Case Studies in Microeconomic Evaluation

Data and methods for learning what works

Difference-in-differences

Competence Centre on Microeconomic Evaluation (CC-ME)
Joint Research Centre
2nd Dec 2020

A Brief Introduction to Difference-in-Differences

Evaluating the effects of an intervention

What would have happened to treated units in absence of the treatment?

- Our goal is to find the **Average Treatment on the Treated (ATT)**
- Ideally, we would like to observe *two parallel worlds*

Problem: We can observe only one of the two parallel worlds!

Evaluating the effects of an intervention

- Consider outcome Y
- We have 2 time periods
  - Time $t = 0$: before the intervention
  - Time $t = 1$: after the intervention
- We can identify 2 groups
  - Treatment group T: receives the intervention
  - Control group C: does not receive the intervention

Feasible but problematic solutions (1)

"Simple differences" estimator

Compares Treated units T and Non-Treated units C in post-intervention period ($t = 1$)

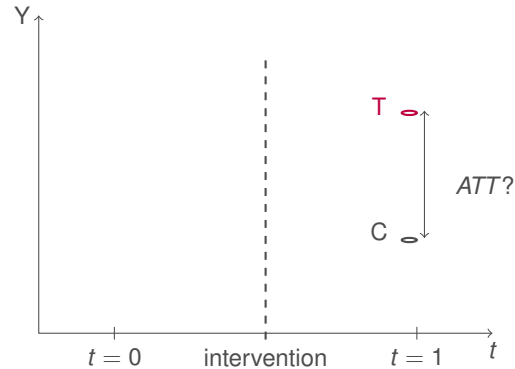

Feasible but problematic solutions (1)

"Simple differences" estimator
Compares Treated units T and Non-Treated units C in post-intervention period ($t = 1$)

Problem: unobserved differences between treated and non-treated units that are correlated with outcomes influence the estimation of the effect

- Simple difference ignores pre-existing heterogeneity between T and C groups

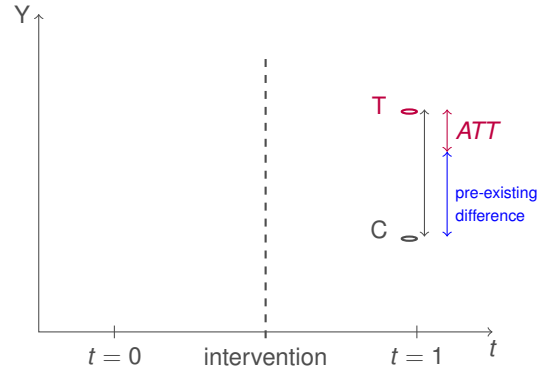

Feasible but problematic solutions (2)

"Before-After" estimator
compares outcomes of treated units T before and after intervention, i.e. $t = 0$ vs $t = 1$

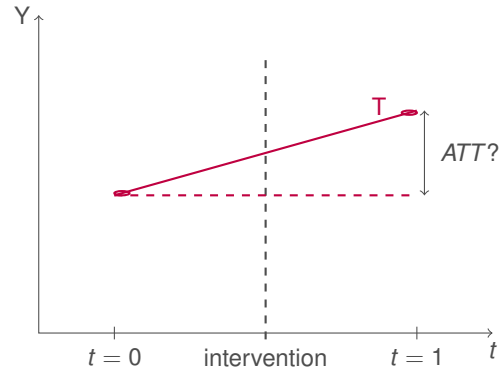

Feasible but problematic solutions (2)

"Before-After" estimator
compares outcomes of treated units T before and after intervention, i.e. $t = 0$ vs $t = 1$

Increasing time-trend
causes the effect of the intervention to be overestimated

- Before-after comparison ignores time-varying factors

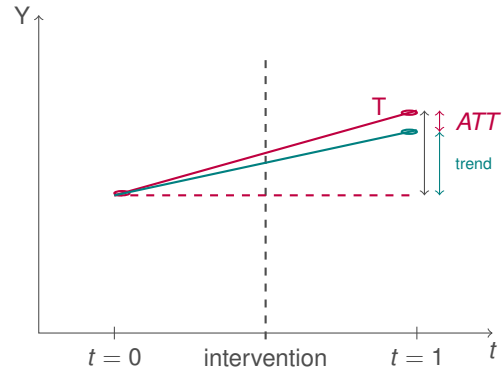

Feasible but problematic solutions (2)

