Adaptive Design of Multiple Stage Experiments using the Propensity Score

Jinyong Hahn (UCLA)

Keisuke Hirano (Arizona)

Dean Karlan (Yale)

Outline

- Background/ Motivation/ Objective $\Leftarrow$
- Summary
- Counterfactual Framework in Observational Studies
- Details of Our Procedure
- Numerical Examples, Extensions

Background 1

- Substantial econometric development for estimating the effects of labor market programs, e.g., since the work by Card and Sullivan (1988), Heckman and Robb (1984), and Lalonde (1986) among others.

- One strand of this literature has developed methods for estimating average treatment effects for a binary treatment under assumptions of exogeneity or unconfoundedness, also known as selection on observables.

- Various approaches to semiparametric estimation have been proposed.

Background 2

- Randomized experiments have now become a standard tool of empirical analysis, but not much econometric development has been made.

- It is quite intuitive to believe that randomized experiment would render econometrics irrelevant, except when experiment has a fundamental limitation.

- Hopefully econometrics may be of help in designing an optimal experiment!

Inspiration

- "What's Psychology Worth? A Field Experiment in the Consumer Credit Market" by Bertrand, Karlan, Mullainathan, Shafir, and Zinman
- Their study was conducted in three waves
- The first wave was a "pilot study"
- The second and third waves correspond to our first and second samplings
- For administrative reasons, the DGP remained fixed over the second and third "waves", but if it weren't, maybe we could have done something to improve efficiency.

Objective

- Apply the cumulative knowledge in the program evaluation literature, and develop a method of designing an optimal experiment.

- Depending on the financial cost in implementing an experiment, it may be desirable to design an experiment in such a way that the resultant inference out of the experimental data set would be as precise as possible.

- The paper establishes that the question of optimal experimental design can be understood within the existing econometric framework of observational studies.

Related Literature

- Biostatistics literature on adaptive treatment rules and sequential design
  - Design of an "optimal" one-stage experiment, or about optimal stopping time (e.g., Chernoff, 1959).
  - Potential confusion: "adaptiveness" means something completely different sometimes. Can we use the information in the middle of the experiment to increase the social welfare while maintaining unbiasedness?
  - Most previous literature does not incorporate covariate information.
- Also related to literature on stratified sampling (e.g. Manski and McFadden, 1981)

Outline

- Background/ Motivation/ Objective
- Summary of Our Proposal $\Leftarrow$
- Counterfactual Framework in Observational Studies
- Details of Our Procedure
- Numerical Examples, Extensions

Summary of Our Proposal for Multi-stage Experimental Design

- Objective of interest: Average treatment effect
- First stage: Pure randomization (Pure randomization not really necessary. More on this later.)
- Second stage: Make treatment probability depend on observed covariates. In other words, data from Stage 1 is used to estimate the optimal treatment probability for Stage 2
- Finally, use all data to estimate effect of treatment

Outline

- Background/ Motivation/ Objective
- Summary of Our Proposal
- Counterfactual Framework in Observational Studies $\Leftarrow$
- Details of Our Procedure
- Numerical Examples, Extensions

Potential Outcome Framework: Notation

- D_i denotes a dummy variable indicating a treatment status $D_i = 1$ when treatment is given to the ith individual, and $D_i = 0$ otherwise.

- Y_{0i} and Y_{1i} denote potential outcomes when $D_i = 0$ and $D_i = 1$.

- Econometrician only observes $Y_i = D_i Y_{1i} + (1 - D_i) Y_{0i}$.

- The treatment causes the outcome variable of the ith individual to increase by $Y_{1i} - Y_{0i}$.

- Average treatment effects $\beta = E\left[Y_{1i} - Y_{0i}\right]$

Potential Outcome Framework: Assumption of Selection on Observables

1. $\exists\ X_i$ such that D_i is ignorable given X_i, *i.e.*, $D_i \perp (Y_{0i}, Y_{1i})|\ X_i$

2. Regularity condition $0 < P\left[D_i = 1|\ X_i\right] < 1,\ \forall\ X_i$

- In "experiments", no need to worry about the regularity conditions.

