Spillover Effects in Empirical Corporate Finance: Choosing the Proxy for the Treatment Intensity

Fabiana Gómez
David Pacini

Discussion Paper 22/766

16 May 2022

School of Economics

University of Bristol
Priory Road Complex
Bristol
BS8 1TU
United Kingdom

University of BRISTOL
School of Economics

Spillover Effects in Empirical Corporate Finance: Choosing the Proxy for the Treatment Intensity

Fabiana Gómez* David Pacini†

University of Bristol *University of Bristol*

May 16, 2022

Abstract. The existing literature indicates that spillovers lead to a complicated bias in the estimation of treatment effects in empirical corporate finance. We show that, under simple random treatment assignment, such a complicated bias is simplified if the proxy chosen for the group-level treatment intensity is the leave-one-out average treatment. This choice brings two advantages: first, it facilitates the diagnosis of the bias and, second, it facilitates the interpretation of the average spillover effect on the treated. These two advantages justify the use of the leave-one-out average treatment as the preferred proxy for the treatment intensity. We illustrate these advantages in the context of measuring the effect of credit supply contractions on firms' employment decisions.

JEL Classification: C21, G30

Keywords: Spillover; Average Direct Effects; Average Spillover Effects

*School of Accounting and Finance, 15/19 Tyndall's Park Road, BS8 1PQ, Bristol, UK. E-mail address: Fabiana.Gomez@bristol.ac.uk

†School of Economics, 8 Woodland Road, BS8 1TN, Bristol, UK. Email address: David.Pacini@bristol.ac.uk

1. Introduction

It is well-known that spillovers arise in corporate finance through firm competition and/or agglomeration. Typically, a firm-level outcome, such as sales or investment, depends not only on the firms' own treatment status, but also on the fraction of firms treated in the same industry and/or region. Berg, Reisinger and Streitz (2021), BRS21 hereon, develop a model to measure spillover effects using the group-level treatment intensity of an intervention. The group-level treatment intensity can be approximated by either the *group-level average treatment* or by the *leave-one-out treatment average*.[1] Little is known about which of the two proxies one should use. This paper compares the implications of using each of these two proxies. This comparison yields, first, a justification for using the *leave-one-out treatment average* proxy and, second, it provides empirical researchers with a simple test for the diagnosis of the bias induced by spillovers.

We show that, when the treatment is randomly assigned, choosing the leave-one-out average treatment proxy has two advantages. First, this choice simplify the formula for bias arising from the spillover effects, thereby, facilitating its diagnosis. Second, it clarifies the definition of the average spillover effect on the treated, thereby, facilitating its interpretation. These advantages justify the use of the leave-one-out average treatment as the preferred proxy, for both theoretical developments and in empirical applications.

We propose a straightforward statistical test for the diagnosis of the bias induced by the spillover effects. It consists on performing a heteroskedatic-robust Wald test for the null hypothesis of equal average spillover effects on the treated and untreated groups versus the alternative hypothesis of different average spillover effects. If this null hypothesis is rejected, the ordinary least square estimator of the treatment

[1] There is a growing literature considering the presence of spillover effects between treated and untreated firms in empirical corporate finance (see e.g. Huber, 2018; Beck, Da-Rocha-Lopes and Silva, 2021; Breuer, Hombach and Muller, 2019; Doerr, 2021; Gopalakrishnan, Jacob and Mohapatra, 2021). These papers use either the *group-level average treatment* or the *leave-one-out average treatment* as a proxy for the treatment intensity. See Section 4 for a summary of these applications.

effect in the model ignoring spillovers is biased. We illustrate the implementation of the test in the context of measuring the effect of credit supply contractions on firms' employment decisions.

The rest of the paper proceeds as follows. In Section 2, we present BRS21's framework to handle spillover effects in empirical research. Section 3 contains the main result and a discussion about the advantages of using the *leave-one-out treatment average* as a proxy for the treatment intensity. Section 4 presents an illustration of the implications of our result. Section 5 concludes. The Appendix contains auxiliary calculations.

