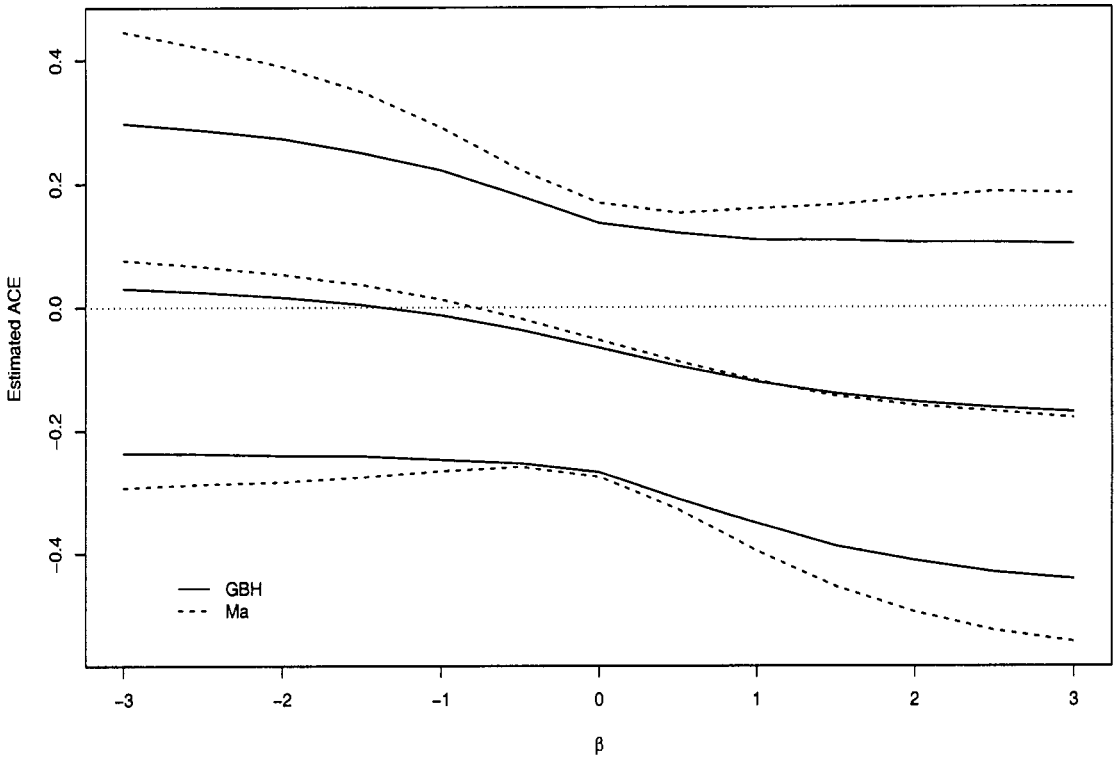

Figure 7.1: Sensitivity analysis estimates and 95% confidence intervals of the ACE for the complete VaxGen cohort using both the method of GBH and parameterizing with $\mathcal{M}_a$ accounting for censoring.

Table 7.1: Estimates (95% confidence intervals*) of relationship between covariates and infection, vaccine efficacy, and viral load, assuming $\mathcal{M}_a$ and model parameterizations (5.1)-(5.4) with $\beta = \beta_v = 0$.

		Covariates†				
	Intercept	White	Age 31-40	Age > 40	Risk Score	Education
Odds/odds ratios of infection in placebo arm (e^{μ})	0.05	0.75	0.88	0.55	1.90	0.84
	(0.03, 0.07)	(0.51, 1.13)	(0.64, 1.22)	(0.38, 0.79)	(1.72, 2.10)	(0.63, 1.12)
Odds/odds ratios of infection in vaccine arm	6.77	6.14	0.19	0.59	0.63	2.34
given infection in placebo arm (e^{α})	(0.50, 91.5)	(1.00, 37.7)	(0.02, 1.71)	(0.02, 14.8)	(0.39, 1.04)	(0.46, 12.0)
Viral load coefficients for infected placebos (λ)	3.93	0.20	0.14	0.36	-0.03	-0.00
	(3.53, 4.32)	(-0.14, 0.55)	(-0.18, 0.45)	(-0.02, 0.74)	(-0.11, 0.06)	(-0.28, 0.28)
Viral load coefficients for infected vaccinees (λ_v)	3.81	0.25	0.03	0.15	0.04	0.01
	(3.43, 4.18)	(-0.09, 0.59)	(-0.24, 0.29)	(-0.16, 0.46)	(-0.04, 0.12)	(-0.23, 0.24)

* Confidence intervals based on the expected information (confidence intervals based on the observed information were similar).

† White compared to non-white, Ages compared to Age $\leq$ 30, Education more than high school compared to high school or less.

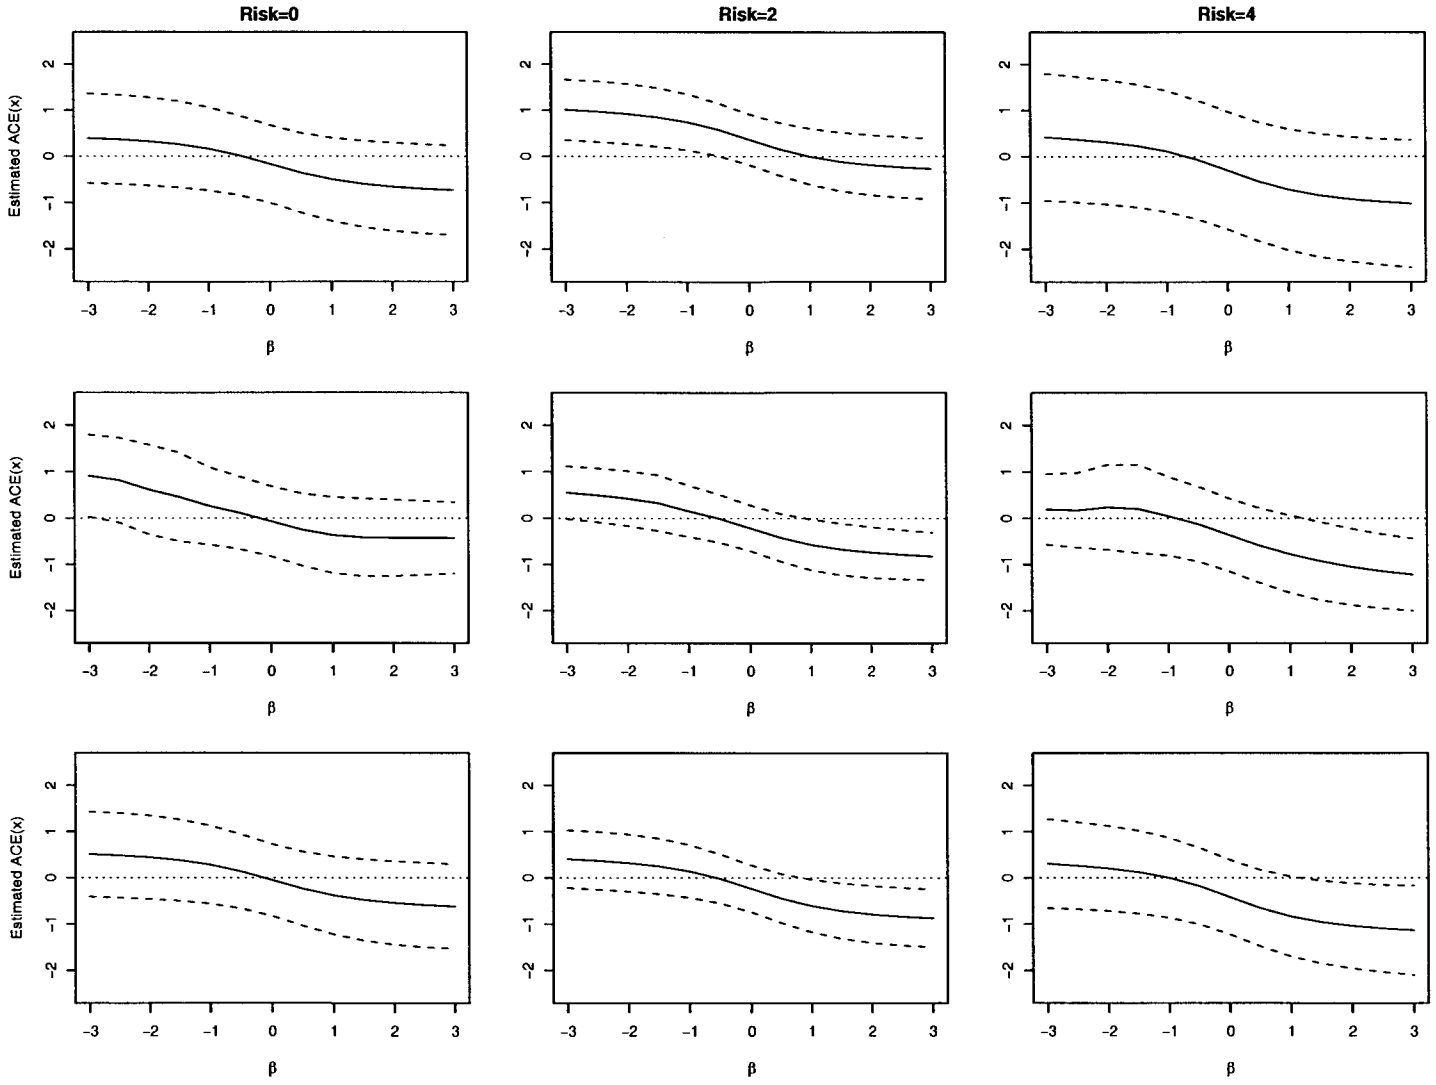

Figure 7.2: Sensitivity analyses of the $ACE(x)$ at different risk behavior levels in the non-white cohort. The first row is under distributional assumptions $\mathcal{M}_a$ using (5.1)-(5.4) for the relevant distributions and probabilities. The second row is under $\mathcal{M}_b$ using (5.1)-(5.2) and (5.5). The third row is under $\mathcal{M}_a$ using (5.1)-(5.4) except fixing $\beta_v = 0$ in (5.3).

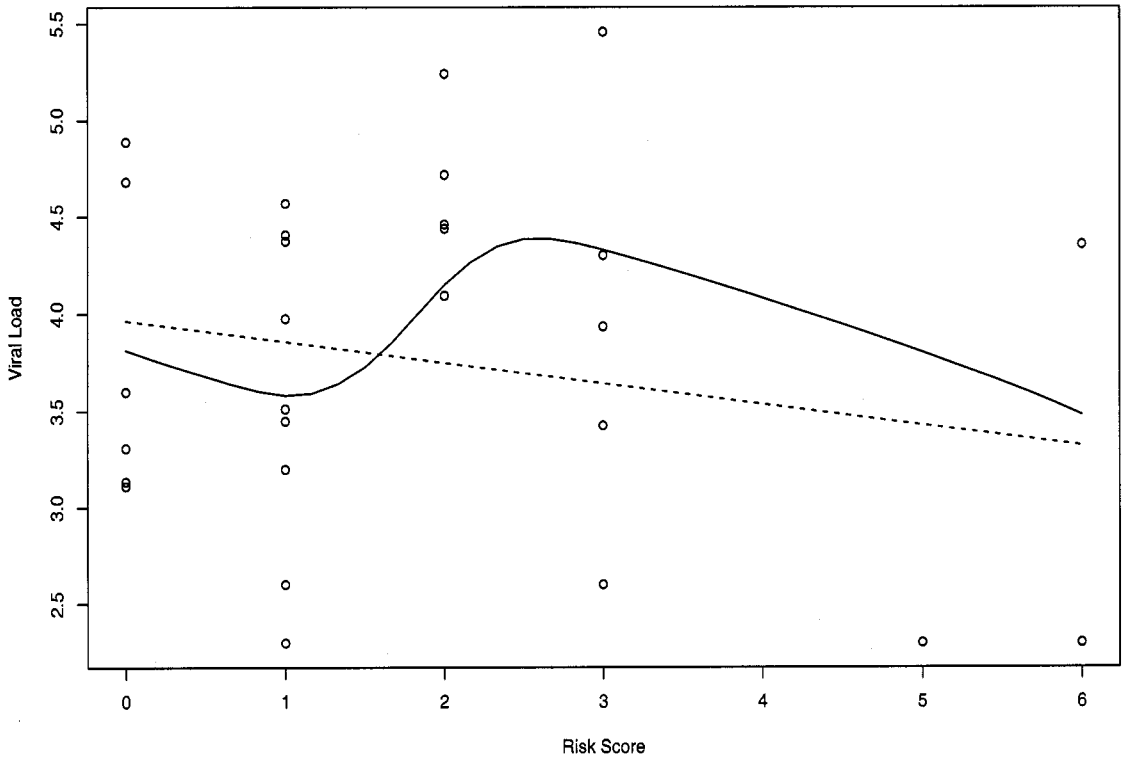

Figure 7.3: Viral loads and estimates for the expected viral load among infected non-whites in the vaccine arm conditional on risk score based on model (5.3) (solid line) and model (5.5) (dashed line).

7.3 *Time-to-Event Analyses*

One of the main post-infection analyses that VaxGen's protocol specified was an analysis looking at the time to a composite endpoint defined as the first occurrence of virologic failure ($\log_{10}$ viral load > 4) or the initiation of ART. The rationale for choosing this endpoint is discussed in Gilbert et al. (2003b). (Because of the great variability in viral load values immediately after infection, month < 1 values were not used for determining composite end points.) The protocol specified comparing the probability of this composite end point occurring in the vaccine and placebo arms 14 months after infection diagnosis. In this example, I also make the comparison 24 months after infection diagnosis. Another post-infection end point that is of clinical interest is the time from infection diagnosis to ART initiation; I also performed analyses investigating this outcome. For both end points, the assumption of independent censoring (6.4) seems reasonable and was made in previous analyses (Gilbert et al., in press, 2005).

My goal is to look at the causal effect of vaccination on each of these end points; specifically investigating whether the probability of the end point occurring by time $t = 14$ and 24 months differs by treatment among those who would have been infected regardless of treatment assignment. This calls for the sensitivity analysis methods described in Chapter 6, testing the hypothesis $H_0 : SCE(t) = 0$ for $t = 1.17$ and 2 years.

The three plots in Figure 7.4 show analyses looking at the time from infection diagnosis to the initiation of ART. Figure 7.4.A shows the Kaplan-Meier estimates for both the vaccine and placebo arms for the probability of not yet starting ART. The plot also includes the estimates of the upper and lower bounds of $F_p^{ai}(t)$, described in Section 6.4. Any semiparametric sensitivity analysis (described in Section 6.5) will result in estimates of $F_p^{ai}(t)$ between these bounds. Figures 7.4.B and 7.4.C are semiparametric sensitivity analyses looking at $SCE(t)$ for $t = 1.17$ and 2 years. The plots contain both the estimate for $SCE(t)$ and 95% Wald confidence intervals (constructed using the asymptotic variance approximation and the bootstrap, with 500 bootstrap replications). In these analyses (and all other sensitivity analyses in this section), I modeled the probability of infection in the vaccine arm given infection in the placebo arm, $w(t; \alpha, \beta)$, with (6.9). Therefore, $exp(\beta)$ has

an odds ratio interpretation: given infection in the placebo arm for a one year increase in the time from infection diagnosis to ART, the odds of infection if randomized to the vaccine arm multiplicatively increase $exp(\beta)$. (Of course, this interpretation is only for times less than τ, which in these analyses was taken to be the last observed placebo failure time, approximately 2 years post-infection diagnosis.) Before performing these analyses, I elicited a plausible range for the sensitivity parameter from a subject matter expert, Dr. Marc Gurwith of VaxGen. (My email eliciting the range and his response is given in Appendix F.) His best "guess" for a range for $exp(\beta)$ was 0.70 to 1.1, corresponding to β from -0.36 to 0.1. Figures 7.4.B and 7.4.C actually show the estimates of $SCE(t)$ over a much larger range, for β from -3 to 3. The open circles (and plus signs) in the plots represent the sharp bounds of $SCE(t)$ (and 95% confidence intervals for these bounds based on 500 bootstrap replications), corresponding to an analysis with $\beta = \pm\infty$. (Confidence intervals based on the analytic variance and bootstrap percentiles were similar.) Regardless of the range, it is clear from the Figure 7.4 that the vaccine is having no causal effect on the initiation of ART.

Figure 7.5 shows a similar analysis looking at the time from infection diagnosis to reaching the composite end point. Again, the vaccine appears to be having no effect on the time to event outcome.

As stated in Section 7.2, it may be of more interest to look at the effect of vaccination on these time-to-event outcomes for the non-white subgroup, where the vaccine appeared to partially protect against HIV infection. These analyses are shown in Figures 7.6 and 7.7 for the ART initiation and composite end points, respectively. Because the estimate for VE is much larger in the non-white cohort (0.469), the bounds for $F_p^{ai}(t)$ are farther apart, as seen in Figures 7.6.A and 7.7A. This is also reflected in plots B and C of both figures, as the estimates for $SCE(t)$ cover a wider range. Of course, the smaller sample size ($N = 914$ non-whites, of which 59 became infected) also inflates the length of the confidence intervals.

Notice that for the analyses looking at the ART end point in the non-white cohort, if $\beta > 0.5$ or 0 for $t = 1.17$ or 2 years, respectively, then $H_0 : SCE(t) = 0$ is rejected at the 0.05-level. Let us specifically consider the analysis for $t = 2$ years. This means that given infection in the placebo arm, if the odds of infection if randomized to vaccine are greater

for someone who has a longer time to the initiation of ART, then there is evidence that the vaccine is causing non-white participants to have a higher probability of starting ART by 2 years post-infection diagnosis. This would imply that among non-whites the vaccine is having a detrimental effect, causing more rapid post-infection progression. Interestingly, this value of β is just inside Dr. Gurwith's plausible range. Therefore, if we use his range, we are unable to make a conclusion about the causal effect of vaccination on this outcome because both the null and the alternative are favored in the range. On the other hand, for $t = 1.17$ years, the null is not rejected for all values of β in Dr. Gurwith's range. Hence, using his range there is insufficient evidence to conclude that the vaccine is having a causal effect on the probability of not yet starting ART 1.17 years after infection diagnosis.

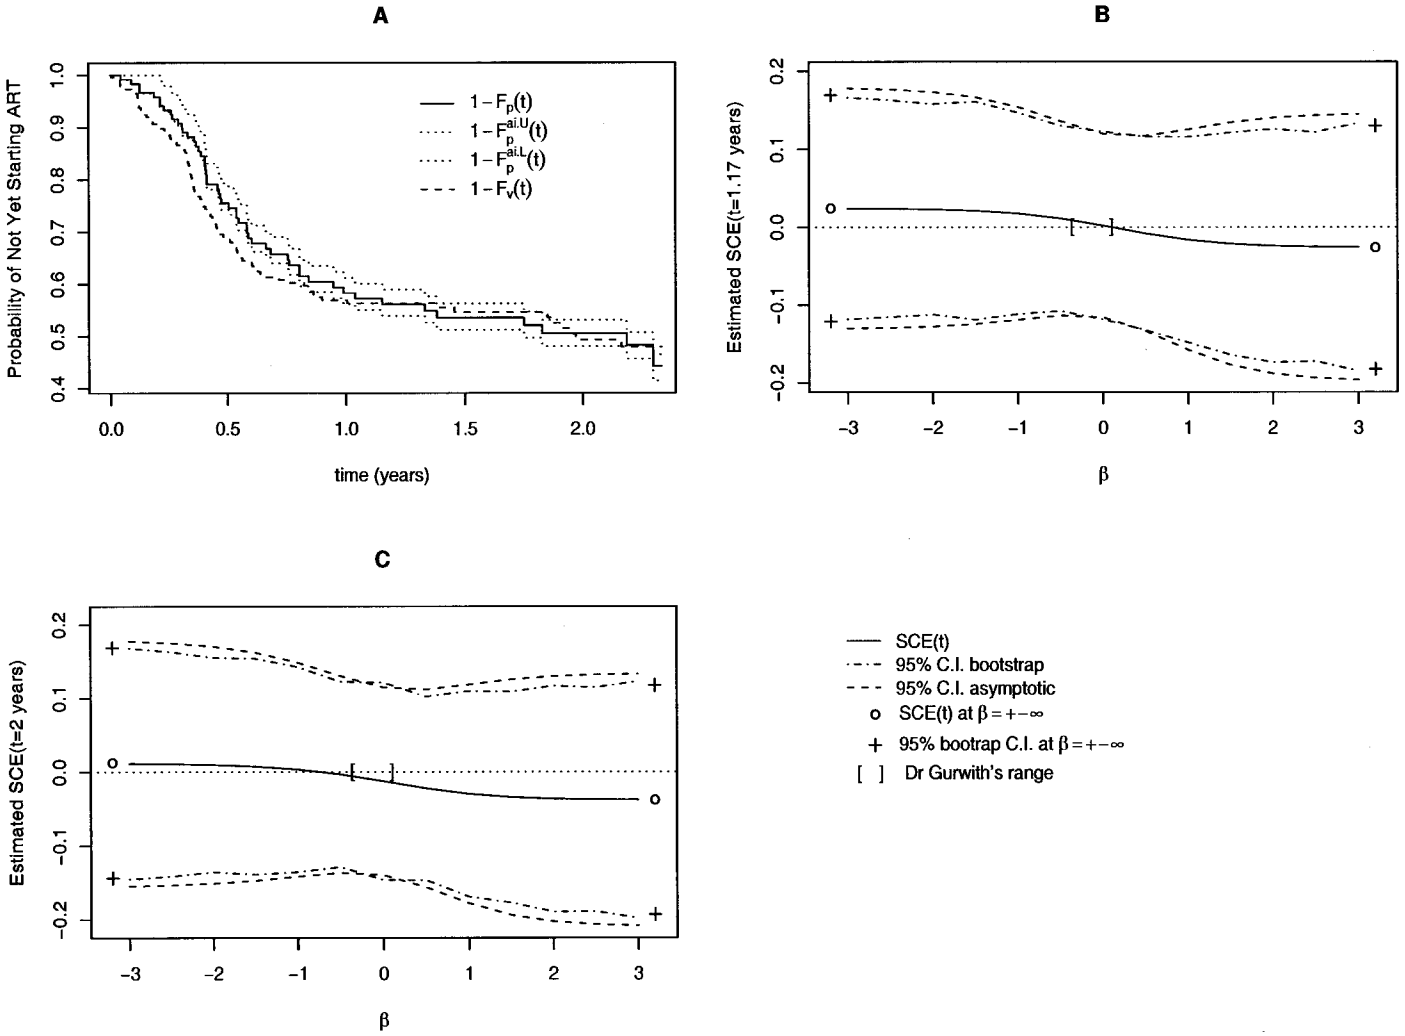

Figure 7.4: Sensitivity analyses of the effect of vaccination on the probability of not yet initiating ART.

