Optional Reading: event study regression, application and theory

(Jing Li, Miami University)

1. Event study or dynamic DID aims to show how treatment effect evolves over time, and is especially useful when timing of treatment is staggered—for instance, different patients receive new medicine at different times, or, states change certain law in different years. Even if treatment timing is not staggered (for example, a new national law starts to affect all states in the same year), event study can be useful for estimating dynamic treatment effect, e.g., the treatment effect one year, two years.... following the enactment of national law.

2. The key regressor in a conventional panel data Fixed Effect DID regression[1] is a dummy variable $post_{it}$ that equals one if a new law applies in state i in year t. $post_{it}$ always equals zero for a state that has no new law; it always equals one if a state changes to the new law before the first year in the sample.

3. Those states with time-invariant $post_{it}$ are in control group since no event (law change) occurs during the sample period. The states with time-varying $post_{it}$ are in treatment group. Under assumption of parallel trends, control group could provide the counterfactual of what would have happened to treatment group in the absence of treatment (law change)

4. The conventional or static DID regression with long form panel data is

$$y_{it} = \beta post_{it} + \alpha_i + \gamma_t + \theta x_{it} + e_{it} \tag{1}$$

   where α_i is state FE (a set of state dummies); γ_t is time FE (a set of time dummies); x_{it} is time-varying control variable or covariate. The DID estimate is β.

5. Implicitly, conventional DID treats a state passing new law in the j-th year equally as a state passing the new law in the k-th year. By contrast, dynamic DID or event study explicitly takes into account the staggered timing of event. In other words, conventional DID reports aggregate before-and-after-treatment difference in outcome, whereas event study reports separately disaggregate j-period-after-and-before-treatment difference.

[1]See `http://www.fsb.miamioh.edu/lij14/311r_or_paneldid.pdf` for discussion of FEDID regression.

6. The dynamic DID or event study regression with long form panel data is

$$y_{it} = \sum_{j \neq -1} \beta_j 1(year - eventyear = j) + \alpha_i + \gamma_t + \theta x_{it} + e_{it} \tag{2}$$

where indicator $1(A) = 1$ if A is true. Note that the index $j < 0$ for pre-treatment periods (lags), while $j > 0$ for post-treatment periods (leads). We drop the dummy if $year - eventyear = -1$ to avoid perfect collinearity. That category is the base category.

7. The parallel trend assumption is supported by the finding that all β_j are statistically insignificant for $j < -1$. The treatment effect is supported by the finding that β_j is statistically significant for some $j \geq 0$. The dynamic DID allows for time-evolving treatment effect, i.e., $\beta_j \neq \beta_k$ for $j > 0, k > 0, j \neq k$

8. We can interpret β_j as follows (see the second section of this note for more discussion): it is the difference between the treatment and control groups, relative to the period -1, j periods from the treatment. Mathematically,

$$\beta_j \approx (\bar{y}^{treatmentgroup}_{post-j-period} - \bar{y}^{controlgroup}_{post-j-period}) - (\bar{y}^{treatmentgroup}_{pre-1-period} - \bar{y}^{controlgroup}_{pre-1-period}) \tag{3}$$

Or after algebra rearrangement:

$$\beta_j \approx (\bar{y}^{treatmentgroup}_{post-j-period} - \bar{y}^{treatmentgroup}_{pre-1-period}) - (\bar{y}^{controlgroup}_{post-j-period} - \bar{y}^{controlgroup}_{pre-1-period}) \tag{4}$$

The focus of event study is the set of β_j.

9. Thus, the reason why we hope β_j coefficients for $j < -1$ be equal to zero (statistically insignificant) is that we want the difference between the treatment and control groups to be the same in all periods leading up to the treatment, i.e., parallel trends.

