Panel Data Methods - Difference in Differences

Applied Microeconometrics

2011/12

References

- Angrist and Pischke, ch.5.
- Woolridge, ch. 10 (basics) and 11 (more advanced).
- Cameron and Trivedi, ch. 21 and 22.
- D. Card and A. Krueger (1994): "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania", American Economic Review, Vol. 84 (4), pp. 772-793.

Overview

1. Introduction
2. The Difference in Differences (DD) Estimator
   - Basic Idea
   - Regression Implementation
3. Extensions
   - Multiple control groups: Difference in Difference in Differences (DDD)
   - Further Extensions
4. Validity of the DD estimator
5. Later: More general panel data estimators
   - Random Effects estimators
   - Fixed Effects estimators
   - Dynamic Models

1. Introduction

- General Problem so far: if treatment is not randomly assigned, coefficients capture the treatment effect plus selection bias.
- Solution Concepts to handle unobserved heterogeneity so far:
  - OLS: control for all factors that determine participation and that are related to potential outcomes.
  - IV: instrument induces quasi-experimental variation in treatment status.
  - RDD: discontinuity induces quasi-experimental variation in treatment status if units have imprecise control over the running variable.
- Panel Data offer another powerful way to tackle issues related to omitted variable bias.
- Specifically, panel data allow to control for **unobserved** but **fixed** factors that drive participation and are related to potential outcomes.
- The trick is to exploit to have several observations which all contain the same unobserved information.

- As an example, assume potential outcomes for individual i at time t can be written as:

$$\begin{aligned} Y_{0it} &= \alpha + \epsilon_{it}, \\ Y_{1it} &= Y_{0it} + \rho. \end{aligned}$$

- Observed outcomes are given by

$$Y_{it} = \alpha + \rho D_{it} + \epsilon_{it},$$

  but treatment is not as good as randomly assigned (D_{it} is not independent of ϵ_{it}).
- The crucial assumption that allows exploiting the panel structure of the data is that the unobserved component of potential outcomes ϵ_{it} can be decomposed.

Assumption

1. *ϵ_{it} can be written as*

$$\epsilon_{it} = \gamma_i + \lambda_t + \eta_{it},$$

   where

   1. *γ_i is specific to individual i and fixed over time;*
   2. *λ_t is a time trend; and*
   3. *η_{it} is a transitory mean zero noise term.*

2. *Selection into treatment only depends on the individual fixed effect γ_i but is independent of λ_t or η_{it}.*

   1. *$E(\lambda_t|D_{it}) = E(\lambda_t)$ and*
   2. *$E(\eta_{it}|D_{it}) = E(\eta_{it}) = 0$; and hence*
   3. *$E(\epsilon_{it}|D_{it}) = E(\gamma_i|D_{it}) + E(\lambda_t)$*

   Hence, treatment and control group differ only in terms of the individual fixed effect, not in terms of the time trend and transitory shocks to outcomes.

- γ_i is also known as a **fixed effect**. Note that the fixed effect enters **additively** and **linearly**!
- Example?

- Under this assumption we can write observed outcomes as

$$Y_{it} = \alpha + \gamma_i + \rho D_{it} + \lambda_t + \eta_{it}.$$

- Under the assumption above we can take advantage of multiple observations on each unit and eliminate the fixed effect by, for example, differencing the equation above:

$$\Delta Y_{it} = \Delta D_{it} \rho + \Delta \lambda_t + \Delta \eta_{it},$$

  where the Δ denotes changes of the variable from $t-1$ to t.
- Note that for differencing to work it is necessary that the fixed effect enter additively and linearly!
- Under rather mild assumptions $\Delta \eta_{it}$ is uncorrelated to ΔD_{it} and running OLS on the differenced outcome equation yields the causal effect.

