computation

Communication

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

Fabiana Gómez [1] and David Pacini [2,*]

[1] School of Accounting and Finance, University of Bristol, 15/19 Tyndall's Park Road, Bristol BS8 1PQ, UK
[2] School of Economics, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK
* Correspondence: david.pacini@bristol.ac.uk

Abstract: The existing literature indicates that spillovers can lead to a complicated bias in the estimation of causal effects in empirical corporate finance. We show that, under the assumption of simple random treatment assignment and when the proxy chosen for the group-level treatment coverage is the leave-one-out average treatment, such a spillover bias exists if and only if the average indirect effects on the treated and untreated groups are different. We quantify the gains in spillover bias reduction using Monte Carlo exercises. We propose a Wald test to statistically infer the presence of bias. We illustrate the application of this test to bear out spillovers in firms' employment decisions.

Keywords: spillover bias; average direct effect; average indirect effect; causal graph

JEL Classification: C21; G30

Citation: Gómez, F.; Pacini, D. Spillover Effects in Empirical Corporate Finance: Choosing the Proxy for Treatment Coverage. *Computation* **2022**, *10*, 149. https://doi.org/10.3390/computation10090149

Academic Editors: Priyantha Wijayatunga, Linbo Wang and Wang Miao

Received: 6 June 2022
Accepted: 25 August 2022
Published: 31 August 2022

1. Introduction

It is well known that spillovers arise in corporate finance through firm competition and geographical agglomeration; see, for example, [1], BRS21 hereon. Firm-level outcomes, such as sales or investments, depend on firms' own treatment assignment in a given intervention and on the fraction of firms treated in the same industry and/or geographical region. As an example of these spillovers, consider the following illustration taken from BRS21. Assume that some coffee shops (i.e., the treated shops) in a given neighborhood are subject to a rise in the price of coffee beans (the treatment). The rise in the input price leads to a rise in the final price per cup of coffee for treated shops, and consequently, to a reduction in their volume of sales, i.e, the direct effect of the treatment. This is not the only effect on sales. It is likely that due to the increase in price, some consumers switch coffee shops. For coffee shops in the same neighborhood whose price has not changed, this implies an increase in the volume of sales, i.e, the spillover effect on the untreated. This spillover is an example of interference in causal inference; see, for example, [2].

Spillovers can lead to a complicated bias in the estimation of causal effects; see, for example, BRS21. Spillover bias arises when the coverage of an intervention is omitted from the analysis. In terms of the above illustration, the coverage of the intervention is the proportion of coffee shops in a neighborhood that were affected by the increase in their costs. The reduction in the volume of sales suffered by shops in the treated group could be lower if more shops in the neighborhood were affected by the increase in costs. Omitting the proportion of affected coffee shops in a given neighborhood would induce a bias in the estimation of the direct effect. The coverage of an intervention can be measured either by the *group-level average* or by the *leave-one-out average* proxy. There is a The group-level average proxy is the average number of firms in a group subject to the treatment including the firm itself. The leave-one-out average proxy is the average number of firms subject to the treatment excluding the firm itself. Little is known about which of the two proxies one should use when controlling for spillovers.

The objective of this paper is to compare the implications for the spillover bias when using the leave-one-out average or the group-level average proxy. We show that choosing the leave-one-out average proxy has two advantages. First, it simplifies the formula for the spillover bias, which facilitates its diagnosis. Second, it clarifies the definition of the average indirect effects, thereby facilitating its interpretation. These advantages justify the use of the leave-one-out average as the preferred proxy.

The leave-one-out average proxy suggests a straightforward statistical test for the diagnosis of the spillover bias. The test is a heteroskedastic-robust Wald test for the null hypothesis of equal average indirect effects on the treated and untreated groups. If this null hypothesis is rejected, the ordinary least square estimator of the average direct effect omitting the spillovers is biased. We illustrate the implementation of this 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 describe the two proxies to model spillovers in empirical research in corporate finance. Section 3 contains the main result and a discussion of the advantages of using the leave-one-out average as the preferred proxy. Section 4 presents results from a Monte Carlo study exploring the bias of alternative estimators of the average direct effect. Section 5 presents an illustration of the implications of our results. Section 6 concludes. Appendix A contains auxiliary calculations.

