

Simple Linear Regression

Part B: Ordinary Least Squares

THOMAS WIEMANN
University of Chicago

Econometrics
Econ 21020

Updated: May 9, 2022

Summary

In Part A, we introduced $\text{BLP}(Y|X)$ as approximation to $E[Y|X]$.

- ▷ Showed that the BLP-coefficients are well-defined when $\text{Var}(X) > 0$;
- ▷ Discussed interpretation using Yitzhaki's Theorem;

The BLP and its coefficients α and β are theoretical concepts.

In Part B, we bridge the gap between BLP and real data using statistics.

- ▷ Develop the *ordinary least squares* estimator;
- ▷ Analyze its statistical properties under an iid sample;
- ▷ Apply it to study the California STAR data.

1. Ordinary Least Squares
2. Estimator Properties
 - ▷ Bias
 - ▷ Consistency
 - ▷ Asymptotic Distribution
3. Case Study: The California STAR Data

1. **Ordinary Least Squares**
2. Estimator Properties
 - ▷ Bias
 - ▷ Consistency
 - ▷ Asymptotic Distribution
3. Case Study: The California STAR Data

Ordinary Least Squares

Let Y and X be two random variables. Throughout, we consider a random sample $(Y_1, X_1), \dots, (Y_n, X_n) \stackrel{iid}{\sim} (Y, X)$.

From Lecture 6A, we know that the BLP-coefficients are given by

$$\beta = \frac{\text{Cov}(Y, X)}{\text{Var}(X)}, \quad \text{and} \quad \alpha = E[Y] - E[X]\beta, \quad (1)$$

whenever $\text{Var}(X) > 0$.

This suggests sample analogue estimators

$$\hat{\beta}_n = \quad (2)$$

and

$$\hat{\alpha}_n = \quad (3)$$

Ordinary Least Squares (Contd.)

The estimators $(\hat{\alpha}_n, \hat{\beta}_n)$ are known as *ordinary least squares* (OLS) estimators. This is because they can also be motivated as solutions to the least-squares sample criterion:

$$(\hat{\alpha}_n, \hat{\beta}_n) = \arg \min_{\alpha, \beta \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n (Y_i - (\alpha + X_i \beta))^2, \quad (4)$$

whenever $\frac{1}{n} \sum_{i=1}^n X_i^2 - \left(\frac{1}{n} \sum_{i=1}^n X_i\right)^2 > 0$. In particular, we have:

Ordinary Least Squares (Contd.)

For our analysis, it's useful to rewrite $\hat{\beta}_n$ using $\varepsilon_i \equiv Y_i - \text{BLP}(Y_i|X_i)$:

$$\hat{\beta}_n =$$

(5)

1. Ordinary Least Squares
2. **Estimator Properties**
 - ▷ **Bias**
 - ▷ Consistency
 - ▷ Asymptotic Distribution
3. Case Study: The California STAR Data

Bias

Our analysis of the OLS estimator begins with its bias.

We assume here that X is continuous to avoid division by zero with positive probability (for $n \geq 2$).

- ▷ Flashbacks? Recall Problem 4d) of Problem Set 2 and the additional complications non-continuous X would introduce.

The bias of $\hat{\beta}_n$ when X is continuous is given by

$$\text{Bias}(\hat{\beta}_n) = E[\hat{\beta}_n] - \beta =$$

(6)

Hence, for continuous X , if $E[\varepsilon_i|X_i] = 0$, then $\text{Bias}(\hat{\beta}_n) = 0$.

- ▷ Does $E[\varepsilon_i|X_i] = 0$ hold generally? No: $E[\varepsilon_i X_i] = 0 \not\Rightarrow E[\varepsilon_i|X_i] = 0$.
- ▷ When do we know that $E[\varepsilon_i|X_i] = 0$? When $E[Y|X]$ is linear.
- ▷ When do we know that $E[Y|X]$ is linear? When X is binary.
- ▷ Binary X is not continuous, so need to worry about division by zero...