Potential Outcome Framework: Semiparametric Efficiency Bound

- If $D_i \perp (Y_{0i}, Y_{1i}) | X_i$, then the semiparametric asymptotic variance bound for β is

$$E\left[\frac{\sigma_1^2(X_i)}{p(X_i)} + \frac{\sigma_0^2(X_i)}{1 - p(X_i)} + (\beta(X_i) - \beta)^2\right] \qquad (1)$$

where

$$\begin{aligned} \sigma_1^2(X_i) &= \text{Var}(Y_{1i} | X_i) \\ \sigma_0^2(X_i) &= \text{Var}(Y_{0i} | X_i) \\ \beta(X_i) &= E[Y_{1i} - Y_{0i} | X_i] \\ p(X_i) &= P[D_i = 1 | X_i] \qquad \text{(Propensity Score)} \end{aligned}$$

Efficient Semiparametric Estimator for ATE

- Efficient estimator for β takes the form of

$$\widehat{\beta} = \frac{1}{n} \sum_{i=1}^{n} \left(\widehat{r}_1 (X_i) - \widehat{r}_0 (X_i) \right) \tag{2}$$

where $\widehat{r}_1 (X_i)$ and $\widehat{r}_0 (X_i)$ are nonparametric analogs of

$$r_1 (X_i) = \frac{E [D_i Y_i | X_i = x]}{E [D_i | X_i = x]}, \quad r_0 (X_i) = \frac{E [(1 - D_i) Y_i | X_i = x]}{E [1 - D_i | X_i = x]}$$

- Slightly different characterization of (2):

$$\widehat{\beta} = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{D_i Y_i}{\widehat{p} (X_i)} - \frac{(1 - D_i) Y_i}{1 - \widehat{p} (X_i)} \right) \tag{3}$$

where $\widehat{p} (x)$ denotes a nonparametric estimator of the propensity score.

- The characterization in (3) is due to Hirano, Imbens, and Ridder (2000).

- In order to simplify notation, it will be assumed that the covariate X is multinomial.

- Although the estimator (3) is only first order asymptotically equivalent to (2) in general, it is numerically identical to (2) under our special case where the covariate X has a finite support.

$$
\begin{aligned}
\sqrt{n}\left(\widehat{\beta}-\beta\right) \approx\; & \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left(r_1\left(X_i\right)-r_0\left(X_i\right)-\beta\right) \\
& +\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left(\frac{D_i\left(Y_i-r_1\left(X_i\right)\right)}{p\left(X_i\right)}-\frac{\left(1-D_i\right)\left(Y_i-r_0\left(X_i\right)\right)}{1-p\left(X_i\right)}\right)
\end{aligned}
$$

Outline

- Background/ Motivation/ Objective
- Summary of Our Proposal
- Counterfactual Framework in Observational Studies
- Details of Our Procedure $\Leftarrow$
- Numerical Examples, Extensions

Ideal Propensity Score

- Asymptotic variance bound

$$E\left[\frac{\sigma_1^2(X_i)}{p(X_i)} + \frac{\sigma_0^2(X_i)}{1-p(X_i)} + (\beta(X_i) - \beta)^2\right]$$

is a function of $\sigma_1^2(X_i), \sigma_0^2(X_i), \beta(X_i)$, and $p(X_i)$.

$$\begin{aligned}
\sigma_1^2(X_i) &= \mathsf{Var}(Y_{1i}|X_i) \\
\sigma_0^2(X_i) &= \mathsf{Var}(Y_{0i}|X_i) \\
\beta(X_i) &= E[Y_{1i} - Y_{0i}|X_i].
\end{aligned}$$

- $\sigma_1^2(X_i), \sigma_0^2(X_i)$, and $\beta(X_i)$ cannot be manipulated by the researcher, and in observational studies, even the propensity score $p(X_i)$ is beyond the control of any econometrician.

- In an experimental framework, we can in principle manipulate the propensity score $p(X_i)$, and try to increase efficiency of the standard semiparametric estimator.

- A nonstandard question. The optimization problem is to minimize the asymptotic variance bound, which is usually taken to be the object that cannot be further minimized. Also, randomized experiments are usually believed to be a method of ascertaining that the treatment D_i is indeed independent of the potential outcomes (Y_{0i}, Y_{1i}).

Ideal Optimization Problem

- Suppose that $\sigma_1^2(X_i)$ and $\sigma_0^2(X_i)$ are known to the researcher.

- The researcher would then like to solve

$$\min_{p(X_i)} E\left[\frac{\sigma_1^2(X_i)}{p(X_i)} + \frac{\sigma_0^2(X_i)}{1-p(X_i)} + (\beta(X_i) - \beta)^2\right] \tag{4}$$

subject to the constraint

$$E[p(X_i)] = p$$

- Here, the p denotes the overall probability of treatment, which is limited by the researcher's budget.