Observation

When the **level** *of potential outcomes differs between treatment and control group due to a linear and additive fixed effect, the* **change** *of potential outcomes over time does not differ.*

Example 1

- Assume you would like to estimate how business taxes impact on FDI.
- Arguably, the attractiveness of a region is partly determined by unobserved factors. Which?
- At least within a short period of time, these unobserved factors are likely to be fix.
- Hence, the unobserved factors influence only the **level** of FDI but not its **change** over time.
- Estimating an equation that relates the change in tax rates to the change in FDI is therefore less likely to suffer from an omitted variable bias.

Example 2

- Wages, health, happiness ... are all outcomes that depend on observable variables but also genetic factors which are unobserved.
- To tackle this kind of unobserved heterogeneity, a lot of studies use data on monozygotic twins.
- The idea is that differences in outcomes between monozygotic twins can solely be attributed to observable factors.

- The Difference in Differences (DD) estimator is the simplest estimator that makes use of data with a time dimension.
- The DD estimator can be interpreted as a fixed effects estimator that uses only aggregate (group level) data.
- That means that the DD estimator uses differencing at the group level, not at the individual level as in our introductory example.
- This can be done if treatment status varies only at the group level (e.g. state, cohort).
- Non random assignment of treatment must therefore come from unobserved variables at the group level.
- The DD approach captures these unobserved variables by a group level fixed effect.
- Since the DD estimator does not use data at the individual level it can also be used in a repeated cross section.
- We will make the following two additional assumptions:
  1. there are only 2 periods: "before treatment" ($t = 0$) and "after treatment" ($t = 1$); and
  2. the treatment variable is binary;

- Card and Krueger (1994) want to estimate the impact of a minimum wage on employment.
  - In 1992, the state of New Jersey increased its minimum wage by roughly 20%.
  - In Pennsylvania (neighboring state) the minimum wage did not change.
  - Card and Krueger have data on employment at fast food restaurants close to the state border for both states.
- There are two naive approaches to estimate the employment effect of the minimum wage.
  1. A comparison of employment levels in New Jersey and Pennsylvania after the introduction of the new minimum wage.
  2. A comparison of employment levels in New Jersey before and after the introduction of the new minimum wage.
- Both naive approaches are unlikely to uncover the causal effect of the minimum wage.
- Why?

- Let us assume a constant treatment effect and abstract from any covariates so we write potential outcomes as

$$Y_{0st} = \epsilon_{st} \text{ and } Y_{1st} = Y_{0st} + \rho,$$

where the index s denotes the state.

- Moreover, we assume that ϵ_{ist} can be decomposed into
  1. a group level (state) effect γ_s (that is the fixed effect);
  2. a time trend λ_t common to all states; and
  3. a transitory mean zero noise term η_{st}.
- While γ_s can be different for the two states, the time trend and the idiosyncratic noise term do not vary systematically between states.
  $\Rightarrow$ Hence, treatment is as good as randomly assigned conditional on the state effect γ_s.
- The observed outcome can be written as

$$Y_{st} = \gamma_s + \lambda_t + \rho D_{st} + \eta_{st},$$

where $E(\eta_{st}|s, t, D_{st}) = E(\eta_{st}) = 0$.

2. Difference in Differences

- Now consider what we would get by comparing average employment in both states before ($t = 0$) and after treatment ($t = 1$).

$$\begin{aligned} E(Y_{st}|s = \text{Penn}, t = 1) - E(Y_{st}|s = \text{Penn}, t = 0) &= \lambda_1 - \lambda_0; \\ E(Y_{st}|s = \text{NJ}, t = 1) - E(Y_{st}|s = \text{NJ}, t = 0) &= \rho + \lambda_1 - \lambda_0. \end{aligned}$$

- Hence, the treatment effect ρ is given by the difference in differences:

$$\begin{aligned} E(Y_{st}|s = \text{NJ}, t = 1) - E(Y_{ist}|s = \text{NJ}, t = 0) &\quad - \\ E(Y_{st}|s = \text{Penn}, t = 1) - E(Y_{ist}|s = \text{Penn}, t = 0) &= \rho. \end{aligned}$$

- This can easily be estimated using sample means.

- Let us summarize the key idea:
  1. The comparison over time within a state eliminates the state fixed effect.

