SEVENTH STATA LAB

The difference-in-differences method (DiD)

Introduction

- The difference-in-differences method is one of the most widely used in labor economics, and more generally in applied economics.
- Also in this setting, we do not have a random assignment of a treatment (as in the RCT).
- However, we do have a so-called "natural experiment": one group of people is affected by a particular intervention (X) for external reasons independent of their own choices (e.g., by the introduction of a law); at the same time, a second group of people whom we may consider similar are not affected by this intervention.
- The idea is simple: we can compare whether the performance of a particular outcome (Y) in the group influenced by the treatment (X=1) is different as a result of the intervention, compared with the same performance observed in the control group (X=0)
- The basic assumption is that, in the absence of the treatment, the X=1 group would have had the same trend in Y that we can observe in the X=0 group. In other words, the X=0 group is an appropriate counterfactual, correctly representing what would have happened to X=1 in the absence of the treatment.

The DiD specification

The DiD model can be estimated with an OLS specification that includes interactions

- Take the Lalonde (1985) example, on STATA as well:

Let X represent participation to the treatment. Let Y_t e income in each year *t*, where *t*=1975, 1978. We can estimate a DiD model with the following specification

$$Y_t = \beta_0 + \beta_1 X + \beta_2 D(t = 1978) + \beta_3 X * D(t = 1978) + e$$

$D(t = 1978)$ is a dummy variable =1 in year 1978, $X * D(t = 1978)$ is the interaction (the product) between the treatment variable X the variable $D(t = 1978)$

β_1: Difference in income between X=1 and X=0 in 1975

β_2: Difference in income between t=1978 and t=1975 for group X=0

$\beta_2 + \beta_3$: Difference in income between t=1978 and t=1975 for group X=1

β_3 is the estimate of the effect of the NWD program using the DiD specification. It is in fact the additional change in income that is observed in X=1 compared to the change observed in X=0!

An application of DiD: the effect of the minimum wage

- One application in which the DiD method has been widely used is the case of studying the effect of the minimum wage:
  - Economic theory suggests that the minimum wage, by increasing actual wages, could also influence several other variables: employment, profits, sales prices...
  - It is a hotly debated topic: a recent European directive mandated the introduction of a minimum wage in all countries where collective bargaining covers less than 80 percent of the workforce (Italy is exempt)...
- Card and Krueger (1994) first applied this approach to the case of raising the minimum wage from $4.25 to $5.05 in New Jersey in 1992...
- Card and Kruger's (1994) idea: compare fast-food employment trends in New Jersey (NJ) and Pennsylvania (PA) before and after 1992
- Fast-food restaurants are a type of business that hires many workers potentially affected by the minimum wage
- NJ and PA are two U.S. states that are geographically close and have very similar economies (PA can be considered a good counterfactual group)

The results in Card and Krueger (1994)

Variable	PA (i)	NJ (ii)	Difference, NJ – PA (iii)
	Stores by state		
1. FTE employment before, all available observations	23.33 (1.35)	20.44 (0.51)	−2.89 (1.44)
2. FTE employment after, all available observations	21.17 (0.94)	21.03 (0.52)	−0.14 (1.07)
3. Change in mean FTE employment	−2.16 (1.25)	0.59 (0.54)	2.76 (1.36)

The results show:

- an average decrease in PA employment (X=0) of 2.16 employees from before to after 1992
- an average increase in employment in NJ (X=1) of 0.59 employees from before to after 1992
- the DiD effect is the difference between growth in NJ (X=1) and growth in PA (X=0)
- thus, according to the results employment in NJ grew by 2.76 more workers than the trend observed in PA
- the effect of minimum wage on employment was positive!