2. Framework and Graphical Representation

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. Following BRS21, we assume that y_{ig} is determined by

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

where $\varphi(\cdot)$ is an unknown function, d_{ig} is a treatment indicator variable, and f_{ig} is the group-level treatment coverage (or intensity). The treatment indicator variable is equal to one if firm i receives the treatment, and is equal to zero otherwise. The group-level treatment coverage takes values between zero and one. The available data are a sample of size n $\{y_{ig}, d_{ig}, s_i\}_{i=1}^{n}$, where the group variable $s_i \in \{1, \ldots, g, \ldots, G\}$ records firm i's group. The object of study in this paper is the empirical specification of f_{ig}.

Causal estimands of interest include the average direct, indirect, total, and overall effects. The average direct effect is the difference between the average outcome for treated and untreated firms given all other things being equal. Following the illustration in the introduction, the average direct effect is the average change in the sales of coffee shops that experienced an increase in the price of coffee beans in the absence of spillovers. Formally, the average direct effect is

$$\Delta_D := E(y_{ig}|d_{ig} = 1, f_{ig} = 0) - E(y_{ig}|d_{ig} = 0, f_{ig} = 0). \tag{2}$$

The average indirect effects are those due to treatment coverage. They can be defined by comparing the outcomes in the treated or untreated firms. Following the illustration, the average indirect effect on the treated firms is the difference in average sales for a treated coffee shop between two hypothetical situations for the group: the group is fully treated vs. the group is not treated at all. Formally, the average indirect effect on the treated is

$$\Delta_T := E(y_{ig}|d_{ig} = 1, f_{ig} = 1) - E(y_{ig}|d_{ig} = 1, f_{ig} = 0). \tag{3}$$

The average indirect effect on the untreated is defined similarly:

$$\Delta_U := E(y_{ig}|d_{ig} = 0, f_{ig} = 1) - E(y_{ig}|d_{ig} = 0, f_{ig} = 0). \tag{4}$$

The average total and overall effects provide summary measures combining direct and indirect effects. The average total effect is the sum of the average direct effect and the average indirect effect on the untreated:

$$\Delta_{tot} := \Delta_D + \Delta_U. \tag{5}$$

The average overall effect is

$$\Delta_{over} := E(y_{ig}|d_{ig} = 1, f_{ig} = 1) - E(y_{ig}|d_{ig} = 0, f_{ig} = 0) = \frac{\Delta_D + \Delta_T + \Delta_U}{2}. \tag{6}$$

Finally, for later reference, we define the average effect at the a-coverage as

$$\Delta_a := E(y_{ig}|d_{ig} = 1, f_{ig} = a) - E(y_{ig}|d_{ig} = 0, f_{ig} = a), \tag{7}$$

where a is a number between zero and one. Following the illustration, the average effect at the a-coverage is the difference in average sales for treated and untreated coffee shops when the group is treated at the coverage level a.

We consider two alternative models to estimate the causal estimands of interest. In the first model, the treatment assignment d_{ig} is independent of the group assignment variable s_i. This model is a special case of the setting delineated by [2].

Spillover Model with Leave-one-out Average:

$$y_{ig} = \gamma_1 + \gamma_2 d_{ig} + \gamma_3 d_{ig}\tilde{d}_{ig} + \gamma_4(1 - d_{ig})\tilde{d}_{ig} + \zeta_{ig}, \tag{8}$$

$$E(\zeta_{ig}|d_{1g}, ..., d_{ng}, s_1, ..., s_n) = 0, \tag{9}$$

$$d_{ig} \text{ and } d_{jg} \text{ are independent and identically distributed for all } i \neq j \in \{1, ..., n\}, \tag{10}$$

$$d_{ig} \text{ and } s_j \text{ are independent for all } i, j \in \{1, ..., n\}, \tag{11}$$

where

$$\tilde{d}_{ig} := (n_g - 1)^{-1} \sum_{j \neq i} d_{jg} 1(s_j = g) 1(s_i = g) \tag{12}$$

