Gov 2002: 8. Panel Data

Matthew Blackwell

October 22, 2015

Repeated measurements

- Up until now, we have assumed that there was either a completely randomized experiment or a randomized experiment within levels of X_i that gave us exogeneous variation in the treatment.
- Today we're going to look to another possible source of variation: repeated measurements on the same unit over time.
- What if selection on the observables doesn't hold, but do have repeated measurements. Can we use this to identify and estimate effects?
- Message: simply having panel data does not identify an effect, but it does allow us to rely on different identifying assumptions.

Basic Idea

- The basic idea is that ignorability doesn't hold, conditional on the observed covariates, $Y_{it}(d) \not\perp\!\!\!\perp D_{it}|X_{it}$, but ignorability might hold conditional on some unobserved, time-constant, variable:

$$Y_{it}(d) \perp\!\!\!\perp D_{it}|X_{it}, U_i.$$

- Within units, effects are identified.
- This is because, even if U_i is unobserved, it is held constant within a unit.
- Thus, by performing analyses within the units, we can control for this unobserved heterogeneity.

Motivation

Is Democracy Good for the Poor?

Michael Ross University of California, Los Angeles

- Relationship between democracy and infant mortality?
- Compare levels of democracy with levels of infant mortality, but...
- Democratic countries are different from non-democracies in ways that we can't measure?
  - they are richer or developed earlier
  - provide benefits more efficiently
  - posses some cultural trait correlated with better health outcomes
- If we have data on countries over time, can we make any progress in spite of these problems?

Ross data

```r
ross <- foreign::read.dta("ross-democracy.dta")
head(ross[, c("cty_name", "year", "democracy", "infmort_unicef")])
```

```
##      cty_name year democracy infmort_unicef
## 1 Afghanistan 1965         0            230
## 2 Afghanistan 1966         0             NA
## 3 Afghanistan 1967         0             NA
## 4 Afghanistan 1968         0             NA
## 5 Afghanistan 1969         0             NA
## 6 Afghanistan 1970         0            215
```

Pooled OLS with Ross data

```
pooled.mod <- lm(log(kidmort_unicef) ~ democracy + log(GDPcur),
                 data = ross)
summary(pooled.mod)
```

```
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)   9.7640     0.3449    28.3   <2e-16 ***
## democracy    -0.9552     0.0698   -13.7   <2e-16 ***
## log(GDPcur)  -0.2283     0.0155   -14.8   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.8 on 646 degrees of freedom
##   (5773 observations deleted due to missingness)
## Multiple R-squared:  0.504,  Adjusted R-squared:  0.503
## F-statistic:  329 on 2 and 646 DF,  p-value: <2e-16
```

Note about terminology

- Generally, we talk about *panel data* and *time-series cross-sectional data* in political science.
- **Panel data**: small T, large N
  - The NES panel is like this: 2000 respondent asked questions at various points in time over the course of an election (or multiple elections).
- **TSCS data**: high T, low medium N.
  - U.S. states over time
  - Western European countries over time.
- For the most part, the issues of causality are the same for these two types of data, so I will refer to them both as panel data.
- But estimation is a different issue. Different estimators work differently under either data types.

1/ Fixed effects estimators

Notation

- Units $i = 1, \ldots, N$
- Time periods $t = 1, \ldots, T$ with $T \geq 2$,
- Y_{it}, D_{it} are the outcome and treatment for unit i in period t
  We have a set of covariates in each period, as well,
- Covariates X_{it}, causally "prior" to D_{it}.

$$
\begin{tikzcd}
 & X_t \arrow[dl] \arrow[dr] & \\
D_t \arrow[rr] & & Y_t
\end{tikzcd}
$$

- $U_i =$ unobserved, time-invariant unit effects (causally prior to everything)
- History of some variable: $\underline{D}_{it} = (D_1, \ldots, D_t)$.
- Entire history: $\underline{D}_i = \underline{D}_{iT}$