- The solution will satisfy

$$-\frac{\sigma_1^2(x)}{p(x)^2} + \frac{\sigma_0^2(x)}{(1-p(x))^2} = \lambda \tag{5}$$

for all x in the support of X, where λ denotes the Lagrangian, which is set by the budget constraint.

Two Stage Study

- Previous minimization problem was an infeasible problem.

- We now consider a feasible problem

- In the preliminary stage (or pilot study), the researcher obtains nonparametric estimates of $\sigma_1^2(X_i)$ and $\sigma_0^2(X_i)$, as well as the ideal propensity score.

- In the second stage, the researcher implements the experiment with the estimated ideal propensity score , and computes a semiparametric estimator of the average treatment effects β.

Some Digression: Why use pure randomization in the first stage?

- In our asymptotic framework, it is not really necessary to use pure randomization in the first stage.

- We propose it for "robustness"

- Suppose that the second stage cannot be executed for some unexpected administrative reason.

- We want to have some justification for pure randomization

- Perhaps a minmax justification?

- Imagine a game between a researcher an the nature. The researcher moves after nature makes the first move. The nature tries to make the problem as difficult as possible by making the asymptotic variance bound as large as possible.

$$E\left[\frac{\sigma_1^2(X_i)}{p(X_i)}+\frac{\sigma_0^2(X_i)}{1-p(X_i)}+(\beta(X_i)-\beta)^2\right]$$

- To make it interesting, assume that it is a cheating for the nature to set $\sigma_1^2(x)=\infty$ and $\sigma_0^2(x)=\infty$. In other words, the rule of the game is such that the nature should set $0\le\sigma_1^2(x)\le M_1<\infty$ and $0\le\sigma_0^2(x)\le M_0<\infty$ for all x.

- Then the variance bound becomes

$$E\left[\frac{M_1}{p(X_i)}+\frac{M_0}{1-p(X_i)}+(\beta(X_i)-\beta)^2\right]$$

- The researcher now minimizes the worst variance bound by manipulating the propensity score. The solution will satisfy

$$-\frac{M_1}{p(x)^2} + \frac{M_0}{(1 - p(x))^2} = 0 \qquad \forall x$$

  In other words, the minmax propensity score is constant over x.

- When the rule of the game is such that $M_1 = M_0$, then we obtain an intuitive 50-50 rule.

Back to the two-stage study. Notation/ Assumption

- Assume that the budget for the experiment is such that only $n = n_1 + n_2$ observations can be taken over all.
- n_1 denotes the number of observations in the pilot study
- n_2 denotes the number of observations in the second stage study.
- Let $\kappa = \frac{n_1}{n}$ denote the proportion of the pilot study.
- We will adopt an asymptotic framework where n_1 and n_2 grow to infinity at the same rate.

- Budget constraint is such that the overall number of the treated is limited, i.e., overall probability of treatment (over the first and second stages combined) is fixed to be p.

- We will further assume that, in the first stage, we will use the strategy of using equal propensity score over the two groups of X. In other words, the first stage design is a classic randomized experiment design.

- Such constant probability of treatment π_1.

- Because the overall probability of treatment combining the first and second stages is fixed to be p, in the second stage, we will be forced to choose the overall probability of treatment π_2 to be such that

$$\pi_1 \kappa + \pi_2 (1 - \kappa) = p$$

- For now, we will assume that κ and π_1 are given, which obviously implies that π_2 should be treated given as well.

- The only thing under the researcher's control is the propensity score $\pi_2(x)$ to be used in the second stage.

- Let $\pi(x) = \kappa\pi_1 + (1-\kappa)\pi_2(x)$ denote the overall propensity score in the combined sample, including both the pilot study and the second stage study.

- We would like to choose $\pi_2(x)$ such that the overall propensity score $\pi(x)$ is ideal, i.e., satisfies the first order condition (5).

- Because we do not know $\sigma_1^2(x)$ and $\sigma_0^2(x)$, we will instead use the estimates based on the pilot study. We will call them $\widehat{\sigma}_1^2(x)$ and $\widehat{\sigma}_0^2(x)$.

- We will denote the estimated ideal overall propensity score $\widehat{\pi}(x)$.

- The corresponding second stage propensity score will be likewise denoted $\widehat{\pi}_2(x)$.

