The Price of Protecting Small Firms: Evidence from Public Procurement*

Darcio Genicolo-Martins[1]

[1] INSPER

Abstract

Governments worldwide restrict public tenders to small and medium-sized enterprises (SMEs), yet the fiscal cost of this policy is virtually unknown. I exploit a natural experiment in Sao Paulo, Brazil, where the policy operates through opt-out costs: public buyers that wish to conduct open tenders must justify their choice to audit courts, effectively penalizing the efficient procurement decision. Until March 2018, medical supplies were uniquely exempt from these costs. When a legal reinterpretation—unrelated to procurement outcomes—ended this exemption, it created a clean difference-in-differences in reverse (DiDiR) design. Open tenders yielded prices 7–13% lower (12–13% nominal, 7–10% in real terms) and attracted 19% more firms, but a Gelbach decomposition reveals that increased participation explains only 6% of the price reduction—most of the savings come from fiercer competition among bidders. The fiscal cost for one product group alone is R$50–85 million over 18 months, depending on the inflation adjustment. These costs are concentrated among standardized goods; for specialized items, restrictions appear less distortionary. The one clear benefit: SME-only tenders successfully promote local sourcing.

Keywords: public procurement, SMEs, policy costs, restricted public tenders.

JEL No.: H32, H57, L26, L53, O12, R12.

* *This version: April 2026 (First version: January 2021)*

Email address: `darciogm1@insper.edu.br` (Darcio Genicolo-Martins)

1. Introduction

Public procurement is increasingly used as an instrument of social and economic policy. Beyond securing goods and services at the best price, governments deploy procurement to promote small firms, encourage local sourcing, advance environmental sustainability, and pursue distributional objectives. SME set-asides are the most widespread of these instruments: in the European Union, SMEs account for nearly 99.8% of all registered firms and receive a large share of public contract value (PwC, 2014); in Brazil, where approximately 97% of firms are SMEs representing half of formal employment (Bastos et al., 2018), federal law mandates exclusive SME tenders for all items valued below R$80,000 (Thai, 2017; Bosio et al., 2022). A growing body of evidence documents the benefits of these policies—firm growth (Mel et al., 2008; Freedman, 2013), employment (Szerman, 2023), supplier network expansion (Cardoza et al., 2016)—though effectiveness depends on management quality (McKenzie and Woodruff, 2015) and capital provision (Fafchamps et al., 2014). What is largely missing from this literature is the other side of the ledger: credible evidence on what these policies *cost* the public sector. Without cost estimates, policymakers cannot evaluate whether the efficiency-equity trade-off justifies the instrument.

There are strong theoretical reasons to expect these costs to be large. Bulow and Klemperer (1996) show that, in a wide class of auction environments, attracting one additional bidder generates greater surplus for the auctioneer than any mechanism design refinement, implying that policies restricting participation impose a first-order efficiency loss. Krasnokutskaya and Seim (2011) formalize this in the procurement context, showing that bid preferences reduce both the number of bidders (extensive margin) and the aggressiveness of their bids (intensive margin). Empirically, Marion (2007) estimates a 3.8% cost increase from small business preferences in California highway procurement, and Athey et al. (2013) document similar trade-offs in US timber auctions. But outside the US, credible estimates of set-aside costs are rare, and in developing economies—where SME policies are most prevalent—they are essentially absent (Nakabayashi, 2013; Szucs, 2024; Decarolis et al., 2024).

This paper provides such estimates.

I exploit a quasi-experimental variation in Sao Paulo, Brazil, where the SME set-aside policy operates through a distinctive institutional mechanism: *opt-out costs*. Public buyer units (PBUs) that wish to conduct open tenders rather than SME-only tenders must justify their choice through a costly formal process, subject to scrutiny by audit courts and potential sanctions. These opt-out costs penalize the efficient procurement choice, shifting the default from open competition to restricted tenders. I exploit the timing of a policy change (March 2018) that extended these opt-out costs to a previously exempt product group (Group 65, medical and hospital supplies), using time and cross-sectional variation to estimate the fiscal cost of this institutional design.

Between August 2014 and February 2018, the state of Sao Paulo exempted Group 65 from the mandatory SME-only tender requirement that applied to all other product groups, based on a joint agreement between the state government and the audit court that health items constituted strategic goods warranting open competition. In March 2018, a legal opinion from a different oversight agency reversed this interpretation, subjecting Group 65 to the same SME-favoring rules as all other groups. This policy change was driven by a legalistic reinterpretation of isonomy principles rather than by changes in procurement outcomes for Group 65, making it plausibly exogenous to the outcomes studied.

The costs of the SME policy can be assessed only indirectly. Instead of a standard difference-in-differences (DiD), I use a variation known as difference-in-differences in reverse (DiDiR), or 'time-reversed DiD' (Kim and Lee, 2019). In standard DiD, the control group is never treated; in DiDiR, the control group is always treated, and the other group—here, Group 65—undergoes a switch.

In Sao Paulo, between August 2014 and February 2018, the government's default was SME-only tenders for items valued at R$80,000 or less, except for Group 65, which was restricted to open tenders. PBUs could avoid restricted tenders by justifying the choice to audit courts, but this costly process constitutes the opt-out costs that discourage open competition. From March 2018 on, Group 65 became subject to these opt-out costs just as any other product group.

DiDiR identifies pre-switch-period effects—it estimates what outcomes would have been for Group 65 if opt-out costs had been in place before March 2018 (Kim and Lee, 2019). The estimated price effects are large: negotiated prices were, on average, 7–13% lower for Group 65 than for other groups in the pre-period (12–13% nominal, 7–10% in IPCA-deflated real terms). The number of participating firms was approximately 19–20% higher in the short-term window, and valid bids were 16–17% more numerous. Winning firms were located 5–11 km further from PBUs, depending on the specification. The price effects are concentrated among high-value items, where the pool of capable non-SME suppliers is broader. The fiscal cost of the SME-only restriction for Group 65 alone ranges from R$50–85 million over the 18-month pre-period (R$50 million in real terms, R$85 million nominal)—about 7–12% of this product group's total procurement value. The gap between nominal and real estimates reflects differential inflation between medical supplies (subject to CMED price ceilings) and other product categories.

The design's statistical power is limited by the single-unit treatment structure: a randomization inference exercise yields a permutation p-value of 0.317, reflecting the inherent difficulty of distinguishing treatment effects from group-specific noise when only one group switches. The evidence for a causal effect rests on three pillars: the clean placebo for the primary outcome (prices show no pre-trends), the consistency of effects across four outcomes, and the institutional argument that the policy change was driven by a legalistic reinterpretation unrelated to procurement conditions. Placebo tests, alternative clustering, winsorization, Lee bounds, quantile DiD, and causal forest analysis provide further corroboration in the Appendix and Online Appendix.

This paper contributes to two strands of the literature. It adds to the evidence on the costs of restricting competition in public procurement (Bosio et al., 2022; Szucs, 2024; Decarolis et al., 2024; Bandiera et al., 2009; Best et al., 2023), providing estimates from a large developing economy where SME policies are especially prevalent. It also contributes to the literature on set-asides and bid preferences (Athey et al., 2013; Krasnokutskaya and Seim, 2011; Marion, 2007) by documenting how these costs vary across item types.

Section 2 provides the institutional background. Section 3 describes the data. Section 4 presents the empirical analysis, heterogeneity, robustness, and fiscal cost quantification. Section 5 concludes with policy implications.

2. Institutional Background

SMEs' participation in public procurement in Sao Paulo operates within a layered regulatory framework. I describe the relevant legislation and then explain how the state of Sao Paulo applied it to health items.

2.1. Public Procurement and SME Law in Brazil

Public procurement constitutes a relevant part of economies worldwide (Bosio et al., 2022). In 2016, OECD countries spent an average of approximately 12% of GDP on public procurement, while in Brazil, this proportion was approximately 10% in the same year. Recent evidence from Brazil shows that procurement institutions and their enforcement have significant effects on firm dynamics and employment (Szerman, 2023; Colonnelli and Prem, 2022).

As in many other countries, Brazilian law establishes as a general rule that all purchases, services, and works hired by the public administration should be subject to a public tender. Federal Law 8,666/1993 institutes a general framework applicable to all public bids in the country, and all three government branches must adhere to this framework.

Entities directly or indirectly controlled by the federal, state, or municipal governments must comply with the government procurement rules. Federal, state, and municipal governments, autonomous government entities, public foundations, regulatory agencies, state-owned companies, and mixed capital companies controlled by the government are subject to these rules. These entities are known as public buyer units (PBUs).

Although the public administration may decide to make purchases centrally, in Brazil, almost all acquisitions are decentralized and made by PBUs. A ministry or bureau may consist of many PBUs that have budgetary autonomy and make purchases from private companies. PBUs may contract a wide variety of products and services from private companies, including engineering and infrastructure work. However, this paper focuses on analyzing the acquisition of common and standardized goods.

The primary purpose of a bidding process conducted by a PBU is to seek the best contract possible for the government. The Brazilian public procurement law provides guidelines on how the procurement process should be organized and executed. In some cases, public tenders for SMEs are subject to different treatment.

The Brazilian federal SME law was enacted in 2006 to regulate favored, simplified and differentiated treatment for this sector, as provided for in the Federal Constitution. This law's explicit goal was to promote SMEs' economic and social development and competitiveness as a strategy for job creation, income distribution, social inclusion, reduced informality, and a strengthened economy.

The Brazilian SME law adopts the following classification for companies, according to their annual gross revenue: (i) microbusiness: annual gross revenue of R$360,000 or less (roughly US$72,000); and (ii) small business: annual gross revenue greater than R$360,000 and less than or equal to R$4,800,000 (between US$72,000 and US$960,000). In this paper, these companies are referred to as SMEs.