Many textbooks state that the OLS estimator $\hat{\beta}_n$ is unbiased for β .

- ▷ Importantly: Strong assumption are made along the way!
- ▷ We showed $\text{Bias}(\hat{\beta}_n) = 0$ if $E[Y|X]$ linear *and* X is continuous.
- ▷ With binary X , showed *conditional* unbiasedness in Problem Set 2.

Generally, little reason to believe $\text{Bias}(\hat{\beta}_n) = 0$ in economic applications:

- ▷ Economic theory rarely implies linear $E[Y|X]$ with continuous X .
- ▷ Horrible news? No: Most estimators are biased in practice...

1. Ordinary Least Squares
2. **Estimator Properties**
 - ▷ Bias
 - ▷ **Consistency**
 - ▷ Asymptotic Distribution
3. Case Study: The California STAR Data

Consistency

Theorem 1 ensures OLS satisfies the minimum requirement: Consistency.

Theorem 1

Let Y and X be random variables such that $\text{Var}(X) > 0$, and let (α, β) denote the BLP($Y|X$)-coefficients. If $(\hat{\alpha}_n, \hat{\beta}_n)$ are the OLS estimators constructed using $(Y_1, X_1), \dots, (Y_n, X_n) \stackrel{iid}{\sim} (Y, X)$, then

$$\begin{bmatrix} \hat{\alpha}_n \\ \hat{\beta}_n \end{bmatrix} \xrightarrow{P} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad (7)$$

Since the OLS estimators are continuous functions of moments of (Y, X) , we can prove this straightforwardly:

- ▷ Show $\hat{\beta}_n \xrightarrow{P} \beta$ using WLLN + CMT;
- ▷ Show $\hat{\alpha}_n \xrightarrow{P} \alpha$ using WLLN + CMT;
- ▷ Then use that marginal convergence in probability implies joint convergence in probability.

Consistency (Contd.)

Proof.

1. Ordinary Least Squares
2. **Estimator Properties**
 - ▷ Bias
 - ▷ Consistency
 - ▷ **Asymptotic Distribution**
3. Case Study: The California STAR Data

Asymptotic Distribution

OLS is asymptotically normal, which is useful for approximating its sampling variance in practice. Theorem 2 focuses on $\hat{\beta}_n$ for conciseness.

Theorem 2

Let Y and X be random variables such that $\text{Var}(X) > 0$, and let (α, β) denote the BLP($Y|X$)-coefficients. If $\hat{\beta}_n$ is OLS estimator for β constructed using $(Y_1, X_1), \dots, (Y_n, X_n) \stackrel{iid}{\sim} (Y, X)$, then

$$\sqrt{n} \left(\hat{\beta}_n - \beta \right) \xrightarrow{d} N \left(0, \frac{E[\varepsilon^2 (X - E[X])^2]}{\text{Var}(X)^2} \right), \quad (8)$$

where $\varepsilon \equiv Y - \text{BLP}(Y|X)$.

Note: Classical treatments of OLS would refer to the variance expression in Theorem 2 as the asymptotic variance under heteroskedasticity. By this, they mean the general case that $\text{Var}(\varepsilon|X)$ is allowed to vary with X . As ε is simply the BLP-residual, there is no obvious reason to think otherwise. We thus don't highlight this as a special case.

Asymptotic Distribution (Contd.)

Proof.

Asymptotic Distribution (Contd.)



Standard Error

Theorem 2 is of no practical use unless we can replace the expression for the asymptotic variance by a consistent estimator. Fortunately, we can.

