Introduction to Fixed Effects

Why One Regression Line Isn't Enough

Jake Anderson

May 16, 2026

The Data

A coach tracks **training hours** vs. **performance** for their trainees.

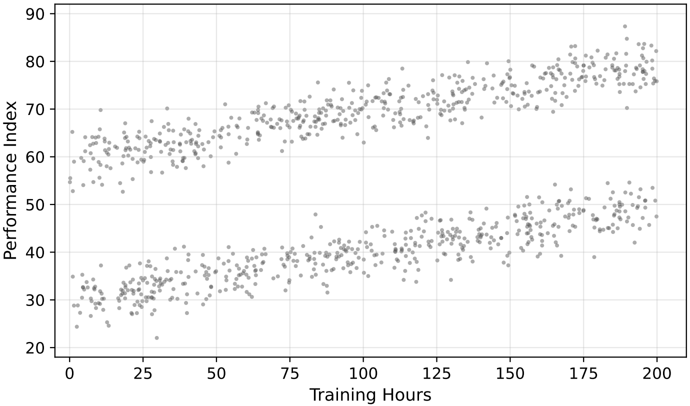

How could this data be generated?

The Setup

It turns out there are **two teams**: Varsity and Junior Varsity (JV). Same training program, different baseline ability.
Let's assume the following:

- JV players have a baseline of 30 "skill points"; Varsity have 60
- Each additional 10 hours of training $\rightarrow +1$ performance point (slope $= 0.1$)

If we ignore team membership and run a single regression:

$$\text{Performance}_i = \underbrace{\beta_0}_{\substack{\text{Assumes the} \\ \text{same intercept} \\ \text{for both teams!}}} + \underbrace{\beta_1}_{\substack{\text{Same slope} \\ \text{for both — OK}}} \text{Hours}_i + \varepsilon_i$$

But JV starts at 30 and Varsity starts at 60 — a single β_0 **cannot be right for both groups**.

Question: What goes wrong when we force one intercept on data with two?

What Happens If We Run Naive OLS?

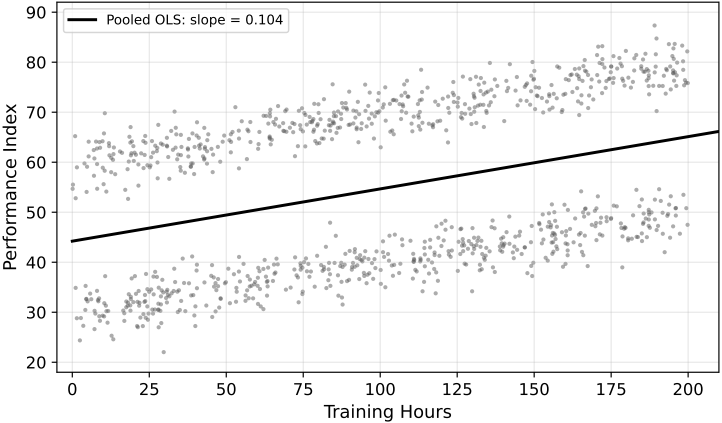

Pooled OLS: slope ≈ 0.10 — looks correct! But this only works because the sample is balanced (50/50) and x is identically distributed across groups. What if that changes?

First Check: Are Groups Sampled Differently?

In practice, the distribution of x often differs across groups.

- Varsity players may train **more** (selection into training)
- Or JV players may train more (catching up)

Key question: Does $\text{Cov}(\text{Group}, x) \neq 0$?

If so, pooled OLS conflates the group effect (α_j) with the treatment effect (β).

This is **omitted variable bias**.

