Introduction to Random Effects

From Fixed Parameters to Random Draws

Jake Anderson

February 27, 2026

The Data

A researcher tracks **days to recovery** vs. **hours of physical therapy per week** across patients at several hospitals.

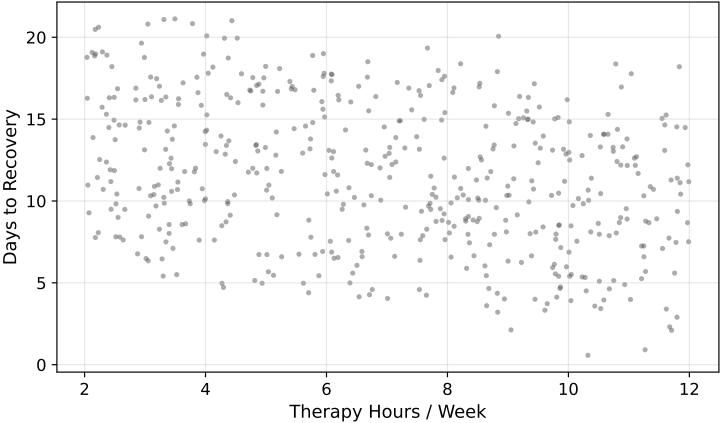

How could this data be generated?

Reveal: Five Hospitals

Patients come from **5 hospitals**, each with a different baseline recovery time.

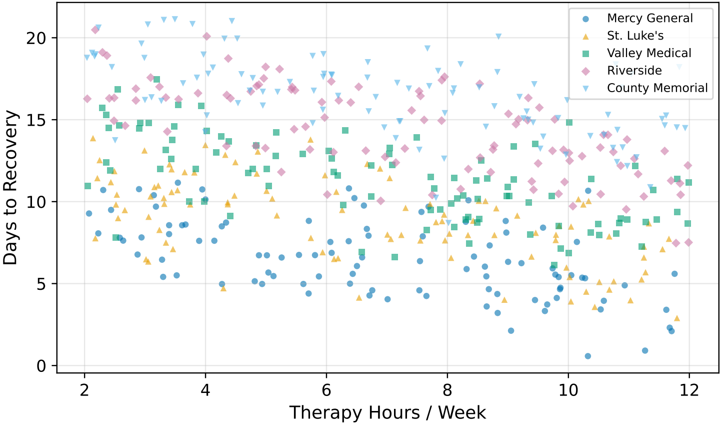

Recall: ignoring groups biases OLS. FE solved this by giving each group its own intercept.

FE Recap: Group-Specific Intercepts

The **fixed effects** model treats each hospital's baseline as a fixed unknown:

$$y_{ij} = \alpha_j + \beta\, x_{ij} + \varepsilon_{ij}$$

FE estimates a **separate intercept** for each hospital:

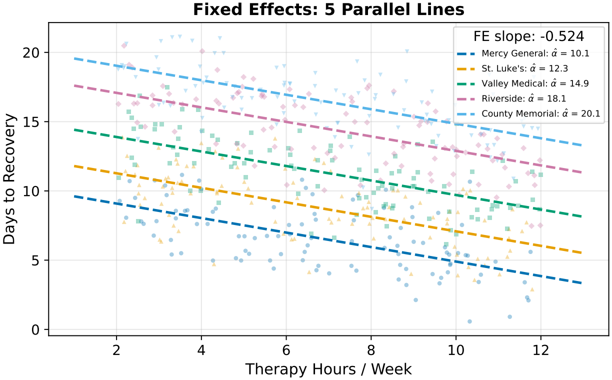

The Cost of FE

With 5 hospitals, FE estimates **5 separate intercepts**. That's manageable.

But what if you have:

- 50 hospitals? $\rightarrow$ 50 intercepts
- 500 hospitals? $\rightarrow$ 500 intercepts
- 5,000 hospitals? $\rightarrow$ 5,000 intercepts

Two problems with FE:

1. Uses up degrees of freedom (one parameter per group)
2. **Cannot** estimate the effect of time-invariant variables

For example: "Do *teaching* hospitals have faster recovery?" Teaching status doesn't change, so FE absorbs it into α_j.

$\Longrightarrow$ What if we could model the group effects with **fewer parameters**?