Theorem 3

Let Y and X be random variables such that $\text{Var}(X) > 0$, and let (α, β) denote the BLP($Y|X$)-coefficients. If $\hat{\beta}_n$ is OLS estimator for β constructed using $(Y_1, X_1), \dots, (Y_n, X_n) \stackrel{iid}{\sim} (Y, X)$, then

$$\frac{(\hat{\beta}_n - \beta)}{\text{se}(\hat{\beta}_n)} \xrightarrow{d} N(0, 1), \quad (9)$$

where

$$\text{se}(\hat{\beta}_n) = \frac{\sqrt{\frac{1}{n} \sum_{i=1}^n \hat{\varepsilon}_i^2 (X_i - \frac{1}{n} \sum_{i=1}^n X_i)^2}}{\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i^2 - \left(\frac{1}{n} \sum_{i=1}^n X_i \right)^2 \right)}, \quad (10)$$

and $\hat{\varepsilon}_i = Y_i - \hat{\alpha}_n - X_i \hat{\beta}_n$.

Standard Error (Contd.)

We prove Theorem 3 for X such that $E[X] = 0$. The more general result holds, of course, but its proof requires even more painful algebra.

Proof.

OLS Estimation in R

This completes the theoretical analysis: Now we need to implement it!

OLS Estimation in R

```
# Compute OLS estimates
beta <- cov(y, x) / var(x)
alpha <- mean(y) - mean(x) * beta

# Compute BLP estimates
blp_yx <- alpha + x * beta

# Compute standard error for beta
epsilon <- y - blp_yx
se_numer <- sqrt(mean(epsilon^2 * (x - mean(x))^2))
se <- (se_numer / var(x)) / sqrt(n)
```

Note: There exists an OLS implementation in R – the `lm`-command. But importantly: Base-R does not implement the standard error of Theorem 3! So have some faith in your abilities and implement OLS yourself. See Problem 7 of Problem Set 3.

1. Ordinary Least Squares
2. Estimator Properties
 - ▷ Bias
 - ▷ Consistency
 - ▷ Asymptotic Distribution
3. **Case Study: The California STAR Data**

The California STAR Data

The California Standardized Testing and Reporting (STAR) data:

- ▷ All 420 K-6 and K-8 districts in California for 1998/99;
- ▷ Average reading and math scores on 5th grade standardized test;
- ▷ Average school characteristics: E.g., student-teacher ratios, income.

Suppose we are interested in assessing the relationship between students' math scores and student-teacher ratios. For this purpose, let (Y, W, U) be random variables, where $Y = g(W, U)$ and

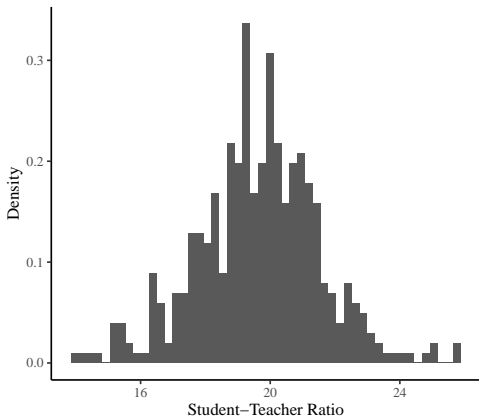
- ▷ $Y \equiv$ students' average math scores in a California district;
- ▷ $W \equiv$ the average student-teacher ratio of the district;
- ▷ $U \equiv$ all determinants of Y other than W .

We assume that the STAR data is the realization of the sample $(Y_1, W_1), \dots, (Y_{420}, W_{420}) \stackrel{iid}{\sim} (Y, W)$.

Note: The data is used as an example in Stock and Watson (2015) and readily available on their website. You can find the dataset on Canvas: `star98.csv`.

The California STAR Data (Contd.)

Figure 1: Histogram of the Student-Teacher Ratios in the STAR Data



Notes. You can find the code generating the figure on GitHub: [example_star98.R](#).

The California STAR Data (Contd.)

Since the student-teacher ratio is (nearly) continuous, need BLP:

- ▷ Conditional expectations can't be estimated directly;
- ▷ Choose to focus on best linear approximation instead for feasibility.

In particular, we consider estimating the BLP($Y|W$)-coefficient β .