10. Next we try to illustrate the main idea of the paper `https://users.nber.org/~jwolfers/papers/bargaining_in_the_shadow_of_the_law.pdf`. The outcome y variable asmrs is female suicide rate; the event is that some states underwent no-fault divorce reforms in different years (staggered treatment). The data are a balanced panel with 49 states from 1964 to 1996.

```
> d = read.csv("311r_esdata.csv")
> head(d)
```

```
  stateid year X_nfd post     asmrs     pcinc     asmrh       cases   weight   copop
1       1 1964   1971    0 35.63988 12406.18 5.007341 0.012312243 1715156 1715156
2       1 1965   1971    0 41.54375 13070.21 4.425367 0.010419407 1715156 1725186
3       1 1966   1971    0 34.25234 13526.66 4.874819 0.009900097 1715156 1735219
4       1 1967   1971    0 34.46502 13918.19 5.362014 0.009974692 1715156 1745250
5       1 1968   1971    0 40.44011 14684.81 4.643759 0.012400659 1715156 1755283
6       1 1969   1971    0 42.49012 15638.88 5.296976 0.015006756 1715156 1765316
> unique(d$stateid)
 [1]  1  4  5  6  8  9 10 11 12 13 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
[33] 38 39 40 41 42 44 45 46 47 48 49 50 51 53 54 55 56
> table(d$stateid)
 1  4  5  6  8  9 ...
33 33 33 33 33 33 ...
> min(unique(d$year))
[1] 1964
> max(unique(d$year))
[1] 1996
```

Note that stateids have gaps (no stateid equals 2 or 3)

11. The $post_{it}$ dummy equals one if state i is treated in year t. It is instructive to examine which state has time-varying or time-invariant $post_{it}$

```
idset = unique(d$stateid)
j = 1
d$group = rep("treated group---post is varying",nrow(d))
while (j<=length(idset)) {
        cat("state id is ",idset[j])
        if (sum(d$post[d$stateid==idset[j]])==0) {
                d$group[d$stateid==idset[j]] = "control group---post is always 0"
        }
        else if (sum(d$post[d$stateid==idset[j]])==33) {
                d$group[d$stateid==idset[j]] = "control group---post is always 1"
        }
        cat("this state is in",d$group[d$stateid==idset[j]][1])
        if (d$group[d$stateid==idset[j]][1] =="treated group---post is varying") {
```

```
        m = lm(asmrs~post+pcinc+asmrh+cases,data=d[d$stateid==idset[j],])
        cat("coefficient of post is ",summary(m)$coef[2,1])
        cat("t value of post is ",summary(m)$coef[2,3])
        }
        j = j + 1
        cat("","")
}
```

```
state id is  1
this state is in treated group---post is varying
coefficient of post is  9.055808
t value of post is  1.746568
 ...
state id is  5
this state is in control group---post is always 0
...
state id is  22
this state is in control group---post is always 1
...
```

(a) stateid 1 is a state that had divorce law reform in 1971. Its post dummy equals zero before 1971, and equals one in and after 1971, implying that the post is time-varying. We then run a time series regression for that state, that is, we hold α_i constant, and regress female suicide rate (asmrs) onto post, per-capita income (pcinc), homicide mortality (asmrh), and rates for a family of four (cases). The insignificant t value 1.7465679 for post implies no significant before-and-after-reform difference in female suicide rate for that state.

(b) stateid 5 is a state that had no divorce law reform during the entire sample periods. Its post dummy always equals zero.

(c) stateid 22 is a state that had reform before 1964. Its post dummy always equals one.

(d) states 5 and 22 are in control group; states 1 is in treatment group

The control groups have the following 13 states

```
> unique(d$stateid[d$group=="control group---post is always 0"])
[1]  5 10 28 36 47
> unique(d$stateid[d$group=="control group---post is always 1"])
[1] 22 24 37 40 49 50 51 54
```

12. Next, we obtain an aggregate treatment effect by running static DID regression (1)

```
> staticdid = lm(asmrs~post+pcinc+asmrh+cases+factor(stateid)+factor(year),data=d)
> summary(staticdid)$coef[1:5,]
                  Estimate   Std. Error   t value     Pr(>|t|)
(Intercept)  5.187839e+01 3.677622e+00 14.106506 1.414887e-42
post        -2.515964e+00 1.099745e+00 -2.287770 2.228663e-02
pcinc       -1.118218e-03 2.064454e-04 -5.416532 7.046295e-08
asmrh        1.184598e+00 3.152737e-01  3.757363 1.781519e-04
cases       -1.785179e+02 5.393950e+01 -3.309595 9.560070e-04
```

The t value -2.287770 implies a significant aggregate before-and-after-reform change in *asmrs*. Recent study such as Callaway and Sant'Anna (2020) titled "Difference-in-Differences with multiple time periods" shows that this static DID estimate may be biased in the presence of staggered treatment.