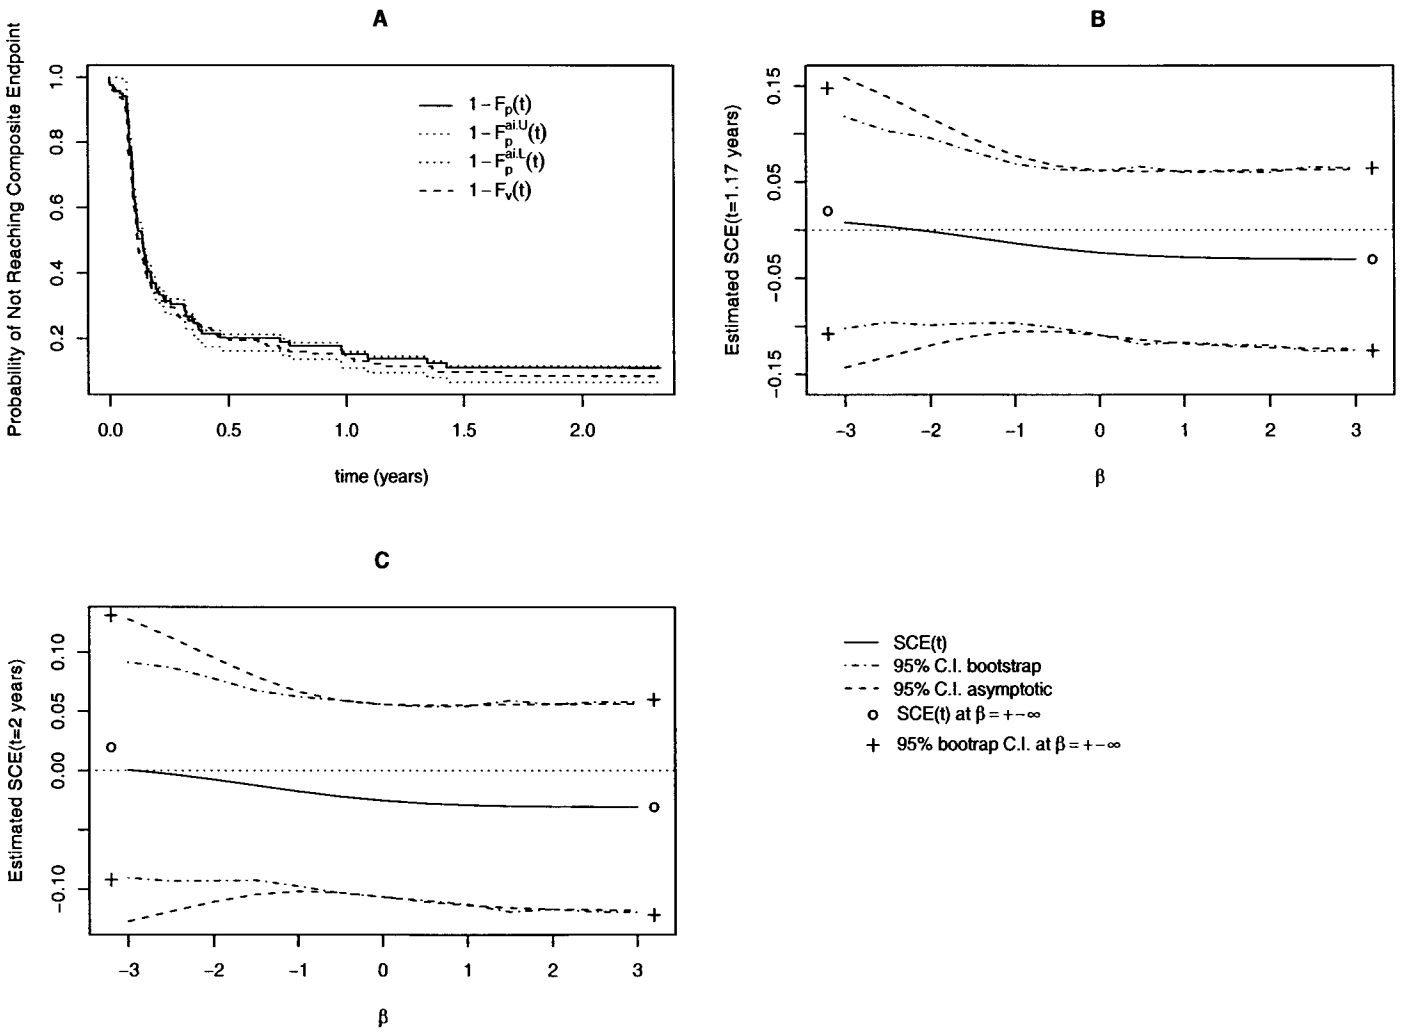

Figure 7.5: Sensitivity analyses of the effect of vaccination on the probability of not yet reaching the composite end point.

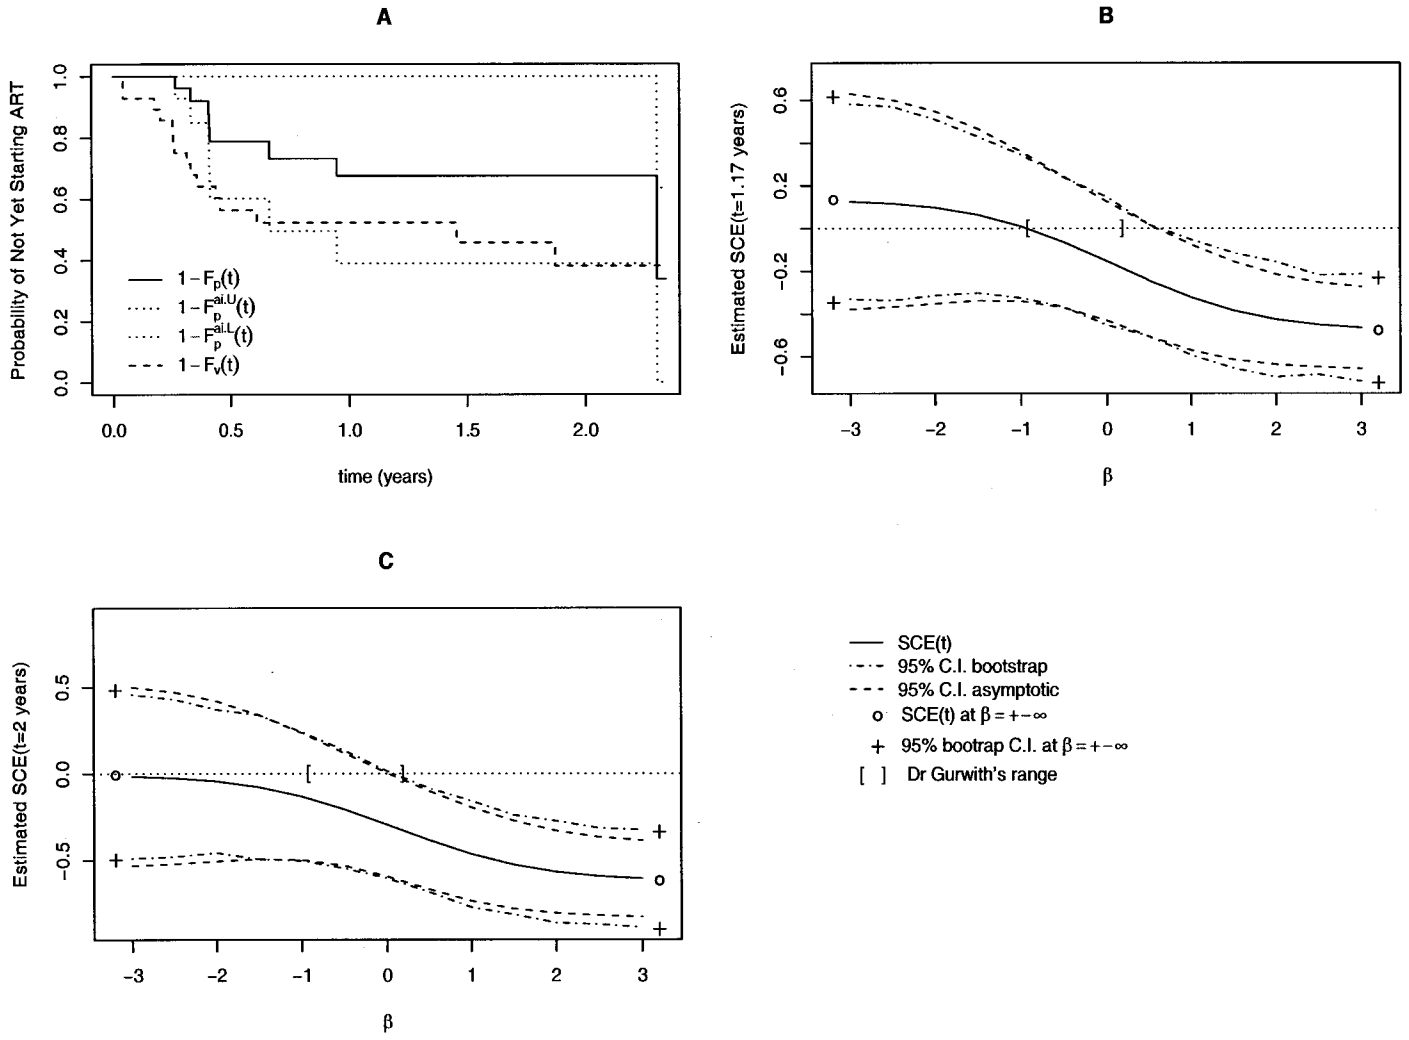

Figure 7.6: Sensitivity analyses of the effect of vaccination on the probability of not yet initiating ART in the non-white cohort.

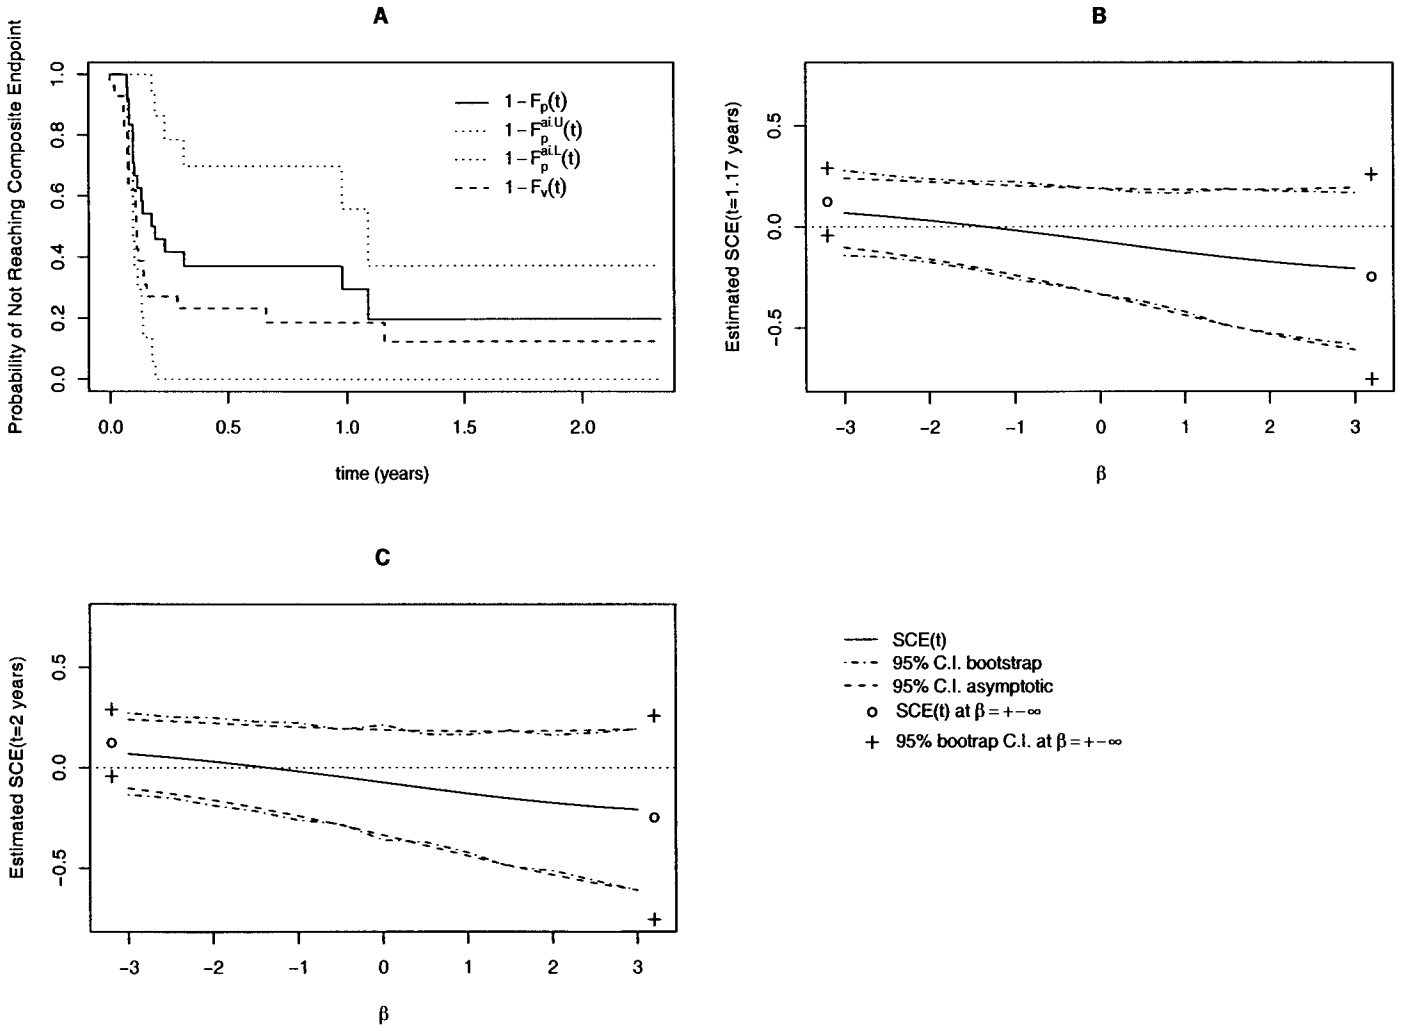

Figure 7.7: Sensitivity analyses of the effect of vaccination on the probability of not yet reaching the composite end point in the non-white cohort.

Chapter 8

DISCUSSION

8.1 Summary

As candidate vaccines continue to be developed and enter the clinical trial phase, there is particular interest in looking at the effect of vaccination on post-infection outcomes. In this dissertation I have proposed sensitivity analysis methods for evaluating the causal effect of vaccination on post-infection outcomes. One can think of this work as an extension of the methods of GBH. Under plausible assumptions on the potential outcomes, I first constructed a likelihood. This likelihood is general: the outcome of interest could be continuous or discrete, and the likelihood is easily extended to handle truncation or censoring. One may also incorporate different models for $w(\cdot)$, the probability of infection if randomized to vaccine given infection in the placebo arm. I then showed that the estimators of GBH in fact maximize this likelihood when one does not specify a distribution for the outcome variable and models $w(\cdot)$ with the expit function. This likelihood was particularly constructed to allow analyses conditioning on baseline covariates. I extensively studied these methods under different parameterizations for a continuous outcome conditioning on covariates. I then turned my attention to time-to-event outcomes, extending the methods of HHS and GBH to handle independently right censored data. I explored asymptotic properties of these estimators and derived an analytic variance expression. Finally, I applied the methods of this dissertation to investigate the causal effect of VaxGen's AIDSVAX B/B vaccine on post-randomization outcomes.

As I have presented this research, I frequently have been informed of other situations where these methods may apply. For example, in cancer research there was a major clinical trial looking at the effect of Finasteride on the development of prostate cancer. Finasteride was shown to lower the risk of prostate cancer. However, among those who developed

prostate cancer, those in the Finasteride arm had more severe levels of cancer (Thompson et al., 2003). For those who would have developed prostate cancer regardless of treatment assignment, does Finasteride cause higher-grade prostate cancer? Other possible applications of these methods include: mother-to-child HIV transmission studies, comparing outcomes in always surviving infants; quality of life studies, comparing some outcome between treatments among those patients who would have survived regardless of treatment assignment; and heart transplantation experiments, comparing liver toxicitiy between standard treatment and standard treatment plus an experimental drug designed to increase the success of transplantation, among those who would have had a successful transplantation in either arm. These are only a few examples. I believe there are many others. In general, these methods are applicable to intervention studies with post-randomization selection criteria where SUTVA and monotonicity are thought to hold.

It should be re-emphasized that the goal of this work is not to replace ITT-type analyses and standard analyses conditioning on post-infection outcomes with the causal analyses of this dissertation. Each type of analysis addresses a different question which may be the question of interest for a particular study. These sensitivity analyses are meant to supplement other analyses, as a tool to address whether or not vaccination directly affects post-randomization variables. In many studies, this may not be of interest. However, quite often this is exactly the question researchers want answered.

8.2 New Developments

We have had the privilege of collaborating with Yannis Jemiai and Andrea Rotnitzky for the past year-and-a-half. They have done some significant work that is closely related to this dissertation that is worth mentioning.

8.2.1 Semiparametric Estimation Conditional on Covariates

In Chapter 5, we studied estimation of the average causal effect of vaccination on viral load, conditional on baseline covariates. Modeling assumptions M.1-M.4 were made to allow the inclusion of covariates. However, we assumed more than was necessary to estimate $ACE(x)$.

Under the basic assumptions (2.1), (2.2), A.1, and A.2, Jemiai et al. (submitted, 2005) show that to estimate $ACE(x)$ one only needs to additionally assume M.2 and forms for the marginal viral load means in the ai stratum, i.e., the form of $m(z, X; \gamma) \equiv E(Y(z)|S(0) = S(1) = 1, X), z \in 0, 1$, where $m(z, x; \gamma)$ is some smooth function of a parameter vector γ. Notice that these assumptions are similar to parameterization $\mathcal{M}_b$ except no modeling assumptions are made with regards to the form of $\theta_p(x)$ (the probability of infection given covariates when assigned placebo), and instead of modeling viral load densities in the ai stratum, one is modeling the conditional expectations of viral loads in the ai stratum. This paper then develops a class of estimating equations, deriving locally semiparametric efficient estimators of γ.

Of course, there are advantages to using this semiparametric approach. The major advantage is that it gives inferences (in particular, tests of the null hypothesis $H_0 : ACE\,(x)$ for all x), that are valid under less stringent assumptions about the distribution of the outcome in the always infected population. Hence, these semiparametric methods are harder to misspecify. Unfortunately, however, these semiparametric methods also suffer from incompatibility problems for certain values of the sensitivity parameter β, due to it being identified (similar to the parametric methods under $\mathcal{M}_b$, as discussed in Chapter 5). But these problems may be less of an issue because one is only specifying conditional means rather than entire distributions. Advantages to the parametric, likelihood-based methods of Chapter 5 include the usual advantages: if properly specified, the parametric methods will have more power; since one has specified the entire distribution, one can make other comparisons between treatment arms (in addition to comparing conditional means); and by making parametric assumptions it is easy to account for the fact that viral load values are actually truncated.

8.2.2 Asymptotic Behavior of ACE Under Monotonicity near the Boundary

Throughout this paper, when discussing asymptotics we have made the assumption that $VE > 0$. Under monotonicity, this is equivalent to assuming that the protected principal stratum is non-empty. Even though we have assumed $VE > 0$, we have provided estimators

for the case when $\widehat{VE} = 1 - (n_v/N_v)/(n_p/N_p) \leq 0$ (which, of course, is entirely possible when $VE > 0$); under the assumption of monotonicity, these estimators are simply the estimators assuming no selection bias.

An important result reported in Jemiai and Rotnitzky (submitted, 2005) is the behavior of the estimator of GBH when VE is close to 0. Jemiai and Rotnitzky show that convergence of GBH's estimator is not uniform for all distributions allowed under the assumption that $VE > 0$. Therefore, for any given sample size, one can find a law such that the distribution of the estimator is not centered at the truth nor is it normal. As a consequence, this implies that near the boundary, Wald-based confidence intervals may not have their prescribed coverage.

This result has implications to the work of this dissertation, as it applies to all estimators presented here, both parametric and semiparametric. In particular, this fact means we must be cautious interpreting the analyses presented in Chapter 7. In the VaxGen trial, using the entire cohort the estimate for VE is 0.048. If VE does in fact equal 0.05, then as the sample size goes to infinity, everything is well-behaved. But is the VaxGen sample large enough? Jemiai and Rotnitzky performed simulations patterned after the VaxGen viral load analysis (shown in Figure 7.1) and demonstrated that for large values of β, coverage might be poor (for $\beta = 1, 2.5$ they reported coverage of 0.909 and 0.827, respectively, for Wald-based 95% confidence intervals of $E(Y(0)|S(0) = S(1) = 1)$ using the bootstrap variance estimate; and 0.948 and 0.891, respectively, using the analytic variance estimate). This result causes one to question the validity of the confidence intervals presented in Figures 7.1, 7.4, and 7.6. I also performed a small simulation study, designed to resemble the VaxGen time to initiation of ART data, with true $VE = 0.05$. This simulation was performed for $\beta = 1.2$ (based on a years time scale) in a manner identical to those simulations described in Section 6.6. From these simulations, Wald-based 95% confidence intervals of $SCE(t = 2$ years) using the bootstrap and analytic variance estimates had good coverage, 0.941 and 0.938, respectively. Of course, in the VaxGen trial we do not know the true VE. It may be misleading to read too much into any of these simulations, as coverage may be better or worse than seen here. In any case, it is important to recognize this potential problem as one approaches the boundary $VE = 0$. (This is not an issue for the analyses in the non-white cohort where

$\widehat{VE} = 0.469$, although one could be concerned about the smaller sample sizes.)

An important area of future research, therefore, is trying to overcome this problem – examining approaches for creating confidence intervals near the boundary. We have toyed with bootstrap-based confidence interval approaches proposed by Efron and Tibshirani (1998); Andrea Rotnitzky believes this problem is similar to one encountered by Robins (2004), where he suggests an approach for computing confidence intervals based on the inversion of efficient semiparametric score tests. This warrants further investigation because as seen from the VaxGen trial, VE close to 0 seems to be the current state of HIV vaccines.

8.3 Future Research

I believe that there are many other interesting questions either stemming from this dissertation research or closely related to it. In this section I briefly mention a few ideas for future study.

First, with regards to time-to-event analyses: This dissertation discusses methods for testing $H_0 : SCE(t) = 0$ for a given t. I would like to develop causal methods that allow one to test this hypothesis over the entire range of t. For example, by showing that $\widehat{SCE}(\cdot)$ converges to a Gaussian process (when $VE > 0$), we had hoped to be able to construct simultaneous confidence bands for $SCE(\cdot)$ using the methods of Parzen et al. (1997). Creating confidence bands for $SCE(t)$, however, is not a simple application of Parzen et al. because their method requires the use of the failure indicator variable (δ_i). In our case, we need δ_i for individuals in the ai stratum, who are not known. We could weight the δ_i by their probability of corresponding to an individual in the ai stratum, but this would presumably have implications on the variance and needs to be further studied. A related problem would be the development of a "log-rank"-type statistic. Another important extension would be to allow the inclusion of baseline covariates in time-to-event analyses, yet to not specify the distribution of the time-to-event outcome (as was done in Section 6.3). A natural first thought would be to use a Cox proportional hazards type model. Solving this problem may also give us a log-rank-type statistic, as the log-rank test is equivalent to a test of no treatment effect under a Cox proportional hazards model with treatment as the

only covariate.