Assumptions

- **Potential outcomes**: $Y_{it}(1) = Y_{it}(d_t = 1)$ is the potential outcome for unit i at time t if they were treated at time t.
    - Here we focus on contemporaneous effects, $Y_{it}(d_t = 1) - Y_{it}(d_t = 0)$
    - Harder when including lags of treatment, $Y_{it}(d_t = 1, d_{t-1} = 1)$
- **Consistency** for each time period:

$$Y_{it} = Y_{it}(1)D_{it} + Y_{it}(0)(1 - D_{it})$$

- **Strict ignorability**: potential outcomes are independent of the entire history of treatment conditional on the history of covariates and the time-constant heterogeneity:

$$Y_{it}(d) \perp\!\!\!\perp \underline{D}_i | \underline{X}_i, U_i$$

Basic linear fixed-effects model

- Assume that the CEF for the mean potential outcome under control is:

$$\mathbb{E}[Y_{it}(0)|\underline{X}_i, U_i] = X_{it}'\beta + U_i$$

- And then assume a constant treatment effects:

$$\mathbb{E}[Y_{it}(1)|\underline{X}_i, U_i] = \mathbb{E}[Y_{it}(0)|\underline{X}_i, U_i] + \tau$$

- With consistency and strict ignorability, we can write this as a CEF of the observed outcome:

$$\mathbb{E}[Y_{it}|\underline{X}_i, \underline{D}_i, U_i] = X_{it}'\beta + \tau D_{it} + U_i$$

Relating to traditional models

- We can now write the observed outcomes in a traditional regression format:

$$Y_{it} = X_{it}'\beta + \tau D_{it} + U_i + \varepsilon_{it}$$

- Here, the error is similar to what we had for regression:

$$\varepsilon_{it} \equiv Y_{it}(0) - \mathbb{E}[Y_{it}(0)|\underline{X}_i, U_i]$$

- In traditional FE models, we skip potential outcomes and rely on a **strict exogeneity** assumption:

$$\mathbb{E}[\varepsilon_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

Strict ignorability vs strict exogeneity

$$Y_{it}(d) \perp\!\!\!\perp \underline{D}_i | \underline{X}_i, U_i$$

- Easy to show to that strict ignorability implies strict exogeneity:

$$\begin{aligned}
\mathbb{E}[\varepsilon_{it} | \underline{X}_i, \underline{D}_i, U_i] &= \mathbb{E}\left[\left(Y_{it}(0) - \mathbb{E}[Y_{it}(0) | \underline{X}_i, U_i]\right) | \underline{X}_i, \underline{D}_i, U_i\right] \\
&= \mathbb{E}[Y_{it}(0) | \underline{X}_i, \underline{D}_i, U_i] - \mathbb{E}[Y_{it}(0) | \underline{X}_i, U_i] \\
&= \mathbb{E}[Y_{it}(0) | \underline{X}_i, U_i] - \mathbb{E}[Y_{it}(0) | \underline{X}_{iT}, U_i] \\
&= 0
\end{aligned}$$

Fixed-effects within estimator

- Define the "within" model:

$$(Y_{it} - \overline{Y}_i) = (X_{it} - \overline{X}_i)'\beta + \tau(D_{it} - \overline{D}_i) + (\varepsilon_{it} - \overline{\varepsilon}_i)$$

- Here, let $\overline{Y}_i$ be the unit averages. Note that:

$$\overline{Y}_i = \overline{X}_i'\beta + \tau\overline{D}_i + U_i + \overline{\varepsilon}_i$$

- Logic: since the unobserved effect is constant over time, subtracting off the mean also subtracts that unobserved effect:

$$U_i - \frac{1}{T}\sum_{t=1}^{T} U_i = U_i - U_i = 0$$

- This also demonstrates why the assumption of the fixed effects being time-constant is so important.

Within Estimator

- Let $\ddot{Z}_{it} = Z_{it} - \overline{Z}_i$ be the **time-demeaned** version of Z_{it}. Then the FE model is:

$$\ddot{Y}_{it} = \ddot{X}_{it}'\beta + \tau\ddot{D}_{it} + \ddot{\varepsilon}_{it}$$

- **Within/FE estimator**, $\widehat{\tau}_{FE}$:

  pooled OLS estimator $\ddot{Y}_{it}$ *on* $\ddot{X}_{it}$ *and* $\ddot{D}_{it}$

- Only uses time variation within each cross section.
- Full rank: $\text{rank}[\sum_{t=1}^{T} \mathbb{E}[\ddot{X}_{it}\ddot{X}_{it}']] = K$
  - Basically: no variables that are constant over time. Why?