13. In order to run the dynamic DID or event study regression (2), which provides disaggregate treatment effect, we need to make use of the variable X_nfd showing eventyear, i.e., the year in which the divorce reform began in a given state. For instance, the eventyear is 1971 for stateid 1. The eventyear is missing (NA) for 13 states in the control group. We need to take special care of those NA values later.

```
> colnames(d)[colnames(d) == "X_nfd"] = "eventyear" #rename X_nfd
> any(is.na(d$eventyear))
[1] TRUE
> sum(is.na(d$eventyear))
[1] 429
> 429/33
[1] 13
> table(d$eventyear)
1969 1970 1971 1972 1973 1974 1975 1976 1977 1980 1984 1985
  66   66  231   99  330   99   66   33   99   33   33   33
```

There are two states (66 divided by 33) for which the treatment took place in 1969. Those two states have the most post-treatment observations. By contrast, only one state started treatment in 1985, and that state has the most pre-treatment observations.

14. A new variable yeargap is created, which is $year - eventyear$ in regression (2).

```
> d$yeargap = d$year-d$eventyear
> table(d$yeargap)
-21 -20 -19 -18 -17 -16 -15 -14 -13 -12 -11 -10  -9  -8  -7  -6  -5  -4  -3  -2  -1
  1   2   2   2   2   3   3   3   6   7   9  12  22  25  32  34  36  36  36  36  36
  4   5   6   7   8   9  10  11  12  13  14  15  16  17  18  19  20  21  22  23  24
 36  36  36  36  36  36  36  36  35  34  34  34  34  33  33  33  30  29  27  24  14
> any(is.na(d$eventyear))
[1] TRUE
> max(d$yeargap,na.rm=T)
[1] 27
> 1996-1969
[1] 27
> min(d$yeargap,na.rm=T)
[1] -21
> 1964-1985
[1] -21
```

(a) yeargap has NA values because of eventyear

(b) the distribution of yeargap is informative by showing it is necessary to downplay values on the two tails. For instance, there is only one observation with yeargap $= -21$, which could lead to noise or imprecise estimate.

(c) To increase statistical power, it is common to combine those tails values (called binning), by considering, say, yeargap ≤ -15. See `https://www.aeaweb.org/articles?id=10.1257/jep.37.2.203` for more discussion.

15. The yeargap is categorical, for which we can generate a set of dummy variables or indicators, one for each value or level of yeargap. Those indicators are the key regressors $1(year - eventyear = j)$ in the event-study regression (2).

```
# creat dummy for each level of yeargap
ygset = sort(unique(d$yeargap))
```

```
j = 1
while (j<=length(ygset)) {
dum = as.integer(d$yeargap==ygset[j])
dum[is.na(dum)]=0
assign(paste0("v", j), dum)
j = j + 1
}
```

Those indicators are named v1, v2, ...v49—v1 equals one if yeargap is -21, v2 equals one if yeargap is -20, and so on. We need to replace NA values with 0 so that the states in control group are taken into account.