Another important research direction is extending these methods to incorporate repeated measures after infection. For example, it is of particular scientific interest to study viral load trajectories post-infection. Several different modeling choices must be made: Are we more interested in a marginal means type model or a conditional model? How should we go about performing sensitivity analyses, i.e., should we model the probability of being in the ai stratum using the initial outcomes $Y_{i1}(0)$ of infected placebos, or the slope of outcomes, or the intercept, etc? Ideally any new method will be flexible enough to handle a variety of choices for $w(\cdot)$. In these HIV vaccine trials, when comparing viral load trajectories there are also other complexities beyond the selection bias issue. One such challenge is that ART has a tendency to lower viral load, so it is not fair to compare viral load trajectories between individuals after one of them has commenced ART. One might therefore remove post-ART observations. However, the missing data mechanism could bias results. (For this reason, it would be helpful to set strict guidelines for ART initiation, for example a viral load above a certain threshold, allowing one to model the missing data mechanism based on observed data.) In practice, any longitudinal data methods accounting for post-randomization selection that investigate viral load trajectories must account for missing data due to the onset of ART.

Throughout this dissertation I have assumed that infection status is known for each individual. A simple (hopefully) extension of these methods would be to consider the case where S is not known, but must be estimated, perhaps using a Kaplan-Meier estimate. For example, S_i could be defined as infection by 2 years after randomization. Some individuals may be censored before this time. In that case one might want to estimate VE as one minus the ratio of infection probabilities at 2 years, estimated using Kaplan-Meier estimates. A more complex extension would be to somehow account for the infection diagnosis time. For instance, those people who become infected early in the study may have different post-infection characteristics than those infected late in the study.

In application there are, of course, many other issues one should at least acknowledge, that may warrant further study. For example, infection is diagnosed at different times for different participants: one participant may be diagnosed 5 days post-infection, another 5

months post-infection. How could one perform sensitivity analyses that account for this unequal diagnosis time as well as incorporating the selection bias techniques discussed here? Another example: The key assumption of SUTVA (that the potential outcomes for trial participants are independent of treatment assignment to any other participant) could be violated in vaccine studies performed when participants are interacting. It might be interesting to develop methods that allow one to investigate the sensitivity of analyses to minor violations in the assumption of SUTVA. And how could one perform causal analyses looking at post-randomization variables in community randomized trials, where SUTVA is clearly violated?

Another related problem deals with investigating a vaccine's efficacy against a particular strain of HIV, say A. If the vaccine partially protects against another circulating strain, strain B, then one might underestimate the efficacy of the vaccine against strain A because those protected from infection with B by the vaccine are retained in the pool of subjects susceptible to A, thereby creating greater exposure to A. How could one adjust for this potential bias in the estimate of the vaccine efficacy against strain A?

I have heard some concerns (not just from my work, but more generally) about eliciting sensitivity parameters. It might be fun to study this a little further. For example, I am curious about the different ranges different subject-matter experts would select for these sensitivity analyses. Also, it would be interesting to randomly assign different techniques for eliciting sensitivity parameters to different experts, seeing how this affects the selected range.

In its fullest sense, a sensitivity analysis includes not just varying β, but also varying the form of $w(x, y; \beta)$. Different models for $w(\cdot)$ or perhaps entirely different modeling approaches could be explored. For example, in addition to models $\mathcal{M}_a$ and $\mathcal{M}_b$, one could imagine a third parametric model, say $\mathcal{M}_c$, based on assumptions (1), (2), and A.1, modeling $f_v(\cdot)$, $f_p^{ai}(\cdot)$, and instead of $w(\cdot)$, modeling $f_p^{prot}(\cdot)$ (the distribution of viral loads for placebos in the protected principal stratum), with a sensitivity parameter specifying some discrepancy between $f_p^{prot}(\cdot)$ and $f_p^{ai}(\cdot)$. (This is very similar to model $\mathcal{M}_b$ (where we specified $f_p^{ai}(\cdot)$ and $w(\cdot)$, inducing $f_p^{prot}(\cdot)$).) There are surely advantages and disadvantages to models of this type. For example, in contrast to $\mathcal{M}_b$, under $\mathcal{M}_c$ there would no longer be the

requirement that $w(\cdot) > 0$ for all x and y. However, similar to $\mathcal{M}_b$, care would still have to be taken to make distributional assumptions compatible with the observed data for the entire range of β. This general type of model could also be considered for semiparametric analyses. To be useful, any modeling approach must use a sensitivity parameter with a meaningful interpretation, allowing one to intuitively choose a range over which to perform the analyses.

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*, 158:596–605, 2003.

P.K. Andersen, O. Borgan, R.D. Gill, and N. Keiding. *Statistical Models Based on Counting Processes.* Springer, Berlin, 1992.

J.D. Angrist, G.W. Imbens, and D.B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91:444–455, 1996.

Anonymous. Approaches to the development of broadly protective HIV vaccines: challenges posed by the genetic, biological and antigenic variability of HIV-1: Report from a meeting of the WHO-UNAIDS Vaccine Advisory Committee Geneva, 21-23 February 2000. *AIDS*, 15:W1–W25, 2001.

D.H. Barouch, J. Kunstman, H.J. Kuroda, J.E. Schmitz, S. Santra, F.W. Peyerl, G.R. Krivulka, K. Beaudry, M.A. Lifton, D.A. Gorgone, D.C. Montefiori, M.G. Lewis, S.M. Wolinsky, and N.L. Letvin. Eventual AIDS vaccine failure in a rhesus monkey by viral escape from cytotoxic T lymphocytes. *Nature*, 415:335–339, 2002.

G. Casella and S.P. Schwartz. Comment on 'Causal inference without counterfactuals' by A.P. Dawid. *Journal of the American Statistical Association*, 95:425–427, 2000.

R.M. Chisholm. The contrary-to-fact conditional. *Mind*, 55(19):290–307, 1946.

M.L. Clements-Mann. Lessons for AIDS vaccine development from non-AIDS vaccines. *AIDS Research and Human Retroviruses*, 14:S197–S203, 1998.

K.A. Crandall (ed). *The Evolution of HIV*. Johns Hopkins University Press, Baltimore, 1999.

A.P. Dawid. Conditional independence in statistical theory. *Journal of the Royal Statistical Society Series B*, 41:1–31, 1979.

A.P. Dawid. Causal inference without counterfactuals. *Journal of the American Statistical Association*, 95:407–424, 2000.

F.J. Dorey and E.L. Korn. Effective sample sizes for confidence intervals for survival probabilities. *Statistics in Medicine*, 6:679–687, 1987.

B. Efron and R. Tibshirani. The problem of regions. *The Annals of Statistics*, 26: 1687–1718, 1998.

T.R. Fleming and D.L. DeMets. Surrogate end points in clinical trials: Are we being misled? *Annals of Internal Medicine*, 125:605–613, 1996.

T.R. Fleming and D.P. Harrington. *Counting Processes and Survival Analysis*. Wiley, New York, 1991.

N.M. Flynn, D.N. Forthal, C.D. Harro, F.N. Judson, K.H. Mayer, M.F. Para, and the rgp120 HIV Vaccine Study Group. Placebo-controlled phase 3 trial of a recombinant glycoprotein 120 vaccine to prevent HIV-1 infection. *Journal of Infectious Diseases*, 191:647–649, 2005.

D.P. Francis, T. Gregory, M.J. McElrath, R.B. Belshe, G.J. Gorse, S. Migasena, D. Kitayaporn, P. Pitisuttitham, T. Matthews, D.H. Schwartz, and P.W. Berman. Advancing AIDSVAX to phase 3. Safety, immunogenicity, and plans for phase 3. *AIDS Research and Human Retroviruses*, pages S325–S331, 1998.

C.E. Frangakis and D.B. Rubin. Principal stratification in causal inference. *Biometrics*, 58:21–29, 2002.

P.B. Gilbert, M.L. Ackers, P.W. Berman, D.P. Francis, V. Popovic, D.J. Hu, W.L. Heyward, F. Sinangil, B.E. Shepherd, and M. Gurwith. HIV-1 virologic and immunologic progression and antiretroviral therapy initiation among HIV-1-infected participants in an efficacy trial of recombinant glycoprotein 120 vaccine. *Journal of Infectious Diseases*, in press, 2005.

P.B. Gilbert, R.J. Bosch, and M.G. Hudgens. Sensitivity analysis for the assessment of causal vaccine effects on viral load in HIV vaccine trials. *Biometrics*, 59:531–541, 2003a.

P.B. Gilbert, V.G. De Gruttola, M.G. Hudgens, S.G. Self, S. Hammer, and L. Corey. What constitutes efficacy for a human immunodeficiency virus vaccine that ameliorates viremia: Issues involving surrogate end points in phase 3 trials. *Journal of Infectious Diseases*, 188:179–193, 2003b.

P.B. Gilbert, G.J. Hanna, V.G. De Gruttola, J. Martinez-Picado, D.R. Kuritzkes, V.A. Johnson, D.D. Richman, and R.T. D'Aquila. Comparative analysis of HIV type 1 genotypic resistance across antiretroviral trial treatment regimens. *AIDS Research and Human Retroviruses*, 16:1325–1336, 2000.

P.B. Gilbert, M.L. Peterson, D. Follmann, M.G. Hudgens, D.P. Francis, M. Gurwith, W.L. Heyward, D.V. Jobes, V. Popovic, S.G. Self, F. Sinangil, D. Burke, and P.W. Berman. Correlation between immunologic responses to a recombinant glycoprotein 120 vaccine and incidence of HIV-1 infection in a phase 3 HIV-1 preventive vaccine trial. *Journal of Infectious Diseases*, 191:666–677, 2005.

B.S. Graham. Clinical trials of HIV vaccines. *Annual Review of Medicine*, 53:207–221, 2002.

R.H. Gray, M.J. Wawer, R. Brookmeyer, N.K. Sewankambo, D. Serwadda, F. Wabwire-Mangen, T. Lutalo, X. Li, T. vanCott, and T.C. Quinn. Probability of HIV-1 transmission per coital act in monogamous, heterosexual, HIV-1 discordant couples in Rakai, Uganda. *Lancet*, 357:1149–1153, 2001.

M.E. Halloran and C.J. Struchiner. Causal inference for infectious diseases. *Epidemiology*, 6:142–151, 1995.

D. Hayden, D.K. Pauler, and D. Schoenfeld. An estimator for treatment comparisons among survivors in randomized trials. *Biometrics*, 61:305–310, 2005.

V.M. Hirsch, T.R. Fuerst, G. Sutter, M.W. Carroll, L.C. Yang, S. Goldstein, M. Piatak, W.R. Elkins, W.G. Alvord, D.C. Montefiori, B. Moss, and J.D. Lifson. Patterns of viral replication correlate with outcome in simian immunodeficiency virus (SIV)-infected macaques: effect of prior immunization with a trivalent SIV vaccine in modified vaccinia virus Ankara. *Journal of Virology*, 70:3741–3752, 1996.

P.W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81:945–960, 1986.

M.G. Hudgens and M.E. Halloran. Causal vaccine effects on binary post-infection outcomes. *Journal of the American Statistical Association*, in revision, 2005.

M.G. Hudgens, A. Hoering, and S.G. Self. On the analysis of viral load endpoints in HIV vaccine trials. *Statistics in Medicine*, 22:2281–2298, 2003.

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*, 95:1032–1043, 2001.

Y. Jemiai and A. Rotnitzky. Asymptotic properties of an estimator of treatment effects on an outcome only existing if a post-randomization event has occurred. *Biometrics*, submitted, 2005.

Y. Jemiai, A. Rotnitzky, B.E. Shepherd, and P.B. Gilbert. Semiparametric estimation of treatment effects on an outcome existing only if a post-randomization event occurs. *Journal of the Royal Statistical Society, Series B*, submitted, 2005.

J.D. Kalbfleish and R.L. Prentice. *The Statistical Analysis of Failure Time Data.* Wiley, New York, 1980.

E. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53:457–481, 1958.

J.D. Lifson and M.A. Martin. One step forward, one step back. *Nature*, 415:272–273, 2002.

R.J.A. Little. Modeling the drop-out mechanism in repeated-measures studies. *Journal of the American Statistical Association*, 90:1112–1121, 1995.

V.V. Lukashov, J. Goudsmit, and W.A. Paxton. The genetic diversity of HIV-1 and its implications for vaccine development. In F. Wong-Staal and R.C. Gallo, editors, *AIDS Vaccine Research*, chapter 3, pages 93–120. Dekker, 2002.

R.H. Lyles, A. Munoz, T.E. Yamashita, H. Bazmi, R. Detels, C.R. Rinaldo, J.B. Margolick, J.P. Phair, and J.W. Mellors. Natural history of human immunodeficiency virus type 1 viremia after seroconversion and proximal to AIDS in a large cohort of homosexual men. Multicenter Aids Cohort Study. *Journal of Infectious Diseases*, 181:872–880, 2000.

J.R. Mascola, B.J. Mathieson, P.M. Zack, M.C. Walker, S.B. Halstead, and D.S. Burke. Summary report: workshop on the potential risks of antibody-dependent enhancement in human HIV vaccine trials. *AIDS Research and Human Retroviruses*, 9:1175–1184, 1993.

J.W. Mellors, C.R. Rinaldo, P. Gupta, R.M. White, J.A. Todd, and L.A. Kingsley. Prognosis in HIV-1 infection predicted by the quantity of virus in plasma. *Science*, 272:1167–1170, 1996.

C.W. Mwachari, B.E. Shepherd, O. Cleopa, J.A. Odhiambo, and C.R. Cohen. Mortality and burden of disease in a cohort of HIV-seropositive adults in Nairobi, Kenya. *International Journal of STD and AIDS*, 15:120–126, 2004.

G.J. Nabel. Challenges and opportunities for development of an AIDS vaccine. *Nature*, 410:1002–1007, 2001.

J. Neyman. On the application of probability theory to agricultural experiments: Essay on principles. *Translated in Statistical Science*, 5:465–480, 1923.

T.R. O'Brien, W.A. Blattner, D. Waters, E. Eyster, M.W. Hilgartner, A.R. Cohen, N. Luban, A. Hatzakis, L.M. Aledort, P.S. Rosenberg, W.J. Miley, B.L. Kroner, and J.J. Goedert. Serum HIV-1 RNA levels and time to development of AIDS in the Multicenter Hemophilia Cohort Study. *Journal of the American Medical Association*, 276:105–110, 1996.

M.I. Parzen, L.J. Wei, and Z. Ying. Simultaneous confidence intervals for the difference of two survival functions. *Scandinavian Journal of Statistics*, 24:309–314, 1997.

J. Pearl. *Causality: models, reasoning, and inference.* Cambridge University Press, New York, 2000.

P. Putkonen, R. Thorstensson, L. Ghavamzadeh, J. Albert, K. Hild, G. Biberfeld, and E. Norrby. Prevention of HIV-2 and SIVsm infection by passive immunization in cynomolgus monkeys. *Nature*, 352:436–438, 1991.

T.C. Quinn, M.J. Wawer, N. Sewankambo, D. Serwadda, C. Li, F. Wabwire-Mangen, M.O. Meehan, T. Lutalo, and R.H. Gray. Viral load and heterosexual transmission of

human immunodeficiency virus type 1. *New England Journal of Medicine*, 342:921–929, 2000.

T. Ready. AIDSVAX flop leaves vaccine field unscathed. *Nature Medicine*, 9:376, 2003.

J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modeling*, 7:1393–1512, 1986.

J.M. Robins. An analytic method for randomized trials with informative censoring: Part I. *Lifetime Data Analysis*, 1:241–254, 1995.

J.M. Robins. Correction for non-compliance in equivalence trials. *Statistics in Medicine*, 17:269–302, 1998.

J.M. Robins. Optimal structural nested models for optimal sequential decisions. In D.Y. Lin and P.J. Heagerty, editors, *Proceedings of the Second Seattle Symposium on Biostatistics*, pages 189–326. Springer-Verlag, New York, 2004.

J.M. Robins and S. Greenland. Comment on 'Causal inference without counterfactuals' by A.P. Dawid. *Journal of the American Statistical Association*, 95:431–435, 2000.

J.M. Robins and L. Wasserman. Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs. In D. Geiger and P. Shenoy, editors, *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, Providence, Rhode Island, August 1-3, 1997*, pages 409–420. Morgan Kaufmann, San Francisco, 1997.

W.E. Robinson, D.C. Montefiori, and W.M. Mitchell. Antibody-dependent enhancement of human immunodeficiency virus type 1 infection. *Lancet*, 1:790–794, 1988.

P.R. Rosenbaum. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society, Series A*, 147: 656–666, 1984.

P.R. Rosenbaum and D.B. Rubin. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society, Series B*, 45:212–218, 1983a.

P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983b.

S.L. Rowland-Jones, T. Dong, K.R. Fowke, J. Kimani, P. Krausa, H. Newell, T. Blanchard, K. Ariyoshi, J. Oyugi, E. Ngugi, J. Bwayo, K.S. MacDonald, A.J. McMichael, and F.A. Plummer. Cytotoxic T cell responses to multiple conserved HIV epitopes in HIV-resistant prostitutes in Nairobi. *Journal of Clinical Investigation*, 102:1758–1765, 1998.

D.B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701, 1974.

D.B. Rubin. Bayesian inference for causal effects: the role of randomization. *The Annals of Statistics*, 6:34–58, 1978.

D.B. Rubin. Statistics and causal inference: which ifs have causal answers. *Journal of the American Statistical Association*, 81:961–962, 1986.

D.B. Rubin. Comment on 'Causal inference without counterfactuals' by A.P. Dawid. *Journal of the American Statistical Association*, 95:435–437, 2000.

D.O. Scharfstein, M.J. Daniels, and J.M. Robins. Incorporating prior beliefs about selection bias into the analysis of randomized trials with missing outcomes. *Biostatistics*, 4:495–512, 2003.

D.O. Scharfstein, A. Rotnitzky, and J.M. Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *Journal of the American Statistical Association*, 94:1096–1146, 1999.

G. Shafer. Comment on 'Causal inference without counterfactuals' by A.P. Dawid. *Journal of the American Statistical Association*, 95:438–442, 2000.

B.E. Shepherd, A.J. Rossini, R.J. Soto, I.L. de Rivera, and J.I. Mullins. Sampling designs for HIV molecular epidemiology with application to Honduras. *AIDS Research and Human Retroviruses*, in press, 2005.

J. Stover, N. Walker, G.P. Garnett, J.A. Salomon, K.A. Stanecki, P.D. Ghys, N.C. Grassly, R.M. Anderson, and B. Schwartlander. Can we reverse the HIV/AIDS pandemic with an expanded response? *Lancet*, 360:73–77, 2002.

W. Stute. The central limit theorem under random censorship. *The Annals of Statistics*, 23:422–439, 1995.

I.M. Thompson, P.J. Goodman, C.M. Tangen, M.S. Lucia, G.J. Miller, L.G. Ford, M.M. Lieber, R.D. Cespedes, J.N. Atkins, S.M. Lippman, S.M. Carlin, A. Ryan, C.M. Szczepanek, J.J. Crowley, and C.A. Jr. Coltman. The influence of Finasteride on the development of prostate cancer. *New England Journal of Medicine*, 349:213–222, 2003.

A.W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, Cambridge, 1998.

J.L. Zhang and D.B. Rubin. Estimation of causal effects via principal stratification when some outcomes are truncated by 'death'. *Journal of Educational and Behavioral Statistics*, 28:353–368, 2003.

Appendix A

AN INTRODUCTION TO HIV AND HIV VACCINES

Human immunodeficiency virus (HIV) is one of the leading causes of death worldwide. Since the virus was first isolated in 1983, well over 20 million people have died from HIV-induced disease. There are approximately 40 million people living with HIV and 5 million people were infected this past year. HIV is especially devestating in Sub-Saharan Africa, with nearly 65% of infected individuals living in this region. However, HIV also appears to be spreading rapidly in other areas, particularly southeast Asia (www.unaids.org).

The most common forms of virus acquisition are through sexual activity with an infected individual, sharing intravenous needles with an infected person, or through an infected mother passing on the virus to a child during birth or through breast feeding.

HIV gradually destroys its host's immune system, killing T-cells used to fight off infections. Hence, as more and more T-cells are killed, the body becomes more susceptible to so-called "opportunisitic infections" and eventually progresses to Acquired Immune Deficiency Syndrome (AIDS) and death. Many researchers have studied the clinical manifestations of HIV and AIDS. (Early in my graduate studies at the University of Washington, I was involved in some of this work, looking at opportunisitic infections among seropositive individuals in Nairobi, Kenya (Mwachari et al, 2004)).

Two common, non-clinical measurements of HIV progression are HIV viral load and CD4 count. HIV infection is characterized by an initial spike in the viral load, because of rapid viral replication shortly after infection. After the initial burst of viral replication, the host immune system is able to control the virus to a certain extent and the viral load decreases and levels off at what is referred to as the "set point" viral load. A higher set point viral load is thought to be associated with more rapid progression to AIDS. As the virus lingers in the host, it gradually kills T-cells, thus lowering the host's CD4 count. Once the CD4 count gets fairly low, the immune system is no longer able to suppress the virus and viral

load again increases. At this time clinical manifestations of AIDS become apparent and an untreated individual progresses towards death. Without treatment, the median time from HIV infection until death is around ten years.

One of the major breakthroughs in HIV science in the 1990s was the introduction of highly active antiretroviral therapy (HAART). Through a combination of drugs, individuals are now generally able to control the progression of the virus. HAART is effective but many people experience drug toxicity problems when on these drug regimens, and unfortunately HAART is unable to completely rid the body from HIV. In addition, HAART is largely unavailable to individuals infected in 3rd world countries where HIV is most rampant.

As with many other viruses, HIV mutates frequently, which in turn means that it rapidly evolves. This makes treatment challenging. The virus is known to "escape" treatment regimens by mutating and favoring resistant strains under selective pressures. HIV's rapid evolution also makes it difficult to design a preventative vaccine. An ideal vaccine would be effective at preventing infection for many different genetic strains of HIV. Much work has been done attempting to characterize different types of HIV strains and looking at the geographic distribution of HIV strains. (I have also done a little work in this area, designing a sampling scheme for studying the molecular epidemiology of HIV in Honduras (Shepherd et al., 2005)).

There are two general types of vaccines currently being studied in clinical trials: 1) vaccines designed to prevent infection by creating neutralizing antibodies and 2) vaccines designed to induce a cellular immune response to destroy infected cells.

The idea behind a neutralizing antibody vaccine is to train the immune system so that it will recognize HIV antigens and destroy the virus before it has a chance to infect cells. In order for this type of vaccine to be broadly effective, it will have to be general enough to recognize many different antigens corresponding to the genetic variation between HIV strains. Viruses similar to HIV have been placed in monkeys and vaccines of this type have been designed, producing antibodies that have protected monkeys from HIV challenges (Putkonen et al., 1991). VaxGen's candidate vaccine, AIDSVAX B/B, was based on this idea. It contains recombinant HIV envelope glycoprotein 120 (rgp120) antigens. The idea behind the vaccine is that the immune system might learn to recognize this antigen, producing

antibodies that would recognize and destroy the virus, preventing infection.