Fixed effects with Ross data

```
library(plm)
fe.mod <- plm(log(kidmort_unicef) ~ democracy + log(GDPcur), data = ross,
    index = c("id", "year"), model = "within")
summary(fe.mod)
```

```
## Oneway (individual) effect Within Model
##
## Call:
## plm(formula = log(kidmort_unicef) ~ democracy + log(GDPcur),
##     data = ross, model = "within", index = c("id", "year"))
##
## Unbalanced Panel: n=166, T=1-7, N=649
##
## Residuals :
##     Min.  1st Qu.   Median  3rd Qu.     Max.
## -0.70500 -0.11700  0.00628  0.12200  0.75700
##
## Coefficients :
##             Estimate Std. Error t-value Pr(>|t|)
## democracy    -0.1432     0.0335   -4.28 0.000023 ***
## log(GDPcur)  -0.3752     0.0113  -33.12  < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Total Sum of Squares:    81.7
## Residual Sum of Squares: 23
## R-Squared      :  0.718
##       Adj. R-Squared :  0.532
## F-statistic: 613.481 on 2 and 481 DF, p-value: <2e-16
```

Time-constant variables

- Pooled model with a time-constant variable, proportion Islamic:

```r
library(lmtest)
p.mod <- plm(log(kidmort_unicef) ~ democracy + log(GDPcur) + islam, data = ross,
    index = c("id", "year"), model = "pooling")
coeftest(p.mod)
```

```
##
## t test of coefficients:
##
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) 10.30608    0.35952   28.67  < 2e-16 ***
## democracy   -0.80234    0.07767  -10.33  < 2e-16 ***
## log(GDPcur) -0.25497    0.01607  -15.87  < 2e-16 ***
## islam        0.00343    0.00091    3.77  0.00018 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Time-constant variables

- FE model, where the islam variable drops out, along with the intercept:

```
fe.mod2 <- plm(log(kidmort_unicef) ~ democracy + log(GDPcur) + islam, data = ross,
    index = c("id", "year"), model = "within")
coeftest(fe.mod2)
```

```
##
## t test of coefficients:
##
##             Estimate Std. Error t value Pr(>|t|)
## democracy    -0.1297     0.0359   -3.62  0.00033 ***
## log(GDPcur)  -0.3800     0.0118  -32.07  < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Fixed-effects within estimator

- Informal proof. We have strict exogeneity:

$$\mathbb{E}[\varepsilon_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

- This implies exogeneity of the time-averaged errors:

$$\mathbb{E}[\overline{\varepsilon}_i|\underline{X}_i, \underline{D}_i, U_i] = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\varepsilon_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

- Mean-differenced errors are uncorrelated with the treatment or regressors from *any* time period:

$$\mathbb{E}[\ddot{\varepsilon}_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

- Thus, the mean-differenced treatment and covariates must also be uncorrelated with the mean-differenced errors:

$$\mathbb{E}[\ddot{Y}_{it}|\underline{X}_i, \underline{D}_i, U_i] = \ddot{X}_{it}'\beta + \tau\ddot{D}_{it}$$

Dummy variable regression

- An alternative way to estimate FE models is using a series of dummy variables for each unit, i.
- Let $W_{it}^k = 1$ if $k = i$ and $W_{it}^k = 0$ otherwise for all $k \in 1, \ldots, N$.
- $W_{it} = (W_{it}^1, \ldots, W_{it}^N)$ is the dummy variable vector.
- **Least Squares Dummy Variable** (LSDV) estimator:

  pooled OLS regression Y_{it} on X_{it}, D_{it}, and W_{it}.

- Algebraically equivalent to the within estimator for estimates.

SE issues

- Let $\boldsymbol{\varepsilon}_i$ be the $(T \times 1)$ vector of errors of the FE model.
- Panel homoskedasticity:

$$\mathbb{E}[\boldsymbol{\varepsilon}_i \boldsymbol{\varepsilon}_i'] = \sigma^2_\varepsilon \mathbf{I}_T$$

- Here, $\mathbf{I}_T$ is a diagonal matrix with T rows and columns and so basically:
  - Homoskedasticity: $\mathbb{V}[\varepsilon_{it} | \underline{X}_i, \underline{D}_i, U_i] = \sigma^2_\varepsilon$
  - No serial correlation: $\text{Cov}[\varepsilon_{it}, \varepsilon_{is} | \underline{X}_i, \underline{D}_i, U_i] = 0$ when $t \neq s$
- $\rightsquigarrow$ FE via within/LSDV are **efficient** estimators.
- Robust/sandwich SEs available via the usual formulas.