16. To avoid dummy variable trap or perfect collinearity, one of v indicators needs to be excluded in the regression. Typically the indicator associated with yeargap of -1, which is v21 in this case, is dropped. Below are results of the event-study regression (2)

```
> # dynamic DID or event study regression
> X = "+pcinc+asmrh+cases+factor(stateid)+factor(year)"
> fo = as.formula(paste("asmrs~",paste(paste0("v", c(1:20,22:49)),collapse ="+"),X)
> esdid = lm(fo, data = d)
> summary(esdid)$coef[1:52,]
                 Estimate   Std. Error     t value     Pr(>|t|)
(Intercept)  5.623195e+01 4.255578e+00 13.21370372 9.046853e-38
v1          -2.292073e+01 1.137384e+01 -2.01521439 4.406208e-02
v2          -1.208418e+01 8.219664e+00 -1.47015479 1.417317e-01
...
v19         -2.312587e+00 2.605016e+00 -0.88774387 3.748223e-01
v20         -5.157395e-01 2.612411e+00 -0.19741899 8.435267e-01
v22          2.507466e-01 2.613315e+00  0.09594962 9.235735e-01
v23         -1.619352e+00 2.602531e+00 -0.62222179 5.338916e-01
...
v49         -4.344756e-01 8.320606e+00 -0.05221682 9.583630e-01
pcinc       -1.104638e-03 2.062957e-04 -5.35463366 9.921363e-08
asmrh        1.080640e+00 3.195947e-01  3.38128461 7.402096e-04
cases       -1.903716e+02 5.498289e+01 -3.46237924 5.506183e-04
```

We see v21 is not included in the regression

17. Usually the event study regression has way too many coefficients including those for v indicators, unit fixed effects, and time fixed effects. Thus, a better way is to visualize only coefficients of v indicators with the coefficient plot.

```
> # coefficient plot for v dummies in event study regression
library(ggplot2)
beta1 = summary(esdid)$coef[2:49,1]
se = summary(esdid)$coef[2:49,2]
output = data.frame(model = c(-21:-2,0:27),beta1,se)
output$min = output$beta1 - 1.96*output$se
output$max = output$beta1 + 1.96*output$se
output$model = factor(output$model,levels=c(-21:-2,0:27))

ggplot(output, aes(x = model, y = beta1)) +
    geom_point(size = 2) +
    geom_errorbar(aes(ymin = min, ymax = max), width = 0.2) +
    geom_hline(yintercept = 0,linetype = "dashed") +
    geom_vline(xintercept = "0",linetype = "dashed",color = "red") +
    labs(title = "Coefficient Plot",
         x = "yeargap",
         y = "coefficients of yeargap indicator") +
    theme_minimal() +
    theme(
    panel.grid.major = element_blank(),
    panel.grid.minor = element_blank(),
    axis.text.x = element_text(angle = 90, vjust = 1, hjust = 1),
    axis.line = element_line(size = 1, color = "black")
  )
```