SMEs enjoy many benefits provided by law, including tax benefits and fewer bureaucratic requirements to adhere to. In addition, public tenders held at the federal, state, and municipal levels can grant differentiated and privileged treatment to SMEs to promote economic and social development, increase public policies' efficiency, and stimulate technological innovation.

The content of the SME law, in its 2006 version, indicated that the public administration *could* create tenders exclusively for the participation of SMEs in purchases in which the item value was up to R$80,000 (approximately US$16,000). Thus, choosing tenders for SMEs only was optional for PBUs.

However, the federal SME law underwent a significant change from SMEs' exclusivity in tenders in 2014. The term "could" was replaced with "must," making it mandatory to execute exclusive public tenders for SMEs up to a value per item of R$80,000.

The law changed PBUs' default choice if the item value fell below the threshold of eighty thousand reais: previously, the default choice was open bids. As of 2014, the standard option for PBUs is to execute public tenders for SMEs only, unless the bid's conditions fall within the exceptions provided for in the updated legislation.

PBUs can avoid restricted bids if at least one of the following conditions is met: (i) there are two or fewer potential competing SME suppliers that are locally or regionally based and able to comply with the notice requirements; or (ii) PBUs consider that the differentiated and simplified treatment for SMEs might not be advantageous for the public administration. Thus, PBUs choose whether the public tender is restricted, but they must justify their choices to their watchdogs, such as audit courts or the judiciary.

On the one hand, this discretion provided for by law can be beneficial since PBUs can more efficiently choose the bidding type to be carried out. However, there are costs involved in the process of avoiding bids restricted to SMEs. For each bidding procedure, PBUs must create an extensive report listing in detail the reasons that justify the use of an open bidding process to the detriment of a bidding process restricted to SMEs.

Additionally, these PBU justifications are subject to scrutiny by both the audit courts and the judiciary. If these bodies consider the arguments unfounded or insufficient, administrative proceedings and punishments may be brought against the public agents responsible for planning and executing the bid in question. Thus, this discretion brings costs to PBUs. I call these costs associated with avoiding SME-only tenders *opt-out costs*.

2.2. SME-only Public Tenders: Group 65 As an Exception

Sao Paulo is the wealthiest and most populous state in Brazil. This state accounts for approximately 23% of the total population and nearly one-third of Brazil's GDP (nearly US$500 billion in 2018). This amount is equivalent to the GDP of countries such as Sweden, Poland, and Belgium and more than twice the GDP of Portugal, Greece, and Finland. The state of Sao Paulo has a diversified economy driven by the automobile, textile, chemical, aeronautical, and computer industries, in addition to services such as finance and agriculture.

Since 2005, all PBUs in the state of Sao Paulo have been required to purchase common goods and services through Bolsa Eletrônica de Compras (BEC), an electronic purchasing platform. BEC operates as a centralized electronic reverse auction system (*pregão eletrônico*), in which registered suppliers submit progressively lower bids in real time until the auction closes, facilitating transparent price competition. This design is comparable to other electronic procurement platforms adopted worldwide, such as Chile's ChileCompra and South Korea's KONEPS, which similarly aim to reduce transaction costs, increase supplier participation, and improve price outcomes through digitized and standardized processes. The BEC figures are revealing: in 2019, approximately

R$13 billion (about US$3 billion) in trade was conducted on this e-platform. Since its implementation in 2005, BEC has moved more than R$105 billion (about US$20 billion) in negotiations; 860,000 purchase offers have been made, and approximately 4.8 million items have been sold.

Despite being subject to federal laws, Brazilian states have the prerogative to regulate or interpret these laws' specific elements. The state of Sao Paulo, for example, has a specific interpretation of how to apply SME law in public procurement.

Since BEC's implementation, the Sao Paulo state government has considered the group of items consisting of health-related products, including medication and hospital supplies (code 65), as a strategic set of items in the public procurement process. For example, bidding procedures related to medication can be carried out only with dynamic reverse auctions (*pregão*) or sealed bidding (*convite*). Direct negotiation (*dispensa de licitação*), which is very common in the purchase of other types of common goods, has always been prohibited for Group 65 in Sao Paulo.

Between 2006 and 2014, when the first version of the SME law was enforced, PBUs located in the state of Sao Paulo had the default choice to hold open tenders; it was optional to set procedures restricted to SMEs. During this period, in accordance with this legal arrangement, there was low adherence to restricted bids; items in these bids accounted for 6–13% of the total number of items bid for in the state of Sao Paulo. However, for Group 65, there were no bids exclusively for SMEs in this same period. The government had an internal orientation to hold open tenders for Group 65, with the explicit agreement of the audit court of the state of Sao Paulo (TCE-SP).

After 2014, with the update of the federal law on SMEs, the default choice was to execute SME-only tenders if the item value was less than or equal to R$80,000. In the state of Sao Paulo, if PBUs consider that any item in a bidding process falls within the exceptions provided for by law, they must justify in detail the reasons for the non-execution of an exclusive bid for SMEs through a report sent to TCE-SP.

This justification offered by PBUs to avoid restricted tenders not only requires excessive work effort for PBUs but also is subject to the scrutiny of the TCE-SP and the judiciary. Public officers can face punishment if there are irregularities or failure to comply strictly with the law.

After opt-out costs were introduced, adherence to this procedure has increased from approximately 13% to almost 70%, on average, in subsequent periods.

However, between August 2014 and February 2018, bids related to Group 65 did not change. The Sao Paulo state government and TCE-SP, in a joint agreement, used a specific interpretation of the new version of the federal SME law to leave Group 65 as an exception. The federal law provides that the requirement of restricted bids to SMEs does not apply when "it is not advantageous to the public administration or represents a loss to public resources." Thus, using this guideline, which allows for a high degree of discretion, the state of Sao Paulo and TCE-SP operated on the interpretation that because Group 65 constitutes a set of strategic items to meet such an essential public policy, it should always be subject to open bids. Thus, in this period, there was no bidding restricted to SMEs involving Group 65.

However, in November 2017, a different control agency in the state of Sao Paulo, PGE-SP, issued a document containing a legal opinion that changed this interpretation. This document reinforces that the principle of isonomy in law enforcement should prevail in public procurement. Therefore, items in Group 65 should be considered for public procurement purposes in the same way as other groups.

Hence, as of March 2018, the opt-out costs also apply to Group 65. Between March 2018 and December 2019, adherence to exclusive bids for SMEs for this group was, on average, 43%. This lower proportion of restricted bids for SMEs for Group 65 compared with other groups may be due to the existence of more oligopolies in this group of items, which include, for example, medication. Thus, in many cases, it is not possible to find at least three potential suppliers that are SMEs.

The trigger for the policy change was a legalistic reinterpretation of the isonomy principle by PGE-SP, not a response to deteriorating procurement outcomes for Group 65. There is no evidence that the decision to extend SME-only tender rules to Group 65 was motivated by price levels, competition patterns, or supplier complaints in this specific product group. The timing of the change was driven by bureaucratic processes within PGE-SP and its communication to PBUs, making it plausibly exogenous to the procurement outcomes studied in this paper. These features provide the institutional basis for the identifying assumption that, absent the policy change, Group 65 and other groups would have followed parallel outcome trajectories in the post-period.

3. Data

I use administrative data on bidding-level public procurement tenders of common goods and services in the state of Sao Paulo, Brazil, from January 2016 to December 2019. All transactions took place on the electronic procurement platform called BEC, available for all PBUs across the state. The SEFAZ/SP (State Finance Secretariat) is responsible for the operational management and centralization of BEC's bidding data. An important advantage of this data source is that BEC records the universe of standardized goods procurement in the state, eliminating selection concerns that arise in settings where data coverage is partial or voluntary.

On BEC, 1,344 PBUs regularly make purchases. These entities include state-level bureaus from the executive, legislative, and judiciary branches in the state of Sao Paulo as well as other affiliated entities, such as municipalities located in the state of Sao Paulo and private organizations. PBUs purchased 82,569 different items (goods) totaling 832,984 successful transactions from 2016 to 2019.

BEC has a very detailed catalog of standardized goods and services organized in three levels of detail: group, class, and item. For instance, health items are classified as Group 65 (medical, dental, and hospital equipment and supplies). The item coded as 110639 refers to the drug 'Furosemide 40 milligrams, coated tablets, units', belonging to class 6531 (Medicines prescribed with or without ANVISA notification/registration) and Group 65.

Data are organized by purchase offer (PO), the electronic document issued by the PBU that identifies and quantifies the goods and services that will be purchased. A PO is defined by a 22-character code and may contain one or more items listed, but each item has its own purchase process. Thus, the purchase of an item is uniquely identified by the combination of the PO and the purchased item codes (POI).

There is a crucial variable for the empirical section, defined from the item group codes. It is a binary variable that assumes the value of 1 if the item belongs to Group 65 and 0 otherwise. The items in Group 65 constitute the 'switched group', i.e., the set of items affected by the purchasing policy change. Group 65 accounted for almost 27% of all purchases from 2016 to 2019. All other groups of items comprise the 'control group.'

There are 75 groups of items, excluding Group 65. Between 2016 and 2019, the most significant groups were groups 89 (food products), 75 (office supplies), 86 (computer products), and 79 (cleaning materials), representing 37.5% of the total purchases in this period. For each POI, there is information about item quantities, bid prices (winners and losers), the number of participant firms, the number of bids, whether the public tender was successful or not, the identification of the PBU, and firms' and PBUs' location, among other variables.