2. Econometric Framework

BRS21 analyze workhorse models of firm interactions leading to the following specification capturing spillovers effects on corporate finance. Let y_{ig} denote an outcome, such as investments, debts, sales or employment, for firm i belonging to group g. Group g typically represents an industry or region. Assume that y_{ig} is determined by

$$y_{ig} = \varphi(d_{ig}, f_{ig}, d_{ig}f_{ig}), \tag{1}$$

where $\varphi(\cdot)$ is an unknown function, d_{ig} is a treatment indicator variable for unit i in group g that is equal to one if treatment is received, and zero otherwise, and f_{ig} is the group-level treatment intensity. The empirical specification of f_{ig} is the object of study in this paper.

The available data are a sample $\{y_{ig}, d_{ig}, s_i\}_{i=1}^{n}$, where the group variable $s_i \in \{1, .., g, .., G\}$ records the group corresponding to firm i. The estimands of interest are:

- the average direct effect:

$$\Delta := E(y_{ig}|d_{ig} = 1, f_{ig} = 0) - E(y_{ig}|d_{ig} = 0, f_{ig} = 0)$$

• the average spillover effect on the treated:

$$\Delta_T := E(y_{ig}|d_{ig} = 1, f_{ig} = 1) - E(y_{ig}|d_{ig} = 1, f_{ig} = 0)$$

• the average effect at the average intensity:

$$\Delta_A := E[y_{ig}|d_{ig} = 1, f_{ig} = E(d_{ig})] - E[y_{ig}|d_{ig} = 0, f_{ig} = E(d_{ig})]$$

There are, at least, two alternative models to estimate these estimands of interest. The first model, see BRS21 (Section 4), is:

Spillover Model with Group-Level Average Treatment:

$$y_{ig} = \beta_1 + \beta_2 d_{ig} + \beta_T d_{ig}\bar{d}_{ig} + \beta_C(1 - d_{ig})\bar{d}_{ig} + \epsilon_{ig}, \tag{2}$$

$$\bar{d}_{ig} := n_g^{-1} \sum_{j=1}^{n} d_{jg} 1(s_j = g) 1(s_i = g) \tag{3}$$

where for $n_g := \sum_{i=1}^{n} 1(s_i = g)$ denotes the number of firms in group g, and $\bar{d}_{ig}$ is the group-level average treatment.[2] This model uses $\bar{d}_{ig}$ as a proxy for the group-level treatment intensity f_{ig}. Under assumptions A1-A4 described in BRS21 and, for the convenience of the reader, replicated in the Appendix, this model delivers the approximations: $\Delta \approx \lim_{c \to 0}\{\beta_1 + \beta_2 + \beta_T c\} - \beta_1 = \beta_2$ and $\Delta_T \approx \beta_T$.

In randomized experiments, with treatment drawn independently from observables and unobservable variables, one has that the treatment indicator d_{ig} is independent of the group variable s_i. This motivates the second model, which is:

[2]Without loss of generality, one can weight d_{jg} in (3) to reflect the market share of treated firm in competition models, the share of R&D expenses of treated firms in agglomeration models, or the share of employment of treated firms in local demand spillover models.

Spillover Model with Leave-one-out Average Treatment:

$$y_{ig} = \gamma_1 + \gamma_2 d_{ig} + \gamma_T d_{ig}\tilde{d}_{ig} + \gamma_C(1 - d_{ig})\tilde{d}_{ig} + \zeta_{ig}, \tag{4}$$

$$E(\zeta_{ig}|d_{1g}, .., d_{ng}, s_1, .., s_n) = 0, \tag{5}$$

$$d_{ig} \text{ and } d_{jg} \text{ are independent and identically distributed for all } i \neq j \in \{1, .., n\}, \tag{6}$$

$$d_{ig} \text{ and } s_j \text{ are independent for all } i, j \in \{1, .., n\}, \tag{7}$$

where

$$\tilde{d}_{ig} := (n_g - 1)^{-1} \sum_{j \neq i} d_{jg} 1(s_j = g) 1(s_i = g) \tag{8}$$