The coefficient plot is

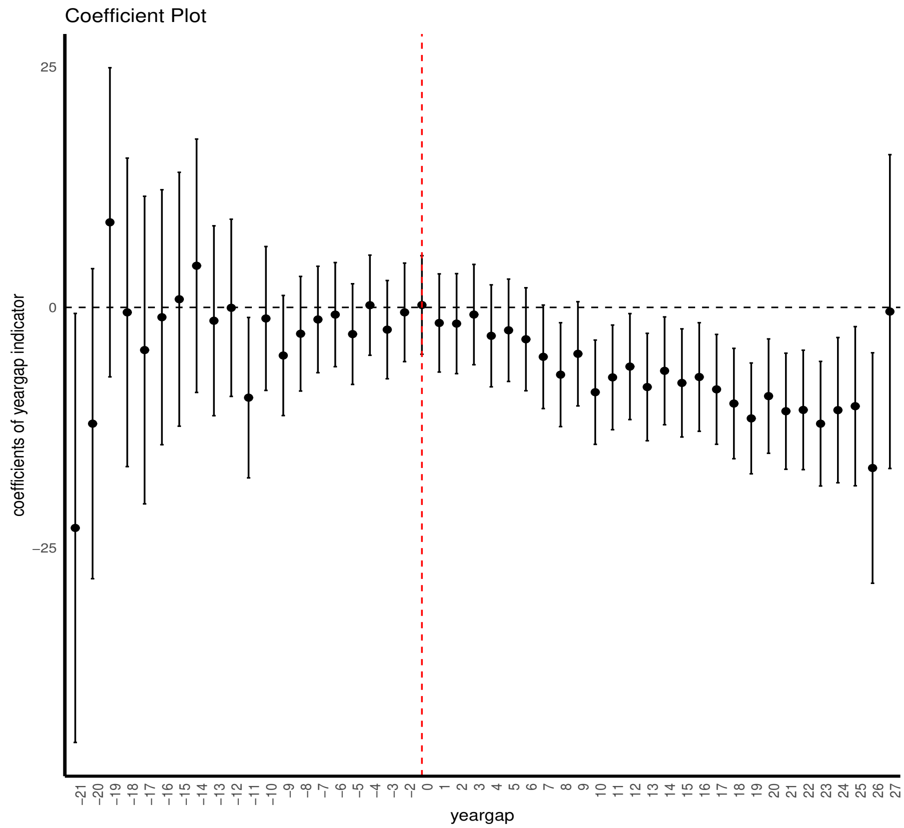

(a) A vertical red dash line is located at yeargap = 0

(b) The value 0, represented by a horizontal black dash line, is inside most confidence intervals before the vertical red dash line. This finding supports the parallel trend assumption

(c) After the vertical red dash line, we see there is a downward trend in point estimates. At yeargap = 8, the DID estimate becomes significant as 0 is outside its confidence interval

(d) The coefficient of indicator for yeargap=8 is -6.991146. Its interpretation is, relative to yeargap=-1, the difference in asmrs between the treated and untreated groups, 8 years after the change in law, is -6.991146.

(e) We conclude that the control group (those 13 states) produces satisfactory counterfactual for the treated group, and the treatment results in a downward trend in asmrh over time. In general, the treatment effects are statistically significant 8 years after the change in law.

Theory of event study regression

18. This section uses toy data shown below to illustrate the basic idea of event study regression.

```
> toy
          y i t post ed gap DM3 DM2 DM1 D0 D1 D2
1  35.63988 1 1    0 NA  NA   0   0   0  0  0  0
2  41.54375 1 2    0 NA  NA   0   0   0  0  0  0
3  34.25234 1 3    0 NA  NA   0   0   0  0  0  0
4  34.46502 1 4    0 NA  NA   0   0   0  0  0  0
5  40.44011 2 1    0  2  -1   0   0   1  0  0  0
6  42.49012 2 2    1  2   0   0   0   0  1  0  0
7  44.23907 2 3    1  2   1   0   0   0  0  1  0
8  49.91466 2 4    1  2   2   0   0   0  0  0  1
9  47.20313 3 1    0  3  -2   0   1   0  0  0  0
10 41.03307 3 2    0  3  -1   0   0   1  0  0  0
11 46.20281 3 3    1  3   0   0   0   0  1  0  0
12 39.24017 3 4    1  3   1   0   0   0  0  1  0
13 48.39191 4 1    0  4  -3   1   0   0  0  0  0
14 40.51945 4 2    0  4  -2   0   1   0  0  0  0
15 33.85831 4 3    0  4  -1   0   0   1  0  0  0
16 36.34627 4 4    1  4   0   0   0   0  1  0  0
```

(a) y is the outcome variable. There are four units, i = 1,2,3,4, and four time periods, t = 1,2,3,4. This is a long form panel data. y has both i and t subscripts.

(b) post is a dummy variable that equals one if a unit is treated in a given period. For instance, unit 1 is never treated since its post always equals zero. post has both i and t subscripts.

(c) ed means event date—the period in which treatment begins for a given unit. Treatments begin (switched on) in the second, third, and fourth periods for the units 2, 3, and 4, respectively. This is staggered treatment. ed only has i subscript.

(d) gap is the difference between t and ed. For instance, gap = -1 means one period before the treatment begins; gap = 2 means two periods after the treatment begins. gap has both i and t subscripts.

(e) DM3 is the indicator (dummy variable) that equals one when gap = -3; DM2 is the indicator for gap = -2,... D2 is the indicator for gap = 2. Notice that there is no missing or NA values in DM3, DM2,...D2.