The next wave of vaccines are designed to induce a cell-mediated response against infected cells. These vaccines are not designed to prevent infection, but to suppress the virus once infections have occurred. There is evidence favoring a vaccine of this type. Certain people with significant exposure to HIV have not become infected, suggesting that perhaps there was "transient infection," eventually leading to protective immunity. Also, some long term non-progressors appear to have T cell activity that has kept the virus under control. Vaccines of this type have been tested in monkeys and have been shown to not protect against infection, but to delay disease progression (Hirsch et al., 1996). (Unfortunately, however, in the monkey vaccine studies the virus has eventually mutated, leading to viral escape (Barouch et al., 2002).) Two current phase III HIV vaccines, HVTN trial 502 being performed by Merck and a trial of ALVAC-HIV in Thailand, are of this type. In these trials, post-infection outcomes are of primary interest for evaluating vaccine efficacy.

It is now generally believed that the most effective vaccine will be a combination of these two vaccine approaches. The most effective vaccine will create neutralizing antibodies to defend cells from infection and then provide a strong cell-mediated response to the virus in cells that may become infected.

Most researchers believe that these vaccines will not increase participants' risk of acquiring HIV; throughout this dissertation I make this assumption. (This assumption, called monotonicity, is formally written as assumption A.1 in Section 2.2.) However, there are hypothesized mechanisms whereby monotonicity could be violated. There is the theoretical possibility that the vaccine might enhance the probability of infection (Robinson et al., 1988). For example, for several infectious diseases, vaccine-induced partial immunity has been observed to increase disease severity post-infection (Mascola et al. (1993); see also references in Gilbert et al. (2005)). However, in all of these situations the vaccine did not increase the probability of infection – only disease severity after infection. (These examples actually demonstrate how important it is to look at the causal effect of vaccination on post-infection outcomes.) In the VaxGen trial, immune response (measured by binding and neutralizing antibody measurements) was correlated to the incidence of HIV infection. Those people with higher antibody levels were less likely to be infected. This could be due

to the vaccine causing both an increased and decreased risk for HIV acquisition. If this were in fact the case, then monotonicity would be violated. However, it is more likely that the antibody levels are simply correlates of susceptibility to HIV (Gilbert et al., 2005).

Of course, this Appendix is only a very brief introduction to HIV/AIDS. The literature in this area is extremely rich. Additional references that I have found helpful include Graham (2002); Lukashov et al. (2002); Crandall (ed) (1999); Lifson and Martin (2002); Clements-Mann (1998); Anonymous (2001); www.hvtn.org; and www.iavi.org.

Appendix B

CAUSAL INFERENCE

The field of causal inference is broad, and overlapping work has been developed in many different disciplines. It is beyond the scope of this dissertation to give a good review of the causal inference literature. Here I give a summary of some basic ideas and then briefly discuss the links between this dissertation and some major areas of causal inference research.

Potential outcomes / counterfactuals were discussed as early as 1923 by Neyman with regards to crop yields: a plot has a "potential yield" if exposed to a particular variety (say, $Z = 0$) and another if exposed to a different variety ($Z = 1$). In my notation, the potential yields are written $Y(0)$ and $Y(1)$. (The term "counter-factual" was apparently introduced by Chisholm, 1946.) Rubin (1974, 1978) defined the individual causal effect of treatment Z on outcome Y for subject i as $Y_i(1) - Y_i(0)$. One could more generally define a causal effect for subject i as any measure of discrepency between $Y_i(0)$ and $Y_i(1)$, for example the ratio. However, the difference is often used because $E(Y_i(1) - Y_i(0)) = E(Y_i(1)) - E(Y_i(0))$. This quantity is referred to as the average causal effect and is identifiable under random treatment assignment and the assumption that potential outcomes for any subject i are unrelated to treatment assignment for other participants. This latter assumption is often referred to as the stable unit treatment value assumption (SUTVA). SUTVA implies consistency, that $Y = Y(z)$ given $Z = z$ (Robins, 1995). A good summary of basic causal inference ideas is given by Holland (1986).

B.1 Propensity Scores

This potential outcomes framework has been used in the analysis of observational studies through the use of propensity scores (Rosenbaum & Rubin, 1983). A propensity score is the probability that a study participant is assigned to the treatment of interest rather than the comparison group given covariates, and possibly the outcome variable. One estimates $P(Z =$

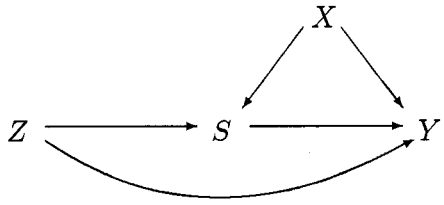

Figure B.1: Directed acyclic graph of HIV vaccine trial

$1|X)$, with X a vector of covariates, and uses this estimate as a weight to estimate $E(Y(1) - Y(0))$. If there are no unmeasured confounders, then the estimate of the corresponding average causal effect is consistent. Although I would not refer to my work as a propensity score method, one could think of $w(\cdot) = P(S(1) = 1|X = x, Y = y, S = 1, Z = 0)$ as a sort of propensity score – propensity to being in the always infected stratum. If one correctly models the "propensity score," i.e. $w(\cdot)$, and if there are no unmeasured "confounders," i.e. $Y(0) \amalg S(1)|S(0) = 1, X$, then one can consistently estimate $E(Y(1) - Y(0))$.

B.2 Directed Acyclic Graphs

A major area of causal inference work deals with directed acyclic graphs (DAGs). These graphs are used to visually express assumptions and to facilitate representation of joint probability functions. Pearl (2000) discusses many of the developments in this area. Our situation is shown in Figure B.1 in the form of a DAG.

Let G denote the graph, and define E as the set of evidence, or the set of information (in addition to Z and Y) to be included in an analysis. The variables Z, S, X, and Y are as defined throughout this document: treatment, infection, baseline covariates, and viral load, respectively. The edges or lines of the graph represent associations between variables and the arrows represent causation. For instance in G, $Z \rightarrow S$ denotes that the value of Z influences the value of S.

In G, there are multiple paths linking Z (treatment) to Y (viral load). I will label path 1 as $Z \rightarrow S \rightarrow Y$, path 2 as $Z \rightarrow S \leftarrow X \rightarrow Y$, and path 3 as $Z \rightarrow Y$. We are interested

in knowing if path 3 exists. In path 2, S is said to be a collider because two arrows are pointing towards it. On the other hand, S is a non-collider in path 1. X is a non-collider in path 2. There is a standard calculus for determining conditional independence using DAGs. Two variables are d-connected through a path if every collider is in E or has a descendent in E and every non-collider on the path is not in E. Otherwise they are d-separated. If Z and Y are d-separated by E in a DAG G, then Z is independent of Y conditional on E in every distribution compatible with G. Conversely, if Z and Y are not d-separated by E in a DAG G, then Z and Y are dependent conditional on E in at least one distribution compatible with G.

In order to determine the presence or absence of path 3, Z and Y must be d-separated through all other paths. In other words, assuming path 3 does not exist and our graph is correct, if Z and Y are d-separated through all paths, then Z and Y are conditionally independent. Then if we do see an association between Z and Y, we know it must be through path 3.

In our situation, if E is the empty set, then Z and Y are d-connected through path 1 and d-separated through path 2. If $E = S$, then Z and Y are d-connected through path 2 and d-separated through path 1. If $E = X$ then Z and Y are d-connected through 1 and d-separated through 2. Finally, if $E = \{X, S\}$ then Z and Y are d-separated through both paths 1 and 2.

Therefore, in order to look at the causal effect of Z on Y, we must condition on both X and S. If we do not condition on anything (the ITT analysis), then we can estimate the effect of Z on Y but this is obviously dependent on S (infection); any observed effect of Z on Y may actually be due to the effect of Z on S. If we condition only on S (the simple or naive analysis) then we cannot estimate the causal effect of Z on Y because this is dependent on X (other variables associated with vaccine efficacy and viral load). If we condition on S and X then a causal comparison can be made.

Another scenario would be the case where the line between Z and S does not exist. In this situation, Z and S are d-separated through paths 1 and 2, no matter what we condition on, so any observed effect of Z on Y must be through path 3.

Notice that the graph G implicitly assumes that there are no unmeasured variables

associated with vaccine efficacy and viral load (i.e., no path 4 with $Z \rightarrow S \leftarrow U \rightarrow Y$, with U not measured). This is typically a large assumption. Hence, the sensitivity analysis methods discussed in this dissertation can be thought as what to do when we believe there may be another path connecting S and Y through unmeasured variables.

B.3 Compliance

The counterfactual framework has also been extensively used to address the problem of noncompliance in randomized trials. The goal of this research is to estimate the effect of taking a treatment, not simply being randomized to a treatment. In much of this work, one conditions on individuals who would comply with their treatment assignment regardless of what it is. These individuals are referred to as compliers. The key assumptions are typically SUTVA, randomization, the exclusion restriction (that treatment assignment has no effect on the outcome other than on what treatment was taken; i.e., any effect of Z on Y must be via the effect of Z on D and then D on Y, where D designates the treatment actually taken), and monotonicity (that there are no defiers; i.e., no individuals who would always take the treatment they were not assigned). Angrist, Imbens, & Rubin (1996) studied noncompliance, linking this general framework to the instrumental variables work in the economics literature.

The DAG describing the non-compliance problem in randomized trials is of the same form as Figure B.1, replacing S (an indicator of infection) with D (an indicator of treatment taken). However, there are fundamental differences between the vaccine and compliance problems which do not allow one to apply previously developed methods for dealing with non-compliance to the vaccine problem, and equivalently make the methods of this dissertation impractical for use in the compliance problem. First, in the compliance literature the exclusion restriction is commonly made. This assumption is that there is no path $Z \rightarrow Y$ directly (no path 3). The effect of Z on Y is, of course, what we are studying. So assuming the exclusion restriction would be assuming our problem away!

Second, the monotonicity assumption typically made in the compliance literature (empty defiers principal stratum) is different from the monotonicity assumption of this work (empty

Table B.1: Vaccine Trial Scenario

	Vaccine ($Z = 1$)	Placebo ($Z = 0$)
Infected ($S = 1$)	always infected ~~harmed~~	always infected protected
Not Infected ($S = 0$)	never infected protected	never infected ~~harmed~~

Table B.2: Compliance Scenario

	Treatment ($Z = 1$)	Control ($Z = 0$)
Took Treatment ($D = 1$)	always taker complier	always taker ~~defier~~
Took Control ($D = 0$)	never taker ~~defier~~	never taker complier

harmed principal stratum). This difference is demonstrated in the tables above. In the vaccine scenario, when one assumes that no individual is in the harmed principal stratum then information is gained with regards to who is in the always infected principal stratum (which is the stratum in which inference is to be made); i.e., everyone infected in the vaccine arm is in the always infected principal stratum. In the compliance setting, however, if one assumes there are no defiers, then no information is gained about those who are compliers. Under this assumption one can estimate the proportion of subjects who are compliers, but without making other assumptions one cannot identify who is a complier in either treatment arm. Therefore the methods discussed in this dissertation do not apply. On the other hand, the methods of this dissertation would apply if one assumes there are no always takers (or similarly no never takers). Then one could say that everyone with Z=1 and D=1 is a complier, and one could then use a single sensitivity parameter to compare Y between compliers in the two treatment arms. However, these assumptions are typically not reasonable.

I choose to ignore compliance issues in this dissertation; these methods look at the effect of being assigned vaccine, not the effect of actually taking the vaccine.

B.4 Frequently Asked Questions / Complaints

As I have presented this work to others, a few questions have been repeatedly asked. They are valid questions, most dealing with causal inference philosophies. In fact, many of the complaints are similar to those given by Dawid, 2000. I have written a few responses, which I feel should be included in my dissertation. After all, a significant amount of my dissertation work has been thinking over these and other issues, trying to understand assumptions, and deciding if I am answering the right question. Some of my responses have been worked into the main text of my dissertation. However, a few of them did not flow with the main text, so I have put them here.

B.4.1 Principal strata are "fatalistic" or "deterministic"

This concern is well stated in Section 7 of Dawid, 2000. In essence, this concern deals with the existence of principal strata. Some argue that it is overly deterministic to group people as belonging to the "always infected" principal stratum. They argue that there are an infinite number of pathways an individual in a clinical trial can take to either end up infected or uninfected and that putting someone in a principal stratum "leaves no scope for introducing realistic stochastic effects of external influences acting between the times of application of treatment and of the response."

That each participant can be categorized into a principal stratum, I would not call fatalism but rather realism (Casella and Schwartz, 2000).

Imagine an experiment where you flip two coins simultaneously, one with your right hand and the other with your left hand. We want to answer the question: what is the probability that both coins will be heads? Or in other words, what is the probability that the coin is always heads. Another way to ask the same question is what is the probability of being in the always heads stratum? ("Always" is referring to an individual getting heads with both his right hand and with his left hand in this experiment - not at all times or in all experiments the individual participates in. Obviously a person could be in the always heads group and then repeat the experiment and not be in the always heads group the next experiment.) There are many ways for each coin to go through the experiment (for example, one coin may bounce twice on the ground before landing, and another coin may roll on the ground), but what we are interested in is the final outcome. In other words, flipping the coin is a stochastic event and there are many pathways to get heads, but we want the probability of being in the always heads stratum (i.e. flipping heads with both hands). Assuming the coin is fair and the results in both hands are independent, this probability is 0.25. You could estimate this by conducting the experiment on many people and taking the proportion of them who came up with both heads. As your sample gets big, your estimate will approach 0.25.

Now suppose each person can only flip one coin at a time, and they are randomized to flip with either their right or left hand. Our question is still the same: what is the probability of

being in the always heads stratum? In other words, what is the probability of being heads if randomized to use left hand and being heads if randomized to use right hand? We do not know if a particular individual is in the always heads group, because they only flip the coin with one hand. But under reasonable assumptions (that the probability of getting heads is identical for both hands for all individuals in the experiment and independent between an individual's hands and between individuals) we can estimate the probability of being in the always heads group as simply the proportion of people who got heads who flipped with their left hand times the proportion of people who got heads who flipped with their right hand. As the sample gets big, this will approach 0.25 with a fair coin. We do not know if an individual who flipped heads with, for example, his left hand is in the always heads stratum. But under our assumptions, we can assume that conditional on getting heads with his left hand, he has a probability of being in the always heads stratum of 0.5.

Similar to the coin toss experiment, in my setting we also have a dichotomous outcome, HIV infection within a pre-specified follow-up period. There are also two treatment arms; instead of left hand/right hand we have vaccine/placebo. People are randomized to get one treatment and we want to know the probability of being in the always infected stratum (similar to the always heads stratum). When I say "always infected" I mean that in this experiment, these individuals would have become infected whether or not they received the vaccine. I am not saying that no matter when they receive the vaccine or placebo they are doomed to be infected, I am just talking about the realization of this experiment. (Perhaps the use of the word "always" that causes confusion. Maybe a better name for the category would be the "infected in this particular experiment regardless of treatment assignment" category.) There may be many ways for each person to go through the experiment, but we are interested in their final outcome. Of course we only observe one outcome for each individual, and unlike the coin toss, it is no longer safe to assume that the infection status if randomized to vaccine is independent of the infection status if randomized to placebo. However, in some situations it may be reasonable to make other assumptions (such as independence between participants and monotonicity) that allow one to identify the probability of being in the always infected stratum.

B.4.2 Monotonicity is Untestable and Easily Violated

Much weight is placed on the monotonicity assumption. This is that one of the principal strata (the "harmed stratum") is empty. This assumption is more than just assuming $P(S_i(1) = 1) \leq P(S_i(0) = 1)$, but is assuming that $S_i(1) \leq S_i(0)$ for all i. This assumption is never completely verifiable. We cannot simultaneously assign someone to receive the placebo and vaccine. Even with a crossover trial with a suitable washout period (which is not possible in our setting), it is impossible to prove that the monotonicity assumption holds.

As stated in the text, assumption A.1 and randomization imply that the rates of infection under vaccine cannot be greater than under placebo. One could perform a test of H_0^* : $P(S = 1|Z = 1) \leq P(S = 1|Z = 0)$ and if H_0^* is rejected, then monotonicity is rejected. However, this test is not a consistent test of monotonicity: if monotonicity is violated but H_0^* holds, no amount of data will allow one to reject monotonicity. One might investigate the reasonableness of monotonicity by performing tests of H_0^* condidional on covariates. If one includes all of the covariates that explain vaccine efficacy and performs tests of H_0^* at each covariate level, then in theory these tests combined consistently test monotonicity. However, of course there is no way to know that all relevant covariates have been collected and even if one did collect all relevant covariates and knew it, sparseness in the data would almost certainly make such tests have inadequate power. Hence, one can look at the data to get an idea of the reasonableness of monotonicity, but in the end, one must realize that monotonicity is never completely testable.

There is another difficulty in making the monotonicity assumption. I call it the "butterfly effect" phenomenon, after the chaos theory idea that whether or not a butterfly decides to flap its wings determines whether or not there will be a great storm (thought to have originated from Edward Lorenz, 1960). This goes to the idea that every little action influences everything else.

Suppose a vaccine does not cause infection in any individual and in some individuals it prevents infection. Then it would be natural to make the monotonicity assumption: no individual has the potential outcome of infection if randomized to vaccine and non-infection

if randomized to placebo. However, even in this situation the monotonicity assumption could be violated.

For example, John did not get infected and he was assigned to the placebo pill. Had he been assigned to the vaccine pill his provider would have read a different number off the random number generator forever changing the course of history. The provider's actions would be slightly different, leading to different actions by John, which could lead to John being at a different place at a different time, being exposed differently to the virus, and perhaps being infected – even though the vaccine did not cause infection. In other words, John could realistically be in the harm principal stratum completely by chance, not causally. In this context, the monotonicity assumption appears to be even more difficult to support. (Shafer hints at this concept in Shafer, 2000.)

But then what statistical assumptions are completely verifiable anyway? For example, independence cannot be completely verified. We make the assumption when independence is plausible or that any dependence is so minimal that it should not alter our conclusions. If the "butterfly effect" is a concern to the validity of monotonicity, then it is also a concern to the validity of independence. The effect of any very minor violation should not alter our conclusions. If one wants to examine the effect of minor violations to the assumption of monotonicity then one could perform a three-parameter sensitivity analysis as discussed in Jemiai's PhD disseration.

Another point worth noting is a sort of risk-benefits assessment of making assumptions. One risk of being afraid of using assumptions is that inferences are more difficult to interpret (e.g., a three-parameter sensitivity analysis). Humans have a limited range / window of perception, and simplifying assumptions are often essential for comprehension.

In many circumstances, however, monotonicity may not be a valid assumption. For example, if one were comparing two vaccines, monotonicity would probably not be reasonable. If monotonicity is a major concern, then the methods of this dissertation should not be implemented. In order to identify $ACE(x)$ one will need to make other assumptions about the joint distribution of $(S(0), S(1))$ given X. One such assumption is explainable non-random infection (Hayden et al, 2005; Robins, 1998). In the absence of covariates and monotonicity, sharp non-parametric bounds have been derived by Jemiai (unpublished Ph.D. dissertation)

and Zhang & Rubin (2003).

B.4.3 SUTVA may be Violated when Dealing with an Infectious Disease

The Stable Unit Treatment Value Assumption (SUTVA) states that the potential outcomes of each trial participant are not influenced by the treatments of other participants. This assumption may be violated when studying an infectious disease. For example, if a vaccine is partially effective then some individuals in the trial will not be infected because they took the vaccine. If they are not infected, then that may make it less likely for other trial participants to become infected, because there are less infected people to infect them. In this situation, the potential outcomes for some trial participants may be influenced by the treatment assignment for other participants – violating SUTVA. Stated another way, one might say the assumption of SUTVA is similar to an assumption of no "herd immunity."

Clearly, SUTVA must be carefully considered before assuming. In our example dataset, SUTVA was thought to hold because the trial was performed in sites all across North America and the Netherlands and there was therefore presumably little, if any, interaction between trial participants. SUTVA would be much less plausible if the trial were performed in a smaller area, with more interaction between participants. SUTVA also might be less plausible for a trial conducted over a longer period of time, allowing more time for the infectious disease to spread between participants. Similarly, SUTVA may be less plausible for a more infectious disease.

I am not aware of causal inference methods that allow one to relax the assumption of SUTVA. Indeed, even the notation adopted in this dissertation implicitly assumes SUTVA. Without this assumption, potential outcomes would need to be written as, for example, $S_i(\mathbf{Z})$, where $\mathbf{Z} = (Z_1, \cdots, Z_N)$ is the treatment assignments for all trial participants. Using this notation, it is difficult to even define interesting estimands. One approach to relaxing SUTVA might be to group interacting individuals together, perhaps defining outcomes at a group level and making the assumption of SUTVA between groups. However, using this type of approach it would still be difficult to define an interesting and identifiable estimand.

Appendix C

TECHNICAL DETAILS FOR CHAPTER 3

C.1 Deriving $ACE(x)$

$$\begin{aligned} ACE(x) &= E(Y(1)-Y(0)|S(0)=S(1)=1, X=x) \\ &= E(Y(1)|S(0)=S(1)=1, X=x) - E(Y(0)S(0)=S(1)=1, X=x) \\ &= \int y f_{Y(1)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1, X=x)dy \\ &\quad - \int y f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1, X=x)dy. \end{aligned}$$

From A.1,

$$\begin{aligned} &\int y f_{Y(1)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1, X=x)dy \\ &\quad = \int y f_{Y(1)|S(0)=S(1)=1,X}(y|S(1)=1, X=x)dy \\ &\quad = \int y f_v(y|x;\eta_1) \qquad \text{by M.3.} \end{aligned}$$

Under $\mathcal{M}_b$,

$$\int y f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1, X=x)dy = \int y f_p^{ai}(y|x;\eta_0^b) \quad \text{by M.4b}$$

which implies

$$ACE(x) = \int y f_v(y|x;\eta_1) - \int y f_p^{ai}(y|x;\eta_0^b) \equiv ACE_b(x;\eta_1,\eta_0^b).$$

Under $\mathcal{M}_b$,

$$\begin{aligned}
&\int y f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1,X=x)dy \\
&\quad = \int y \frac{P(S(1)=1|S(0)=1,Y(0)=y,X=x)P(Y(0)=y|S(0)=1,X=x)}{\int P(S(1)=1|S(0)=1,Y(0)=t,X=x)P(Y(0)=t|S(0)=1,X=x)dt} dy \\
&\qquad \text{(by Bayes Rule)} \\
&\quad = \frac{\int y w(x,y;\beta,\alpha) f_p(y|x;\eta_0^a)dy}{\int w(x,t;\beta,\alpha) f_p(t|x;\eta_0^a)dt} \qquad \text{by A.2, M.2, and M.4a,}
\end{aligned}$$

which implies $ACE(x) = ACE_a(x;\eta_1,\eta_0^a)$.