Within vs LSDV

- Within estimator and LSDV give exactly the same estimates, but SEs will differ slightly.
- SEs from vanilla OLS on the **within estimator** will be slightly off due to incorrect degrees of freedom.
  - OLS doesn't account for you calculating the time-means.
  - Smart software (`plm()` in R, `areg` in Stata) will correct.
- LSDV estimator gets the correct SEs because time-means are calculated by OLS $\rightsquigarrow$ correct degrees of freedom.
  - Downside: can be computationally demanding

Example with Ross data

```r
library(lmtest)
lsdv.mod <- lm(log(kidmort_unicef) ~ democracy + log(GDPcur) + as.factor(id),
    data = ross)
coeftest(lsdv.mod)[1:6, ]
```

```
##                    Estimate Std. Error t value Pr(>|t|)
## (Intercept)           13.76      0.266    51.8 1.0e-198
## democracy             -0.14      0.033    -4.3  2.3e-05
## log(GDPcur)           -0.38      0.011   -33.1 3.5e-126
## as.factor(id)AGO       0.30      0.168     1.8  7.4e-02
## as.factor(id)ALB      -1.93      0.190   -10.2  4.4e-22
## as.factor(id)ARE      -1.88      0.170   -11.0  2.4e-25
```

```r
coeftest(fe.mod)[1:2, ]
```

```
##             Estimate Std. Error t value Pr(>|t|)
## democracy      -0.14      0.033    -4.3  2.3e-05
## log(GDPcur)    -0.38      0.011   -33.1 3.5e-126
```

First differences

- Because the U_i are time-fixed, first-differences are an alternative to mean-differences.
- For some variable, Z_{it}, let $\Delta Z_{it} = Z_{it} - Z_{i,t-1}$
- The **first difference model** is the following:

$$\Delta Y_{it} = \Delta X_{it}' \beta + \tau \Delta D_{it} + \Delta \varepsilon_{it}$$

- This follows from the fact that $\Delta U_i = 0$
- By the same logic as above, strict ignorability implies strict exogeneity which implies $\mathbb{E}[\Delta \varepsilon_{it} | \underline{X}_i, \underline{D}_i, U_i = 0]$, so

$$\mathbb{E}[\Delta Y_{it} | \underline{X}_i, \underline{D}_i, U_i] = \Delta X_{it}' \beta + \tau \Delta D_{it}$$

First differences estimation

- First differences estimator: pooled OLS regression of ΔY_{it} on ΔX_{it} and ΔD_{it}.
- If $\Delta \varepsilon_{it}$ are homoskedastic and without serial correlation, usual OLS SEs work just fine.
- $\varepsilon_{it} = \varepsilon_{i,t-1} + \Delta \varepsilon_{it}$ implies original errors have serial correlation.
- $\rightsquigarrow$ more efficient than FE when there is serial correlation exists in the errors.
- Robust/sandwich SEs available here too.

First differences in R

```
fd.mod <- plm(log(kidmort_unicef) ~ democracy + log(GDPcur), data = ross,
    index = c("id", "year"), model = "fd")
summary(fd.mod)
```

```
## Oneway (individual) effect First-Difference Model
##
## Call:
## plm(formula = log(kidmort_unicef) ~ democracy + log(GDPcur),
##     data = ross, model = "fd", index = c("id", "year"))
##
## Unbalanced Panel: n=166, T=1-7, N=649
##
## Residuals :
##    Min. 1st Qu.  Median 3rd Qu.    Max.
## -0.9060 -0.0956  0.0468  0.1410  0.3950
##
## Coefficients :
##             Estimate Std. Error t-value Pr(>|t|)
## (intercept)  -0.1495     0.0113  -13.26   <2e-16 ***
## democracy    -0.0449     0.0242   -1.85    0.064 .
## log(GDPcur)  -0.1718     0.0138  -12.49   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Total Sum of Squares:    23.5
## Residual Sum of Squares: 17.8
## R-Squared      :  0.246
##       Adj. R-Squared :  0.244
## F-statistic: 78.1367 on 2 and 480 DF, p-value: <2e-16
```

