Introductory Econometrics

Lecture 26: Treatment Effect Model

Instructor: Ma, Jun

Renmin University of China

December 16, 2021

Introduction

- In this class, we consider the problem of estimating the causal effect of a binary explanatory variable, which is referred as the treatment effect in the literature. The treatment effect model is different from the linear regression model.
- In econometrics, the treatment effect model is very often used for evaluating social program/experiment.
- Example 1: Suppose that a selected set of individuals receive training or education initiated by the government with a view to enhancing their employment prospects. Suppose that the government has collected the earnings data for the individuals who received the training and for the individuals who did not. The main purpose of methods of program evaluations is to quantify and estimate the effect of the training program.

- Example 2: Suppose that an education program required high schools to agree to assign teachers and students to small (13 to 17 students) or large (22 to 26 students) classes. The government is interested in the effect of class size on student achievement.
- Such a question can arise in various other situations. A medical experiment studies on the effects of new treatment ask similar questions. One group of patients has received new treatment, and the other group has not.

Potential outcome variables

- Y_i: outcome variable; $D_i \in \{0, 1\}$: the binary explanatory variable; $X_{i1}, ..., X_{ik}$: other observed explanatory variables; ϵ_i: unobserved explanatory factors.
- The variable D_i is a binary variable taking 1 if the individual has gone through the treatment and 0 otherwise. The treatment here represents the actual treatment. The econometrician usually observes the treatment status for each individual D_i.
- $(X_{i1}, ..., X_{ik})$ represents a vector of various demographic characteristics for individual i. E.g., the variables can be annual income, age, gender, status of marriage, the number of children, education, etc. These represent all the observable characteristics of individual i.
- Suppose that Y_i is generated by $Y_i = g(D_i, X_{i1}, ..., X_{ik}, \epsilon_i)$.
- g is unknown and in the treatment effect model, we do not assume g is linear.

- The outcome variable $Y_i(1) = g(1, X_{i1}, ..., X_{ik}, \epsilon_i)$ represents a potential outcome of an individual i in the treatment state (e.g. training is received or studying in a reduced-size class). The variable $Y_i(0) = g(0, X_{i1}, ..., X_{ik}, \epsilon_i)$ represents a potential outcome of the same individual i in the control state (e.g. training is received or studying in a normal-size class).
- Thus, each individual has a random vector $(Y_i(1), Y_i(0))$ that represents potential outcomes depending on the state (treatment or control). Certainly, $(Y_i(1), Y_i(0))$ are correlated.
- The econometrician cannot observe the random vector $(Y_i(1), Y_i(0))$ jointly, because for each individual, only one potential outcome ($Y_i(1)$ or $Y_i(0)$) is realized, depending on whether the individual i has gone through the treatment or not.

The relationship between D_i and $(Y_i(1), Y_i(0))$

- In a medical experiment, the individual is chosen to be in the treatment group through some randomization device or a lottery. In these cases, $D_i \perp\!\!\!\perp (Y_i(1), Y_i(0))$ (i.e., D_i is independent of $(Y_i(1), Y_i(0))$).
- For evaluating social experiment/program with observational data, it may not be convincing to assume $D_i \perp\!\!\!\perp (Y_i(1), Y_i(0))$.

Treatment effects

- The individual treatment effect (ITE) for each individual i is defined as:

$$Y_i(1) - Y_i(0).$$

- The ITE is the difference between the potential outcomes in two different states for the same person.
- The ITE is a counterfactual quantity, in the sense that in the actual world, we cannot observe the vector $(Y_i(1), Y_i(0))$.
- There are mainly two quantities of interest: ATE (average treatment effect)

$$\text{ATE} = \mathrm{E}\left[Y_i(1) - Y_i(0)\right],$$

and ATT (average treatment effect on the treated)

$$\text{ATT} = \mathrm{E}\left[Y_i(1) - Y_i(0) \mid D_i = 1\right].$$

- The average treatment effect on the treated is the treatment effect of the people who have gone through the treatment.