"Before-After" estimator
compares outcomes of treated units T before and after intervention, i.e. $t = 0$ vs $t = 1$

Decreasing time-trend
causes the effect of the intervention to be underestimated

- Before-after comparison ignores time-varying factors

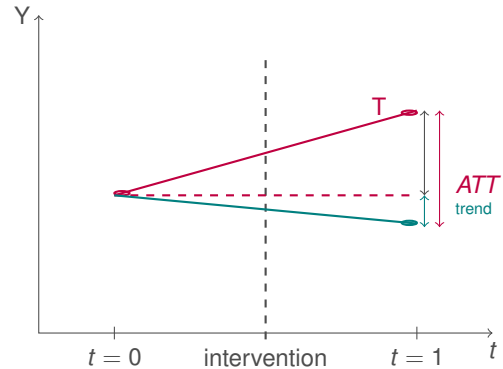

What then?

Combine the two: Difference-in-Differences (DiD)

- Take the mean value of each group's outcome **before** and **after** the intervention
- Compute the 'difference-in-differences' of the means

	Treatment Group (T)	Control Group (C)	Δ
Pre ($t = 0$)	T_0	C_0	
Post ($t = 1$)	T_1	C_1	
Change over time	$T_1 - T_0$	$C_1 - C_0$	$T_1 - T_0 - (C_1 - C_0)$ or, equivalently $T_1 - C_1 - (T_0 - C_0)$

Difference-in-Differences

Definition
Difference-in-differences compares the changes in outcomes over time between units that are subject to the intervention (the **treatment group**) and units that are not (the **comparison or control group**).

This allows to correct for:

- pre-existing time-invariant differences across groups, and
- common time-trends

Can also include covariates, i.e. the effect can be netted out of other factors

Difference-in-Differences

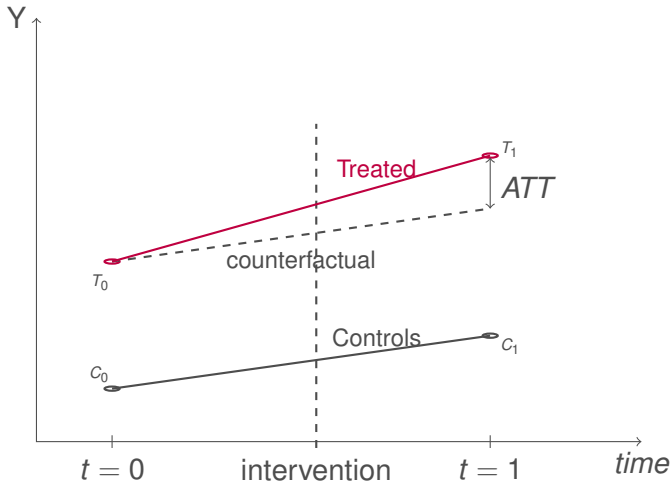

It’s all about assumptions

DiD is a wonderful tool,
but crucially depends on the **credibility of the assumptions** in the specific case

- We are “creating” a parallel world, it needs to make sense!
- The fundamental assumption is the **common trend**
  - Visual inspection of the evolution of *Y* in the two groups over time helps
  - Relatedly: the more periods you have (especially in the “pre” period) ... the better!

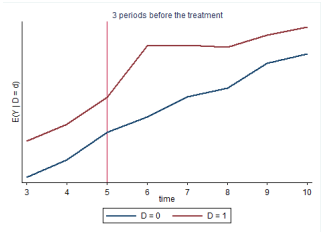

It’s all about assumptions

DiD is a wonderful tool,
but crucially depends on the **credibility of the assumptions** in the specific case

- We are “creating” a parallel world, it needs to make sense!
- The fundamental assumption is the **common trend**
  - Visual inspection of the evolution of *Y* in the two groups over time helps
  - Relatedly: the more periods you have (especially in the “pre” period) ... the better!