19. The event study regression is

$$y_{it} = \beta_0 + \alpha_{-3}DM3 + \alpha_{-2}DM2 + \alpha_{-1}DM1 + \alpha_0 D0 + \alpha_1 D1 + \alpha_2 D2 \tag{5}$$

in which the error term, time fixed effect, and unit fixed effect are absent for expository purpose. There is no dummy variable trap for this regression since $DM3 + DM2 + ...D2 \neq 1$. Let us discuss

- For the group of $y_{11}, y_{12}, y_{13}, y_{14}$, we can show all the regressors in regression (5) are zero. As a result, the intercept term represents this group, i.e.,

$$\beta_0 = base\ group = \frac{y_{11} + y_{12} + y_{13} + y_{14}}{4}$$

Notice that this base group consists of the <u>never treated</u> unit 1 $(y_{11}, y_{12}, y_{13}, y_{14})$.

- DM3 equals one only for y_{41}. So α_{-3} is the difference between y_{41} and the base group

$$y_{41} = \beta_0 + \alpha_{-3} * 1 \Rightarrow$$

$$\alpha_{-3} = y_{41} - \beta_0 = y_{41} - base\ group$$

- Similarly we can show

$$\alpha_{-2} = \frac{y_{31} + y_{42}}{2} - base\ group$$

$$\alpha_{-1} = \frac{y_{21} + y_{32} + y_{43}}{3} - base\ group$$

$$\alpha_0 = \frac{y_{22} + y_{33} + y_{44}}{3} - base\ group$$

$$\alpha_1 = \frac{y_{23} + y_{34}}{2} - base\ group$$

$$\alpha_2 = y_{24} - base\ group$$

20. It follows that under the <u>assumption</u> of the treated and untreated groups being similar

in the pre-treatment periods (Assumption of Parallel Trend, AOPT), we have

$$\alpha_{-3} = 0, \alpha_{-2} = 0, \alpha_{-1} = 0 \quad (Parallel\ Trend)$$

An F test can be applied to check AOPT

21. If Assumption of Parallel Trend holds, then the base group could provide the counterfactual or potential outcome of what would have happened to the treated group in the absence of treatment. Thus, α_0 measures the instantaneous treatment effect; and α_1 and α_2 measure the subsequent treatment effects. Testing no treatment effect amounts to testing

$$\alpha_0 = 0, \alpha_1 = 0, \alpha_2 = 0 \quad (No\ Treatment\ Effect)$$

22. The results of regression (5) are

```
# event study regression applied to toy data
> summary(lm(y~DM3+DM2+DM1+D0+D1+D2),data=toy)$coef
             Estimate Std. Error    t value     Pr(>|t|)
(Intercept) 36.475250   2.051269 17.7818008 2.553968e-08
DM3         11.916661   4.586776  2.5980474 2.883041e-02
DM2          7.386040   3.552901  2.0788757 6.738787e-02
DM1          1.968581   3.133365  0.6282644 5.454416e-01
D0           5.204486   3.133365  1.6609897 1.310834e-01
D1           5.264371   3.552901  1.4817104 1.725588e-01
D2          13.439415   4.586776  2.9300353 1.675605e-02
> basegroup = mean(toy$y[toy$i==1])
> basegroup
[1] 36.47525
> mean(toy$y[toy$gap==-3],na.rm=T) - basegroup
[1] 11.91666
> mean(toy$y[toy$gap==-2],na.rm=T) - basegroup
[1] 7.38604
> mean(toy$y[toy$gap==-1],na.rm=T) - basegroup
[1] 1.968581
> mean(toy$y[toy$gap==0],na.rm=T) - basegroup
[1] 5.204486
```

```
> mean(toy$y[toy$gap==1],na.rm=T) - basegroup
[1] 5.264371
> mean(toy$y[toy$gap==2],na.rm=T) - basegroup
[1] 13.43942
```

The empirical results are consistent with our logical reasoning.

23. From statistics we learn that the standard error of sample average is inverse to sample size
$$se\left(\bar{y}\right)=\frac{\sigma}{\sqrt{n}},$$
or, bigger n (denominator of the ratio) can lead to smaller se and more precise estimate. This implies that the preciseness of α_0 is greater than α_1, and α_{-3} and α_2 are least precise. Intuitively, α_0 is precise because it uses all three treated units, whereas α_1 only uses two of them. Notice that the se of α_0 and α_{-1} (in green color) are smallest

24. In terms of confidence interval, α_0 has the narrowest one. It is not uncommon to see widest confidence intervals in the two tails of coefficient plot.