- Note that the expectation in the definition of ATE involves the joint distribution of $(Y_i(1), Y_i(0))$, and the expectation in the definition of ATT involves the joint distribution of $(Y_i(1), Y_i(0), D_i)$, which are both unobserved.
- ATE or ATT can not be estimated accurately merely by collecting a large size of samples.

The observed information

- The econometrician observes the treatment status D_i and covariates X_i. She also observes the outcome variable:

$$Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0).$$

- The observed outcome variable Y_i is not the same as the potential outcomes $Y_i(1)$ or $Y_i(0)$. It is a realized outcome for an individual i depending on whether she has received treatment (Y_i is realized to be $Y_i(1)$) or not (Y_i is realized to be $Y_i(0)$).
- Identification of these parameters is concerned with the following question: can we uniquely determine the value of these parameters once we know the joint distribution of the observable random variables?

Randomized experiments

- In medical experiments, the treatment is performed using a randomization device. More specifically, for patient i, a lottery is run, and the patient is selected into the treated group with the design probability p, and stays in the control group with the design probability $1-p$.
- In these cases, we have $D_i \perp\!\!\!\perp (Y_i(1), Y_i(0), X_{i1}, ..., X_{ik})$. Randomized experiment assumption requires that knowing whether patient i is treated or not gives one no informational advantage in predicting the potential outcomes of i over another who does not know whether patient i is treated or not.
- This assumption is still possibly violated in medical studies if only those patients who have higher potential treatment effect are selected into treatment among all the patients in the study on purpose.
- In this case, observing D_i will give information about the treatment effect $(Y_i(1) - Y_i(0))$ for individual i.

- We use the following result from probability theory: if $V \perp\!\!\!\perp W$, then for any function f,

$$\mathrm{E}\left[f\left(V,W\right) \mid W = w\right] = \mathrm{E}\left[f\left(V,w\right)\right]. \tag{1}$$

- By (1) and the randomized experiment assumption, $D_i \perp\!\!\!\perp \left(Y_i\left(1\right), Y_i\left(0\right)\right)$, we have

$$\begin{aligned}
\text{ATE} &= \mathrm{E}\left[Y_i\left(1\right) - Y_i\left(0\right)\right] \\
&\quad \mathrm{E}\left[Y_i\left(1\right)\right] - \mathrm{E}\left[Y_i\left(0\right)\right] \\
&= \mathrm{E}\left[D_i Y_i\left(1\right) + \left(1 - D_i\right) Y_i\left(0\right) \mid D_i = 1\right] \\
&\quad -\mathrm{E}\left[D_i Y_i\left(1\right) + \left(1 - D_i\right) Y_i\left(0\right) \mid D_i = 0\right] \\
&= \mathrm{E}\left[Y_i \mid D_i = 1\right] - \mathrm{E}\left[Y_i \mid D_i = 0\right].
\end{aligned}$$

- By LIE,

$$\begin{aligned}
\mathrm{E}\left[Y_i D_i\right] &= \mathrm{E}\left[\mathrm{E}\left[Y_i D_i \mid D_i\right]\right] \\
&= \Pr\left[D_i = 1\right] \mathrm{E}\left[Y_i D_i \mid D_i = 1\right] \\
&\quad + \Pr\left[D_i = 0\right] \mathrm{E}\left[Y_i D_i \mid D_i = 0\right] \\
&= \mathrm{E}\left[D_i\right] \mathrm{E}\left[Y_i \mid D_i = 1\right],
\end{aligned}$$

  where

$$\begin{aligned}
\mathrm{E}\left[Y_i D_i \mid D_i = 0\right] &= \mathrm{E}\left[\left(D_i Y_i\left(1\right) + \left(1 - D_i\right) Y_i\left(0\right)\right) D_i \mid D_i = 0\right] \\
&= 0
\end{aligned}$$

  follows from (1).

