Sensitivity analysis for inferences on causal treatment effects on endpoints measured only in a subset selected post-randomization, with application to HIV vaccine trials

Biostat 578A Lecture 5

Contributors: Bryan Shepherd
Yannis Jemiai
Andrea Rotnitzky
Devan Mehrotra
Xiaoming Li

Outline

I. Introduction and Motivation: HIV vaccine efficacy trials

II. Likelihood Methods for Estimating Causal Effects

- Nonparametric, semiparametric, parametric
- Binary, quantitative, time-to-event endpoints
- Baseline covariates

III. Example

Example: HIV Vaccine Trials

Primary Question:

Does the vaccine prevent HIV infection?

Vaccine Efficacy (VE_S)

$$\mathrm{VE}_S = 1 - \frac{P(infected|vaccine)}{P(infected|placebo)}$$

Co-primary/Secondary Question:

Does vaccine ameliorate HIV progression post-HIV-infection?

Y=$\log_{10}$ viral load

Surrogate endpoint for time-to-AIDS and secondary transmission

Viral Load Surrogate Endpoint:

- Feasible approach to assessing the vaccine's effect on disease progression and infectiousness: use viral load as a surrogate marker
  - Viral load is prognostic for rate of disease progression and level of infectiousness
    - E.g., Mellors et al. (1997, *Annals of Internal Medicine* **126**:946-954); Quinn et al. (2000, *New England Journal of Medicine* **342**:921-929)
  - A correlate does not a surrogate make
  - Difficult to validate viral load as an accurate surrogate endpoint for the vaccine's effect on the true clinical endpoints

Analysis of Pre-Treatment Viral Load:

- **Collected data:** viral load measurements from vaccine trial participants who become HIV infected
  - sampled from plasma, genital secretions, lymph nodes, and/or elsewhere
  - Example sampling schedule: months
    0, 1, 2, 3, 6, 12, 18, 24, 30, 36, 42, 48, 54, 60
- Let Y be some univariate summary measure of the pre-treatment viral loads from a subject
  - sample average, slope estimate,
    estimate of a viral dynamics parameter

Other Post-Infection Outcomes:

Time from infection detection to AIDS or death

Secondary transmission

Other biomarkers

Genetic distance

CD4 count

Time from infection Dx until viral load exceeds specified level

Post-Randomization Outcomes in Other Types of Trials:

Disease severity in preventative trials (e.g., cancer)

Quality of Life

Time to HIV transmission through breastmilk in infants alive and HIV-free at 1 month of age

Simple Analysis:

Estimand $E(Y|infected, vaccine) - E(Y|infected, placebo)$

Question Answered:

Do people randomized to vaccine who became infected have a different viral load than people randomized to placebo who became infected?

Not a Causal Estimand:

Conditions on a post-randomization variable, infection

Measures a mixture of the causal vaccine effect and any differences in characteristics among the infected subgroups

Partial vaccine efficacy ($\text{VE}_S > 0$) is a selection mechanism that may lead to different characteristics

Example of selection bias from $\mathbf{VE}_S$:

	Immune System		
	Strong	Weak	Average
Vaccine	–	Y=5	$\bar{Y}$=5
Placebo	Y=3	Y=5	$\bar{Y}$=4

- Vaccine effect defined as average difference in viral load in infected subgroups $= 4 - 5 = -1$
- This result suggests the vaccine is increasing viral load
- Incorrect inference, which could adversely affect vaccine development

Background: What is a Causal Effect?

- Potential outcomes framework (Rubin, 1974)
  - $Y_i(1)$ = Outcome if subject i is assigned treatment 1 (vaccine)
  - $Y_i(0)$ = Outcome if subject i is assigned treatment 0 (placebo)
    - $Y_i(1)$ is observed for subjects in arm 1, and is *counterfactual* for subjects in arm 0
    - Similarly $Y_i(0)$ is observed for subjects in arm 0, and is *counterfactual* for subjects in arm 1
  - Early use of *counterfactuals* in philosophy:

    Blaise Pascal, *The Pensees*, 1656-1661:

    "Had Cleopatra's nose been shorter, the whole face of the world would have been different"

Background: What is a Causal Effect?

- $Y_i(1) - Y_i(0)$ is a *causal effect*, because it compares the outcome under the two conditions *for the same individual*
- $Y_i(1) - Y_j(0)$ for $i \neq j$ is *not* a causal effect
- Definition of a causal effect: the comparison of potential outcomes on a common set of units
  - Causal effects are comparisons of apples and apples, not apples and oranges

Causal vs Non-Causal Effects (see Rubin 2004, Scand J Stat)

- Z = Treatment assignment;, Y = Time to treatment failure

		Potential Outcomes		Individual causal
	Z_i	$Y_i(0)$	$Y_i(1)$	effects $Y_i(1) - Y_i(0)$
	1	19	20	+1
	1	10	13	+3
	1	9	10	+1
	0	8	1	-7
	0	6	0	-6
	0	9	2	-7
	0	5	1	-4
True Averages		9.4	6.7	-2.7
Observed Averages		7	14.3	

- On average treatment 0 is better (average causal effect = -2.7)
- All patients receive the best treatment for themselves
- The observed data suggest treatment 1 is better, the **wrong conclusion**

Causal Inference needs an Assignment Mechanism:

- Assignment mechanism is a model for

$$Pr(Z|Y(0), Y(1))$$

- Randomized experiment:

$$Pr(Z|Y(0), Y(1)) \quad \text{independent of} \quad Y(0), Y(1)$$

  - E.g., completely randomized experiment with 3 treated and 4 control

$$\begin{aligned} Pr(Z|Y(0), Y(1)) &= 1/35 \quad \text{if} \quad \sum_{i=1}^{7} Z_i = 3 \\ &= 0 \quad \text{otherwise} \end{aligned}$$

    * Here Z is a 7-vector of assignments, and 35 is the number of ways to choose 3 objects from 7 w/o replacement

Causal Inference needs an Assignment Mechanism:

- The reason randomized clinical trials are the gold standard for evaluating interventions is that causal effects can be estimated under tenable assumptions
- Consider a 2 arm trial with n_1 and n_0 randomized subjects, $Z = 1, 0$
- The standard estimand, $E(Y_i|Z_i = 1) - E(Y_i|Z_i = 0)$, can be estimated by the difference in sample means, $\bar{Y}_1 - \bar{Y}_0$
- However, the scientific goal is to estimate a causal estimand, for example the *average causal effect (ACE)*:

$$ACE = E(Y_i(1) - Y_i(0))$$

- **Randomization assumption** (also called *ignorable treatment assignments*):

$$Z \amalg Y(0), Y(1) \qquad (i.e., Z \text{ is independent of } Y(0), Y(1))$$

- **Key result**: Under SUTVA (described later) and randomization, $E(Y_i|Z_i = 1) - E(Y_i|Z_i = 0) = E(Y_i(1) - Y_i(0))$, so that $\bar{Y}_1 - \bar{Y}_0$ is an unbiased estimate of the average causal effect

Causal Inference needs an Assignment Mechanism:

- Demonstration of the key result:

$$\begin{aligned} E(Y_i|Z_i=1) - E(Y_i|Z_i=0) &= E(Y_i(1)|Z_i=1) - E(Y_i(0)|Z_i=0) && \text{by SUTVA} \\ &= E(Y_i(1)) - E(Y_i(0)) && \text{by randomiz.} \\ &= E(Y_i(1) - Y_i(0)) \\ &= \text{Average Causal Effect} \end{aligned}$$

- The SUTVA assumption is explained later in these slides

Causal Inference needs an Assignment Mechanism:

- Complementary explanation of the key result:

$$\begin{aligned}
\bar{Y}_1 - \bar{Y}_0 &= \frac{1}{n_1}\sum_{i=1}^{n} Z_i Y_i - \frac{1}{n_0}\sum_{i=1}^{n}(1-Z_i)Y_i \\
&= \frac{1}{n_1}\sum_{i=1}^{n} Z_i Y_i(1) - \frac{1}{n_0}\sum_{i=1}^{n}(1-Z_i)Y_i(0) \quad \text{by SUTVA}
\end{aligned}$$

Causal Inference needs an Assignment Mechanism, Continued:

- By randomization

  $\frac{1}{n_1}\sum_{i=1}^{n} Z_i Y_i(1)$ estimates the same thing as $\frac{1}{n}\sum_{i=1}^{n} Y_i(1)$

  and

  $\frac{1}{n_2}\sum_{i=1}^{n}(1-Z_i) Y_i(0)$ estimates the same thing as $\frac{1}{n}\sum_{i=1}^{n} Y_i(0)$

  Therefore $\bar{Y}_1 - \bar{Y}_0$ estimates the same thing as $\bar{Y}(1) - \bar{Y}(0)$

Intent-to-Treat Approach to Estimating the Average Causal Effect on Viral Load:

"Burden-of-illness" Approach, Chang, Guess, and Heyse, 1994

Compare all randomized vaccine recipients to all randomized placebo recipients

Uninfected participants assigned a viral load value of $Y = 0$

Estimate the average causal effect $E(Y_i(1) - Y_i(0))$

Limitations:

Low power (documented in Lecture 4)

Aggregates two kinds of vaccine effects

Data: N iid Observations $O_i = (Z_i, X_i, S_i, Y_i)$:

$$Z_i = \begin{cases} 0 \text{ if randomized to placebo} \\ 1 \text{ if randomized to vaccine} \end{cases}$$

$X_i =$ vector of baseline covariates

$$S_i = \begin{cases} 0 \text{ if not infected} \\ 1 \text{ if infected} \end{cases}$$

$Y_i = \log_{10}$ viral load if infected

Potential Outcomes/Counterfactuals:

$S_i(z) =$ infection status indicator if subject i is assigned treatment z

$Y_i(z) =$ viral load if subject i is assigned treatment z

If $S_i(z) = 0$, we define $Y_i(z) = *$

Principal Stratification (Frangakis and Rubin, 2002)

The basic principal stratification with respect to infection is a partitioning of units according to the same vector $(S_i(0), S_i(1))$. A comparison made within a principal stratum is a causal effect.

