

Linear D-in-D with linear covariate adjustment, heterogenous treatment effects and compositional change

Arek Szydlowski*

October 23, 2026

1 Setup

Assume cross-sectional data in two periods $T \in \{0, 1\}$. Consider a standard difference-in-differences regression:

$$Y = \alpha + \gamma D + \lambda T + \delta(D \times T) + X'\beta + \varepsilon \quad (1)$$

where $D \in \{0, 1\}$ denotes treatment status and X is a vector of control covariates. The covariates are usually included for two related reasons: 1) the parallel trends assumption is believed to hold only conditional on covariates; 2) to control for changes in the composition of covariates between the cross-sections among treated and control groups (compositional change).

Let $Y_T(1), Y_T(0)$ denote a potential treated and untreated outcome in period T . We are interested in learning about the average treatment effect on the treated in the second period (ATT):

$$\tau = \mathbb{E}[Y_1(1) - Y_1(0)|D = 1, T = 1].$$

Often one assumes constant treatment effect across time so $\tau = \mathbb{E}[Y_1(1) - Y_1(0)|D = 1]$ as well.

2 What does δ identify with heterogeneous treatment effects and composition change?

It is well known that with homogeneous treatment effects (under conditional parallel trends), OLS estimator of δ in (1) gives a consistent estimator of τ even if we have compositional change in covariates. This note explains what happens if treatment effects are heterogeneous.

To analyse this case assume correct linear specification of the outcome model:

*School of Economics, Politics and IR; University of Kent, LASE University Group, Canterbury, UK

Assumption 1 (Conditional linear DiD representation). *There exist finite dimensional parameters $\alpha_0, \gamma_0, \lambda_0$, and a treatment-effect function $\tau(x)$ such that*

$$\mathbb{E}[Y \mid D = d, T = t, X = x] = \alpha_0 + \gamma_0 d + \lambda_0 t + \tau(x)(d \times t) + x' \beta_0. \quad (2)$$

Thus, we abstract from bias coming from misspecification of the linear model. See e.g. Deb et al. (2026) for recent methods to flexibly control for covariates in D-in-D regressions and Sant'Anna and Xu (2026) for methods to deal with compositional changes in this more general setup.

Define $\mu_{dt} = \mathbb{E}[X \mid D = d, T = t]$ and $\Delta_X = (\mu_{11} - \mu_{10}) - (\mu_{01} - \mu_{00})$. A straightforward calculation shows that:

$$\delta = \mathbb{E}[\tau(X) \mid D = 1, T = 1] + \Delta_X'(\beta_0 - \beta) = \tau + \Delta_X'(\beta_0 - \beta), \quad (3)$$

under Assumption 1. This means that even under the correctly specified linear outcome model including covariates linearly does not fully deal with compositional change unless the treatment effects are homogenous ($\beta_0 = \beta$) or we have **parallel trends in covariates** ($\Delta_X = 0$).

Checking parallel trends in covariates is straightforward in applications. Our results suggest that such check should be run in D-in-D applications with compositional changes. If parallel trends in covariates cannot be rejected, then the estimated ATT is robust to heterogeneous treatment effects. If not, one can resort to more advanced methods to deal with compositional changes like Sant'Anna and Xu (2026).

Remark 1. *One can relate the observation above to results in Słoczyński (2022). Assuming $\mathbb{P}(D = d, T = t) > 0$ for all (d, t) , we can derive that:*

$$\delta = \frac{\mathbb{E}[r_{11}(X)\tau(X) \mid D = 1, T = 1]}{\mathbb{E}[r_{11}(X) \mid D = 1, T = 1]} \quad (4)$$

where:

$$\begin{aligned} r_{11}(x) &= \frac{1}{H} \left[\frac{1}{s_{11}} - (x - \Pi(X \mid 1, 1))' \Sigma_X^{-1} \Delta_x \right] \\ s_{dt} &= \mathbb{P}(D = d, T = t), \\ H &= \sum_d \sum_t \frac{1}{s_{dt}}, \\ \Sigma_X &= \mathbb{E}[(X - \Pi(X \mid D, T))(X - \Pi(X \mid D, T))'] \end{aligned}$$

and $\Pi(X \mid D, T)$ is a linear projection of X onto $(1, D, T)$.

With $\Delta_X = 0$, the weights become $r_{11}(X) = \frac{1}{H} \frac{1}{s_{11}}$, hence are constant, and we obtain $\delta = \mathbb{E}[\tau(X) \mid D = 1, T = 1]$. Intuitively, with parallel trends in covariates the linear projection of $D \times T$ onto $(1, D, T, X)$ does not depend on X (coefficient on X in this projection is equal to $\frac{1}{H} \Sigma_X^{-1} \Delta_X$) and, thus, variance of this projection conditional on $D \times T$ is zero and we are in the degenerate case

discussed extensively in Słoczyński (2018).

Remark 2. Observation above accords with the observations made in the recent literature that full stationarity of (D, X) is not necessary for D -in- D methods to pin down the relevant treatment effect without bias. Sant’Anna and Xu (2026) show that in semiparametric D -in- D compositional change would be inconsequential if $\mathbb{E}[\tau(X)|D = 1] = \mathbb{E}[\tau(X)|D = 1, T = 1]$, i.e. when ATT does not vary with time. Our result provides an alternative condition for a simple linear D -in- D .