- Similarly, we have

$$\mathrm{E}\left[Y_i \mid D_i = 0\right] = \frac{\mathrm{E}\left[Y_i\left(1 - D_i\right)\right]}{\mathrm{E}\left[1 - D_i\right]}.$$

- We can write

$$\text{ATE} = \frac{\text{E}\left[Y_i D_i\right]}{\text{E}\left[D_i\right]} - \frac{\text{E}\left[Y_i \left(1 - D_i\right)\right]}{\text{E}\left[1 - D_i\right]},$$

  where the right hand side depends on the joint distribution of the observed random variables.
- For estimation, we replace the population mean by the sample mean (this is sometimes called the analogue principle):

$$\widehat{\text{ATE}} = \frac{\frac{1}{n}\sum_{i=1}^{n} Y_i D_i}{\frac{1}{n}\sum_{i=1}^{n} D_i} - \frac{\frac{1}{n}\sum_{i=1}^{n} Y_i \left(1 - D_i\right)}{\frac{1}{n}\sum_{i=1}^{n} \left(1 - D_i\right)}.$$

- We can check its consistency by using LLN and Slutsky's lemma.
- This randomization assumption is not convincing when the individuals in the social experiments are people who may select into the treatment or not.

Comparison with the linear regression

- It seems that D_i is nothing but a dummy variable. Can we run a regression of Y_i on D_i and $X_{i1}, ..., X_{ik}$ to estimate the ATE? Can the parameter of interest, the ATE, be formulated as a coefficient in a regression model.
- One possible assumption is that

$$Y_i = g\left(D_i, X_{i1}, ..., X_{ik}, \epsilon_i\right) = \gamma_0 + \gamma_1 D_i + \sum_{j=1}^{k} \beta_j X_{ij} + \epsilon_i .$$

  In this case, the ITE $Y_i(1) - Y_i(0) = \gamma_1$ is constant. This is very unrealistic. We investigate alternative model assumptions.
- We first consider the following model assumption

$$\begin{aligned} Y_i(0) &= \mu_0 + U_i(0) \\ Y_i(1) &= \mu_1 + U_i(1), \end{aligned}$$

  where μ_0 and μ_1 are constants common across individuals and assumed to be nonstochastic and $(U_i(0), U_i(1))$ are stochastic components.

- We denote $\boldsymbol{X}_i = (X_{i1}, ..., X_{ik})^\top$ for the vector of observed covariates.
- We assume $\mathrm{E}[U_i(0) \mid \boldsymbol{X}_i] = \mathrm{E}[U_i(1) \mid \boldsymbol{X}_i]$, which implies

$$\mathrm{E}[Y_i(1) - Y_i(0) \mid \boldsymbol{X}_i] = \mu_1 - \mu_0,$$

  i.e., the ITE is mean independent of $\boldsymbol{X}_i$ but it can be random. And by LIE,

$$\mathrm{ATE} = \mathrm{E}[Y_i(1) - Y_i(0)] = \mu_1 - \mu_0.$$

- We assume $\mathrm{E}[Y_i(1) \mid D_i, \boldsymbol{X}_i] = \mathrm{E}[Y_i(1) \mid \boldsymbol{X}_i]$ and $\mathrm{E}[Y_i(0) \mid D_i, \boldsymbol{X}_i] = \mathrm{E}[Y_i(0) \mid \boldsymbol{X}_i]$, i.e., the conditional mean independence of potential outcomes with treatment status, conditional on demographic status $\boldsymbol{X}_i$.
- When we focus on a sub-population of indivdiuals with specific demographic status $\boldsymbol{X}_i$, $Y_i(1)$ and $Y_i(0)$ are both mean independent of D_i.