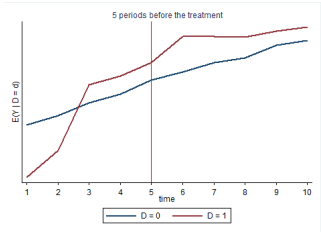

It’s all about assumptions

DiD is a wonderful tool,
but crucially depends on the **credibility of the assumptions** in the specific case

- We are “creating” a parallel world, it needs to make sense!
- The fundamental assumption is the **common trend**
  - Visual inspection of the evolution of Y in the two groups over time helps
  - Relatedly: the more periods you have (especially in the “pre” period) ... the better!

- Additionally, **no other change should occur** that systematically affects either group (treated or control)

It’s all about assumptions

Unfortunately these assumptions cannot be formally tested, but...

- **Event-study analysis** can provide some informal testing
  - Tells you whether treated units behave differently from control units at each point in time (especially before the treatment)
  - Evidence of significant differences before the treatment are bad news for the common trend assumption

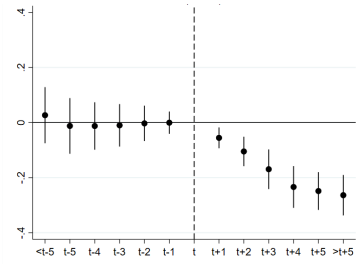

It’s all about assumptions

Unfortunately these assumptions cannot be formally tested, but...

- **Event-study analysis** can provide some informal testing
  - Tells you whether treated units behave differently from control units at each point in time (especially before the treatment)
  - Evidence of significant differences before the treatment are bad news for the common trend assumption
- **Unit-specific trend**
  - More time periods are required
  - We estimate a unit-specific trend (linear quadratic etc.)
  - Similar results (with vs without) are reassuring
- **Placebo tests**
  - “Move” artificially the intervention in time
  - Check the effect on similar but unaffected outcomes
  - Check the effect on a fake treatment group

Standard DiD

- A relatively identifiable group (*T*) receives the intervention ("treatment") at time *y*
- Need to find a reasonable control group (*C*)
  - Control units expected to behave similarly to treated units in the absence of the treatment
  - Control units not subject to any type of intervention in the same period
- Need to gather data on both *T* and *C* units, before and after the treatment

Matching DiD

- Sometimes you can select the control group (*C*) using a “matching procedure”

- Matching methods allow identifying the set of non-treated units that look more similar to the treated ones, based on the available observable characteristics
- The matched non-treated units become the control group
- A good match for each treated requires a large and complete set of data

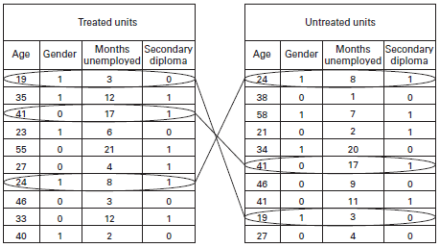

Treated units

Age	Gender	Months unemployed	Secondary diploma
19	1	3	0
35	1	12	1
41	0	17	1
23	1	6	0
55	0	21	1
27	0	4	1
24	1	8	1
46	0	3	0
33	0	12	1
40	1	2	0

Untreated units

Age	Gender	Months unemployed	Secondary diploma
24	1	8	1
38	0	1	0
58	1	7	1
21	0	2	1
34	1	20	0
41	0	17	1
46	0	9	0
41	0	11	1
19	1	3	0
27	0	4	0

Staggered DiD

- You may have that everyone is eventually treated (*T*), i.e. receives the intervention ("treatment")
- As long as the treatment is staggered over time, you can identify the control group (*C*)
  - (Groups of) units are treated at different points in time
  - When a unit becomes treated, their control will be the units who are not yet treated
- Here, you may have "always treated", "never treated" and "switchers"
- Need to gather data on a sufficient number of "switches"