(0,0) never infected
(0,1) infected with vaccine, not with placebo (harmed)
(1,0) infected with placebo, not with vaccine (protected)
(1,1) always infected (ai)

Viral load comparisons in the strata (0,0), (0,1), and (1,0) are undefined

Average Causal Effect

$$ACE = E[Y(1) - Y(0)|(S(0), S(1)) = (1, 1)] = \int y dF_v^{ai}(y) - \int y dF_p^{ai}(y)$$

Causal Question

Does the vaccine alter viral loads for those who would have been infected regardless of vaccine assignment?

Hudgens, Hoering, Self (2003, Statistics in Medicine)

Gilbert, Bosch, Hudgens (2003, Biometrics)

- Nonparametric and semiparametric estimation of ACE without covariates
- Key Assumption (Monotonicity):

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

For each trial participant, assignment to vaccine does not increase the risk of infection compared to assignment to placebo

$\Rightarrow$ There are no participants in the harmed stratum

Possible Principal Strata of Participants

	Vaccine	Placebo
Infected $S = 1$	always infected ~~harmed~~ (n_v)	always infected protected (n_p)
Not Infected $S = 0$	never infected protected $(N_v - n_v)$	never infected ~~harmed~~ $(N_p - n_p)$
	N_v	N_p

$F_v^{ai}(y)$ is identified: $F_v^{ai}(y) = F_v(y)$

$F_p^{ai}(y)$ is not identified

$F_p^{ai}(y)$ is not identified

$$f_p(y) = (1 - VE_S) f_p^{ai}(y) + VE_S f_p^{prot}(y)$$

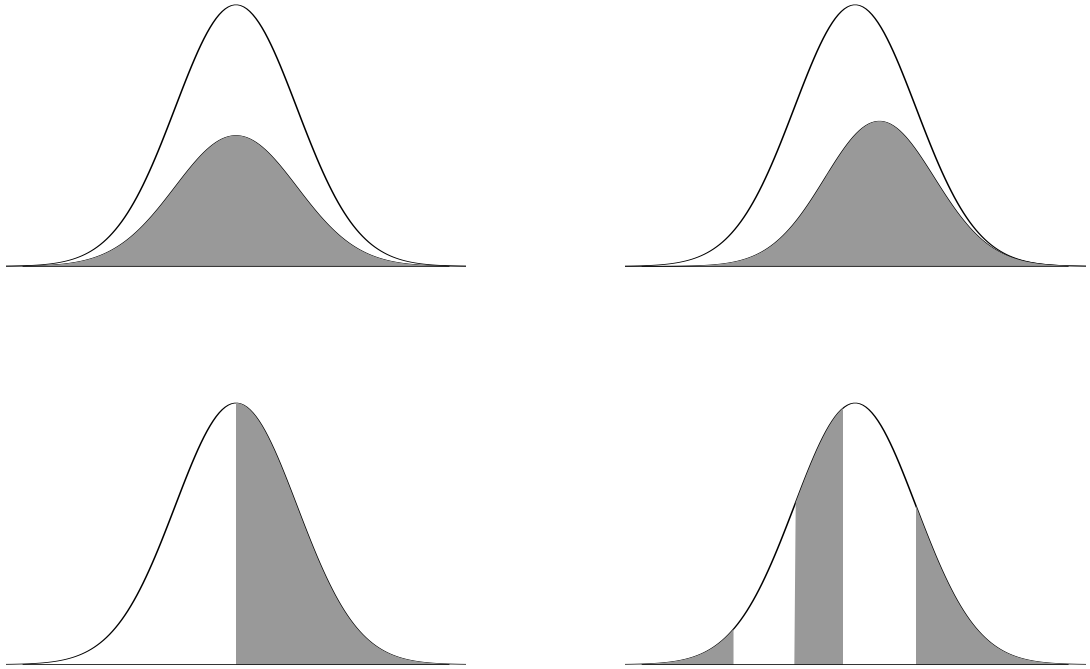

Shaded [unshaded] area represents those in *ai* [*prot*] principal stratum

Unshaded area $= VE_S$

Hudgens, Hoering, Self (2003) (Nonparametric Bounds):
Considered the 2 extreme cases wherein the *protected* have the highest viral loads (left plot) and the *protected* have the lowest viral loads (right plot)

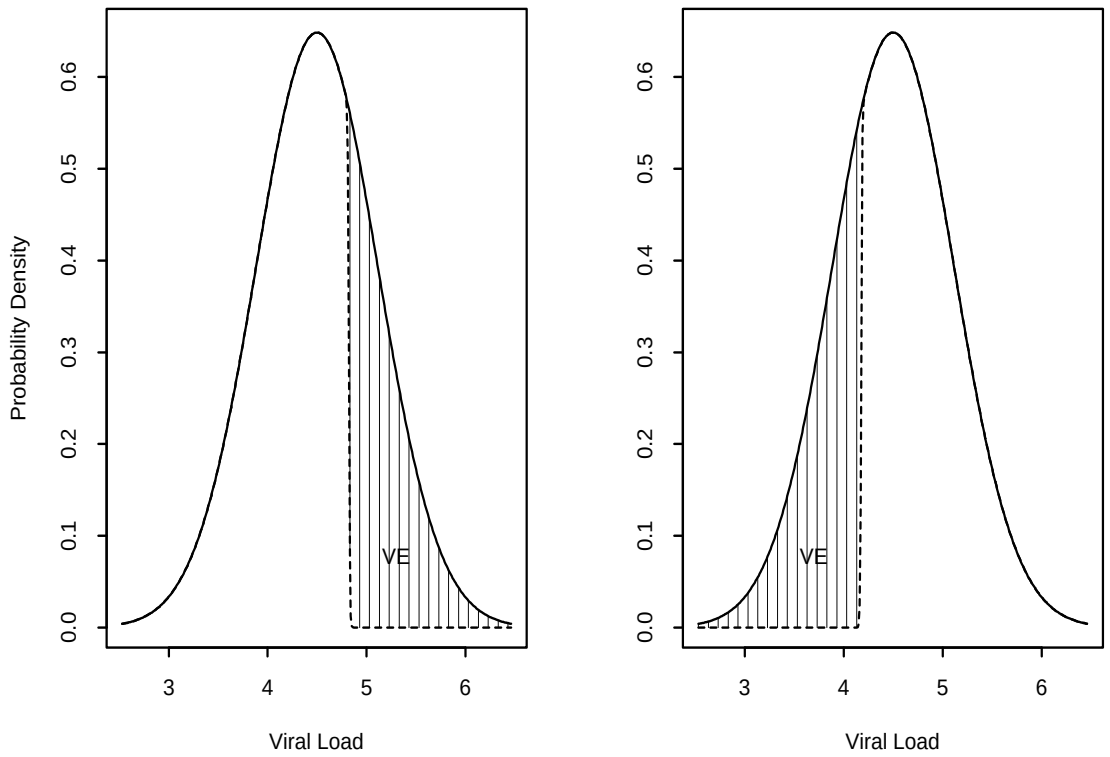

The density $f_p(y)$ (total area) partitions into the sub-densities $f_p^{ai}(y)$ (hatchmarked area $= VE$) and $f_p^{prot}(y)$ (unshaded area $= 1 - VE$)

Gilbert, Bosch, Hudgens (2003) (Semiparametric):

Selection bias model:

Mixture of $f_p(\cdot)$ can be re-written as

$$f_p^{ai}(y) = \frac{w(y)}{1 - VE_S} f_p(y)$$

$$\text{where } \int w(y) f_p(y) dy = 1 - VE_S$$

and $w(y) \equiv P(S(1) = 1 | S(0) = 1, Y(0) = y)$

Perform sensitivity analysis by specifying $w(\cdot)$, e.g.

$$w(y) = \frac{exp(\alpha + \beta y)}{1 + exp(\alpha + \beta y)}$$

where β is an unidentifiable sensitivity parameter

About β:

$\beta = 0 \quad \Leftrightarrow \quad S(1) \amalg Y(0) | S(0) = 1$

$\beta = \pm\infty \quad \Rightarrow$ Reconstruct the nonparametric bounds $F_p^{ai,L}(\cdot)$ and $F_p^{ai,U}(\cdot)$ for $F_p^{ai}(\cdot)$ given by Hudgens, Hoering, Self (2003)

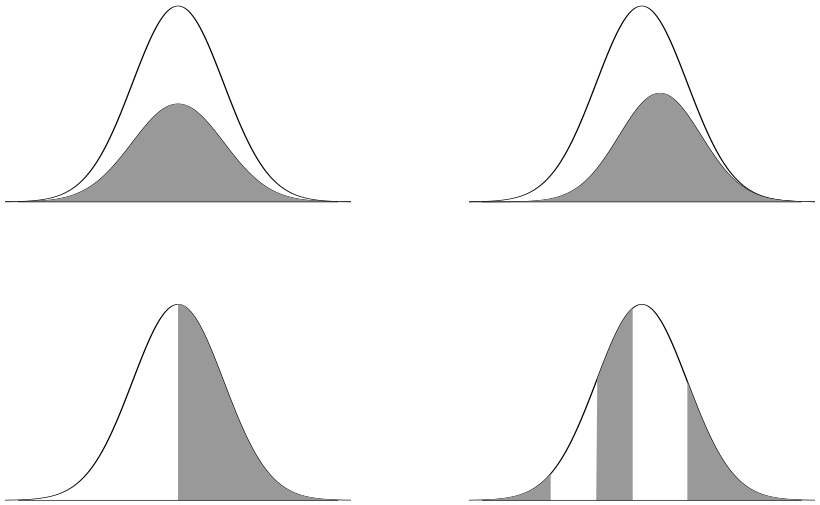

Odds Ratio Interpretation: e^{β} = odds ratio an infected placebo subject would be infected if assigned vaccine for a one-unit higher viral load $Y(0)$