- Let us write

$$\begin{aligned}\mathrm{E}\left[Y_i \mid D_i, \boldsymbol{X}_i\right] &= D_i \mathrm{E}\left[Y_i(1) \mid D_i, \boldsymbol{X}_i\right] + (1 - D_i)\,\mathrm{E}\left[Y_i(0) \mid D_i, \boldsymbol{X}_i\right] \\ &= D_i \mathrm{E}\left[Y_i(1) - Y_i(0) \mid D_i, \boldsymbol{X}_i\right] + \mathrm{E}\left[Y_i(0) \mid D_i, \boldsymbol{X}_i\right] \\ &= D_i \mathrm{E}\left[Y_i(1) - Y_i(0) \mid \boldsymbol{X}_i\right] + \mathrm{E}\left[Y_i(0) \mid \boldsymbol{X}_i\right],\end{aligned}$$

  where the last equality follows from the conditional mean independence assumption.
- By the assumption $\mathrm{E}\left[U_i(0) \mid \boldsymbol{X}_i\right] = \mathrm{E}\left[U_i(1) \mid \boldsymbol{X}_i\right]$, we have

$$\begin{aligned}&D_i \mathrm{E}\left[Y_i(1) - Y_i(0) \mid \boldsymbol{X}_i\right] + \mathrm{E}\left[Y_i(0) \mid \boldsymbol{X}_i\right] \\ &\quad = D_i\left(\mu_1 - \mu_0\right) + \mathrm{E}\left[Y_i(0) \mid \boldsymbol{X}_i\right] \\ &\quad = \mu_0 + D_i\left(\mu_1 - \mu_0\right) + g\left(X_{i1}, \ldots, X_{ik}\right),\end{aligned}$$

  where we denote $g\left(X_{i1}, \ldots, X_{ik}\right) = \mathrm{E}\left[U_i(0) \mid \boldsymbol{X}_i\right]$.

- Therefore, we have

$$\mathrm{E}\left[Y_i \mid D_i, \boldsymbol{X}_i\right] = \mu_0 + (\mu_1 - \mu_0) D_i + g\left(X_{i1}, \ldots, X_{ik}\right).$$

- Define

$$V_i = Y_i - \mathrm{E}\left[Y_i \mid D_i, \boldsymbol{X}_i\right]$$

and now we have the following regression model:

$$Y_i = \mu_0 + (\mu_1 - \mu_0) D_i + g\left(X_{i1}, \ldots, X_{ik}\right) + V_i.$$

- We have $\mathrm{E}\left[V_i \mid D_i, \boldsymbol{X}_i\right] = 0$ by definition.
- We assume g is linear in $X_{i1}, \ldots, X_{ik}$:

$$g\left(X_{i1}, \ldots, X_{ik}\right) = \sum_{j=1}^{k} \beta_j X_{ij},$$

and then

$$Y_i = \mu_0 + (\mu_1 - \mu_0) D_i + \sum_{j=1}^{k} \beta_j X_{ij} + V_i.$$

- A multiple linear regression of Y_i on D_i and $X_{i1}, \ldots, X_{ik}$ consistently estimates $\mathrm{ATE} = (\mu_1 - \mu_0)$.

- We assume $\mathrm{E}\left[U_i\left(0\right) \mid \boldsymbol{X}_i\right] = \mathrm{E}\left[U_i\left(1\right) \mid \boldsymbol{X}_i\right]$, which implies

$$\mathrm{E}\left[Y_i\left(1\right) - Y_i\left(0\right) \mid \boldsymbol{X}_i\right] = \mu_1 - \mu_0.$$

- This assumption implies that the conditional average treatment effect given $\boldsymbol{X}_i$ does not depend on $\boldsymbol{X}_i$, the characteristics of individual i.
- This assumption can be unrealistic. E.g., Average treatment of the class-size is the same between students from high-income family and students from low-income family.