2/ Random effects

Random effects

$$Y_{it} = X'_{it}\beta + \tau D_{it} + U_i + \varepsilon_{it}$$

- With fixed effects, we have:

$$\mathbb{E}[\varepsilon_{it} | \underline{X}_i, \underline{D}_i, U_i] = 0$$

- "Random effects" models make an additional assumption:

$$\mathbb{E}[U_i | \underline{X}_i, \underline{D}_i] = \mathbb{E}[U_i] = 0$$

- Unit-level effects are uncorrelated with treatment and covariates.
- **Important**: implies that ignorability holds without conditioning on $U_i \rightsquigarrow$ no unmeasured confounding.

Why random effects?

- So why do people use random effects? Standard errors!
- Under the RE assumption, we have the following:

$$Y_{it} = X_{it}'\beta + \tau D_{it} + \nu_i$$

  where $\nu_i = U_i + \varepsilon_{it}$.
- Now, notice that

$$\text{cov}[Y_{i1}, Y_{i2} | \underline{X}_{it}, \underline{D}_{it}] = \sigma_u^2$$

  where σ_u^2 is the variance of the U_i.
- This violates the assumption of no autocorrelation for OLS. What's the problem with this?
- Random effects models gets us consistent standard error estimates.

Quasi-demeaning

- Random effects models usually transform the data via what is called **quasi-demeaning** or **partial pooling**:

$$(Y_{it} - \theta\overline{Y}_i) = (X_{it} - \theta\overline{X}_i)'\beta + \tau(D_{it} - \theta\overline{D}_i) + (\nu_{it} - \theta\overline{\nu}_i)$$

- Here θ is between zero and one, where $\theta = 0$ implies pooled OLS and $\theta = 1$ implies fixed effects. Doing some math shows that
$$\theta = 1 - \left[\sigma_u^2/(\sigma_u^2 + T\sigma_\varepsilon^2)\right]^{1/2}$$
- the **random effect estimator** runs pooled OLS on this model replacing θ with an estimate $\hat{\theta}$.

Example with Ross data

```
re.mod <- plm(log(kidmort_unicef) ~ democracy + log(GDPcur), data = ross,
    index = c("id", "year"), model = "random")
coeftest(re.mod)[1:3, ]
```

```
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)    12.31      0.255    48.3 1.6e-216
## democracy      -0.19      0.034    -5.6  2.4e-08
## log(GDPcur)    -0.36      0.011   -32.8 1.5e-139
```

```
coeftest(fe.mod)[1:2, ]
```

```
##             Estimate Std. Error t value Pr(>|t|)
## democracy      -0.14      0.033    -4.3  2.3e-05
## log(GDPcur)    -0.38      0.011   -33.1 3.5e-126
```

```
coeftest(pooled.mod)[1:3, ]
```

```
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)     9.76      0.345      28 2.9e-115
## democracy      -0.96      0.070     -14  1.2e-37
## log(GDPcur)    -0.23      0.015     -15  1.2e-42
```

- More general random effects models using `lmer()` from the `lme4` package

Hausman tests

- Can we test the assumption that $\mathbb{E}[U_i|\underline{X}_i, \underline{D}_i] = \mathbb{E}[U_i]$?
  - If true (and all the RE assumptions hold), then RE and FE are consistent, but RE is efficient.
  - If false, then RE is inconsistent, but FE is consistent.
- A **Hausman test** uses these facts to develop a hypothesis test of the assumption:
  - If FE and RE estimates are similar $\rightsquigarrow$ assumption plausible.
  - If FE and RE very different $\rightsquigarrow$ assumption perhaps not plausible.
- Limitations:
  1. We must maintain strict exogeneity for null and alternative.
  2. Must maintain that U_i is homoskedastic (not required for FE)
  3. Limited to comparing coefficients on variables that vary in i and t.

