Introduction to Dynamic Treatment Regimes
Homework 2, Spring 2019

1. Consider the antihypertensive medication observational point exposure study introduced in Problem 3 of Homework 1; the data are in the file `hyper.txt` on the course website. There were $n = 1000$ subjects diagnosed as hypertensive, with baseline characteristics $SBP0$, systolic blood pressure (mmHg); W, weight (kg); K, potassium level (mg/dl); Cr, creatinine level (mg/dl); and Ch, total cholesterol (mg/dl) recorded on each, so that the observed history at baseline is $H_1 = (SBP0, W, K, Cr, Ch)$. The observed treatment indicator A is equal to 1 if the subject received the new medication and equal to 0 if s/he received no treatment ($A = 0$), where this choice was based on patient/physician discretion. Systolic blood pressure at six months, $SBP6$, was measured, and the outcome of interest is $Y = SBP0 - SBP6$, where larger values are considered better.

   The goal in this problem is to estimate, based on these data, an optimal regime $d^{opt} \in \mathcal{D}$ or an optimal restricted regime $d_\eta^{opt} \in \mathcal{D}_\eta$ as defined in (3.42), where $\mathcal{D}_\eta$ is a restricted class of regimes, along with the associated value $\mathcal{V}(d^{opt})$ or $\mathcal{V}(d_\eta^{opt})$, respectively.

   (a) Estimate $d^{opt} \in \mathcal{D}$ using the regression-based estimator $\widehat{d}_Q^{opt} = \{\widehat{d}_{Q,1}^{opt}(h_1)\}$ in (3.26), and estimate $\mathcal{V}(d^{opt})$ using the regression-based estimator $\widehat{\mathcal{V}}_Q(d^{opt})$.

   (b) Consider the restricted class of regimes $\mathcal{D}_\eta$ comprising regimes d_η with rules of the form

$$d_1(h_1; \eta_1) = \mathrm{I}(Ch > \eta_1).$$

   Estimate $d_\eta^{opt} \in \mathcal{D}_\eta$ and its value $\mathcal{V}(d_\eta^{opt})$ using value/policy search estimation based on the IPW estimator $\widehat{\mathcal{V}}_{IPW}(d_\eta)$ in (3.43).

   (c) For the restricted class $\mathcal{D}_\eta$ in (b), estimate $d_\eta^{opt} \in \mathcal{D}_\eta$ and its value $\mathcal{V}(d_\eta^{opt})$ using value/policy search estimation based on the AIPW estimator $\widehat{\mathcal{V}}_{AIPW}(d_\eta)$ in (3.44).

   (d) Estimate $d_\eta^{opt} \in \mathcal{D}_\eta$ and its value $\mathcal{V}(d_\eta^{opt})$ based on the classification analogy for the AIPW estimator in Section 4.1 of the notes, where the restricted class $\mathcal{D}_\eta$ is dictated by the classification and regression trees (CART) approach.

   Except for perhaps (a), your best bet is to use the R package `DynTxRegime`, available at `http://www2.cscc.unc.edu/impact7/DynTxRegime`. At this site, you will also find a link to an online interactive tutorial demonstrating use of the package in detail. You are also free to program the estimators yourself. Modeling choices are up to you!

   Feel free to try additional methods if you want!

2. In the file `cd4.dat` on the course website you will find data from an observational longitudinal study of $n = 5000$ individuals infected with the human immunodeficiency virus (HIV). We also will be using this data set in the next homework assignment, and a detailed description is given here.

   HIV infection compromises the immune system, and a routine measure of the immunological status of an HIV-infected individual is CD4 T-cell count (cells/mm^3), where larger counts are associated with better immunological status. HIV-infected individuals are ordinarily treated with a combination of antiretroviral agents; if antiretroviral therapy is efficacious, CD4 count is expected to increase over time. In fact, greater increases are expected to be associated with lower CD4 counts. With no treatment, CD4 counts are expected to decrease over time.

For some individuals, the virus can develop resistance to the therapy, where the greater the accumulation of therapy over time, the greater the likelihood is that resistance will develop. Once an individual's virus is resistant, therapy is no longer efficacious, although some individuals insist on continuing therapy regardless, in which case CD4 count is expected to decline even more than if therapy were discontinued.

