Differences-in-Differences and Regression Discontinuity

Before we face harsh reality again, let us think of an academic paradise where we could simply conduct the experiments we deem to be interesting. How would we do this? And how would we evaluate the results? It seems like a good idea to split your test subjects in two groups, the treated and the untreated. Selection into groups should be random. Then we apply some treatment and check the value of the outcome of interest for the two groups. How do we evaluate the impact of our treatment? Well, random assignment means that the treatment is uncorrelated (actually independent) of all the characteristics of the participants and we can simply take the difference between the outcome of the treatment group and the control group.

Now in real life we are rather constraint when it comes to experiments. But sometimes it is possible to find "natural" experiments, i.e. constellations where by chance or design some changes affected groups at different points in time or simply only applied to certain groups. An example due to Steve Pischke (2007) is that German federal states changed the school year to start in fall, in the mid 1960s. This required two short school years. Bavaria's schools already started in fall, so students in this state were unaffected. So we have a treatment and control group. How can we proceed?

The problem is that assignment is not totally random[1]We can use the so called differences-in-differences estimator (Diff-in-Diff). To apply Diff-in-Diff we need panel data and some (exogenous) change that affects a share of the observations in our sample, but not all of them, or at least not all at the same time. The classical example for Diff-in-Diff strategies are changes in state legislation, where some states either do not change their legislation or states change their legislation but at different points in time.

We separate our sample into treatment (those with changes at a certain point in time) and control group (no change) and take a double difference. First we take the difference across time but within group. This eliminates any group specific unobserved but time fixed effects, for the state example that means any unobserved non-time varying state fixed effects. Afterwards we take the difference of the differences. This will get rid of any time trends in the results (assuming that the treated group would have followed the same time trend as the untreated group).

Let's consider another example. The question is whether minimum wages have an impact on employment, and if so, how large an effect they have. David Card and Alan Krueger (1994) collected data from fast food restaurants at the border of New Jersey and Pennsylvania. They did this a two months before and seven months after a major increase in the minimum wage in New Jersey. So Pennsylvania, or rather the fast food restaurants in Pennsylvania are their untreated control group and the fast food restaurants in New Jersey are the treatment group. We can simply take the differences as explained above and get our point estimate. But for inference we would like to have standard errors.

[1] We might be confident that assignment is random conditional on some observable factors. E.g. where oil is found in a country is random conditional on the landside characteristics in an area.

Luckily we can simply implement the difference in difference strategy as an OLS estimator. Let us have a look at what we want to estimate first. We have an outcome of interest y, an indicator whether an observation belongs to the treatment or control group: $d = 1$ treated, $d = 0$ control and (at least) two time periods denoted by $t = 0$ and $t = 1$. Where the treatment, i.e. the legislation changes takes place between the two periods. The Diff-in-Diff estimator is simply:

$$\beta_1^{DD} = [E(y|d = 1, t = 1) - E(y|d = 1, t = 0)] - [E(y|d = 0, t = 1) - E(y|d = 0, t = 0)]$$

We can estimate the Diff-in-Diff parameter using the following regression model:

$$y_i = \beta_0 + \beta_1 d_i * t_i + \beta_3 d_i + \beta_4 t_i + \varepsilon_i$$

To check that β_1, i.e. the coefficient of the interaction between time and treatment dummy in the regression model is actually equal to β_1^{DD} we can simply take the expectations of the regression equation and calculate the difference:

$$E(y_i|d = 1, t = 1) = \beta_0 + \beta_1 + \beta_3 + \beta_4$$

$$E(y_s|d = 1, t = 0) = \beta_0 + \beta_3$$

$$E(y_i|d = 0, t = 1) = \beta_0 + \beta_4$$

$$E(y_i|d = 0, t = 0) = \beta_0$$

We can easily see that 1. the Diff-in-Diff estimator is equal to the coefficient of the interaction and 2. that the first difference takes out the state fixed effect β_3 and taking the difference of the differences gets rid of the time trend β_4.

The crucial assumption that we make is that the time trend in New Jersey (the treatment group) would have been the same as the time trend in the control group (Pennsylvania), hadn't the treatment group been treated. Similarly for the example by Pischke, where we have to assume that the trends in Bavaria did not differ from the trend in other federal states. So basically we assume away any unobserved time and state specific effects.

Another way to use institutional features is the use of a Regression Discontinuity (RD) design. We use RD when passing a certain threshold induces a change in the RHS variable of interest. The idea is that observations just below and just above the threshold are fairly comparable. Let's have a look at some examples. David Lee (2008) asks the question whether representatives of incumbent political parties have a higher probability of being elected. Incumbency is determined by winning the previous election, so we can use the share of votes by which a party won as our RHS variable. The threshold is zero, positive numbers indicating the party won the election, negative numbers that they lost. John DiNardo and David Lee (2004) use the same idea on union representation in the workplace. A workplace in the US is unionized if the union wins an election.

How do we estimate such a model. We fit a model, but allow for a discontinuity at the threshold. This is achieved by using a dummy variable which is one for all observations above the threshold and zero otherwise and introducing this dummy in level as completely interacted with all regressors on the RHS. The effect is then simply the coefficient on the dummy variable.