Unconfoundedness assumption

- Unconfoundedness is the key assumption of the basic treatment effect model.
- Unconfoundedness assumption: $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i \mid \boldsymbol{X}_i$, i.e., $(Y_i(1), Y_i(0))$ and D_i are conditionally independent given $\boldsymbol{X}_i$.
- Unconfoundedness can be thought of as an assumption that the decision to take the treatment is purely random for individuals with similar values of the covariates.
- Suppose that we have three random vectors V, W and X, where (V, W) is a continuous random vector. Then we say V and W are conditionally independent given X, if for all possible values of v, w and x,

$$f_{(V,W)|X}(v, w \mid x) = f_{V|X}(v \mid x) f_{W|X}(w \mid x).$$

- Unconfoundedness is satisfied if (Y_i, D_i) are generated by the model

$$\begin{aligned} Y_i &= g\left(D_i, X_{i1}, ..., X_{ik}, \epsilon_i\right) \\ D_i &= m\left(X_{i1}, ..., X_{ik}, \eta_i\right) \end{aligned}$$

and $\epsilon_i \perp\!\!\!\perp \eta_i \mid X_{i1}, ..., X_{ik}$.

More on conditional independence

- When V and W are conditionally independent given X, one can easily see that for any function φ,

$$\mathrm{E}\left[\varphi\left(V\right) \mid W, X\right] = \mathrm{E}\left[\varphi\left(V\right) \mid X\right].$$

  I.e., once we observe X, knowledge of W does not give us any further advantage in predicting the value of $\varphi\left(V\right)$.

- We notice that

$$\begin{aligned} f_{(V,W)|X}\left(v, w \mid x\right) &= \frac{f_{(V,W,X)}\left(v, w, x\right)}{f_X\left(x\right)} \\ &= \frac{f_{(V,W,X)}\left(v, w, x\right)}{f_{(W,X)}\left(w, x\right)} \frac{f_{(W,X)}\left(w, x\right)}{f_X\left(x\right)} \\ &= f_{V|(W,X)}\left(v \mid w, x\right) f_{W|X}\left(w, x\right). \end{aligned}$$

- Therefore, we have $f_{V|X}(v \mid x) = f_{V|(W,X)}(v \mid w, x)$, if (V, W) are conditionally independent given X. Hence,

$$\begin{aligned} \mathrm{E}[\varphi(V) \mid W = w, X = x] &= \int \varphi(v) f_{V|(W,X)}(v \mid w, x) \, \mathrm{d}v \\ &= \int \varphi(v) f_{V|X}(v \mid x) \, \mathrm{d}v \\ &= \mathrm{E}[\varphi(V) \mid X = x]. \end{aligned}$$

- Therefore, the unconfoundedness assumption $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i \mid \boldsymbol{X}_i$ implies the conditional mean independence assumption:

$$\begin{aligned} \mathrm{E}[Y_i(1) \mid D_i, \boldsymbol{X}_i] &= \mathrm{E}[Y_i(1) \mid \boldsymbol{X}_i] \\ \mathrm{E}[Y_i(0) \mid D_i, \boldsymbol{X}_i] &= \mathrm{E}[Y_i(0) \mid \boldsymbol{X}_i]. \end{aligned}$$

- We can also show: if $V \perp\!\!\!\perp W \mid X$,

$$\mathrm{E}[\eta(V, W) \mid X, W = w] = \mathrm{E}[\eta(V, w) \mid X]. \tag{2}$$

The unconfoundedness and randomization assumptions

- It can be shown that the randomization assumption $(Y_i(1), Y_i(0), \boldsymbol{X}_i) \perp\!\!\!\perp D_i$ implies the unconfoundedness assumption $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i \mid \boldsymbol{X}_i$.
- The randomized experiment assumption does not allow $X_{i1}, \ldots, X_{ik}$ to be correlated with D_i,
- The unconfounded condition allows D_i to be affected by $X_{i1}, \ldots, X_{ik}$, while the randomized experiment assumption does not.