25. To increase preciseness or statistical power, we may consider "combining" or "binning" D1 and D2 by defining a new D1 that equals one if gap is greater than 0. Then we can show the coefficient of this "combined" dummy is
$$new\ \alpha_1=\frac{y_{23}+y_{34}+y_{24}}{3}-base\ group$$
Below are results of using new D1

```
# event study regression with combined indicator
> newD1 = as.integer(toy$gap>0)
> newD1[is.na(newD1)] = 0
> toy = data.frame(toy,newD1)
> summary(lm(y~DM3+DM2+DM1+D0+newD1),data=toy)$coef
             Estimate Std. Error    t value     Pr(>|t|)
(Intercept) 36.475250   2.213772 16.4765126 1.413931e-08
DM3         11.916661   4.950146  2.4073355 3.685080e-02
DM2          7.386040   3.834366  1.9262741 8.294407e-02
DM1          1.968581   3.381593  0.5821462 5.733580e-01
D0           5.204486   3.381593  1.5390633 1.548093e-01
```

```
newD1          7.989386   3.381593  2.3626099 3.977573e-02
> mean(toy$y[toy$gap>0],na.rm=T) - basegroup
[1] 7.989386
```

Notice the se decreases from 3.552901 to 3.381593 (blue color)

26. In reality, the event study regression usually includes unit and time fixed effects. To avoid the dummy variable trap, it is common to drop one of the gap-level indicators. For instance, we may drop DM1, but then the intercept (base group) becomes the never-treated unit 1 plus y_{21}, y_{32}, y_{43}, three treated units in their pre-treatment period. Other coefficients change accordingly, see results below

```
# event study regression with one gap-level indicator being dropped
> summary(lm(y~DM3+DM2+D0+D1+D2),data=toy)$coef
             Estimate Std. Error   t value     Pr(>|t|)
(Intercept) 37.318927   1.502953 24.830406 2.565864e-10
DM3         11.072984   4.250992  2.604800 2.627421e-02
DM2          6.542363   3.188244  2.052027 6.727964e-02
D0           4.360808   2.744004  1.589214 1.430962e-01
D1           4.420693   3.188244  1.386560 1.957074e-01
D2          12.595738   4.250992  2.963011 1.421472e-02
> basegroup = mean(toy$y[toy$i==1|toy$gap==-1],na.rm=T)
> basegroup
[1] 37.31893
> mean(toy$y[toy$gap==-3],na.rm=T) - basegroup
[1] 11.07298
```

Mathematically

$$
\begin{aligned}
new\ \alpha_{-3} = y_{41} - new\ base\ group = y_{41} - \frac{y_{11}+y_{12}+y_{13}+y_{14}+y_{21}+y_{32}+y_{43}}{7} = \\
\left(y_{41} - \frac{y_{11}+y_{12}+y_{13}+y_{14}}{4}\right) - \\
\left(\frac{y_{11}+y_{12}+y_{13}+y_{14}+y_{21}+y_{32}+y_{43}}{7} - \frac{y_{11}+y_{12}+y_{13}+y_{14}}{4}\right) \\
\approx \left(y_{41} - \frac{y_{11}+y_{12}+y_{13}+y_{14}}{4}\right) - \left(\frac{y_{21}+y_{32}+y_{43}}{3} - \frac{y_{11}+y_{12}+y_{13}+y_{14}}{4}\right) \\
= (difference\ at\ gap = -3) - (difference\ at\ gap = -1) = old\ \alpha_{-3} - old\ \alpha_{-1} \quad (6)
\end{aligned}
$$

Thus, 11.072984 (and other coefficients) has a difference-in-difference interpretation—it measures the treated-and-untreated difference at gap = -3, relative to the difference at gap = -1.

27. (Exercise) How to modify the data and event-study regression if the treatment terminates (switched off) at $t = 3$ for $i = 2$