     ⇒ We can remove differences in employment levels in the two states that have nothing to do with the minimum wage, by considering the difference in employment levels before and after the introduction of the new minimum wage.

     ⇒ In case of NJ, this difference captures the general time trend in employment and the treatment effect.
  2. To net out the general time trend we compare the development in employment levels in NJ with the development of employment in the control state which consists only of the time trend.
- The following assumption is therefore crucial.

Assumption 1

The time trend $\lambda_1 - \lambda_0$ is the same in both states.

- The following figure illustrates the workings of the DD estimator.

2. Difference in Differences

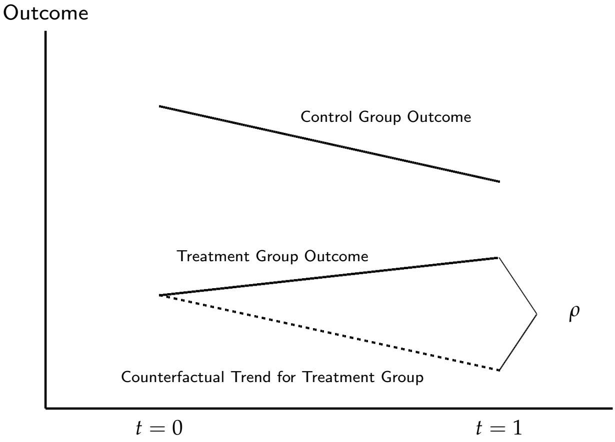

- Some of Card's results relating to the average employment levels in fast-food restaurants are shown below (with standard errors in parentheses).

	Before Increase	After Increase	Difference
New Jersey	20.44	21.03	0.59
(Treatment)	(0.51)	(0.52)	(0.54)
Pennsylvania	23.33	21.17	-2.16
(Control)	(1.35)	(0.94)	(1.25)
Difference	-2.89	-0.14	2.76
	(1.44)	(1.07)	(1.36)

- The difference in difference estimator shows a small increase in employment in New Jersey where the minimum wage increased.
- The study has been very controversial but helped to change the common presupposition that a small change in the minimum wage from a low level was bound to cause a significant decrease in employment.

- The DD estimator can easily be implemented using regression. This is a convenient way to obtain estimates and the corresponding standard errors in one step.
- Let
  - NJ denote a dummy equal to one for restaurants in New Jersey and let
  - d_1 be a dummy variable equal to one for observations after the introduction of the new minimum wage.
- Then the DD estimate equals the coefficient ρ from the following regression:

$$Y_{st} = \alpha + \gamma NJ + \lambda d_1 + \rho(NJ \cdot d_1) + \eta_{st}.$$

- Note that
  - the variable $NJ \cdot d_1$ equals D_{st} and
  - this is a fully saturated model since there are four possible state-year combinations and four parameters.

- To show that ρ equals the DD estimator, we make use of the fact that in a saturated model the CEF is necessarily linear in the explanatory variables. To see this, define

$$
\begin{aligned}
\alpha^* &= E(Y_{st}|s = \text{Penn}, t = 0); \\
\gamma^* &= E(Y_{st}|s = \text{NJ}, t = 0) - E(Y_{st}|s = \text{Penn}, t = 0); \\
\lambda^* &= E(Y_{st}|s = \text{Penn}, t = 1) - E(Y_{st}|s = \text{Penn}, t = 0); \text{and} \\
\rho^* &= \big(E(Y_{st}|s = \text{NJ}, t = 1) - E(Y_{st}|s = \text{NJ}, t = 0)\big) \\
&- \big(E(Y_{st}|s = \text{Penn}, t = 1) - E(Y_{st}|s = \text{Penn}, t = 0)\big).
\end{aligned}
$$

- The CEF can be written as

$$E(Y_{st}|s, t) = \alpha^* + \gamma^* NJ + \lambda^* d_1 + \rho^* (NJ \cdot d_1).$$

- Since this is a linear function in the explanatory variables NJ, d_1 and $NJ \cdot d_1$ the regression function equals the CEF (recall the linear CEF theorem from the chapter on OLS). Thus the coefficients from the linear regression recover the parameters of the CEF and we have

$$(\alpha^*, \gamma^*, \lambda^*, \rho^*) = (\alpha, \gamma, \lambda, \rho)$$

- Hence, ρ is identical to the DD estimator.