Scenario: Varsity Trains More

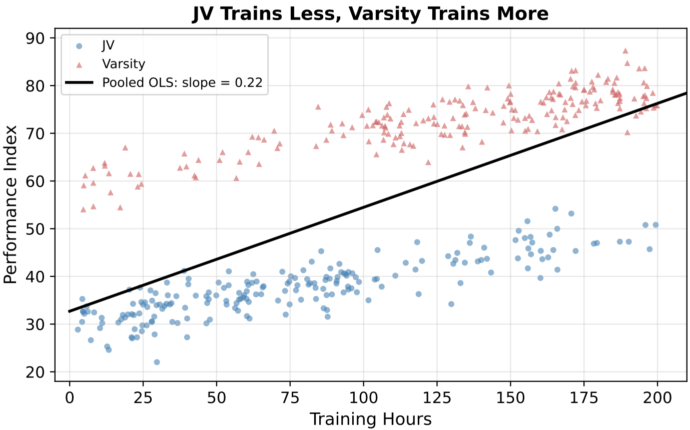

OLS slope = **0.22** (true = 0.10). Bias is **positive**: OLS attributes Varsity's higher baseline to their greater training hours.

Scenario: JV Trains More (Catching Up)

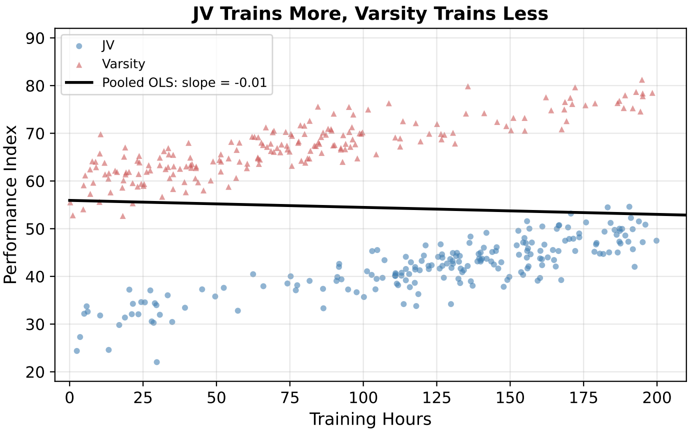

OLS slope $= -\mathbf{0.02}$ (true $= 0.10$). Bias is **negative**: OLS thinks training has nearly *zero effect* because the high-training group has lower baseline ability.

The Omitted Variable Bias

Let's map this to the OVB framework you already know. Let X_2 = Group / team membership (OV)

Short regression (what we run $\rightarrow$ omits X_2):

$$Y_i = \beta_0 + \beta_1 X_{1i} + u_i$$

Long regression (what we should run $\rightarrow$ includes X_2):

$$Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \varepsilon_i$$

Auxiliary regression (relationship between included and omitted):

$$X_{2i} = \delta_0 + \delta_1 X_{1i} + v_i$$

OVB formula: (effect of X_2 on Y) $\times$ (relationship of X_2 to X_1)

$$\hat{\beta}_1^{\text{short}} = \hat{\beta}_1^{\text{long}} + \hat{\beta}_2 \times \hat{\delta}_1$$

OVB Applied to Our Example

$$\hat{\beta}_1^{\text{short}} = \underbrace{\hat{\beta}_1^{\text{long}}}_{=0.10} + \underbrace{\hat{\beta}_2}_{\substack{\text{effect of group} \\ \text{on Performance}}} \times \underbrace{\hat{\delta}_1}_{\substack{\text{relationship of} \\ \text{group to Hours}}}$$

Varsity has higher baseline performance, so $\hat{\beta}_2 > 0$.

Scenario	$\hat{\delta}_1$	**Bias** $(\hat{\beta}_2 \times \hat{\delta}_1)$	**OLS slope**
Varsity trains more	> 0	$+$	0.22
JV trains more	< 0	$-$	-0.02
Equal training	≈ 0	≈ 0	0.10

Same data, same true effect. The OLS estimate swings from -0.02 to $+0.22$ just by changing which group trains more.

What About Class Imbalance?

Suppose x is distributed the same across groups ($\delta \approx 0$), but the **sample composition** is unbalanced.

- Does the slope change?
- What about the intercept?