In the study, at baseline (Decision 1), CD4 count was measured on each individual, and an indicator of whether or not the individual then received antiretroviral therapy (coded as 1) or not (coded as 0) at Decision 1 was recorded. All individuals returned at 6 months (Decision 2), at which time CD4 count was recorded for each individual. In addition, the resistance status of the individual's virus was ascertained (0 = not resistant, 1 = resistant), so that the additional information available at Decision 2 are 6 month CD4 count and resistance status. Whether or not the individual received antiretroviral therapy at 6 months was then also recorded. All individuals returned at 12 months (1 year) Decision 3) and at 18 months (Decision 4); at each decision point, CD4 count and resistance status were measured, and whether or not antiretroviral therapy then was administered was recorded.

CD4 counts are well known to have skewed distributions, so it is customary to consider analyses of them on a transformed scale on which they are approximately normally distributed with constant variance; here, we will consider log CD4 counts (so you will need to take logs of the CD4 counts in the data set in all analyses you do).

The primary outcome of interest is log CD4 count at 2 years (24 months).

Summarizing, the 12 columns of the data set are:

Column	Variable
1	CD4 count at baseline
2	Antiretroviral therapy at baseline (0=no, 1=yes)
3	CD4 count at 6 months
4	Resistance status at 6 months (0=not resistant, 1=resistant)
5	Antiretroviral therapy at 6 months (0=no, 1=yes)
6	CD4 count at 12 months
7	Resistance status at 12 months (0=not resistant, 1=resistant)
8	Antiretroviral therapy at 12 months (0=no, 1=yes)
9	CD4 count at 18 months
10	Resistance status at 18 months (0=not resistant, 1=resistant)
11	Antiretroviral therapy at 18 months (0=no, 1=yes)
12	CD4 count at 24 months

All individuals were not resistant to antiretroviral therapy at baseline, so resistance status is equal to 0 for all individuals at Decision 1 and thus is not included as a variable in the data set. An individual's virus cannot become resistant to therapy until s/he has received therapy; and, once an individual's virus becomes resistant, it is resistant henceforth.

In the framework in the notes, $K = 4$. Let C_k = log CD4 count at Decision k, $k = 1, \ldots, K+1$, where Decision $K+1$ corresponds to 2 years, so that the outcome of interest $Y = C_{K+1}$. Let R_k be the indicator of resistance status at Decision k, $k = 1, \ldots, K$, with $R_1 \equiv 0$ for all individuals, and let A_k be the indicator of whether or not antiretroviral therapy was administered at Decision k, $k = 1, \ldots, K$. Then the baseline and intervening information and outcome are

$$X_1 = C_1, \quad X_2 = (C_2, R_2), \quad X_3 = (C_3, R_3), \quad X_4 = (C_4, R_4), \quad Y = C_5,$$

where we do not include R_1 in X_1 as it is a known constant for all individuals.

We take SUTVA (5.10), the SRA (5.11), and the positivity assumption (5.15) to hold for these data.

The goal of this problem is to estimate, based on these data, the value of some given, fixed treatment regimes (defined below) using the g-computation simulation algorithm of Robins (1986), which is described on Slide 281. In doing so, you will use the following posited parametric models as on Slide 279.

First, take

$$C_1 \sim \mathcal{N}(\beta_1, \sigma^2),$$

so that $p_{X_1}(x_1;\zeta_1)$ is the normal density with mean β_1 and variance σ^2, and $\zeta_1 = (\beta_1, \sigma^2)^T$.

For $k = 2, \ldots, K$, assume C_k and R_k are conditionally independent given $\overline{X}_{k-1} = (\overline{C}_{k-1}, \overline{R}_{k-1})$ and $\overline{A}_{k-1}$, so that

$$\begin{aligned} &p_{X_k|\overline{X}_{k-1},\overline{A}_{k-1}}(x_k|\overline{x}_{k-1}, \overline{a}_{k-1}; \zeta_k) \\ &\qquad = p_{C_k|\overline{X}_{k-1},\overline{A}_{k-1}}(c_k|\overline{x}_{k-1}, \overline{a}_{k-1}; \zeta_{k1})\, p_{R_k|\overline{X}_{k-1},\overline{A}_{k-1}}(r_k|\overline{x}_{k-1}, \overline{a}_{k-1}; \zeta_{k2}), \end{aligned}$$

$\zeta_k = (\zeta_{k1}^T, \zeta_{k2}^T)^T$. We specify models for each component on the right hand side as follows.