Some Notes

- Because $\widehat{\sigma}_1^2(x)$ and $\widehat{\sigma}_0^2(x)$ are solely based on the pilot study, and because the observations in the second study are independent of those in the pilot study, the $\widehat{\sigma}_1^2(x)$ and $\widehat{\sigma}_0^2(x)$ are independent of the second sample.

- The same comment applies to $\widehat{\pi}(x)$ and $\widehat{\pi}_2(x)$.

- $\widehat{\sigma}_1^2(x)$ and $\widehat{\sigma}_0^2(x)$ are $\sqrt{n_1}$-consistent in general because X is assumed to be multinomial.

- As a consequence, $\widehat{\pi}(x)$ and $\widehat{\pi}_2(x)$ are also $\sqrt{n_1}$-consistent.

- Denote the probability limit of $\widehat{\pi}_2(x)$ by $\pi_2^*(x)$.

- Although we will in the end advocate the practice of setting $\pi_2^*(x)$ equal to the ideal propensity score that satisfies (5), our asymptotic theory is about a more general case where $\pi_2^*(x)$ is arbitrary.

- Assume that the researcher uses an efficient estimator

$$\widehat{\beta} = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{D_i Y_i}{\widehat{p}\left(X_i\right)} - \frac{\left(1 - D_i\right) Y_i}{1 - \widehat{p}\left(X_i\right)} \right)$$

- Recall that $\widehat{p}\left(X_i\right)$ denotes the estimated propensity score, whereas $\widehat{\pi}\left(x\right)$ denotes the propensity score that is chosen by (hence known to) the researcher.

- The only reason why the propensity score $\widehat{\pi}\left(x\right)$ has the "hat" notation is because it is based on estimated quantities $\widehat{\sigma}_1^2\left(x\right)$ and $\widehat{\sigma}_0^2\left(x\right)$.)

Some Technicality in Asymptotics

$$\begin{aligned}\widehat{\beta} - \beta &= \frac{1}{n}\sum_{i=1}^{n}\left(r_1(X_i) - r_0(X_i) - \beta\right) \\ &+ \frac{1}{n}\sum_{i=1}^{n}\left(\frac{D_i(Y_i - r_1(X_i))}{\widehat{\pi}(X_i)} - \frac{(1 - D_i)(Y_i - r_0(X_i))}{1 - \widehat{\pi}(X_i)}\right) \\ &+ \text{Garbage}\end{aligned}$$

The "Garbage" is the sum of

$$\sum_{x}\left(r_{1}(x)-\widehat{r}_{1}(x)\right)\left(\frac{\widehat{p}(x)-\widehat{\pi}(x)}{\widehat{\pi}(x)}\right)\left(\frac{1}{n}\sum_{i=1}^{n}1\left(X_{i}=x\right)\right) \quad (6)$$

and

$$-\sum_{x}\left(r_{0}(x)-\widehat{r}_{0}(x)\right)\left(\frac{\widehat{p}(x)-\widehat{\pi}(x)}{1-\widehat{\pi}(x)}\right)\left(\frac{1}{n}\sum_{i=1}^{n}1\left(X_{i}=x\right)\right) \quad (7)$$

- We would like to prove that (6) and (7) are $o_p\left(\frac{1}{\sqrt{n}}\right)$ by establishing $\widehat{p}(x)-\widehat{\pi}(x)$, $r_1(x)-\widehat{r}_1(x)$, and $r_0(x)-\widehat{r}_0(x)$ are all $O_p\left(n^{-1/2}\right)$.

- This is not trivial, because the second sample has a different DGP from the first sample, and the DGP of the second sample is conditional on the first sample.

Example of Some Details: $\widehat{p}(x)$

Recall that

$$\begin{aligned}\widehat{p}(x) &= \frac{\sum_{i=1}^{n} D_i \mathbf{1}(X_i = x)}{\sum_{i=1}^{n} \mathbf{1}(X_i = x)} \\ &= \frac{\sum_{i=1}^{n_1} D_i \mathbf{1}(X_i = x)}{\sum_{i=1}^{n} \mathbf{1}(X_i = x)} + \frac{\sum_{i=n_1+1}^{n} D_i \mathbf{1}(X_i = x)}{\sum_{i=1}^{n} \mathbf{1}(X_i = x)}\end{aligned}$$

The first piece is easy to analyze. By the law of large numbers and the CLT, we expect to have

$$\frac{1}{n_1} \sum_{i=1}^{n_1} D_i \mathbf{1}(X_i = x) = \pi_1 \Pr(X_i = x) + O_p\left(\frac{1}{\sqrt{n}}\right)$$

The second piece is not so easy because of dependence of $\widehat{\pi}_2(x)$ on the first sample.