C.2 Deriving Restrictions on the Observed Data

Under $\mathcal{M}_a$,

$$\begin{aligned}
&f_{Y|S,Z,X}(y|S=1,Z=0,X=x) \\
&\quad = f_{Y(0)|S(0)=1,Z,X}(y|S(0)=1,Z=0,X=x) \qquad \text{by (2.1)} \\
&\quad = f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x) \qquad \text{by (2.2)} \\
&\quad = f_p(y|x;\eta_0^a) \qquad \text{by M.4a.}
\end{aligned}$$

Under $\mathcal{M}_b$,

$$\begin{aligned}
&f_{Y|S,Z,X}(y|S=1,Z=0,X=x) \\
&\quad = f_{Y(0)|S(0)=1,Z,X}(y|S(0)=1,Z=0,X=x) \qquad \text{by (2.1)} \\
&\quad = f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x) \qquad \text{by (2.2)} \\
&\quad = \frac{w^{-1}(x,y;\beta,\alpha)w(x,y;\beta,\alpha)f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x)}{\int w(x,y;\beta,\alpha)f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x)dy} \Bigg/ \\
&\qquad \frac{\int w^{-1}(x,y;\beta,\alpha)w(x,y;\beta,\alpha)f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x)dy}{\int w(x,y;\beta,\alpha)f_{Y(0)|S(0)=1,X}(y|S(0)=1,X=x)dy} \qquad \text{if } w(\cdot)>0 \\
&\quad = \frac{w^{-1}(x,y;\beta,\alpha)f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1,X=x)}{\int w^{-1}(x,y;\beta,\alpha)f_{Y(0)|S(0)=S(1)=1,X}(y|S(0)=S(1)=1,X=x)dy} \qquad \text{by A.2, M.2} \\
&\quad = \frac{w^{-1}(x,y;\beta,\alpha)f_p^{ai}(y|x;\eta_0^b)}{\int w^{-1}(x,y;\beta,\alpha)f_p^{ai}(y|x;\eta_0^b)dy} \qquad \text{by M.4b} \\
&\quad \equiv f_p^*(y|x;\alpha,\eta_0^b).
\end{aligned}$$

Under $\mathcal{M}_a$ or $\mathcal{M}_b$,

$$
\begin{aligned}
& f_{Y|S,Z,X}\left(y|S=1,Z=1,X=x\right) \\
&= f_{Y(1)|S(1)=1,Z,X}\left(y|S\left(1\right)=1,Z=1,X=x\right) && \text{by (2.1)} \\
&= f_{Y(1)|S(1)=1,X}\left(y|S\left(1\right)=1,X=x\right) && \text{by (2.2)} \\
&= f_{v}(y|x;\eta_1) && \text{by M.3.}
\end{aligned}
$$

$$
\begin{aligned}
& P_{S|Z,X}\left(S=1|Z=0,X=x\right) \\
&= P_{S(0)|Z,X}\left(S\left(0\right)=1|Z=0,X=x\right) && \text{by (2.1)} \\
&= P_{S(0)|Z,X}\left(S\left(0\right)=1|X=x\right) && \text{by (2.2)} \\
&= \theta_p\left(x;\mu\right) && \text{by M.1.}
\end{aligned}
$$

$$
\begin{aligned}
& P_{S|Z,X}\left(S=1|Z=1,X=x\right) \\
&= P_{S(1)|Z,X}\left(S\left(1\right)=1|Z=1,X=x\right) && \text{by (2.1)} \\
&= P_{S(1)|Z,X}\left(S\left(1\right)=1|X=x\right) && \text{by (2.2)} \\
&= P_{S(0)|X}\left(S\left(0\right)=1|X=x\right) P_{S(1)|S(0),X}\left(S\left(1\right)=1|S\left(0\right)=1,X=x\right) && \text{by A.1} \\
&= \theta_p\left(x;\mu\right) \times \int P_{S(1)|S(0),Y(0),X}\left(S\left(1\right)=1|S\left(0\right)=1,Y\left(0\right)=y,X=x\right) \\
&\quad f_{Y(0)|S(0)=1,X}\left(y|S\left(0\right)=1,X=x\right)dy && \text{by M.1} \\
&= \begin{cases} \theta_p(x;\mu)\int w(x,y;\beta,\alpha) f_p\left(y|x;\eta_0^a\right)dy & \text{by A.2, M.2, M.4a} \\ \theta_p\left(x;\mu\right)\int w\left(x,y;\beta,\alpha\right) f_p^*(y|x;\alpha,\eta_0^b)dy & \text{by A.2, M.2, M.4b} \end{cases}
\end{aligned}
$$

C.3 *Asymptotic Distribution of the ML Estimator of* $ACE(x)$

Define $I(\rho^a) \equiv -E(\frac{\partial^2 l(\rho^a|O)}{\partial(\rho^a)^2})$ where $l(\rho^a|O) = \log(\mathcal{L}_a(\rho^a|O))$, the log-likelihood, with $\mathcal{L}_a(\rho^a|O)$ given by (3.1). I also define $I_0(\alpha, \eta_0^a) \equiv -E(\frac{\partial^2 l(\alpha,\eta_0^a|O,Z=0)}{\partial(\alpha,\eta_0^a)^2})$ where $l(\alpha, \eta_0^a|O, Z = 0)$=$\log(\mathcal{L}_a(\rho^a|O, Z = 0))$ and similarly, $I_1(\eta_1) \equiv -E(\frac{\partial^2 l(\eta_1|O,Z=1)}{\partial(\eta_1)^2})$ where $l(\eta_1|O, Z = 1)$ =$\log(\mathcal{L}_a(\rho^a|O, Z = 1))$. It is easily seen from (3.1) that

$$I(\rho^a) = \begin{bmatrix} I_0(\alpha, \eta_0^a) & 0 \\ 0 & I_1(\eta_1) \end{bmatrix},$$

so we will consider the vaccine and placebo arms separately.

Let $h_0(\alpha, \eta_0^a) \equiv E(Y(0)|S(0) = S(1) = 1, X = x)$ and $h_1(\eta_1) \equiv E(Y(1)|S(1) = 1, X = x)$. Therefore, $ACE(x) = h_1(\eta_1) - h_0(\alpha, \eta_0^a)$ and $\widehat{ACE}(x) = h_1(\hat{\eta}_1) - h_0(\hat{\alpha}, \hat{\eta}_0^a)$ where $\hat{\eta}_1$ denotes the MLE of η_1, etc.

In addition to the assumptions used to construct the likelihood and the assumption that the protected principal stratum is non-empty ($VE(x) > 0$), we need the following standard conditions (van der Vaart, 1998): (i) identifiability, (ii) the region of positive probability does not depend on ρ^a, and in open neighborhoods of the parameter space (iii) the log-likelihood is twice continuously differentiable, (iv) second and third order derivatives exist and are bounded, and (v) $I(\rho^a)$ is positive definite. If these conditions hold and $h_0(\cdot)$ and $h_1(\cdot)$ are differentiable with $h_0'(\cdot)$ and $h_1'(\cdot)$ representing the first derivatives, then applying standard likelihood theory and the delta method,

$$\sqrt{N_p}\left(h_0(\hat{\alpha}, \hat{\eta}_0^a) - h_0(\alpha, \eta_0^a)\right) \rightarrow^d h_0'(\alpha, \eta_0^a)^T \mathcal{N}\left(0, I_0(\alpha, \eta_0^a)^{-1}\right) \tag{C.1}$$

$$\sqrt{N_v}\left(h_1(\hat{\eta}_1) - h_1(\eta_1)\right) \rightarrow^d h_1'(\eta_1)^T \mathcal{N}\left(0, I_1(\eta_1)^{-1}\right). \tag{C.2}$$

Therefore, asymptotically

$$\widehat{ACE}(x) \sim \mathcal{N}\left(ACE(x), \frac{h_1'(\eta_1)^T I_1(\eta_1)^{-1} h_1'(\eta_1)}{N_v} + \frac{h_0'(\alpha, \eta_0^a)^T I_0(\alpha, \eta_0^a)^{-1} h_0'(\alpha, \eta_0^a)}{N_p}\right).$$

Of course, one can estimate $h_1'(\eta_1)$ with $h_1'(\hat{\eta}_1)$ and $I_1(\eta_1)$ with either $I_1(\hat{\eta}_1)$ (the estimated expected information) or with $-N_v^{-1}\sum_{i=1}^{N_v} \frac{\partial^2 l_1(\eta_1)}{\partial \eta_1^2}$ with η_1 evaluated at $\hat{\eta}_1$ (observed information).

Note that when we assume (5.3) with β_v unknown, the standard condition (i), identifiability, does not hold in certain regions of the parameter space. Therefore (C.2) does not hold, although (C.1) may still be valid.

C.4 Hudgens & Halloran's work when Y is Binary

In a paper that is currently under revision, Hudgens & Halloran consider the case where Y is a post-infection binary outcome. I have been asked how my work relates to their work. This section of the appendix essentially works through notation and algebra to show that the likelihood they present is a special case of the more general likelihood given by (3.1).

Let Y be a binary outcome that only exists in those who become infected. Define $\pi_p = P(Y(0) = 1|S(0) = 1)$ and $\pi_v = P(Y(1) = 1|S(1) = 1)$. Therefore modeling assumptions M.3 and M.4a are $f_z(y) = \pi_z^y(1-\pi_z)^{1-y}$ for $z = p, v$. The integral $\int w(y;\alpha,\beta) f_p(y;\eta_0^a)dy$ is written here as $1 - VE(\alpha,\beta,\pi_p)$. Hence writing the likelihood (3.1) without covariates I get

$$\begin{aligned}\mathcal{L}_a(\rho^a|\mathbf{O}) \propto \prod_{i=1}^{N} \Big\{ &\left[\theta_p\left(1 - VE\left(\alpha,\beta,\pi_p\right)\right)\pi_v^{y_i}(1-\pi_v)^{1-y_i}\right]^{S_i} \\ &\times \left[1 - \theta_p\left(1 - VE\left(\alpha,\beta,\pi_p\right)\right)\right]^{1-S_i}\Big\}^{Z_i} \\ &\times \left\{\left[\theta_p \pi_p^{y_i}(1-\pi_p)^{1-y_i}\right]^{S_i}\left[1-\theta_p\right]^{1-S_i}\right\}^{1-Z_i}.\end{aligned}$$

I will adopt the notation of Hudgens & Halloran to count the number of subjects in each group. Specifically,

$$\begin{aligned} n_{0*}(v) &= \sum z_i(1-s_i) \\ n_{10}(v) &= \sum z_i s_i(1-y_i) \\ n_{11}(v) &= \sum z_i s_i y_i \\ n_{0*}(p) &= \sum (1-z_i)(1-s_i) \\ n_{10}(p) &= \sum (1-z_i)s_i(1-y_i) \\ n_{11}(p) &= \sum (1-z_i)s_i y_i, \end{aligned}$$

Table C.1: Translating Hudgens & Halloran's notation to my notation.

My notation	Hudgens and Halloran
θ_p	$1-\theta^{00}$
$1-VE(\alpha,\beta,\pi_p)$	$\theta^{11}/(1-\theta^{00})$
π_p	$\phi^{\cdot 1}\theta^{11}/(1-\theta^{00})+\gamma^1(1-\theta^{11}/(1-\theta^{00}))$
π_v	$\phi^{1\cdot}$

with all summations from $i=1,\cdots,N$. The log-likelihood can then be written as

$$\begin{aligned} l(\rho^a;\mathbf{O}) =& n_{0*}(p)log(1-\theta_p)+n_{10}(p)log(\theta_p(1-\pi_p))+n_{11}(p)log(\theta_p\pi_p) \\ &+n_{0*}(v)log(1-\theta_p(1-VE(\alpha,\beta,\pi_p))) \\ &+(n_{10}(v)+n_{11}(v))log(\theta_p(1-VE(\alpha,\beta,\pi_p))) \\ &+n_{11}(v)log\pi_v+n_{10}(v)log(1-\pi_v). \end{aligned}$$

Hudgens & Halloran make the same assumptions that I make throughout this dissertation: SUTVA (2.1), randomization (or rather the conditional independence given by (2.2) which randomization ensures), monotonicity (A.1), and various assumptions that can be translated into using different forms of $w(\cdot)$ in assumption A.2. They define $\theta^{00}=P(ni)$, $\theta^{01}=P(prot)$, $\theta^{11}=P(ai)$, $\gamma^0=P(Y(0)=0|prot),\gamma^1=P(Y(0)=1|prot),\phi^{00}=P(Y(0)=0,Y(1)=0|ai),\phi^{10}=P(Y(0)=1,Y(1)=0|ai),\phi^{01}=P(Y(0)=0,Y(1)=1|ai),\phi^{11}=P(Y(0)=1,Y(1)=1|ai),\phi^{\cdot 1}=\phi^{01}+\phi^{11}$, and $\phi^{1\cdot}=\phi^{10}+\phi^{11}$. My parameters can be converted into their notation as shown in Table C.1.

Substituting the notation of Hudgens & Halloran into our log-likelihood and performing some minor algebra, we get

$$\begin{aligned} l(\rho^a;\mathbf{O}) =& n_{0*}(p)log(\theta^{00})+n_{10}(p)log(\theta^{11}(1-\phi^{\cdot 1})-\theta^{01}\gamma^0)+n_{11}(p)log(\phi^{\cdot 1}\theta^{11}+\gamma^1\theta^{01}) \\ &+n_{0*}(v)log(\theta^{00}+\theta^{01})+n_{10}(v)log(\theta^{11}(1-\phi^{1\cdot}))+n_{11}(v)log(\theta^{11}\phi^{1\cdot}), \end{aligned}$$

the log-likelihood reported by Hudgens & Halloran (2005).

Appendix D

TECHNICAL DETAILS FOR CHAPTER 5

D.1 Identification of η_1 under $\mathcal{M}_a$

Under $\mathcal{M}_a$ using (5.2)-(5.3) and (5.4) for relevant probabilities and distributions, $\eta_1 = (\alpha_v, \beta_v, \eta_v)$ is identifiable if and only if

- A. $\beta_v \neq 0$ and
- B. $\alpha_v \neq (-\beta_v\lambda_{v0} - \beta_v^2\sigma_v^2/2, -\beta_v\lambda_{v1}, \cdots, -\beta_v\lambda_{vq})$.

Proof. η_1 is identifiable if distinct values of η_1 correspond to distinct pdfs. That is, if $\eta_1 \neq \eta_1'$, then $f_v(y|x;\eta_1) \neq f_v(y|x;\eta_1')$, where

$$f_v(y|x;\eta_1) = \frac{\left(1+exp(-x^T\alpha_v - \beta_v y)\right)^{-1} \frac{1}{\sqrt{2\pi}\sigma_v} exp\left(-\frac{1}{2\sigma_v^2}(y - x^T\lambda_v)^2\right)}{\int \left(1+exp(-x^T\alpha_v - \beta_v y)\right)^{-1} \frac{1}{\sqrt{2\pi}\sigma_v} exp\left(-\frac{1}{2\sigma_v^2}(y - x^T\lambda_v)^2\right) dy}.$$

A. If $\beta_v = 0$ then $\eta_1 = (\alpha_v, 0, \lambda_v, \sigma_v) \neq (\alpha_v', 0, \lambda_v, \sigma_v) = \eta_1'$ but $f_v(y|x;\eta_1) = f_v(y|x;\eta_1')$. Therefore, α_v is not identifiable if $\beta_v = 0$.

B. For $\beta_v \neq 0$, $f_v(y|x;\eta_1) = f_v(y|x;\eta_1')$ if and only if the numerators of these densities are equal, or equivalently,

$$\frac{exp(x^T\alpha_v + \beta_v y)(1+exp(-x^T\alpha_v' - \beta_v' y))}{1+exp(x^T\alpha_v + \beta_v y)} =$$
$$\frac{\sigma_v}{\sigma_v'} exp\left[-\frac{1}{2}\left(\frac{1}{\sigma_v'^2}(y - x^T\lambda_v')^2 - \frac{1}{\sigma_v^2}(y - x^T\lambda_v)^2\right)\right].$$

Without loss of generality, consider the case where there is only one covariate. The components of this equation that inhibit equality are the denominator on the left side of

the equation and x_1^2, y^2, and $x_1 y$ on the right side of the equation. In other words, with the denominator remaining in the left side of the equation or with x_1^2, y^2, or $x_1 y$ remaining in the right side of the equation, equality is impossible for all x_1 and y. Notice that y^2 disappears only if $\sigma_v = \sigma_v'$, x_1^2 disappears only if $\lambda_{v1}'/\sigma_v' = |\lambda_{v1}|/\sigma_v$, and $x_1 y$ disappears only if $\lambda_{v1}'/\sigma_v' = \lambda_{v1}/\sigma_v$. Considering the left side of the equation, the denominator disappears only if $(\alpha_v', \beta_v') = -(\alpha_v, \beta_v)$. Therefore, if $(\alpha_v', \beta_v', \sigma_v', \lambda_{v1}') = (-\alpha_v', -\beta_v', \sigma_v', \lambda_{v1}')$, there may be an identifiability problem. In this parameter space, the equation becomes

$$\begin{aligned} exp(x^T\alpha_v + \beta_v y) &= exp\left[-\frac{1}{2\sigma_v^2}\left((y - x^T\lambda_v)^2 - (y - x^T\lambda_v)^2\right)\right] \\ &= exp\left[\frac{1}{2\sigma_v^2}\left(\lambda_{v0}^2 - \lambda_{v0}'^2\right) + \frac{1}{\sigma_v^2}\lambda_{v1}\left(\lambda_{v0} - \lambda_{v0}'\right)x_1 + \frac{1}{\sigma_v^2}\left(\lambda_{v0}' - \lambda_{v0}\right)y\right]. \end{aligned}$$

This implies that

$$\begin{aligned} \eta_1 &= \left(\frac{\lambda_{v0}^2 - \lambda_{v0}'^2}{2\sigma_v^2}, \frac{\lambda_{v1}(\lambda_{v0} - \lambda_{v0}')}{\sigma_v^2}, \frac{\lambda_{v0}' - \lambda_{v0}}{\sigma_v^2}, \lambda_{v0}, \lambda_{v1}, \sigma_v\right) \\ &\neq \left(-\frac{\lambda_{v0}^2 - \lambda_{v0}'^2}{2\sigma_v^2}, -\frac{\lambda_{v1}(\lambda_{v0} - \lambda_{v0}')}{\sigma_v^2}, -\frac{\lambda_{v0}'\lambda_{v0}}{\sigma_v^2}, \lambda_{v0}', \lambda_{v1}, \sigma_v\right) = \eta_1', \end{aligned}$$

and $f_v(y|x;\eta_1) = f_v(y|x;\eta_1')$. Therefore, at certain values of η_1 there is exactly one other $\eta_1' \neq \eta_1$ such that $f_v(y|x;\eta_1) = f_v(y|x;\eta_1')$, implying that η_1 is not identifiable. A little algebra shows that this occurs only when η_1 that can be written as

$$\eta_1 = (-\beta_v\lambda_{v0} - \beta_v^2\sigma_v^2/2, -\beta_v\lambda_{v1}, \beta_v, \lambda_{v0}, \lambda_{v1}, \sigma_v).$$

Therefore η_1 is identifiable if

$$\alpha_v \neq (-\beta_v\lambda_{v0} - \beta_v^2\sigma_v^2/2, -\beta_v\lambda_{v1}, \cdots, -\beta_v\lambda_{vq}).$$

□

As previously stated, regardless of the choices for M.2 and M.4a, at the value of β_v that makes $w(x, y; \beta_v, \alpha_v)$ a function of x only, the parameter α_v is not identified. I just showed that under distributional assumptions (5.2)-(5.3) and (5.4), there are other regions of unidentifiability. These additional identification problems are also seen with other choices

for M.2 and M.4a. For example, suppose we let $w(x, y; \alpha, \beta)$ be the function, $exp(x^T\alpha+\beta y)$. (For our purposes, let us assume that this is bounded between 0 and 1, making it a valid cdf.) With $f_p(y|x;\eta_0^a)$ being the normal density defined by (5.4), then it is easily shown by completing the square that the density $f_p^{ai}(y|x;\alpha,\eta_0^a)$, is normal with mean $x^T\lambda+\beta\sigma^2$ and variance σ^2, is independent of α, and $1-VE(x;\alpha,\beta,\eta_0^a) = exp\left(x^T\lambda\beta+\beta^2\sigma^2/2+x^T\alpha\right)$. Hence, with $f_v(y|x;\eta_1)$ defined by (5.3) there is no need to estimate α_v. However, there is still unidentifiability: If $\sigma_v = \sigma_v'$ and $\lambda_{v1} = \lambda_{v1}'$ then $f_v(y|x;\eta_1) = f_v(y|x;\eta_1')$ iff $\lambda_{v0}+\beta_v\sigma_v^2 = \lambda_{v0}'+\beta'\sigma_v^2$.

D.2 Score Equations and Information under $\mathcal{M}_a$ using (5.1)-(5.4).

Because $f_v(y|x;\eta_1)$ factors out of the likelihood, it is considered separately later.