For $k = 2, \ldots, K+1$, let

$$\begin{aligned} &\mu(C_{k-1}, R_{k-1}, A_{k-1}; \underline{\beta}_2) \\ &\quad = C_{k-1} + \{\beta_2 A_{k-1} + \beta_3 A_{k-1} C_{k-1} + \beta_4(1 - A_{k-1})\}(1 - R_{k-1}) + (\beta_5 + \beta_6 A_{k-1}) R_{k-1}, \end{aligned}$$

$\underline{\beta}_2 = (\beta_2, \beta_3, \beta_4, \beta_5, \beta_6)^T$, and take the conditional distribution

$$C_k|\overline{X}_{k-1}, \overline{A}_{k-1} \sim \mathcal{N}\{\mu(C_{k-1}, R_{k-1}, A_{k-1}; \underline{\beta}_2), \sigma^2\}, \tag{1}$$

so that $p_{C_k|\overline{X}_{k-1},\overline{A}_{k-1}}(c_k|\overline{x}_{k-1}, \overline{a}_{k-1}; \zeta_{k1})$ is the normal density with mean $\mu(c_{k-1}, r_{k-1}, a_{k-1}; \underline{\beta}_2)$ and variance σ^2, and $\zeta_{k1} = (\underline{\beta}_2^T, \sigma^2)^T$. Model (1) allows for the features noted above, as follows. The model implies that, for an individual whose virus is not resistant at Decision $k-1$ and who received therapy at Decision $k-1$, there is an expected increase/decrease in log CD4 count at Decision k of $\beta_2 + \beta_3 C_{k-1}$; whereas if s/he did not receive therapy, there is an expected increase/decrease of β_4. If an individual's virus is resistant at Decision $k-1$, there is an expected increase/decrease in log CD4 count of $\beta_5 + \beta_6 A_{k-1}$.

For $k = 2, \ldots, K$, let $\text{cum}(\overline{A}_{k-1}) = \sum_{j=1}^{k-1} A_j$, the cumulative therapy received up to and including Decision $k-1$. According to the above requirements, an individual's virus cannot be resistant at Decision k unless s/he has already received therapy (so that $\text{cum}(\overline{A}_{k-1}) > 0$). If $R_{k-1} = 1$, it must be that $\text{I}\{\text{cum}(\overline{A}_{k-1}) > 0\} = 1$, and $R_k = 1$ with certainty. If $\text{I}\{\text{cum}(\overline{A}_{k-1}) > 0\} = 0$, then it must be that $R_{k-1} = 0$ and $R_k = 0$ with certainty. Accounting for these restrictions, we posit for $k = 2, \ldots, K$ the logistic regression model

$$\begin{aligned} &P(R_k = 1|\overline{X}_{k-1}, \overline{A}_{k-1}) = \\ &\quad = \left[\frac{\exp\{\psi_1 + \psi_2 \text{cum}(\overline{A}_{k-1}) + \psi_3 \text{cum}(\overline{A}_{k-1}) C_{k-1}}{1 + \exp\{\psi_1 + \psi_2 \text{cum}(\overline{A}_{k-1}) + \psi_3 \text{cum}(\overline{A}_{k-1}) C_{k-1}}\right]^{(1-R_{k-1})} \text{I}\{\text{cum}(\overline{A}_{k-1}) > 0\}. \end{aligned} \tag{2}$$

Thus, with these restrictions and the convention that $0^0 = 1$, from (2), $p_{R_k|\overline{X}_{k-1},\overline{A}_{k-1}}(r_k|\overline{x}_{k-1},\overline{a}_{k-1};\zeta_{k2})$, $k = 2,\ldots,K$, is given by

$$\left(\frac{\exp[\{\psi_1 + \psi_2\text{cum}(\overline{a}_{k-1}) + \psi_3\text{cum}(\overline{a}_{k-1})c_{k-1}\}r_k]}{1+\exp[\{\psi_1 + \psi_2\text{cum}(\overline{a}_{k-1}) + \psi_3\text{cum}(\overline{a}_{k-1})c_{k-1}\}r_k]}\right)^{(1-r_{k-1})\text{I}\{\text{cum}(\overline{a}_{k-1})>0\}},$$

and $\zeta_{k2} = (\psi_1,\psi_2,\psi_3)^T$.