- Our trick is to write $D_i = 1(U_i \leq \pi_1)$ for the first sample, and $D_i = 1(U_i \leq \widehat{\pi}_2(X_i))$ for the second sample

- Here the U_i denotes the iid uniform $(0, 1)$ random variable independent of Ys and Xs.

- It is trivial to show that the Ds generated in such a way will have the desired probability of being equal to one.

- We tackle the second piece by defining the empirical process

$$\xi_2(\cdot, \pi_2) = \frac{1}{\sqrt{n_2}} \sum_{i=n_1+1}^{n} \left(1(U_i \le \pi_2(x))\, 1(X_i = x) - E\left[1(U_i \le \pi_2(x))\, 1(X_i = x)\right]\right)$$

- By stochastic equicontinuity, we have

$$\xi_2(\cdot, \widehat{\pi}_2) = \xi_2(\cdot, \pi_2^*) + o_p(1)$$

where π_2^* denotes the probability limit of $\widehat{\pi}_2$ as discussed before.

- Using stochastic equicontinuity, we obtain that

$$\frac{1}{n_2} \sum_{i=n_1+1}^{n} 1(U_i \le \widehat{\pi}_2(x))\, 1(X_i = x) = \pi_2^*(x) \Pr(X_i = x) + O_p\left(\frac{1}{\sqrt{n}}\right)$$

Likewise we obtain

$$\begin{aligned} r_1(x) - \widehat{r}_1(x) &= O_p\left(n^{-1/2}\right) \\ r_0(x) - \widehat{r}_0(x) &= O_p\left(n^{-1/2}\right). \end{aligned}$$

which implies that

$$\begin{aligned} \sqrt{n}\left(\widehat{\beta} - \beta\right) \approx & \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left(r_1(X_i) - r_0(X_i) - \beta\right) \\ & + \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left(\frac{D_i\left(Y_i - r_1(X_i)\right)}{\widehat{\pi}(X_i)} - \frac{(1 - D_i)\left(Y_i - r_0(X_i)\right)}{1 - \widehat{\pi}(X_i)}\right) \end{aligned}$$

Compare it with

$$\sqrt{n}\left(\widehat{\beta}-\beta\right) \approx \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left(r_1\left(X_i\right)-r_0\left(X_i\right)-\beta\right) + \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left(\frac{D_i\left(Y_i-r_1\left(X_i\right)\right)}{p\left(X_i\right)}-\frac{\left(1-D_i\right)\left(Y_i-r_0\left(X_i\right)\right)}{1-p\left(X_i\right)}\right)$$

in observational studies context.

Using similar tricks, we obtain

$$
\begin{aligned}
& \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \frac{D_i \left(Y_i - r_1 \left(X_i\right)\right)}{\widehat{\pi} \left(X_i\right)} \\
= \quad & \frac{\sqrt{n_1}}{\sqrt{n}} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} \sum_{x} \frac{1 \left(U_i \leq \pi_1\right) \left(Y_{1i} - r_1 (x)\right)}{\pi^* (x)} 1 \left(X_i = x\right) \\
& + \frac{\sqrt{n_2}}{\sqrt{n}} \frac{1}{\sqrt{n_2}} \sum_{i=n_1+1}^{n} \sum_{x} \frac{1 \left(U_i \leq \pi_2^* (x)\right) \left(Y_{1i} - r_1 (x)\right)}{\pi^* (x)} 1 \left(X_i = x\right)
\end{aligned}
$$

and

$$
\begin{aligned}
&\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\frac{(1-D_i)\left(Y_i-r_0\left(X_i\right)\right)}{1-\widehat{\pi}\left(X_i\right)}\\
=\;&\frac{\sqrt{n_1}}{\sqrt{n}}\frac{1}{\sqrt{n_1}}\sum_{i=1}^{n_1}\sum_{x}\frac{1\left(U_i>\pi_1\right)\left(Y_{0i}-r_0\left(x\right)\right)}{1-\pi^*\left(x\right)}1\left(X_i=x\right)\\
&+\frac{\sqrt{n_2}}{\sqrt{n}}\frac{1}{\sqrt{n_2}}\sum_{i=n_1+1}^{n}\sum_{x}\frac{1\left(U_i>\pi_2^*\left(x\right)\right)\left(Y_{0i}-r_0\left(x\right)\right)}{1-\pi^*\left(x\right)}1\left(X_i=x\right)
\end{aligned}
$$