Identification of ATE

- By LIE, we have

$$\begin{aligned} \text{ATE} &= \mathrm{E}\left[Y_i(1) - Y_i(0)\right] \\ &= \mathrm{E}\left[\mathrm{E}\left[Y_i(1) \mid \boldsymbol{X}_i\right]\right] - \mathrm{E}\left[\mathrm{E}\left[Y_i(0) \mid \boldsymbol{X}_i\right]\right], \end{aligned} \tag{3}$$

and

$$\begin{aligned} \mathrm{E}\left[Y_i D_i \mid \boldsymbol{X}_i\right] &= \mathrm{E}\left[\mathrm{E}\left[Y_i D_i \mid \boldsymbol{X}_i, D_i\right] \mid \boldsymbol{X}_i\right] \\ &= \Pr\left[D_i = 1 \mid \boldsymbol{X}_i\right] \mathrm{E}\left[Y_i D_i \mid \boldsymbol{X}_i, D_i = 1\right] \\ &\quad + \Pr\left[D_i = 0 \mid \boldsymbol{X}_i\right] \mathrm{E}\left[Y_i D_i \mid \boldsymbol{X}_i, D_i = 0\right]. \end{aligned}$$

- By the unconfoundedness assumption: $(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i \mid \boldsymbol{X}_i$, the result (2) and the relation $Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$, we have

$$\begin{aligned} &\mathrm{E}[Y_i D_i \mid \boldsymbol{X}_i, D_i = 1] \\ =&\mathrm{E}[(D_i Y_i(1) + (1 - D_i) Y_i(0)) D_i \mid \boldsymbol{X}_i, D_i = 1] = \mathrm{E}[Y_i(1) \mid \boldsymbol{X}_i] \end{aligned}$$

and

$$\mathrm{E}[Y_i D_i \mid \boldsymbol{X}_i, D_i = 0] = 0.$$

- Therefore, we have

$$\mathrm{E}[Y_i D_i \mid \boldsymbol{X}_i] = \Pr[D_i = 1 \mid \boldsymbol{X}_i] \mathrm{E}[Y_i(1) \mid \boldsymbol{X}_i] \tag{4}$$

and similarly,

$$\mathrm{E}[Y_i(1 - D_i) \mid \boldsymbol{X}_i] = \Pr[D_i = 0 \mid \boldsymbol{X}_i] \mathrm{E}[Y_i(0) \mid \boldsymbol{X}_i]. \tag{5}$$

- Now (3), (4), (5) and LIE imply

$$\begin{aligned} \text{ATE} &= \mathrm{E}\left[\frac{\mathrm{E}\left[Y_i D_i \mid \boldsymbol{X}_i\right]}{\Pr\left[D_i = 1 \mid \boldsymbol{X}_i\right]}\right] - \mathrm{E}\left[\frac{\mathrm{E}\left[Y_i \left(1 - D_i\right) \mid \boldsymbol{X}_i\right]}{\Pr\left[D_i = 0 \mid \boldsymbol{X}_i\right]}\right] \\ &= \mathrm{E}\left[\frac{Y_i D_i}{\Pr\left[D_i = 1 \mid \boldsymbol{X}_i\right]} - \frac{Y_i \left(1 - D_i\right)}{\Pr\left[D_i = 0 \mid \boldsymbol{X}_i\right]}\right]. \end{aligned}$$

  Now the right hand side depends only on the joint distribution of observed random variables.

- Denote

$$p\left(\boldsymbol{x}\right) = \Pr\left[D_i = 1 \mid \boldsymbol{X}_i = \boldsymbol{x}\right].$$

  This function is called propensity score. It is the probability of the event that the individual belongs to the treatment group, given that the observed characteristics are $\boldsymbol{x} \in \mathbb{R}^k$.

Estimation of ATE

- Let $\hat{p}(\boldsymbol{x})$be an estimator of the propensity score, then we can estimate the ATE:

$$\widehat{\text{ATE}} = \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{Y_i D_i}{\hat{p}(\boldsymbol{X}_i)} - \frac{Y_i(1-D_i)}{1-\hat{p}(\boldsymbol{X}_i)}\right\}.$$

- It is straightforward to construct $\hat{p}(\boldsymbol{x})$ if $\boldsymbol{X}_i$ is discrete:

$$\hat{p}(\boldsymbol{x}) = \frac{\sum_{i=1}^{n} 1(D_i = 1, \boldsymbol{X}_i = \boldsymbol{x})}{\sum_{i=1}^{n} 1(\boldsymbol{X}_i = \boldsymbol{x})}.$$