Calculate the Hausman test

- Let $\widehat{SE}[\widehat{\tau}_{FE}]$ and $\widehat{SE}[\widehat{\tau}_{RE}]$ be the estimated SEs of the estimators.
  - Under the null that RE is correct, $\widehat{SE}[\widehat{\tau}_{FE}] > \widehat{SE}[\widehat{\tau}_{RE}]$
- **Hausman test statistic**:

$$H = \frac{\widehat{\tau}_{FE} - \widehat{\tau}_{RE}}{\left(\widehat{SE}[\widehat{\tau}_{FE}]^2 - \widehat{SE}[\widehat{\tau}_{RE}]^2\right)^{1/2}}$$

- Under the null hypothesis that RE is correct, H is asymptotically normal.
- When $\widehat{\tau}_{FE}$ and $\widehat{\tau}_{RE}$ are very different relative to their uncertainty, H will be big in absolute value and we will reject the null.

Hausman test in R

```r
phtest(fe.mod, re.mod)
```

```
##
##  Hausman Test
##
## data:  log(kidmort_unicef) ~ democracy + log(GDPcur)
## chisq = 70, df = 2, p-value = 8.041e-16
## alternative hypothesis: one model is inconsistent
```

3/ Fixed effects with heterogeneous treatment effects

Potential outcomes in the general setting

- Let's allow for heterogeneous treatment effects:

$$\tau_{it} = Y_{it}(1) - Y_{it}(0)$$

- Keeping the old linearity in X_{it} assumption:

$$Y_{it} = X_{it}'\beta + \tau_{it}D_{it} + U_i + \varepsilon_{it}$$

- Add and substract τD_{it}, where $\tau = \mathbb{E}[\tau_{it}]$:

$$Y_{it} = X_{it}'\beta + \tau D_{it} + U_i + \eta_{it}$$

- Where the combined error is:

$$\eta_{it} = \underbrace{(\tau_{it} - \tau)D_{it}}_{\text{non-constant effects}} + \underbrace{Y_{it}(0) - E[Y_{it}(0)|\underline{X}_i, U_i]}_{\text{typical errors, } \varepsilon_{it}}$$

Assumptions

- Earlier we showed that strict ignorability implied strict exogeneity for ε_{it}. What about η_{it}?

$$\mathbb{E}[\eta_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

- Since $\eta_{it} = (\tau_{it} - \tau)D_{it} + \varepsilon_{it}$ and we showed that $\mathbb{E}[\varepsilon_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$, it suffices to show:

$$\mathbb{E}[(\tau_{it} - \tau)D_{it}|\underline{X}_i, \underline{D}_i, U_i] = 0$$

Non-constant effects errors

- How does the non-constant effect error work here?

$$\begin{aligned}
\mathbb{E}[(\tau_{it} - \tau)D_{it}|\underline{X}_i, \underline{D}_i, U_i] &= D_{it}(\mathbb{E}[\tau_{it} - \tau|\underline{X}_i, \underline{D}_i, U_i]) \\
&= D_{it}(\mathbb{E}[\tau_{it}|\underline{X}_i, \underline{D}_i, U_i] - \tau) \\
&= D_{it}(\mathbb{E}[\tau_{it}|\underline{X}_i, U_i] - \tau) \\
&= D_{it}(\mathbb{E}[\tau_{it}|\underline{X}_i, U_i] - \mathbb{E}[\tau_{it}])
\end{aligned}$$

- Thus, we can see that the combined error will only satisfy the strict exogeneity assumption of fixed effects when

$$E[\tau_{it}|\underline{X}_i, U_i] = E[\tau_{it}]$$

- This is when the treatment effects are independent of the unit effects and the covariates.

Regression bias?

- We've seen this before: it's a general problem with regression and varying treatment effects.

$$\eta_{it} = \underbrace{D_{it}(\tau_{it} - \tau)}_{\text{non-constant effects}} + \underbrace{Y_{it}(0) - E[Y_{it}(0)|\underline{X}_i, U_i]}_{\text{typical errors}}$$

- Generally the issue here is that non-constant effects induce correlation between the treatment and the error term.
- Distinct from confounding bias since we could, in principle, estimate $E[\tau_{it}|\underline{X}_i, U_i]$ to then calculate $E[\tau_{it}]$
- Overall ATE still nonparametrically identified, even if the FE regression doesn't estimate it.