The asymptotic distribution of $\sqrt{n}\left(\widehat{\beta}-\beta\right)$ is normal with mean zero and variance

$$E\left[\frac{\kappa\pi_1+(1-\kappa)\,\pi_2^*(X_i)}{\pi^*(X_i)^2}\sigma_1^2(X_i)\right]$$
$$+E\left[\frac{1-\left(\kappa\pi_1+(1-\kappa)\,\pi_2^*(X_i)\right)}{(1-\pi^*(X_i))^2}\sigma_0^2(X_i)\right]$$
$$+E\left[(\beta(X_i)-\beta)^2\right]$$

Because

$$\kappa \pi_1 + (1 - \kappa)\, \pi_2^* (X_i) = \pi^* (X_i)$$

by construction, the asymptotic distribution of $\sqrt{n}\left(\widehat{\beta} - \beta\right)$ is

$$N\left(0, E\left[\frac{1}{\pi^* (X_i)} \sigma_1^2 (X_i) + \frac{1}{1 - \pi^* (X_i)} \sigma_0^2 (X_i) + (\beta (X_i) - \beta)^2\right]\right)$$

Some Notes

- The error in estimation of $\widehat{\pi}_2(X_i)$ is asymptotically negligible.

- As long as we set $\pi^*(X_i)$ equal to the ideal propensity score, ...

Outline

- Background/ Motivation/ Objective
- Summary of Our Proposal
- Counterfactual Framework in Observational Studies
- Details of Our Procedure
- Numerical Examples, Extensions $\Leftarrow$

Numerical Example 1: "Fake" Data

Covariate: $X = 0, 1$ with $1/2$ probability.

Assume $p = \pi_1 = 2/3$, $\kappa = 1/2$

<u>Estimates from Stage 1:</u>

$$\begin{aligned}\hat{\sigma}_0^2(0) &= 9 \\ \hat{\sigma}_0^2(1) &= 120\end{aligned}$$

$$\begin{aligned}\hat{\sigma}_1^2(0) &= 20 \\ \hat{\sigma}_1^2(1) &= 130\end{aligned}$$

Overall optimal assignment probabilities:

$$\hat{p}(0) = 0.80$$
$$\hat{p}(1) = 0.56$$

2nd Stage optimal assignment probabilities:

$$\hat{\pi}_2(0) = 0.93$$
$$\hat{\pi}_2(1) = 0.45$$

Numerical Example 2: Gerber, Karlan, Bergan (2006), "Does The Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions"

- Studied impact of exposure to media on political views.
- Randomly assigned some individual to receive Washington Post (left-leaning), other the Washington Times (right-leaning).
- Covariate: Gender, other baseline characteristics.
- Outcome: measure of political leaning.

- We find that optimal propensity score is very close to $p(x) = 0.5$.

- Turns out that all conditional variances are about the same $\ddot\smile$

Extension to LATE

- Perhaps the treatment cannot be randomized, because of non-compliance.

- If the intention to treat can be randomized, and if monotonicity as well as exclusion restriction is satisfied, then we may want to use the IV estimator and estimate the LATE.

- It can be shown that we can design the two-stage experiments in a similar manner, and increase efficiency.

Potential Technical Extension

- Case where the covariate vector X_i has a continuous distribution but the conditional variance functions $\sigma_1^2(X_i)$ and $\sigma_0^2(X_i)$ are parametrically specified.
  - Such parametric restriction allows one to assume that the ideal propensity score is $\sqrt{n}$-consistently estimated. As a result, we may continue to use the empirical process theory and the stochastic equicontinuity argument.

- Case where X_i is allowed to have a continuous distribution and $\sigma_1^2(X_i)$ and $\sigma_0^2(X_i)$ are nonparametrically specified.

- Under this scenario, the usual empirical process theory may not be immediately applicable. It is because $\widehat{\pi}(X_i)$ has a nonparametric specification as a consequence, which makes it difficult to deal with objects such as the indicator function $1(U_i \leq \widehat{\pi}_2(x))$ that contains a nonparametric estimator as an argument.

- A different asymptotics (n_1-fixed asymptotics?) may be considered as an alternative.