The empirical analysis focuses on four main outcome variables. First, the *negotiated price* (in logs) captures the final price at which the item was transacted; this variable is available only for completed items. Second, the *number of participant firms* (in logs) measures how many distinct firms submitted at least one bid, regardless of whether the item was ultimately purchased. Third, the *number of valid bids* (in logs) counts the total number of qualifying bids received per item. Fourth, *distance* measures the geographic distance in kilometers between the winning firm's location and the PBU. The regressions control for the log quantity ordered, the type of tender (sealed bid vs. reverse auction), and include item-level fixed effects that absorb all time-invariant heterogeneity across items—including differences in markets, pricing norms, and typical supplier pools.

To ensure that the estimates are not driven by compositional changes in the sample, the analysis restricts attention to items that appear in both the pre- and post-periods within each time window. The three time windows (6, 12, and 18 months around the March 2018 cutoff) provide a check on the stability of the estimates as the sample expands. The price and distance regressions are estimated on the subsample of completed items (where a winning bid exists), while the participation and bid regressions use all items, including those where no firm was selected. An important implication of conditioning on completion is that the price and distance estimates may be affected by sample selection if the treatment itself influences whether items are successfully transacted; the Lee bounds analysis in the Appendix directly addresses this concern. Table 1 presents descriptive statistics for the key variables used in the empirical analysis, disaggregated by group (Group 65 vs. Others) and period (Pre vs. Post the policy shift in March 2018).

Table 1: Descriptive Statistics (18-month window)

	Group 65, Pre		Group 65, Post		Others, Pre		Others, Post	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price (levels)	1,402.09	20,563.54	1,507.92	25,872.31	1,465.85	50,385.93	1,291.37	68,042.55
Log price	1.72	2.84	1.91	2.81	2.85	2.12	2.78	2.07
Number of firms	3.17	2.52	3.02	2.40	4.66	3.27	4.69	3.28
Log number of firms	0.88	0.73	0.84	0.71	1.30	0.72	1.30	0.72
Number of valid bids	11.10	15.49	10.99	16.51	10.13	15.74	10.47	16.87
Log number of valid bids	1.64	1.24	1.61	1.23	1.69	1.06	1.70	1.08
Distance (km)	238.54	270.42	215.08	239.63	164.78	182.23	170.28	186.07
N (completed / all)	65,606	77,206	66,822	85,181	246,277	287,500	271,011	323,836

Notes: Price and distance statistics computed on completed items (status = 1). Firm and bid statistics computed on all items.

4. Empirical Strategy and Results

4.1. Main Specification and Identification Strategy

The identification strategy exploits the timing of the policy change (March 2018) that affected only Group 65, using time and cross-sectional variation in a difference-in-differences in reverse (DiDiR) design (Kim and Lee, 2019). Unlike a standard DiD, where the control group is never treated, the DiDiR control group is *always* treated: all product groups except Group 65 faced opt-out costs throughout the sample period, while Group 65 switched from no opt-out costs to positive opt-out costs in March 2018. This framework—also employed by Autor et al. (2006), Shapiro and Gentzkow (2008), —avoids the staggered-adoption concerns that may bias two-way fixed effects estimators with heterogeneous treatment effects (de Chaisemartin and D'Haultfoeuille, 2020; Roth et al., 2023).

As described in Section 3, from August 2014 to February 2018, PBUs faced opt-out costs to avoid tendering for SMEs only, with the exception of Group 65 (with mandatory open tenders). From March 2018 to December 2019, PBUs were subject to opt-out costs for all groups of items purchased. Thus, the always-treated group here consists of all groups of items but Group 65, which comprises the switched group. Table 2 summarizes the above description.

Table 2: Description of groups in DiDiR

Groups	$t = 1$ (before March 2018)	$t = 2$ (after March 2018)
Group 65 (switched group)	Opt-out costs $= 0$	Opt-out costs > 0
Others (always-treated group)	Opt-out costs > 0	Opt-out costs > 0

DiDiR identifies pre-switch-period effects—it estimates what Group 65 outcomes would have been if opt-out costs had been imposed before March 2018 (Kim and Lee, 2019). Using subindex p for purchase offer, i for items, t for months, and g for groups, the estimating equation is:

$$y_{pigt} = \eta_i + \gamma\,\mathrm{Pre}_t + \beta\left(g65_{pgt} \times \mathrm{Pre}_t\right) + x_{pigt}\,\delta + \varepsilon_{pigt} \tag{1}$$

where y is an outcome, η_i represents item fixed effects, and *Pre* is a dummy variable with value 1 if it is a month before March 2018 and 0 otherwise. The variable $g65$ is binary with a value of 1 if it belongs to Group 65 and 0 otherwise. The covariates are represented by x_{pigt}. Standard errors are clustered at the item level.

The coefficient of interest, β, captures the pre-switch-period effect of the shift in the SME tender policy on the outcomes for the switched group. Under the DiDiR framework, β measures the difference in outcomes for Group 65 relative to the counterfactual scenario in which Group 65 had always faced the same opt-out costs as other groups. A negative β for prices, for instance, indicates that prices were lower under open tenders than they would have been under SME-restricted tenders.

The central identifying assumption behind the empirical model is that, in the period after the shift in the tenders policy, the outcomes for Group 65 and the set of all other groups of items would have followed a similar trajectory. This assumption is testable in the post-period (where both groups face the same policy regime) and is supported by the institutional argument that the policy change was driven by a legalistic reinterpretation rather than by changes in the outcomes themselves. Thus, it is necessary to check whether the outcome paths of the always treated and switched groups are parallel in the post-switch period.

Anticipation is another potential threat. The PGE-SP legal opinion was issued in November 2017, roughly four months before its effective implementation in March 2018. If PBUs anticipated the policy change and adjusted their procurement behavior—for instance, by accelerating purchases of Group 65 items under open-tender rules or by pre-stocking medical supplies—the pre-period estimates could be contaminated. If anticipation led PBUs to front-load procurement under open tenders, this would increase the volume of open-tender transactions in the months immediately before the cutoff, *attenuating* rather than inflating the estimated price difference between Group 65 and other groups. In other words, anticipation biases the estimates toward zero, making them conservative. The stability of the DiDiR coefficients across the three time windows (6, 12, and 18 months) suggests that the estimates are not driven by behavior in the months immediately surrounding the policy change. If anticipation effects were quantitatively important, the 6-month window—which places the most weight on the months closest to the cutoff—would yield systematically different estimates than the 18-month window. The coefficients are stable across windows.

The validity of this assumption of parallel trends in the post-switch period can be partially assessed by estimating the following nonparametric regression, similar to that of Naritomi (2019) and Gallego et al. (2020):

$$y_{pgt} = \eta_g + \text{Semester}_t + \sum_{\substack{k=-3 \\ k\neq -1}}^{3} \beta_k \left(g65_{pgt} \times \mathbf{1}[t = k] \right) + \mu_{pgt} \tag{2}$$

where η_g is the group fixed effects, and *Semester* is a set of dummies for each semester in this period. The error μ_{pgt} is clustered by the group of items. Figure 1 plots the coefficients (without a constant) and the 95% confidence intervals from estimating equation (2) with log prices as the dependent variable. The graphs for the other outcomes of interest are included in the Appendix.

Although there are few periods observed after the change in the SME purchasing policy, it is possible to observe that the difference between the two groups narrows dramatically after the change in the purchasing policy. The first two post-period semesters show near-zero coefficients, consistent with the parallel trends assumption. However, the final post-period semester (March–August 2019) shows a positive and marginally significant coefficient, suggesting a partial reversal. Three interpretations are possible: PBU learning

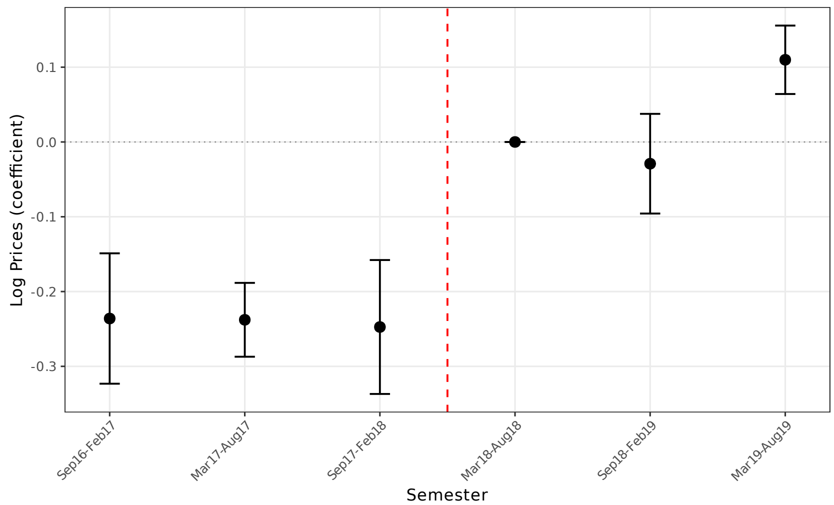

Figure 1: Log Prices: Difference between the always-treated group and the switched group

(buyers gradually mitigate efficiency costs as they adapt to SME-only rules), a confounding Group-65-specific shock in 2019, or the policy effect dissipating over time. The first interpretation is most consistent with the data: the main DiDiR coefficients are stable across the 6-, 12-, and 18-month windows (Table 4), and re-estimating while excluding this final semester yields virtually identical results (Online Appendix). If the effect were dissipating, the 18-month estimate would be systematically smaller than the 6-month estimate—it is not. Figures A.1–A.3 (Appendix) report equivalent event study results for distance, participating firms, and valid bids.

4.2. Results

I perform item-level regressions in a two-period DiDiR for which t is collapsed into pre and post periods. The estimations refer to the pre-switch-period effect for four distinct outcomes: negotiated prices, the number of participant firms, the number of valid bids, and the distance from PBUs to winner firms.