Score Equations:

$$\frac{\partial l}{\partial \mu_j} = \sum_{i=1}^{N} z_i s_i x_{ij}(1-\theta_i) - \sum_{i=1}^{N} z_i(1-s_i)x_{ij}\frac{J_i A_i}{1-\theta_i A_i} + \sum_{i=1}^{N}(1-z_i)s_i x_{ij}(1-\theta_i) - \sum_{i=1}^{N}(1-z_i)(1-s_i)x_{ij}\theta_i$$

$$\frac{\partial l}{\partial \alpha_j} = \sum_{i=1}^{N} z_i s_i x_{ij}\frac{M_i}{A_i} - \sum_{i=1}^{N}(1-z_i)(1-s_i)x_{ij}\frac{\theta_i M_i}{1-\theta_i A_i}$$

$$\frac{\partial l}{\partial \lambda_j} = \sum_{i=1}^{N} z_i s_i x_{ij}\frac{B_i - x_i^T\lambda A_i}{\sigma^2 A_i} - \sum_{i=1}^{N} z_i(1-s_i)x_{ij}\theta_i\frac{B_i - x_i^T\lambda A_i}{\sigma^2(1-\theta_i A_i)} + \sum_{i=1}^{N}(1-z_i)s_i x_{ij}\frac{y_i - x_i^T\lambda}{\sigma^2}$$

$$\frac{\partial l}{\partial \sigma} = \sum_{i=1}^{N} z_i s_i\frac{Q_i}{A_i} + \sum_{i=1}^{N} z_i(1-s_i)\frac{\theta_i Q_i}{1-\theta_i A_i} + \sum_{i=1}^{N}(1-z_i)s_i\left[\frac{(y_i - x_i^T\lambda)^2}{\sigma^3} - \frac{1}{\sigma}\right],$$

where

$$\theta_i = \frac{exp(x_i^T\mu)}{1+exp(x_i^T\mu)}$$
$$A_i = \int w(x_i,y) f_p(y|x_i)dy$$
$$B_i = \int y w(x_i,y) f_p(y|x_i)dy$$
$$C_i = \int y^2 w(x_i,y) f_p(y|x_i)dy$$
$$D_i = \int (y-x_i^T\lambda)^2 w(x_i,y) f_p(y|x_i)dy$$
$$J_i = \frac{\theta_i}{1+exp(x_i^T\mu)}$$
$$M_i = \int \frac{w(x_i,y)}{1+exp(x_i^T\alpha+\beta y)} f_p(y|x_i)dy$$
$$Q_i = \frac{1}{\sigma^3}D_i - \frac{1}{\sigma}A_i.$$

In the presence of censoring, with lower and upper detection limits L and U, respectively, the last terms of $\frac{\partial l}{\partial \lambda_j}$ and $\frac{\partial l}{\partial \sigma}$ are multiplied inside the summation by $I(L \le y_i \le U)$, and the following is added to the score equations:

$$\text{to } \frac{\partial l}{\partial \lambda_j}, \sum_{i=1}^{N}(1-z_i)x_{ij}s_i I(y_i<L)\frac{G(L)}{F_p(L)} - \sum_{i=1}^{N}(1-z_i)x_{ij}s_i I(y_i>U)\frac{G(U)}{1-F_p(U)};$$
$$\text{to } \frac{\partial l}{\partial \sigma}, \sum_{i=1}^{N}(1-z_i)x_{ij}s_i I(y_i<L)\frac{R(L)}{F_p(L)} - \sum_{i=1}^{N}(1-z_i)x_{ij}s_i I(y_i>U)\frac{R(U)}{1-F_p(U)},$$

where

$$F_p(L) = \int_{-\infty}^{L} f_p(y|x_i)dy$$
$$G(L) = \int_{-\infty}^{L} \frac{(y-x_i^T\lambda)}{\sigma^2} f_p(y|x_i)dy$$
$$R(L) = \int_{-\infty}^{L} \left(\frac{(y-x_i^T\lambda)^2}{\sigma^2} - 1\right)\frac{1}{\sigma} f_p(y|x_i)dy.$$

Expected Information:

$$
\begin{aligned}
I_{\mu_j\mu_k} &= \frac{1}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\left[N_v\left(J_i\theta_i A_i + \frac{\theta_i A_i\left(1+exp\left(x_i^T\mu\right)\left(\theta_i A_i - 1\right)\right)}{\left(1+exp\left(x_i^T\mu\right)\right)^2\left(1-\theta_i A_i\right)}\right) + N_p J_i\right] \\
I_{\alpha_j\alpha_k} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\frac{\theta_i M_i^2}{A_i\left(1-\theta_i A_i\right)} \\
I_{\lambda_j\lambda_k} &= \frac{1}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\frac{\theta_i}{\sigma^2}\left[N_v\frac{\left(B_i - x_i^T\lambda A_i\right)^2}{A_i\left(1-\theta_i A_i\right)\sigma^2} + N_p\right] \\
I_{\sigma\sigma} &= \frac{1}{N}\sum_{i=1}^{N}\left[\frac{2N_p\theta_i}{\sigma^2} + \frac{N_v\theta_i Q_i^2}{\left(1-\theta_i A_i\right)A_i}\right] \\
I_{\mu_j\alpha_k} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\frac{J_i M_i}{1-\theta_i A_i} \\
I_{\mu_j\lambda_k} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\frac{J_i\left(B_i - x_i^T\lambda A_i\right)}{\left(1-\theta_i A_i\right)\sigma^2} \\
I_{\alpha_j\lambda_k} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}x_{ik}\frac{\theta_i M_i\left(B_i - x_i^T\lambda A_i\right)}{\sigma^2 A_i\left(1-\theta_i A_i\right)} \\
I_{\mu_j\sigma} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}\frac{J_i Q_i}{1-\theta_i A_i} \\
I_{\alpha_j\sigma} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}\frac{\theta_i M_i Q_i}{\left(1-\theta_i A_i\right)A_i} \\
I_{\lambda_j\sigma} &= \frac{N_v}{N}\sum_{i=1}^{N} x_{ij}\frac{\theta_i Q_i\left(B_i - x_i^T\lambda A_i\right)}{A_i\left(1-\theta_i A_i\right)\sigma^2},
\end{aligned}
$$

for $j = 0, 1, \cdots, q$ and $k = 0, 1, \cdots, q$.

The expected information was not calculated in the case where there is censoring. In practice, when there is censoring I recommend simply using the observed information, which can be computationally obtained simply from the score equations.

Score Equations for η_1:

$$\frac{\partial l}{\partial \alpha_{vj}} = \sum_{i=1}^{N} z_i s_i x_{ij} \left(\frac{1}{1 + exp(x_i^T \alpha_v - \beta_v y_i)} - \frac{M_i}{A_i} \right),$$

$$\frac{\partial l}{\partial \beta_v} = \sum_{i=1}^{N} z_i s_i y_i \left(\frac{1}{1 + exp(x_i^T \alpha_v - \beta_v y_i)} - \frac{P_i}{A_i} \right),$$

$$\frac{\partial l}{\partial \lambda_{vj}} = \frac{1}{\sigma_v^2} \sum_{i=1}^{N} z_i s_i x_{ij} \left(y_i - x_i^T \lambda_v - \frac{B_i - x_i^T \lambda_v A_i}{A_i} \right),$$

$$\frac{\partial l}{\partial \sigma_v} = \frac{1}{\sigma_1^3} \sum_{i=1}^{N} z_i s_i \left[(y_i - x_i^T \lambda_v)^2 - \frac{C_i - 2 x_i^T \lambda_v B_i + (x_i^T \lambda_v)^2 A_i}{A_i} \right],$$

where A_i, B_i, C_i, and M_i are as previously defined except indexed by η_1 rather than $(\alpha, \beta, \eta_0^a)$.

Besides multiplying all of the above formulas by $I(L \leq y_i \leq U)$, in the presence of censoring the following is added to the score equations:

$$\text{to } \frac{\partial l}{\partial \alpha_{1j}}, \sum_{i=1}^{N_v} z_i s_i x_{ij} I(y_i < L) \left(\frac{M_i(L)}{A_i(L)} - \frac{M_i}{A_i} \right) - \sum_{i=1}^{N} z_i s_i x_{ij} I(y_i > U) \left(\frac{M_i(U) A_i - A_i(U) M_i}{A_i (A_i - A_i(U))} \right);$$

$$\text{to } \frac{\partial l}{\partial \beta_v}, \sum_{i=1}^{N_v} z_i s_i x_{ij} I(y_i < L) \left(\frac{P_i(L)}{A_i(L)} - \frac{P_i}{A_i} \right) - \sum_{i=1}^{N} z_i s_i x_{ij} I(y_i > U) \left(\frac{P_i(U) A_i - A_i(U) P_i}{A_i (A_i - A_i(U))} \right);$$

$$\text{to } \frac{\partial l}{\partial \lambda_j}, \sum_{i=1}^{N_v} z_i s_i x_{ij} I(y_i < L) \left(\frac{V_i(L)}{A_i(L)} - \frac{V_i}{A_i} \right) - \sum_{i=1}^{N} z_i s_i x_{ij} I(y_i > U) \left(\frac{V_i(U) A_i - A_i(U) V_i}{A_i (A_i - A_i(U))} \right);$$

$$\text{to } \frac{\partial l}{\partial \sigma}, \sum_{i=1}^{N_v} z_i s_i x_{ij} I(y_i < L) \left(\frac{T_i(L)}{A_i(L)} - \frac{T_i}{A_i} \right) - \sum_{i=1}^{N} z_i s_i x_{ij} I(y_i > U) \left(\frac{T_i(U) A_i - A_i(U) T_i}{A_i (A_i - A_i(U))} \right);$$

where

$$P_i = \int (y \frac{w(x_i, y)}{1 + exp(x_i^T \alpha + \beta y)} f_p(y|x_i) dy$$

$$V_i = \frac{B_i - x_i^T \lambda_v A_i}{\sigma_v^2}$$

$$T_i = \frac{C_i - 2 x_i^T \lambda_v B_i + (x_i^T \lambda_v)^2 A_i}{\sigma_v^2} - \frac{A_i}{\sigma_v};$$

and $A_i(L), M_i(L), P_i(L)$, etc. are A_i, M_i, P_i, etc. except the integrals are evaluated from $-\infty$ to L (as opposed to $-\infty$ to ∞).

In B.2 I outlined the use of the delta method. The function

$$\begin{aligned} h_0(\alpha, \eta_0^a) &\equiv E(Y(0)|S(0) = S(1) = 1, X = x) \\ &= \frac{\int yw\,(x, y; \beta, \alpha)\, f_p(y|x; \eta_0^a)dy}{\int w\,(x, y; \beta, \alpha)\, f_p(y|x; \eta_0^a)dy(0)}. \end{aligned}$$

Under models (5.1)-(5.2) and (5.4), then

$$h_0'(\alpha, \eta_0^a) = \left[\frac{\partial h_0}{\partial \alpha_0}, \cdots, \frac{\partial h_0}{\partial \alpha_q}, \frac{\partial h_0}{\partial \lambda_0}, \cdots, \frac{\partial h_0}{\partial \lambda_q}, \frac{\partial h_0}{\partial \sigma}\right]^T,$$

where

$$\begin{aligned} \frac{\partial h_0}{\partial \alpha_j} &= x_j \frac{AP - BM}{A^2} \\ \frac{\partial h_0}{\partial \lambda_j} &= x_j \frac{A(C - x^T\lambda B) - B(B - x^T\lambda A)}{\sigma^2 A^2} \\ \frac{\partial h_0}{\partial \sigma} &= \frac{AK - BQ}{A^2} \end{aligned}$$

for $j = 0, \cdots, q$ and where A, C, D, M, P, and Q are as previously defined for $x_i = x$, and

$$K = \frac{1}{\sigma^3}\int y(y - x^T\lambda)^2 w(x, y) f_p(y|x)dy - \frac{1}{\sigma}B.$$

Programs:

/research/gilbert/simulations/newsims/programs

D.3 Closed form for $VE(x; \cdot)$ under (5.2) and (5.5)

Proof.

$$\begin{aligned} f_p^*(y|x; \cdot) &= \frac{w(x, y; \alpha, \beta)^{-1} f_p^{ai}(y|x; \eta_0^b)}{\int w(x, y; \alpha, \beta)^{-1} f_p^{ai}(y|x; \eta_0^b)dy} \\ &= \frac{(1 + exp(-x^T\alpha_0 - \beta y)) f_p^{ai}(y|x; \cdot)}{\int (1 + exp(-x^T\alpha_0 - \beta u)) f_p^{ai}(u|x; \cdot)du} \\ &= k(x; \beta, \alpha, \eta_0^b)\left[f_p^{ai}(y|x; \cdot) + \frac{1}{\sqrt{2\pi}\sigma_p} exp\left(\frac{-1}{2\sigma_p^2}(y - x^T\gamma_p)^2 - x^T\alpha - \beta y\right)\right] \\ &= k(x; \cdot)\left[f_p^{ai}(y|x; \cdot) + \frac{1}{\sqrt{2\pi}\sigma_p} e^{\frac{-1}{2\sigma_p^2}(y - (x^T\gamma_p - \beta\sigma_p^2))^2} e^{\beta(\frac{1}{2}\beta\sigma_p^2 - x^T\gamma_p) - x^T\alpha}\right] \\ &= k(x; \cdot)\left[f_p^{ai}(y|x; \cdot) + f_p^{prot}(y|x; \cdot) e^{\beta(\frac{1}{2}\beta\sigma_p^2 - x^T\gamma_p) - x^T\alpha}\right] \\ &= (1 - VE(x; \cdot)) f_p^{ai}(y|x; \cdot) + VE(x; \cdot) f_p^{prot}(y|x; \cdot), \end{aligned}$$

where

$$\begin{aligned} k(x;\beta,\alpha,\eta_0^b) &= 1 - VE((x;\beta,\alpha,\eta_0^b) \\ &= \left[\int f_p^{ai}(u|x;\cdot) + f_p^{prot}(u|x;\cdot)e^{\beta(\frac{1}{2}\beta\sigma_p^2 - x^T\gamma_p) - x^T\alpha}du\right]^{-1} \\ &= \left[1 + e^{\beta(\frac{1}{2}\beta\sigma_p^2 - x^T\gamma_p) - x^T\alpha}\right]^{-1}. \end{aligned}$$

□

D.4 Log-likelihood and Score Equations under $\mathcal{M}_b$

For computational convenience, instead of parameterizing as written by (5.5), I chose the following equivalent parameterization:

$$f_p^{ai}(y|x;\eta_0^b) = \phi(y;x^T\gamma,\sigma_p^2) \text{ and } f_v(y|x;\eta_1) = \phi(y;x^T\gamma + zx^T\gamma_v,\sigma_v^2), \tag{D.1}$$

where $\eta_0^b = (\gamma,\sigma_p)$, $\eta_1 = (\gamma,\gamma_v,\sigma_v)$, and γ_v therefore represents the interaction terms. This parameterization allows flexibility in the number of interaction terms. To shorten notation, $A_i \equiv 1 - VE(x_i)$, $VE_i \equiv VE(x_i)$, $\theta_i \equiv \theta_p(x_i)$, $f_p^{prot} \equiv f_p^{prot}(y_i|x_i)$, $\int_{-\infty}^{L} f_p^{prot}(y|x_i)dy \equiv F_p^{prot}(L)$, etc.

The log-likelihood is

$$\begin{aligned} l(\mu,\alpha,\eta_0^b,\eta_1|H) = &\sum_{i=1}^{N} s_i log(\theta_p(x_i)) + \sum_{i=1}^{N} z_i(1-s_i)log(1-\theta_p(x_i)(1-VE(x_i)) \\ &+ \sum_{i=1}^{N} z_i s_i log(1-VE(x_i)) + \sum_{i=1}^{N}(1-z_i)(1-s_i)log(1-\theta_p(x_i)) \\ &+ \sum_{i=1}^{N} z_i s_i(-log(\sigma_v) - (y_i - \gamma^T x_i - z_i\gamma_v^T x_i)^2/(2\sigma_v^2)) \\ &+ \sum_{i=1}^{N}(1-z_i)s_i log[(1-VE(x_i))f_p^{ai}(y_i|x_i) + VE(x_i)f_p^{prot}(y_i|x_i)] + k. \end{aligned} \tag{D.2}$$

In the presence of censoring, the last two terms of (D.2) are multiplied by $I(L \le y_i \le U)$

inside the summation, and the following is added to the log-likelihood:

$$\sum_{i=1}^{N} z_i s_i I(y_i < L) log(F_v(L)) + \sum_{i=1}^{N} z_i s_i I(y_i > U) log(1 - F_v(U))$$
$$+ \sum_{i=1}^{N} (1 - z_i) s_i I(y_i < L) log(A_i F_p^{ai}(L) + VE_i F_p^{ai}(U))$$
$$+ \sum_{i=1}^{N} (1 - z_i) s_i I(y_i > U) log(1 - A_i F_p^{ai}(L) - VE_i F_p^{ai}(U)).$$

The score equations are

$$\frac{\partial l}{\partial \mu_j} = \sum_{i=1}^{N} z_i s_i x_{ij} (1 - \theta_i) - \sum_{i=1}^{N} z_i (1 - s_i) x_{ij} \frac{\theta_i (1 - \theta_i) A_i}{1 - \theta_i A_i} + \sum_{i=1}^{N} (1 - z_i) x_{ij} (s_i - \theta_i)$$

$$\frac{\partial l}{\partial \alpha_j} = \sum_{i=1}^{N} z_i s_i x_{ij} VE_i - \sum_{i=1}^{N} z_i (1 - s_i) x_{ij} \frac{VE_i A_i \theta_i}{1 - \theta_i A_i} + \sum_{i=1}^{N} (1 - z_i) s_i x_{ij} \frac{VE_i A_i (f_p^{ai} - f_p^{prot})}{A_i f_p^{ai} + VE_i f_p^{prot}}$$

$$\frac{\partial l}{\partial \gamma_j} = \beta \sum_{i=1}^{N} z_i s_i x_{ij} VE_i - \beta \sum_{i=1}^{N} z_i (1 - s_i) x_{ij} \frac{VE_i A_i \theta_i}{1 - \theta_i A_i} + \sum_{i=1}^{N} z_i s_i x_{ij} \frac{y_i - \gamma^T x_i - z_i \gamma_v^T x_i}{\sigma_v^2}$$
$$+ \sum_{i=1}^{N} (1 - z_i) s_i x_{ij} \frac{1}{A_i f_p^{ai} + VE_i f_p^{prot}} \left\{ A_i f_p^{ai} \left[\beta VE_i + \frac{1}{\sigma_p^2} (y_i - \gamma^T x_i) \right] \right.$$
$$\left. - VE_i f_p^{prot} \left[\beta A_i - \frac{1}{\sigma_p^2} (y_i - \gamma^T x_i + \beta \sigma_p^2) \right] \right\}$$

$$\frac{\partial l}{\partial \gamma_{vj}} = \sum_{i=1}^{N} z_i s_i x_{ij} \frac{y_i - \gamma^T x_i - \gamma_v^T x_i}{\sigma_v^2}$$

$$\frac{\partial l}{\partial \sigma_p} = \beta^2 \sigma_p \sum_{i=1}^{N} \frac{(1 - z_i) s_i}{A_i f_p^{ai} + VE_i f_p^{prot}} \left\{ A_i f_p^{ai} \left[\frac{(y_i - \gamma^T x_i)^2 - \sigma_p^2}{\beta^2 \sigma_p^4} - VE_i \right] \right.$$
$$+ VE_i f_p^{prot} \left[\frac{(y_i - \gamma^T x_i + \beta \sigma_p^2)^2 - \sigma_p^2}{\beta^2 \sigma_p^4} \right.$$
$$\left. \left. + A_i - \frac{2\beta (y_i - \gamma^T x_i + \beta \sigma_p^2)}{\beta^2 \sigma_p^2} \right] \right\}$$
$$- \beta^2 \sigma_p \sum_{i=1}^{N} z_i s_i VE_i + \beta^2 \sigma_p \sum_{i=1}^{N} z_i (1 - s_i) \frac{VE_i A_i \theta_i}{1 - \theta_i A_i}$$

$$\frac{\partial l}{\partial \sigma_v} = \sum_{i=1}^{N} z_i s_i \left[\frac{(y_i - \gamma^T x_i - \gamma_v^T x_i)^2}{\sigma_v^3} - \frac{1}{\sigma_v} \right]$$

for $j = 0, \cdots, q$, with $x_i = \{1, x_{i1}, \cdots, x_{iq}\}$.

These score equations can be easily modified to account for models assuming $\sigma_p = \sigma_v$. In the presence of censoring, $I(L \leq y_i \leq U)$ is multiplied inside all summations which include y_i (e.g., the last term in $\frac{\partial l}{\partial \alpha_j}$, the last two terms in $\frac{\partial l}{\partial \gamma_j}$, etc.). In the presence of censoring, the following is also added:

to $\dfrac{\partial l}{\partial \alpha_j}$,

$$\sum_{i=1}^{N}(1-z_i)s_i x_{ij} V E_i A_i \left[\frac{I(y_i < L)(F_p^{ai}(L) - F_p^{prot}(L))}{A_i F_p^{ai}(L) + V E_i F_p^{prot}(L)} - \frac{I(y_i > U)(F_p^{ai}(U) - F_p^{prot}(U))}{1 - A_i F_p^{ai}(U) - V E_i F_p^{prot}(U)} \right];$$

to $\dfrac{\partial l}{\partial \gamma_j}$,

$$\sum_{i=1}^{N} z_i s_i x_{ij} \left[I(y_i > U)\frac{f_v(U|x_i)}{1-F_v(U)} - I(y_i < L)\frac{f_v(L|x_i)}{F_v(L)} \right]$$
$$+ \sum_{i=1}^{N} \frac{(1-z_i)s_i x_{ij} I(y_i < L)}{A_i F_p^{ai}(L) + V E_i F_p^{prot}(L)} \left[\beta V E_i A_i \left(F_p^{ai}(L) - F_p^{prot}(L) \right) - A_i f_p^{ai}(L|x_i) - V E_i f_p^{prot}(L|x_i) \right]$$
$$- \sum_{i=1}^{N} \frac{(1-z_i)s_i x_{ij} I(y_i > U)}{1 - A_i F_p^{ai}(U) - V E_i F_p^{prot}(U)} \left[\beta V E_i A_i \left(F_p^{ai}(U) - F_p^{prot}(U) \right) - A_i f_p^{ai}(U|x_i) - V E_i f_p^{prot}(U|x_i) \right];$$

to $\dfrac{\partial l}{\partial \gamma_{vj}}$,

$$\sum_{i=1}^{N} z_i s_i x_{ij} \left[I(y_i > U)\frac{f_v(U|x_i)}{1-F_v(U)} - I(y_i < L)\frac{f_v(L|x_i)}{F_v(L)} \right];$$

to $\dfrac{\partial l}{\partial \sigma_v}$,

$$\sum_{i=1}^{N} z_i s_i I(y_i < L)\frac{\frac{1}{\sigma_v^3}R_v(L) - \frac{1}{\sigma_v}F_v(L)}{F_v(L)} - \sum_{i=1}^{N} z_i s_i I(y_i > U)\frac{\frac{1}{\sigma_v^3}R_v(U) - \frac{1}{\sigma_v}F_v(U)}{1-F_v(U)};$$

to $\frac{\partial l}{\partial \sigma_p}$,

$$\sum_{i=1}^{N}(1-z_i)s_iI(y_i<L)\frac{\beta^2\sigma_pA_iVE_i\left(F_p^{prot}(L)-F_p^{ai}(L)\right)+A_i\left(\frac{1}{\sigma_p^3}R_p^{ai}(L)-\frac{1}{\sigma_p}F_p^{ai}(L)\right)}{A_iF_p^{ai}(L)+VE_iF_p^{prot}(L)}$$

$$+\sum_{i=1}^{N}(1-z_i)s_iI(y_i<L)\frac{VE_i\left(\frac{1}{\sigma_p^3}R_p^{prot}(L)-\frac{1}{\sigma}F_p^{prot}(L)+2\beta\sigma_pf_p^{prot}(L|x_i)\right)}{A_iF_p^{ai}(L)+VE_iF_p^{prot}(L)}$$

$$-\sum_{i=1}^{N}\frac{(1-z_i)s_iI(y_i>U)}{1-A_iF_p^{ai}(U)-VE_iF_p^{prot}(U)}\left[A_i\left(\frac{1}{\sigma_p^3}R_p^{ai}(U)-\frac{1}{\sigma_p}F_p^{ai}(U)\right)\right.$$

$$\left.+\beta^2\sigma_pA_iVE_i\left(F_p^{prot}(U)-F_p^{ai}(U)\right)\right]$$

$$-\sum_{i=1}^{N}(1-z_i)s_iI(y_i>U)\frac{VE_i\left(\frac{1}{\sigma_p^3}R_p^{prot}(U)-\frac{1}{\sigma}F_p^{prot}(U)+2\beta\sigma_pf_p^{prot}(U|x_i)\right)}{1-A_iF_p^{ai}(U)-VE_iF_p^{prot}(U)};$$

where

$R_p^{ai}(L)\equiv\int_{-\infty}^{L}(y-\gamma^Tx_i)^2f_p^{ai}(y|x_i)dy=\frac{\sigma_p^2}{2}(1-(-1)^{I(L>\gamma^Tx_i)}GAM((L-\gamma^Tx_i)^2,3/2,2\sigma_p^2)$, $R_p^{prot}(L)\equiv\int_{-\infty}^{L}(y-\gamma^Tx_i+\beta\sigma_p^2)^2f_p^{prot}(y|x_i)dy=\frac{\sigma_p^2}{2}(1-(-1)^{I(L>\gamma^Tx_i-\beta\sigma_p^2)}GAM((L-\gamma^Tx_i+\beta\sigma_p^2)^2,3/2,2\sigma_p^2)$, etc., with $GAM(a,b,c)$ denoting the cdf of a gamma distribution evaluated at a with shape parameter $=b$ and scale parameter $=c$.