is the group-level average treatment *excluding firm i itself*.[3] We call $\tilde{d}_{ig}$ the leave-one-out treatment average. In the Spillover Model with Leave-one-out Average Treatment, the treatment is allocated as in the simple randomization procedure, i.e., ζ_{ig} is mean-independent of (d_{jg}, s_j) for any i, j, d_{ig} is independent of s_j for any i, j, and d_{jg} and d_{ig} are independent and identically distributed for any $i \neq j$. In particular, Assumption (7) restricts the dependence between the treatment indicator variable d_{ig} and the group variable s_i. Since d_{ig} and s_i are both observed, this restriction is testable and hence should not be taken as a disadvantage with respect to the Spillover Model with Group-Level Average Treatment. The Spillover Model with Leave-one-out Average Treatment uses $\tilde{d}_{ig}$ as a proxy for f_{ig} and it delivers the approximations

$$\Delta \approx \gamma_2, \; \Delta_T \approx \gamma_T \text{ and } \Delta_A \approx \gamma_2 + (\gamma_T - \gamma_C)E(d_{ig}).$$

Note that the approximations for Δ and Δ_A coincide if the average spillover effects

[3] The usual rank restriction, e.g., $rank E(x_i x_i^\top) = 3$ for $x_i^\top = (1, d_{ig}, d_{ig}\tilde{d}_{ig}, (1 - d_{ig})\tilde{d}_{ig})$, is assumed to hold.

are homogeneous, i.e., $\gamma_T = \gamma_C$.

Comparison. Two differences arise when comparing the two Spillover Models. First, the proxies for the group-level treatment intensity, and consequently the approximations to the estimands of interest, do not necessarily coincide. Second, while the treatment in model with the Leave-one-out Average Treatment is allocated as in the simple random assignment procedure, in the model with the Group-Level Average Treatment it is not clear whether treatment is allocated as in the simple random assignment procedure or as in a more sophisticated randomization procedure. Little is known about whether these differences are relevant and, if indeed they are, whether one should use the Group-Level Average Treatment or the Leave-one-out Average Treatment. The next section spells out two advantages of using the Leave-one-out Average Treatment over the Group-Level Average Treatment when treatment is allocated as in the simple randomization procedure described in (5)-(7). These advantages illustrate, first, the relevance of the choice of the proxy for treatment intensity and, second, the benefits obtained from the rigorous modeling of the treatment allocation procedure.

3. Main Result

Consider the model ignoring spillovers:

Baseline Model:

$$y_{ig} = \alpha_1 + \alpha_2 d_{ig} + \xi_{ig},\ E(\xi_{ig}|d_{1g},..,d_{ig},..,d_{ng}) = 0,\ \ i = 1,..,n. \tag{9}$$

To describe the advantages of using the Leave-one-out Average Treatment, we now compare the biases arising from estimating γ_2 and β_2 using the ordinary least squares estimator $\hat{\alpha}_2$ of the coefficient α_2 in the Baseline Model.

BRS21 (Proposition 1) prove the following result:

Lemma 1 *The bias of the baseline estimator $\hat{\alpha}_2$ for the estimand β_2 defined in the*

Spillover Model with Group-Level Average Treatment is:

$$E(\hat{\alpha}_2) - \beta_2 = (\beta_T - \beta_C)E(d_{ig}) + \beta_T \frac{V(\bar{d}_{ig})}{E(d_{ig})} + \beta_C \frac{V(\bar{d}_{ig})}{1 - E(d_{ig})}. \tag{10}$$

We show in the Appendix that:

Proposition 1 *The bias of the baseline estimator $\hat{\alpha}_2$ for the estimand γ_2 defined in the Spillover Model with Leave-one-out Average Treatment is:*

$$E(\hat{\alpha}_2) - \gamma_2 = (\gamma_T - \gamma_C)E(d_{ig}). \tag{11}$$

The expression in Proposition 1 is simpler to interpret than the one in Lemma 1: $\hat{\alpha}_2$ is an unbiased estimator of γ_2 (and, on the proviso that $\Delta = \gamma_2$, an unbiased estimator of the average direct effect) if and only if the spillover effects on the treated and the untreated groups are homogeneous, i.e., $\gamma_T = \gamma_C$. This is the first advantage from choosing the Leave-one-out Average Treatment as a proxy for Group-level Treatment Intensity. Consequently, from the characterization of the bias in Proposition 1, the following statistical test can be performed to confirm that the baseline estimator is a biased estimator of the average direct effect.