For each outcome, I run regressions for three different time windows considering different pre- and post-change periods: (i) a 6-month window where the pre-change period is from September 2017 to February 2018 and the post-change period is from

March 2018 to August 2018; (ii) a 12-month window where the pre-change period is from March 2017 to February 2018 and the post-change period is from March 2018 to February 2019; and (iii) an 18-month window where the pre-change period is from September 2016 to February 2018 and the post-change period is from March 2018 to August 2019. For each time window, I estimate a baseline specification with item fixed effects and a more saturated specification that adds buyer-unit (PBU) fixed effects to control for time-invariant heterogeneity across purchasing entities.

Table 3: Number of Participant Firms (log): Pre-switch-period effect on group 65

	(1) 6-month	(2) 6-month	(3) 12-month	(4) 12-month	(5) 18-month	(6) 18-month
$g65 \times Pre$	0.1776***	0.1821***	0.1495***	0.1540***	0.0926***	0.1004***
	(0.0079)	(0.0081)	(0.0062)	(0.0063)	(0.0059)	(0.0060)
sealed-bids	0.3651***	0.3730***	0.3377***	0.3443***	0.3431***	0.3566***
	(0.0082)	(0.0091)	(0.0075)	(0.0088)	(0.0073)	(0.0083)
lquantity	0.1944***	0.1763***	0.1952***	0.1782***	0.1934***	0.1774***
	(0.0024)	(0.0023)	(0.0021)	(0.0021)	(0.0020)	(0.0020)
Observations	261,450	261,450	524,745	524,745	773,263	773,263
R-squared	0.2135	0.1668	0.2097	0.1653	0.2058	0.1632
Item Fixed Effects	YES	YES	YES	YES	YES	YES
Controlling for PBU	NO	YES	NO	YES	NO	YES

Notes: Standard errors clustered at the item level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Constant absorbed by fixed effects.

Table 3 shows a consistent increase in companies participating in Group 65 bids when PBUs used open tenders. In the short term (6-month window), the number of participants is approximately 19–20% higher (coefficient of 0.178, implying $e^{0.178} - 1 = 19.4\%$), attenuating to approximately 10% at 18 months (coefficient of 0.093)—consistent with PBUs gradually adapting to the new regime. The pattern for valid bids is similar (Table B.1, Appendix).

These competition effects should be interpreted with caution: the placebo tests (Section 4.4) show significant pre-trends for both firms and bids, likely reflecting the secular rollout of SME restrictions across other groups in earlier periods. The price placebo passes cleanly, so the competition results are best read as suggestive evidence consistent with the price findings rather than independently identified causal effects.

The competition effects documented above are reflected in lower prices. Table 4 presents the price results.

Table 4: Prices (log): Pre-switch-period effect on group 65

	(1) 6-month	(2) 6-month	(3) 12-month	(4) 12-month	(5) 18-month	(6) 18-month
$g65 \times Pre$	-0.1311***	-0.1441***	-0.1370***	-0.1369***	-0.1309***	-0.1330***
	(0.0121)	(0.0116)	(0.0107)	(0.0104)	(0.0096)	(0.0094)
Sealed bids	-0.1461***	-0.2216***	-0.1430***	-0.2025***	-0.1501***	-0.2048***
	(0.0088)	(0.0131)	(0.0079)	(0.0123)	(0.0079)	(0.0116)
lquantity	-0.2598***	-0.2972***	-0.2611***	-0.2956***	-0.2597***	-0.2934***
	(0.0082)	(0.0089)	(0.0075)	(0.0080)	(0.0073)	(0.0077)
Observations	219,535	219,535	439,054	439,054	649,714	649,714
R-squared	0.1909	0.2124	0.1928	0.2108	0.1934	0.2108
Item Fixed Effects	YES	YES	YES	YES	YES	YES
Controlling for PBU	NO	YES	NO	YES	NO	YES

Notes: Standard errors clustered at the item level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Constant absorbed by fixed effects.

Across all specifications, the coefficients range from -0.131 to -0.144 log points, implying that negotiated prices are, on average, between 12% and 13% lower for Group 65 than for other groups before March 2018 (using the exact transformation $e^{\hat{\beta}} - 1$). The estimates are stable across the three time windows, suggesting that the effect is not an artifact of a particular sample period but rather a persistent feature of the procurement environment. These magnitudes are larger than those found in the set-aside literature in other contexts: Marion (2007) estimates a 3.8% cost increase from small business bid preferences in California highway procurement, and Krasnokutskaya and Seim (2011) document similar participation-cost trade-offs in US highway auctions. The larger estimates from Sao Paulo's procurement system are consistent with the idea that set-aside costs may be amplified in settings with thinner markets or greater supplier heterogeneity. The Gelbach decomposition in the Extensions section reveals that the competition channel documented above—more firms and a shift in winner composition—

accounts for only a small fraction (approximately 6%) of the total price effect, suggesting that most of the price reduction operates through fiercer bidding behavior rather than simply through more bidders.

The distance from PBUs to winner firms provides evidence on the distributional dimension of the policy: before the policy change, winning suppliers for Group 65 were located approximately 5–11 km further from PBUs, depending on the time window and specification (Table B.2, Appendix). The 6-month estimate (4.9 km) is not statistically significant, but the 12-month (9.7 km) and 18-month (10.9 km) baseline estimates are highly significant. With PBU fixed effects, the distance effect ranges from 4.6 to 5.3 km in the 12- and 18-month windows. This indicates that the SME restriction does encourage participation from more local suppliers.

Several additional analyses probe the robustness and heterogeneity of the price effect. In the Appendix, Lee bounds (Lee, 2009) correct for potential sample selection: after trimming 8.82% of the outcome distribution, the bounds range from -0.131 to -0.123, both highly significant, indicating that selection plays a negligible role (Table B.6). The HonestDiD sensitivity analysis (Rambachan and Roth, 2023) demonstrates that the main price effect remains statistically significant under substantial violations of the parallel trends assumption (Figure A.4).

The Online Appendix reports several further analyses. Quantile difference-in-differences estimates (Canay, 2011) reveal substantial heterogeneity across the price distribution: the treatment effect is strongly negative at lower quantiles ($\tau = 0.10$: -0.623; $\tau = 0.25$: -0.543) but turns positive at the upper tail ($\tau = 0.90$: $+0.445$). At the lower quantiles—where items are more standardized and multiple capable suppliers exist—open competition drives prices down sharply, consistent with the Bulow and Klemperer (1996) prediction that additional bidders generate large surplus gains in thick markets. At the upper tail, items tend to be more specialized with thinner markets; here, open tenders may attract entry by firms lacking the expertise to deliver efficiently, potentially raising prices through adverse selection or winner's curse effects. Uniform set-aside rules impose the largest efficiency costs precisely where competitive bidding is most effective, while offering little benefit for specialized procurement. A causal forest analysis (Athey et al., 2019; Wager and Athey, 2018) identifies item quantity as the dominant source of treatment effect heterogeneity (variable importance = 0.486). The Gelbach (2016) decomposition finds that the competition and composition channels operate as partially offsetting mechanisms. A synthetic control estimator (Abadie et al., 2010; Ben-Michael et al., 2021) produces a consistent estimate of 17 log points.

4.3. Heterogeneous Effects by Item Value

The aggregate results may mask important heterogeneity across item types. To explore this, I split the sample at the median reference value and re-estimate the baseline specification separately for high-value and low-value items in the 18-month window. Table 5 reports the results.

Table 5: Heterogeneous Effects by Item Value (18-month window, base specification)

	Log prices	Log firms	Log bids	Distance
Panel A: High-value items (above median)				
$g65 \times Pre$	-0.1369***	0.1075***	0.0456***	12.0849***
	(0.0103)	(0.0072)	(0.0093)	(2.9018)
Observations	383,949	445,903	445,903	383,949
Panel B: Low-value items (below median)				
$g65 \times Pre$	-0.0981***	0.0592***	0.0577***	2.8312
	(0.0086)	(0.0075)	(0.0110)	(2.2944)
Observations	265,765	327,360	327,360	265,765
Item FE	YES	YES	YES	YES

Notes: 18-month window, base specification. Sample split at median reference value. Standard errors clustered at the item level. *** p<0.01, ** p<0.05, * p<0.1.

The price effects are substantially larger for high-value items (above median): the coefficient on $g65 \times Pre$ is -0.1369 (Panel A) compared to -0.0981 for low-value items

(Panel B), a difference of nearly 40%. Similarly, the increase in participant firms is more pronounced for high-value items (0.1075 vs. 0.0592). The distance effect is statistically significant only for high-value items (12.08 km, $p < 0.01$), while the point estimate for low-value items (2.83 km) is not statistically distinguishable from zero. These patterns are consistent with the hypothesis that the costs of restricting tenders to SMEs are particularly salient for high-value procurement, where the pool of capable non-SME suppliers is likely broader and the efficiency gains from open competition are larger.

4.4. Robustness Checks

A natural concern with the DiDiR design is that the estimated pre-period effects could reflect pre-existing differential trends rather than the causal impact of the policy regime. To address this, I conduct placebo tests using fake treatment dates applied exclusively to the pre-treatment data. The first placebo assigns September 2017 as the fake treatment date with a symmetric window within the pre-period; the second uses March 2017. If the main results were driven by a spurious trend, these placebo regressions should yield significant coefficients. Table 6 reports the results.