The RE Assumption

Fixed effects: Each α_j is a fixed, unknown parameter to estimate.

$$y_{ij} = \alpha_j + \beta\, x_{ij} + \varepsilon_{ij}$$

Random effects: Decompose α_j into a common mean plus a random deviation:

$$\alpha_j = \bar{\alpha} + u_j \qquad \text{where } u_j \sim (0, \sigma_u^2)$$

- $\bar{\alpha}$ = average baseline across *all* hospitals
- u_j = hospital j's **deviation** from that average (same role as α_j, but a random draw instead of a free parameter)
- σ_u^2 = how spread out hospital baselines are

Instead of estimating 5 (or 500) separate α_j's, we estimate $\bar{\alpha}$ and **one variance** σ_u^2.

The critical assumption:

$$\boxed{\mathrm{Cov}(u_j, x_{ij}) = 0}$$

i.e., patients at a **better** (worse) hospital can't tend to receive **more** (less) therapy.

The Error Components Model

Substitute $\alpha_j = \bar{\alpha} + u_j$ into the FE model:

$$y_{ij} = \bar{\alpha} + \beta\, x_{ij} + \underbrace{u_j + e_{ij}}_{v_{ij}}$$

where:

- u_j = hospital-specific random component (same for all patients at hospital j)
- e_{ij} = idiosyncratic error (patient-specific noise)
- $v_{ij} = u_j + e_{ij}$ = **composite error**

Problem: v_{ij} is **not** iid.

Two patients at the same hospital share u_j, so their errors are **correlated**. OLS ignores this.

The Correlation Structure

Within hospital j (patients $i \neq k$):

$$\text{Corr}(v_{ij},\, v_{kj}) = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_e^2} > 0$$

Across hospitals ($j \neq l$):

$$\text{Corr}(v_{ij},\, v_{kl}) = 0$$

What this means:

- Patients at the same hospital have positively correlated errors
- The correlation equals the share of total variance due to hospital effects
- OLS standard errors are **too small** because they count within-hospital observations as independent

$\Longrightarrow$ Even if the OLS slope is OK, the inference is wrong.

Why This Correlation Is Useful

Define the **intraclass correlation**:

$$\rho = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_e^2}$$

$\rho =$ the share of total variance explained by **which hospital** a patient is in.

- $\rho \approx 0$: hospitals are basically the same $\rightarrow$ grouping doesn't help much
- $\rho \approx 1$: almost all variation is between hospitals $\rightarrow$ knowing the hospital tells you almost everything

This is where RE gets its power. Instead of estimating each hospital in isolation, RE **borrows strength from the ensemble** (Tukey, 1970; Efron & Morris, 1973): it pulls each hospital's estimate toward the overall mean, especially when a hospital has few patients.

$\implies$ A hospital with only 5 patients gets a better estimate by "learning" from the other 495. The more hospitals look alike (small ρ), the more we can borrow.

GLS Intuition: Transforming the Data

OLS ignores the correlation in v_{ij}. GLS accounts for it by **transforming** the data.

Introduce the parameter $\hat{\alpha}_j$:

$$\hat{\alpha}_j = 1 - \frac{\sigma_e}{\sqrt{N_j \sigma_u^2 + \sigma_e^2}}$$

where $N_j =$ observations in group j (patients per hospital in our example).

The RE transformation is a partial demeaning:

$$y_{ij} - \hat{\alpha}_j \, \bar{y}_j \qquad \text{and} \qquad x_{ij} - \hat{\alpha}_j \, \bar{x}_j$$

Then run OLS on the transformed data. This is **feasible GLS** (FGLS).

What $\hat{\alpha}$ Does: A Spectrum

The RE estimator lives on a spectrum between pooled OLS and FE:

$\hat{\alpha}$	**Transformation**	**Equivalent to**
0	$y_{ij} - 0 \cdot \bar{y}_j = y_{ij}$	Pooled OLS (no group effect)
$0 < \hat{\alpha} < 1$	$y_{ij} - \hat{\alpha}\,\bar{y}_j$	**RE: weighted average**
1	$y_{ij} - \bar{y}_j$	FE (full demeaning)

When does $\hat{\alpha} \to 1$?

- σ_u^2 is large relative to σ_e^2 (strong group effects)
- N_j is large (many obs per group)

$\implies$ With strong group effects or large panels, RE $\approx$ FE.