Interpretation of Sensitivity Parameter:

- e^{β} = odds ratio of infection if assigned vaccine comparing two infected placebo recipients with viral loads Y and $Y - 1$
  - $\beta = 0$: reflects zero selection bias
  - $\beta > 0$ and finite: reflects selection bias towards higher viral loads in vaccinees
  - $\beta < 0$ and finite: reflects selection bias towards lower viral loads in vaccinees
  - the magnitude of presumed selection bias increases with $|\beta|$

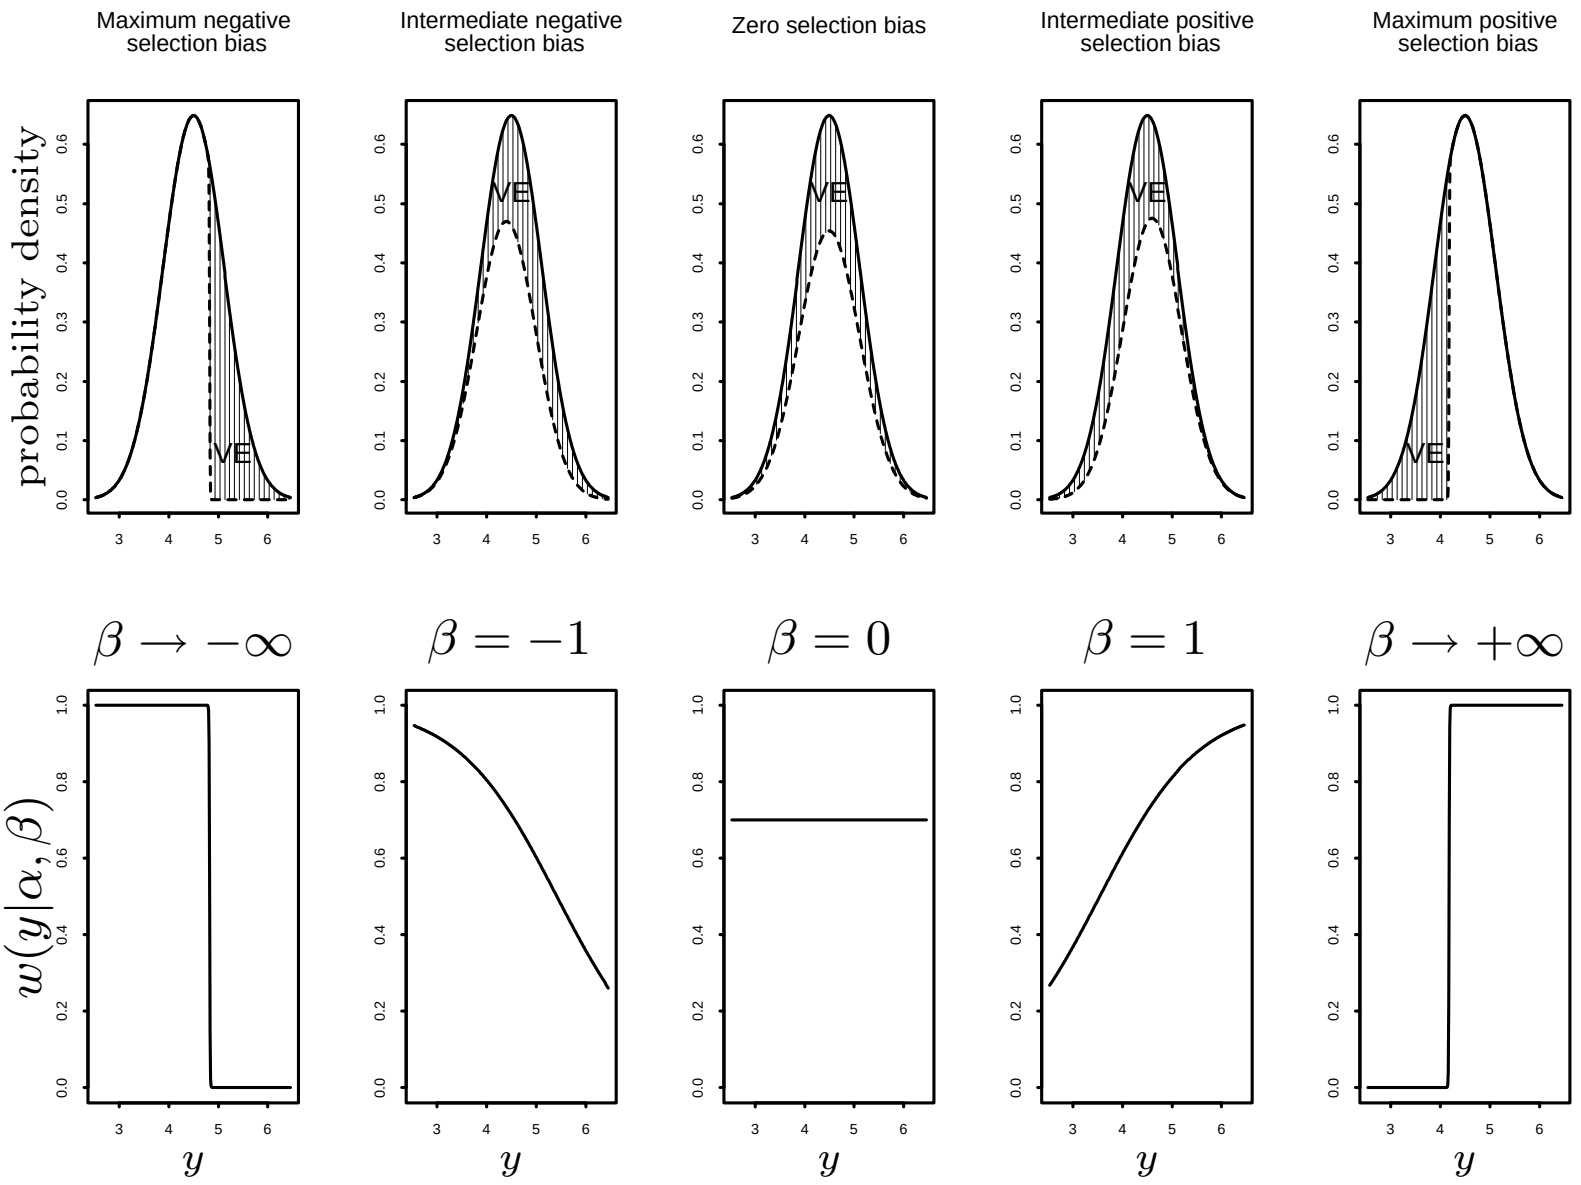
Maximum negative selection bias
Intermediate negative selection bias
Zero selection bias
Intermediate positive selection bias
Maximum positive selection bias
probability density
VE
$\beta \to -\infty$
$\beta = -1$
$\beta = 0$
$\beta = 1$
$\beta \to +\infty$
$w(y|\alpha, \beta)$
y

Assumptions to Identify ACE:

A.1: Stable Unit Treatment Value Assumption (SUTVA)

The potential outcomes for each trial participant are unrelated to the treatment assignment of other subjects (Rubin 1978, 1986)

A.2: Randomization (Ignorable Treatment Assignments):

$$Z \amalg (S(0), S(1), Y(0), Y(1))$$

A.3: Monotonicity: $S_i(1) \leq S_i(0)$

A.4: $w(Y;\beta) = \Phi\{\alpha + g(Y;\beta)\}$,

where β is fixed and known; $\Phi(u)$ is a known cdf; α is an unknown constant parameter to be estimated, and $g(Y;\beta)$ is a known function of Y and β

SUTVA in More Detail:

- $\mathbf{Z}$ = vector of vaccination assignments for the N randomized subjects, with i^{th} element Z_i
- $\mathbf{S}(\mathbf{Z})$ = N-vector with ith element $S_i(\mathbf{Z})$
  - $S_i(\mathbf{Z})$ = indicator of whether the i^{th} subject would be infected given $\mathbf{Z}$
- SUTVA states that potential outcomes $(S_i(1), S_i(0), Y_i(1), Y_i(0))$ for each subject i are unrelated to the assignment Z_j of other subjects, and allows $S_i(\mathbf{Z})$ and $Y_i(\mathbf{Z}, \mathbf{S})$ to be written as $S_i(Z_i)$ and $Y_i(Z_i)$, respectively
- Therefore, under SUTVA each subject has two potential infection outcomes ($S_i(1)$, $S_i(0)$) and at most two potential viral load outcomes $(Y_i(1), Y_i(0))$
- SUTVA implies *consistency*:

$$Y_i(z) = Y_i \text{ if } Z_i = z,\ z = 0, 1$$

Concept of Sensitivity Analysis of ACE:

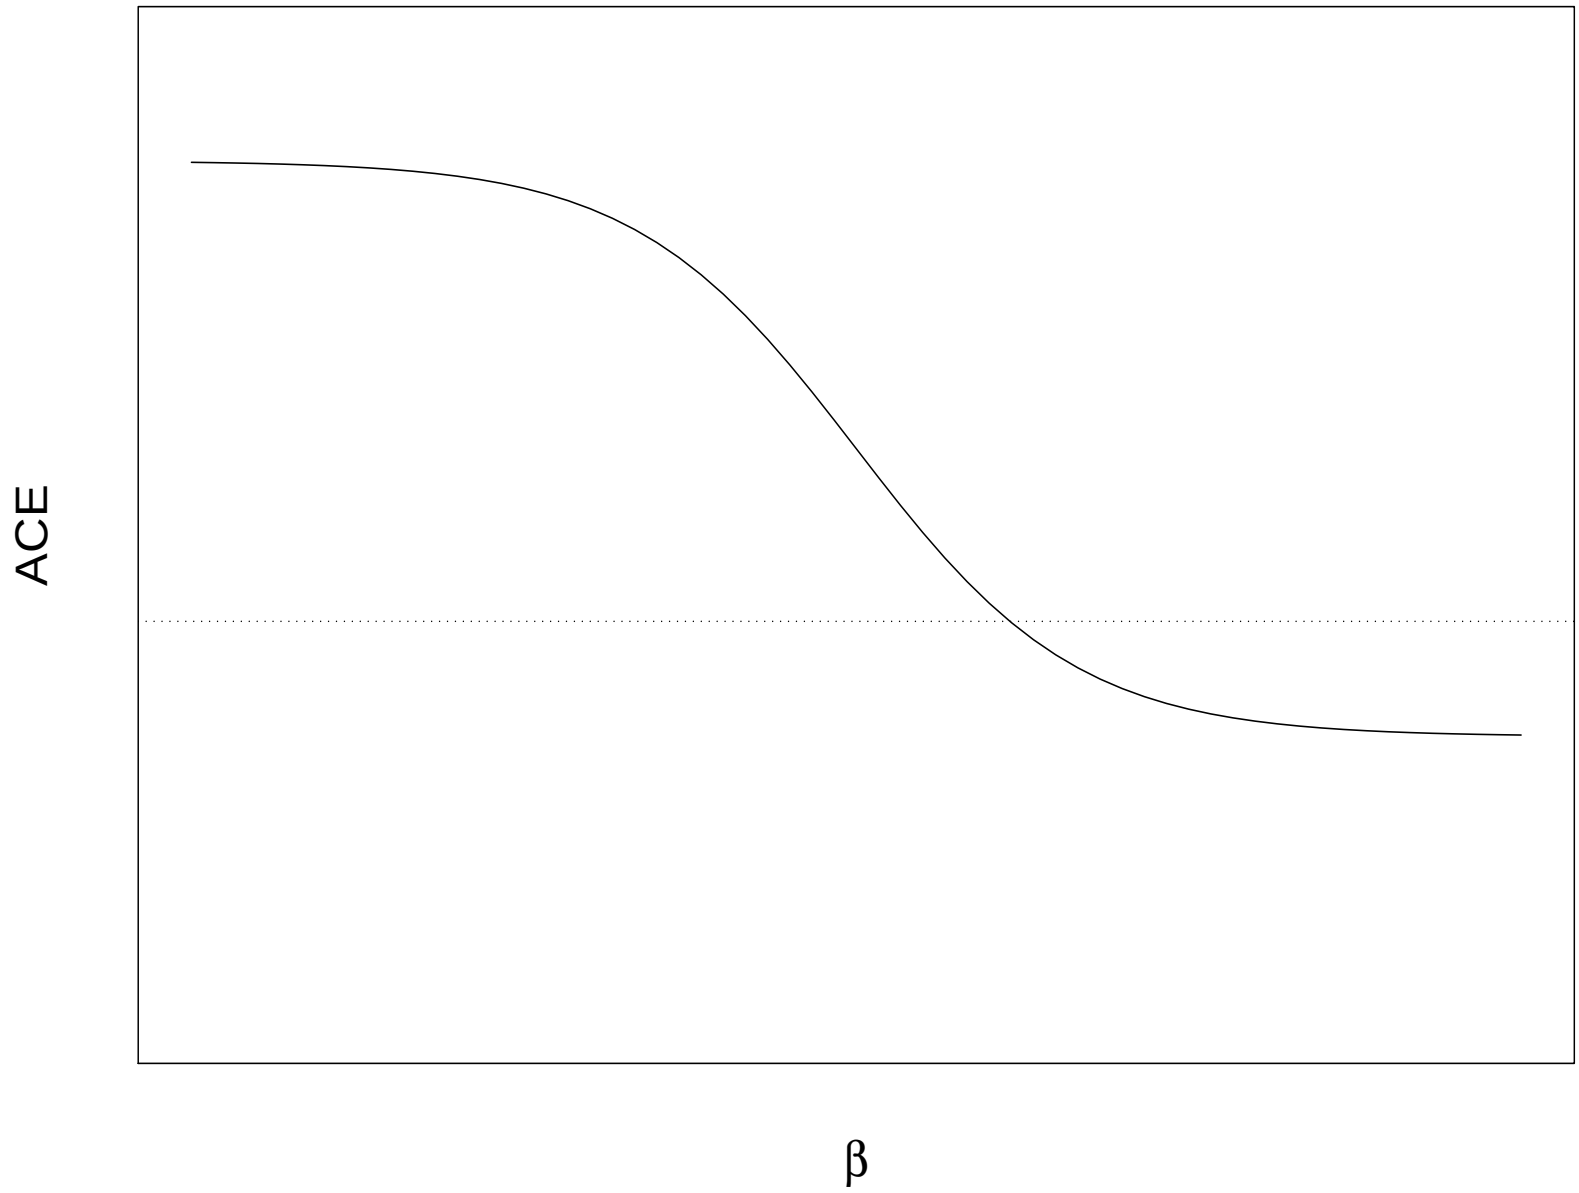

Gilbert, Bosch, Hudgens (GBH) (2003), Semiparametric Maximum Likelihood Estimation:

Solve for $\widehat{\alpha}$ using the equation

$$1 - \widehat{VE_S} = \int_{-\infty}^{\infty} \frac{exp(\widehat{\alpha} + \beta y)}{1 + exp(\widehat{\alpha} + \beta y)} d\widehat{F}_p(y)$$

where $1 - \widehat{VE_S} = min[1, (n_v/N_v)/(n_p/N_p)]$ and $\widehat{F}_p$ is the NPMLE of F_p

$\widehat{\alpha}$ is the value of α that makes the area between $\widehat{F}_p^{ai}(\cdot)$ and $\widehat{F}_p^{prot}$ equals $\widehat{VE_S}$

Then

$$\widehat{F}_p^{ai}(y) = (1 - \widehat{V}E_S)^{-1} \int_{-\infty}^{y} \frac{exp(\widehat{\alpha} + \beta z)}{1 + exp(\widehat{\alpha} + \beta z)} d\widehat{F}_p(z).$$

Causal inferences: $\widehat{ACE}(\beta) = \widehat{E}(Y(1)|ai) - \widehat{E}(Y(0)|ai; \beta)$

Bootstrap Procedure for Estimating $Var(\widehat{ACE}(\beta))$ and Obtaining a Confidence Interval for $ACE(\beta)$:

- Step 1. Sample values $n_v^\star$ and $n_p^\star$ from Poisson distributions with infection rate parameters $\frac{n_v}{N_v}$ and $\frac{n_p}{N_p}$
- Step 2. Calculate $\widehat{\text{VE}}_S^\star = max(0, 1 - \frac{n_v^\star}{N_v} / \frac{n_p^\star}{N_p})$ and calculate $\widehat{\alpha}^\star$ as the solution of

$$1 - \widehat{\text{VE}}_S^\star = \int_{-\infty}^{\infty} \frac{\exp\{\alpha + \beta y\}}{1 + \exp\{\alpha + \beta y\}} d\widehat{F}_p(y)$$

- Step 3. Draw a bootstrap sample of size $n_p^\star$ from $\widehat{F}_p(\cdot)$, and independently draw a bootstrap sample of size $n_v^\star$ from $\widehat{F}_v(\cdot)$, and compute the bootstrap estimate $\widehat{ACE}^\star(\beta)$
- Step 4. Estimate $Var(ACE(\beta))$ by the sample variance of

$$\widehat{ACE}^\star(\beta)_1, \cdots, \widehat{ACE}^\star(\beta)_B$$

Obtain a $(1-\gamma) \times 100\%$ percentile confidence interval for $ACE(\beta)$ as the $\gamma/2$ and $(1-\gamma/2)$ quantiles of

$$\widehat{ACE}^\star(\beta)_1, \cdots, \widehat{ACE}^\star(\beta)_B$$

Simulation Study:

- Data simulated from a repeated measures model:

$$Y_{pij} = \mu_{pi} + \epsilon_{pij}, \qquad i = 1, \cdots, n_p; j = 1, \cdots, K \tag{1}$$

  with $\mu_{pi} \sim N(\mu_p, \sigma^2_{\mu_p})$ and $\epsilon_{pij} \sim N(0, \sigma^2_\epsilon)$ independent

  - $Y_{p1}, \cdots, Y_{pn_p}$ simulated from (1) with
    $\mu_p = 4.50, \sigma^2_{\mu_p} = 0.58, \sigma^2_\epsilon = 0.068$
  - $Y_{v1}, \cdots, Y_{vn_v}$ simulated under an assumed (true) selection bias model

$$f^{ai}(y; F_p, VE_S, \beta) = W(\alpha, \beta, F_p)^{-1} \frac{\exp\{\alpha + \beta y\}}{1 + \exp\{\alpha + \beta y\}} f_p(y)$$

    with $\beta = 0, 1,$ or ∞

- Simulation parameters:
  - Expected number infected in placebo group (45)
  - Vaccine efficacy VE against infection (50%)
  - Number of measurements of viral load from each participant (2)
  - Size of location departure Δ (in $\log_{10}$s) from H_0 over and above any selection bias induced by the true β $(0, 1/3, 1/2)$

Table 1. Power $\times$ 100% for detecting a 0, 1/3, and 1/2 $\log_{10}$ mean-shift alternative, 1-sided 5% level test (from GBH)

True	Specified	Shift Alternative		
β	β	0	1/3	1/2
0	0	**6.6**	**70.0**	**93.2**
0	1	1.0	30.6	64.2
0	∞	0.2	3.4	10.2
1	0	25.8	92.2	99.0
1	1	**5.0**	**62.4**	**88.2**
1	∞	0.4	9.8	30.6
∞	0	94.8	100	100
∞	1	56.8	99.2	100
∞	∞	**5.2**	**59.2**	**84.2**