Corollary 1 *Empirical researchers can check that the baseline estimator $\hat{\alpha}_2$ is biased for the average direct effect Δ by performing a heteroskedastic-robust Wald test for the null hypothesis $H_0 : \gamma_T - \gamma_C = 0$ versus the alternative $H_1 : \gamma_T - \gamma_C \neq 0$ based on the ordinary least squares estimator of γ_T, γ_C from the Spillover Model with the Leave-one-out Average Treatment.*

This check complements, by providing a rigorous statistical test for confirming the presence of bias, the three-step heuristic guidance suggested by BRS21 (Section 5) for the diagnosis the bias of the baseline estimator. One could use the homoskedastic-only Wald test in this check if, instead of the mean-independence assumption (5), one would assume that ζ_{ig} is independent of the treatment indicator variables $d_1, .., d_n$ and the group indicator variables $s_1, .., s_n$.

Since $E[y_{ig}|d_{ig} = 1, \tilde{d}_{ig} = E(d_{ig})] = \gamma_1 + \gamma_2 + \gamma_T E(d_{ig})$ and $E[y_{ig}|d_{ig} = 0, \tilde{d}_{ig} = E(d_{ig})] = \gamma_1 + \gamma_C E(d_{ig})$, one has $E(\hat{\alpha}_2) = \Delta_A$ and the following Corollary holds.

Corollary 2 *The baseline estimator $\hat{\alpha}_2$ is unbiased for the average effect at the average intensity Δ_A.*

Note that, in general, the average effect at the average intensity is not equal to the sum of the average direct effect and the average spillover effect on the treated, which should prevent one from interpreting the baseline estimator as an unbiased estimator of the aggregation of the average direct effect and the average spillover effects (see e.g., Biswas and Zhai, 2021).

The second advantage from choosing the Leave-one-out average Treatment proxy comes from the interpretation of the approximation γ_T to the average spillover effect on the treated Δ_T. Consider the case of a group g with two firms and only i is treated so $\bar{d}_{ig} = 1/2$ and $\tilde{d}_{ig} = 0$. In this case, there is no spillover effect on the treated, which is not reflected in the difference $E(y_{ig}|d_{ig} = 1, \bar{d}_{ig} = 1/2) - E(y_{ig}|d_{ig} = 1, \bar{d}_{ig} = 0) = \beta_T/2$ obtained from the Spillover Model with the Group-Level Average Treatment. Compare this result with $E(y_{ig}|d_{ig} = 1, \tilde{d}_{ig} = 0) - E(y_{ig}|d_{ig} = 1, \tilde{d}_{ig} = 0) = 0$ obtained from the Spillover Model with the Leave-one-out Average Treatment. This suggests that γ_T approximates the average spillover effect on the treated Δ_T that we are looking for, while β_T approximates something else. Another way of interpreting this difference is that Spillover Model with the Group-Level Average Treatment counts 'twice' the effect of d_{ig} (by including it in $\beta_2 d_{ig}$ and in $\bar{d}_{ig}$ in $\beta_T d_{ig}\bar{d}_{ig}$), while the Spillover Model with the Leave-one-out Average Treatment counts only once the effect of d_{ig} (by including it in $\gamma_2 d_{ig}$ and excluding it from $\tilde{d}_{ig}$ in $\gamma_T d_{ig}\tilde{d}_{ig}$).

4. Illustration

We now illustrate the use of the previous results in the context of applications investigated in the empirical literature. The aim is to show the advantages of using the leave-one-out average treatment proxy to diagnose the bias on the baseline estimator induced by the spillover effects.

There is a growing empirical literature seeking to incorporate spillovers to their baseline models. These papers differ on their modeling of spillovers in two dimensions. They either use the group-level average treatment or the leave-one-out treatment average as a proxy for the treatment intensity, and they either assume homogeneous or heterogeneous effects for the group of treated and untreated firms. The table below summarizes these differences among already published papers.

Table I: Proxies Employed in Applications

Spillover Effects ↓ / Proxy →	Group-level Average	Leave-one-out Average
Homogeneous	Doerr (2021); Breuer et al. (2021)	Beck et al. (2021); Huber (2018)
Heterogeneous	Gopalakrishnan et al. (2021)	BRS21

Our results can be apply to any of these papers. We choose the application in BRS21 because the careful execution of the study lends itself to extension by applying the result in Proposition 1 (and its Corollaries).