- Of course, the control group can consist of more than a single control state.
- Including several states as controls is beneficial since it provides a hedge against idiosyncratic shocks in a control state which might impair the common trend assumption.
- Assume we also had data on Conneticut in the example above. We could still work with the same regression function

$$Y_{st} = \alpha + \pi Conn + \gamma NJ + \lambda d_1 + \rho(NJ \cdot d_1) + \eta_{st}.$$

- Now λ would capture an average time trend for Pennsylvania and Conneticut. Precisely, λ captures average employment differences between
    - establishments which are either in Pennsylvania or Conneticut in $t = 1$ with
    - establishments which are either in Pennsylvania or Conneticut in $t = 0$.
- The treatment effect ρ would now be obtained by using the average of Pennsylvania and Conneticut as a "control" state.

- A still more convincing analysis than just using multiple control groups would be possible if we could define a "treatment" and a "control" group **within** each state.
- In the minimum wage example, assume we also had data on employment in sectors not affected by minimum wage legislation.
- Then we could think about two possible DD strategies:
  1. we could use employment in the non affected sector in the treatment state as the control group; or
  2. We would use employment in the fast food sector in a control state as the control group (approach so far).
- There is a pro and a con for each approach:
  1. The first strategy would be immune to different time trends across states but would hinge on the assumption that the time trend in employment is the same for different sectors.
  2. The second strategy would control for employment trends in the fast food sector but would be vulnerable to different time trends across the treatment and the control state.

- The DDD approach combines both strategies and computes 2 DD estimators:
  - the DD estimator using the non affected sector in the same state as control group:

$$\begin{aligned} DD_{NJ} &= \big(E(Y_{st}|s=\text{NJ},t=1,\text{affected})-E(Y_{st}|s=\text{NJ},t=0,\text{affected})\big) \\ &- \big(E(Y_{st}|s=\text{NJ},t=1,\text{unaffected})-E(Y_{st}|s=\text{NJ},t=0,\text{unaffected})\big) \end{aligned}$$

  - and, in order to control for different time trends in the affected versus the non affected sector:

$$\begin{aligned} DD_{Penn} &= \big(E(Y_{st}|s=\text{Penn},t=1,\text{affected})-E(Y_{st}|s=\text{Penn},t=0,\text{affected})\big) \\ &- \big(E(Y_{st}|s=\text{Penn},t=1,\text{unaffected})-E(Y_{st}|s=\text{Penn},t=0,\text{unaffected})\big). \end{aligned}$$

- The DDD estimator is given by the difference between the two DD estimators:

$$DDD = DD_{NJ} - DD_{Penn}$$

- By simply rearranging the expression above, we see that the DDD estimator could also be calculated as the difference between
  - the DD estimator using the affected sector in the control state as control group:

$$
\begin{aligned}
DD_{affected} \quad = \quad & \big(E(Y_{st}|s = \text{NJ}, t = 1, \text{affected}) - E(Y_{st}|s = \text{NJ}, t = 0, \text{affected})\big) \\
- \quad & \big(E(Y_{st}|s = \text{Penn}, t = 1, \text{affected}) - E(Y_{st}|s = \text{Penn}, t = 0, \text{affected})\big)
\end{aligned}
$$

  - and, in order to control for different time trends between the treatment and the control state:

$$
\begin{aligned}
DD_{nonaffected} \quad = \quad & \big(E(Y_{st}|s = \text{NJ}, t = 1, \text{unaffected}) - E(Y_{st}|s = \text{NJ}, t = 0, \text{unaffected})\big) \\
- \quad & \big(E(Y_{st}|s = \text{Penn}, t = 1, \text{unaffected}) - E(Y_{st}|s = \text{Penn}, t = 0, \text{unaffected})\big).
\end{aligned}
$$