is the leave-one-out average, $n_g := \sum_{i=1}^{n} 1(s_i = g)$ is the number of firms in group g, and $1(\cdot)$ is the indicator function, taking a value of one when the condition in parentheses is satisfied, and zero otherwise. In this model, the treatment is allocated as in the simple random treatment assignment assumption, 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 (11) 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 of the model. This model uses $\tilde{d}_{ig}$ as a proxy for the coverage f_{ig}. It delivers the approximations

$$\Delta_D \approx \gamma_2, \Delta_T \approx \gamma_3, \Delta_U \approx \gamma_4 \text{ and } \Delta_a \approx \gamma_2 + (\gamma_3 - \gamma_4)a.$$

The approximations for Δ and Δ_a coincide if the average indirect effects are homogeneous, i.e., $\gamma_3 = \gamma_4$.

In the second model, the treatment indicator d_{ig} and the group variable s_i can be related. Treatment in this model may not be assigned as in the simple random treatment assignment assumption. This model has been postulated by BRS21.

Spillover Model with Group-Level Average:

$$y_{ig} = \beta_1 + \beta_2 d_{ig} + \beta_3 d_{ig}\bar{d}_{ig} + \beta_4(1 - d_{ig})\bar{d}_{ig} + \epsilon_{ig}, \tag{13}$$

where

$$\bar{d}_{ig} := n_g^{-1} \sum_{j=1}^{n} d_{jg} 1(s_j = g) 1(s_i = g) \tag{14}$$

is the group-level average treatment. This model uses $\bar{d}_{ig}$ as a proxy for f_{ig}. BRS21 shows, under additional assumptions replicated in the appendix, that this model delivers the approximations:

$$\Delta_D \approx \lim_{x \to 0} (\beta_1 + \beta_2 + \beta_3 x) - \beta_1 = \beta_2, \ \Delta_T \approx \beta_3, \text{ and } \Delta_U \approx \beta_4.$$

To facilitate the comparison between the two models, we now represent them using causal graphs.

Causal Graph I: Spillover Model with Leave-one-out Average

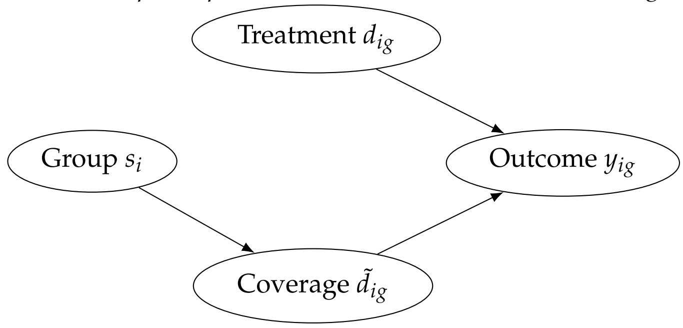

Causal Graph II: Spillover Model with Group-Level Average

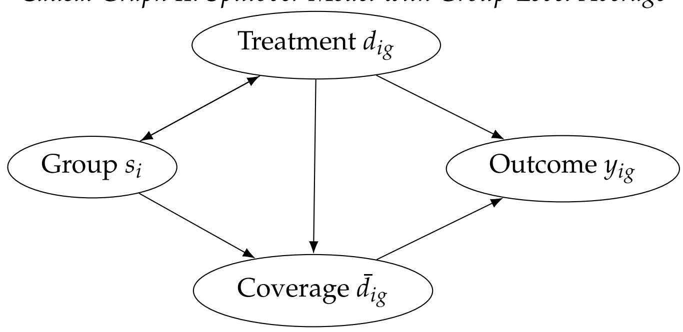

Two differences arise when comparing the models. First, the proxies for coverage, and consequently, the approximations of the estimands of interest, do not coincide. While d_{ig} and the group-level average $\bar{d}_{ig}$ are correlated, d_{ig} and the leave-one-out average $\tilde{d}_{ig}$ are not. Notice the absence of an arrow connecting the nodes 'Treatment' and 'Coverage' in Causal Graph I. Second, while the treatment variable in the model with the leave-one-out average is assumed to follow the simple random assignment assumption, in the model with the group-level average, it is not clear whether treatment is allocated as in a more sophisticated experimental procedure. Notice the presence of the bi-directed arrow connecting the nodes 'Treatment' and 'Group' in Causal Graph II. Little is known about whether these differences are relevant and, if they are, whether one should use the group-level average or the leave-one-out average proxy. The next section spells out two advantages of using the leave-one-out proxy. These advantages illustrate, first, the relevance of the choice of proxy for the coverage and, second, the benefits obtained from the rigorous modeling of the treatment allocation procedure.