The estimand of interest is the average direct effect of a bank-lending cut (the bank in the database is Commerzbank) on German firms' employment growth. Here y_{ig} is the symmetric growth employment rate over the 2008 to 2012 period for firm i located in county g; d_{ig} is a dummy variable that equals one if the fraction of the firm's relationship banks that are Commerzbank branches is greater or equal than .5, and zero otherwise ($\text{CBdep}(0/1)_{ic}$ in BRS21's notation); $\tilde{d}_{ig}$ is the average Commerzbank dependence calculated based on d_{ig} of all other firms in the county (g), excluding firm i itself ($\overline{\text{CBdep}(0/1)}_{ic}$ in BRS21's notation). For the convenience of the reader, we reproduce the estimates in the table below (see BRS21, Table 5, Columns (4) and (6)).

Table II: Estimates from BRS21

	(1)	(2)
d_{ig}	-.028	-.053
	(.006)	(.017)
$d_{ig}\tilde{d}_{ig}$		.025
		(.068)
$(1-d_{ig})\tilde{d}_{ig}$		-0.115
		(0.038)

Note: The dependent variable is the symmetric growth rate of firm employment from 2008 to 2012. Robust standard errors, clustered at the county level, are in parenthesis. Source: BRS21 (Table 5).

By comparing the baseline estimate $\hat{\alpha}_2 = -.028$ with $\hat{\gamma}_2 = -.053$ in Table II, BRS21 infer that ignoring spillover effects makes the baseline estimator $\hat{\alpha}_2$ biased for the direct treatment effect Δ. This comparison, however, does not take into account sampling variability, which, as we are going to show below, can change the above inference.

Corollary 1 proposes a Wald test to rigorously confirm, by taking into account sampling variability, that the baseline estimator is biased for the direct treatment effect. Performing this test is straightforward. It requires to compute the Wald test statistic:

$$W = \frac{(\hat{\gamma_T} - \hat{\gamma_C})^2}{se^2_{\hat{\gamma_T}} + se^2_{\hat{\gamma_C}} - 2c\hat{o}v(\hat{\gamma_T}, \hat{\gamma_C})},$$

where $\hat{\gamma_T}$ and $\hat{\gamma_C}$ are the OLS estimators for the parameters γ_T and γ_C in Equation (4), $se_{\hat{\gamma_C}}$ and $se_{\hat{\gamma_T}}$ are their respective standard errors and $c\hat{o}v(\hat{\gamma_T}, \hat{\gamma_C})$ is the covariance estimator. The asymptotic null distribution of the Wald statistic is a chi-squared distribution with one degree of freedom, from which we can compute critical values. The Wald test will suggest to reject the null hypothesis (and confirm that the baseline estimator is biased for the average direct effect) if the realized value of the Wald test statistic is greater or equal than the critical value.

Table II has all the values to compute the realized value of the Wald test statistic except for $c\hat{o}v(\hat{\gamma_T}, \hat{\gamma_C})$. For illustrative purposes, we take two values: a lower bound of zero and an upper bound from the Cauchy-Schwarz Inequality.[4] In the case of the upper bound, the realized value of the Wald statistic is

$$W = \frac{[.025 - (-.115)]^2}{.068^2 + .038^2 - 2 \times .00258} = \frac{.0196}{.0009} = 21.77,$$

while the critical value at the 99% confidence level is $cv_{.90} = 6.63$. Since the realized value of the statistic ($W = 21.77$) is greater than the 99% critical value ($cv_{.90} = 6.63$), the test indicates that the baseline estimator is biased for the average direct effect. However, the baseline estimator still is an unbiased estimator for the average effect at the average intensity (Corollary 2). In the case of the lower bound, the realized value of the statistic ($W = 3.23$) is smaller than the 99% critical value ($cv_{.90} = 6.63$). In such a case, Proposition 1 indicates that there is no evidence that the baseline estimator is a biased estimator of the average direct effect. We conclude that, from the estimates in Table II, one cannot infer that ignoring spillover effects makes the baseline estimator biased for the average direct effect. We remark that these results are not available if one chooses the Group-Level Average Treatment as a proxy for the Group-Level Treatment Intensity.