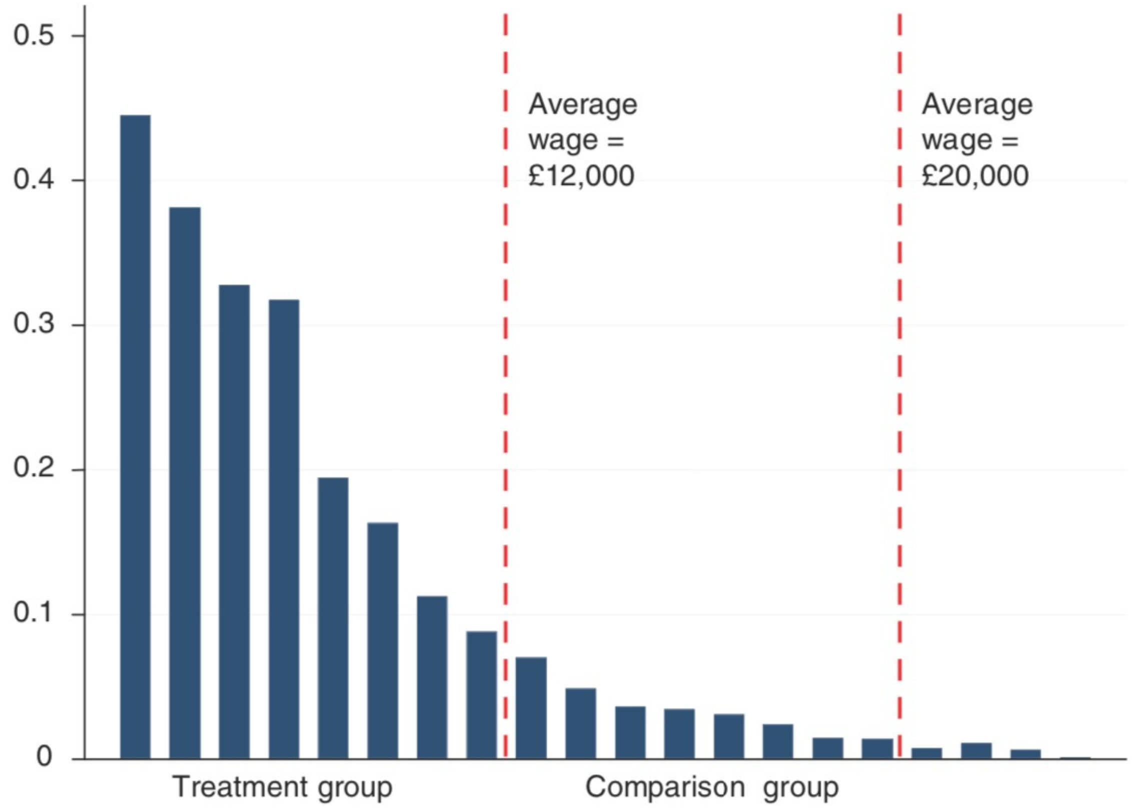

FIGURE 1. VALIDATION OF AVERAGE WAGE DATA
(*Comparison of proportion of low-wage workers and establishment average wages, WERS 1998*)

Notes: The *y*-axis shows the proportion of workers paid below the minimum wage (£3.60 per hour) in the establishment. The *x*-axis shows the average annual wage at the workplace. This is

Draca et al. (2011)

- Draca, Macin and van Reenen (2011) study the effect of the introduction of minimum wage in the United Kingdom in 1999 **on firms' profits**.

- These authors observe that firms paying lower average wages are more affected by the minimum wage because they hire a larger share of workers paid less than the minimum wage before the reform, and can therefore be considered "treated" (X=1).

- In contrast, firms with higher average wages can be used as a control group (X=0) because they are less affected by the reform.

Draca et al. (2011): DiD estimate of the effect of the MW on profits

```
. reg net_pcm post_treat post treat if pp==1,cluster(regno)

Linear regression                               Number of obs     =      4,112
                                                F(3, 950)         =       9.81
                                                Prob > F          =     0.0000
                                                R-squared         =     0.0199
                                                Root MSE          =     .15237

                                 (Std. err. adjusted for 951 clusters in regno)
```

net_pcm	Coefficient	Robust std. err.	t	P>|t|	[95% conf.	interval]
post_treat	-.0273999	.0137241	-2.00	0.046	-.054333	-.0004668
post	-.0118679	.0041968	-2.83	0.005	-.020104	-.0036318
treat	.0580316	.013659	4.25	0.000	.0312263	.0848368
_cons	.0698972	.0048096	14.53	0.000	.0604586	.0793358

- We define as treat=1 the enterprises with an average wage < than £12000 per year
- We define as post=1 the years 2000,2001,2002 (post-introduction of minimum wage)
- Treat=1 firms had 2.7% lower profit growth than treat=0 firms in the post=1 period

The DiD hypothesis and the placebo test

- The fundamental hypothesis of the DD model is that group X=0 represents a valid counterfactual for group X=1: without treatment intervention, the trend of Y in group X=1 would have been the same as the trend observed in group X=0
- It is a hypothesis that fundamentally cannot be tested directly
- There is, however, the possibility of conducting "placebo tests"
- In the years/periods <u>prior to the introduction of the policy/intervention on X=1</u>, was the trend of Y in group X=0 similar to the trend observed in group X=1? Ideally we would like it to be so
  - The failure of a placebo test is an indication that group X=0 might not be a good counterfactual group
  - The failure of a placebo test, however, could also be the consequence of "anticipatory effects"...if I know that next year there will be a higher minimum wage, I might begin to anticipate some of the choices induced by this policy

```
. reg net_pcm i.year##i.treat if pp==1, cluster(regno)

Linear regression                               Number of obs     =      4,112
                                                F(11, 950)        =       3.98
                                                Prob > F          =     0.0000
                                                R-squared         =     0.0215
                                                Root MSE          =     .15239

                                (Std. err. adjusted for 951 clusters in regno)
```