The $\hat{\alpha}$ Spectrum: Visualized

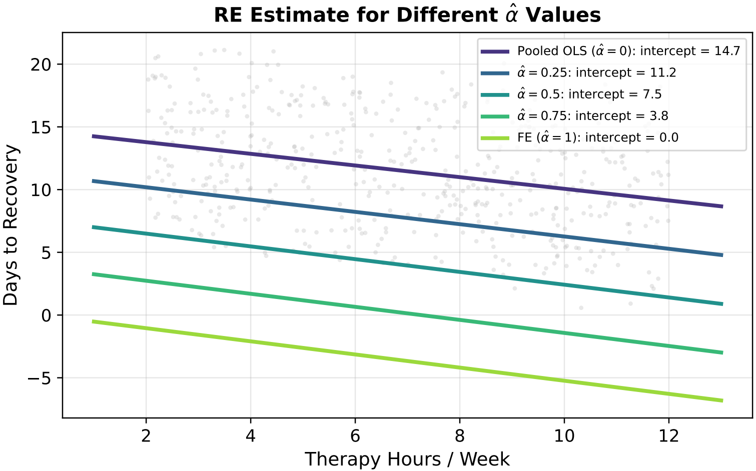

As $\hat{\alpha}$ increases from 0 to 1, we demean more aggressively, trusting each hospital's own data rather than pulling it toward the overall mean.

The RE Assumption Revisited

RE requires:

$$\mathrm{Cov}(u_j, x_{ij}) = 0$$

In our hospital example: hospital quality must be **uncorrelated** with therapy hours.

Is that realistic?

- Better hospitals might prescribe *more* therapy (better protocols)
- Or *less* therapy (patients recover faster anyway)
- Patient selection: sicker patients go to better hospitals

If $\mathrm{Cov}(u_j, x_{ij}) \neq 0$:

- FE is still consistent (it eliminates u_j entirely)
- RE is **inconsistent** (the partial demeaning doesn't fully remove u_j)

When RE Is Appropriate

Use RE when:

1. Groups are **random draws** from a larger population
   $\rightarrow$ Sampled hospitals from all hospitals in the country
2. No reason to think group effects correlate with regressors
   $\rightarrow$ Random assignment, natural experiment
3. You want to estimate effects of **time-invariant variables**
   $\rightarrow$ Teaching hospital status, rural vs. urban
4. **Efficiency**: RE uses both within and between variation
   $\rightarrow$ Smaller standard errors than FE

When FE Is Safer

Use FE when:

1. Groups are **specific entities** you care about
   - $\rightarrow$ These particular 5 hospitals, not a random sample
2. Plausible that $\text{Cov}(u_j, x_{ij}) \neq 0$
   - $\rightarrow$ Better hospitals may assign more/less therapy
3. Micro data (individuals, firms) $\implies$ FE is **almost always safer**
   - $\rightarrow$ Unobserved ability, management quality, etc.
4. You only care about **within-group** effects
   - $\rightarrow$ "Among patients at the same hospital, does more therapy help?"

$\implies$ When in doubt, FE is the conservative choice. But how do we decide formally?

Motivation: What Difference Does It Make?

If RE assumption holds: Both FE and RE are consistent, but RE is more efficient.

If RE assumption fails: Only FE is consistent. RE gives the wrong answer.

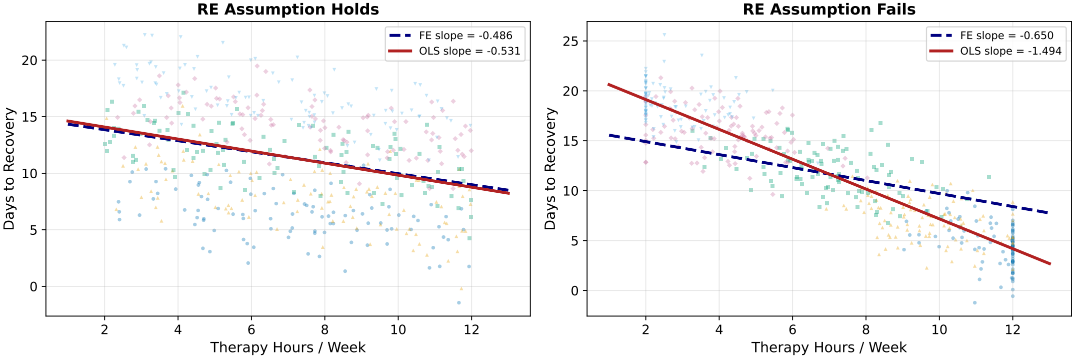

When the assumption holds, slopes are similar. When it fails, they **diverge**.