Finally, when $k = K+1$, $Y = C_{K+1}$,

$$\mu(C_K,R_K,A_K;\underline{\beta}_2) = C_K + \{\beta_2 A_K + \beta_3 A_K C_K + \beta_4(1-A_K)\}(1-R_K) + (\beta_5+\beta_6 A_K)R_K,$$

and (1) becomes

$$Y|\overline{X},\overline{A} \sim \mathcal{N}\{\mu(C_K,R_K,A_K;\underline{\beta}_2),\sigma^2\},$$

so that $p_{Y|\overline{X},\overline{A}}(y|\overline{x},\overline{a};\zeta_{K+1})$ is the normal density with mean $\mu(c_K,r_K,a_K;\underline{\beta}_2)$ and variance σ^2, and $\zeta_{K+1} = (\underline{\beta}_2^T,\sigma^2)^T$.

From above, the parameters ζ_k, $k = 1,\ldots,K+1$, overlap across k, and in fact $\zeta = (\beta_1,\beta_2,\ldots,\beta_6,\sigma^2,\psi_1,\ldots,\psi_3)^T$.

The partial likelihood for all parameters can now be specified as on Slide 279 as

$$\prod_{i=1}^{n}\Bigg[p_{Y|\overline{X},\overline{A}}(y|\overline{x},\overline{a};\underline{\beta}_2,\sigma^2) \times \prod_{k=2}^{K}\left\{p_{C_k|\overline{X}_{k-1},\overline{A}_{k-1}}(c_k|\overline{x}_{k-1},\overline{a}_{k-1};\underline{\beta}_2,\sigma^2)\,p_{R_k|\overline{X}_{k-1},\overline{A}_{k-1}}(r_k|\overline{x}_{k-1},\overline{a}_{k-1};\psi_1,\psi_2,\psi_3)\right\} \times p_{X_1}(x_1;\beta_1,\sigma^2)\Bigg],$$

where we show explicitly the parameters on which each component depends. It is straightforward to observe that $(\beta_1,\beta_2,\ldots,\beta_6,\sigma^2)^T$ separate from $(\psi_1,\ldots,\psi_3)^T$ in the partial likelihood, and thus each set of parameters can be estimated separately.

(a) It is straightforward to show (you may want to try it or take our word for it) that maximizing the portion of the partial likelihood depending on $(\beta_1,\beta_2,\ldots,\beta_6,\sigma^2)^T$ in these parameters is equivalent to finding the maximum likelihood (ML) estimators in the linear model

$$\mathcal{Y} = \mathbb{X}\underline{\beta} + \epsilon, \quad \underline{\beta} = (\beta_1,\beta_2,\ldots,\beta_6)^T. \tag{3}$$

Here,

$$\mathcal{Y} = (\mathcal{Y}_1^T,\ldots,\mathcal{Y}_n^T)^T,$$

where

$$\mathcal{Y}_i = \begin{pmatrix} C_{1i} \\ C_{2i}-C_{1i} \\ C_{3i}-C_{2i} \\ C_{4i}-C_{3i} \\ Y_i - C_{4,i} \end{pmatrix},$$

is the $\{(K+1)n \times 1\}$ "outcome vector. " The "design matrix" $\mathbb{X}$ is such that the rows that correspond to C_{1i}, $i = 1,\ldots,n$, are

$$(1,0,0,0,0,0),$$

and the rows corresponding to $(C_{ki} - C_{k-1,i})$,$i = 1, \ldots, n$, are

$$[0, A_{k-1,i}(1 - R_{k-1,i}), A_{k-1,i}C_{k-1,i}(1 - R_{k-1,i}), (1 - A_{k-1,i})(1 - R_{k-1,i}), R_{k-1,i}, R_{k-1,i}A_{k-1,i}],$$