Table 6: Placebo Tests: Fake Treatment Dates

	Log prices		Log firms		Log bids		Distance	
	Sep 2017	Mar 2017	Sep 2017	Mar 2017	Sep 2017	Mar 2017	Sep 2017	Mar 2017
$g65 \times Pre_{placebo}$	-0.0145	0.0206*	-0.1038***	-0.1346***	-0.1363***	-0.1343***	3.7807	-1.2018
	(0.0149)	(0.0112)	(0.0063)	(0.0070)	(0.0092)	(0.0106)	(2.4452)	(2.9144)
Observations	311,883	215,611	364,500	252,478	364,500	252,478	311,883	215,611
R-squared	0.1952	0.1962	0.2170	0.2274	0.2289	0.2352	0.0007	0.0005
Item FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Placebo tests using fake treatment dates (Sep 2017 and Mar 2017) on pre-treatment data only. Standard errors clustered at the item level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

The placebo coefficients for log prices are small in magnitude. The September 2017 placebo (-0.015) is statistically insignificant, while the March 2017 placebo (0.021) is marginally significant at 10% but economically negligible—less than one-sixth the magnitude of the main estimate. These results are consistent with the parallel trends assumption. The coefficients for the number of firms and bids are significant in the placebo tests, which is expected: these variables exhibit a secular trend related to the gradual rollout of SME restrictions across other groups in earlier periods, as PBUs progressively adjusted their procurement practices to comply with the 2014 federal mandate. The price outcome—the primary measure of procurement cost—shows no evidence of differential pre-trends, confirming that the estimated price effects are not an artifact of group-specific time trends.

The main results survive several additional robustness checks reported in the Appendix. First, the baseline standard errors are clustered at the item level. However, treatment is assigned at the item-group level (76 groups), which is a coarser level than the clustering unit. To assess whether this understates standard errors, I re-estimate the main specifications clustering at the item group level, at the PBU level, and with two-way clustering at the item and PBU levels simultaneously (Table B.3). Group-level clustering roughly doubles the standard errors (e.g., 0.0178 vs. 0.0096 for the 18-month price coefficient), but all primary results remain statistically significant at conventional levels. The t-statistic for the price effect under group-level clustering is approximately 7.3, well above conventional thresholds. A score cluster bootstrap with 9,999 Rademacher replications confirms $p < 0.001$ for all four outcomes (Table B.5, Appendix). Second, winsorizing the dependent variables at the 1st/99th and 5th/95th percentiles produces estimates of similar magnitude and significance (Table B.4), ruling out the possibility that the results are driven by extreme values in the tails of the price or distance distributions. Third, I conduct a randomization inference exercise by randomly permuting the treatment assignment (Group 65 vs. other groups) 500 times and re-estimating the main price regression. The permutation p-value is 0.317 (Figure A.8), which reflects the inherently low power of permutation tests with a single treated unit rather than evidence against the main findings: with 76 product groups and only one assigned to treatment, many permutations assign the "fake treatment" to large groups with strong idiosyncratic price trends, generating a fat-tailed permutation distribution. The primary reassurance against group-specific confounds comes instead from the four-outcome consistency test and the placebo analysis discussed above.

As a complementary estimator, I construct a synthetic control for Group 65 using the 69 product groups that appear in both the pre- and post-periods as donors (Abadie et al., 2010; Ben-Michael et al., 2021). The synthetic control matches Group 65 prices almost exactly in the pre-period and estimates that, after losing the open-tender exemption, Group 65 prices rose by approximately 17 log points relative to the synthetic—consistent in direction and magnitude with the DiDiR estimate of 13 log points. The in-space placebo rank p-value is 0.103 (Online Appendix).

A limitation of the DiDiR design—and the central threat to identification—is that Group 65 is the sole switched unit. Any group-65-specific shock coinciding with March 2018 could confound the estimates. Two candidate confounders deserve explicit consideration. First, ANVISA (Brazil's pharmaceutical regulator) could have changed registration requirements or supply conditions for medical products around the same date. A review of ANVISA regulatory actions in 2017–2018 reveals no major policy change affecting the supply or pricing of Group 65 items in this period; the CMED annual price adjustment follows a predictable calendar (published each April) and applies uniformly across pharmaceutical products regardless of the procurement regime. Second, supply chain disruptions specific to medical supplies could generate a contemporaneous shock. However, such disruptions would plausibly raise prices and reduce participation for Group 65 in the post-period, which is the *opposite* of the convergence pattern observed in the event study.

Three observations bear on this. The randomization inference exercise (Figure A.8) provides partial reassurance, though the permutation p-value of 0.317 reflects the low power inherent in single-unit permutation tests (see Section 4.4 for discussion). The consistency of the results across four outcomes—prices, participation, bids, and distance—imposes a stringent test: a single confounding shock would need to simultaneously produce all four effects. No single supply- or demand-side shock plausibly generates this pattern. The estimates are therefore most directly informative about the costs of SME restrictions in the specific context of medical and hospital supplies, and extrapolation to other product groups requires the (untested) assumption of comparable treatment effects.

A related concern is cross-group spillovers. When Group 65 became subject to SME-only rules in March 2018, non-SME firms that had previously competed for Group 65 contracts may have redirected their bidding activity toward other product groups, increasing competition in those groups and thereby lowering prices in the always-treated category. Such spillovers would *narrow* the estimated gap between Group 65 and other groups, biasing the DiDiR coefficient toward zero. The estimates should therefore be interpreted as a lower bound on the true cost of the SME restriction, to the extent that spillovers are present.

4.5. Extensions

Real prices. A potential concern is that the nominal price effects could partly reflect differential inflation trends across product groups. To address this, I deflate prices using the IPCA (Brazil's consumer price index) and re-estimate the full set of specifications (Online Appendix, Table OA.1). The IPCA-deflated estimates are qualitatively similar but smaller in magnitude, and the gap widens with the time window: the 6-month real coefficient (-0.108) is 18% smaller than the nominal estimate (-0.131), while the 18-month real coefficient (-0.076) is 42% smaller (-0.131). This divergence likely reflects differential inflation between medical supplies—Group 65, which includes pharmaceuticals subject to Câmara de Regulação do Mercado de Medicamentos (CMED) price ceilings—and other product categories. While the real-price specifications partially address this concern, residual differential inflation cannot be fully excluded. The real-price estimates provide a lower bound on the procurement cost of SME restrictions, and the fiscal cost calculation should be interpreted accordingly (see Section 4.6).

Winner composition. If the main channel operates through increased competition from non-SME firms, the composition of winners should shift accordingly. Indeed, the probability that the winning firm is an SME decreases under open tenders (Online Appendix, Table OA.4), directly supporting the competition mechanism: open tenders attract larger firms that outbid SMEs on price.

Heterogeneity by PBU type. The effects may differ across buyer types if, for example, direct administration entities (state bureaus) face different institutional constraints than indirect administration entities (foundations, municipalities). I test for this by interacting the treatment indicator with a direct administration dummy. The price effects are roughly twice as large for direct administration PBUs ($-0.072 + (-0.066) = -0.138$) as for indirect administration entities (-0.072), suggesting that institutional constraints or procurement practices within the state bureaucracy amplify the costs of SME restrictions (Online Appendix, Table OA.5). The firm participation interaction is also positive and significant for direct administration ($+0.118$).

Mechanism evidence. Three additional analyses in the Online Appendix shed light on the channels driving the price results. First, splitting the sample by item quantity reveals a dose-response pattern: the increase in participating firms is five times larger for high-quantity items (12.4% vs. 2.4%), consistent with the causal forest finding that item quantity is the dominant moderator. Second, decomposing Group 65 into medications (class 6531) and other medical supplies reveals that the price effect is 67% larger for medications (-0.166 vs. -0.099), which could reflect either greater market power in the oligopolistic pharmaceutical sector or residual CMED-driven inflation.

Open tenders attract more firms, including larger non-SME suppliers, and this additional competition drives prices down. But how much of the 12–13% price reduction do these observable channels actually explain? A Gelbach (2016) mediation analysis shows that measured channels—firm entry and winner composition—account for only approximately 6% of the total price effect. The vast majority of the price reduction under open tenders operates through channels *beyond* the simple fact that more firms show up or that non-SME firms win more often. The remaining 94% likely reflects two channels identified by Bulow and Klemperer (1996) and Krasnokutskaya and Seim (2011) : *bid aggressiveness conditional on entry*: in a thicker market, each participating firm faces greater competitive pressure and submits more aggressive bids, even holding the number of bidders constant. This intensive-margin channel is predicted by standard auction theory but is not captured by the number-of-firms mediator, which measures only the extensive margin. and *supplier matching efficiency*: open tenders may attract not just more firms but *better-matched* firms—suppliers whose cost structures are particularly well suited to the item being procured. Under SME-only rules, the lowest-cost eligible supplier may be a large firm that is excluded from bidding, leaving the contract to a less efficient SME. Table 7 provides direct evidence for the bid aggressiveness channel. Columns (1)–(3) sequentially add competition mediators to the baseline price regression: controlling for the log number of firms and winner SME status reduces the treatment effect by only 6.3%, confirming the Gelbach result. Columns (4)–(6) provide a more direct test: hold the extensive margin *exactly* constant by restricting the sample to items with a fixed number of participating firms. Among items with exactly 2 firms, the price effect is -0.108 ($p < 0.01$); with exactly 3 firms, -0.151; with 5 or more, -0.122. The larger point estimate for the 3-firm subsample likely reflects sample composition: items that attract exactly 3 bidders under both policy regimes tend to be those where competition is most intense, amplifying the intensive-margin channel. The persistence of the price effect across all subsamples—where the number of bidders is identical between Group 65 and other groups—demonstrates that the cost reduction under open tenders operates primarily through the intensive margin: firms bid more aggressively when non-SME competitors are eligible, even when the same number of firms ultimately participate.