Proof.

$$\begin{aligned}
R_p^{ai}(L)&\equiv\int_{-\infty}^{L}(y-\gamma^Tx_i)^2f_p^{ai}(y|x_i;\cdot)dy\\
&=\frac{1}{\sqrt{2\pi}\sigma_p}\int_{-\infty}^{L}\frac{1}{2}2(y-\gamma^Tx_i)(y-\gamma^Tx_i)exp\left(-\frac{1}{2\sigma_p^2}(y-\gamma^Tx_i)^2\right)dy\\
&=\frac{1}{2\sqrt{2\pi}\sigma_p}\int u^{1/2}exp\left(-\frac{1}{2\sigma_p^2}u\right)du\\
&=\frac{\Gamma\left(\frac{3}{2}\right)\left(2\sigma_p^2\right)^{3/2}}{2\sqrt{2\pi}\sigma_p}\int\frac{u^{3/2-1}e^{-u/(2\sigma_p^2)}}{\Gamma\left(\frac{3}{2}\right)\left(2\sigma_p^2\right)^{3/2}}du,\\
&\qquad\text{where } u \text{ is integrated over } ((L-\gamma^Tx_i)^2,\infty) \text{ if } L<\gamma^Tx_i,\\
&\qquad\text{else over } (0,\infty)+(0,(L-\gamma^Tx_i)^2);\\
&=\frac{\sigma_p^2}{2}(1-(-1)^{I(L>\gamma^Tx_i)}GAM((L-\gamma^Tx_i)^2,3/2,2\sigma_p^2).
\end{aligned}$$

□

D.5 Identification of Sensitivity Parameters

If in assumption A.2 the parameter β of model (2.3) is regarded as unknown, it is not identified under (2.1), (2.2), and assumptions A.1 and A.2. However, β is identified if the distributional forms imposed by M.1-M.3 and M.4b are also assumed. Consequently, under $\mathcal{M}_b$ rather than regarding β as fixed and known, one could estimate it. (This also comes up in the missing data literature – see Little (1995); Scharfstein et al. (1999)). In response to pieces of this dissertation submitted for publication, one referee stated, "estimating parameters identifiable from the data is not a stupid thing to do." On this point I disagree. Under $\mathcal{M}_b$, β is only identified because of models imposed to reduce the data dimensionality.

Consider the case where there are no covariates. In this case, one only needs to assume A.1 and A.2. Viral load distributions can be left unspecified and sensitivity analyses can be performed over a plausible range for β. This was the approach of GBH. Alternatively, one could additionally assume a form for the distribution of viral loads in the always infected principal stratum in the placebo arm (distributional assumption M.4b). This is model $\mathcal{M}_b$ without covariates. Under this model, some of the selected β might yield a distribution for the viral loads among infected placebos that is not compatible with the observed viral load data. It is true that one could argue that either A.2 is misspecified (possibly a poor choice of β) or M.4b is misspecified. We do not know which one. However, since M.4b was an optional assumption (in this case with no covariates), rather than ruling out certain values of β that we a priori thought were plausible, we would instead simply discard M.4b and perform the analysis using GBH. The range for β is something that we want chosen independent of the data.

When multiple or continuous covariates are included, it becomes necessary to make some modeling assumptions. One choice is $\mathcal{M}_b$. However, if certain model choices and plausible values for β are incompatible with the observed data, then similar to the situation without covariates we take this to imply that the distributional assumptions are misspecified rather than β. Our distributional assumptions were only made so that we could incorporate covariates in the analysis.

It is important that we do not solve the inherent problem of identification of $ACE(x)$ by

assuming it away in the technical details of the model. This is also important in the Bayesian context. It is conceivable that one might want to put a distribution on the sensitivity parameter β – integrating over this distribution to obtain a single estimate for $ACE(x)$. However, if one puts a prior distribution on β, one would want the posterior distribution of β to be equal to the prior. If the two are not equal (as would be the case under $\mathcal{M}_b$) then as stated before, this imples that information on β is obtained because of modeling assumptions.

Appendix E

TECHNICAL DETAILS OF CHAPTER 6

E.1 Additional Restriction on Observed Data

$$
\begin{aligned}
& f_{C|T,S,Z,X}(y|T, S=1, Z=z, X) \\
& \quad = f_{C(z)|T,S(z)=1,Z,X,T}(y|T(0)=t, S(z)=1, Z=z, X) && \text{by (2.1)} \\
& \quad = f_{C(z)|S(z)=1,Z,X}(y|S(z)=1, Z=z, X=x) && \text{by (6.4)} \\
& \quad = f_{C(z)|S(z)=1,X}(y|S(z0=1, X=x) && \text{by (6.1)} \\
& \quad \equiv g_z(y|x).
\end{aligned}
$$

As stated in chapter 2, consistency (2.1) follows from SUTVA.

E.2 GBH Extension is Semi-parametric MLE

Consider the case where there are no covariates and the distribution of the censored time-to-event outcome is left unspecified. Assume $\mathcal{A}$ with $w(\cdot)$ specified so that for $y > \tau$, $w(\cdot)$ is constant. With $w(\cdot)$ specified by (6.9) or (6.10), the semiparametric ML estimators for $F_p^{ai}(y)$ are equivalent to the estimators of GBH, only estimating $F_p(y)$ with Kaplan-Meier estimates.

Proof. Under the scenario described above, the likelihood can be written as

$$
\begin{aligned}
\mathcal{L}(\rho) \quad \propto \quad & \prod_{i=1}^{N} \left\{ \left[f_v(y_i)^{\delta_i} \theta_p (1 - F_v(y_i))^{1-\delta_i} \int w(y;\beta,\alpha) f_p(y) dy \right]^{s_i} \right. \\
& \times \left. \left[1 - \theta_p \int w(y;\beta,\alpha) f_p(y) dy \right]^{1-s_i} \right\}^{z_i} \\
& \times \left\{ \left[\theta_p f_p(y_i)^{\delta_i} (1 - F_p(y_i))^{1-\delta_i} \right]^{s_i} (1-\theta_p)^{1-s_i} \right\}^{1-z_i}. \qquad \text{(E.1)}
\end{aligned}
$$

Again, since

$$\left[f_v(y_i)^{\delta_i}\left(1-F_v(y_i)\right)^{1-\delta_i}\right]^{s_i z_i}$$

factors out of the likelihood we can treat this part separately. It is well known that the Kaplan-Meier estimates of $F_v(y)$ are the nonparametric ML estimates.

Similar to Chapter 4, to make things simpler I am going to order the data in the following manner: 1) those randomized to placebo and infected, with Y_i ordered; 2) those randomized to placebo and not infected; 3) those randomized to vaccine. My notation is not going to reflect this change of ordering (just to keep things simple) because it has no effect on the likelihood. The part of the likelihood we are interested in is now written as

$$\begin{aligned}\mathcal{L} \;\propto\; & \prod_{i=1}^{n_p}\left[f_p(y_i)^{\delta_i}\left(1-F_p(y_i)\right)^{1-\delta_i}\right]^{s_i(1-z_i)} \times \prod_{i=1}^{N_p}\left[\theta_p^{s_i}(1-\theta_p)^{1-s_i}\right] \\ & \times \prod_{i=N_p+1}^{N}\left\{\left[\theta_p\int w(y;\beta,\alpha)f_p(y)dy\right]^{s_i}\left[1-\theta_p\int w(y;\beta,\alpha)f_p(y)dy\right]^{1-s_i}\right\}^{z_i}.\end{aligned}$$

Consider estimation of f_p. Probabilitiy mass will only be put on the set of points $\{y_{k_1},\cdots,y_{k_{m+1}}\}$ where $y_{k_{m+1}}$ is some time later than y_{n_p} if $\delta_{n_p}=0$ and y_{n_p} if $\delta_{n_p}=1$. This implies that there are m distinct failure times. These points, y_{k_j}, correspond to the non-censored failure times where $(1-z_i)s_i=1$. Let $(p_{k_1},\cdots,p_{k_m})$ be the parameters we want to estimate where $p_{k_i}=f_p(y_{k_i})$. Those that are censored will also have their probability mass put on these points. We use an EM-type approach: The "complete" data would tell us when censored observations would have had their event. So the "complete" likelihood could be written as follows:

$$\begin{aligned}\mathcal{L}^c \;\propto\; & \prod_{j=1}^{m+1} p_{k_j}^{n_j} \times \prod_{i=1}^{N_p}\left[\theta_p^{s_i}(1-\theta_p)^{1-s_i}\right]^{1-z_i} \\ & \times \prod_{i=N_v+1}^{N}\left\{\left[\theta_p\sum_{j=1}^{m+1} w(y_{k_j};\beta,\alpha)p_{k_j}\right]^{s_i}\left[1-\theta_p\sum_{j=1}^{m+1} w(y_{k_j};\beta,\alpha)p_{k_j}\right]^{1-s_i}\right\}^{z_i},\end{aligned}$$

where n_j is the unknown number of events that happen at time point k_j. It then follows

that the "complete" log-likelihood is

$$\begin{aligned} l^c &= C + \sum_{j=1}^{m+1} n_j log\left(p_{k_j}\right) + n_p log(\theta_p) + (N_p - n_p) log(1-\theta_p) + n_v log(\theta_p) \\ &\quad + n_v log\left(1 - VE\right) + (N_v - n_v) log\left(1 - \theta_p(1-VE)\right), \end{aligned} \tag{E.2}$$

where $VE \equiv 1 - \sum_{j=1}^{m+1} w(y_{k_j}; \beta, \alpha) p_{k_j}$.

Notice that the "complete" log-likelihood is of the same form as (4.2) in Chapter 4, the log-likelihood in the uncensored case. With $w(\cdot)$ specified by either (6.9) or (6.10), one can follow the same steps shown in Chapter 4 to show that the ML estimators based on this "complete" data log-likelihood are as follows:

For $n_v/N_v > n_p/N_p$:

$$\begin{aligned} &\hat{\theta}_p = \frac{n_p + n_v}{N_p + N_v} \\ &\hat{p}_{k_i} = \frac{n_i}{n_p} \qquad \text{for } i = 1, \cdots, m+1 \\ &\hat{\alpha} = \infty \Rightarrow \widehat{VE} = 0. \end{aligned}$$

For $n_v/N_v \leq n_p/N_p$:

$$\begin{aligned} &\hat{\theta}_p = \frac{n_p}{N_p} \\ &\hat{p}_{k_i} = \frac{n_i}{n_p} \qquad \text{for } i = 1, \cdots, m+1 \\ &\hat{\alpha} \text{ such that } \sum_{i=1}^{n_p} w(y_i; \hat{\alpha}, \beta) \hat{p}_{pi} = \frac{n_v N_p}{N_v n_p} = 1 - \widehat{VE}. \end{aligned}$$

We have just performed the maximization step of the EM-algorithm. Since we do not know n_i, we estimate it with its expectation [the following steps are taken almost directly from my Stat 582 notes]:

$$E(n_i) = \sum_{j=1}^{n_p} P(X_j = y_{k_i} | \delta_j, y_j),$$

where X_j is the observed time (minimum of time until censoring or the event). Therefore,

$$\begin{aligned}
\hat{n}_1 &= 1 + (k_1 - 1)\hat{p}_{k_1} \\
\hat{n}_2 &= 1 + (k_1 - 1)\hat{p}_{k_1} + (k_2 - k_1 - 1)\frac{\hat{p}_{k_2}}{1 - \hat{p}_{k_1}} \\
&\vdots \\
\hat{n}_j &= 1 + (k_1 - 1)\hat{p}_{k_j} + \cdots + (k_j - k_{j-1} - 1)\frac{\hat{p}_{k_j}}{1 - \hat{p}_{k_1} - \cdots - \hat{p}_{k_j}}
\end{aligned}$$

It can then be shown that

$$\frac{\hat{p}_{k_j}}{\sum_{i=j}^{m+1} \hat{p}_{k_i}} = \frac{1}{n_p - k_j + 1},$$

for $j = 1, \cdots, m$ result in a stationary point of the EM algorithm. Hence, the ML estimator for $F_p(t)$ is

$$\hat{F}_p(t) = \sum_{i:y_{k_i} \leq t} \hat{p}_{k_i},$$

which can be equivalently written as the Kaplan-Meier estimator

$$1 - \hat{F}_p(t) = \prod_{i:y_{k_i} \leq t} \left(1 - \frac{1}{n_p - k_j + 1}\right).$$

□

As stated at the end of Section 4.1, in the uncensored case maximum likelihood estimators based on more general $w(\cdot)$ can be constructed from the score equations (4.3)-(4.5). The same is true in the censored case. Notice that the above proof was written in terms of a general $w(\cdot)$ up until after the "complete" log-likelihood was written. To construct the ML estimator for a different $w(\cdot)$, one writes the "complete" log-likelihood (E.2), differentiates with respect to the parameters to obtain score equations of the "complete" data, and then solves to obtain MLEs in the "complete" data scenario. Then one performs the E-step (computes the expectation of this "complete" data MLE) in a manner similar to that described in the preceding proof.

E.3 Asymptotic Normality of GBH Extension

Suppose that data are collected over a finite interval $[0, \tau]$ where τ is fixed as the sample size $N \to \infty$. Assume $\mathcal{A}$, $0 < VE < 1, p_0 > 0$, and let $t \in [0, \tau]$. (Recall that $VE = 1 - Pr(S_i Z_i = 1)/Pr(S_i(1 - Z_i) = 1)$.) Furthermore, assume $0 < w(t; \alpha, \beta) < 1$ for all $t > 0, \alpha$, and β; is constant for $t > \tau$; and is twice continuously differentiable with respect to α, with a bounded second derivative. Then $\hat{F}_p^{ai}(t)$ is consistent and asymptotically normal.

Proof. The estimates $(\hat{p}_0, \hat{\alpha})$ are the solutions to the estimating equation

$$\sum_{i=1}^{N} \psi_i(p_0, \alpha) = 0,$$

where

$$\psi_i(p_0, \alpha) = \begin{cases} (1 - Z_i)(S_i - p_0) \\ Z_i \left(S_i - p_0 \int_0^\infty w(t; \alpha, \beta) d\hat{F}_p(t) \right). \end{cases}$$

It is easily seen that the unique solution to the first equation is $\hat{p}_0 = n_p/N_p$. The second equation can therefore be written as

$$U_N(\alpha) \equiv \int_0^\infty w(t; \alpha, \beta) d\hat{F}_p(t) - \left(1 - \widehat{VE}\right) = 0.$$

(Recall that $\widehat{VE} = 1 - n_v N_p/(N_v n_p)$.)

The process $\hat{F}_p(\cdot)$ in $D[0, \tau]$ and $\widehat{VE}$ are jointly asymptotically normal, where $D[0, \tau]$ is the space of cadlag functions (i.e., right continuous whose limits from the left exist) on $[0, \tau]$. This is shown by looking at the empirical distribution of the observed data

$O_i = (Z_i, S_i, Y_i, \Delta_i)$. Consider the following pieces of the empirical cdf:

$$
\begin{aligned}
a &= \mathbb{P}_N(S_i > 0, Z_i < 1, Y_i < \infty, \Delta_i \leq 1), \\
b &= \mathbb{P}_N(S_i \leq 1, Z_i < 1, Y_i < \infty, \Delta_i \leq 1), \\
c &= \mathbb{P}_N(S_i > 0, Z_i > 0, Y_i < \infty, \Delta_i \leq 1), \\
d &= \mathbb{P}_N(S_i \leq 1, Z_i > 0, Y_i < \infty, \Delta_i \leq 1), \\
e &= \mathbb{P}_N(S_i > 0, Z_i < 1, Y_i \leq t, \Delta_i > 0), \\
f &= \mathbb{P}_N(S_i > 0, Z_i < 1, Y_i \leq t + h, \Delta_i > 0), \\
g &= \mathbb{P}_N(S_i > 0, Z_i < 1, Y_i \leq t, \Delta_i \leq 1).
\end{aligned}
$$

These pieces of the empirical cdf are jointly asymptotically normal. The estimate $\widehat{VE}$ and the Nelson-Aalen estimator of $F_p(t)$, labeled as $\hat{F}_p^*(t)$, can be written as Hadamard (compact)-differentiable maps of these pieces of the empirical cdf:

$$
\begin{aligned}
\widehat{VE} &= \frac{cb}{da}, \\
\hat{F}_p^*(t) &= 1 - exp\left[- \int_0^t \frac{\lim_{h \to 0} \frac{f-e}{h}}{1-g} \right].
\end{aligned}
$$

That the integration is Hadamard (compact)-differentiable is a result of Lemma 20.10 of van der Vaart (1998). Hence $(\widehat{VE}, \hat{F}_p(\cdot))$, which is asymptotically equivalent to $(\widehat{VE}, \hat{F}_p^*(\cdot))$, converges weakly to a Gaussian process.

Now consider the following maps:

$$
\begin{aligned}
\left(\hat{F}_p(\cdot), \widehat{VE}\right) &\mapsto \left(\int_{[0,\tau]} w(t;\alpha,\beta) d\hat{F}_p(t), w(\tau;\alpha,\beta)\left(1 - \hat{F}_p(\tau)\right), 1 - \widehat{VE} \right) \\
&\mapsto \int_{[0,\tau]} w(t;\alpha,\beta) d\hat{F}_p(t) + w(\tau;\alpha,\beta)\left(1 - \hat{F}_p(\tau)\right) - \left(1 - \widehat{VE}\right) \\
&\equiv U_N(\alpha).
\end{aligned}
$$

Each of these maps is Hadamard differentiable; that the first map is Hadamard-differentiable follows from Problem 7, Chapter 20, van der Vaart (1998). (It is also a consequence of Proposition II.8.6 of Andersen, Borgan, Gill, & Keiding (1993)). Therefore, by the functional delta method, $U_N(\alpha)$ is asymptotically normal.

Define the following terms:

$$\begin{aligned} U(\alpha) &\equiv \int_0^\infty w(t;\alpha,\beta)dF_p(t) - (1-VE) \\ &= \int_{[0,\tau]} w(t;\alpha,\beta)dF_p(t) + w(\tau;\alpha,\beta)\left(1-F_p(\tau)\right) - (1-VE), \text{ and} \\ U'(\alpha) &\equiv \int_0^\infty w'(t;\alpha,\beta)dF_p(t) \\ &= \int_0^t w'(t;\alpha,\beta)dF_p(t) + w'(\tau;\alpha,\beta)\left(1-F_p(\tau)\right), \end{aligned}$$

where $w'(t;\alpha,\beta)$ is the partial derivative of $w(\cdot)$ taken with respect to α. Because $\widehat{VE} \to^P VE$ and $\hat{F}_p(t) \to^P F_p(t)$ for all $t \in [0,\tau]$, therefore $U_N(\alpha) \to^P U(\alpha)$ and $U'_N(\alpha) \to^P U'(\alpha)$ by the continuous mapping theorem. Notice that the map $\alpha \mapsto U_N(\alpha)$ is continuous and non-decreasing with $U_N(\hat{\alpha}) = 0$. Let α_0 be the point such that $U(\alpha_0 - \varepsilon) < 0 < U(\alpha_0 + \varepsilon)$ for every $\varepsilon > 0$. Therefore, by Lemma 5.10 of van der Vaart (1998), $\hat{\alpha} \to^P \alpha_0$.

From here on, for notational convenience we will set $\alpha = \alpha_0$. Consider the Taylor expansion:

$$0 = U_N(\hat{\alpha}) = U_N(\alpha) + U'_N(\alpha)(\hat{\alpha}-\alpha) + \frac{1}{2}U''_N(\alpha^*)(\hat{\alpha}-\alpha)^2,$$

which implies

$$\sqrt{N}(\hat{\alpha}-\alpha) = \frac{-\sqrt{N}U_N(\alpha)}{U'_N(\alpha) + \frac{1}{2}U''_N(\alpha^*)(\hat{\alpha}-\alpha)^2},$$

where α^* is some value between α and $\hat{\alpha}$, and $U''_N(\alpha^*) = \int_0^\infty w''(t;\alpha^*,\beta)d\hat{F}_p(t)$. Notice that $U''_N(\alpha^*)$ is bounded because $|w''(t;\alpha,\beta)|$ is bounded. Therefore, because $U'_N(\alpha) \to^P U'(\alpha)$, $\hat{\alpha} \to^P \alpha$, $U''_N(\alpha^*)$ is bounded, and $U_N(\alpha)$ is asymptotically normal; from the Taylor expansion $\hat{\alpha}$ is asymptotically normal.

Now consider the following maps:

$$\begin{aligned} \left(\hat{\alpha}, \hat{F}_p(\cdot)\right) &\mapsto \left(\int_0^\infty w(t;\hat{\alpha},\beta)d\hat{F}_p(t), \int_0^t w(s;\hat{\alpha},\beta)d\hat{F}_p(s)\right) \\ &\mapsto \frac{\int_0^t w(s;\hat{\alpha},\beta)d\hat{F}_p(s)}{\int_0^\infty w(t;\hat{\alpha},\beta)d\hat{F}_p(t)} \equiv \hat{F}_p^{ai}(t). \end{aligned}$$

That the process $\hat{F}_p^{ai}(\cdot)$ in $D[0,\tau]$ is asymptotically normal follows from Hadamard differentiability of each of these maps, the chain rule, and the functional delta method.