$k = 2, \ldots, K + 1$, $i = 1, \ldots, n$; and the "deviation"

$$\epsilon \sim \mathcal{N}(0, \sigma^2 I_{(K+1)n},$$

where $I_{(K+1)n}$ is a $\{(K + 1)n \times (K + 1)n\}$ identity matrix. The ML estimators for $\underline{\beta}$ are then the ordinary least squares estimators in the model (3), and the ML estimator for σ^2 can be obtained as the usual sum of squared deviations divided by the "sample size" $(K + 1)n$.

Estimate $(\beta_1, \beta_2, \ldots, \beta_6, \sigma^2)^T$ using this approach.

Note: You can construct $\mathcal{Y}$ and $\mathbb{X}$ yourself and do this manually, or you can use linear model software. If you do the latter, be sure that in specifying your model you invoke a "no intercept" option so that the usual automatic intercept term is not included.

(b) Likewise, it is straightforward to show that the part of the partial likelihood involving $(\psi_1, \ldots, \psi_3)^T$ can be written as

$$\prod_{i=1}^{n} \prod_{\{k:\, \text{cum}(\overline{A}_{k-1,i})>0 \text{ and } R_{k-1,i}=0\}} \frac{\exp[\{\psi_1 + \psi_2 \text{cum}(\overline{A}_{k-1,i}) + \psi_3 \text{cum}(\overline{A}_{k-1,i})C_{k-1,i}\}R_{k,i}]}{1 + \exp\{\psi_1 + \psi_2 \text{cum}(\overline{A}_{k-1,i}) + \psi_3 \text{cum}(\overline{A}_{k-1,i})C_{k-1,i}\}}. \quad (4)$$

Here, there is no contribution to (4) at Decision k unless an individual has already received therapy and has virus that is not yet resistant (as the contribution is equal to 1 otherwise). The objective function (4) has the form of the likelihood for standard logistic regression, where each individual i contributes terms for each decision point k at which therapy has been administered previously and resistance has not yet taken place. The ML estimator for $(\psi_1, \psi_2, \psi_3)^T$ can be found by maximizing (4), which is what is done in standard logistic regression software, where, here, the data comprise all pairs $(R_{ki},\ \text{cum}(\overline{A}_{k-1,i})$ for each decision point k at which therapy has been administered previously and resistance has not yet taken place for each individual $i = 1, \ldots, n$.

Estimate $(\psi_1, \psi_2, \psi_3)^T$ using this approach.

(c) Now that you have your fitted models, it is straightforward to implement the Monte Carlo g-computation algorithm on Slide 281 for any given regime d. In this and the next two parts of this problem, you will estimate the value $\mathcal{V}(d) = E\{Y^*(d)\}$, where $Y^*(d)$ is the potential log CD4 count at two years under regime d, for a bunch of given, fixed regimes d. (*Hint:* Code your algorithm so that it calls a function to evaluate $\overline{d}_{k-1}(\overline{x}_{k-1})$, $k = 2, \ldots, K+1$. Then all you need do is define that function to compute the value for any fixed regime d.)

Here, consider the eight static regimes encoded as

$d^{(1)}$	(0,0,0,0)
$d^{(2)}$	(0,0,0,1)
$d^{(3)}$	(0,0,1,1)
$d^{(4)}$	(0,1,1,1)
$d^{(5)}$	(1,0,0,0)
$d^{(6)}$	(1,1,1,0)
$d^{(7)}$	(1,1,0,0)
$d^{(8)}$	(1,0,0,0),

where, for example, $d^{(1)}$ is the regime whose 4 decision rules always do not administer antiretroviral therapy, regardless of past CD4, resistance status, or treatment history; and $d^{(6)}$ is the regime whose first three rules alway administer therapy regardless of history and whose final rule always does not.

For each, estimate the value using your g-computation algorithm. It is up to you what to choose for M, the number of simulations. Also, use a nonparametric bootstrap to obtain an approximate standard error to accompany each estimate. 50 bootstrap data sets should be adequate, but feel free to use more (you will see that all of this is not computationally intensive).

