

Economics 696F, Causal Inference and Program Evaluation

Lecture Note 7: More on Multivalued Treatments; Dynamic Treatments and Sequential Randomization

1 GPS Revisited

The generalized propensity score (GPS) notation developed in the previous note might be difficult to follow, especially the idea of quantities like $r(t, X)$, where t is a fixed number and X is regarded as a random variable. Here is a slight reworking which makes the connection to the binary-treatment case clearer (I hope):

Recall that the generalized propensity score is defined as:

$$r(t, x) = Pr(T = t | X = x),$$

where $t \in \{t_1, \dots, t_M\}$ is a particular treatment value. We want to estimate quantities like

$$E[Y(t_j)]$$

for given treatment t_j .

To be more concrete, let us focus on the first (out of M) possible treatments t_1 . So $E[Y(t_1)]$ is the average outcome if everyone in the population had been given treatment t_1 .

In the sample, we observe individuals with various treatments. We define $D(t_1)$ to be the indicator for receiving treatment t_1 . We can think of dividing the sample into those who receive t_1 , and those who receive any other treatment, based on $D(t_1)$.

We assume that

$$D(t_1) \perp Y(t_1) \mid X.$$

So, conditional on X , individuals who receive treatment t_1 are not systematically different (in their $Y(t_1)$ potential outcome) from individuals who receive some other treatment.

Let

$$p_1(x) = P(T = t_1 | X = x) = r(t_1, x).$$

This is interpreted as the probability of an individual receiving t_1 given that $X = x$. Since we are focusing on a “binary” treatment $D(t_1)$, we can think of this as an ordinary propensity score for whether or not the individual receives t_1 .

Therefore, by the standard (binary treatment) argument, it follows that

$$D(t_1) \perp Y(t_1) \mid p_1(X)$$

Let

$$\begin{aligned}\beta(t_1, r) &= E[Y \mid T = t_1, p_1(X) = r] \\ &= E[Y(t_1) \mid T = t_1, p_1(X) = r] \\ &= E[Y(t_1) \mid D(t_1) = 1, p_1(X) = r] \\ &= E[Y(t_1) \mid p_1(X) = r]\end{aligned}$$

By iterated expectations,

$$\begin{aligned}E[Y(t_1)] &= E[E[Y(t_1) \mid p_1(X)]] \\ &= E[\beta(t_1, p_1(X))].\end{aligned}$$

So we can estimate $E[Y(t_1)]$ by first estimating the function $\beta(t_1, r)$, then average with respect to the marginal distribution of $p_1(X)$.

The weighting argument works similarly. Because we have turned the problem into one with a binary treatment (by focusing on t_1 vs. “any other treatment”), we have by our argument in the binary case:

$$\begin{aligned}E[Y(t_1)] &= E\left[\frac{Y \cdot D(t_1)}{p_1(X)}\right] \\ &= E\left[\frac{Y \cdot D(t_1)}{r(t_1, X)}\right]\end{aligned}$$

But since $D(t_1) = 1$ if and only if $T = t_1$, this equals

$$E\left[\frac{Y \cdot D(t_1)}{r(T, X)}\right].$$

We can repeat the argument for different possible treatments; for example if we focus on t_2 , we use $D(t_2)$ and define

$$p_2(x) = P(T = t_2 \mid X = x) = r(t_2, x).$$

2 Dynamic Treatments and Sequential Randomization

This section is based on Robins (1998), and Robins, Hernan, and Brumback (2000).

2.1 Model Setup – Dynamic Treatment Regimes and Outcome Process

Treatments are administered over periods $s = 1, 2, \dots, S$.

Let \mathcal{T}_s denote the possible treatment levels in period s .

So a particular treatment regime is an S -vector

$$(t_1, t_2, \dots, t_S) \in \mathcal{T}_1 \times \mathcal{T}_2 \times \dots \times \mathcal{T}_S.$$

(Note: we are using the subscript s different from the previous note – now it refers to time periods.)

Let

$$T = (T_1, T_2, \dots, T_S)$$

denote the vector of treatments received, which we call a dynamic treatment regime.

To define the outcome process, let $L_s = (Y_s, X_s)$ for $s = 1, \dots, S$. Here we interpret Y_s as the main outcome of interest in time period s , and X_s to be other time-varying variables. Both Y_s and X_s could be affected by treatment up to time s , so we can think of potential outcomes

$$L_s(t_1, \dots, t_s) = (Y_s(t_1, \dots, t_s), X_s(t_1, \dots, t_s)).$$

These are connected to observed outcomes by

$$L_s = L_s(T_1, \dots, T_s).$$

Here we are implicitly ruling out effects of future treatments (e.g. t_{s+1} does not affect L_s).

We will focus on estimating effects of treatments on terminal outcomes Y_S . Our goal is to estimate

$$E[Y_S(t_1, \dots, t_S)],$$

for a given treatment regime (t_1, \dots, t_S) .

2.2 Sequential Randomization Assumption

Assume that

$$Y_s(t_1, \dots, t_s) \perp T_s \mid T_1, \dots, T_{s-1}, L_1, \dots, L_{s-1}.$$

This allows treatment at time s to depend on the intermediate outcomes Y_1, \dots, Y_{s-1} .

2.3 Marginal Structural Model

$$E[Y_S(t_1, \dots, t_S)] = f(t_1, \dots, t_S; \beta).$$

For example, if Y_S is binary, we could use a logistic specification

$$E[Y_S(t_1, \dots, t_s)] = P[Y_S(t_1, \dots, t_s)] = \Lambda \left(\beta_0 + \beta_1 \sum_{s=1}^S t_s \right),$$

where

$$\Lambda(z) = \frac{\exp(z)}{1 + \exp(z)}.$$

This makes the final outcome depend on the sum of the treatments over the S time periods.

2.4 Weighting by Inverse Probability of Treatment

Constructing weights: specify a parametric model for

$$T_s \mid T_1, \dots, T_{s-1}, L_1, \dots, L_{s-1},$$

for $s = 1, \dots, S$.

For example, if T_s is binary we could use a sequence of logit models

$$P[T_s = 1 \mid T_1, \dots, T_{s-1}, L_1, \dots, L_{s-1}] = \Lambda(\gamma_0 + \gamma_1 s + \gamma_2 T_{s-1} + \gamma_3' L_{s-1}).$$

(This assumes that treatment depends on only the last period's variables, but in principle we could use more flexible models.)

We could then estimate the parameters $(\gamma_0, \dots, \gamma_3)$ by MLE.

Similarly, also model and estimate

$$T_s \mid T_1, \dots, T_{s-1}.$$

Weights:

$$W_i := \frac{\prod_{s=1}^S \widehat{P}[T_s = T_{si} \mid T_{1i}, \dots, T_{s-1,i}, L_{1i}, \dots, L_{s-1,i}]}{\prod_{s=1}^S \widehat{P}[T_s = T_{si} \mid T_{1i}, \dots, T_{s-1,i}]}$$

Then run a weighted regression of Y_S on (T_1, \dots, T_S) with observations weighted by W_i , to estimate β .

As with the propensity score weighting estimators in non-dynamic settings, the use of estimated weights generally reduces the asymptotic variance of the estimator, relative to using the true conditional treatment probabilities. Thus, ignoring the estimation of the propensity score in calculating standard errors for β will lead to conservative confidence intervals and tests. Alternatively, with some work we could obtain “correct” standard errors.