The OLS intercept is a **weighted average** of the group intercepts:

$$\begin{aligned}
\hat{\beta}_0 &= \frac{\sum_{i=1}^{n} \nVdash\{i \in \text{JV}\}}{n} \cdot \beta_{0,\text{JV}} + \frac{\sum_{i=1}^{n} \nVdash\{i \in \text{Var}\}}{n} \cdot \beta_{0,\text{Var}} \\
&= \frac{n_{\text{JV}}}{n} \cdot \beta_{0,\text{JV}} + \frac{n_{\text{Var}}}{n} \cdot \beta_{0,\text{Var}} \\
&= (\text{Share JV}) \cdot \beta_{0,\text{JV}} + (\text{Share Var}) \cdot \beta_{0,\text{Var}}
\end{aligned}$$

80% JV, 20% Varsity

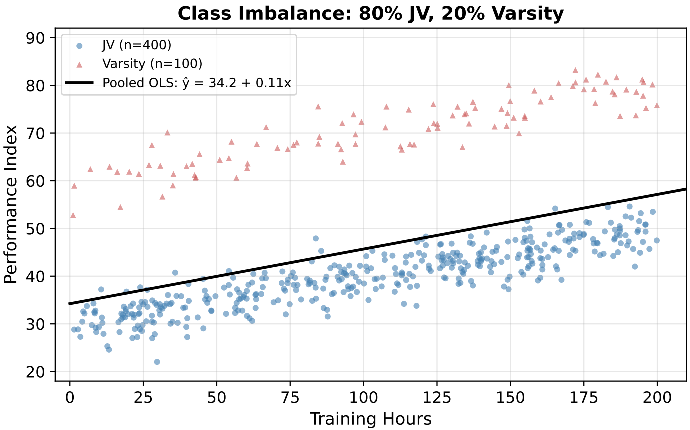

Slope $\approx$ correct. Intercept $=$ **34** — pulled toward JV's true intercept (30). Predictions are wrong for most Varsity players.

20% JV, 80% Varsity

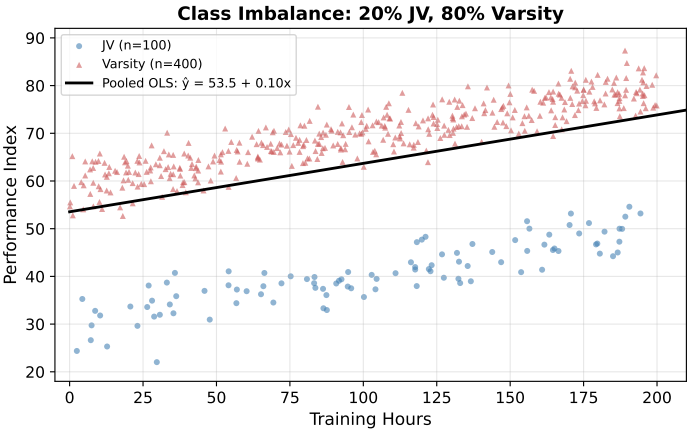

Slope ≈ correct. Intercept = **54** — pulled toward Varsity's true intercept (60). Now predictions are wrong for most JV players.

Class Imbalance: The Intercept Shifts

Sample	OLS intercept	True JV ($\alpha = 30$)	True Var ($\alpha = 60$)
80% JV	34	close	off by 26
50/50	44	off by 14	off by 16
80% Varsity	54	off by 24	close

$\Longrightarrow$ Even when the slope is approximately correct, pooled OLS uses a **single intercept** that is wrong for every subgroup. The error depends on sample composition, which the researcher may not control.

A Separate Intercept per Unit

Every "naive" example we just saw used the same model:

$$y_i = \beta_0 + \beta_1 x_i + \varepsilon_i$$

A single β_0 forces **one baseline** for all groups. That's where the bias comes from.

The fix: let each group j have its own intercept α_j, where $j = \text{JV}$ or $j = \text{Varsity}$:

$$y_i = \alpha_j + \beta_1 x_i + \varepsilon_i$$

The subscript j on the intercept is doing all the work:

- $\beta_0 \rightarrow$ one number, shared by everyone
- $\alpha_j \rightarrow$ a **different number for each group**

This is the core idea behind **fixed effects**.