- If $\boldsymbol{X}_i$ is continuous, we specify a parametric model for the propensity score:

$$\Pr[D_i = 1 \mid \boldsymbol{X}_i] = \Phi(\beta_0 + \beta_1 X_{i1} + \cdots + \beta_k X_{ik})$$

  as what we did for the Probit model. This gives a parametric model for the propensity score. $(\beta_0, ..., \beta_k)$ can be estimated by MLE (denoted by $(\hat{\beta}_0, ..., \hat{\beta}_k)$).

- The estimated propensity score is

$$\hat{p}\left(\boldsymbol{X}_i\right) = \Phi\left(\hat{\beta}_0 + \hat{\beta}_1 X_{i1} + \cdots + \hat{\beta}_k X_{ik}\right).$$

- This estimator is known to be consistent and asymptotically normally distributed, if our propensity score model is correct.
- This approach has the drawback that if our model for the propensity score is wrong, the ATE estimator is inconsistent.
- Actually, $p\left(\boldsymbol{x}\right) = \mathrm{E}\left[D_i \mid \boldsymbol{X}_i = \boldsymbol{x}\right]$ can be estimated without specifying a parametric model for it.

k-NN estimator

- The k-nearest neighbor (k-NN) estimator is the simplest nonparametric estimator of $p(\boldsymbol{x})$.
- Fix $\boldsymbol{x}_0$ and suppose that we want to estimate $p(\boldsymbol{x}_0)$ at this point. Assume that p is a smooth function, which means that its graph does not change too much.
- $p(\boldsymbol{x})$ should be close to $p(\boldsymbol{x}_0)$ when $\boldsymbol{x}$ is close enough to $\boldsymbol{x}_0$. $p(\boldsymbol{X}_i)$ would be close to $p(\boldsymbol{x}_0)$ for observations $\boldsymbol{X}_i$ close to $\boldsymbol{x}_0$.
- We simply average these $p(\boldsymbol{X}_i)$ for observations $\boldsymbol{X}_i$ close to $\boldsymbol{x}_0$. We do not observe $p(\boldsymbol{X}_i)$ but use D_i instead.
- Let
$$d_i(\boldsymbol{x}_0) = \|\boldsymbol{X}_i - \boldsymbol{x}_0\| = \sqrt{(\boldsymbol{X}_i - \boldsymbol{x}_0)^\top (\boldsymbol{X}_i - \boldsymbol{x}_0)}$$
denote the distance of $\boldsymbol{X}_i$ to $\boldsymbol{x}_0$.
- After computing the distance for all n observations in the sample, we sort them in the increasing order
$$d_{(1)}(\boldsymbol{x}_0) \leq d_{(2)}(\boldsymbol{x}_0) \leq \cdots \leq d_{(n)}(\boldsymbol{x}_0).$$

- Let $N_k(\boldsymbol{x}_0)$ denote the identities of the k-nearest neighbors of $\boldsymbol{x}_0$:

$$N_k(\boldsymbol{x}_0) = \left\{ i : d_i(\boldsymbol{x}_0) \leq d_{(k)}(\boldsymbol{x}_0) \right\}.$$

- The k-NN nonparametric estimator of $p(\boldsymbol{x}_0)$ is

$$\hat{p}_{kNN}(\boldsymbol{x}_0) = \frac{1}{k} \sum_{i \in N_k(\boldsymbol{x}_0)} D_i.$$

- The k-NN estimator is simply an average of the values of D_i across the k closest observations in terms of $\boldsymbol{X}_i$.
- There is a data-driven procedure to select k in practical applications.
- The nonparametric ATE estimator using $\hat{p}_{kNN}(\boldsymbol{X}_i)$ is consistent and asymptotically normal. It does not require a parametric model for the propensity score.