Staggered DiD: Alignment around zero

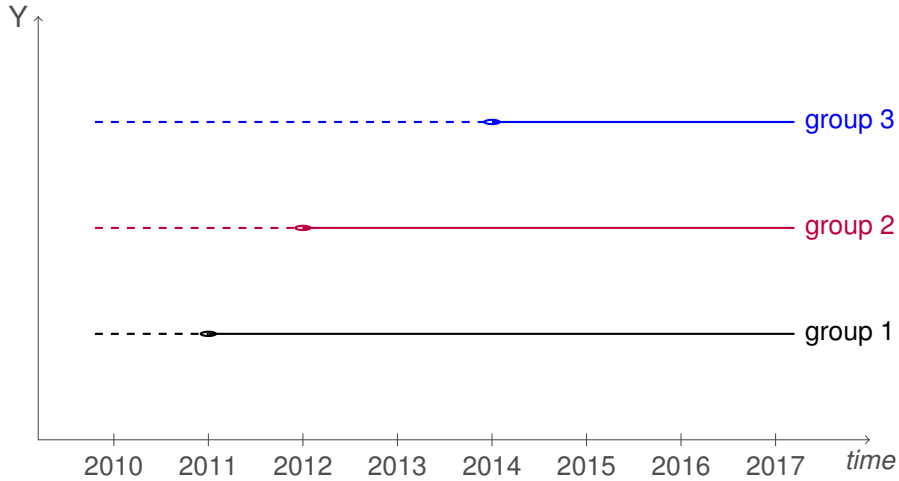

Staggered DiD: Alignment around zero

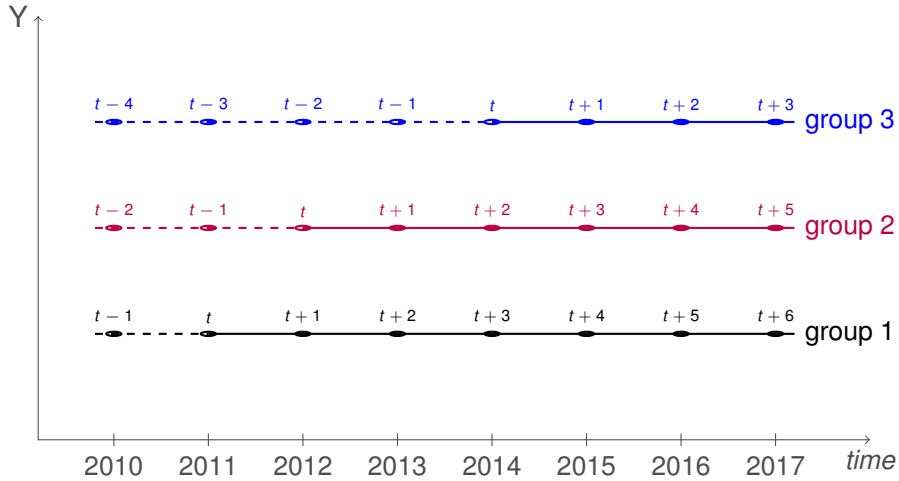

Staggered DiD: Alignment around zero

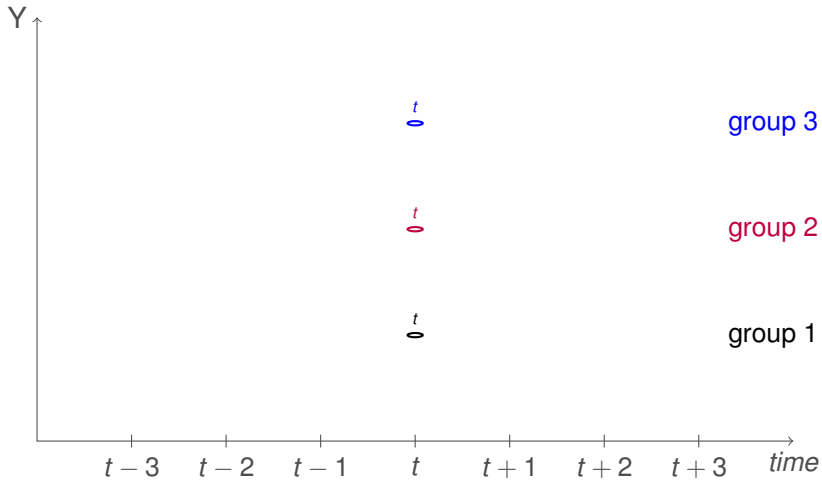

Staggered DiD: Alignment around zero

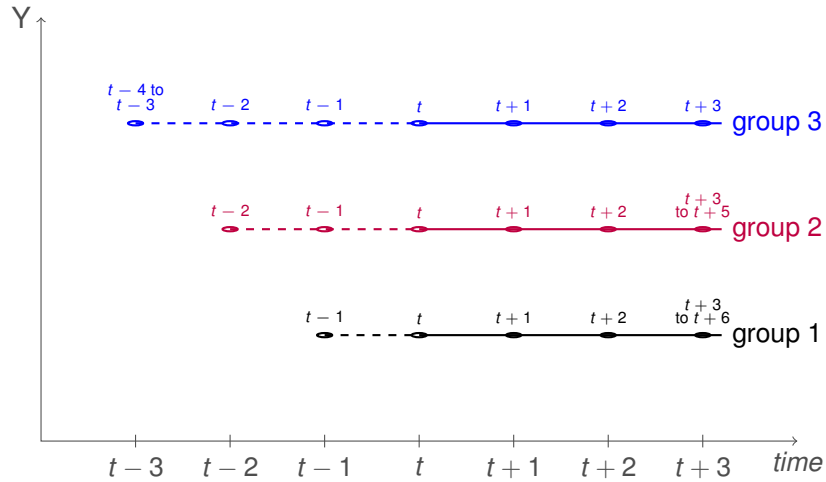

Indirect Cost Compensation under the EU ETS: A Firm-level Analysis

Antonella Ferrara and Ludovica Giua

European Commission

EU Emission Trading System

- Largest trading system of emission allowances worldwide
- Creates a carbon price signal based on a cap-and-trade mechanism:
  - targets energy-intensive installations and aircraft operators
  - businesses must buy CO_2 certificates equivalent to their industrial emissions (**direct costs**)
  - firms incur in an additional cost for the electricity they consume (**indirect costs**) because also their energy suppliers are subject to direct costs
- Drawback: risk of **carbon leakage**
- The ETS Directive provides for **compensation** of direct and indirect costs

EU Emission Trading System

Widespread consensus on the effectiveness of the ETS in abating emissions

But what about firm performance? Evidence is ambiguous:

- negative (Abrell et al., 2011; Wagner et al., 2013; Laukkanen et al., 2019)
- positive (Klemetsen et al., 2020; Chan et al., 2013)
- non significant (Abrell et al., 2011; Commins et al., 2011; Martin et al., 2014; Flues and Lutz, 2015)

Data

- Records on **beneficiaries of indirect cost compensation**
  - DE, NL, ES, BE, FI, UK
  - Info on name of firm and amount received
- **Orbis database** (Bureau Van Dijk)
  - Competitiveness: labour productivity and assets per employee
  - Performance: operating revenues (turnover), number of employees and total assets

The two sources are matched via **probabilistic matching** on company names

We take all 4-digit NACE-coded aided sectors in aided and non-aided countries

- 80% of the original pool of beneficiaries correctly matched to Orbis
- Final sample: unbalanced panel of **3,706 firms**, of which **319** are funded under State Aid measures for ETS indirect cost compensation