Table 7: Bid Aggressiveness: Price Effects Conditional on Competition

	Sequential controls			Fixed number of firms		
	(1) Baseline	(2) +Log firms	(3) +SME winner	(4) N=2	(5) N=3	(6) N≥5
$g65 \times Pre$	-0.1309***	-0.1395***	-0.1227***	-0.1077***	-0.1509***	-0.1223***
	(0.0096)	(0.0099)	(0.0101)	(0.0138)	(0.0153)	(0.0191)
Observations	625,627	624,592	624,592	89,803	87,783	246,590
R-squared	0.9142	0.9148	0.9149	0.9173	0.9176	0.9299
Item Fixed Effects	YES	YES	YES	YES	YES	YES

Notes: 18-month window, completed items. Standard errors clustered at the item level in parentheses. Columns (1)–(3) sequentially add competition mediators to the baseline price regression. Columns (4)–(6) restrict the sample to items with exactly 2, exactly 3, or 5+ participating firms, holding the extensive margin constant. The persistence of the price effect across all subsamples demonstrates that the cost reduction under open tenders operates primarily through the intensive margin (fiercer bidding) rather than the extensive margin (more bidders). *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Additional extensions in the Online Appendix examine the extensive margin—open tenders increase item completion rates by 10.7–12.7 percentage points—and procurement efficiency ratios relative to government reference prices.

Raw monthly trends for all four outcomes (Online Appendix, Figures OA.1–OA.4) visually confirm the parallel trends assumption in the post-treatment period and the divergence in the pre-period for the switched group.

4.6. Fiscal Cost Quantification

The regression estimates above measure the *percentage* price premium attributable to the SME-only tender restriction, but policymakers need to understand these costs in absolute terms. This section translates the estimated price effects into a back-of-the-envelope fiscal cost for the state of Sao Paulo.

Let $\hat{\beta}$ denote the estimated DiDiR coefficient on log prices from the 18-month window specification. Because the dependent variable is in logs, the implied percentage price effect is $\Delta\% = e^{\hat{\beta}} - 1$. Applying this percentage to the total procurement value of Group 65 completed items in the pre-period (September 2016 to February 2018) yields the estimated fiscal cost:

$$\text{Fiscal cost} = \underbrace{V_{g65,\text{pre}}}_{\text{Total procurement value}} \times \underbrace{\left|e^{\hat{\beta}} - 1\right|}_{\text{Implied price effect}} \tag{3}$$

Table 8 presents the results for both the baseline specification (item fixed effects only) and the saturated specification (item and PBU fixed effects).

Table 8: Fiscal Cost of SME-only Tender Restrictions: Group 65

	Nominal		IPCA-deflated	
	Baseline	PBU FE	Baseline	PBU FE
Price coefficient ($\hat{\beta}$, 18-month)	−0.1309	−0.1330	−0.0762	−0.0795
Implied price effect ($e^{\hat{\beta}} - 1$)	−12.27%	−12.45%	−7.33%	−7.64%
G65 pre-period total value (R$ millions)	689.0	689.0	689.0	689.0
Estimated fiscal saving (R$ millions)	**84.5**	**85.8**	**50.5**	**52.6**

Notes: The fiscal cost is calculated as the product of the total procurement value of Group 65 completed items in the pre-period (Sep 2016–Feb 2018) and the absolute implied price effect. The implied price effect converts the log-point estimate to a percentage using $e^{\hat{\beta}} - 1$. Both specifications use the 18-month time window. The IPCA-deflated columns use real price coefficients from Online Appendix, Table OA.1. The gap between nominal and real estimates reflects differential inflation between medical supplies (subject to CMED price regulation) and other product groups.

The estimated fiscal cost of the SME-only tender restriction for Group 65 alone ranges from R$84.5 million to R$85.8 million in nominal terms (approximately US$17 million) over the 18-month pre-period, representing about 12% of the total procurement value for this product group. Using IPCA-deflated prices, the fiscal cost falls to R$50.5–52.6 million, suggesting that between one-third and one-half of the nominal gap reflects differential inflation rather than the direct effect of competition restrictions. Both estimates point to a substantial inefficiency arising from a single policy instrument applied to a single set of items.

Several considerations frame the interpretation of this estimate. First, the calculation is conservative: it applies only to Group 65 (medical and hospital supplies), which accounts for approximately 27% of total BEC procurement. Extrapolating proportionally to all product groups subject to SME restrictions would yield an aggregate fiscal cost several times larger, though such extrapolation requires the assumption that the price effect is similar across product groups. Second, the heterogeneity results in Table 5 suggest that the fiscal cost is disproportionately concentrated among high-value items, where the per-item price premium is approximately 40% larger. Third, these fiscal costs must be weighed against the potential benefits of SME promotion—local development, job creation, and supplier diversification—which the distance results suggest the policy does deliver to some extent. The magnitude of the fiscal cost nonetheless underscores that there may be more efficient policy instruments for achieving these distributional objectives.

5. Conclusion

Restricting public tenders to SMEs costs the state of Sao Paulo an estimated R$50–85 million for one product group over 18 months (R$50 million in real terms, R$85 million nominal), with a rough extrapolation to all groups on the order of R$310 million—though this assumes comparable treatment effects across product groups, which the heterogeneity results suggest may not hold. This estimate, derived from a natural experiment in which medical supplies were exempt from mandatory SME-only tenders until a legal reinterpretation in March 2018, is robust to placebo tests, alternative clustering, sample selection correction, and sensitivity analysis. Open tenders yielded prices 7–13% lower (depending on inflation adjustment), attracted 19–20% more firms, and generated 16–17% more valid bids. Most of the price reduction operates through the intensive margin—fiercer bidding—rather than through additional entry, as a Gelbach decomposition and a direct bid aggressiveness test confirm.

These costs are not uniform. The price effects are approximately 40% larger for high-value items, and the distance effects are statistically significant only for the high-value subsample—consistent with the idea that restricting competition matters most where a broad pool of capable non-SME suppliers exists. The one dimension along which the policy delivers on its stated objectives is geographic proximity: winning firms for Group 65 were located 5–11 km further from PBUs before the restriction, indicating that SME-only tenders do encourage local sourcing.

The fiscal cost for Group 65 alone amounts to R$84.5–85.8 million in nominal terms over 18 months (R$50.5 million in real terms); extrapolating to all product groups yields approximately R$310 million. This extrapolation assumes comparable treatment effects across groups, so it should be read as an order of magnitude.

The mechanisms documented here—reduced competition, weaker bidding, worse supplier matching—are generic consequences of restricting the supplier pool (Bulow and Klemperer, 1996; Krasnokutskaya and Seim, 2011) and are likely to operate wherever set-asides bind. The results have several policy implications.

The heterogeneity and quantile DiD results suggest that *uniform* set-aside rules are suboptimal. Efficiency losses are concentrated among high-value, standardized items at the lower end of the price distribution, where competitive bidding is most effective. At the upper tail—specialized items with thin supplier markets—open tenders may raise prices through adverse selection. A differentiated policy that exempts high-value standardized items from SME-only rules while maintaining preferences for smaller or specialized procurement would better balance efficiency and distributional objectives.

The institutional design of *opt-out costs* provides a direct policy lever. The current system penalizes PBUs that wish to conduct open tenders by imposing costly justification requirements. Reducing these costs—for instance, by allowing PBUs to default to open tenders above a value threshold without audit court approval—would lower procurement costs while preserving the SME preference where it is less distortionary.

More broadly, alternative instruments to promote SME participation—direct subsidies, simplified registration, or capacity-building programs—may achieve similar developmental goals at lower fiscal cost (Athey et al., 2013; Krasnokutskaya and Seim, 2011; Marion, 2007). The distance results confirm that SME-only rules succeed in promoting local sourcing, but the 7–13% price premium suggests that this geographic benefit comes at considerable fiscal cost.

These findings speak directly to the Brazilian context. Federal Law 14,133/2021—Brazil's new public procurement framework, gradually replacing Law 8,666/1993—maintains mandatory SME preferences but introduces new provisions for electronic procurement and framework agreements. The quantile DiD results offer a concrete guide for implementation: exempting high-value standardized items from SME-only rules would capture most of the efficiency gains (the price effect is -0.62 at $\tau = 0.10$) while preserving preferences for specialized procurement where open tenders may actually raise prices ($+0.45$ at $\tau = 0.90$). The opt-out cost mechanism documented here also suggests a specific reform lever: allowing PBUs to default to open tenders above a value threshold—without audit court justification—would reduce procurement costs where restrictions are most distortionary.

Several avenues for future research remain open. Longer post-treatment data would permit a sharper evaluation of whether PBUs learn to mitigate the costs of the restriction over time. Firm-level data would allow researchers to trace the dynamic effects on SME growth, entry, and exit—the supply-side benefits these policies are designed to foster. Structural estimation of the auction model could quantify the welfare consequences of SME preferences under alternative market structures. The present analysis shows that procurement restrictions in favor of small firms carry a measurable fiscal cost—one that policymakers should weigh against the distributional benefits these policies deliver.

References

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association 105*(490), 493–505.

Athey, S., D. Coey, and J. Levin (2013). Set-asides and subsidies in auctions. *American Economic Journal: Microeconomics 5*(1), 1–27.

Athey, S., J. Tibshirani, and S. Wager (2019). Generalized random forests. *The Annals of Statistics 47*(2), 1148–1178.

Autor, D. H., J. Donohue III, and S. J. Schwab (2006). The costs of wrongful-discharge laws. *The Review of Economics and Statistics 88*, 211–231.

Bandiera, O., A. Prat, and T. Valletti (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review 99*(4), 1278–1308.