□

The $w(\cdot)$ defined by (6.9) and (6.10) meet the necessary conditions for asymptotic normality. That $|w''(t;\alpha,\beta)|$ is bounded for $w(\cdot)$ defined by (6.9) (or (6.10)) is shown below:

$$\begin{aligned}
|w''(t;\alpha,\beta)| &= \left| w(t;\alpha,\beta)\frac{1-exp(\alpha+\beta t)}{(1+exp(\alpha+\beta t))^2} \right| \\
&\leq \left| \frac{1-exp(\alpha+\beta t)}{(1+exp(\alpha+\beta t))^2} \right| \\
&= \left| \frac{1}{1+exp(\alpha+\beta t)} \left(\frac{1}{1+exp(\alpha+\beta t)} - \frac{exp(\alpha+\beta t}{1+exp(\alpha+\beta t)} \right) \right| \\
&\leq \left| \frac{1}{1+exp(\alpha+\beta t)} - \frac{exp(\alpha+\beta t)}{1+exp(\alpha+\beta t)} \right| \\
&\leq 1.
\end{aligned}$$

E.4 Asymptotic Variance

Our goal is to estimate the variance of

$$\hat{F}_p^{ai}(t) \equiv \frac{\int_0^t w(s;\hat{\alpha},\beta)d\hat{F}_p(s)}{\int_0^\infty w(t;\hat{\alpha},\beta)d\hat{F}_p(t)}.$$

Recall that $\hat{F}_p(t)$ is the Kaplan-Meier estimate. Our estimating equations have been written as

$$\psi_i(p_0,\alpha) = \begin{cases} (1-Z_i)(S_i - p_0) \\ Z_i\left(S_i - p_0\int_0^\infty w(t;\alpha,\beta)d\hat{F}_p(t)\right). \end{cases}$$

When computing the variance, the tricky part is to account for the variance of the Kaplan-Meier estimates and any correlation between these estimates and the estimate of α. From Stute (1995), we learn that $\int \phi(t)d\hat{F}_p(t)$ can be written as a sum of i.i.d. terms plus some remainder term, R_{n_p}, where

$$|R_{n_p}| = o(n_p^{-1/2}),$$

and $\phi(t)$ is some well-behaved function of t. Therefore, using this result one could include additional estimating equations with zero expectation in the limit which allow us to incorporate the variance of $\hat{F}_p(\cdot)$. Define k as the number of distinct failure times, and let $t_1,\cdots,t_k$ represent the distinct ordered failure times. Because we are not making any distributional assumptions on $F_p(\cdot)$, one could think of there being $k+2$ parameters to estimate:

$(p_0, \alpha, F_p(t_1), F_p(t_2), \cdots, F_p(t_k)) \equiv \theta$. One could therefore add an additional k estimating equations as follows:

$$\psi_i(\theta) = \begin{cases} (1 - Z_i)(S_i - p_0) \\ Z_i \left(S_i - p_0 \int_0^\infty w(s; \alpha, \beta) dF_p(s)\right) \\ (1 - Z_i) S_i \left(V_{1i} - F_p(t_1)\right) \\ \qquad \vdots \\ (1 - Z_i) S_i \left(V_{ki} - F_p(t_k)\right), \end{cases}$$

with V_{ji} for $j = 1, \cdots, k$ and $i = 1, \cdots, N$ defined as

$$V_{ji} = \phi_j(Y_i)\gamma_0(Y_i)\delta_i + \gamma_{j1}(Y_i)(1 - \delta_i) - \gamma_{j2}(Y_i),$$

where

$$\begin{aligned} \phi_j(Y_i) &= I_{\{Y_i \le t_j\}}, \\ \gamma_0(Y_i) &= exp\left(\int_{-\infty}^{Y_i -} \frac{H^0(dz)}{1 - H(z)}\right), \\ \gamma_{j1}(Y_i) &= \frac{1}{1 - H(Y_i)} \int I_{\{Y_i < \omega\}} \phi_j(\omega)\gamma_0(\omega) H^1(d\omega), \\ \gamma_{j2}(Y_i) &= \int\int \frac{I_{\{\nu < Y_i, \nu < \omega\}} \phi_j(\omega)\gamma_0(\omega)}{(1 - H(\nu))^2} H^0(d\nu) H^1(d\omega), \\ H^0(y) &= \mathbb{P}_N(Y \le y, \delta = 0) = \frac{\sum (1 - z_i) s_i (1 - \delta_i) I_{\{Y_i \le y\}}}{\sum (1 - z_i) s_i}, \\ H^1(y) &= \mathbb{P}_N(Y \le y, \delta = 1) = \frac{\sum (1 - z_i) s_i \delta_i I_{\{Y_i \le y\}}}{\sum (1 - z_i) s_i}, \\ H(y) &= \mathbb{P}_N(Y \le y) = \frac{\sum (1 - z_i) s_i I_{\{Y_i \le y\}}}{\sum (1 - z_i) s_i}. \end{aligned}$$

(Details are given by Stute (1995). Note that if $(1 - Z_i)S_i = 0$, then V_{ji} can be assigned any arbitrary value since it has no contribution to the estimating equations.) To be clear and to simplify further notation:

$$\int_0^\infty w(s; \alpha, \beta) dF_p(s) = \sum_{j=1}^{k+1} w_j(\alpha) \left(F_p(t_j) - F_p(t_{j-1})\right),$$

where $w_j(\alpha) = w(t_j; \alpha, \beta), w_{k+1}(\alpha) = w(\tau; \alpha, \beta), F_p(t_0) = 0$, and $F_p(t_{k+1}) = 1$.

From our work in Section E.3, we know that

$$\sqrt{N}\left(\hat{\theta}-\theta\right) \rightarrow^{d} \mathcal{N}\left(0,\Psi\right).$$

Therefore, from Theorem 20.8 of van der Vaart (1998) (the functional delta method), the asymptotic variance matrix of $\hat{\theta}$, Ψ, can be written as

$$\Psi = E\left[\frac{\partial}{\partial\theta}\psi(\theta)\right]^{-1} E\left[\psi(\theta)\psi(\theta)^{T}\right] E\left[\left(\frac{\partial}{\partial\theta}\psi(\theta)\right)^{T}\right]^{-1}.$$

Taking derivatives, one can write

$$\frac{\partial}{\partial\theta}\psi_i(\theta) = \begin{bmatrix} -(1-Z_i) & -Z_i\sum w_j(\alpha)(F_p(t_j)-F_p(t_{j-1})) & 0 & \cdots & 0 \\ 0 & -Z_i p_0 \sum w_j'(\alpha)(F_p(t_j)-F_p(t_{j-1})) & 0 & & \\ 0 & -Z_i p_0(w_1(\alpha)-w_2(\alpha)) & -(1-Z_i)S_i & & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & -Z_i p_0(w_k(\alpha)-w_{k+1}(\alpha)) & 0 & \cdots\ 0 & -(1-Z_i)S_i \end{bmatrix}.$$

(The summations are from $j = 1, \cdots, k+1$.) One can then estimate Ψ with

$$\hat{\Psi} = \left[\frac{1}{N}\sum\frac{\partial}{\partial\theta}\psi_i(\hat{\theta})\right]^{-1}\frac{1}{N}\sum\left[\psi_i(\hat{\theta})\psi_i(\hat{\theta})^{T}\right]\left[\frac{1}{N}\sum\left(\frac{\partial}{\partial\theta}\psi_i(\hat{\theta})\right)^{T}\right]^{-1}.$$

(The summations here are from $i = 1, \cdots, N$.)

Of course, we are interested in the asymptotic variance of $g(\hat{\theta})$ where

$$\begin{aligned} g(\theta) &= \frac{\sum_{j=1}^{l} w_j(\alpha)(F_p(t_j)-F_p(t_{j-1}))}{\sum_{j=1}^{k+1} w_j(\alpha)(F_p(t_j)-F_p(t_{j-1}))} \\ &= F_p^{ai}(t), \end{aligned}$$

and t_l =sup(t_j) such that $t_j < t$. By the delta method,

$$\sqrt{N}\left(g(\hat{\theta})-g(\theta)\right) \rightarrow^{d} \mathcal{N}\left(0, g'(\theta)\Psi g'(\theta)^{T}\right),$$

where

$$g'(\theta) = \left(\frac{\partial g(\theta)}{\partial p_0}, \frac{\partial g(\theta)}{\partial\alpha}, \frac{\partial g(\theta)}{\partial F_p(t_1)}, \cdots, \frac{\partial g(\theta)}{\partial F_p(t_k)}\right).$$

These partial derivatives are

$$\frac{\partial g(\theta)}{\partial p_0} = 0$$
$$\frac{\partial g(\theta)}{\partial \alpha} = \frac{(1-VE)D - BC}{(1-VE)^2}$$
$$\frac{\partial g(\theta)}{\partial F_p(t_j)} = \frac{w_j(\alpha)I_{\{j\leq l\}} - w_{j+1}(\alpha)I_{\{j+1\leq l\}}}{1-VE} - \frac{C(w_j(\alpha) - w_{j+1}(\alpha))}{(1-VE)^2},$$

where

$$B = \sum_{j=1}^{k} w_j'(\alpha)(F_p(t_j) - F_p(t_{j-1}))$$
$$C = \sum_{j=1}^{l} w_j(\alpha)(F_p(t_j) - F_p(t_{j-1}))$$
$$D = \sum_{j=1}^{l} w_j'(\alpha)(F_p(t_j) - F_p(t_{j-1})), \text{ and of course}$$
$$1 - VE = \sum_{j=1}^{k} w_j(\alpha)(F_p(t_j) - F_p(t_{j-1})).$$

One estimates $g'(\theta)$ and Ψ with $g'(\hat{\theta})$ and $\hat{\Psi}$.

Wald confidence intervals based on this variance estimate may be constructed in the usual manner. Confidence intervals may also be constructed using a log-log transformation. If $\hat{\sigma}(t)$ is defined as the estimated standard error of $\hat{F}_p^{ai}(t)$ divided by $\hat{F}_p^{ai}(t)$, then a 95% confidence interval for $F_p^{ai}(t)$ using the log-log transformation is

$$\hat{F}_p^{ai}(t)^{exp\pm 1.96\hat{\sigma}(t)/log(1-\hat{F}_p^{ai}(t))}.$$

The log-log transformation restricts the resulting confidence interval within $(0, 1)$ and tends to improve small-sample performance.

Appendix F

ELICITING SENSITIVITY PARAMETERS

F.1 VaxGen Example

In order for these sensitivity analysis methods to be used, it is important to elicit meaningful ranges for the sensitivity parameter. Before doing the VaxGen time-to-event analyses, I emailed Dr. Marc Gurwith, the Chief Medical Officer at VaxGen, asking him for a biologically plausible range of values for β. I felt it might be helpful to include my email to him. It is given here:

"We are trying to do a sensitivity analysis, looking at the causal effect of vaccination on the time from infection diagnosis to ART. (This is essentially an extension of Figure 2b in the recent JID paper [Gilbert et al. (in press, 2005)].) We need some help choosing the range over which to perform our sensitivity analyses. Let us explain:

'From our primary analysis, we obtained an estimate for vaccine efficacy to prevent infection using the entire cohort. This estimate was $\widehat{VE} = 0.057$ and can be considered as an estimate of the probability that an infected participant in the placebo arm would have been protected if randomized to the vaccine arm. However, this estimate of vaccine efficacy is probably not uniform across all participants. Perhaps the vaccine is more or less effective for people with certain characteristics. (For example, the vaccine may have been more effective for non-whites and those with higher baseline risk scores.) However, there may be some factors in addition to these baseline covariates that help predict whether or not the vaccine may be more or less effective for a particular person. Consider two people in the placebo arm who became infected during the course of the trial. Suppose one of them (John) began ART two years after infection diagnosis, whereas the other (Bill) began ART one year after infection diagnosis. Which one of these individuals do you believe would have been more likely to be infected if randomized to vaccine?

'We would actually like you to put a numerical range on this belief, translating your belief into a range of possible odds ratios. Recall that the odds of infection is defined as

$$\text{odds of infection} = \frac{\text{probability of infection}}{1 - \text{(probability of infection)}}.$$

The odds ratio is simply the odds of infection for individuals with certain characteristics divided by the odds of infection for individuals with other characteristics. We want to elicit a range for the odds ratio using the time from infection diagnosis to ART. For individuals in the placebo arm who become infected, how does the time that they initiate ART affect the odds of them being infected if they had been assigned the vaccine?

'Returning to our example of John and Bill. Suppose the odds ratio is 2: This would imply that the odds that John would have been infected if randomized to vaccine are *twice* the odds that Bill would have been infected if assigned vaccine. On the other hand, an odds ratio of 1/2 would imply that the odds of John being infected if randomized to vaccine are *half* the odds of Bill being infected. An odds ratio of one implies that the odds of infection are equal for John and Bill.

'A couple thoughts that we have had that may be helpful: An individual with a shorter time from infection diagnosis to ART would be more likely to be infected if assigned vaccine if individuals with relatively strong immune systems tend to be slower to initiate ART and if the vaccine is more likely to protect these individuals from infection. If this is thought to be the case, then the odds ratio would be less than 1. On the other hand, the individual with a longer time from infection diagnosis to ART would be more likely to be infected if assigned vaccine if it is believed that the vaccine prevents infection with relatively strong/virulent viruses better than it prevents infection with weaker/a-virulent viruses. This would be a situation where the odds ratio would be greater than 1.

'Of course, we realize that this is simply an educated guess. Please just give us a range for the odds ratio that you feel is scientifically plausible."

I believe Dr. Gurwith's response reflects his understanding of both the problem and of the sensitivity parameter. With his permission, I have included his response here:

"My 'educated guesses' are based on the following reasoning. Let me know if my reasoning was incorrect. If the vaccine was effective in preventing infection only in those with a

strong immune system, then the time to ART would likely be longer in a vaccinee compared to placebo, since the strong immune system would delay progression of disease. Therefore, an individual in the placebo group with a longer time to ART would have a lower odds ratio of becoming infected if he had received the vaccine, since the longer time to ART would be a sign of a strong immune system. If the vaccine was effective only in preventing less virulent viruses, then the breakthrough infections in the placebo group would have a greater proportion of less virulent viruses and the time to ART would be longer in the placebo group; and therefore the placebo recipient with the longer time to ART would also have a lower odds ratio of being infected if he had received vaccine because the longer time to ART would be a sign of infection with a less virulent virus. However, if there were both effects, the vaccine was effective only in those with a strong immune system, and even in those with a strong immune system, it only protected against less virulent viruses, then there would be even a lower odds ratio of being infected if the placebo recipient with a longer time to ART received vaccine. Finally, the more likely case is the straightforward one: the vaccine is equally effective against all viruses, regardless of their virulence, and within limits of normal immune function, there are no meaningful differences in strength of immune system among vaccinees. In that case, the major determining factors for whether an individual is infected after exposure are factors extraneous to the vaccine and without effect on subsequent course of infection, such as type of exposure and/or number of virus particles to which the person is exposed. In this last scenario, I would say that placebo recipients who had longer time to ART and those with shorter time to ART would have similar risk of infection if they had been in the vaccine group.

'Therefore, in our overall trial where vaccine efficacy was only slightly greater (and nowhere near statistically significant) than 0, I think the range of odds ratios for becoming infected for the placebo recipient with longer time to ART, if he had received vaccine, would be 0.70 to 1.1.

'For the non-white cohort, with an apparent vaccine efficacy $> 50\%$, the bias would be the same, the range of odds ratios would be greater, such as 0.40 to 1.20.

'Please let me know if I am providing you the type of educated guesses you wanted, and the emphasis still should be on 'guess'."

Admittedly, Dr. Gurwith is somewhat familiar with this work, and probably not all subject-matter experts would agree with his selected sensitivity analysis range – Dr. Gurwith himself is the first to emphasize that it is a guess. However, this exercise has convinced me that it is possible to elicit meaningful ranges for sensitivity analyses.

F.2 Interpretation

In most of my examples, $w(\cdot)$ is modeled with an expit function with sensitivity parameter β. Under this model β has a standard log-odds ratio interpretation, i.e. given infection if assigned placebo (and possibly covariates) for a one-unit increase in Y the odds of being infected if assigned vaccine multiplicatively increase $exp(\beta)$. Most medical researchers are at least generally familiar with odds ratios. For example, they could say that an odds ratio of 1.1 is small, 3 is big, and 10 is very big.

For some people, however, it may be difficult to choose a specific sensible range for β. To many, odds ratios are not as intuitive as probabilities and relative risks. These intuitive difficulties may be compounded in our situation because the odds ratio is defined at the counterfactual level and is not directly identified.

One alternative approach to eliciting a sensible range for β could be to map choices of the sensitivity parameter into probabilities (or relative risks). Consider the situation without covariates and with the outcome variable being viral load. Using previous viral load data (such as data from the MACS cohort (Lyles et al., 2000)) and a particular VE, one can construct a table containing the probabilies of being infected if assigned vaccine, given assigned placebo for different viral load values and different choices of β. Table F.1 shows these probabilities and relative risks (with respect to the mean viral load) of being in the ai stratum for $VE = 0.1, 0.3$, and 0.6 using placebo viral load quantiles based on the MACS data.

The probabilities and relative risks shown in Table F.1 were easily calculated: Assuming a value of VE and at a fixed choice for β, $\hat{\alpha}$ was first computed using the MACS viral load data. (From Lyles et al. (2000), the first quartile, median, and third quartile values for viral loads at the first seropositive visit were 3.98, 4.49, and 5.01 respectively. Since I did not have

the raw data, I had to make a distributional assumption in order to estimate $\hat{\alpha}$. Assuming normality of the viral loads in the MACS dataset – which based on the median and quartiles appears quite reasonable – the mean equals approximately 4.5, and the standard deviation can be calculated as approximately 0.75, defining the viral load distribution.) Once there is an estimate for α, one can calculate the probability of infection if randomized to the vaccine which is simply the estimate $w(y;\hat{\alpha},\beta) = exp(\hat{\alpha}+\beta y)/(1+exp(\hat{\alpha}+\beta y))$. From this quantity one can easily calculate relative risks.

Hence, another approach to eliciting values of β would be to elicit relative risks comparing some y with y_0 and converting them to β. This could be done in a manner similar to that based on odds ratios, only now asking questions such as the following: Do you believe an infected participant in the placebo arm with a viral load one-log below average would have been twice as likely to be infected? Three times as likely? Half as likely? It should be recognized that there is not a perfect concordance between odds ratios and relative risks. For example, a relative risk of rr may result in an odds ratio of $exp(\beta_1)$ between y and $y+1$, but an odds ratio of $exp(\beta_2)$ between $y-1$ and y, where $\beta_1 \neq \beta_2$.

Table F.2 shows relative risks converted to β for certain values of VE. This table was also created using a viral load distribution based on the MACS dataset (assuming $Y(0)|S(0)=1 \sim \mathcal{N}(4.5, 0.75^2)$). Relative risks in this table are defined as the probabilites being in the ai stratum given the mean placebo viral load divided by the probability of being in the ai stratum given the mean placebo viral load minus one. Notice that some levels of the relative risk might not correspond to any odds ratio for a given VE and viral load distribution. For example, if vaccine efficacy is 0.1, it is not possible to obtain a value of β corresponding to a relative risk of 0.5 between viral load values of 4.5 and 3.5.

An alternative way to elicit sensitivity parameters could be to convert choices of β into some measure of discrepancy between $f_p^{ai}(\cdot)$ and $f_p^{prot}(\cdot)$ – for example, the difference in means. Values using this approach also based on the MACS viral load data are shown in Table F.3.

Finally, another approach could be to "invert" the sensitivity parameter. In other words, the data analyst performs the sensitivity analysis and sees at what values of the sensitivity parameter the study's conclusions change. Then one could determine if these critical sensi-

tivity parameter values are plausible. For example, consider a study where $H_0 : ACE = 0$ is rejected for all $\beta < 2$. Subject-matter experts may look at the results and decide that an odds ratio greater than $exp(2)$ is very unlikely. Hence, scientists may be confident concluding $ACE \neq 0$, that the vaccine has a causal effect on the post-infection outcome. With this approach, however, it should be recognized that researchers may have the tendency to change what they think is plausible based on the results of the analysis.

In short, there are many alternative ways to interpret or elicit a range for the sensitivity parameter based on the expit function. One could of course use other models for $w(\cdot)$: the probit function immediately comes to mind; Jemiai has proposed and studied a different weight function; and there are many other possibilities.

Table F.1: Probabilities (relative risks with respect to the median) of infection if assigned vaccine for given β, VE, and viral load.

β	$Y(0)$						
	3.27 (5%)	3.54 (10%)	3.99 (25%)	4.5 (50%)	5.01 (75%)	5.46 (90%)	5.73 (95%)
				$VE = 0.1$			
0.1	0.89 (0.99)	0.89 (0.99)	0.90 (0.99)	0.90 (1)	0.90 (1.00)	0.91 (1.01)	0.91 (1.01)
0.5	0.84 (0.92)	0.85 (0.94)	0.88 (0.97)	0.90 (1)	0.92 (1.02)	0.94 (1.04)	0.95 (1.05)
1	0.76 (0.83)	0.81 (0.88)	0.87 (0.95)	0.92 (1)	0.95 (1.03)	0.97 (1.05)	0.97 (1.06)
2	0.62 (0.65)	0.74 (0.78)	0.87 (0.92)	0.95 (1)	0.98 (1.03)	0.99 (1.04)	1.00 (1.05)
3	0.49 (0.50)	0.69 (0.70)	0.90 (0.92)	0.98 (1)	0.99 (1.02)	1.00 (1.02)	1.00 (1.02)
				$VE = 0.3$			
0.1	0.67 (0.96)	0.68 (0.97)	0.69 (0.98)	0.70 (1)	0.71 (1.02)	0.72 (1.03)	0.73 (1.04)
0.5	0.56 (0.80)	0.60 (0.85)	0.65 (0.92)	0.71 (1)	0.76 (1.07)	0.79 (1.13)	0.82 (1.16)
1	0.43 (0.60)	0.50 (0.69)	0.61 (0.84)	0.72 (1)	0.81 (1.12)	0.87 (1.21)	0.90 (1.25)
2	0.22 (0.28)	0.32 (0.42)	0.54 (0.71)	0.77 (1)	0.90 (1.17)	0.96 (1.25)	0.97 (1.27)
3	0.10 (0.12)	0.20 (0.24)	0.49 (0.60)	0.82 (1)	0.95 (1.17)	0.99 (1.21)	0.99 (1.22)
				$VE = 0.6$			
0.1	0.37 (0.93)	0.38 (0.94)	0.39 (0.97)	0.40 (1)	0.41 (1.03)	0.42 (1.06)	0.43 (1.07)
0.5	0.26 (0.66)	0.29 (0.73)	0.34 (0.85)	0.40 (1)	0.46 (1.16)	0.52 (1.30)	0.55 (1.38)
1	0.16 (0.40)	0.20 (0.50)	0.28 (0.71)	0.39 (1)	0.51 (1.32)	0.62 (1.61)	0.69 (1.77)
2	0.05 (0.13)	0.08 (0.21)	0.17 (0.47)	0.36 (1)	0.61 (1.69)	0.79 (2.20)	0.87 (2.41)
3	0.01 (0.04)	0.03 (0.08)	0.10 (0.29)	0.33 (1)	0.69 (2.10)	0.90 (2.74)	0.95 (2.90)

Table F.2: Values of β for specific relative risks and VE.

VE	RR							
	0.33	0.5	0.67	0.91	1.1	1.5	2.0	3.0
0.1	–	–	–	–	0.75	3.2	8.5	> 8.5
0.3	–	–	–	-0.37	0.29	1.0	1.6	2.3
0.6	-3.5	-1.6	-0.79	-0.16	0.15	0.60	0.97	1.4

Table F.3: Difference in mean viral loads in the placebo arm between the ai and protected principal strata for a given VE and β, based on the MACS data.

VE	β				
	0.1	0.5	1.0	2.0	3.0
0.1	0.06	0.28	0.53	0.91	1.13
0.3	0.06	0.27	0.51	0.82	0.98
0.6	0.06	0.27	0.50	0.80	0.96

VITA

Bryan Shepherd was born August 12, 1975 in Salt Lake City, Utah to Earl and Anna Clare Shepherd. He obtained a B.S. in Statistics from Brigham Young University in 1999 and an M.S. in Biostatistics from the University of Washington in 2001.