Data: number of firms

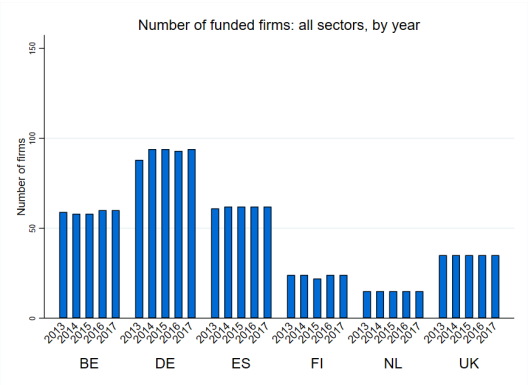

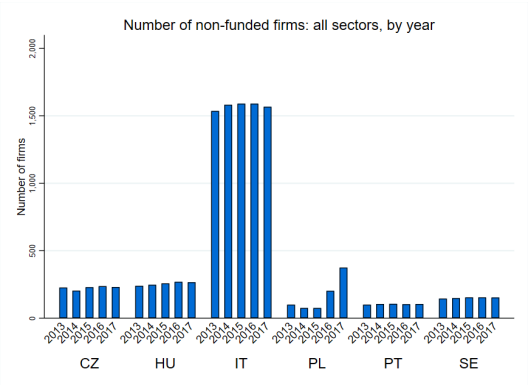

European Commission

Data: number of firms

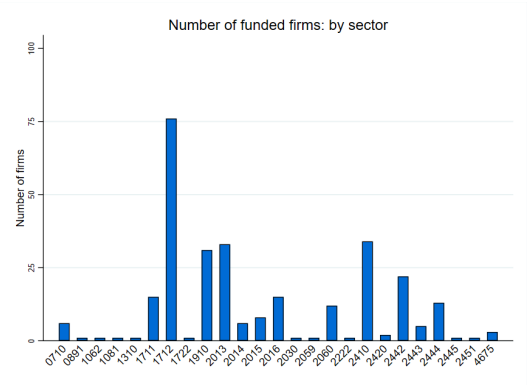

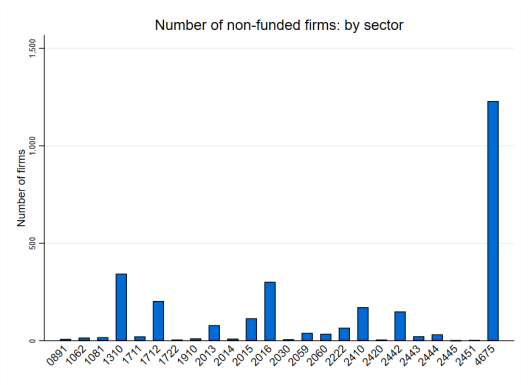

Data: amount of subsidy

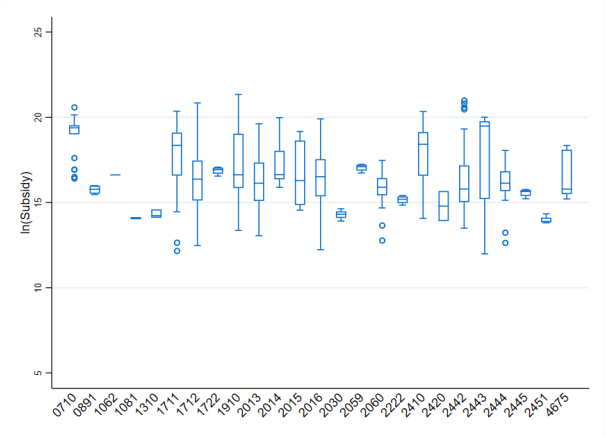

Model

We apply a Difference-in-Differences approach:

- **Treated group**: firms in aided sectors in aided countries (DE, NL, ES, BE, FI, UK)
- **Control group**: firms in aided sectors in non-aided countries (CZ, HU, IT, PL, PT, SE)
- **Time span**: 2009-2017; **Treatment** starts in year 2013 (FI in 2016, Wallonia in 2017)

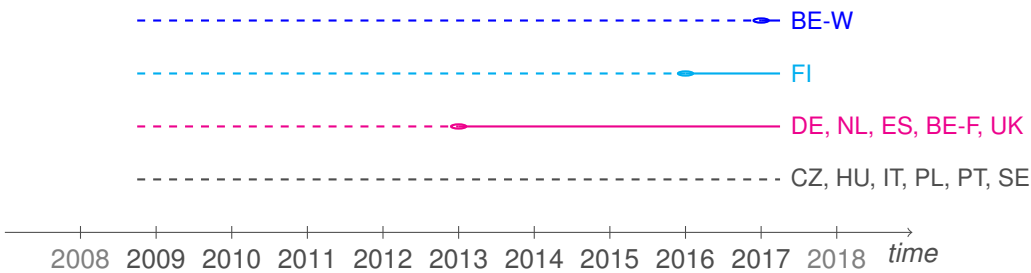

Model

We apply a Difference-in-Differences approach:

- **Treated group**: firms in aided sectors in aided countries (DE, NL, ES, BE, FI, UK)
- **Control group**: firms in aided sectors in non-aided countries (CZ, HU, IT, PL, PT, SE)
- **Time span**: 2009-2017; **Treatment** starts in year 2013 (FI in 2016, Wallonia in 2017)

- We account for time, firm and sector-year unobservable characteristics and for time-varying country-specific variables (GDP per capita, debit-to-GDP ratio, etc.)