5. Conclusion

This paper discusses the choice between two alternative proxies for the Group-Level Treatment Intensity in the framework introduced by BRS21 to estimate spillover effects in empirical corporate finance. We show that this choice is relevant for diagnosing the existence of bias induced by spillover effects. We highlights that the Leave-one-out Average Treatment proxy has two advantages over the Group-Level Average Treatment proxy. First, it simplifies the formula for bias of the baseline estimator arising from the spillover effects, thereby, facilitating its diagnosis. The

[4]By the Cauchy-Schwarz Inequality, $\hat{cov}(\hat{\gamma_T}, \hat{\gamma_C}) \leq se_{\hat{\gamma_T}} \times se_{\hat{\gamma_C}}$. Note that any other possible value for $\hat{cov}(\hat{\gamma_T})$ delivers a lower value of the the statistic. The paper by BRS21 does not report the value for $\hat{cov}(\hat{\gamma_T})$. Empirical researchers have the estimated value of the covariance available since standard statistical software, such as STATA, produces this value.

baseline estimator is unbiased for the average direct effect if and only if the spillover effects are homogeneous in the treated and untreated groups. Second, it clarifies the definition of the average spillover effect on the treated, thereby, facilitating its interpretation. These advantages justify the use of the Leave-one-out Average Treatment as the preferred proxy and suggest a straightforward rule to confirm the existence of the bias of the estimator in the model ignoring spillovers.

References

Beck, T., S. Da-Rocha-Lopes and A.F. Silva (2020): "Sharing the Pain? Credit Supply and Real Effects of Bank Bail-ins", *The Review of Financial Studies.*

Berg, T., M. Reisinger and D. Streitz (2021): "Spillover Effects in Empirical Corporate Finance", *Journal of Financial Economics.*

Biswas, S. and W. Zhai (2021): "Economic policy uncertainty and cross-border lending", *Journal of Corporate Finance.*

M. Breuer, K. Hombach and M.A. MÃ¼ller (2019): "When you talk, I remain silent: Spillover effects of peers' mandatory disclosures on firms' voluntary disclosures", *The Accounting Review.*

Chesher, A. (1989): "Hajek Inequalities, Measures of Leverage and the Size of Heteroskedastic Robust Wald Tests", *Econometrica.*

Doerr, S. (2021): "Stress tests, entrepreneurship, and innovation", *Review of Finance.*

Gopalakrishnan, B., J. Jacob and S. Mohapatra (2021): "Risk-sensitive Basel regulations and firms' access to credit: Direct and indirect effects", *Journal of Banking & Finance.*

Huber, K (2018): "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties", *American Economic Review.*

Appendix

Assumptions in BRS21. For the sake of completeness, we now replicate Assumptions A1-A4 in BRS21:

A.1 Treatment status fulfills the conditional independence assumption (CIA).

A.2 Outcomes not only depend on the treatment status of an individual firm, but also on the treatment intensity in an industry (in the case of competition models) or a region (in the case of spatial models).

A.3 Spillovers occur within industries/regions, but not across industries/regions (i.e., we abstract from general equilibrium effects).

A.4 We assume a linear relation throughout the paper.

Auxiliary Calculations. We now derive the formula for the bias in (10). The bias of the estimator $\hat{\alpha}_2$ is (see Berg et al., 2021, Display (21)):

$$E(\hat{\alpha}_2) = \gamma_2 + \gamma_T \frac{C(d_{ig}, d_{ig}\tilde{d}_{ig})}{V(d_{ig})} + \gamma_C \frac{C[d_{ig}, (1-d_{ig})\tilde{d}_{ig}]}{V(d_{ig})}, \tag{12}$$

where $V(d_{ig} := E(d_{ig})[1-E(d_{ig})]$ is the variance of d_{ig} and, for any random variables a and b, $C(a,b)$ denotes their covariance. We now observe that