Recent Work with 2005 U of W Graduate Bryan Shepherd:

- Include baseline covariates: Study $ACE(x)$ (See Shepherd et al., 2006, Biometrics, which is on the course webpage)
- Asymptotics
- Extend to binary, count, and failure time outcomes

Including Covariates: Assumptions to Identify ACE(x)

A.1: Stable Unit Treatment Value Assumption (SUTVA)

The potential outcomes for each trial participant are unrelated to the treatment assignment of other subjects (Rubin 1978, 1986)

A.2: Randomization:

$$Z \amalg (S(0), S(1), Y(0), Y(1)) | X$$

A.3: Monotonicity: $S_i(1) \leq S_i(0)$

A.4: $w(X, Y; \beta) = \Phi\{m(X) + g(X, Y; \beta)\}$,

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

Distributional Assumptions

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

M.2: The function $m(x)$ in A.4 follows a parametric model $m(X) = m(X; \alpha)$, where α is an unknown parameter vector and for each $\alpha, m(x; \alpha)$ is a known function

M.3: $f_v(y|x; \eta_1)$ is known up to a finite dimensional parameter η_1

M.4: One of the following:

- a: $f_p(y|x; \eta_0^a)$ is known up to a finite dimensional parameter η_0^a
- b: $f_p^{ai}(y|x; \eta_0^b)$ is known up to a finite dimensional parameter η_0^b

Likelihood (using M.4a)

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

Notes on the Likelihood

Likelihood is identical if Y is discrete outcome

Easily modified to handle censoring

Viral load assay had lower and upper quantitative limits of 400 and 750,000 copies/ml, respectively

Easily modified if Y is time-to-event outcome in the presence of independent right-censoring

Maximum Likelihood Estimation

$$\widehat{E}(Y(1)|ai,x) = \int y dF_v(y|x;\widehat{\eta}_1)dy.$$

$$\begin{aligned}\widehat{E}(Y(0)|ai,x;\beta) &= \int y dF_p^{ai}(y|x;\beta,\widehat{\alpha},\widehat{\eta}_0)dy \\ &= \frac{\int y w(x,y;\beta,\widehat{\alpha})d\widehat{F}_p(y|x;\cdot)dy}{\int w(x,y;\beta,\widehat{\alpha})d\widehat{F}_p(y|x;\cdot)dy}\end{aligned}$$

$$\widehat{ACE}(x;\beta) = \widehat{E}(Y(1)|ai,x) - \widehat{E}(Y(0)|ai,x;\beta)$$

Variance:

Under sufficiently smooth parameterizations, for $VE_S > 0$, the variance of parameter estimates can be estimated as the inverse of the information

Delta method can be used to estimate variance of $\widehat{ACE}(x;\beta)$

Parametric Approach 1: Specify $\mathbf{f_p(y|x;\cdot)}$, e.g.

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

$$\Rightarrow f_p^{ai}(y|x;\cdot) = \frac{w(x,y;\beta,\alpha)\phi(y;x^T\lambda,\sigma^2)}{\int w(x,y;\beta,\alpha)\phi(y;x^T\lambda,\sigma^2)dy}$$

To make $H_0 : ACE(x) = 0$ possible for all x and β, must specify $f_v(y|x;\cdot)$ as

$$f_v(y|x;\eta_1) = \frac{w(x,y;\beta_v,\alpha_v)\phi(y;x^T\lambda_v,\sigma_v^2)}{\int w(x,y;\beta_v,\alpha_v)\phi(y;x^T\lambda_v,\sigma_v^2)dy}$$

Problem: η_1 may not be identifiable $\Rightarrow$ Standard asymptotic theory does not hold

Solution: Bootstrapping provided accurate coverage probabilities in simulations

Parametric Approach 2: Specify $\mathbf{f_p^{ai}(y|x;\cdot)}$:

$$\begin{aligned} f_p^{ai}(y|x;\eta_0^b) &= \phi(y; x^T\gamma_p, \sigma_p^2), \\ f_v(y|x;\eta_1) &= \phi(y; x^T\gamma_v, \sigma_v^2), \end{aligned}$$

where $\gamma = (\gamma_0, \gamma_1, \cdots, \gamma_p)^T$

Advantages:

- Does not force $ACE(x) \neq 0$ for some x
- Easily interpretable parameter estimate: $ACE(x) = x^T(\gamma_v - \gamma_p)$
- Simple tests of interaction ($H_0^{inter} : \gamma_{vj} = \gamma_{pj}$) and global tests of the average causal effect of vaccine on viral load ($H_0^{global} : \gamma_v = \gamma_p$)
- Computationally easier

Disadvantages:

- Requires $w(\cdot) > 0$ for all x and y
- Every time β changes, $f_p(y|x;\cdot)$ takes a different form

Distribution of $\mathbf{Y(0)|X}$ changes with β

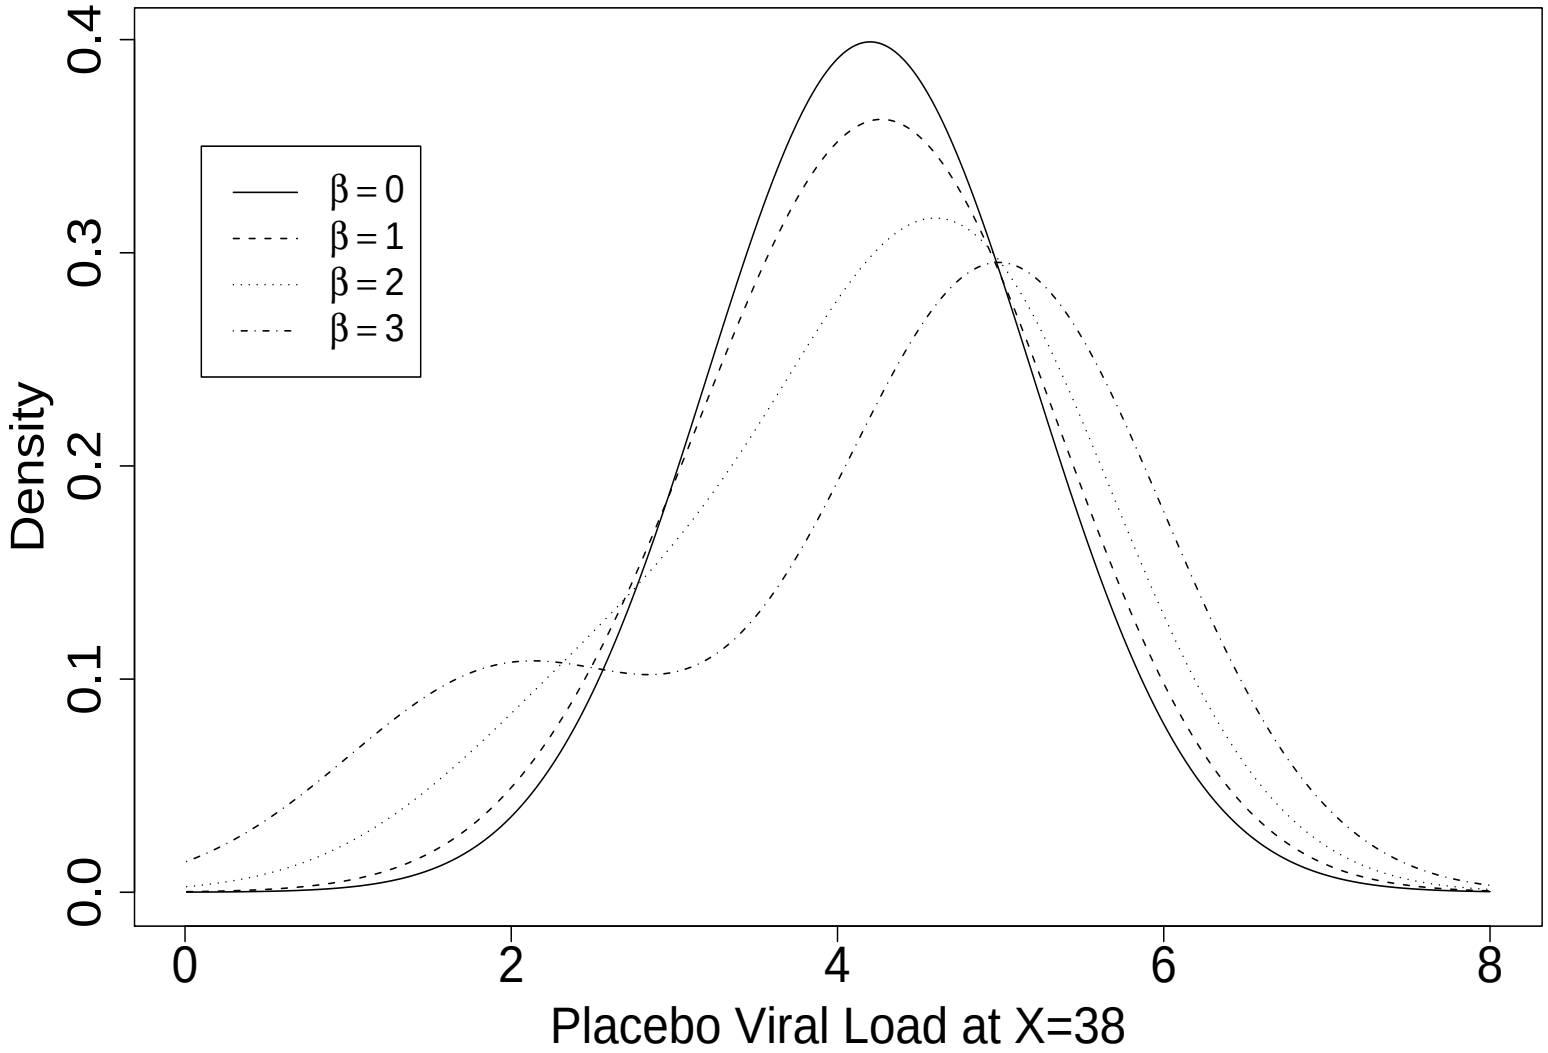

Table 2. Power × 100% for detecting a 0, 1/3, and 1/2 $\log_{10}$ mean-shift alternative, using M.4a, 1-sided 5% level test (from Shepherd et al., 2006)