- The DDD estimator is given by the difference between the two DD estimators:

$$DDD = DD_{affected} - DD_{nonaffected}$$

- Note that DDD is different from just adding a control group since now we define an affected and non affected group within each state:

$$\text{Additional Control group: } T, C \Rightarrow \ T, (C_1, C_2)$$

$$\text{DDD: } T, C \Rightarrow \ (T_{affected}, T_{non\ affected}), (C_{affected}, C_{non\ affected}).$$

- The DDD estimator thus controls at the same time for a state specific and a sector specific trend.
- It can also be implemented via a regression function. Let AF be a dummy equal to one if the sector is affected. Note that the following regression function contains eight parameters, one for each group (NJ affected, NJ non affected, Penn affected, Penn non affected) - time combination.

$$\begin{aligned} Y_{st} &= \alpha + \gamma_0 NJ + \gamma_1 AF + \gamma_2 (NJ \cdot AF) + \\ &\quad \lambda_0 d_1 + \lambda_1 (d_1 \cdot NJ) + \lambda_2 (d_1 \cdot AF) + \rho (d_1 \cdot NJ \cdot AF). \end{aligned}$$

- The regression function is thus saturated and equals the CEF. The coefficient ρ equals the DDD estimator which can be proved analogously to the DD case (do it!).

3. Extensions

The Difference in Differences in Differences Estimator

- A DDD estimator has been applied to examine tax incidence. As you know, who has to bear the economic burden of a tax is unrelated from who has to pay the tax.
- Tax incidence analysis is thus important whenever taxes have redistributive objectives.
- Gruber (1994)[1] studies state mandates for employer-provided health insurance to cover pregnancy costs. The research question is whether these mandates had an impact on wages (i.e. whether firms can pass the costs through to employees).
- In the 1970ies, expected cost for pregnancy were about $\$500$ per year for married women aged 20 to 40. State law changes to mandate coverage of pregnancy costs in 1976.
- Gruber has 3 treatment (IL, NJ and NY) and 5 nearby control states but was concerned about different wage time trends in the treated and the control states.
- He therefore employs a DDD approach, taking employees over 40 and single males as additional (non affected) control group.

[1] J. Gruber (1994), "The Incidence of Mandated Maternity Benefits", American Economic Review, 84(3), 622-641.

TABLE 3—DDD ESTIMATES OF THE IMPACT OF STATE MANDATES ON HOURLY WAGES

Location/year	Before law change	After law change	Time difference for location
A. *Treatment Individuals: Married Women, 20–40 Years Old:*			
Experimental states	1.547 (0.012) [1,400]	1.513 (0.012) [1,496]	−0.034 (0.017)
Nonexperimental states	1.369 (0.010) [1,480]	1.397 (0.010) [1,640]	0.028 (0.014)
Location difference at a point in time:	0.178 (0.016)	0.116 (0.015)	
Difference-in-difference:	−0.062 (0.022)		
B. *Control Group: Over 40 and Single Males 20–40:*			
Experimental states	1.759 (0.007) [5,624]	1.748 (0.007) [5,407]	−0.011 (0.010)
Nonexperimental states	1.630 (0.007) [4,959]	1.627 (0.007) [4,928]	−0.003 (0.010)
Location difference at a point in time:	0.129 (0.010)	0.121 (0.010)	
Difference-in-difference:	−0.008: (0.014)		
DDD:	**−0.054** **(0.026)**		

- The top panel of the table shows
  - that affected women in treatment states experienced a 3.4% decline in wages
  - while the wages of eligible women in control states increased by 2.8%.
  - The DD estimate is thus that there was a significant 6.2% relative fall in wages for women in treatment states.
- However, if the labor markets in treatment states experienced a distinct time trend in the observation period, the 2.8% increase in wages in the control group does not form a valid couterfactual.
- To check for different time trends in the treatment and control states, Gruber repeats the DD exercise for the group of non affected individuals in the bottom panel of the table.
  - This exercise reveals a 0.8% relative wage decline of non affected individuals in the treatment vs the control state.
  - This suggests that the treatment and the control state might indeed have experienced different time trends in the relevant period.
- The overall effect of the mandate is therefore a decline in affected women's wages by 5.4%.
- Annual women's earnings in the mid 70ies were around 10,000$. The results thus imply that wages decreased by roughly the cost of the mandate. Hence, affected women effectively bear the costs of the mandate.