Bastos, A., S. Guimarães, K. C. Medeiros De Carvalho, L. Andrés, and R. Paixão (2018). Micro, pequenas e médias empresas: Conceitos e estatísticas. *IPEA Radar: Tecnologia, Produção e Comércio Exterior 55*, 21–26.

Ben-Michael, E., A. Feller, and J. Rothstein (2021). The augmented synthetic control method. *Journal of the American Statistical Association 116*(536), 1789–1803.

Best, M. C., J. Hjort, and D. Szakonyi (2023). Individuals and organizations as sources of state effectiveness. *American Economic Review 113*(8), 2121–2167.

Bosio, E., S. Djankov, E. Glaeser, and A. Shleifer (2022). Public procurement in law and practice. *American Economic Review 112*(4), 1091–1117.

Bulow, J. and P. Klemperer (1996). Auctions versus negotiations. *American Economic Review 86*(1), 180–194.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics 90*(3), 414–427.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *The Econometrics Journal 14*(3), 368–386.

Cardoza, G., G. Fornes, V. Farber, R. Gonzalez Duarte, and J. Ruiz Gutierrez (2016). Barriers and public policies affecting the international expansion of latin american smes: Evidence from brazil, colombia, and peru. *Journal of Business Research 69*(6), 2030–2039.

Colonnelli, E. and M. Prem (2022). Corruption and firms. *Review of Economic Studies 89*(2), 695–732.

de Chaisemartin, C. and X. D'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Decarolis, F., R. Fisman, P. Pinotti, and S. Vannutelli (2024). Rules, discretion, and corruption in procurement: Evidence from italian government contracting. *Journal of Political Economy: Microeconomics 3*(2), 213–254.

Fafchamps, M., D. McKenzie, S. Quinn, and C. Woodruff (2014). Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in ghana. *Journal of Development Economics 106*, 211–226.

Freedman, M. (2013). Targeted business incentives and local labor markets. *Journal of Human Resources 48*(2), 311–344.

Gallego, J. A., M. Prem, and J. F. Vargas (2020). Corruption in the times of pandemia. *Universidad del Rosario 1*(18178), 1–36.

Kim, K. and M. j. Lee (2019). Difference in differences in reverse. *Empirical Economics 57*(3), 705–725.

Krasnokutskaya, E. and K. Seim (2011). Bid preference programs and participation in highway procurement auctions. *American Economic Review 101*(6), 2653–2686.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies 76*(3), 1071–1102.

Marion, J. (2007). Are bid preferences benign? the effect of small business subsidies in highway procurement auctions. *Journal of Public Economics 91*(7–8), 1591–1624.

McKenzie, D. and C. Woodruff (2015). Business practices in small firms in developing countries. Working paper or report.

Mel, S. d., D. McKenzie, and C. Woodruff (2008). Returns to capital in microenterprises: Evidence from a field experiment. *Quarterly Journal of Economics 123*(4), 1329–1372.

Nakabayashi, J. (2013). Small business set-asides in procurement auctions: An empirical analysis. *Journal of Public Economics 100*, 28–44.

Naritomi, J. (2019). Consumers as tax auditors. *American Economic Review 109*(9), 3031–3072.

PwC (2014). Smes' access to public procurement markets and aggregation of demand in the eu. Retrieved from http://ec.europa.eu/internal_market/publicprocurement/docs/modernising_rules/smes-access-and-aggregation-of-demand_en.pdf.

Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *The Review of Economic Studies 90*(5), 2555–2591.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Shapiro, J. and M. Gentzkow (2008). Preschool television viewing and adolescent test scores: Historical evidence from the coleman study. *Quarterly Journal of Economics 123*(1), 279–323.

Szerman, C. (2023). The employee costs of corporate debarment in public procurement. *American Economic Journal: Applied Economics 15*(1), 411–441.

Szucs, F. (2024). Discretion and favoritism in public procurement. *Journal of the European Economic Association 22*(1), 117–160.

Thai, K. V. (2017). *Global Public Procurement: Theories and Practices.* Springer International Publishing.

Wager, S. and S. Athey (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association 113*(523), 1228–1242.

Appendix

A. Event Study Coefficients

Figures A.1–A.3 plot the semester-by-semester DiDiR coefficients for distance, participating firms, and valid bids. The vertical dashed line marks the policy change in March 2018.

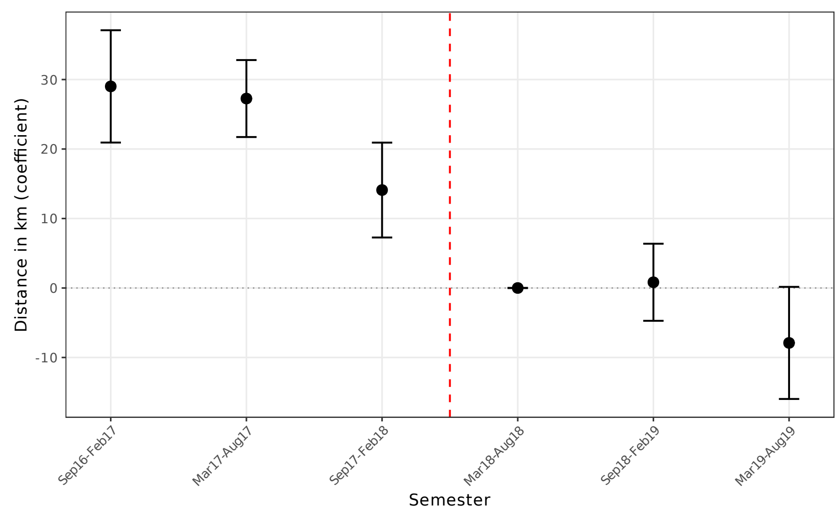

Figure A.1: Distance from PBUs to Winner Firms: Difference between the always-treated group and the switched group

B. Additional Outcome Variables

Tables B.1 and B.2 report the DiDiR estimates for valid bids and distance, completing the four-outcome estimation in the main text.

C. Robustness Checks

Tables B.3 and B.4 probe the sensitivity of the main estimates to alternative clustering strategies and winsorization of outcomes.

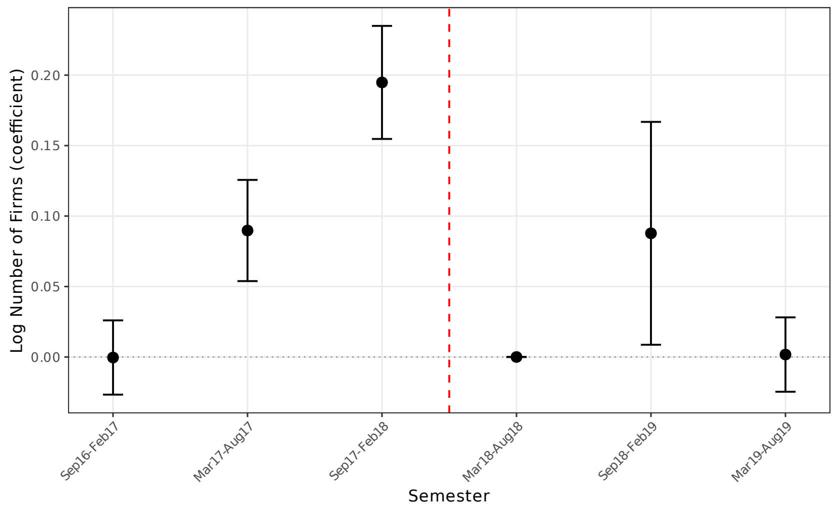

Figure A.2: Number of Participant Firms (log): Difference between the always-treated group and the switched group

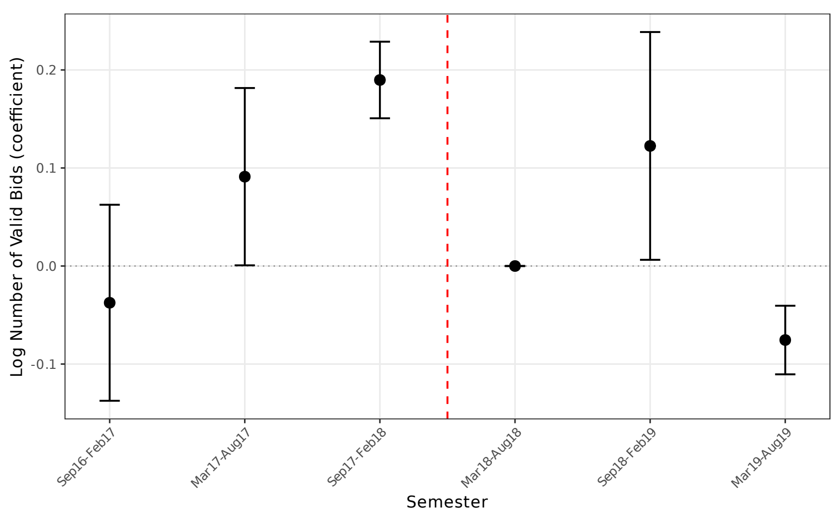

Figure A.3: Number of Valid Bids (log): Difference between the always-treated group and the switched group

Table B.1: Number of Valid Bids (log): Pre-switch-period effect on group 65

	(1) 6-month	(2) 6-month	(3) 12-month	(4) 12-month	(5) 18-month	(6) 18-month
$g65 \times Pre$	0.1524***	0.1533***	0.1011***	0.1078***	0.0511***	0.0603***
	(0.0108)	(0.0110)	(0.0082)	(0.0083)	(0.0077)	(0.0078)
sealed-bids	-0.6431***	-0.6213***	-0.6708***	-0.6516***	-0.6394***	-0.6161***
	(0.0109)	(0.0129)	(0.0103)	(0.0124)	(0.0096)	(0.0112)
lquantity	0.2303***	0.2078***	0.2299***	0.2090***	0.2297***	0.2095***
	(0.0034)	(0.0033)	(0.0029)	(0.0028)	(0.0028)	(0.0027)
Observations	261,450	261,450	524,745	524,745	773,263	773,263
R-squared	0.2239	0.1575	0.2279	0.1593	0.2186	0.1525
Item Fixed Effects	YES	YES	YES	YES	YES	YES
Controlling for PBU	NO	YES	NO	YES	NO	YES

Notes: Standard errors clustered at the item level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Constant absorbed by fixed effects.