Strict exogeneity/ignorability

$$Y_{it}(d) \perp\!\!\!\perp \underline{D}_i | \underline{X}_i, U_i$$

- Strict ignorability is **very strong**.
- Rules out the following:
  - D_{it} affects Y_{it} which then affects $D_{i,t+1}$
  - Basically, any feedback between treatment and the outcome
- Can we weaken this? Yes! **Sequential ignorability**:

$$Y_{it}(d) \perp\!\!\!\perp D_{it} | \underline{X}_{it}, \underline{D}_{i,t-1}, U_i$$

- Note here that the we only condition up to t so that the errors can be correlated with future $D_{i,t+1}$ and so on.
- This implies **sequential exogeneity** of the errors:

$$\mathbb{E}[\varepsilon_{it} | \underline{X}_{it}, \underline{D}_{it}, U_i] = 0.$$

Strict ignorability example

- Example: economic interdependence between countries ($D_{it} = 1$ if county-dyad i is interdependent in period t) and conflict severity (Y_{it}) between countries.
- Strict ignorability assumption implies shocks to conflict severity at t uncorrelated with:
  - future values of conflict severity
  - economic interdendence
  - any other time-varying covariate

Lagged dependent variables

$$Y_{it} = \beta Y_{i,t-1} + \tau D_{it} + U_i + \varepsilon_{it}$$

- Fixed effects models with lagged dependent variables is much harder.
- Easiest to see with first differences:

$$(Y_{it} - Y_{i,t-1}) = \beta(Y_{i,t-1} - Y_{i,t-2}) + \tau(D_{it} - D_{i,t-1}) + (\varepsilon_{it} - \varepsilon_{i,t-1})$$

- Obviously, $Y_{i,t-1}$ is correlated with the $\varepsilon_{i,t-1}$.
- This is sometimes called a **dynamic panel model**, where we can't rely on the exogeneity assumption alone.
- $\rightsquigarrow$ need an instrumental variable approach (coming up in a few weeks).

4/ Cumulative effects

Contemporaneous vs Cumulative effects

- Another assumption we've been making is that there is only a contemporaneous effect: τD_{it}.
- Implicitly or explicitly fixing the past history of the treatment.
- What if we want to estimate the cumulative effects?
- Very difficult, if not impossible with fixed effects models.
- Why?
  - For cumulative effects, we need to consider the effects of treatment on time-varying confounders, $X_{it}(\underline{d}_{i,t-1})$.
  - Those pathways might be hard to identify

New notation

- Two-period effects: $Y_{it}(d_{t-1}, d_t)$
- New consistency assumption:

$$Y_{it} = Y_{it}(D_{i,t-1}, D_{it})$$

- In general, we will be interested in average treatment effects:

$$\mathbb{E}[Y_{it}(d_{t-1}, d_t) - Y_{it}(d^*_{t-1}, d^*_t)].$$

- Let $\underline{d} = (d_1, \ldots, d_T)$ be one entire history of $\underline{D}$.
- Partial history: $\underline{d}_t = (d_1, \ldots, d_t)$.

Fixed effects causal models

- Need a causal model:

$$Y_{it}(d_{t-1}, d_t) = X'_{it}(d_{t-1})\beta_c + \tau_{i,t-1}d_{t-1} + \tau_{it}d_t + U_i + \varepsilon_{it}$$

- β_c have c subscript here to denote difference from above fixed effect regressions.
- Allows for heterogeneous effects in each unit-period.

$$\mathbb{E}[Y_{it}(1,1) - Y_{it}(0,0)] = \mathbb{E}[\underbrace{\tau_{i,t-1} + \tau_{it}}_{\text{direct effects}} + \underbrace{(X_{it}(1) - X_{it}(0))'\beta_c}_{\text{effect of } D_{i,t-1} \text{ through } X_{it}}]$$

Cumulative effects notes

- Sobel paper shows that under fixed effects-style confounding can only estimate contemporaneous effect, where d_{t-1} is the same for the comparison:

$$\mathbb{E}[Y_{it}(d_{t-1}, 1) - Y_{it}(d_{t-1}, 0)] = \mathbb{E}[\tau_{it}]$$

- β_c is very difficult to identify! Need more restrictions.
- Exception: X_{it} is unaffected by $D_{i,t-1}$ so that $X_{it}(1) = X_{it}(0)$ and so:

$$\mathbb{E}[Y_{it}(1, 1) - Y_{it}(0, 0)] = \mathbb{E}[\tau_{i,t-1} + \tau_{it}]$$