- The regression formulation of DD also allows to include additional control variables. For example, you could estimate

$$Y_{st} = \gamma_s + \lambda d_1 + X'_{st}\beta + \rho(NJ \cdot d_1) + \eta_{st},$$

  where

  - γ_s is a separate dummy for each state; and
  - X_{st} are observable characteristics for each state (e.g. industry structure).
- In this specification
  - λ would capture an average time trend (across all states); and
  - the inclusion of X_{st} would allow for differences in the time trend across states based on observables X_{st}.
- Hence, the estimate of ρ would isolate the treatment effect from a general time trend and state specific trends due to observable differences.

- The DD estimator can also be used when several groups were treated with differing intensity.
- In the minimum wage example, there might be two reasons for that:
  1. The minimum wage changes could be different in each state.
  2. Even if the minimum wage changes are the same we might expect a different impact across states if, for example, states differ in the fraction of individuals earning minimum wages before the increase.
- In the former case we could use a continous minimum wage regressor w_{st} instead of the binary treatment D_{st}.
- In the latter case, a natural specification would be

$$Y_{st} = \gamma_s + \lambda d_1 + \rho(d_1 \cdot FA_s) + \eta_{st},$$

  where
  - FA_s is a variable measuring the fraction of individuals likely to be affected by the change in minimum wage laws; and
  - the interaction $d_1 \cdot FA_s$ is the treatment variable that accounts for differing treatment intensities.

- One advantage of more than two time periods is that it is possible to shed light on the validity of the common trend assumption.
- For example, in a reply to their critics, Card and Krueger (2000) provided the following figure. Why is it useful?

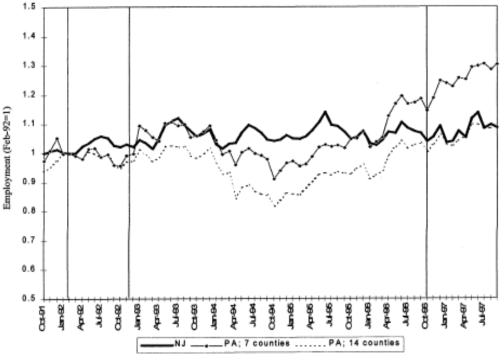

- If the common trend assumption does not hold exactly, a longer time horizon allows to control for different time trends across groups.
- One possibility would be to include linear, state specific time trends into the model and estimate

$$Y_{st} = \alpha + \gamma NJ + \lambda_t + \lambda_s \cdot t + \rho D_{st} + \eta_{st},$$

  where
  - D_{st} is the treatment indicator;
  - $\lambda_s \cdot t$ are separate linear time trends for each state $s \in \{Penn, NJ\}$; and
  - λ_t is a set of time dummies (one for each period) which capture any non linear component of the average time trend in all states.
- The estimate for the counterfactual time trend in NJ for the period surrounding the introduction of the minimum wage would now consist of
  - the linear trend specific to NJ; and
  - the relevant time dummy λ_t which captures the part of the employment development in Penn that is **not** caused by the linear Penn trend (e.g. a transitory component).
- The crucial identifying assumption now is that the transitory component is the same for both states.