$$\begin{aligned} E(d_{ig}\tilde{d}_{ig}) &\underset{(i)}{=} E[E(d_{ig}\tilde{d}_{ig}|s_1,..,s_n)] \underset{(ii)}{=} E[\tilde{d}_{ig}E(d_{ig}|s_1,..,s_n)] \\ &\underset{(iii)}{=} E[\tilde{d}_{ig}E(d_{ig})] \underset{(iv)}{=} E(d_{jg})E(d_{ig}) \underset{(v)}{=} E(d_{ig})^2, \end{aligned}$$

where (i) follows from the Law of Iterated Expectations, (ii) follows from observing that $\tilde{d}_{ig}$ is a function of $s_1,..,s_n$ and the assumption that d_{ig} and d_{jg} are independent, (iii) follows from the assumption that d_{ig} and $s_1,...,s_n$ are independent, (iv) follows from observing that $E(\tilde{d}_{ig}) = E[E(\tilde{d}_{ig}|s_1,..,s_n)] = E[(n_g-1)^{-1}\sum_{j\neq i} E(d_{jg}1(s_i = g)1(s_j = g)|s_1,..,s_n] = E(d_{jg})$, and (v) follows from the assumption that d_{ig} and d_j

are identically distributed. Hence,

$$
\begin{aligned}
C(d_{ig}, d_{ig}\tilde{d}_{ig}) &= E(d_{ig}\tilde{d}_{ig}) - E(d_{ig})E(d_{ig}\tilde{d}_{ig}) = E(d_{ig}\tilde{d}_{ig})[1 - E(d_{ig})] = E(d_{ig})^2[1 - E(d_{ig})] \\
&= E(d_{ig})V(d_{ig}) \\
C[d_{ig}, (1 - d_{ig})\tilde{d}_{ig}] &= E[d_{ig}(1 - d_{ig})\tilde{d}_{ig}] - E(d_{ig})E[(1 - d_{ig})\tilde{d}_{ig}] \\
&= E(d_{ig}\tilde{d}_{ig}) - E(d_{ig}\tilde{d}_{ig}) - E(d_{ig})E(\tilde{d}_{ig}) + E(d_{ig})E(d_{ig}\tilde{d}_{ig}) \\
&= E(d_{ig})E(d_{ig})^2 - E(d_{ig})^2 \\
&= E(d_{ig})E(d_{ig})[E(d_{ig}) - 1] = -E(d_{ig})V(d_{ig}).
\end{aligned}
$$

Replacing these expressions back in (11) one obtains:

$$
E(\hat{\alpha}_2) = \gamma_2 + \gamma_T \frac{E(d_{ig})V(d_{ig})}{V(d_{ig})} + \gamma_C \frac{[-E(d_{ig})]V(d_{ig})]}{V(d_{ig})} = \gamma_2 + (\gamma_T - \gamma_C)E(d_{ig}).
$$

Wald Test. We now describe the Wald test for confirming the presence of bias in $\hat{\alpha}_2$ when estimating γ_2. Let $\hat{\gamma}$ denote the ordinary least squares estimator obtained from specification (4). Let $v\hat{a}r(\hat{\gamma})$ denote a consistent estimator of the variance of $\hat{\gamma}$. Define $\gamma = (\gamma_1, \gamma_2, \gamma_T, \gamma_C)^\top$ and the vector $Q = (0, 0, 1, -1)$. Rewrite now the null hypothesis $H_0 : \gamma_T - \gamma_C = 0$ as $H_0 : Q\gamma = 0$. The Wald statistic is:

$$
W = (Q\hat{\gamma})^\top [Qv\hat{a}r(\hat{\gamma})Q^\top]^{-1} Q\hat{\gamma} = \frac{(Q\hat{\gamma})^2}{Qv\hat{a}r(\hat{\gamma})Q^\top}.
$$

Under standard regularity conditions, the distribution of W, when the null hypothesis holds, is approximately a chi-squared distribution with one degree of freedom.[5] The Wald test confirms (with significance level α) the presence of bias in $\hat{\alpha}_2$ for estimating γ_2 when the Wald statistic W is above the $(1 - \alpha)$ -quantile of a chi-square distribution with one degree of freedom.

[5]This approximation applies when the data do not contain points of high leverage (see, e.g., Chesher (1989) for a definition of the leverage of points in regression designs). If the data contain points of high leverage, the discrepancy between the exact and nominal size of the Wald test can be substantial and the test can deliver misleading inferences.