- We consider two measures of *treatment*
  - **Aid**: takes value one if the firm is deemed to receive indirect cost compensation, because it operates in a country where the sector is eligible to funding in that year, and value zero otherwise
  - **Subsidy**: (log of) amount received by the firm in a given year

Effect of the aid on the extensive margin

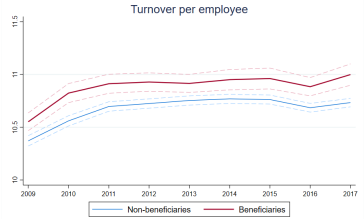

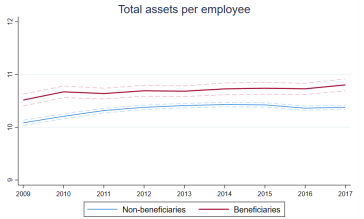

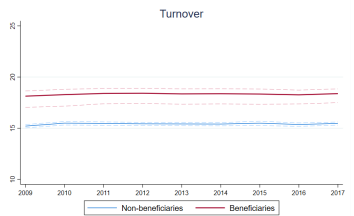

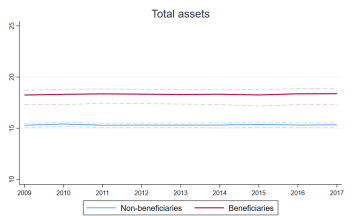

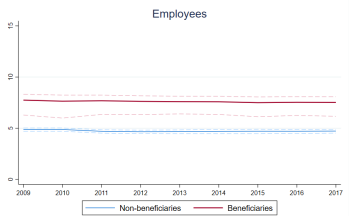

European Commission

Effect of the aid on the extensive margin

Dependent variable	(1) Turnover per employee	(2) Total assets per employee	(3) Turnover	(4) Total assets	(5) Employees
Aid	-0.012	-0.020	-0.056**	-0.063**	-0.044**
	(0.024)	(0.030)	(0.025)	(0.027)	(0.020)
Observations	23,277	23,277	23,277	23,277	23,277
R-squared	0.924	0.919	0.975	0.981	0.974
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
NACE-Year FE	✓	✓	✓	✓	✓
Country-specific controls	✓	✓	✓	✓	✓