The Hausman Test: Setup

Idea: If RE is valid, then $\hat{\beta}_{FE}$ and $\hat{\beta}_{RE}$ should be close. If they diverge, something is wrong with RE.

Hypotheses:

- H_0: $\text{Cov}(u_j, x_{ij}) = 0 \iff$ RE is consistent (and more efficient)
- H_1: $\text{Cov}(u_j, x_{ij}) \neq 0 \iff$ only FE is consistent

Test statistic (single regressor):

$$t = \frac{\hat{\beta}_{FE} - \hat{\beta}_{RE}}{\sqrt{\widehat{\text{Var}}(\hat{\beta}_{FE}) - \widehat{\text{Var}}(\hat{\beta}_{RE})}}$$

With multiple regressors, the test generalizes to a χ^2_K statistic (where $K =$ number of regressors). Software handles this automatically.

The Hausman Test: Why Does It Work?

Under H_0: Both estimators converge to the true β, so the difference $\hat{\beta}_{FE} - \hat{\beta}_{RE}$ is small.

Under H_1: RE is biased, so the difference is systematically large.

Why is the denominator well-defined?

Under H_0, FE is less efficient than RE:

$$\mathrm{Var}(\hat{\beta}_{FE}) > \mathrm{Var}(\hat{\beta}_{RE})$$

so $\widehat{\mathrm{Var}}(\hat{\beta}_{FE}) - \widehat{\mathrm{Var}}(\hat{\beta}_{RE}) > 0$ and the square root exists.

Interpreting the Hausman Test

Decision rule:

Result	**Conclusion**	**Action**
$p < 0.05$	Reject H_0	Use **FE**
$p \geq 0.05$	Fail to reject H_0	Can use **RE**

Intuition:

- Reject $\rightarrow$ FE and RE give **different** answers $\rightarrow$ RE is biased $\rightarrow$ use FE
- Fail to reject $\rightarrow$ FE and RE give **similar** answers $\rightarrow$ RE is OK and more efficient

$\Longrightarrow$ The Hausman test is a **specification test**: it checks whether $\text{Cov}(u_j, x_{ij}) = 0$ is reasonable.

Worked Example: Hospital Data

Using our hospital recovery data:

	FE	**RE**
Slope ($\hat{\beta}$)	-0.500	-0.497
$\widehat{\text{Var}}(\hat{\beta})$	0.00180	0.00165

Hausman statistic:

$$t = \frac{-0.500 - (-0.497)}{\sqrt{0.00180 - 0.00165}} = \frac{-0.003}{\sqrt{0.00015}} = \frac{-0.003}{0.0122} = -0.25$$

$|t| = 0.25 < 1.96 \implies$ **Fail to reject** H_0.

$\implies$ RE and FE agree. RE is appropriate here, and more efficient.

Note: These are illustrative values from our simulated data. In practice, use software output.

FE vs. RE: Decision Flowchart

Step 1: Is there unobserved group heterogeneity?

- No $\rightarrow$ Pooled OLS is fine | Yes $\rightarrow$ Go to Step 2

Step 2: Do you need to estimate effects of time-invariant variables?

- Yes $\rightarrow$ Must use RE (FE absorbs them) | No $\rightarrow$ Go to Step 3

Step 3: Is $\text{Cov}(u_j, x_{ij}) = 0$ plausible?

- Clearly no $\rightarrow$ Use FE | Maybe $\rightarrow$ Run the Hausman test

Step 4: Hausman test result?

- Reject H_0 $\rightarrow$ Use FE | Fail to reject $\rightarrow$ Use RE (more efficient)

$\Longrightarrow$ When in doubt, FE is the sa**fe** default. RE is the **re**ward for being able to argue $\text{Cov}(u_j, x_{ij}) = 0$.

Thank you!

jakeanderson@g.ucla.edu