The Fixed Effects Model

Allow each group its own intercept:

$$\text{Performance}_{ij} = \alpha_j + \beta\, \text{Hours}_{ij} + \varepsilon_{ij}$$

Equivalently, add a group dummy:

$$y_i = \beta\, x_i + \alpha_{\text{JV}} \cdot \nVdash\{\text{JV}\} + \alpha_{\text{Var}} \cdot \nVdash\{\text{Var}\} + \varepsilon_i$$

Key idea: FE estimates β using only *within-group* variation in x.
It asks: "Among JV players, do those who train more perform better?"

FE Estimation: The Result

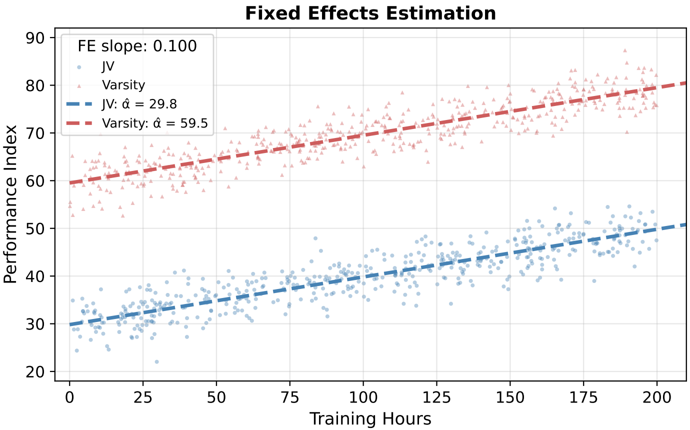

FE slope = **0.100** (true = 0.10). Intercepts: JV = **29.8** (true 30), Varsity = **59.5** (true 60).

Pooled OLS vs. Fixed Effects

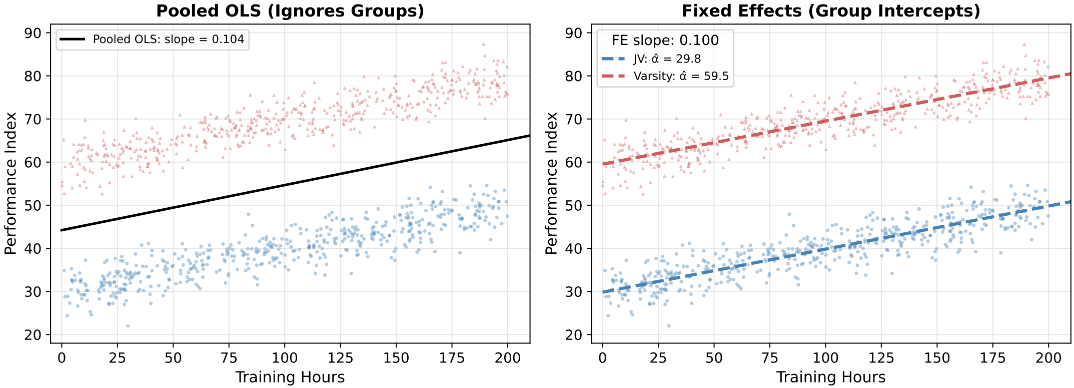

Why Does This Work? The Rank Condition

Basic idea: You need at least **two points to draw a line** (one intercept + one slope = 2 unknowns).

Pooled OLS: 2 unknowns (β_0, β_1) $\implies$ need ≥ 2 observations total. Easy.

Fixed effects with 2 groups: 3 unknowns (α_{JV}, α_{Var}, β_1).

- You need ≥ 2 observations *per group* (to pin down each group's line)
- So at minimum: **4 observations** (2 JV + 2 Varsity)

In general: with J groups, you have $J + 1$ unknowns (J intercepts + 1 slope), so you need at least 2 observations per group to be identified.

$\implies$ More groups = more unknowns = more data required.