Effect of the aid on the extensive margin

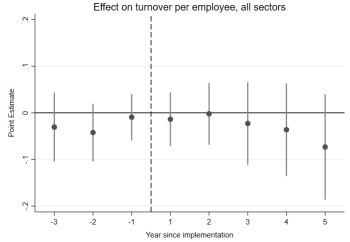

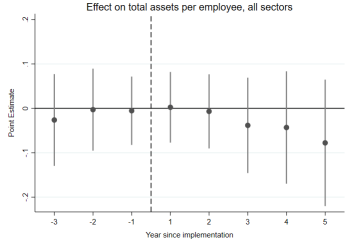

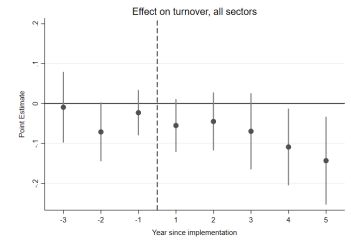

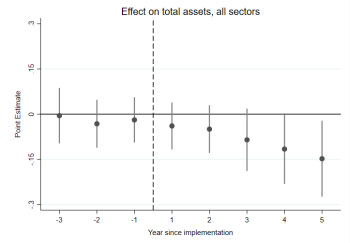

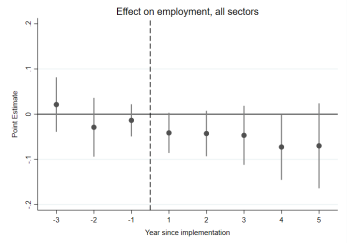

European Commission

Effect of the aid on the extensive margin

- Robust to exclusion of **countries** or sectors one-by-one

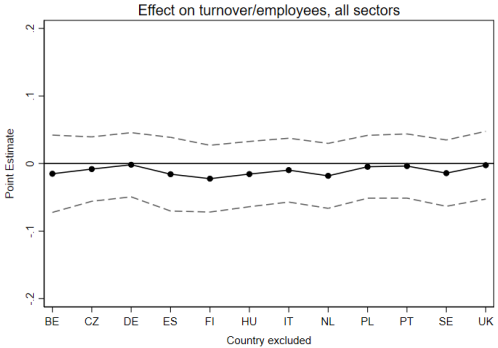

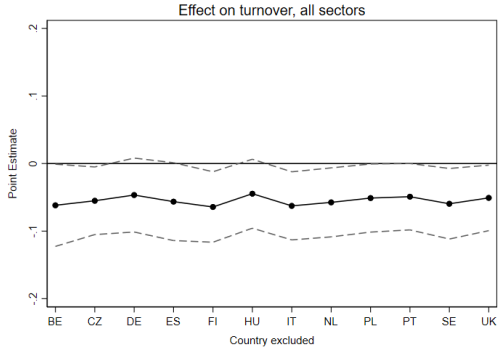

European Commission

Effect of the aid on the extensive margin

- Robust to exclusion of countries or **sectors** one-by-one

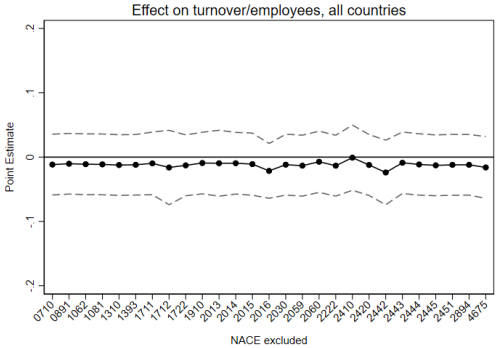

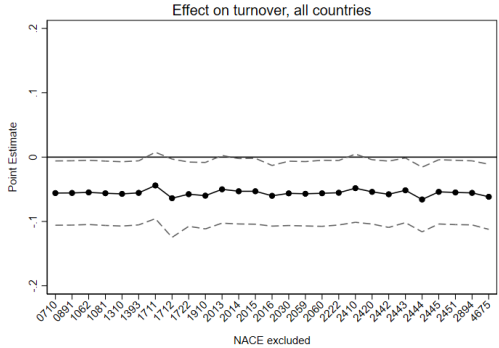

European Commission

Effect of the aid on the extensive margin

- Robust to exclusion of countries or sectors one-by-one
- Robust to considering firms belonging to same group

Are firms receiving compensation experiencing a worse performance compared to those who do not receive funding?

- Firms in the treated group are substantially larger than controls
- There are elements that enter the production function which we cannot observe, e.g. the cost of the inputs (electricity)
  - It could be a problem if these factors changed differently for beneficiaries and for firms in the comparison group!

Effect of the subsidy on the intensive margin

Dependent variable	(1) Turnover per employee	(2) Total assets per employee	(3) Turnover	(4) Total assets	(5) Employees
ln(subsidy)	0.025	0.026	0.098**	0.098**	0.073**
	(0.048)	(0.049)	(0.047)	(0.048)	(0.036)
Observations	617	617	617	617	617
R-squared	0.959	0.957	0.989	0.985	0.990
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
NACE-Year FE	✓	✓	✓	✓	✓
Country-specific controls	✓	✓	✓	✓	✓

Conclusions

- First analysis on the impact of the ETS indirect cost compensation at firm level and EU-wide coverage
- Results suggest that the aid has not had a significant impact on average relative competitiveness, measured as turnover per worker and the value of total assets per employees
- But beneficiaries seem to perform worse in terms of turnover, value of total assets and number of employees
- This might be due to systematic differences across aided and non-aided countries such as changes in electricity costs
- When we focus only on beneficiaries for which the amount of the compensation is known, we find that for each 1% increase in the amount of the subsidy received (i.e. around 1,000 EUR), firms expand their turnover and their assets value by 0.01%, and their workforce by 0.07%

Lessons learned

- Estimating the impact of the ETS indirect cost compensation is not straightforward
- Lack of accurate data is a big issue
- The EC could take advantage of reporting obligations
  - Adopt standardised forms
  - Make it easier to link data
  - Exploit the existence of other data sources (e.g. E-PRTR)

Thank you