True	Specified		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	**0**	**4.7**	**5.2**	**5.0**	**25.1**	**63.4**	**12.9**	**50.6**	**93.6**	**23.0**
0	1	26.0	34.2	7.4	5.3	9.8	7.3	9.6	36.0	14.5
0	3	52.0	58.8	8.6	13.7	9.7	6.5	6.6	8.0	10.6
0	0	38.3	47.9	4.8	86.8	99.4	11.5	95.6	100	23.7
1	**1**	**5.6**	**5.6**	**5.9**	**26.4**	**60.5**	**8.6**	**47.7**	**91.9**	**15.9**
1	3	14.6	13.2	6.0	6.6	17.3	7.5	12.3	46.0	12.8
3	0	83.2	95.1	6.7	99.3	100	16.1	99.8	100	28.8
3	1	17.8	27.5	7.4	65.3	94.9	11.4	83.8	99.8	19.7
3	**3**	**4.8**	**7.3**	**8.3**	**18.4**	**52.1**	**8.9**	**37.9**	**86.0**	**16.1**

Parametric Methods

Usual Advantages and Disadvantages

An Additional Disadvantage:

- Compatability of models to H_0 and to data

Recent Semiparametric Work

Jemiai, Rotnitzky, Shepherd, Gilbert (2006, under revision) developed less restrictive semiparametric methods that leave $f_v(\cdot), f_p(\cdot), f_p^{ai}(\cdot)$ unspecified (analogous to GEE)

"always-infected marginal structural mean model":

$$E[Y(Z)|S(0) = S(1) = 1, Z, X] = \mu(Z, X) = \mu(Z, X; \gamma), Z = 0, 1$$

Recent Semiparametric Work

Estimating equation:

$$\sum_{i=1}^{n} d(X_i) q(O_i; \alpha, \gamma) = 0,$$

$$q(O; \alpha, \gamma) = (q_1(O; \gamma), q_2(O; \alpha, \gamma), q_3(O; \alpha), Z - E(Z|X))^T, \qquad (2)$$

$q_1(O; \gamma) = SZ\{Y - \mu(Z, X; \gamma)\}$,

$q_2(O; \alpha, \gamma) = S(1 - Z)w(r(X; \alpha) + g(Y, X))\{Y - \mu(Z, X; \gamma)\}$, and

$q_3(O; \alpha) =$
$\Big[Sw(r(X; \alpha) + g(Y, X))^{(1-Z)} - E\{Sw(r(X; \alpha) + g(Y, X))^{(1-Z)}|X\}\Big]$
$\times\{Z - E(Z|X)\}$

Derived the $d(X)$ that provides a semiparametric efficient estimator

Article available upon request

Example: VaxGen's Phase III Trial of AIDSVAX B/B

Randomized, double-blind, placebo controlled

North America and Netherlands, 1998-2003

$N = 5403$ HIV negative high risk individuals

$N_v = 3598$, $N_p = 1805$, $n_v = 241$, $n_p = 127 \Rightarrow \widehat{V}E_S = 0.048$

Covariates:

- Race (white or non-white)
- Risk behavior score (assumed continuous)

Viral Load:

- Median of all measurements taken by Month 2 post infection visit
- Lower and upper assay quantitation limits of 2.60 and 5.88

VaxGen Trial: Pre-ART Viral Loads

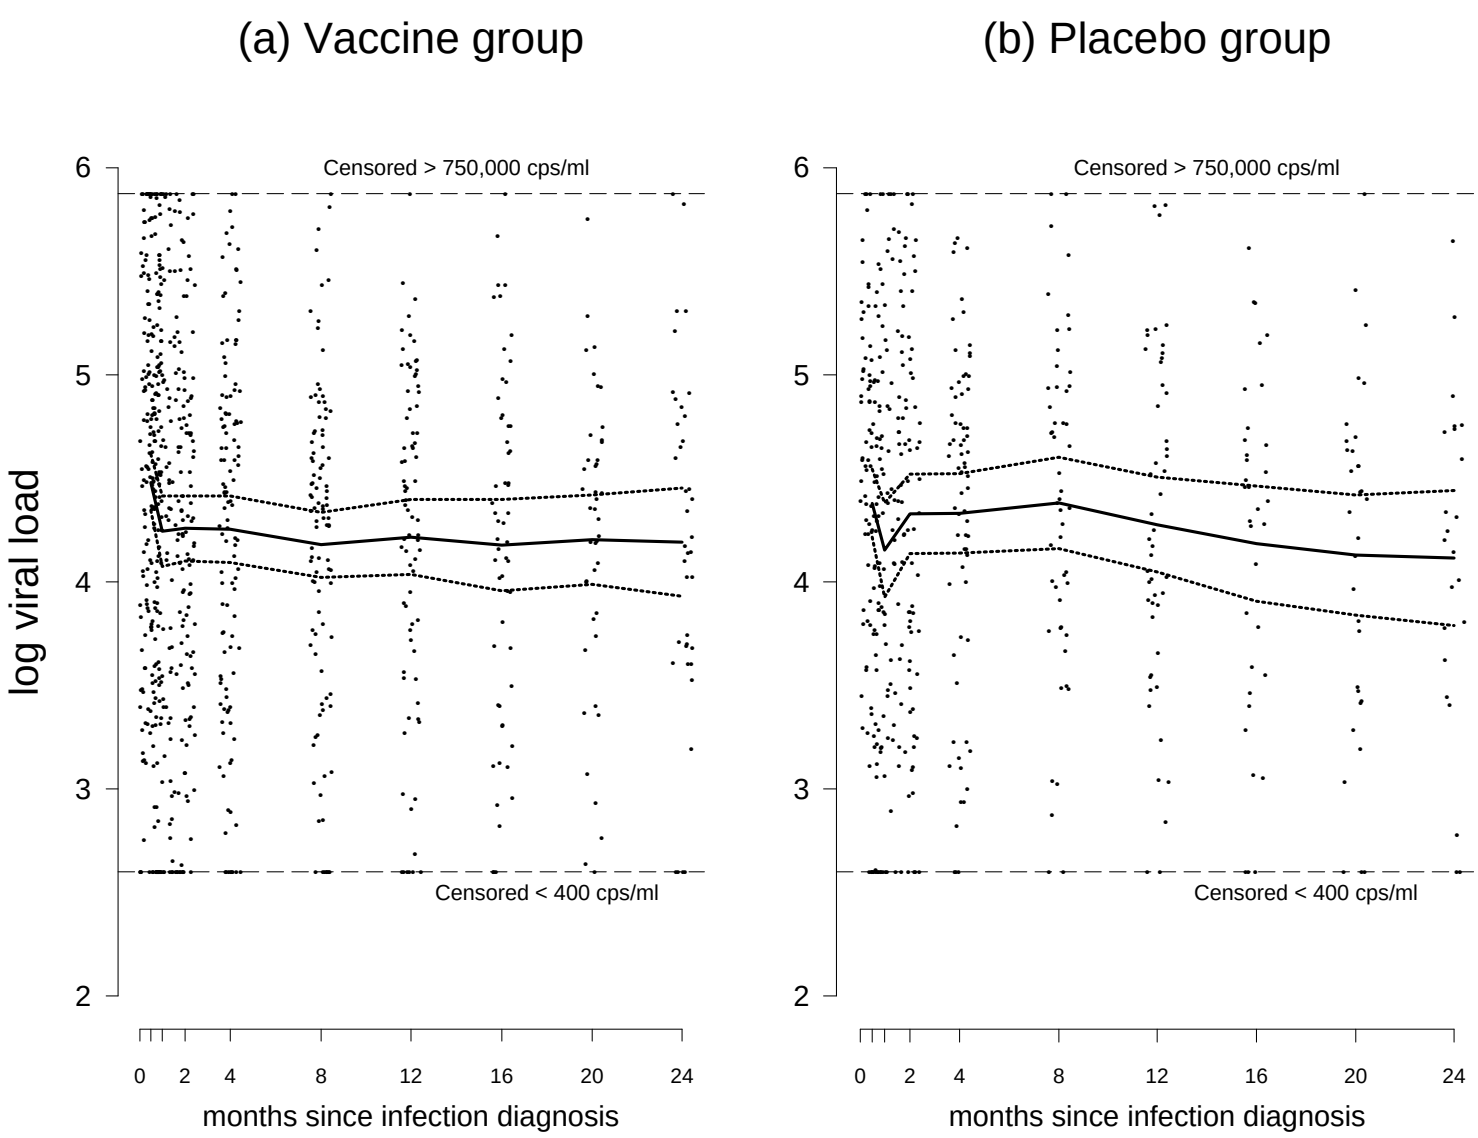

Estimation of ACE with 95% pointwise CIs, using GBH and Shepherd et al. (2006) parametric (M.4a) method to account for assay-censoring