Connection to Panel Data

Our training example maps directly to the panel data framework:

$$y_{it} = \alpha_i + \beta x_{it} + \varepsilon_{it}$$

- i = individual (trainee $\rightarrow$ firm, person, country)
- t = time period
- α_i = **individual fixed effect** (unobserved, time-invariant)

Two equivalent estimation approaches:

1. **Least Squares Dummy Variable** (when we want all of the individual fixed effects):

$$y_{it} = \beta x_{it} + \sum_i \alpha_i D_i + \varepsilon_{it}$$

2. **Within / Demeaning Estimator:** Subtract individual means

$$(y_{it} - \bar{y}_i) = \beta(x_{it} - \bar{x}_i) + (\varepsilon_{it} - \bar{\varepsilon}_i)$$

Both give the **same** $\hat{\beta}$. Let's walk through the within estimator step by step.

The Within Estimator: Step 1 $\rightarrow$ Compute Group Means

Start with the model:

$$y_{it} = \alpha_i + \beta\, x_{it} + \varepsilon_{it}$$

For each individual i, take the **time average** of both sides:

$$\bar{y}_i = \alpha_i + \beta\, \bar{x}_i + \bar{\varepsilon}_i$$

where:

- $\bar{y}_i = \frac{1}{T} \sum_{t=1}^{T} y_{it}$ (mean outcome for individual i)
- $\bar{x}_i = \frac{1}{T} \sum_{t=1}^{T} x_{it}$ (mean regressor for individual i)

Notice that α_i survives averaging because it **doesn't vary over time**.

The Within Estimator: Step 2 $\to$ Subtract

Subtract the individual mean equation from the original:

$$y_{it} - \bar{y}_i = (\alpha_i - \alpha_i) + \beta(x_{it} - \bar{x}_i) + (\varepsilon_{it} - \bar{\varepsilon}_i)$$

The fixed effect cancels:

$$\ddot{y}_{it} = \beta\,\ddot{x}_{it} + \ddot{\varepsilon}_{it}$$

where $\ddot{y}_{it} \equiv y_{it} - \bar{y}_i$ is the **demeaned** variable.

$\implies$ We have a simple regression with **no intercept** and no α_i.

Just run OLS on the demeaned data.

The Within Estimator: Step 3 $\rightarrow$ Estimate β

OLS on the demeaned regression gives:

$$\hat{\beta} = \frac{\text{Cov}(\ddot{x}_{it}, \ddot{y}_{it})}{\text{Var}(\ddot{x}_{it})} = \frac{\sum_i \sum_t (x_{it} - \bar{x}_i)(y_{it} - \bar{y}_i)}{\sum_i \sum_t (x_{it} - \bar{x}_i)^2}$$

$\Longrightarrow$ The covariance of the demeaned variables equals the covariance of the original variables:

$$\text{Cov}(\ddot{x}_{it}, \ddot{y}_{it}) = \text{Cov}(x_{it} - \bar{x}_i, y_{it} - \bar{y}_i) = \text{Cov}(x_{it}, y_{it})$$

Why? Let's show it briefly.

Proof: Demeaning Doesn't Change Covariance

Let $c_i = \bar{x}_i$ and $d_i = \bar{y}_i$ (constants within group i). Then:

$$\begin{aligned} \text{Cov}(x_{it} - c_i,\ y_{it} - d_i) = \text{Cov}(x_{it},\ y_{it}) - \text{Cov}(x_{it},\ d_i) \\ - \text{Cov}(c_i,\ y_{it}) + \text{Cov}(c_i,\ d_i) \end{aligned}$$

Within each group i, the means c_i and d_i are **constants**, so:

$$\text{Cov}(x_{it},\ d_i) = 0, \quad \text{Cov}(c_i,\ y_{it}) = 0, \quad \text{Cov}(c_i,\ d_i) = 0$$

Therefore:

$$\text{Cov}(x_{it} - \bar{x}_i,\ y_{it} - \bar{y}_i) = \text{Cov}(x_{it},\ y_{it}) \qquad \blacksquare$$

$\implies$ The within estimator uses **only within-group variation**. All between-group differences are "absorbed" by the fixed effects.

Thank you!

jakeanderson@g.ucla.edu