3. Main Results

To proceed, we compare the spillover bias arising from estimating γ_2 using a baseline model ignoring spillovers.

Baseline Model:

$$y_{ig} = \alpha_1 + \alpha_2 d_{ig} + \xi_{ig}, \; E(\xi_{ig}|d_{1g}, ..., d_{ig}, ..., d_{ng}) = 0, \; i = 1, ..., n. \tag{15}$$

If there are no spillovers, the OLS estimator $\hat{\alpha}_2$ of α_2 is an unbiased estimator of the average direct effect. BRS21 (Proposition 1) proves the following result:

Lemma 1. *The spillover bias of the baseline estimator $\hat{\alpha}_2$ for the estimand β_2 is :*

$$E(\hat{\alpha}_2) - \beta_2 = (\beta_3 - \beta_4)E(d_{ig}) + \beta_3 \frac{V(\bar{d}_{ig})}{E(d_{ig})} + \beta_4 \frac{V(\bar{d}_{ig})}{1 - E(d_{ig})}. \tag{16}$$

We show in Appendix A that:

Proposition 1. *The spillover bias of the baseline estimator $\hat{\alpha}_2$ for the estimand γ_2 is:*

$$E(\hat{\alpha}_2) - \gamma_2 = (\gamma_3 - \gamma_4)E(d_{ig}). \tag{17}$$

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 if and only if the indirect effects on the treated and the untreated groups are homogeneous, i.e., $\gamma_3 = \gamma_4$. This is the first advantage of choosing the leave-one-out average proxy.

From the characterization of the spillover bias in Proposition 1, the following statistical test can statistically infer if 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_D by performing a heteroskedastic-robust Wald test for the null hypothesis $H_0 : \gamma_3 - \gamma_4 = 0$ versus the alternative $H_1 : \gamma_3 - \gamma_4 \neq 0$ based on the ordinary least squares estimator of γ_3, γ_4.*

This check complements the heuristic guidance suggested by BRS21 by providing a test for statistically inferring the presence of spillover bias. If ζ_{ig} is independent of the treatment indicator variables and the group indicator variables, the homoskedastic-only Wald test is an alternative to perform this check.

What is the baseline estimator unbiased for? Since

$$E(y_{ig}|d_{ig} = 1, \tilde{d}_{ig} = a) = \gamma_1 + \gamma_2 + \gamma_3 a \tag{18}$$

$$E(y_{ig}|d_{ig} = 0, \tilde{d}_{ig} = a) = \gamma_1 + \gamma_4 a, \tag{19}$$

one has $E(\hat{\alpha}_2) = \Delta_{a=E(d_{ig})}$ and the following corollary holds.

Corollary 2. *The baseline estimator $\hat{\alpha}_2$ is unbiased for the average effect at the average coverage $\Delta_{a=E(d_{ig})}$.*

The average effect at the average coverage is not equal to the sum of the average direct effect and the average indirect effect on the treated firms, which should prevent one from interpreting the baseline estimator as an unbiased estimator of the aggregation of the average direct effect and the average indirect effects (see, for example, [3]).