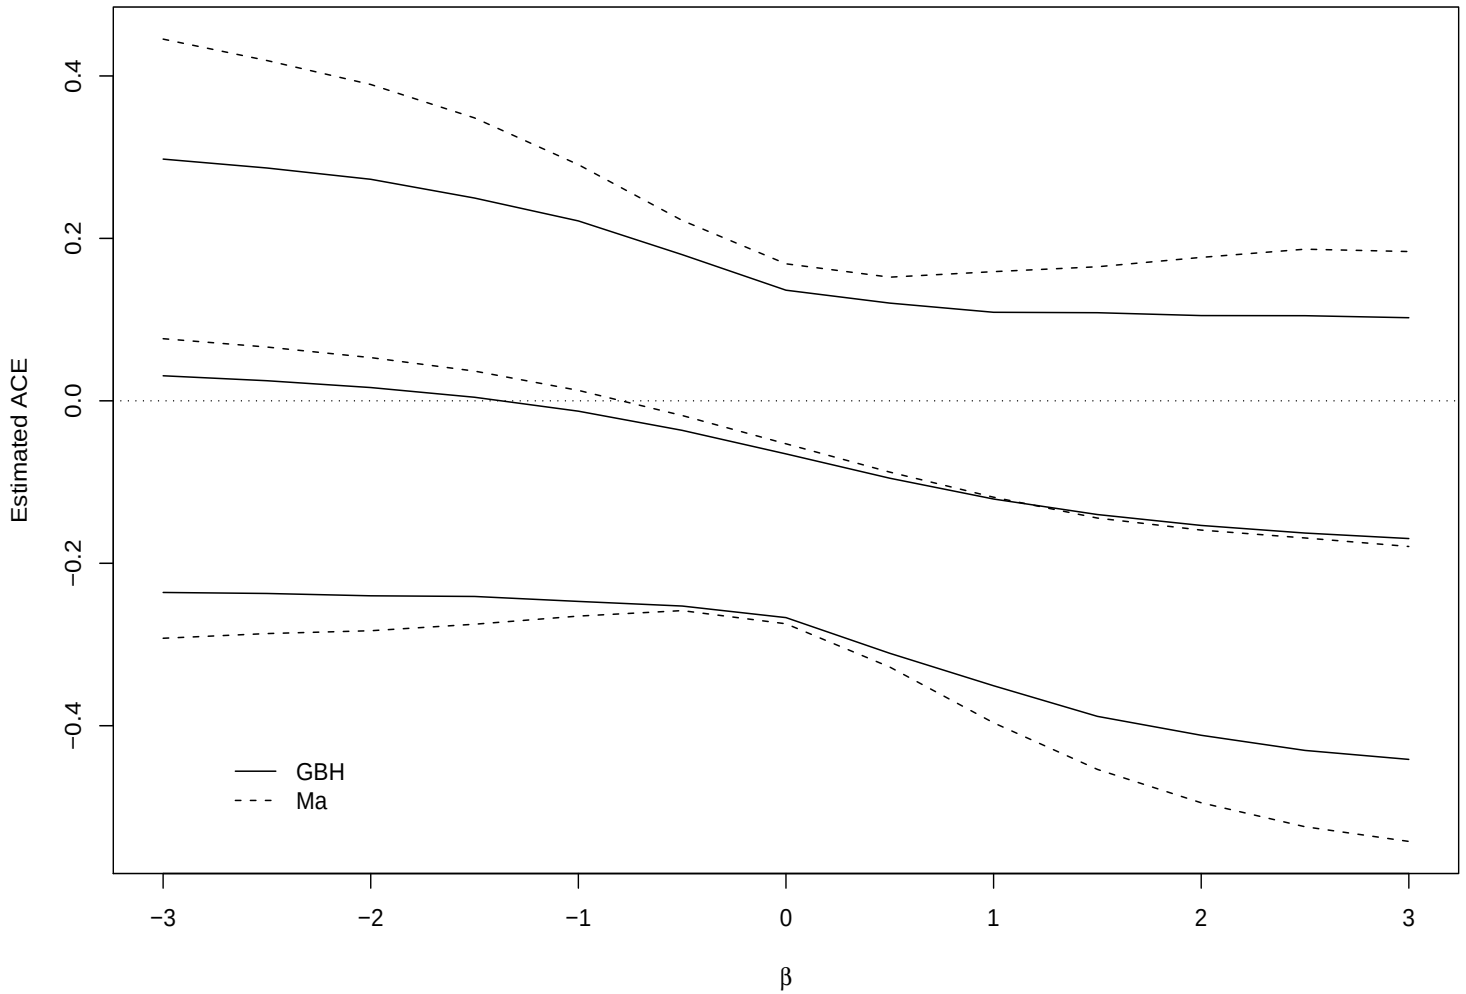

$\widehat{ACE}$ in Non-White Subgroup Conditional on Baseline Risk

$(N_v = 604,\ N_p = 310,\ n_v = 30,\ n_p = 29, \widehat{VE_S} = 0.469)$

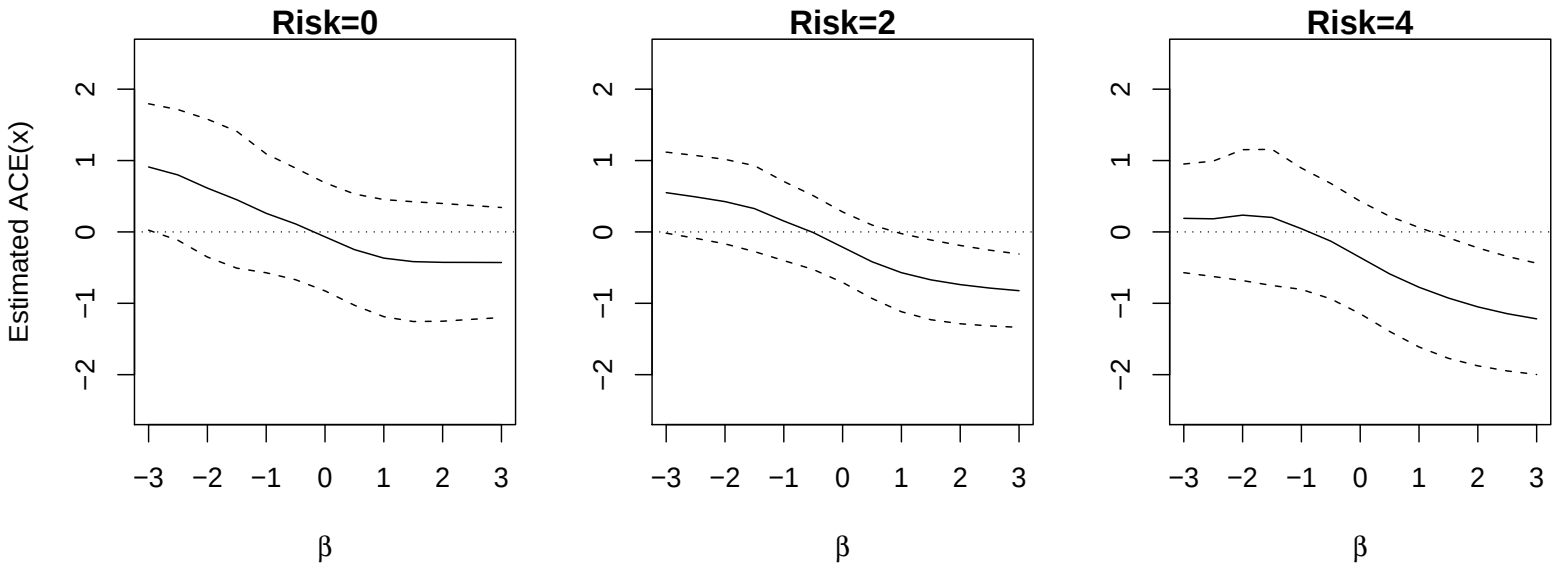

Time-to-Event Outcome

Extend GBH to handle right-censoring:

Solve for $\widehat{\alpha}$ using the equation

$$1 - \widehat{VE_S} = \int_0^\infty \frac{exp(\widehat{\alpha} + \beta y)}{1 + exp(\widehat{\alpha} + \beta y)} d\widehat{F}_p(y)$$

where $1 - \widehat{VE_S} = min[1, (n_v/N_v)/(n_p/N_p)]$ and $\widehat{F}_p$ is the Kaplan-Meier estimate of F_p Then

$$\widehat{F}_p^{ai}(t) = (1 - \widehat{VE_S})^{-1} \int_0^t \frac{exp(\widehat{\alpha} + \beta y)}{1 + exp(\widehat{\alpha} + \beta y)} d\widehat{F}_p(y).$$

$\widehat{F}_v(t) - \widehat{F}_p^{ai}(t)$ is the MLE of the "survival causal effect" parameter

$SCE(t) = \Pr(T(0) > t | (S(0), S(1)) = (1, 1)) - \Pr(T(1) > t | (S(0), S(1)) = (1, 1))$

$= SCE(t) = F_v^{ai}(t) - F_p^{ai}(t)$

Asymptotics:

$$\left(\widehat{F}_p, \widehat{VE}_S\right) \mapsto \left(\int_{[0,\infty]} w(t;\alpha,\beta)d\widehat{F}_p(t) - \left(1 - \widehat{VE_S}\right)\right) \equiv U_n(\alpha)$$

$$0 = U_n(\widehat{\alpha}) = U_n(\alpha) - U_n'(\alpha)(\widehat{\alpha}-\alpha) + \frac{1}{2}U_n''(\alpha^*)(\widehat{\alpha}-\alpha)^2$$

$$\Rightarrow \sqrt{N}(\widehat{\alpha}-\alpha) = \frac{-\sqrt{N}U_n(\alpha)}{U_n'(\alpha) + \frac{1}{2}U_n''(\alpha^*)(\widehat{\alpha}-\alpha)^2},$$

where

$U_n'(\alpha) \to^P U'(\alpha), \widehat{\alpha} \to^P \alpha$, and $U_n''(\alpha^*)$ is bounded

Asymptotic normality of $\widehat{F}_p^{ai}(t)$ established through Hadamard differentiability of the following map:

$$\left(\widehat{\alpha}, \widehat{F}_p(t)\right) \mapsto \frac{\int_{[0,t]} w(y;\alpha,\beta)d\widehat{F}_p(t)}{\int_{[0,\infty]} w(y;\alpha,\beta)d\widehat{F}_p(t)} \equiv \widehat{F}_p^{ai}(t)$$

Table 3. Bias of estimates and coverage probability of Wald-based 95% confidence intervals for $SCE(t)$ with $t = 24$ months (from Shepherd et al., 2006)