(d) As noted above, individuals whose virus has become resistant should not be treated. Thus, consider the eight regimes $d^{(1)*}, \ldots, d^{(8)*}$, where regime $d^{(\ell)*}$ administers therapy or not according to regime $d^{(\ell)}$ in (c) as long an individual's virus is not resistant but discontinues therapy when the virus becomes resistant. Thus, $d^{(4)*}$ would not administer therapy at Decision 1, would administer therapy at Decision 2 if $R_2 = 0$ but would not administer it at Decision 2 and henceforth if $R_2 = 1$, would administer therapy at Decision 3 if $R_3 = 0$ but would not administer it at Decision 3 and henceforth if $R_3 = 1$, and so on.

For each of these regimes, estimate the value using your g-computation algorithm and obtain accompanying standard errors using a nonparametric bootstrap.

(e) Now consider regimes with rules of the form

$$\begin{aligned} d_k(\overline{x}_k, \overline{a}_k; \eta) &= 1 \quad \text{if } c_k < \log(\eta) \text{ and } r_k = 0 \\ &= 0 \quad \text{otherwise,} \end{aligned}$$

$k = 1, \ldots, 4$, for a fixed choice of the threshold η. In words, at Decision k, the rule dictates administering therapy at Decision k if the current CD4 count is $< \eta$ as long as the individual's virus is not yet resistant; otherwise, do not administer therapy. Thus, such a regime administers therapy only if CD4 count becomes sufficiently low (with no resistance), reflected by the choice of the threshold η.

Using your g-computation algorithm, estimate the value of regimes d with rules of this form for $\eta = 100, 200, 300$, and 400 cells/mm^3 and obtain accompanying standard errors using a nonparametric bootstrap.

(f) Comment on the results you obtained in (c)-(e).

3. Consider the HIV study data in Problem 2, where $K = 4$. Now let's estimate the value of some regimes using the simple inverse probability weighted estimators in (5.33) and (5.34) of the notes. As in the previous problem, take SUTVA (5.10), the SRA (5.11), and the positivity assumption (5.15) to hold.

   There are two treatment options at each decision point. Assume that the probability of receiving antiretroviral therapy each Decision conditional on past history, the propensity score

$$\pi_k(h_k) = P(A_k = 1 | H_k = h_k) = \pi_k(\overline{x}_k, \overline{a}_{k-1}), \quad k = 1, \ldots, K,$$

   depends on the past history only through the most recent CD4 count and resistance status measures, i.e., at Decision k, for $k = 2, 3, 4$. At Decision 1, where all individuals have non-resistant virus at baseline, assume that the propensity score depends only on baseline CD4

count. Model the propensity scores under these conditions using logistic regression models; thus, take

$$\pi_1(x_1) = P(A_1 = 1|X_1 = x_1) = \frac{\exp(\gamma_{11} + \gamma_{12}C_1)}{1 + \exp(\gamma_{11} + \gamma_{12}C_1)},$$
$$\pi_k(\overline{x}_k, \overline{a}_{k-1}) = P(A_k = 1|\overline{X}_k = \overline{x}_k, \overline{A}_{k-1} = \overline{a}_{k-1}) = \frac{\exp(\gamma_{k1} + \gamma_{k2}C_k + \gamma_{k3}R_k)}{1 + \exp(\gamma_{k1} + \gamma_{k2}C_k + \gamma_{k3}R_k)}, \quad k = 2, 3, 4. \tag{5}$$

(a) Fit the propensity score models in (5).

(b) Using your fitted models, estimate the value of regime d using the estimator $\widehat{\mathcal{V}}_{IPW}(d)$ in (5.33), where d is each of the regimes in (c), (d), and (e) of Problem 2. Use a nonparametric bootstrap to obtain an approximate standard error to accompany each estimate.

(c) Using your fitted models, estimate the value of regime d using the estimator $\widehat{\mathcal{V}}_{IPW*}(d)$ in (5.34), where d is each of the regimes in (c), (d), and (e) of Problem 2. Use a nonparametric bootstrap to obtain an approximate standard error to accompany each estimate.

(d) Comment on the results in (b) and (c) and how they compare to those you obtained using the g-computation algorithm.