The second advantage of choosing the leave-one-out average proxy comes from the interpretation of the approximation $\Delta_T \approx \gamma_2$. Consider the case of a group g with two firms. Only i is treated, so $\bar{d}_{ig} = 1/2$ and $\tilde{d}_{ig} = 0$. In this case, there is no indirect effect on the treated firm, which is not reflected in the difference $E(y_{ig}|d_{ig} = 1, \bar{d}_{ig} = 1/2) - E(y_{ig}|d_{ig} = 1, \bar{d}_{ig} = 0) = \beta_3/2$. 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 using the leave-one-out average. This suggests that γ_3 approximates the average indirect effect on the treated Δ_T that we are looking for, while β_3 approximates something else. Another way of interpreting this difference is that the group-level average counts 'twice' the effect of d_{ig}: by including it, first, in $\beta_2 d_{ig}$, and, second, in $\bar{d}_{ig}$ in $\beta_3 d_{ig}\bar{d}_{ig}$. The leave-one-out average counts only once the effect of d_{ig}: by including it in $\gamma_2 d_{ig}$ and excluding it from $\tilde{d}_{ig}$ in $\gamma_3 d_{ig}\tilde{d}_{ig}$.

4. Monte Carlo Exercises

To explore the finite sample properties of the estimator using the leave-one-out average proxy, we carry out a Monte Carlo study. We consider the specification:

$$y_{ig} = 10 - 4d_{ig} + \gamma_3 d_{ig}\tilde{d}_{ig} + 3.6(1-d_{ig})\tilde{d}_{ig} + \zeta_{ig}, \text{ for } i = 1, ..., N \in \{100, 400, 1600\}, \quad (20)$$

where $\gamma_3 \in \{-3.6, 3.6\}$ and ζ_{ig} has a normal distribution with mean 0 and variance 2. The design $\gamma_3 = 3.6$ has homogeneous average indirect effects and $\gamma_3 = -3.6$ has heterogeneous average indirect effects. These values are taken from the illustration in BRS21. The treatment variable d_{ig} follows a Bernoulli distribution with mean $E(d_{ig}) = 0.5$. The group variable s_i follows from the specification

$$s_i = \sum_{j=1}^{6} j1(c_{j-1} < z_i^\star \leq c_j), \quad (21)$$

where $c_0 = -\infty, c_1 = -1, c_2 = -0.5, c_3 = 0, c_4 = 0.5, c_5 = 1, c_6 = \infty$, and $z_i^\star$ is a standard normal random variable. The group variable s_i is independent of the treatment variable d_{ig}. The disturbance term ζ_{ig} is independent of the covariates.

Table 1 reports the results for the bias of different estimators of $\gamma_2 = -4$. 'Baseline' labels the ordinary least squares estimator from the specification (15) and 'leave-one-out' the ordinary least squares estimator from (8). As predicted by the theory for these experiments, the bias of the baseline estimator of the average direct effect is approximately $E(d_{ig})(\gamma_3 - \gamma_4) = 0.5(\gamma_3 - 3.6)$, which is half of the difference between the average indirect effect on the treated and untreated firms.

Table 1. Bias comparison of estimators.

		Heterogeneous: $\gamma_3 = -3.6$		**Homogeneous: $\gamma_3 = 3.6$**	
N		**Baseline**	**Leave-One-Out**	**Baseline**	**Leave-One-Out**
100	d_{ig}	−3.60	−0.0367	−0.0439	−0.0049
	$d_{ig}\tilde{d}_{ig}$		0.007		0.0433
	$(1-d_{ig})\tilde{d}_{ig}$		0.063		0.0569
400	d_{ig}	−3.60	0.0120	−0.0163	0.0195
	$d_{ig}\tilde{d}_{ig}$		−0.1152		−0.0502
	$(1-d_{ig})\tilde{d}_{ig}$		−0.0947		0.0007
1600	d_{ig}	−3.60	0.0053	0.0003	0.0108
	$d_{ig}\tilde{d}_{ig}$		0.0031		0.0182
	$(1-d_{ig})\tilde{d}_{ig}$		0.0155		0.0388

Note: The number of simulations is 5000.

5. Illustration

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

There is a growing body of empirical literature seeking to incorporate spillovers in baseline models. These papers differ in their modeling of spillovers in two dimensions. They either use the group-level average or the leave-one-out average as a proxy for the treatment coverage, and they either assume homogeneous or heterogeneous average indirect effects. Table 2 below summarizes these differences among already published papers.

Table 2. Proxies employed in applications.

Spillover Effects ↓/Proxy →	**Group-Level Average**	**Leave-One-Out Average**
Homogeneous	[4,5]	[6,7]
Heterogeneous	[8]	BRS21

Our results 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 0.5, and is 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)).