VE_S	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	–

Sensitivity Analysis of SCE(1.17 yrs) and SCE(2 yrs)

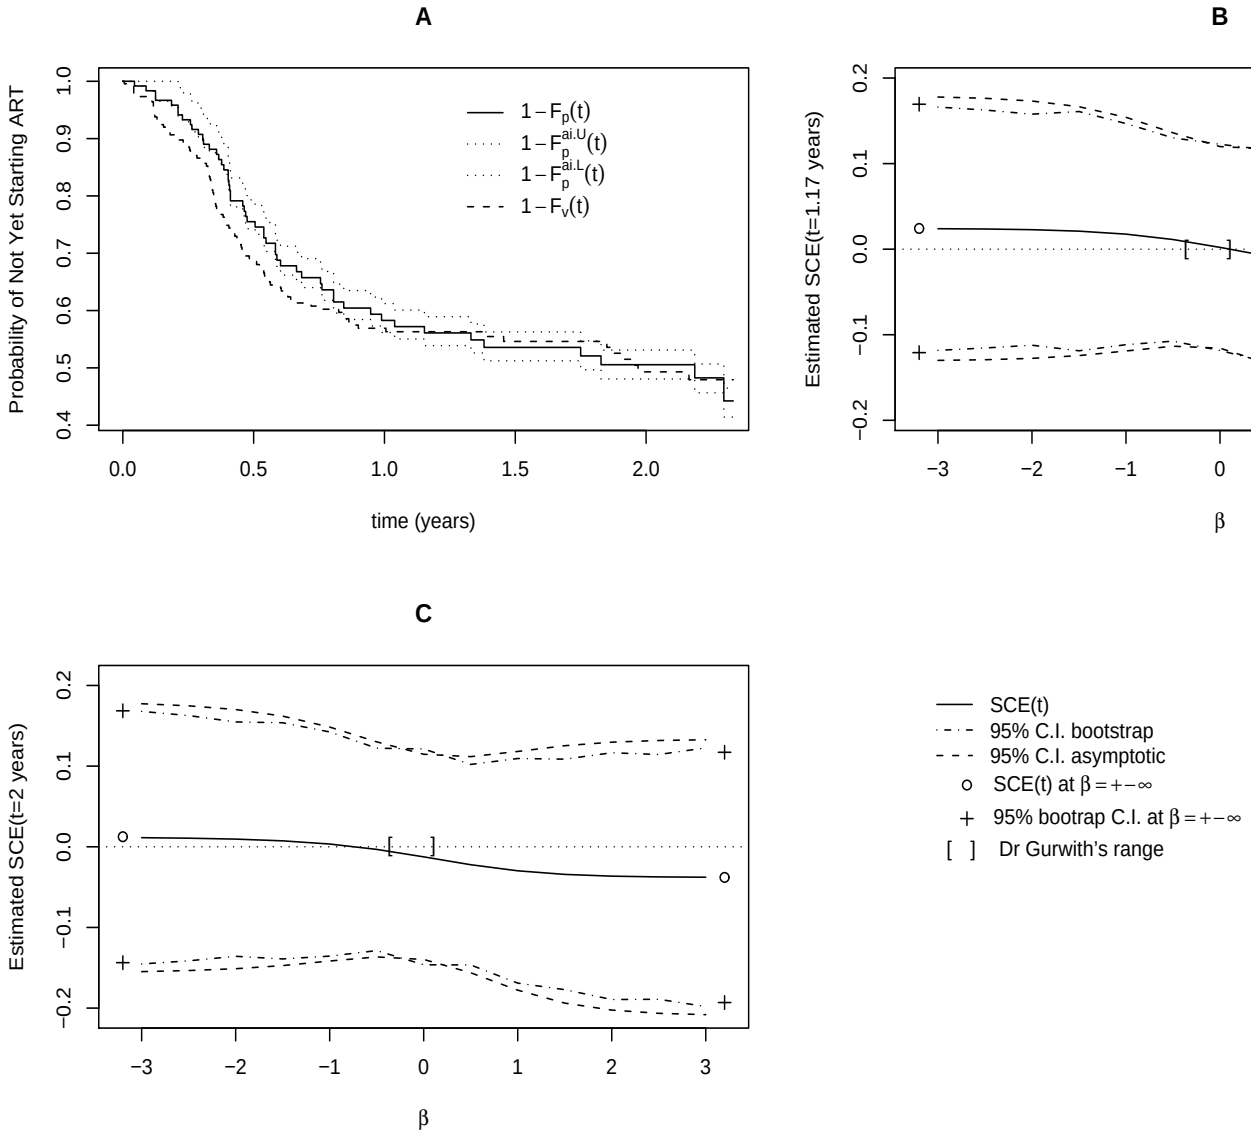

Sensitivity Analysis of $SCE(1.17)$ and $SCE(2)$ for Non-Whites

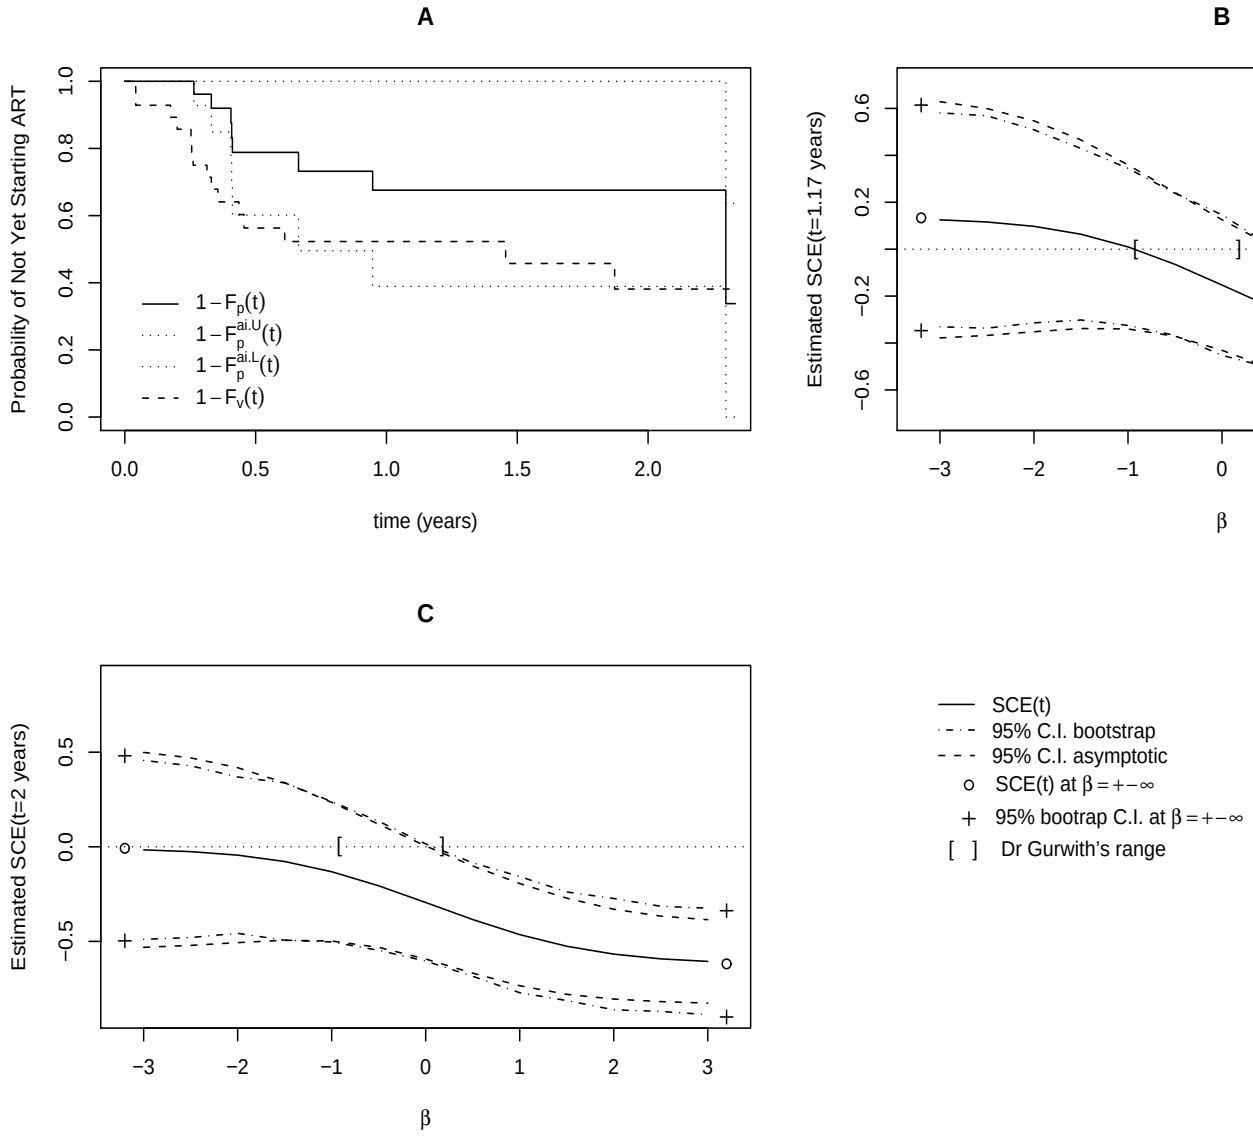

Some Conclusions

Making inferences in the always infected subpopulation useful because:

1. Avoid being misled by estimating a non-causal estimand
2. Assess mechanisms of vaccine effect
3. Can provide far greater power than the ITT approach

Remaining challenges include diagnosing/relaxing monotonicity and SUTVA

Open Research Problems

1. Extend method for survival endpoint (Shepherd et al., 2006) to handle dependent censoring by ART initiation

2. Extend method for longitudinal data (Jemiai et al., 2006) to handle dependent censoring by ART initiation