net_pcm	Coefficient	Robust std. err.	t	P>|t|	[95% conf.	interval]
year						
1998	.0036753	.0048508	0.76	0.449	-.0058442	.0131947
1999	-.0094913	.0060185	-1.58	0.115	-.0213024	.0023198
2000	-.0067129	.006691	-1.00	0.316	-.0198438	.006418
2001	-.0173143	.0064151	-2.70	0.007	-.0299037	-.0047249
2002	-.0203263	.0073264	-2.77	0.006	-.0347041	-.0059485
1.treat	.040096	.0155143	2.58	0.010	.0096498	.0705422
year#treat						
1998 1	.0203926	.0137357	1.48	0.138	-.0065631	.0473484
1999 1	.0296416	.0165639	1.79	0.074	-.0028644	.0621476
2000 1	-.0132727	.0198344	-0.67	0.504	-.052197	.0256515
2001 1	-.0104384	.0209385	-0.50	0.618	-.0515294	.0306527
2002 1	-.0030203	.0199008	-0.15	0.879	-.0420749	.0360342
_cons	.0722012	.0063045	11.45	0.000	.0598288	.0845736

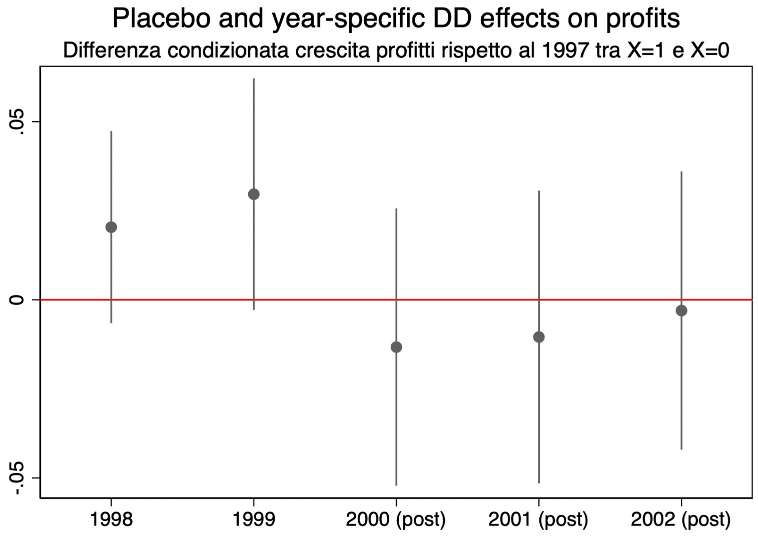

«Event study» specifications for the placebo test

- The placebo test can be estimated by interacting the treat variable with the year variable.
- In this way we estimate the difference in the growth of Y (relative to a base year) between group X=1 and group X=0 separately for each year.
- Ideally, in years prior to the intervention policy on X=1 the coefficients of the interactions should be statistically = to 0.
- In years after the policy, the coefficients can be interpreted as the effect of the policy on X=1 in each year after its introduction (short-run to long-run effects)

The placebo test proposed by Draca et al. (2014)

```
. reg net_pcm placpost_treat placpost plactreat,cluster(regno)

Linear regression                               Number of obs     =      4,689
                                                F(3, 1087)        =       0.92
                                                Prob > F          =     0.4315
                                                R-squared         =     0.0008
                                                Root MSE          =     5.5326

                                (Std. err. adjusted for 1,088 clusters in regno)
```

net_pcm	Coefficient	Robust std. err.	t	P>|t|	[95% conf.	interval]
placpost_treat	.2760885	.4393673	0.63	0.530	-.5860155	1.138192
placpost	.1161078	.0898231	1.29	0.196	-.0601384	.2923541
plactreat	-.3337987	.4310618	-0.77	0.439	-1.179606	.5120088
_cons	.0411154	.0439367	0.94	0.350	-.045095	.1273258

- Draca et al (2014) propose a different type of placebo than the "event-study" estimate (which is usually the most common approach):
- We keep the same group X=0, but define group X=1 as firms that paid less than £12000 on average in 1996
- We estimate whether there are differences in the growth of Y between the 1994-1996 period and the 1997-1999 periods by comparing group X=0 and group X=1
- Since there was no change in the minimum wage after 1996, this "fake treatment" should ideally have no significant effect

A final note on treatment effect heterogeneity and «staggered treatment» designs

- We say that a treatment is «staggered» when:
  - X=1 occurs with a different timing across observations (i.e. the MW is increased in NJ in 1992, in NY in 1993, and we want to estimate both treatment effects together)
  - X=1 can turn back to X=0 (i.e. the MW is increased in NJ in 1992, but then it is reduced to its original level in 1993, and we want to use all the variation across all years to estimate the treatment effect).
- In such cases, if the treatment effect is heterogeneous, we could run into problems because the weight of individual treatment effects used to compute the aggregate DiD estimate can be negative or have bad properties:
  - There are solutions to this problem proposed in the recent literature, which are mostly based on the idea of restricting the treatment-control comparisons used in order to compute the DiD estimator (see paper by de Chaisemartin and D'Haultfœuille 2024)