By comparing the baseline estimate $\hat{\alpha}_2 = -0.028$ with $\hat{\gamma}_2 = -0.053$ in Table 3, BRS21 infers that ignoring spillovers causes the baseline estimator to be biased for the average direct effect. This comparison, however, does not take into account sampling variability, which, as we are going to show below, can change the above inference.

Table 3. Estimates from BRS21.

	(1)	**(2)**
d_{ig}	−0.028 (0.006)	−0.053 (0.017)
$d_{ig}\tilde{d}_{ig}$		0.025 (0.068)
$(1 - d_{ig})\tilde{d}_{ig}$		−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 parentheses. Source: BRS21 (Table 5).

The estimate $\hat{\gamma}_3$ is 0.025, whereas the estimate $\hat{\gamma}_4$ is -0.115. To verify that this difference is not only due to sampling variability, Corollary 1 proposes a Wald test. Performing this test is straightforward. It requires the Wald test statistic to be computed:

$$W = \frac{(\hat{\gamma}_3 - \hat{\gamma}_4)^2}{se^2_{\hat{\gamma}_3} + se^2_{\hat{\gamma}_4} - 2\hat{cov}(\hat{\gamma}_3, \hat{\gamma}_4)},$$

where $\hat{\gamma}_3$ and $\hat{\gamma}_4$ are the OLS estimators for the estimands γ_3 and γ_4, $se_{\hat{\gamma}_3}$ and $se_{\hat{\gamma}_4}$ are their respective standard errors, and $\hat{cov}(\hat{\gamma}_3, \hat{\gamma}_4)$ 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 suggests rejecting the null hypothesis (and statistically inferring that the baseline estimator is biased for the average direct effect) if the realized value of the Wald test statistic is greater than or equal to the critical value.

Table 3 contains all of the values to compute the realized value of the Wald test statistic, except for $\hat{cov}(\hat{\gamma}_3, \hat{\gamma}_4)$. For illustrative purposes, we take two values: a lower bound of zero and an upper bound from the Cauchy–Schwarz Inequality. In the case of the upper bound, the realized value of the Wald statistic is

$$w = \frac{[0.025 - (-0.115)]^2}{0.068^2 + 0.038^2 - 2 \times 0.00258} = \frac{0.0196}{0.0009} = 21.77,$$

while the critical value at the 99% confidence level is $cv_{0.99} = 6.63$. Since the realized value of the statistic ($w = 21.77$) is greater than the 99% critical value ($cv_{0.99} = 6.63$), the test indicates that the baseline estimator is biased for the average direct effect. However, the baseline estimator is still an unbiased estimator for the average effect at the average coverage (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 ($cv_{0.99} = 6.63$). 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, from the estimates in Table 3, that one cannot infer that ignoring spillovers causes the baseline estimator to be biased for the average direct effect. We remark that these results are not immediately available if one chooses the group-level average as a proxy for the coverage.

6. Conclusions

Competitive interactions and agglomeration among firms generate spillovers after a shock, a change in regulation, or any kind of intervention affecting firms. Ignoring these spillovers when estimating causal effects leads to biased estimation. This paper discusses the choice between two alternative proxies for modeling spillovers. We show that this choice is relevant for diagnosing the existence of spillover bias. The leave-one-out average proxy has two advantages over the group-level average proxy. First, it simplifies the formula for the spillover bias, thereby facilitating its diagnosis. The baseline estimator is unbiased for the average direct effect if and only if the average indirect effects are homogeneous. Second, it clarifies the definition of the average indirect effect on the treated firms, thereby facilitating its interpretation. These advantages justify the use of the leave-one-out average as the preferred proxy and suggest a straightforward test to statistically infer the existence of spillover bias.

One natural extension is to investigate how to define the coverage proxy when the treatment is continuous instead of binary. This extension is outside of the scope of this paper and is left for future research.

Author Contributions: All authors have contributed equally. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable..