- In addition, many periods offer the opportunity to examine lagged or anticipatory effects of treatment.
- Assume treatment D_{st} changes at different times in different groups and consider

$$Y_{st} = \gamma_s + \lambda_t + \sum_{\tau=0}^{m} \delta_{-\tau} D_{s,t-\tau} + \sum_{\tau=1}^{q} \delta_{+\tau} D_{s,t+\tau} + \eta_{st},$$

  where γ_s and λ_t are again full sets of state and time dummies.
- $D_{s,t-\tau}$ equals one if the treatment occurred exactly τ periods in the past.
- The term $\sum_{\tau=0}^{m} \delta_{-\tau} D_{s,t-\tau}$ captures lagged effects of the treatment.
  - It would capture either a growing treatment effect over time or (alternatively) a fading out.
- Analogously ,$D_{s,t+\tau} = 1$ if treatment occurs exactly τ periods in the future.
- The term $\sum_{\tau=1}^{q} \delta_{+\tau} D_{s,t+\tau}$ captures anticipatory effects of the treatment.
  - For example, the mere announcement of a higher future minimum wage might have an impact on employment in affected sectors.

- The inclusion of lagged and anticipatory effects allows running a causality test, more precisely, a test for **Granger Causality**.
- The idea of Granger Causality is to check whether causes happen before consequences, i.e. whether (treatment) effects follow treatment and not vice versa.
- If the anticipatory effects $(\delta_{+1}, \ldots, \delta_{+q})$ are different from zero, future treatment would predict current outcomes, suggesting that causality also runs from outcomes to treatment.
- As we have already seen in the chapter on OLS, under reverse causality coefficients fail to have a causal interpretation.
- The reason is that if $(\delta_{+1}, \ldots, \delta_{+q})$ are different from zero, there are systematic employment trends prior to changes in treatment status which make it hard to find suitable control groups.

1. The most important condition for the validity of DD is the common trend assumption. We have just seen, how data over a longer time horizon can be used to assess (or weaken in case of state specific trends) this assumption.
2. We have said in the beginning that DD can be applied in repeated cross sections as well since all we need are group averages. A Caveat:
    - The composition of treatment and control groups must not change! If it does, the group "fixed" effect changes over time and can no longer be differenced out.
    - Example: A higher minimum wage induces more able and motivated individuals to work in the fast food industry which makes it more attractive to hire more workers.
    - As long as the composition changes along observable dimensions, one can control for it.
    - However, if observable group characteristics change by a large amount, we might suspect the same for unobservable characteristics as well.
    - If group composition changes over time it is thus a good idea to examine observable group characteristics pre- and post-treatment in practice (see e.g. Gruber (1994), table 2).
    - It might also help to examine observable characteristics across groups. If those are similiar one can be more confident that the time trend of both groups is similiar as well.

- DD also fails to uncover the causal effect if treatment and control group differ in their idiosyncratic (transitory) shocks prior to treatment. Formally, if the transitory component η_{ist} of the error

$$\epsilon_{ist} = \gamma_s + \lambda_t + \eta_{ist},$$

  differs between the treatment and the control group, the DD estimator has no causal interpretation.
- An Example is Ashenfelter's famous study: Evaluation of a job training program where participants entered the program (or were selected) when earnings were particularly low.
- That is, there is a dip in earnings prior to treatment but we would expect earnings to recover anyway (since the dip is transitory) even without the program.
- Ashenfelter's dip would correspond to a different expected value of η_{ist} for the treatment and the control group in the period before treatment.

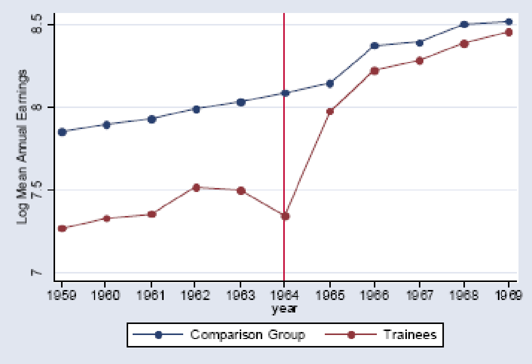

- What is the problem caused by the dip?
- Ashenfelter's Dip can often be detected graphically. If you see a dip, dynamic models (see below) are more appropriate.