2.1 Adjusted outcomes regression

One may wonder if re-weighting the outcomes first, in order to correct for compositional change, would avoid the bias present above. Define $q_d(x) = \mathbb{P}(T = 1|D = d, X = x)$ and consider the new outcome $\tilde{Y} = \left[\frac{T}{q_D(X)} - \frac{1-T}{1-q_D(X)} \right] Y$ (see e.g. Słoczyński, Uysal, and Wooldridge (2025)). Suppose that we run a simple D-in-D of the form:

$$\tilde{Y} = \tilde{\alpha} + \tilde{\delta}D + \varepsilon. \quad (5)$$

Under Assumption 1 we have:

$$\mathbb{E}[\tilde{Y}|D = d, X = x] = \mathbb{E}[Y|D = d, T = 1, X = x] - \mathbb{E}[Y|D = d, T = 0, X = x] = \lambda_0 + \tau(x)d \quad (6)$$

which implies:

$$\tilde{\delta} = \mathbb{E}[\tilde{Y}|D = 1] - \mathbb{E}[\tilde{Y}|D = 0] = \mathbb{E}[\tau(x)|D = 1] \equiv \tau_1$$

and, assuming $0 < q_1(x) < 1$ for all x , by Bayes rule:

$$\tilde{\delta} = \frac{\mathbb{E}[q_1^{-1}(x)\tau(x)|D = 1, T = 1]}{\mathbb{E}[q_1^{-1}(x)|D = 1, T = 1]}$$

This shows that correcting for compositional change by re-weighting works. It provides unbiased estimate of the average treatment effect on the treated, τ_1 , and, even, if the target is the average treatment effect on the treated in the second period, τ , it provides a weighted average of the individual treatment effects with positive weights, which are lower for X ’s that are over-represented in the second period (i.e. $q_1(x)$ high).

A Derivation of (3)

For simplicity let $\mathbb{E}[\cdot|D = d, T = t]$ be denoted by $\mathbb{E}[\cdot|dt]$. From OLS we get:

$$E[Y|dt] = a_{dt} + \mu'_{dt}\beta \quad (7)$$

where a_{dt} is a dt -specific regression intercept, e.g. $a_{11} = \alpha + \gamma + \lambda + \delta$. Now using the fact that $\delta = a_{11} - a_{10} - a_{01} + a_{00}$, we obtain:

$$\delta = \mathbb{E}[Y|11] - \mathbb{E}[Y|10] - \mathbb{E}[Y|01] + \mathbb{E}[Y|00] - \Delta'_X \beta \quad (8)$$

Under Assumption 1: $E[Y|11] - E[Y|10] - E[Y|01] + E[Y|00] = E[\tau(X)|11] + \Delta'_X \beta_0$, and substituting this into equation (8) we obtain the main result in (3).

B Derivation of (4)

First residualise $W = DT$ on $A = (1, D, T)$. Because A has three linearly independent cell-level regressors over four (D, T) cells, the residual is one-dimensional. Orthogonality to $1, D, T$ implies that, for a scalar c ,

$$u_{00} = \frac{c}{s_{00}}, \quad u_{10} = -\frac{c}{s_{10}}, \quad u_{01} = -\frac{c}{s_{01}}, \quad u_{11} = \frac{c}{s_{11}}.$$

Since $u = M_A W$ is a projection residual, $\mathbb{E}[uW] = \mathbb{E}[u^2]$. The left-hand side equals c , while the right-hand side equals $c^2 H$. Hence $c = 1/H$ and

$$u_{dt} = \frac{(-1)^{d+t}}{s_{dt} H}. \quad (9)$$

Now apply Frisch-Waugh-Lovell theorem to the auxiliary regression of W on (A, X) . Let $\tilde{X} = M_A X$. The coefficient on $\tilde{X}$ is

$$\pi_X = \Sigma_{\tilde{X}}^{-1} \mathbb{E}[\tilde{X} u].$$

Because u is orthogonal to every linear combination of A , replacing $\tilde{X}$ by X inside $\mathbb{E}[\tilde{X} u]$ does not change the expectation. Using (9),

$$\begin{aligned} \mathbb{E}[X u] &= \sum_{d,t} s_{dt} u_{dt} \mu_{dt} \\ &= \frac{1}{H} (\mu_{00} - \mu_{10} - \mu_{01} + \mu_{11}) \\ &= \frac{1}{H} \Delta \mu_X. \end{aligned}$$

The final residual is

$$r_R = u - \tilde{X}' \pi_X.$$

In the treated-post cell, $u_{11} = 1/(s_{11} H)$ and $\tilde{X} = x - \mu_{11}^A$, yielding (4).

References

- Deb, Partha, Edward C. Norton, Jeffrey M. Wooldridge, and Jeffrey E. Zabel. 2026. “A Flexible, Heterogeneous Treatment Effects Difference-in-Differences Estimator for Repeated Cross-Sections.” NBER Working Paper No. 33026, revised April 2026. doi:10.3386/w33026.
- Sant’Anna, Pedro H. C., and Qi Xu. 2026. “Difference-in-Differences with Compositional Changes.” *Journal of Econometrics* 253: 106147. doi:10.1016/j.jeconom.2025.106147.
- Słoczyński, Tymon. 2018. “A General Weighted Average Representation of the Ordinary and Two-Stage Least Squares Estimands.” IZA Discussion Paper No. 11866. See especially Theorem 4 and the section “Difference-In-Differences.”
- Słoczyński, Tymon. 2022. “Interpreting OLS Estimands When Treatment Effects Are Heterogeneous: Smaller Groups Get Larger Weights.” *Review of Economics and Statistics* 104(3): 501–509.
- Słoczyński, Tymon, Derya Uysal, and Jeffrey M. Wooldridge. 2025. “Covariate Balancing and the Equivalence of Weighting and Doubly Robust Estimators of Average Treatment Effects.” IZA Discussion Paper No. 18147.