- ▷ β captures the approximate expected change in students' math score associated with an additional student per teacher.

Our OLS estimate and the corresponding standard error are

$$\hat{\beta}_n \approx \quad , \text{ and } se(\hat{\beta}_n) \approx$$

Hence, a 95% confidence interval is given by

$$c_n \approx$$

Notes. You can replicate these estimates with the code on GitHub: [example_star98.R](#).

The California STAR Data (Contd.)

W/o additional assumptions, a correct interpretation of $\hat{\beta}_n$ is:

- ▷ An additional student per teacher is associated with an approximate expected change of students' math score by -1.939 points.

It's important to qualify your interpretation:

- ▷ “approximate” to highlight that $\text{BLP}(Y|W) \neq E[Y|W]$;
- ▷ “associated” to emphasize that β is a descriptive parameter.

The California STAR Data (Contd.)

Suppose that a peer who has not yet taken Econ 21020 sees your estimate $\hat{\beta}_n$ and concludes:

- ▷ “Just as I expected. Low student-teacher ratios are causing disastrous learning outcomes. We should hire more math teachers! ”

What's wrong with this interpretation?

- ▷ It's is a *causal* interpretation for a *descriptive* parameter: Not good!

From Lecture 6A, we know that under random assignment (RA), a causal interpretation of β would be appropriate. What does RA imply here?

- ▷ $W \perp U$, i.e., that the student-teacher ratio is independent of all other determinants of students' average math scores.

Do you think random assignment is a plausible assumption here?

- ▷ Probably not. E.g., wealthier districts may hire both 1) better teachers, and 2) more teachers per student.

The California STAR Data (Contd.)

To convince your peer that RA is likely to be implausible in the context of the California STAR data, you consider conducting a balance test.

Let the random variable X denote districts' average parental income.

- ▷ RA implies $W \perp U \Rightarrow W \perp X$;
- ▷ Hence, $W \perp X \Rightarrow \text{Cov}(W, X) = 0 \Rightarrow \beta_X = 0$.

Here: β_X is the BLP($X|W$)-coefficient.

A balance test may thus consider testing

$$H_0 : \beta_X = 0 \quad \text{versus} \quad H_1 : \beta_X \neq 0.$$

Rejecting H_0 would provide evidence that $W \not\perp X$.

- ▷ Potential cause for worry (but of course: Type I errors exist!).

Failure to reject H_0 would *not* provide evidence that $W \perp X$.

- ▷ May be because of $W \perp X$ or low power of the test!

The California STAR Data (Contd.)

Our OLS estimate of β_X and the corresponding standard error are

$$\hat{\beta}_{X,n} \approx \quad , \text{ and } se(\hat{\beta}_{X,n}) \approx$$

The test statistic for a two-sided hypothesis test is

$$T_n \approx$$

The corresponding p -value is

$$p\text{-value} \approx$$

Hence, on a 5% significance level, we reject H_0 .

- ▷ On a 5% significance level, there is sufficient evidence to reject that the approximate expected change in average parental income associated with a unit-change in the student-teacher ratio is zero.

As a consequence, we may deem RA implausible for the STAR context.

Summary

Today, we introduced OLS as an estimator for the BLP($Y|X$).

- ▷ Showed that it is consistent and asymptotically normal;
- ▷ Considered the California STAR data as an example.

We're now well-equipped for causal analysis under random assignment:

- ▷ Defined interesting causal parameters using the all causes model;
- ▷ Showed identification of the ATE;
- ▷ Concluded that if W is discrete, may use the binning estimator;
- ▷ If W is continuous, can leverage OLS to obtain approximate results.

But as we saw, random assignment is not always plausible or desired:

- ▷ Rarely applicable in non-experimental settings;
- ▷ Does not allow to study selection into treatment;

Next, we study a new identifying assumption: *Selection on Observables*.

References

Stock, J. H. and Watson, M. W. (2015). Introduction to econometrics.