Table B.2: Distance from PBUs to winner firms: Pre-switch-period effect on group 65

	(1) 6-month	(2) 6-month	(3) 12-month	(4) 12-month	(5) 18-month	(6) 18-month
$g65 \times Pre$	4.8999	1.0072	9.7156***	4.6172*	10.8570***	5.2865**
	(2.9860)	(2.9948)	(2.3938)	(2.3967)	(2.2326)	(2.2277)
convite	-10.0185***	-3.8438**	-8.6811***	-1.7095	-9.0559***	-1.6071
	(1.6623)	(1.7136)	(1.4358)	(1.4242)	(1.3238)	(1.2758)
lquantidade	1.7198***	3.6041***	1.9968***	4.1204***	1.5738***	4.2666***
	(0.6450)	(0.5207)	(0.5693)	(0.4237)	(0.5389)	(0.3709)
Observations	219,535	219,535	439,054	439,054	649,714	649,714
R-squared	0.0008	0.0008	0.0008	0.0010	0.0008	0.0010
Item Fixed Effects	YES	YES	YES	YES	YES	YES
Controlling for PBU	NO	YES	NO	YES	NO	YES

Notes: Standard errors clustered at the item level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Constant absorbed by fixed effects.

Table B.3: Alternative Clustering (18-month window, base specification)

	Log prices	Log firms	Log bids	Distance
Panel A: Cluster by item group				
$g65 \times Pre$	-0.1309***	0.0926***	0.0511***	10.8570***
	(0.0178)	(0.0024)	(0.0028)	(0.4079)
Panel B: Cluster by PBU				
$g65 \times Pre$	-0.1309***	0.0926***	0.0511***	10.8570***
	(0.0137)	(0.0188)	(0.0178)	(4.1277)
Panel C: Two-way clustering (item × PBU)				
$g65 \times Pre$	-0.1309***	0.0926***	0.0511***	10.8570**
	(0.0153)	(0.0186)	(0.0179)	(4.2771)
Observations	649,714	773,263	773,263	649,714

Notes: 18-month window, base specification (item FE only). Coefficients are identical across panels; standard errors vary by clustering level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.4: Winsorized Regressions (18-month window)

	Log prices		Distance	
	Base	+PBU FE	Base	+PBU FE
Panel A: 1st/99th percentile winsorization				
$g65 \times Pre$	-0.1218***	-0.1239***	7.4267***	2.3852
	(0.0086)	(0.0085)	(1.8574)	(1.8531)
Panel B: 5th/95th percentile winsorization				
$g65 \times Pre$	-0.0923***	-0.0931***	5.4278***	0.8886
	(0.0068)	(0.0067)	(1.5113)	(1.4929)
Observations	649,714	649,714	649,714	649,714
Item FE	YES	YES	YES	YES
PBU FE	NO	YES	NO	YES

Notes: 18-month window. Dependent variables winsorized at indicated percentiles. Standard errors clustered at the item level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

D. Score Cluster Bootstrap

Table B.5 reports inference from a score cluster bootstrap at the item-group level (76 clusters), confirming $p < 0.001$ for all four outcomes.

Table B.5: Score Cluster Bootstrap Inference (Group-Level, 76 Clusters)

	Log prices	Log firms	Log bids	Distance
$g65 \times Pre$	-0.1309***	0.0926***	0.0511***	10.8570***
Group-clustered SE	(0.0178)	(0.0024)	(0.0028)	(0.4079)
Bootstrap p-value	0.0000	0.0000	0.0000	0.0000
Bootstrap 95% CI	[-0.1487, -0.1131]	[0.0902, 0.0950]	[0.0482, 0.0539]	[10.4491, 11.2649]
Observations	649,714	773,263	773,263	649,714

Notes: 18-month window, baseline specification (item FE). Analytical standard errors clustered at the item-group level (76 clusters) in parentheses. Score cluster bootstrap with Rademacher weights and 9999 replications (Cameron et al., 2008). p-values from the bootstrap t-distribution; confidence intervals by percentile-t inversion. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E. Sensitivity to Parallel Trends Violations

Figure A.4 displays the Rambachan and Roth (2023) sensitivity analysis (HonestDiD), showing that the main price effect remains statistically significant even under substantial deviations from strict parallel trends.

F. Sample Selection Correction (Lee Bounds)

Table B.6 applies the Lee (2009) bounding procedure. The bounds for log prices range from -0.131 to -0.123, both highly significant, indicating that sample selection has a negligible impact on the main estimates.

Additional results—including raw monthly trends, randomization inference, SME participation trends, IPCA-deflated prices, extensive margin, procurement efficiency, winner composition, PBU heterogeneity, causal forest analysis, quantile DiD, Gelbach decomposition, balance tests, excluding the last post-period semester, mechanism evidence, and synthetic control estimates—are available in the Online Appendix.

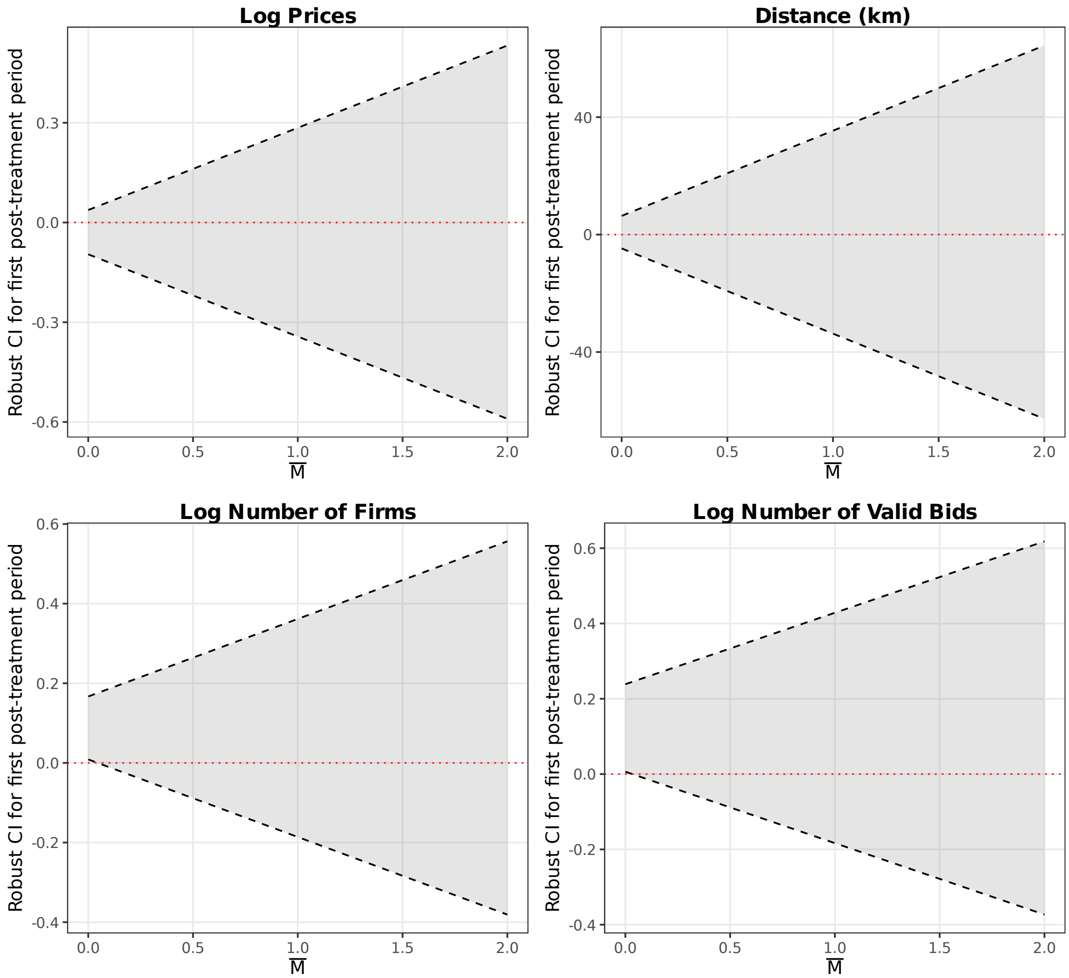

Figure A.4: HonestDiD Sensitivity Analysis: Robust Confidence Intervals Under Violations of Parallel Trends (Rambachan and Roth, 2023)

Table B.6: Lee (2009) Bounds: Sample Selection Correction

	Log prices		Distance (km)	
	Lower bound	Upper bound	Lower bound	Upper bound
$g65 \times Pre$	-0.1306***	-0.1227***	10.7750***	10.8042***
	(0.0096)	(0.0085)	(2.2322)	(2.2328)
Observations	625,816	625,814	625,814	625,814
R-squared	0.1870	0.1941	0.0009	0.0013
Trimming proportion		0.0882		
Item FE	YES	YES	YES	YES

Notes: Lee (2009) bounds correct for sample selection arising from treatment effects on completion rates. 18-month window. Lower/upper bounds obtained by trimming the outcome distribution in the excess-selected cell. DiD in completion rate: -0.0882. Standard errors clustered at the item level. *** p<0.01, ** p<0.05, * p<0.1.