Acknowledgments: We thank Cecilia Dassatti and Klaus Schaeck for their comments and suggestions.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

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

Assumption A1. *Treatment status fulfills the conditional independence assumption (CIA).*

Assumption A2. *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).*

Assumption A3. *Spillovers occur within industries/regions, but not across industries/regions (i.e., we abstract from general equilibrium effects).*

Assumption A4. *We assume a linear relationship throughout the paper.*

Auxiliary Calculations. We now derive the formula for the bias in (17). The bias of the estimator $\hat{\alpha}_2$ is (see Berg et al., 2021, Display (21)):

$$E(\hat{\alpha}_2) = \gamma_2 + \gamma_3 \frac{C(d_{ig}, d_{ig}\tilde{d}_{ig})}{V(d_{ig})} + \gamma_4 \frac{C[d_{ig}, (1 - d_{ig})\tilde{d}_{ig}]}{V(d_{ig})}, \tag{A1}$$

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 $s_1, \ldots, s_n$ and the assumption that d_{ig} and d_{jg} are independent, (iii) follows from the assumption that d_{ig} and $s_1, \ldots, s_n$ are independent, (iv) follows from observing that $E(\tilde{d}_{ig}) = E[E(\tilde{d}_{ig}|s_1, \ldots, s_n)] = E[(n_g - 1)^{-1} \sum_{j \neq i} E(d_{jg} 1(s_i = g) 1(s_j = g)|s_1, \ldots, s_n] = E(d_{jg})$, and (v) follows from the assumption that d_{ig} and d_j are identically distributed. Hence,

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

Substituting these expressions back into (16), one obtains:

$$E(\hat{\alpha}_2) = \gamma_2 + \gamma_3 \frac{E(d_{ig})V(d_{ig})}{V(d_{ig})} + \gamma_4 \frac{[-E(d_{ig})]V(d_{ig})]}{V(d_{ig})} = \gamma_2 + (\gamma_3 - \gamma_4)E(d_{ig}).$$

Wald Test. We now describe the Wald test for statistically inferring 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 $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4)^\top$ and the vector $Q = (0, 0, 1, -1)$. Then, rewrite the null hypothesis $H_0 : \gamma_3 - \gamma_4 = 0$ 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, when the null hypothesis holds, the distribution of W is approximately a chi-squared distribution with one degree of freedom. This approximation applies when the data do not contain points of high leverage (see, for example, [9] 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. The Wald test statistically infers (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.

References

1. Berg, T.; Reisinger, M.; Streitz, D. Spillover Effects in Empirical Corporate Finance. *J. Financ. Econ.* **2021**, *142*, 1109–1127. [CrossRef]
2. Hudgens, M.; Halloran, M. Towards Causal Inference with Interference. *J. Am. Stat. Assoc.* **2008**, *103*, 832–842. [CrossRef] [PubMed]
3. Biswas, S.; Zhai, W. Economic policy uncertainty and cross-border lending. *J. Corp. Financ.* **2021**, *67*, 101867. [CrossRef]
4. Breuer, M.; Hombach, K.; Müller, M. When you talk, I remain silent: Spillover effects of peers' mandatory disclosures on firms' voluntary disclosures. *Account. Rev.* **2019**, *97*, 155–186. [CrossRef]
5. Doerr, S. Stress tests, entrepreneurship, and innovation. *Rev. Financ.* **2021**, *25*, 1609–1637. [CrossRef]
6. Huber, K. Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties. *Am. Econ. Rev.* **2018**, *108*, 868–898. [CrossRef]
7. Beck, T.; Da-Rocha-Lopes, S.; Silva, A.F. Sharing the Pain? Credit Supply and Real Effects of Bank Bail-ins. *Rev. Financ. Stud.* **2020**, *34*, 1747–1788. [CrossRef]
8. Gopalakrishnan, B.; Jacob, J.; Mohapatra, S. Risk-sensitive Basel regulations and firms' access to credit: Direct and indirect effects. *J. Bank. Financ.* **2021**, *126*, 106101. [CrossRef]
9. Chesher, A. Hajek Inequalities, Measures of Leverage and the Size of Heteroskedastic Robust Wald Tests. *Econometrica* **1989**, *57*, 971–977. [CrossRef]